GARCH Based Risk Estimation in Emerging Market Foreign Exchange Rates

Henry Dyer

Applied Mathematics Undergraduate Honors Thesis

University of Colorado, Boulder

April 9, 2024

Thesis Advisor: Assistant Professor Yu-Jui Huang, Applied Mathematics

Honors Council Representative: Professor Anne Dougherty, Applied Mathematics

Outside Reader: Professor Carlos Martins-Filho, Economics

Abstract

We examine four autoregressive conditional heteroskedasticity (ARCH) type models, including one long memory and two asymmetric models, to assess their usefulness in Conditional Value at Risk and Conditional Expected Shortfall estimation. Alongside the four ARCH type models, we consider three additional models: historical simulation, standard parametric, and RiskMetrics. Estimation is performed on the five foreign exchange rates of the BRICS (Brazil, Russia, India, China, South Africa) emerging economies. We find that there is no single best model but that model selection for risk analysis should be done on an case by case basis. Furthermore, while the four ARCH type models produce similar results when estimating risk measures, we find that the standard GARCH model typically outperforms the asymmetric and long memory models when applied to out of sample data.

Contents

1 Introduction

Understanding volatility associated with financial assets plays a crucial role across many areas of the industry, including but not limited to pricing, asset allocation, and risk management. Because of this, volatility modeling has been extensively studied by academia, industry, and policymakers. This work is concerned with the modeling of foreign exchange rate volatility, and particularly in relation to risk management. Understanding foreign exchange rate volatility is critical as it can play a major role in production decisions for firms [Vania Stavrakeva and Tang, 2023] and foreign direct investment [Goldberg, 2006] among others. We focus on the foreign exchange rates of the original five emerging economies making up the BRICS intergovernmental organization – Brazil, Russia, India, China, and South Africa.

Modeling the volatility of financial assets is in itself a mathematical exercise. Many methods and models have been proposed, from simple historical simulation models to much more complex local or stochastic volatility models. Perhaps the most well known models, which we primarily focus on in this paper, are variants of the autoregressive conditional heteroskedasticity (ARCH) model of [Engle, 1982]. Specifically, we consider the popular generalized ARCH (GARCH) model of [Bollerslev, 1986], as well as a long memory and two asymmetry models. Each aim to build on the short comings of the ARCH model by incorporating additional stylized facts of the financial time series literature.

In addition to concerning ourselves with the modeling of foreign exchange volatility, we extend this work to evaluate the performance of certain volatility models in different risk management frameworks. The global financial crisis of 2008 prompted a wealth of new regulation related to risk management, including the Basel Accords and the Dodd-Frank Act, which has increased the need for better risk measures. Generating accurate risk measures is important both from the viewpoint of both regulators as well as portfolio managers. The importance of calculating an accurate downside loss is obvious, but making this as tight of a bound as possible is desirable as capital reserve requirements are frequently based on this estimate. We estimate two risk measures, namely Conditional Value at Risk (CVaR) and Conditional Expected Shortfall (CES), using the four ARCH variants previously mentioned as well as the RiskMetrics model [Morgan, 1996] a historical simulation model, and a standard parametric model. A multitude of other methods, including some based on extreme value theory (see [Genaçay and Selçuk, 2004] and [Martin Filho et al., 2018]), or conditional quantile estimation [Engle and Manganelli, 2004] have also been proposed but are not considered.

The rest of this paper is organized as follows. Section 2 introduces financial time series. Section 3 formulates the mathematics underlying the considered volatility. Section 4 describes the methods of risk estimation as well as methods used for evaluating their performance. Section 5 describes the data and presents the results of our exchange rate risk estimation. Section 6 discusses the results, and section 7 concludes.

2 Financial Time Series

Financial time series often deal with unique phenomena not present in the other time series data. This required the development of new models compared to traditional univariate time series models such as the Autoregressive Integrated Moving Average (ARIMA) model. As previously mentioned, this paper utilizes multiple variations of the ARCH model, each of which succeeds in modeling different portions of the stylized facts below. All estimations are done on the log returns

$$r_t = 100 * (\log r_t - \log r_{t-1})$$

which are scaled by a factor of 100 to assist with model parameter estimation.

2.1 Volatility

Traditional time series models, like the ARIMA model, assume a constant variance term over the lifetime of the process. In financial time series though, one need not look further than the price of the S&P 500 over the last decade to determine constant variance for financial assets is an unrealistic assumption. Successfully modeling the volatility of financial returns plays a key role in many areas across the financial industry. This paper is concerned with the volatility modeling directly related to risk management though volatility modeling plays an intricate role in other areas of financial mathematics. Due to this importance, much time has been spent researching different techniques to best capture volatility.

Likely the simplest method way to define volatility and the center piece of the standard parametric VaR and CVaR models described later is from the historical variance of returns y_t over the previous n periods

$$\sigma^2(y_t) = \frac{1}{n-1}\sum_{i=0}^{n}(y_{t-n} - \bar{y})^2$$

where $\bar{y}$ is the average return over the previous n periods. However, a major issue with this definition of volatility is that as previous periods of high volatility among asset prices drop out of the n period window, volatility calculations can change significantly even if the current market conditions appear to remain unchanged. Certain modifications to this historical variance measure, such as adding an exponentially declining weight to the previous terms have been well studied.

Another common way to calculate volatility is through back-calculation of the famous Black-Scholes model [Black and Scholes, 1973]. It calculates the price of a European call option, C, through the following formula

$$C = N(d_1)S_t - N(d_2)Ke^{-rt}$$

$$d_1 = \frac{\ln\left(\frac{S_t}{K}\right) + \left(r + \frac{\sigma^2}{2}\right)t}{\sigma\sqrt{t}} \qquad d_2 = d_1 - \sigma\sqrt{t}$$

where S_t is the current price of the underlying asset, K the options strike price, t the time left to maturity, σ the underlying asset's volatility, and N is the standard normal distribution. At any time, the current market price of the option is readily available, and all inputs into the equation besides volatility are known. From this, one can solve for volatility by working backwards from the Black-Scholes pricing equations. This is the basic idea behind volatility indices such as the VIX.

For the remainder of this paper, we consider the ARCH class of volatility models which are described in further detail in the following section.

2.2 Fat Tails

Normal distributions have often been used to model the unconditional returns of a financial asset. This would suggest that nearly all returns fall within three standard deviations of the mean, however, empirical results are not in line with this. It has become a stylized fact in the financial econometrics literature that returns are leptokurtic. This means that financial returns see large gains or losses more frequently than would be predicted by a normal. Distributions such as student-t or generalized error distribution are frequently employed to account for these fat tails. Additionally, conditional heteroskedasticity models, such as the ARCH class models considered in this paper, generate a fat tailed unconditional, even while sampling from a normal distribution.

2.3 Volatility Asymmetry & Clustering

Volatility asymmetry, also known as the leverage effect, claims that a negative volatility shock at time $t-1$ is likely to have a larger effect on the conditional variance at time t than a positive shock would at time $t-1$. This idea was formalized by [Engle and NG, 1993], who showed negative news, such as a poor inflation report, has a stronger effect on volatility than positive news.

The clustering of volatility within financial asset returns is also well documented. High volatility in previous periods is often associated with higher volatility in the current period, and low volatility is often preceded by additional periods of low volatility. Volatility clustering is possibly the largest benefit of ARCH models in financial time series as the variance in the current period is conditional on the variance of previous periods.

2.4 Long Memory

In the modeling of volatility in asset returns, long memory relates to the question of how long a period of low or high volatility will effect the current volatility. The ARCH and GARCH models have auto correlation functions which follow a power law, so in many cases this is likely leads to too quick of a decay of past shocks. On the other hand, integrated models, such as the IGARCH or RiskMetrics model we consider, have indefinite volatility persistence, which in many cases may seem to extreme. A possible solution considered in this paper is with fractionally integrated models. In particular, we consider the Fractionally Integrated GARCH model (FIGARCH) of [Baillie et al., 1996], which allows for a level of volatility persistence d such that $0 < d < 1$, where $d = 0$ would represent a traditional GARCH model with no long memory or $d = 1$ would represent an integrated model with infinite volatility persistence.

3 Conditional Heteroskedasticity Models

[Engle, 1982], proposed the autoregressive conditional heteroskedasticity (ARCH) model to assist with estimating the nonconstant variance process commonly found in asset returns. Since then, numerous variations of the ARCH model have been introduced to improve the capturing of the volatility process. Each model successfully captures different elements commonly found in financial time series, as described in the previous section, and does so differently, so it is appropriate to test multiple models.

This study explores a total of four variations of the ARCH model. We consider the Generalized ARCH (GARCH) model [Bollerslev, 1986], the Exponential GARCH (EGARCH) model [Nelson, 1991], Glosten-Jagannathan-Runkle GARCH (GJR-GARCH) model [Glosten et al., 1993], and finally the Fractionally Integrated GARCH (FIGARCH) model [Baillie et al., 1996]. The mathematical formulation underlying these models is described below, where each model assumes a sequence of log returns $\{r_t\}_{t=0}^{n}$, for $r_t = \mu + \epsilon_t$, where $\epsilon_t = \sigma_t z_t$ denotes the residual term at time t, σ_t is the time varying conditional variance, and z_t is a sequence of I.I.D $N(0,1)$ random variables.

3.1 ARCH

[Engle, 1982] introduced the autoregressive conditional heteroskedasticity (ARCH) model to assist with estimating the nonconstant variance process commonly found in asset returns. For the sequence of returns $\{r_t\}_{t=0}^{n}$ previously described, The conditional variance, σ_t is formulated as

$$\sigma_t^2 = \omega + \sum_{i=1}^{q} \alpha_i \epsilon_{t-i}^2$$

meaning the variance at period t is conditional on the squared returns of the previous q periods. ω and α_i are assumed to be > 0 to ensure a positive variance. We can solve for the unconditional variance of the ARCH(1) process assuming $\alpha_1 < 1$ for stationarity

$$\begin{aligned} Var(\epsilon_t) &= E[\epsilon_t^2] - E[\epsilon_t]^2 \\ &= E[\epsilon_t^2] \\ &= E[\sigma_t^2 z_t^2] \\ &= E[\sigma_t^2] \\ &= E[\omega + \alpha_1 \epsilon_{t-1}^2] \end{aligned}$$

Since the process is stationary, $Var(\epsilon_t^2) = Var(\epsilon_{t-1}^2)$ so substituting in, we recover

$$Var(\epsilon_t) = \frac{\omega}{1 - \alpha_1}$$

As previously mentioned, for an ARCH(1) model, we have a stationary process and finite unconditional variance when $\alpha_1 < 1$. This can be extended simply to an ARCH(q) model when $\sum_{i=1}^{q} \alpha_i < 1$. In the case when $\sum_{i=1}^{q} \alpha_i = 1$ the process is known as integrated, this is described in more detail in the following section.

3.2 GARCH & IGARCH

The Generalized ARCH (GARCH) model, [Bollerslev, 1986], builds upon the seminal work of [Engle, 1982] and allows for both longer memory and a more flexible lag structure within the process. A GARCH(p, q) process is given by

$$\sigma_t^2 = \omega + \sum_{i=1}^{q} \alpha_i \epsilon_{t-i}^2 + \sum_{i=1}^{p} \beta_i \sigma_{t-i}^2$$

Setting p = 0, the process reduces to the familiar ARCH(q) process. Compared to the ARCH(q) model, which only treats conditional variance as a linear combination of previous squared residuals (often referred to as the ARCH terms), the GARCH(p, q) model also allows for conditional variance to be a function of the previous conditional variances (often referred to as the GARCH term). This flexibility allows for more accurate volatility modeling. In fact, with minimal work, it can be shown the GARCH(1, 1) process is equivalent to an ARCH(∞) process

Proof.

$$\begin{aligned}
\sigma_t^2 &= \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \\
&= \omega + \alpha\epsilon_{t-1}^2 + \beta(\omega + \alpha\epsilon_{t-2}^2 + \beta\sigma_{t-2}^2) \\
&= \omega + \alpha\epsilon_{t-1}^2 + \beta(\omega + \alpha\epsilon_{t-2}^2 + \beta(\omega + \alpha\epsilon_{t-3}^2 + \beta\sigma_{t-3}^2)) \\
&= \omega\sum_{i=0}^{\infty}\beta^i + \sum_{i=0}^{\infty}\alpha\beta^i\epsilon_{t-i-1}^2 \\
&= \frac{\omega}{1-\beta} + \sum_{i=0}^{\infty}\alpha\beta^i\epsilon_{t-i-1}^2
\end{aligned}$$

Let $\phi_0 = \frac{\omega}{1-\beta}$ and $\phi_i = \alpha\beta^i$ and we recover the the formulation for an ARCH(∞) model. □

This paper considers only the GARCH(1, 1) model

$$\sigma_t^2 = \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2$$

as opposed to a GARCH(p, q) model as previous related work has found $p, q = 1$ to best fit financial time series. As can be shown in a similar method as above, the GARCH(1,1) process is covariance stationary when $\alpha+\beta < 1$, and furthermore, a GARCH(p, q) model is covariance stationary when $\sum_{i=1}^{q}\alpha_i + \sum_{i=1}^{p}\beta_i < 1$. Bollerslev and Engle also consider the possibility when the process contains a unit root, $\alpha + \beta = 1$, meaning changes in variance are persistent, which is referred to as the Integrated GARCH (IGARCH) model

$$\sigma_t^2 = \omega + \beta\epsilon_{t-1}^2 + (1-\beta)\sigma_{t-1}^2$$

when $\beta = 0.94$, the RiskMetrics model considered in this paper is derived.

3.3 Asymmetric Models

We consider two asymmetric GARCH models. First, the Exponential GARCH (EGARCH) model of [Nelson, 1991], and second, the GJR-GARCH model of [Glosten et al., 1993]. While the GARCH model above succeeds in capturing the well-documented fat tails in financial time series, it treats positive and negative shocks to the variance process the same, failing to capture the asymmetry/leverage effect commonly found in asset returns. Both the EGARCH and GJR-GARCH models allow for these asymmetries while simultaneously capturing the excess kurtosis, though the two of them achieve this through different methods.

An EGARCH(p, q) model is formulated as follows

$$\ln(\sigma^2) = \omega + \sum_{i=1}^{q}(\alpha_i(|z_{t-i}| - \mathbb{E}(|z_{t-i}|)) + \gamma_i z_{t-i}) + \sum_{i=1}^{p}\beta_i \ln(\sigma_{t-i}^2)$$

Similar to the GARCH model, we only consider the EGARCH(1, 1) model

$$\ln(\sigma^2) = \omega + \alpha(|z_{t-1}| - \mathbb{E}(|z_{t-1}|)) + \gamma z_{t-1} + \beta \ln(\sigma_{t-1}^2)$$

Notice that the residual in the previous period, z_{t-1} has coefficient $\alpha - \gamma$ for a negative shock and coefficient $\alpha + \gamma$ for a positive shock. This introduces the previously mentioned asymmetry in volatility. Furthermore, previous literature frequently finds $\gamma < 0$ at a statistically significant level, which is in line with the news impact curve discussed earlier, meaning that negative shocks in the previous period have a larger effect on conditional variance than positive shocks.

On the other hand, the GJR-GARCH(p, q) model assumes the following conditional volatility structure

$$\sigma_t^2 = \omega + \sum_{i=1}^{q}(\alpha_i + \gamma I_{t-i})\epsilon_{t-i}^2 + \sum_{i=1}^{p}\beta_i\sigma_{t-i}^2$$

where I is an indicator variable taking value 1 if the previous return, $\epsilon_{t-1} < 0$, and 0 if $\epsilon_{t-1} \geq 0$. Again we consider only the GJR-GARCH(1, 1) model

$$\sigma_t = \omega + (\alpha + \gamma I_{t-1})\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2$$

Compared this to the EGARCH model, which captures asymmetry through the difference in $\alpha - \gamma$ and $\alpha + \gamma$, the GJR-GARCH model allows for asymmetry by introducing an additional γ variable when the previous residual, $\epsilon_{t-1} < 0$.

3.4 Long Memory Models

We consider one long memory model, the Fractionally Integrated GARCH (FIGARCH) model of [Baillie et al., 1996], and follow their convention for our mathematical formulation. We first consider the GARCH(1, 1) model expressed as an ARMA(1, 1) process on the squared residuals ϵ_t^2, by letting $v_t = \epsilon_t^2 - \sigma_t^2$.

$$\begin{aligned}
\sigma_t^2 &= \omega + \alpha\epsilon_{t-1}^2 + \beta\sigma_{t-1}^2 \\
\epsilon_t^2 - v_t &= \omega + \alpha\epsilon_{t-1}^2 + \beta(\epsilon_{t-1}^2 - v_{t-1}) \\
\epsilon_t^2 - (\alpha+\beta)\epsilon_{t-1}^2 &= \omega + v_t - \beta v_{t-1} \\
(1 - \alpha(L) - \beta(L))\epsilon_t^2 &= \omega + (1 - \beta(L))v_t
\end{aligned}$$

Where L is the lag operator, $\alpha(L) = \alpha_1 L + \cdots + \alpha_n L^n$ and $L^k\epsilon_t^2 = \epsilon_{t-k}^2$. If the process contains a unit root, $\alpha + \beta = 1$, and therefore is integrated, the process is then represented as

$$\phi(L)(1 - L)\epsilon_t^2 = \omega + (1 - \beta(L))v_t$$

where $\phi(L) = [1 - \alpha(L) - \beta(L))](1 - L)^{-1}$. From this representation, formulating the Fractionally Integrated, FIGARCH, model is done by adding a fractionally differencing operator $(1 - L)^d$ into the previous equation where $d \in (0, 1)$.

$$\phi(L)(1 - L)^d\epsilon_t^2 = \omega + (1 - \beta(L))v_t$$

If $d = 1$, we recover the IGARCH model, and for $d = 0$, the GARCH model. In order to acquire a form where the conditional variance σ_t^2 can easily be estimated, we first substitute $\epsilon_t^2 - \sigma_t^2 = v_t$ back into the above equation.

$$\begin{aligned}
\phi(L)(1 - L)^d\epsilon_t^2 &= \omega + (1 - \beta(L))(\epsilon_t^2 - \sigma_t^2) \\
\phi(L)(1 - L)^d\epsilon_t^2 &= \omega + (1 - \beta(L))\epsilon_t^2 - (1 - \beta(L))\sigma_t^2 \\
(1 - \beta(L))\sigma_t^2 &= \omega + (1 - \beta(L) - \phi(L)(1 - L)^d)\epsilon_t^2 \\
\sigma_t^2 &= \frac{\omega}{1 - \beta} + (1 - (1 - \beta(L))^{-1}\phi(L)(1 - L)^d)\epsilon_t^2
\end{aligned}$$

setting $\lambda(L) = 1 - (1 - \beta(L))^{-1}\phi(L)(1 - L)^d$, we see that the conditional variance of the FIGARCH model can be represented, similar to a GARCH(1, 1) model, as an ARCH(∞) model

$$\sigma_t^2 = \frac{\omega}{1 - \beta} + \sum_{i=1}^{\infty} \lambda_i \epsilon_{t-i}^2$$

. Like previous models, we consider the most basic FIGARCH(1, d, 0) model. , and follow their convention for our mathematical formulation. We show that for this model, the λ coefficients required for estimation can be derived as follows

$$\lambda_k = \left(1 - \beta - \frac{1 - d}{k}\right) \frac{\Gamma(k + d - 1)}{\Gamma(k)\Gamma(d)}$$

where $\Gamma(x)$ represents the Gamma function. We use this method for our conditional variance estimation with the fractionally integrated GARCH model.

4 Risk Estimation

Effective risk management in finance is essential from the perspective of all parties involved. Investors and portfolio managers care about the potential losses of their investments while regulators often use these metrics to determine the required capital reserves should losses entail. While Many risk methods have been proposed, we focus on two widely used models, Conditional Value at Risk (CES) and Conditional Expected Shortfall (CES). Both CVaR and CES, along with the majority of risk measures, are a mathematical exercise in estimating a potential loss distribution. They differ in how this distribution is evaluated. CVaR is a measure of the potential loss at some percentile while CES is a measurement of the expected value of all possible profits or losses below some level α conditioned on all previous information. Selecting a particular risk measurement is based on preference as there are pros and cons to calculating and interpreting all models.

4.1 Conditional Value at Risk

Following the definition given by [Martin Filho et al., 2018] and [Sarykalin et al., 2008], we define CVaR_α for $\alpha \in [0, 1]$, as

$$\text{CVaR}_\alpha = \min\{x | F_X(x) \geq \alpha\}$$

where X is a random variable with CDF $F_X(x)$. For our work, we assume a normal distribution of returns for period $t \sim N(\mu, \sigma_t^2)$ so,

$$\alpha = \int_{-\infty}^{\text{CVaR}_\alpha} \frac{1}{\sigma_t\sqrt{2\pi}} exp\left(-\frac{1}{2}\left(\frac{x-\mu}{\sigma_t}\right)^2\right) dx$$

A common issue with CVaR estimates is that many portfolios can sustain large losses below the confidence level α which make it unrepresentative of the associated risk. Take for instance, a portfolio containing a single deep out of the money option. This option may payout \$1 99% of the time, and the other 1% of the time, require a payment of \$99. A fifth percentile CVaR would measure this portfolio to contain zero risk, which obviously is not the case.

4.2 Conditional Expected Shortfall

One solution to this issue is using a Conditional Expected Shortfall as an alternative risk measure. Similar to CVaR, we integrate over the assumed risk distribution (a normal for this work) on the interval $(-\infty, \text{CVaR}_\alpha)$. However for CES, we calculate an expected value by introducing a function $p(x)$ which represents the profit or loss of the portfolio for an outcome x and adding a normalizing constant $\frac{1}{\alpha}$

$$\text{CES}_\alpha = \frac{1}{\alpha}\int_{-\infty}^{\text{CVaR}_\alpha} p(x)\frac{1}{\sigma_t\sqrt{2\pi}} exp\left(-\frac{1}{2}\left(\frac{x-\mu}{\sigma_t}\right)^2\right) dx$$

By considering all possible losses below a confidence level α, CES protects against major losses that are hidden by a single CVaR metric. However, there are many instances where

CVaR can be a better metric than CES. This "blindness" to extreme tails of distributions can lead to better predictions as data surrounding the tails is not readily available and can end up weighting potential outcomes more heavily than they are in reality leading CES to take to conservative of a position. Finally, note that as CVaR is only concerned with the loss at the upper bound of the interval $(-\infty, \mathrm{CVaR}_\alpha)$, it is a lower bound on CES.

4.3 The Historical Simulation Model

The historical simulation model is the simplest of the models we consider, however this simplicity continues to make it appealing for reasons both computational and conceptual. The model estimates CVaR and CES based off a selected n previous payouts. This paper uses $n = 252$, which corresponds to the 252 trading days in a calendar year. For CVaR, the historical simulation for a confidence level α is the α-percentile corresponding to the previous returns being considered. For example, the $\alpha = 5\%$ CVaR over the previous 100 periods would simply be the loss incurred on the fifth worst day over the period. As expected, CES is the average of all returns incurred below α.

Another advantage of the historical simulation model is that it is nonparametric so no assumptions regarding the shape of the underlying distribution are required. On the other hand, historical CVaR and CES have no forward looking properties and assumes volatility in asset returns are related to returns over the previous n periods. Additionally, as data points drop out of the dataset as time moves on, historical simulation models can produce greatly varying measurements even in constant market conditions.

4.4 Parametric Models

The remaining estimation techniques considered in this paper all fall under the umbrella of parametric models, meaning we must make an assumption about the distribution underlying data sample. As previously mentioned we assume a normal distribution of returns, other works consider 'fat tailed' distributions such as the Student-t or Generalized Error Distribution.

The most basic parametric model considered in this paper again uses data over the previous 252 periods to calculate the CVaR or CES in the next period. This model, known as the Parametric model, gives equal weights to each data point and uses them to calculate the parameters (mean and variable in our case) for the assumed underlying distribution. Similar to the historical simulation model, a shortfall here is that as previous returns fall out of the window T, this can have large effects on the parameter estimates even if market conditions remain the same. A common solution to this is adding a exponential weight to the previous returns so that previous returns are factored in more than older returns, for this Exponentially Weighted Moving Average (EMWA) model variable for an exponential decay factor $\lambda \in (0,1)$, is calculated as follows

$$\sigma_t^2 = (1-\lambda)\sum_{i=0}^{t} \lambda^i \epsilon_{t-i-1}$$

One can follow a the same procedure of the proof demonstrating GARCH(1, 1) is equivalent to ARCH(∞) to show that the EMWA model is the same as an IGARCH(1, 1) model as $t \to \infty$. For this reason, we set $\lambda = 0.94$ as specified by the RiskMetrics model [Morgan, 1996] and treat the above EMWA as our RiskMetrics model.

Finally, for risk estimation with GARCH models, the variance in calculated by estimating the model parameters on in sample data and used to forecast volatility (variable) for the following period.

4.5 Evaluating Risk Estimation Methods

For a multitude of reasons, producing accurate risk measurements is of high importance for banks, regulators, or portfolio managers. If the goal in risk management was simply to find a bound such that no returns would fall below it, one could set the bound at zero and sleep well. However, many regulations set capital reserve requirements among other things depends based on the produced risk estimate. It is in the best interest of all parties to produce as tight of risk measurement as possible. In the following section, we evaluate CVaR and CES models based on two metrics. First off is how often their risk estimates are broken, we would expect that a well calculated fifth quantile CVaR would be broken roughly five percent of the time. If instead the CVaR estimates are broken 10% of the time, then the estimate was clearly to aggressive. On the other hand, if the CVaR estimate was only broken 1% of the time, it was to conservative which will harm the financial institution. The second metric used to evaluate CVaR and CES performance is a goodness of fit or accuracy test, where models are rewarded for estimating a tight yet accurate bound and penalized for being overly conservative. This metric, $s(x_{o,t}, x_{\alpha,t})$ is defined as follows:

$$s(x_{o,t}, x_{\alpha,t}) \begin{cases} \delta(x_{o,t} - x_{\alpha,t})^2 & : x_{o,t} < x_{\alpha,t} \\ \frac{1}{\delta}(x_{o,t} - x_{\alpha,t})^2 & : x_{o,t} \geq x_{\alpha,t} \end{cases}$$

where $x_{o,t}$ is the observed return for an asset in period t, and $x_{\alpha,t}$ is the estimated risk parameter for a level α in period t and δ is a chosen parameter to best fit a specific situation, the rest of this paper uses $\delta = 4$. The better a risk metric is, the closer its score will be to zero. Also notice that the above function punishes metrics which are overly aggressive more than metrics which are less aggressive. While neither of these risk measure evaluation metrics are end all be all's, considering them in tandem can help risk managers make appropriate choices of differing models.

5 Data & Empirical Work

The empirical section of this paper focuses on evaluating the performance of Conditional Value at Risk and Conditional Expected Shorfall estimation through four GARCH methods, RiskMetrics, historical simulation, and a standard parametric model. The analysis is performed on the exchange rates of the five countries making up the BRICS emerging economies – Brazil, Russia, India, China, and South Africa.

All exchange rate data is retrieved from the *Bank of International Settlements (BIS)* over the period from the beginning of 2001 to the end of 2022, with the exception of China, where data is only considered starting January 1st 2006 as its foreign exchange rate was fixed until mid 2005. Over this time period we estimated parameters for the GARCH models on the first 3335 of 4335 measurements and tested their estimation for CVaR and CES on the final 1000 out of sammple data points. The Data is summarized in the table below.

	Brazil	China	India	Russia	South Africa
Mean	0.01397	-0.00789	0.00781	0.01616	0.01143
Minimum	-10.93549	-1.20956	-3.08915	-10.06088	-8.15293
Maximum	11.37772	1.84037	4.32598	19.322483	8.52607
Std. Dev.	1.13925	0.141208	0.423510	0.78474	1.09986
Skewness	0.11482	0.70120	0.35608	2.94645	0.42191
Kurtosis	9.90491	117.41076	8.34125	102.15976	4.40992
Jarque-Bera	17730(0.0)	39040(0.0)	12758(0.0)	1891387(0.0)	3641(0.0)

Table 1: Summary of in sample BRICS foreign exchange rate data.

As evident from the Jarque-Bera test for normality (the P-Value for the JB test is in parenthesis) at the bottom of the table as well as the kurtosis greater than three, none of the exchange rates fit a normal distribution. We also observe a range of standard deviations making these assets interesting to investigate the performance of differing CVaR and CES methods.

We now present estimates for the parameters of the four GARCH models for each asset. parameters marked with a * represent significance at 1% level while ** represents significance at the 5% level.

	Brazil	China	India	Russia	South Africa
GARCH					
ω	0.0290*	0.3868**	0.0001	0.0004**	0.0138**
α	0.1483*	0.1968*	0.0750*	0.0829*	0.0610*
β	0.8348*	0.6237*	0.9250*	0.9171*	0.9285*
EGARCH					
ω	0.0004	0.0108*	-0.0169	0.0018	0.0051
α	0.2135*	0.0422**	0.2055*	0.2058*	0.1300*
γ	0.0809*	0.0183**	0.0289**	0.0261**	0.0387*
β	0.9652*	0.995*	0.9853*	0.9918*	0.9841*
GJR-GARCH					
ω	0.0278*	0.3487	0.0001	0.0002	0.0139**
α	0.1858*	0.1618*	0.1138**	0.1412*	0.0723*
γ	-0.1142*	-0.0158	-0.0332	-0.0471**	-0.0326**
β	0.8550*	0.6377*	0.8995*	0.8844*	0.9324*
FIGARCH					
ω	0.0447*	0.3938	0.0028	0.0013	0.0493**
ϕ	0.1126	0.4233	0.3015**	0.2212*	0.2434*
d	0.5775*	0.1534	0.3971*	0.5576*	0.3754*
β	0.5183*	0.3762	0.5587*	0.6746*	0.5482*

Table 2: Estimates of GARCH Parameters

Different from expected, we observe a positive γ for the EGARCH models and a negative γ in the GJR-GARCH model, with the exception. This means that positive returns have a larger effect on volatility than negative returns, thought the effect is minor for all countries besides Brazil. The γ estimations for both models agree on this. A possible reason for this is that exchange rates, unlike stock prices, are bilateral so one could also interpret these coefficients as negative shocks to the US Dollar (the forex denomination) increasing the volatility in the current period. Again for all currencies except China, we notice a very high GARCH term β meaning that the conditional variance in period $t-1$ has a large effect on the conditional variance in period t. The ARCH term α is also much different from zero in all cases so the previous returns play a large role in the current conditional variance.

Finally, for the FIGARCH models, we observe differencing parameters d loosely around 0.5 for all models except China, who's lack of data likely contributed to less statistically significant parameter estimations overall, meaning that volatility persistence was captured in the returns to a higher degree than considered by GARCH models and a lower degree than the IGARCH/RiskMetrics model assumes.

With parameters generated for each model, we apply these models to out of sample data for the following 1000 days to assess their performance in CVaR and CES modeling. Both CVaR and CES are estimated to confidence levels of $1\%, 2.5\%,$ and 5%, all common levels among practitioners. Each table shows the percentage of time the estimated CVaR or CES level was broken for the corresponding confidence levels.

	GARCH	EGARCH	GJR	FIGARCH	Historical	Parametric	RiskMetrics
$\alpha = 1.0\%$							
Brazil	3.6	3.8	4.0	3.5	1.9	2.1	1.2
China	1.9	1.8	1.9	1.4	1.3	1.7	2.2
India	0.9	0.9	0.9	0.9	1.4	1.3	1.2
Russia	0.0	0.0	0.2	0.0	1.4	0.6	1.3
South Africa	2.9	2.8	3.0	2.6	1.1	1.1	1.1
$\alpha = 2.5\%$							
Brazil	5.6	5.1	5.2	4.9	3.3	2.9	2.6
China	2.1	2.1	2.1	1.8	2.4	2.2	2.6
India	1.2	1.2	1.2	1.2	3.0	2.2	1.8
Russia	0.2	0.1	0.3	0.1	2.9	1.8	2.2
South Africa	5.2	5.1	5.5	5.1	2.4	2.1	2.3
$\alpha = 5.0\%$							
Brazil	7.5	6.9	7.3	7.0	5.6	5.1	5.0
China	2.2	2.2	2.2	2.1	4.6	4.2	4.4
India	1.5	1.4	1.6	1.5	6.1	4.2	4.1
Russia	0.3	0.3	0.4	0.3	5.5	3.3	3.3
South Africa	7.5	7.7	7.7	7.2	5.3	4.1	4.7

Table 3: Conditional Value at Risk Estimation

We save formal discussions of the performances of the differing models but we quickly note that there is a large disparity among performance of the same model depending on the currency being considered. We now present estimation for the CES of each currency following the same format as above.

	GARCH	EGARCH	GJR	FIGARCH	Historical	Parametric	RiskMetrics
$\alpha = 1.0\%$							
Brazil	2.7	2.8	2.8	2.7	0.7	0.8	0.5
China	1.5	1.6	1.5	1.3	0.6	1.0	1.8
India	1.0	0.8	0.9	0.9	0.6	0.7	0.9
Russia	0.0	0.0	0.0	0.0	0.3	0.4	0.7
South Africa	2.1	2.2	1.9	1.9	0.3	0.4	0.6
$\alpha = 2.5\%$							
Brazil	3.4	3.5	3.5	2.9	1.7	1.3	1.1
China	1.8	2.0	1.9	1.7	0.1	1.8	2.3
India	1.2	0.9	1.2	1.1	1.3	1.1	1.1
Russia	0.0	0.0	0.1	0.0	0.8	0.5	0.9
South Africa	2.5	2.9	2.6	2.3	1.0	0.7	1.0
$\alpha = 5.0\%$							
Brazil	4.2	4.2	4.2	3.9	2.4	2.1	1.6
China	2.2	2.1	2.2	1.8	1.7	2.2	2.5
India	1.3	1.1	1.3	1.3	1.9	1.5	1.6
Russia	0.0	0.0	0.1	0.0	1.7	1.3	1.9
South Africa	4.5	4.5	4.7	4.1	1.7	1.6	1.9

Table 4: Conditional Expected Shortfall Estimation

The final table here provides a summary of overall model performance. The accuracy row ranks the average position of the models with respect to the goodness of fit score $s(x_{o,t}, x_{\alpha,t})$ previously discussed. The coverage column denotes the number of estimates (out of 15 for both CVaR and CES) in which confidence level was not broken by the foreign exchange returns. It is important to note that while this data can provide a general summary of model performance, general performance, particularly among the GARCH variants was highly dependent on the underlying currency being modeled. We now discuss the results.

	GARCH	EGARCH	GJR	FIGARCH	Historical	Parametric	RiskMetrics
CVaR							
Accuracy	3.2	5.3	4.7	6.3	2.2	3.4	3.1
Coverage	8	8	8	8	3	9	8
CES							
Accuracy	2.9	4.8	4.3	5.7	4.3	3.5	2.8
Coverage	11	10	10	11	15	15	15

Table 5: Summary of CVaR and CES Estimates

6 Discussion

In short, the GARCH model and its variants produce widely differing CVaR and CES results depending on the currency being modeled while the historical simulation, standard parametric, and RiskMetrics models produce more consistent results across the board. For example, the GARCH variants produce far to conservative CVaR estimates for Russia and India while estimating a much to liberal CVaR for Brazil and South Africa. Interestingly, when the in sample GARCH parameters are used to estimate out of sample CVaR and CES, the standard GARCH model tends to perform the best of the four variants, though there is little difference in coverage across the GARCH type models. Because of this though, the standard GARCH model may be preferred as it has the fewest parameters and therefore the least likelihood of over fitting. However, since the GARCH model does not perform significantly better (and often is worse) than the historical simulation, standard parametric, and RiskMetrics models, we cannot conclude it is the best fit. A potential issue with GARCH models not faced by the other three is that the underlying market fundamentals can change so parameters estimated on historical data may not be representative of the volatility process today. Parameter estimation for GARCH models can also be sensitive to the number of observations as well as the starting point. When examining the results between estimation of CVaR and CES, we find CES converge to be higher across the board which is inline with what we would expect as CVaR is a lower bound for CES. Another interesting observation is that risk levels tend to be underestimated for low α levels and overestimated for higher α levels. A possible reason for this is that the data points which break a $\alpha = 5\%$ level also break $\alpha = 1\%$, which is inline with what we observe in the data. For example the CES coverage for Indian foreign exchange is relatively similar for both $\alpha = 1\%$ and $\alpha = 5\%$, and the days that they are broken are nearly identical leading to the close coverage numbers. In summary, there is no definitively best volatility model and one should be diligent when deciding on a risk measure.

7 Conclusion

In this work we introduce a multitude of autoregressive conditional heteroskedasticity models and developed the mathematical formulation behind two common practice risk metrics, Conditional Value at Risk and Conditional Expected Shortfall. With these tools at our disposal, we produced four different volatility models for the foreign exchange rates of the BRICS emerging economies. These volatility models allowed us to analyze their power to estimate both CCVaR and CES and to compare these results against historical simulation, standard parametric, and RiskMetrics models. Among GARCH type models, we found a high sensitivity to the underlying asset in risk estimation, which was less prevalent among the three other models considered. We also found that performance among the GARCH models was quite similar, which makes the standard GARCH model an attractive choice due to its simplicity compared to the others. However, there is no conclusive evidence to their being an overall 'best' model for estimating CVaR and CES, specifically for BRICS foreign exchange, and one should be careful when selecting a risk model.

Future directions for this work could include analyzing a larger basket of currencies and not limiting ourselves to emerging markets, as well as developing a more robust metric to evaluate the overall performance of risk models. Additionally, more complex multivariate volatility models which account for other factors shown to influence exchange rates, such as interest rates and the capital account, could be considered.

References

[Baillie et al., 1996] Baillie, R. T., Bollerslev, T., and Mikkelsen, H. O. (1996). Fractionally integrated generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, 74:3–30.

[Bank of International Settlements, 2024] Bank of International Settlements (2024). Bilateral Exchange Rates.

[Black and Scholes, 1973] Black, F. and Scholes, M. (1973). The Pricing of Options and Corporate Liabilities. *The Journal of Political Economy*, 81(3):637–654.

[Bollerslev, 1986] Bollerslev, T. (1986). GENERALIZED AUTOREGRESSIVE CONDITIONAL HETEROSKEDASTICITY. *Journal of Econometrics*, 31:307–327.

[Cont, 2001] Cont, R. (2001). Empirical properties of asset returns: stylized facts and statistical issues. *Quantitative Finance*, 1:223–236.

[Engle, 1982] Engle, R. F. (1982). AUTOREGRESSIVE CONDITIONAL HETEROSCEDASTICITY WITH ESTIMATES OF THE VARIANCE OF UNITED KINGDOM INFLATION. *Econometrica*, 50(2).

[Engle and Manganelli, 2004] Engle, R. F. and Manganelli, S. (2004). CAViaR: Conditional Autoregressive Value at Risk by Regression Quantiles. *Journel of Business & Economic Statistics*, 22(4):367–381.

[Engle and NG, 1993] Engle, R. F. and NG, V. K. (1993). Measuring and Testing the Impact of News on Volatility. *The Journel of Finance*, 48(5):1749–1778.

[Genaçay and Selçuk, 2004] Genaçay, R. and Selçuk, F. (2004). Extreme value theory and Value-at-Risk: Relative performance in emerging markets. *International Journal of Forecasting*, 20:287–303.

[Glosten et al., 1993] Glosten, L. R., Jagannathan, R., and Runkle, D. E. (1993). On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks. *The Journal of Finance*, 48(5):1779–1801.

[Goldberg, 2006] Goldberg, L. S. (2006). Exchange Rates and Foreign Direct Investment.

[Martin Filho et al., 2018] Martin Filho, C., Yao, F., and Torero, M. (2018). Nonparametric Estimation of Conditional Value-at-Risk and Expected Shortfall based on Extreme Value Theory. *Econometric Theory*, 34:23–67.

[Morgan, 1996] Morgan, J. (1996). RiskMetrics.

[Nelson, 1991] Nelson, D. B. (1991). CONDITIONAL HETEROSKEDASTICITY IN ASSET RETURNS: A NEW APPROACH. *Econometrica*, 59(2):347–370.

[Ozturk, 2006] Ozturk, I. (2006). EXCHANGE RATE VOLATILITY AND TRADE: A LITERATURE SURVEY. *International Journal of Applied Econometrics and Quantitative Studies*, 3(1).

[Sarykalin et al., 2008] Sarykalin, S., Serraino, G., and Uryasev, S. (2008). Value-at-Risk vs. Conditional Value-at-Risk in Risk Management and Optimization. *Tutorials in Operations Research*.

[Sheppard, 2024] Sheppard, K. (2024). ARCH.

[So and Yu, 2006] So, M. K. and Yu, P. L. (2006). Empirical Analysis of GARCH Models in Value at Risk Estimation. *International Financial Markets, Institutions, & Money*, 16:180–197.

[Tsay, 2010] Tsay, R. S. (2010). *Analysis of Financial Time Series*. Third edition edition.

[Vania Stavrakeva and Tang, 2023] Vania Stavrakeva and Tang, J. (2023). Explaining the Great Moderation Exchange Rate Volatility Puzzle.