The Forecast Ability of Option-Implied Densities from Emerging Markets Currencies*

José Renato Haas Ornelas**

Abstract

This paper empirically evaluates risk-neutral densities (RNDs) and real-world densities (RWDs) as predictors of emerging markets currencies. The dataset consists of volatility surfaces from 11 emerging markets currencies, with approximately six years of daily data, using options with one month to expiration. Therefore, there is a data overlapping issue, which is tackled with specific econometric techniques. Results of the out-of-sample assessment show that both RND and RWD underweight the tails of the actual distribution. This is probably due to the lack of options with extreme strikes. Although RWDs outperform RNDs in terms of the Kolmogorov distance, they still have problems in fitting the tails of actual data. Thus, risk aversion adjustment may improve the forecast ability, but it does not solve the misspecification of the distribution tails.

Keywords: Relative Risk Aversion, Risk-Neutral Density, Exchange Rates.

JEL Codes: C53, C13, E47, G17, F31.

*Submitted in February 2015. Revised in June 2015. I would like to thank the useful comments from Marcelo Bittencourt dos Santos, Sandro Andrade, José Fajardo, Aquiles Farias, and two anonymous referees. Part of this research was conducted while the author was a postdoctoral fellow at the University of Miami.
The views expressed in the paper are those of the author and do not necessarily reflect those of the Banco Central do Brasil.
**Banco Central do Brasil, Brazil. E-mail: `jrenato.ornelas@bcb.gov.br`

1. Introduction

Risk-neutral densities (henceforth RNDs) calculated from option prices have been used to infer market beliefs about future distributions of asset prices, such as exchange rates, interest rates, and stock indexes. Market participants often use these RNDs to make decisions about asset allocation and risk management. In fact, any empirical application in finance that requires density forecasts may also take advantage of RNDs. Regulators are also influenced by RNDs when making policy decisions.

The empirical extraction of RNDs started about 20 years ago with the papers of Shimko (1993) and Rubinstein (1994). During the 2000s, the relationship between RNDs and real-world densities[1] (henceforth RWDs) attracted the interest of researchers. An interesting point in this relationship is that we may calculate a relative risk aversion (henceforth RRA) measure by comparing both densities as originally done by Jackwerth (2000). This risk aversion measure can be then used to transform RNDs into RWDs as done by Bliss and Panigirtzoglou (2004) and Liu et al. (2007), among others. Empirical evidence has shown that these transformed densities have a better forecast ability than the pure RND.

Although the initial literature on RND and RWD focused on stock index options, other kinds of instruments have been analyzed, such as interest rates, commodities, and currencies. This paper focuses on emerging markets currencies.

The goal of this paper is to evaluate RNDs and RWDs as predictors of future outcomes of emerging markets currencies. A set of 11 emerging markets currencies is used in the empirical investigation. I use six years of daily data from over-the-counter (OTC) currency options with one month to expiration. Therefore, there are overlapping data. Most of previous papers filtered the data to get only non-overlapping time series in order to avoid econometric problems arising from autocorrelation. After filtering, the number of observations is reduced.

In this paper, I use econometric techniques to deal with overlapping data. Specifically, I tackle this issue by using the stationary bootstrap of Politis and Romano (1994) to adjust t-statistics and p-values for hypothesis testing regarding RRA. This way, I am able to use overlapping time series and then use all the available data without having to discard the overlapping data. To evaluate the performance of out-of-sample density forecasts, I follow Christoffersen and Mazzotta (2005) and test the moments of the normal transformed variable Z using Newey-West correction for standard errors.

Many papers have already extracted RNDs from emerging markets exchange rates (for instance, Abe et al., 2007, León and Casanova, 2004). Regarding RWD, the article of Fajardo et al. (2012) is the only one that estimates RRA and RWD for an emerging markets currency, but only for the Brazilian Real. Thus, this

[1]In the literature, "real-world," "risk-adjusted," "physical," "subjective," and "historical" refer to the same concept.

article contributes to the literature in two ways: first by empirically analyzing RNDs, RRA, and RWDs in a set of 11 emerging markets currencies; and second, by employing econometric methods that allow the use of overlapping data when dealing with risk aversion estimation.

Results of the out-of-sample assessment show that both RND and RWD place less weight on the tails of the actual distribution for all emerging markets currencies. A similar result was found by Christoffersen and Mazzotta (2005) for a sample of currency options from developed countries. This misspecification of the tail distribution is probably due to the lack of options with extreme strikes.[2] Although RWDs outperform RNDs in terms of the Kolmogorov distance, they still have problems in fitting the tails. Therefore, risk aversion adjustment may improve the forecast ability, but it does not solve the misspecification of distribution tails.

The paper is organized as follows. Section 2 briefly reviews the literature on RND and risk transformation methods. Section 3 shows the methodology used to estimate the RRA parameter and the RND and RWD distributions. Section 4 provides an overview of the dataset used. Section 5 shows the results of the RRA estimation. Section 6 presents the results of the out-of-sample density forecasts. Finally, Section 7 concludes the paper.

2. Risk-Neutral Density and Risk Transformation Methods

By assuming the existence of a representative agent, the real-world asset returns dynamics, some risk aversion functions, and risk-neutral probabilities are mutually related. The first two imply the third one. According to Jackwerth (2000), in each state of the world, the following relationship holds:

Risk-Neutral Probability=Real-World Probability×Risk Aversion Adjustment

The real-world probability is the representative investor's assessment of how likely that state is likely to occur. Risk aversion adjustment is useful when the investor values the dollar more highly in the "bad" states. By adjusting risk aversion over the real-world probability, we get the risk-neutral probability, which is the appropriate measure for derivative pricing. This risk aversion adjustment is also called a change of measure, specifically a change from the real world to the risk-neutral one. When investors are indifferent to risk, these probabilities are the same, and there is no need of risk aversion adjustment.

While the real-world probability distribution can be gathered from historical realized returns, the risk-neutral distribution is obtained from option prices. By comparing both distributions, we can empirically estimate the risk preferences embedded in the change of measure.

[2]While Christoffersen and Mazzotta (2005)'s data had strikes with a minimum delta of 25, my deltas are as low as 10.

One way to view this change of measure is to consider a specific utility function to describe risk preferences, as in Bliss and Panigirtzoglou (2004), and then use the following relationship between risk-neutral and objective densities:

$$h(x) = \frac{\frac{g(x)}{u'(x)}}{\int_0^\infty \frac{g(y)}{u'(y)} dy} \tag{1}$$

where $h(x)$ is the RWD, $g(x)$ is the RND, and $u(x)$ is the utility function.

Bliss and Panigirtzoglou (2004) use two types of utility function, the power and exponential utility functions, which contain only one parameter ρ. In the case of the power utility, this parameter is the RRA. Therefore, for a power utility function of the form $u(x) = \frac{x^{1-\rho}}{1-\rho}$, we would have:

$$h(x) = \frac{x^\rho g(x)}{\int_0^\infty y^\rho g(y) dy} \tag{2}$$

Once we have the RND $g(x)$, we can estimate the risk-aversion parameter by using the actual realizations of the underlying asset. In this paper, I use the approach of Liu et al. (2007), which uses power utility and calculates an RRA parameter from historical returns and options.

3. Distribution Estimation Methodology

A wide range of methods is available for estimating the RND $g(x)$. Jackwerth (1999) reviews this literature. Parametric or nonparametric methods can be used. There is a natural trade-off between the flexibility and stability of functions. Obviously, the higher the flexibility, the higher the in-sample goodness-of-fit. Micu (2005) analyzes several methods to extract RND for emerging markets currencies and suggests "... that there is a large scope for selection between these methods without essentially sacrificing the accuracy of the analysis."

In this paper, I use a mixture of lognormal distributions to model the RNDs. This is a choice for stability of the distributions. Also, Liu et al. (2007) provide an elegant transformation of RND into RWD, showing the relationship between the two sets of parameters and risk aversion.

Therefore, I model the exchange rate by using a mixture of two lognormal densities:

$$g\left(x|w, F_1, \sigma_1, F_2, \sigma_2\right) = w^* pdf_{LN}\left(x|F_1, \sigma_1\right) + (1-w)^* pdf_{LN}\left(x|F_2, \sigma_2\right) \tag{3}$$

with

$$pdf_{LN}\left(x|F, \sigma\right) = \left(x\sigma\sqrt{2\pi T}\right)^{-1} \exp\left(-\frac{1}{2}\left[\frac{\log(x) - \left(\log(F) - 0.5\sigma^2 T\right)}{\sigma\sqrt{T}}\right]^2\right) \tag{4}$$

where:
x is the exchange rate expressed as emerging market currency units per USD;
T is the time to maturity (one month in our case) expressed in years;
$\sigma, \sigma_1, \sigma_2$ are the volatility parameters;
F, F_1, F_2 are the future exchange rate parameters.

I use the forward exchange rate to reduce the number of free parameters of the mixture of lognormal distributions. This is done by making the expectation of the distribution equal to the forward exchange rate: $F = wF_1 + (1-w)F_2$. This way, F is the forward exchange rate quoted in the market. Therefore, out of five overall parameters, only four are free parameters. The parameters F_1 and F_2 are the expectation of the two lognormal distributions of the mixture, while the σ's set their volatility. The price of a European call option is the weighted average of two of Black (1976)'s call option formulas. The parameter estimation of the mixture of lognormals was done using an adaptation of the algorithm of Jondeau and Rockinger[3] for the characteristics and data of emerging markets exchange rates. One special issue when dealing with a mixture of lognormals is to avoid shapes with "spikes" (see Anderssona and Lomakka, 2005), i.e., one of the lognormals with a σ much lower than the other. This issue was avoided by allowing σ_i parameters to be at most twice as high as the other. Regarding optimization criteria, the algorithm estimates the parameters by minimizing the squared errors of the theoretical and actual option prices.

Once the RND is obtained, one can calculate the RRA parameter by following Liu et al. (2007)'s parametric risk transformation. As seen in Section 2, they consider the RWD h defined by (1) when there is a representative agent who has a constant RRA equal to ρ. If g is a single lognormal density, so is h. The volatility parameters for functions g and h are then equal, but their expected values are respectively F and $F\exp(\rho\sigma^2 T)$ when g is defined by (3).

Thus, a transformed mixture of two lognormals is also a mixture of two lognormals. For a mixture of lognormals $g(x|w, F_1, \sigma_1, F_2, \sigma_2)$ given by (3), it is shown by Liu et al. (2007) that the RWD h is also a mixture of lognormals with the following density:

$$\tilde{g}\left(x|w, F_1, \sigma_1, F_2, \sigma_2, \rho\right) = h\left(x|w', F_1', \sigma_1', F_2', \sigma_2'\right) \tag{5}$$

With the new set of transformed parameters given by:

$$\begin{aligned} F_1' &= F_1 \exp\left(\rho\sigma_1^2 T\right) \\ F_2' &= F_2 \exp\left(\rho\sigma_2^2 T\right) \\ \left(\frac{1}{w'}\right) &= 1 + \left(\frac{1-w}{w}\right)\left(\frac{F_2}{F_1}\right)^{\rho} \exp\left((1/2)\, T\left(\rho^2 - \rho\right)\left(\sigma_2^2 - \sigma_1^2\right)\right) \end{aligned}$$

[3]The original algorithm of Jondeau and Rockinger is available at `http://www.hec.unil.ch/MatlabCodes/rnd.html`.

I calculated an RRA for the full sample using the log-likelihood function as in Liu et al. (2007). For estimation of the RND parameters $(w, F_1, asigma_1, F_2, \sigma_2)$, I minimize the squared errors of the actual option price and of the theoretical option. For the RRA parameter (ρ), I maximize the log-likelihood function, using RRA as the only free parameter:

$$\sum_{i=1}^{N} \log\left(\tilde{g}\left(S_{i+1}|\hat{\theta}_i, \rho\right)\right) \tag{6}$$

where n is the number of days with volatility surface data and S_i is the spot exchange rate at time i.

It is worth noting that all quotes are in terms of emerging markets currency denominated in U.S. dollars, i.e., I am quoting the U.S. dollar instead of the risky asset. In this context, a negative ρ means investors demand a premium to hold the EM currency. This would be the RRA if the U.S. dollar were our risky asset. In order to have an RRA for the emerging markets currency, I change the sign[4] of ρ. To avoid ambiguity, I call this emerging market RRA parameter $\rho^{EM} = -\rho$.

4. Sample of Over-the-Counter Options

The exchange rate option prices used in this paper are over-the-counter (OTC). These data can be obtained from main data providers such as Thomson Reuters and Bloomberg. Both providers conduct a daily pool with market participants asking for their estimates of the volatility surface of OTC currency options. Thomson Reuters data consist of 17 strikes with deltas varying from 10% to 45% for calls and puts. Bloomberg data consist of quotes for nine fixed moneyness options. Therefore, Reuters data have more cross-sectional granularity than Bloomberg's. However, Bloomberg data were available for more currencies and with a longer time series. As the scope of this study is to calculate risk aversion, Bloomberg data seem to be more appropriate, given the longer time series.

The original data from Bloomberg consist of four risk reversals, four butterflies, besides the at-the-money volatility. Risk reversals and butterflies have four different deltas: 10, 15, 25, and 35. So, I have the following data for each day for each currency:

- At-the-money implied volatility (ATMV). This comes from a delta-neutral straddle implied volatility. A straddle is a set of call option and put option with the same strike. This quote corresponds to a strike that makes the Garman-Kohlhagen delta of the straddle equal to zero.

[4]This can be viewed if we think in terms of log-returns. The log-return of the conventional quote exchange rate (EM currency per unit of USD) is equal to minus the log-return of the inverted quote exchange rate (USD per EM currency unit). Therefore, their returns distributions are mirrored.

- Four different delta risk reversals. The risk reversal measures the difference in implied volatilities between an out-of-the-money call option with a specific delta and an out-of-the-money put option with the same delta. Option traders use risk reversal quotes to quantify the asymmetry of the implied volatility smile, which reflects the skewness of the risk-neutral currency return distribution. Mathematically, the risk reversals (RR) for each delta are calculated as follows:

$$\begin{aligned} RR10 &= IV(10c) - IV(10p) \\ RR15 &= IV(15c) - IV(15p) \\ RR25 &= IV(25c) - IV(25p) \\ RR35 &= IV(35c) - IV(35p) \end{aligned}$$

  where RR10 is the 10-delta risk reversal; IV(10c) is the implied volatility for a call with a delta of 10; IV(10p) is the implied volatility for a put with a delta of 10, and so on.

- Four different delta butterfly spreads. Butterfly spreads are defined as the average difference between out-of-the-money implied volatilities and the delta-neutral straddle implied volatility.

  Mathematically, the risk reversals (RR) for each delta are calculated as follows:

$$\begin{aligned} BF10 &= (IV(10c) + IV(10p))/2 - ATMV \\ BF15 &= (IV(15c) + IV(15p))/2 - ATMV \\ BF25 &= (IV(25c) + IV(25p))/2 - ATMV \\ BF35 &= (IV(35c) + IV(35p))/2 - ATMV \end{aligned}$$

  where BF10 is the 10-delta butterfly; ATMF is the at-the-money implied volatility; IV(10c) is the implied volatility for a call with a delta of 10; IV(10p) is the implied volatility for a put with a delta of 10, and so on.

Therefore, with the at-the-money implied volatility, four risk reversals, and four butterflies, we can recover the implied volatilities for calls and puts with the four different deltas. With the ATMV, we have nine data points in total.

The above data describe the volatility surface for each day and for each currency. I used only options with one month to expiration. The data include 11 emerging markets currencies: Brazilian Real (BRL), Chilean peso (CLP), Colombian peso (COP), Malaysian ringgit (MYR), Mexican peso (MXN), Indonesian Rupiah (IDR), Israeli Shekel (ILS), Philippine peso (PHP), Thai baht (THB), Turkish lira (TRY), and South African rand (ZAR). The surfaces cover the period from July 2007 to July 2013. There are some missing data during the period for some currencies, so we have an unbalanced panel of surfaces. Table 1 shows the number of surfaces for each currency. Other emerging markets currencies have data available from Bloomberg, but starting later.

One practical issue when dealing with OTC foreign exchange options regards the delta convention used. As mentioned by Reiswich and Wystup (2010), many academics overlook this issue. For emerging markets forex options, the premium-adjusted forward delta is often used, so I follow this convention.

Besides the data on options, I also collected data on the spot exchange rate, one-month forward exchange rate, and one-month deposit rate. All spot and forward exchange rates are from Bloomberg. For non-convertible currencies, in which offshore delivery is not possible, I used either on-shore forward or futures,[5] or non-deliverable forwards.

Data from the one-month interbank deposit rate were taken from several sources. For the Brazilian and Mexican markets, I used data from the Swap market. For IDR, ILS, MYR, TRY, THB, PHP, and ZAR, the data were taken from Libor-like deposits. For CLP, the Nominal Average Interbank Rate from "*Asociación Nacional de Bancos*" was used, while the "*Tasa Básica de la Superintendencia Bancaria*" was the deposit rate for COP.

It is worth noting that all quotes in those markets are obtained in terms of emerging markets currency per U.S. Dollar, which means that an appreciation (depreciation) of the EM currency decreases (increases) the exchange rate.

Table 1 shows the mean values for the deposit rate and volatilities. The one-month mean deposit rate was 5.2%. Brazil and Turkey had the highest rates, around 10.1%. Israel and Thailand had the lowest rates, slightly above 2%. For the same period, the USD Libor rate averaged 0.91%.

The at-the-money volatilities range from 6.7% for THB to 18.3% for ZAR and CLP, averaging 12.9%. All currencies had call volatilities higher than put volatilities, suggesting a skewness risk against emerging markets currencies.

Table 2 shows the descriptive statistics for the exchange rate returns. Again, it is worth recalling that exchange rate quotes are expressed in terms of emerging markets currency in U.S. dollars. Therefore, a positive total return means that the emerging markets currency has depreciated. This is the case of 8 out of 11 currencies. The largest depreciation is for TRY. The period of the sample starts before the financial crisis of 2008, when emerging markets currencies suffer considerably. Positive skewness is consistent with investors expecting more negative than positive surprises regarding these EM currencies. The kurtosis data show fatter tails than the Normal distribution in all cases, with an average kurtosis of 11.6. The historical volatility is lower than the implied ATM volatility (Table 2) for 7 out of 11 currencies. On average, the implied volatility is 0.9 percentage

[5] For the Brazilian Real, I used the dollar-Real futures contract traded at BM&F Bovespa exchange. This contract is the most liquid instrument on the Brazilian Real currency market. In order to have a constant expiration of one month, I interpolated the exchange rates of the first two contracts.

Table 1
Volatility and Deposit Rates – Descriptive Statistics

	#Days	Deposit rate	ATM volatility	Mean call volatility				Mean put volatility			
				10-delta	15-delta	25-delta	35-delta	10-delta	15-delta	25-delta	35-delta
BRL	1,467	10.1	15.8	22.1	20.8	18.7	17.2	14.3	14.3	14.4	14.9
CLP	1,462	4.9	18.3	17.3	15.7	14.6	13.6	12.9	12.7	12.7	13.0
COP	1,448	2.7	14.7	20.0	18.9	17.0	15.8	14.0	13.7	13.8	14.1
IDR	1,331	6.3	11.4	17.1	16.1	14.2	12.7	10.4	10.1	10.2	10.7
ILS	1,518	2.2	9.9	12.1	11.6	10.8	10.3	10.2	10.1	9.8	9.7
MYR	1,493	2.8	8.0	10.1	9.6	8.9	8.4	7.9	7.8	7.7	7.8
MXN	1,519	5.1	13.6	18.6	17.5	15.8	14.6	12.6	12.5	12.6	12.9
PHP	1,378	4.4	11.5	10.9	9.7	9.0	8.3	8.5	8.2	8.0	8.1
THB	1,453	2.4	6.7	8.9	8.3	7.5	7.0	7.6	7.1	6.7	6.7
TRY	1,548	9.6	13.4	17.9	17.0	15.5	14.4	12.5	12.4	12.5	12.8
ZAR	1,386	6.7	18.3	23.9	22.7	20.9	19.5	17.1	17.1	17.1	17.5
Mean	1,455	5.2	12.9	16.3	15.3	13.9	12.9	11.6	11.4	11.4	11.6

This table shows the descriptive statistics for 11 emerging markets currency options, namely: Brazilian Real (BRL), Chilean peso (CLP), Colombian peso (COP), Malaysian ringgit (MYR), Mexican peso (MXN), Indonesian Rupiah (IDR), Israeli Shekel (ILS), Philippine peso (PHP), Thai baht (THB), Turkish lira (TRY), and South African rand (ZAR). Data are from Bloomberg. The sample period goes from July 2007 to July 2013. For each currency, the number of days with available data is shown in the "# Days" column. The deposit rate with one-month maturity is shown on an annualized basis in percentage points. The ATM (at-the-money) volatility and the call and put volatilities are also expressed on an annualized basis and in percentage points. Call and put deltas follow the premium-adjusted forward delta convention. The maturity of the options is one month. All options are quoted considering an exchange rate expressed as emerging markets currency in U.S. dollars.

point higher than the historical one. Some currencies like CLP and PHP have a much higher implied volatility than the historical one.

5. Relative Risk Aversion Estimation Results

The main issue when estimating the RRA daily observations of one-month-ahead volatility surfaces is the overlapping nature of the data. Although the point estimate might remain the same, we cannot use the traditional likelihood ratio test approach to test if the RRA is statistically different from zero. When using non-overlapping data, the log-likelihood function from equation (6) has an asymptotic chi-square distribution. However, the use of overlapping data induces an autocorrelation in the monthly returns.

This kind of problem has already occurred in other research settings with overlapping data. For instance, Patton and Timmermann (2010) test the monotonicity of the term premium using overlapping data. This is possible thanks to the stationary bootstrap method proposed by Politis and Romano (1994). This approach is based on resampling blocks of random length, which breaks the autocorrelation structure of the data. This way, the issue of testing with overlapping data is overcome, and this makes the stationary bootstrap a suitable way to perform tests on the RRA estimates.

The idea is to generate several simulations using the stationary bootstrap in my data. Then, the RRA is calculated for each simulation. Finally, I sort the simulated RRAs in order to get a p-value for the desired hypothesis, i.e., an RRA different from zero. This is a novel approach for RRA testing with options, which is necessary because of the overlapping data. All similar previous studies used non-overlapping data (see Liu et al., 2007, Fajardo et al., 2012).

Table 2
Returns – Descriptive Statistics

	Total return	Historical volatility	Historical skewness	Historical kurtosis	Implied ATM-historical volatility
BRL	19.1	19.0	0.05	10.8	-3.2
CLP	0.7	12.0	0.57	7.5	6.3
COP	-4.1	13.6	-0.39	16.9	1.1
IDR	11.4	9.8	-0.08	18.7	1.6
ILS	-14.4	9.9	0.12	5.5	0.0
MYR	1.1	6.7	-0.10	5.8	1.3
MXN	16.0	14.0	0.71	12.8	-0.4
PHP	4.3	6.4	0.20	3.8	5.0
THB	-8.1	5.0	0.13	7.3	1.7
TRY	48.5	15.0	0.13	14.5	-1.6
ZAR	26.0	19.6	1.66	23.4	-1.4
Mean	9.14	11.9	0.27	11.6	0.9

This table shows the descriptive statistics for the returns of 11 emerging markets exchange rates. All exchange rates are quoted as emerging-market currency denominated in U.S. dollars. The currencies are: Brazilian Real (BRL), Chilean peso (CLP), Colombian peso (COP), Malaysian ringgit (MYR), Mexican peso (MXN), Indonesian Rupiah (IDR), Israeli Shekel (ILS), Philippine peso (PHP), Thai baht (THB), Turkish lira (TRY), and South African rand (ZAR). Data are from Bloomberg. The sample period goes from July 2007 to July 2013. The total return is the percentage return throughout the sample period, which may vary for each currency. The historical volatility is calculated based on daily continuously compounded returns, and then expressed on an annualized basis in percentage points. Skewness and kurtosis are calculated based on daily continuously compounded returns.

The results for the RRA estimation and the stationary bootstrap p-value are presented in Table 3. Point estimates show the presence of risk aversion for 10 out of 11 currencies. The relative risk aversion coefficient of 3.1 is in line with the previous literature. For instance, Bliss and Panigirtzoglou (2004) and Liu et al. (2007) found RRA coefficients between 2 and 4 using UK stock index options. Fajardo et al. (2012) found a 2.7 RRA coefficient using exchange-traded BRL options.

Nevertheless, the stationary bootstrap p-values fail to reject the hypothesis that these coefficients are different from zero, except for the case of IDR. Most of the p-values are around 20%, showing very weak statistical evidence of a risk premium against these currencies. However, as mentioned by Liu et al. (2007), it can be the case that these test conclusions are type II errors, reflecting the challenges of estimating risk premium accurately. As the risk premium is small relative to volatility, a large number of observations are needed to capture its true value. In fact, previous studies used more than 100 observations of volatility surfaces to obtain risk aversion estimates.

It is worth comparing our results with those of Fajardo et al. (2012) using BRL with exchange-traded options over 143 non-overlapping months from 1999 to 2011. Their RRA estimate of 2.7 is statistically different from zero with a p-value equal to 7.45%. Although the RRA is similar to our average across all EM currencies, it is approximately twice my estimate for the BRL, which is not statistically different from zero. I believe their longer time period could explain this result, but not the fact that I am using OTC instead of exchange-traded options.

The key issue concerning RRA estimation is the size of the sample. A way to increase sample size, keeping the same time length, is to pool several similar EM currencies into a single maximum likelihood estimation of the RRA parameter. This means using shorter time series and taking advantage of cross-sectional information. The idea behind this approach is that foreign investors in specific regions (e.g., Latin America or Asia) have similar behavior and risk attitudes, so that we could assume a common RRA. This would be also a measure of regional risk aversion.

Bearing this in mind, we may think of estimating a Latin American relative risk aversion using BRL, CLP, COP, and MXN together; or a Southeast Asian RRA using IDR, MYR, PHP, and THB together. The Latin American RRA estimate is 2.15, while the Southeast Asian RRA is 5.97.

Table 3
Relative Risk Aversion Estimates for the Full Sample

	Relative risk aversion (ρ^{EM})	Stationary bootstrap p-value
BRL	1.34	0.32
CLP	3.02	0.26
COP	3.78	0.20
IDR	8.74	0.03
ILS	5.43	0.21
MYR	1.23	0.44
MXN	0.93	0.40
PHP	5.33	0.20
THB	4.31	0.29
TRY	-0.43	0.52
ZAR	0.77	0.42
Mean	3.13	0.30

This table shows the relative risk aversion estimates of 11 emerging markets exchange rates for Brazilian Real (BRL), Chilean peso (CLP), Colombian peso (COP), Malaysian ringgit (MYR), Mexican peso (MXN), Indonesian Rupiah (IDR), Israeli Shekel (ILS), Philippine peso (PHP), Thai baht (THB), Turkish lira (TRY), and South African rand (ZAR). The sample period goes from July 2007 to July 2013. The relative risk aversion (ρ^{EM}) is estimated using a risk-neutral distribution with mixture of lognormal density and the parametric risk transformation of Liu et al. (2007). A positive risk aversion means investors charge a risk premium to carry the currency. The stationary bootstrap p-value is calculated using Politis and Romano (1994)'s stationary bootstrap in the time series of volatility surfaces, and then calculating the RRA coefficient for each simulation. The p-value is the percentile of $RRA = 0$ of the sorted simulated RRAs.

6. Density Forecast Evaluation

In the previous sections, both the RND and RWD were estimated for a set of emerging markets currencies. In this section, I assess the out-of-sample goodness-of-fit of these densities. The goal is to evaluate if these densities provide useful information about the future outcomes of the exchange rates. In fact, only the RND is fully out-of-sample, since the RWD uses an in-sample RRA.

I follow the density forecast evaluation literature and use the probability integral transformation in order to generate a time series U as follows:

$$U = \{U_i\} = \left\{ \tilde{g}_{CDF}^{-1} \left(S_{i+30} | \hat{\theta}_i, \rho \right) \right\} \tag{7}$$

Note that the density forecast U is built at time i with options that expire in one month, i.e., it is a forecast for $i + 30$. Then, we compare this forecast with the spot exchange rate one month ahead, i.e., S_{i+30}. As this is done for every business day, we have here an overlapping setting.

If the forecast densities are appropriate, U must be a uniform distribution with a domain in the range $[0, 1]$. The left side of Figure 1 shows the distribution of U for all 11 RND currencies. The right side shows U for the RWDs using the RRA calculated in the last section. In all cases, except for IDR, the tails of the distribution are higher than the mid-range. This means that option price densities are underweighting the tails when compared to the actual distribution, since the actual frequency of returns for that range was higher than the forecasted distribution.[6] This happens for both RND and RWD. However, in many cases, RWD reduces this underweighting, especially for the left tail. Therefore, although parametric risk transformation is not able to correct the thin tails of the RND, it minimizes the problem a little bit.

This pattern in which RNDs are not able to fit well the fatter tails of the actual distribution has already been identified previously in the literature. Christoffersen and Mazzotta (2005) found a similar problem with developed countries' foreign exchange rate options, and Castren (2004) with Eastern European currencies.

Christoffersen and Mazzotta (2005) mention that the fatter tails of the actual distribution compared to the RND could be attributable to the lack of very out-of-the-money options. Their sample had options with strikes with a minimum delta of 25, and RND forecast ability was appropriate only in the middle 70% range of the distribution. In fact, it is not possible to directly draw information about the distribution tails from options with strikes located after the most distant strikes. Some kind of extrapolation is needed.

In my sample, options are available with strikes as low as 10 deltas, going deeper into the tails of the distribution than did Christoffersen and Mazzotta (2005). The question that arises is whether these deltas would be far enough to accurately model, for instance, the first and last deciles of the distributions, which show poor goodness-of-fit, as seen in Figure 1. The strikes of the 10-delta puts and 10-delta

[6] In equation (7), if we have a higher number of observations in a specific range of S than the forecast by $\tilde{g}_{CDF}$, then the $\tilde{g}_{CDF}$ is underweighting the actual realizations of the exchange rate. Thus, in this case of underweighting, this range in the U function will be above average.

Figure 1 Out-of-sample Density Forecast Evaluation – U_i

Currency | Risk-Neutral Densities | Real-World Densities

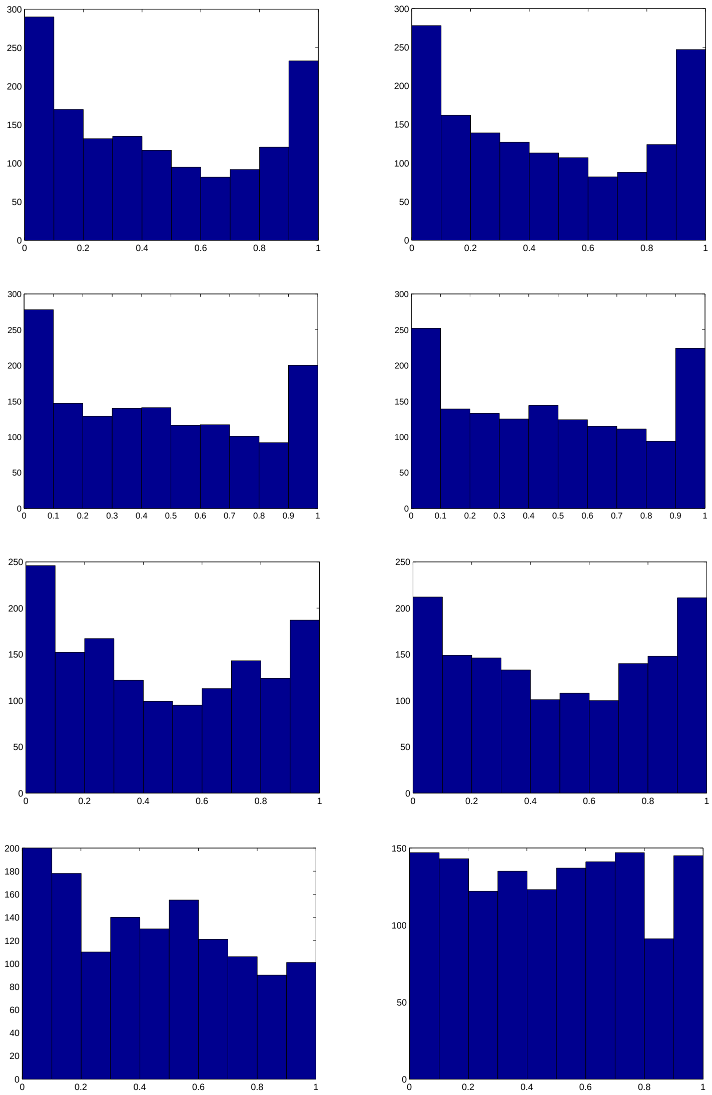

BRL

CLP

COP

IDR

Figure 1 Out-of-sample Density Forecast Evaluation – U_i (cont.)

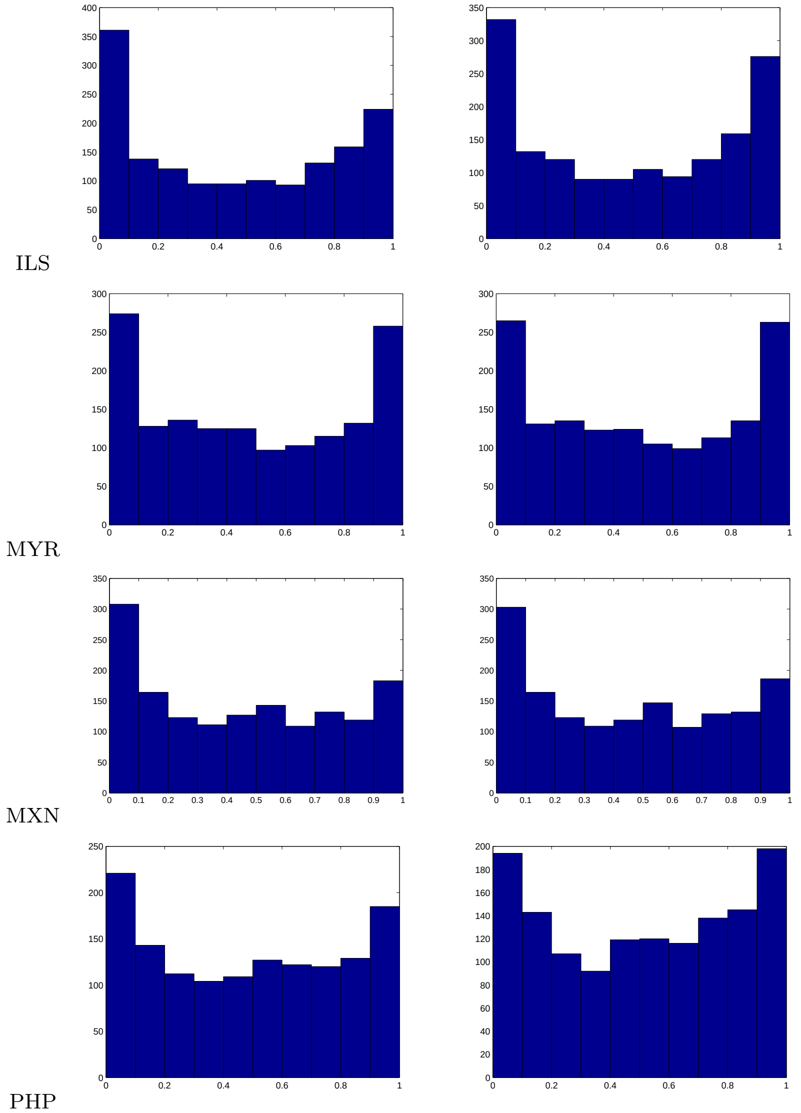

Figure 1 Out-of-sample Density Forecast Evaluation – U_i (cont.)

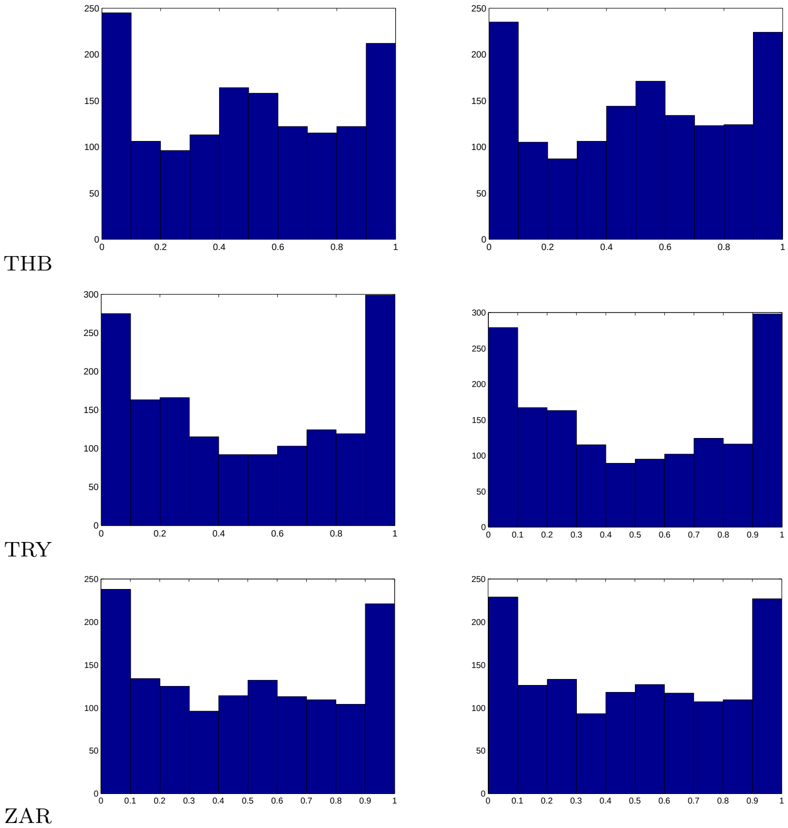

calls are located, on average, in the 10.3% and 93.2% cumulative distribution, respectively. This means that the first and last deciles of the CDF are estimated based (almost) on extrapolation. So, in this specific case, the functional parametric form of the distribution is relevant. Non-parametric density estimations could not fit these deciles, since they are not able to extrapolate. Parametric functions should have heavier tails than those estimated by the mixture of lognormals in this paper. I have also tried to estimate the parameters of the lognormal mixture by using a fitting method that gives emphasis on the tails of the distribution (see, for instance, Prause (1999), or Fajardo et al. (2005), so that fitting errors in the tails of the distribution would have more importance. However, this could not fix the problem. Other possible solutions would be to use time series data to shape the tails of the distribution or to use a more flexible parametric distribution such as the generalized hyperbolic distribution.

The pattern of strikes used to build the RND shows to be an important factor for tail accuracy. OTC option data usually have fixed strikes, while exchange-traded option strikes vary over time. The empirical evidence seems to favor exchange-traded options when assessing the accuracy of the tails. While Christoffersen and Mazzotta (2005), Castren (2004) and my paper use OTC forex options and find poor tail properties, Fajardo et al. (2012) and Craig and Keller (2005) use exchange-traded options and obtain good results.

There are also analytical ways to further evaluate the forecast. Berkowitz (2001) goes further and "normalizes" this U series using the inverse of the standard normal distribution, generating a Z series:

$$Z = \{Z_i\} = \left\{\phi^{-1}\left(U_i\right)\right\} \tag{8}$$

If the forecast density models are good, this Z series should follow a standard normal distribution. Then, it is possible to use a statistical test to check if the Z series are normally distributed. I calculate the Kolmogorov distance in these Z series in order to assess the quality of the density forecast (see Table 4). In all cases, the RWD provided a lower distance than the RND. Nevertheless, the Kolmogorov-Smirnov test rejects normality in all cases, except for the RWD of IDR. Therefore, the only case in which the parametric risk transformation was able to improve the RND forecast enough to pass the Kolmogorov-Smirnov test was in the IDR.

One issue with the above test is the autocorrelation in the series caused by data overlapping. Thus, I follow Christoffersen and Mazzotta (2005) to test if the Z moments are the same of a standard normal distribution, accounting for autocorrelation. So, the first and third Z moments should be zero, the second should be one, and the fourth should be three. This test can be done using the following set of regressions, with Newey-West standard errors to deal with data

Table 4
Kolmogorov Distances

	Risk-neutral densities (RND)	Real-world densities (RWD)
BRL	12.0%	10.8%
CLP	9.8%	8.0%
COP	9.2%	5.8%
IDR	8.9%	2.7%
ILS	14.4%	12.5%
MYR	8.4%	8.2%
MXN	12.0%	11.5%
PHP	7.3%	6.2%
THB	7.8%	7.1%
TRY	9.6%	9.5%
ZAR	7.5%	7.0%
Mean	9.7%	8.1%

This table shows the Kolmogorov distance of the normal transform variable Z for RND and RWD of 11 emerging markets exchange rates. The variable Z is calculated as in equation (8). If the density forecast is appropriate, Z should be normally distributed. The Kolmogorov distance measures the discrepancies of the Z variable from a normal distribution for each currency and for the risk-neutral densities and real-world densities. The currencies are: Brazilian Real (BRL), Chilean peso (CLP), Colombian peso (COP), Malaysian ringgit (MYR), Mexican peso (MXN), Indonesian Rupiah (IDR), Israeli Shekel (ILS), Philippine peso (PHP), Thai baht (THB), Turkish lira (TRY), and South African rand (ZAR). The sample period goes from July 2007 to July 2013.

overlapping:

$$\begin{aligned} Z_i &= a_1 + \varepsilon_{1,t} \\ Z_i^2 - 1 &= a_2 + \varepsilon_{2,t} \\ Z_i^3 &= a_3 + \varepsilon_{3,t} \\ Z_i^4 - 3 &= a_4 + \varepsilon_{4,t} \end{aligned} \tag{9}$$

The idea behind these regressions is that, if the forecast is appropriate, all coefficients a should be zero so that Z moments match those of a normal distribution. This can shed some light on why the Kolmogorov-Smirnov test rejected normality and what could be done to improve it. The results are shown in Table 5.

The point estimates of the first-moment coefficients (a_1) are mostly negative for the RND, reflecting a risk premium against EM currencies ($\rho^{EM} = -\rho > 0$). The parametric risk transformation brings this coefficient closer to zero in most of the cases, as can be seen in coefficients a_1 of the RWD. On average, coefficient a_1 is -0.06 for the RNDs and $+0.01$ for the RWD. However, in all cases, the coefficient is statistically equal to zero. Thus, although RWDs may be better centered than RNDs, the first moment does not seem to be a major problem for the accuracy of the RNDs.

Nevertheless, the second-moment tests show evidence that densities derived from option prices do not fit well the actual data. All coefficients a_2 are positive

Table 5
Tests of Z Moments

	First Moment (a_1)		Second Moment (a_2)		Third Moment (a_3)		Fourth Moment (a_4)	
	RND	RWD	RND	RWD	RND	RWD	RND	RWD
BRL	-0.06	-0.02	1.43	1.43	3.35	3.61	30.25	30.54
	(-0.38)	(-0.12)	(2.95)	(2.93)	(1.14)	(1.22)	(1.65)	(1.65)
CLP	-0.09	-0.01	1.46	1.46	3.67	4.15	38.06	38.84
	(-0.56)	(-0.06)	(2.80)	(2.74)	(1.08)	(1.20)	(1.72)	(1.71)
COP	-0.08	0.02	0.62	0.62	0.21	0.74	5.00	5.28
	(-0.64)	(0.19)	(3.32)	(3.15)	(0.32)	(1.07)	(2.94)	(2.79)
IDR	-0.16	0.03	0.33	0.29	0.49	1.30	5.26	5.83
	(-1.37)	(0.25)	(1.45)	(1.16)	(0.54)	(1.28)	(1.52)	(1.38)
ILS	-0.17	-0.06	0.92	0.92	-0.92	-0.30	6.73	6.70
	(-1.30)	(-0.48)	(4.74)	(4.74)	(-1.22)	(-0.39)	(3.27)	(3.13)
MYR	0.00	0.02	0.89	0.89	0.61	0.71	7.21	7.28
	(0.02)	(0.16)	(4.05)	(4.04)	(0.75)	(0.88)	(3.45)	(3.44)
MXN	-0.07	-0.04	1.55	1.55	5.90	6.04	51.93	52.20
	(-0.41)	(-0.26)	(2.27)	(2.26)	(1.18)	(1.20)	(1.39)	(1.39)
PHP	-0.02	0.06	0.68	0.69	0.49	0.91	7.46	7.73
	(-0.18)	(0.51)	(2.99)	(2.97)	(0.52)	(0.95)	(2.16)	(2.14)
THB	-0.06	0.00	0.88	0.89	-1.22	-0.93	10.22	10.17
	(-0.44)	(-0.03)	(3.06)	(3.10)	(-1.08)	(-0.83)	(2.68)	(2.73)
TRY	0.02	0.01	1.15	1.15	2.61	2.54	22.54	22.49
	(0.15)	(0.07)	(3.11)	(3.12)	(1.11)	(1.09)	(1.44)	(1.44)
ZAR	0.05	0.08	1.09	1.09	3.89	4.05	27.09	27.45
	(0.33)	(0.52)	(2.52)	(2.50)	(1.43)	(1.47)	(1.50)	(1.49)

This table shows the results for the set of regressions described in equation (9). This set of regressions is estimated for the RND and RWD of 11 emerging markets currency options, namely: Brazilian Real (BRL), Chilean peso (CLP), Colombian peso (COP), Malaysian ringgit (MYR), Mexican peso (MXN), Indonesian Rupiah (IDR), Israeli Shekel (ILS), Philippine peso (PHP), Thai baht (THB), Turkish lira (TRY), and South African rand (ZAR). All options are quoted considering an exchange rate expressed as emerging market currency in U.S. dollars. The point estimates of coefficients **a** are in bold and 21-lag Newey-West t-statistics are in parentheses.

and statistically greater than zero, except for the IDR. The fourth moment also shows performance problems with the forecasted densities. All coefficients a_4 are positive and most of them are statistically greater than zero. This evidence is consistent with fatter tails observed in Figure 1 and with the findings of Christoffersen and Mazzotta (2005) for major exchange rates, and those of Castren (2004) for Eastern European currencies.

Finally, the third moment shows mainly positive coefficients a_3, but none of them are statistically different from zero. This means that, on the negative EM currency return distribution side, the actual data show a higher probability than that forecasted by RND and RWD. This is a puzzling result, since it would be consistent with a negative skewness risk premium. Anyway, as the coefficients are not significant, this may not be a serious problem.

By comparing RNDs and RWDs, we see that the results are very similar, except for the first moment. Thus, the use of the risk transformation method could be justified just to fix the drift, but it seems useless in tackling tail misspecification.

Overall, the results in Table 5 show evidence against RND and RWD, especially regarding even moments. This analysis focuses on one moment at a time, i.e., it tests each coefficient **a** individually. It is also possible to test if coefficients **a** are jointly different from zero using a Wald test. I do it by testing all coefficients **a** at the same time and the odd and even moments separately. The results can be seen in Table 6.

Considering tests for odd (a_1 and a_3) and even (a_2 and a_4) moments separately, we see that the problems actually come from even moments, since none of the tests for odd moments for RWD was rejected at the 10% level. A good fit is observed only in the IDR obtained for even moments.

It is also possible to test if all coefficients **a** are jointly different from zero by using a Wald test as in Christoffersen and Mazzotta (2005). The results are also shown in Table 6. As in the Kolmogorov-Smirnov tests, the RWD of the IDR is the only distribution that does not reject normality at a 10% significance level of Z. However, by lowering the significance level to 1%, we would have six currencies being rejected and five not rejected.

Table 6
Joint Test of Z Moments

	All moments		Odd moments		Even moments	
	$(a_1) = (a_2) = (a_3) = (a_4) = 0$		$(a_1) = a_3) = 0$		$(a_2) = (a_4) = 0$	
	RND	RWD	RND	RWD	RND	RWD
BRL	23.14	21.93	4.53	3.71	18.06	17.72
	(0.000)	(0.000)	(0.104)	(0.167)	(0.000)	(0.000)
CLP	17.57	14.74	5.45	3.51	13.32	12.83
	(0.002)	(0.005)	(0.066)	(0.173)	(0.001)	(0.002)
COP	12.73	10.56	3.96	3.46	11.17	10.09
	(0.013)	(0.032)	(0.138)	(0.177)	(0.004)	(0.007)
IDR	10.33	4.64	9.98	3.93	3.45	3.39
	(0.035)	(0.326)	(0.007)	(0.140)	(0.178)	(0.183)
ILS	33.37	33.97	1.74	0.23	32.53	32.79
	(0.000)	(0.000)	(0.419)	(0.890)	(0.000)	(0.000)
MYR	18.09	17.93	2.96	2.96	17.30	17.31
	(0.001)	(0.001)	(0.228)	(0.228)	(0.000)	(0.000)
MXN	18.31	18.51	4.82	4.15	13.85	13.75
	(0.001)	(0.001)	(0.090)	(0.125)	(0.001)	(0.001)
PHP	10.37	10.22	1.43	1.11	10.16	10.17
	(0.035)	(0.037)	(0.490)	(0.573)	(0.006)	(0.006)
THB	9.72	10.3	2.05	2.52	9.53	9.81
	(0.045)	(0.036)	(0.358)	(0.283)	(0.009)	(0.007)
TRY	23.95	24.1	1.83	1.91	22.63	22.68
	(0.000)	(0.000)	(0.400)	(0.384)	(0.000)	(0.000)
ZAR	13.22	12.7	2.96	2.70	11.79	11.56
	(0.010)	(0.013)	(0.228)	(0.259)	(0.003)	(0.003)

This table shows the GMM estimation results for the regression system described in equation (9). The first two columns show the results for the joint (Wald) test, according to which all coefficients **a** are equal to zero for both RND and RWD. The Wald test statistics are in bold and the p-values of the Bartlett kernel with a 21-day bandwidth are in parentheses. The other columns show the results for the joint (Wald) test, indicating that the coefficients of the odd moments are zero ($(a_1) = (a_3) = 0$) and zero ($(a_2) = (a_4) = 0$) for the even moments. This set of regressions is estimated for the RND and RWD of 11 emerging markets currency options, namely: Brazilian Real (BRL), Chilean peso (CLP), Colombian peso (COP), Malaysian ringgit (MYR), Mexican peso (MXN), Indonesian Rupiah (IDR), Israeli Shekel (ILS), Philippine peso (PHP), Thai baht (THB), Turkish lira (TRY), and South African rand (ZAR). All options are quoted considering an exchange rate expressed as emerging markets currency in U.S. dollars.

7. Final Remarks

This paper evaluates the forecast performance of option-implied densities for emerging markets currencies. Results show that both RND and RWD fail to correctly forecast the tails of the realized distribution, specifically the first and last deciles. The reason is probably the lack of option data with strikes in this region. The use of a parametric risk transformation to build RWD from RND and a relative risk aversion parameter was not able to properly address tail misspecification. The relative risk aversion estimation shows weak evidence that investors are willing to charge a premium to invest in emerging markets currencies. However, this weak statistical evidence may be due to the small size of the premium, if compared with its volatility. Therefore, the use of RWD instead of RND brings only a limited advantage.

Despite this tail underweighting problem, these option-implied distributions could be used in applications in which the tails are not so important. For instance, they could be used in a mean-variance optimization process. The paper of Kostakis et al. (2011) uses S&P500 implied distributions to build optimal portfolios and then evaluate the performance of this procedure. This approach could be used for a multicurrency portfolio of emerging markets using the data in this paper.

In order to tackle tail misspecification, there exist two possible solutions: to use time series data to shape the tails of the distribution or to use a more flexible parametric distribution such as the Generalized Hyperbolic. Thus, an estimation process that blends time series and option-implied data through the use of a heavy-tailed distribution, such as Generalized Hyperbolic, is suggested as further research.

References

Abe, M. M., Chang, E. J., & Tabak, B. M. (2007). Forecasting exchange rate density using parametric models: The case of Brazil. *Brazilian Finance Review*, 5:29–39.

Anderssona, M. & Lomakka, M. (2005). Evaluating implied RNDs by some new confidence interval estimation techniques. *Journal of Banking and Finance*, 6:1535–1557.

Berkowitz, J. (2001). Testing density forecasts with applications to risk management. *Journal of Business and Economic Statistics*, 19:465–474.

Black, F. (1976). The pricing of commodity contracts. *Journal of Financial Economics*, 3:167–179.

Bliss, R. R. & Panigirtzoglou, N. (2004). Option implied risk aversion estimate. *Journal of Finance*, 59:407–446.

Castren, O. (2004). Do options-implied RND functions on G3 currencies move around the times of intervention on JPY-USD? Working papers series, European Central Bank.

Christoffersen, P. & Mazzotta, S. (2005). The accuracy of density forecasts from foreign exchange options. *Journal of Financial Econometrics*, 3:578–605.

Craig, B. & Keller, J. (2005). The forecast ability of risk-neutral densities of foreign exchange. Discussion Paper. Series 2: Banking and Financial Studies 05/2005, Deutsche Bundesbank.

Fajardo, J. S., Ornelas, J. R. H., & Farias, A. R. (2005). Analyzing the use of generalized hyperbolic distributions to value at risk calculations. *Brazilian Journal of Applied Economics*, 9:25–38.

Fajardo, J. S., Ornelas, J. R. H., & Farias, A. R. (2012). Estimating risk aversion. risk-neutral and real-world densities using Brazilian real currency options. *Brazilian Journal of Applied Economics*, 16:665–675.

Jackwerth, J. C. (1999). Option-implied risk-neutral distributions and implied binomial trees: A literature review. *Journal of Derivatives*, 7:66–82.

Jackwerth, J. C. (2000). Recovering risk aversion from option prices and realized returns. *Review of Financial Studies*, 13:433–451.

Kostakis, A., Panigirtzoglou, N., & Skiadopoulos, G. (2011). Market timing with option-implied distributions: A forward-looking approach. *Management Science*, 57:1231–1249.

León, A. D. d. & Casanova, M. E. (2004). Expectativas del mercado implícitas en los precios de instrumentos derivados: Aplicaciones al mercado cambiario y petrolero. Documento de Investigación 2004-01, Dirección General de Investigación Económica. Banco de México.

Liu, X., Shackleton, M., Taylor, S., & Xu, X. (2007). Closed-form transformations from risk-neutral to real-world distributions. *Journal of Banking and Finance*, 31:1501–1520.

Micu, M. (2005). Extracting expectations from currency option prices: A comparison of methods. *Computing in Economics and Finance*, 226. Society for Computational Economics.

Patton, A. & Timmermann, A. (2010). Monotonicity in asset returns: New tests with applications to the term structure, the CAPM, and portfolio sorts. *Journal of Financial Economics*, 98:605–625.

Politis, D. N. & Romano, J. R. (1994). The stationary bootstrap. *American Statistical Association*, 89:1303–1313.

Prause, K. (1999). *The Generalized Hyperbolic Model: Estimation, Financial Derivatives, and Risk Measures.* PhD thesis, University of Freiburg.

Reiswich, D. & Wystup, U. (2010). A guide to FX options quoting conventions. *Journal of Derivatives*, Winter:58–68.

Rubinstein, M. (1994). Implied binomial trees. *Journal of Finance*, 49:771–818.

Shimko, D. (1993). Bounds of probability. *Risk*, 6:33–37.