# Bayesian Extreme Value Analysis of Risk in Selected Foreign Exchange Markets

by

Ziming Wang
and
David E. Giles

Department of Economics, University of Victoria, Canada

*Draft Version - February, 2017*
*(Not to be Quoted Without Permission)*

**Abstract**

Risk management is crucial to financial institutions. It aims to minimize the potential loss or maximize the possible gain from unusual events. Extreme value theory (EVT) is commonly applied to analyze the extreme behavior in financial markets. In this paper, we use EVT to investigate the volatility of the exchange rates between the U.S. dollar and the Euro and Great British Pound, respectively. The peak over threshold (POT) method from EVT uses the result that exceedances over a predetermined threshold follow the generalized Pareto distribution (GPD). A major contribution of this paper is that we use a Bayesian approach to directly estimate the threshold as well as the GPD parameters. Specifically, Markov Chain Monte Carlo (MCMC) simulation with Metropolis-Hastings sampling is used to model the tail distributions of our exchange rate data. The two risk measures, value at risk (VaR) and expected shortfall (ES), are constructed based on the Bayesian estimation results. They both indicate that trading the British pound is less risky than trading the Euro. By comparing our method with the traditional two-step EVT approach, which determines the threshold graphically, we conclude that our method is superior. Also, some sensitivity tests show that the choice of prior distributions for the threshold parameter and the different starting value for MCMC analysis have little impact on the results.

*Key words:* Extreme value theory, Bayesian inference, Markov chain Monte Carlo, Value at risk, Expected shortfall

**Contact Author:** David E. Giles, Department of Economics, University of Victoria, P.O. Box 1700, STN CSC, Victoria, B.C., Canada, V8W 2Y2; email, dgiles@uvic.ca; voice, +1-250-721-8540.

## 1. Introduction.

The management of risk is crucial to financial institutions. It aims to minimize the potential loss or maximize the possible gain from unusual events. After the financial crisis in 2008, the new Basel III Accord (2010) raised the capital requirement to strengthen the risk resistance capacity. It requires bank and financial institutions to set aside a minimum of capital as risk control to cover predicted losses from daily portfolio trading. Thus, those institutions are interested in the tail behavior—extreme events—of their portfolios. The extreme value theory (EVT) is commonly used to estimate the extreme distribution of financial time series.

De Dieu *et al*. (2014) indicate that EVT is a well-developed and significant theory in explaining sample extreme and it provides the opportunity to study tail distributions to predict losses from sudden crashes, which is important to financial institutions. It can help financial institutions to accurately reserve the minimum amount of capital required to cover the maximum losses that will occur with low probabilities. Haile and Pozo (2006) note that EVT is an appropriate and useful approach to investigate and analyze statistical modeling with extreme and unusual cases, and EVT need not only be applied in the financial field, but also can be adopted and applied to other fields that deal with risk management and unusual variations in data. Gilli and Këllezi (2006) discuss the use of two risk measures, Value at Risk (VaR) and Expected Shortfall (ES) under the Generalized Pareto Distribution (GDP), and estimate the tail behavior of stock indices using EVT. They suggest that EVT is a well-founded method to build up the model and determine the distribution of risk or unusual events.

The VaR is interpreted as the maximum value of the loss or gain of a portfolio given a pre-assigned very small probability over a certain period, and the ES as the expected value of the loss or gain that exceeds the VaR. Fretheim and Kristiansen (2015) point out that the benefit of using EVT is that investigators can focus only on the tail of the series, and EVT tolerates non-normality, which allows for skewed data that may also exhibit excess kurtosis. De Jesus *et al*. (2013) investigate the Mexican peso and U.S. dollar exchange rate, and they find that the EVT is a useful methodology which can identify the distribution of the extreme values of such data. In addition, Iglesias (2012) investigates the extreme movements of exchanges rates for several major currencies.

Since EVT is a commonly selected method to apply in risk management in the financial field, it is a good opportunity for financial institutions to apply the method and analyze its results in trading portfolio selection. In this paper, we apply the EVT to compare the risk behavior of two major currencies, the Euro and Great British Pound (against the U.S. dollar), to determine which of the two is less volatile—in other words, which currency is less risky to invest in. The two risk measures, VaR and ES are calculated to indicate the potential loss or gain for exchange returns under in extreme situations.

McNeil an d Frey (2000) explore a method of combining GARCH modeling and EVT with maximum likelihood estimation to investigate the volatility of financial return series over a short period of time. They suggest a two-step EVT method which can capture two main characteristics of financial data: the fat-tailed distribution and the volatility clustering. However, the crucial part of the modelling procedure, namely the choice of the threshold itself, is based on two graphs which are difficult to interpret in practice. This graphical approach to determining the location of the extreme region of the data is standard in the vast majority of the literature on EVT. It remains a weakness of much of the associated empirical literature.

In contrast, Behrens *et al.* (2004) suggest an alternative method which involves estimating the threshold (that marks the start of the extreme range in the data) directly, using a Bayesian approach. The Markov Chain Monte Carlo (MCMC) simulation with Metropolis-Hastings sampling is introduced to model the tail distribution of financial time series. Following their ideas, we apply a modified Bayesian approach to evaluate the risk associated with daily changes in exchange rate data. Details of the model are discussed in the following section. We also apply the two-step EVT method based on Maximum Likelihood estimation for comparison purpose.

The remainder of this paper is structured as follows: Section 2 introduces the background of the extreme value theory as well as the Bayesian analysis. Section 3 provides a detailed description of our methodology. Section 4 discusses the properties of the daily returns for the exchange rates and provides descriptive statistics as well as some preliminary tests. Section 5 summarizes our findings, and also presents a comparison of results from different methods. Section 6 discusses the sensitivity tests that we have conducted. Section 7 concludes the paper, and offers some suggestions for future research.

## 2. Background

### 2.1 Extreme Value Theory

In preparation for the potential loss during the extreme events, financial institutions are interested in estimating the tail distribution of the returns for an investment portfolio. Three large families of estimation methods are commonly applied: the nonparametric simulation approach, the parametric ARCH/GARCH estimation, and the method based on extreme value theory (EVT). However, each of those three techniques has drawbacks. It is well-known that to estimate extreme situations beyond the historical data period is difficult; and the extreme value approach within the sample period is inefficient. The estimator from the simulation approach, based on empirical data, has high variance. As for the ARCH/GARCH approach, the assumption that the data are normally distributed is not likely to hold for financial time series, which usually have fatter tails than normal data. The extreme value theory is designed to describe the asymptotic distribution—especially the tail part—of a random variable under unusual situations. Although this method provides a parametric estimation of the tail distribution, it cannot capture the stochastic volatility commonly exhibited by financial data.

To overcome these shortfalls, McNeil and Frey (2000) suggest a two-step EVT method that combines the ideas of all three approaches. They first apply a GARCH filter to the data, and then historical simulation and EVT are used to estimate the distribution of the residuals. McNeil and Frey (2000) state that their approach captures the two main properties of financial return series: the stochastic volatility and the fat-tailed distribution. The estimation of the tail distribution is based on a predetermined threshold which is chosen graphically. However, the associated plots are often hard to interpret. Thus, the choice of the threshold is arbitrary. Nonetheless, the basic ideas of modeling extreme events with EVT and the GARCH filter for financial data are adopted in this paper.

### 2.2 Block Maxima *vs*. Peak Over Threshold

There are two methodologies that are commonly used in EVT to model and estimate the distributions of extrema: the block maxima method, and the peak over threshold (POT) approach. Relying on the

Fisher-Tippett-Gnedenko theorem (Fisher and Tippett, 1928; Gnedenko, 1943), the block maxima method uses the result that the extreme values for an i.i.d. random variable *must follow* one of the Generalized Extreme Value (GEV) distributions. The data are divided into several equal length blocks and the extreme values within each block are used to fit one of the Fréchet, Gumbel, or Weibull distributions. In contrast, the POT method studies the data values beyond a predetermined high threshold value, $u$. It can be shown that, if the threshold $u$ is large enough, the exceedances above $u$ will follow the Generalized Pareto distribution (GPD) (Pinkands, 1975; Balkema and de Haan, 1974).

Many studies in the literature (*e.g.*, Ren and Giles, 2010; Chen and Giles, 2016) indicate that there is a major drawback with the block maxima method. Instead of the whole data set, it only considers the maximum values in each blocks, which may lead the estimation to be less efficient due to the different features in each block. In particular, some blocks might have smaller maximum values than others. The distribution estimation would be biased when the data contains those outliers. To avoid such undesirable situations, we have chosen to apply the POT approach, because it takes all of the individual sample values in the tail distribution into account. In addition, the POT approach has been used quite widely and successfully in analyzing other high frequency financial time series.

## 2.3 The Generalized Pareto Distribution.

The POT approach uses the fact that the data "exceedances" above a sufficiently high threshold follow the GPD. The three-parameter GPD density function for a random variable, $X$, is:

$$G_{\xi,\beta,u}(x) = \begin{cases} 1-\left(1+\xi\left(\frac{x-u}{\beta}\right)\right)^{-\frac{1}{\xi}} \; if\ \xi \neq 0 \\ 1-\exp\left(-\frac{x-u}{\beta}\right) \quad if\ \xi = 0 \end{cases} \tag{2.1}$$

where $\xi$ and $\beta$ are the shape and scale parameters, while as noted before, $u$ is the predetermined threshold. It is effectively the location parameter for the distribution. The shape parameter determines the form of the tail of the distribution. If $\xi > 0$, the distribution has a fat tail compared with the Normal distribution. Gilli and Këllezi (2006) point out that this situation is likely to apply to "returns" data for financial assets. A negative value of $\xi$ implies that the distribution has a thinner tail than the Normal distribution. If $\xi = 0$, the tail of the distribution dies down exponentially (McNeil and Frey 2000).

The choice of the threshold is crucial. It determines the sample of “exceedances” that are then used to estimate the shape and scale parameters of the underlying distribution, so it also affects the corresponding estimates of VaR and ES because these are functions of these parameters. As the result, different choices of the threshold can influence decisions relating to risk capital reserves. Moreover, implied by the EVT, the exceedances would only converge to the GPD if the threshold is large enough. Thus, to ensure the accuracy of the estimation, it is important to choose a proper threshold. Choosing the threshold too low will result in a distribution that is not generalized Pareto, and the parameter estimates will be beiased. However, a very high threshold will result in very few exceedances. The small number of observations will then result in high estimation variance. The choice of the threshold is a trade-off between the bias and the variance of the parameter estimator (Gilli and Këllezi, 2006; Ren and Giles, 2010).

In general, the threshold is determined using the assistance of the Mean Excess (ME) plot, and the so-called parameter plot. The ME function is defined as the expected value of the exceedances over a certain threshold. For an underlying GPD, the (theoretical) ME function is linear, and upward-sloping. Thus, the threshold is selected where the (empirical) ME graph starts to exhibit this linearity. The parameter plot illustrates the stability of the estimated shape and scale parameters under the different choice of the threshold. The two GPD parameters should be stable above a certain high threshold, as the exceedances should approach the GPD asymptotically. Thus, the threshold would be chosen where the two parameters begin to vary noticeably.

Behrens *et al.* (2004) argue that although the POT method considers the full information above the threshold, the data below the threshold are still ignored. In addition, it is difficult to interpret the ME and parameter plots in practice. A more effective method of choosing the threshold is required.

## 2.4 Bayesian Estimation of the Threshold.

Behrens *et al*. (2004) propose an alternative way to determine the value of the threshold. Instead of locating the threshold with the help of two graphs, they treat the threshold as an unknown parameter of the underlying parametric model. Therefore, its value can be estimated directly. More specifically, the whole data set is fitted to a designated distribution, which (in their case) combines the gamma distribution below the threshold and the GPD afterward. They construct the prior distributions for each parameter in the model, and the Markov Chain Monte Carlo (MCMC) method with Metropolis-Hastings sampling is used to draw the posterior inferences.

Based on this general approach, we apply a modified model to estimate the threshold on the exchange return data. Instead of the gamma distribution which is used by Behrens *et al.*, we assume that the data are normally distributed below the threshold. That is, a truncated normal distribution is used. The full details of the model are discussed in later sections of this paper.

### 2.5 Estimation of VaR and ES

It will be recalled that the VaR measures the potential loss or gain for a portfolio under some extreme event. Santamaría *et al.* (2016) point out that the estimation methods for VaR would achieve their best performances at the high quantiles, such as the 99th percentile. So, we calculate the VaR at the 1% probability level over a one-day period. Mathematically, the VaR can be obtained by inverting the tail estimator of an underlying distribution given a probability level $p$. The tail estimator of the GPD is given as follows (McNeil and Frey, 2000):

$$\hat{F}(x) = 1 - \frac{Nu}{n}(1 + \hat{\xi}\frac{x-u}{\hat{\beta}})^{-\frac{1}{\hat{\xi}}} \quad , \tag{2.2}$$

where $N_u$ is the number of observations above a given threshold, $u$, and $n$ is the total number of observations in the original sample of data. By inverting the tail estimator for a given probability level $p$, we obtain the estimator for the VaR as:

$$\widehat{VaR_p} = u + \frac{\hat{\beta}}{\hat{\xi}}((\frac{n}{Nu}p)^{-\hat{\xi}} - 1) \quad . \tag{2.3}$$

ES provides the information above the VaR. In particular, it reports the conditional expected value above the VaR given a probability level $p$. The ES is defined as:

$$ES_p = E(x | x > VaR_p) \quad . \tag{2.4}$$

Equation (2.4) can be expressed as:

$$ES_p = VaR_p + E(x - VaR_p | x > VaR_p) \tag{2.5}$$

The $E(x - VaR_p | x > VaR_p)$ term in equation (2.5) is defined as the mean excess function of exceedances beyond the threshold: $\text{VaR}_p$. From Ren and Giles (2010), and Chen and Giles (2016), the mean excess for an underlying GPD is:

$$E(u) = \frac{\beta + \xi u}{1 - \xi} \quad . \tag{2.6}$$

From the Pickands-Balkema-de Haan results, if the value of threshold is large enough, the mean excess function can be constructed on the basis of equation (2.6) replacing the threshold with ($VaR_p - u$). (Ren and Giles, 2010; Chen and Giles, 2016):

$$\frac{\beta + \xi(VaR_p - u)}{1 - \xi}$$

Substituting this expression into equation (2.5), the formula for ES can be written as:

$$ES_p = VaR_p + \frac{\beta + \xi(VaR_p - u)}{1 - \xi} \tag{2.7}$$

Simplifying equation (2.7), we obtain the estimator of the expected shortfall as:

$$\widehat{ES}_p = \frac{\widehat{VaR}_p + \hat{\beta} - \hat{\xi}u}{1 - \hat{\xi}} \tag{2.8}$$

By substituting the proper threshold $u$ and the estimated shape and scale parameters for the GPD into the equation (2.3) and (2.8), we can acquire point estimates of the risk measurements: VaR and ES.

Interval estimators for the VaR and the ES are also reported in a handful of papers (Ren and Giles, 2010, and Chen and Giles 2016). These estimators are formed by applying the delta method (Oehlert, 1992) to approximate the asymptotic variances of those risk measurements. However, the delta method requires the information about the covariance matrix of the GPD parameter estimators, and this is not readily available from our Bayesian analysis. Thus, at this stage we report only point estimates of the VaR and the ES below.

## 3. Methodology

### 3.1 Model

Recall from the last section that a modified version of Behrens *et al.* (2004) method is applied in this paper. The distribution for our data combines the (truncated) normal distribution and the GPD. We have a random variable $X$, which is assumed to be normally distributed with the distribution function $H(x \mid \mu, s^2)$ below the threshold, $u$. Let $h(x \mid \mu, s^2)$ denote the corresponding density function. On the other hand, the exceedances above the threshold follow the GPD, denoted as $G$ ($x$ | $\xi$, $\beta$, $u$). Hence, the distribution function for $X$ can be constructed as the following mixture:

$$F(x|\mu,s^2,\xi,\beta,u)=\begin{cases} H(x|\mu,s^2) & x<u \\ H(x|\mu,s^2)\mathbf{1}_{(x<u)}+[1-H(u|\mu,s^2)]\mathbf{1}_{(x\geq u)}G(x|\xi,\beta,u) & x\geq u \end{cases} \quad (3.1)$$

where $\mu$ and $s^2$ are the population mean and the variance below the threshold. The **1** is an indicator function which takes the value of 1 when $X$ is within the pre-specified range, and 0 elsewhere. From equation (3.1) above it is obvious that the threshold, $u$, is the discrete transition point between two distributions. Behrens *et al.* (2004) argue that the estimation for the threshold would be more difficult if the transition between the distributions is modelled as being a smooth one.

Assuming independent sampling, the corresponding likelihood function can be written as:

$$L(x|\mu,s^2,\xi,\beta,u)=\begin{cases} \prod_{(i,x_i<u)} h(x_i|\mu,s^2)\prod_{(i,x_i\geq u)}\big(1-H(x_i|\xi,\beta,u)\big)\left(\frac{1}{\beta}\left(1+\frac{\xi(x_i-u)}{\beta}\right)^{-\frac{(1+\xi)}{\xi}}\right) & \xi\neq 0 \\ \prod_{(i,x_i<u)} h(x_i|\mu,s^2)\prod_{(i,x_i\geq u)}\big(1-H(x_i|\xi,\beta,u)\big)\left(\frac{1}{\beta}\exp\left(\frac{x_i-u}{\beta}\right)\right) & \xi=0 \end{cases}$$

(3.2)

## 3.2 Parameter Priors

There are five parameters in the likelihood function, namely the mean, $\mu_u$, and the variance, $s_u^2$, from the Normal distribution, the GPD shape, $\xi$ and scale, $\beta$ parameters, and the threshold $u$. Instead of setting up a prior for each parameter, we consider prior distributions only for the GPD parameters and the threshold, $u$. Values are assigned directly to the parameters for the distribution below the threshold, in order to simplify matters. This part of the analysis is being extended in work currently in progress. In the current analysis the mean of the Normal distribution below the threshold is set to a suitably high quantile of the data. From the results of some sensitivity testing, we have found that the choice of the mean and the variance from the truncated Normal distribution has very little impact on the estimation of the threshold itself.

As for the GPD parameters, we use Jeffreys' "invariant" prior which is derived for this distribution by Castellanos and Cabras (2007). It is constructed from the information matrix for the shape and the scale parameters. Specifically, it takes the form:

$$\pi(\xi,\beta)\propto\beta^{-1}(1+\xi)^{-1}(1+2\xi)^{-\frac{1}{2}},\ \xi>-0.5,\beta>0 \quad . \quad (3.3)$$

There are several ways to model the prior distribution of the threshold. Following the steps of Behrens *et al.* (2004), we assume that the prior distribution of the threshold, $u$, is a truncated normal distribution:

$$\pi(u|\mu_u, s_u^2, t) \propto \frac{1}{\sqrt{2\pi s_u^2}} \frac{\exp(-\frac{1}{2}\frac{(u-\mu_u)^2}{s_u^2})}{\Phi(\frac{-(t-\mu_u)}{s_u})} \quad , \tag{3.4}$$

where $t$ is the truncation point. The $\mu_u$ and $s_u^2$ are the mean and the variance of the normal distribution of the threshold. As suggested by Behrens *et al.* (2004), the value of $\mu_u$ is assigned using the 75$^{\text{th}}$ percentile of the data. Choosing a relatively large value for $s_u^2$ ensures that the prior distribution is relatively flat. In our case, after conducting a sensitivity test, different choices of the $\mu_u$ and $s_u^2$ are found to have low impact on the Bayes estimate of the threshold. Bermudez *et al.* (2001) propose an alternative method to describe the threshold. Instead of setting the prior with respect to $u$, they construct a discrete prior for the upper order statistics of the data. The threshold is determined indirectly based on the values of the order statistics. In our own study we consider two possible prior distributions for the threshold, and discuss the possible impact of the choice of the prior on the threshold estimate as well as the estimates of the risk measures in the later sections of this paper.

## 3.3 Posterior distribution

By combining the likelihood function (3.2) and the priors of the parameters (3.3), and (3.4), we obtain the posterior distribution function in logarithmic form as follows:

$$\begin{aligned}
&\log P(x|\xi, \beta, u, \mu_u, s_u) \\
&= C + \sum_{i=1}^{n} I_{(x_i<u)} \left(-\log \mu_u - \left(\frac{\frac{1}{2}(x_i-\mu_u)^2}{s_u^2}\right)\right. \\
&+ \sum_{i=1}^{n} I_{(x_i\geq u)} \log\left(1 - \left(\frac{1}{2}\left(1 + \operatorname{erf}\left(\frac{u-\mu_u}{s_u\sqrt{2}}\right)\right)\right)\right) + \sum_{i=1}^{n} I_{(x_i\geq u)} (-\log \beta) \\
&+ \sum_{i=1}^{n} I_{(x_i\geq u)} \left(-\frac{1+\xi}{\xi}\log\left(1 + \frac{\xi(x_i-\mu_u)}{\beta}\right)\right) - \log(\beta) - \log(1+\xi) \\
&- \frac{1}{2}\log(1+2\xi) - \frac{1}{2}(\frac{u-\mu_u}{s_u})^2
\end{aligned} \tag{3.5}$$

where $C$ is a constant term, and erf ( . ) is the so-called error function. The MCMC method with Metropolis-Hastings sampling is applied to evaluate the posterior distribution using the R package *MHadaptive* (Chivers, 2015). Starting values of each parameter are required by the simulation process. The starting value of the threshold is chosen first third quartile of the sample. Once the hyper-threshold is determined, $\mu_u$ and $s_u^2$ are calculated as the mean and the variance of the data before the hyper-threshold. As for the GPD parameters, the starting values are based on the Maximum Likelihood estimates given the starting value of the threshold. The results are detailed in the later sections.

## 4. Data Characteristics

### 4.1 Overview

In this paper we focus on the daily "returns" for two exchange rates: EUR/USD, and GBP/USD. By considering the potential structural change during the 2008 financial crisis, we restrict our analysis to the period from 2nd January 2009 to 13th November 2015 to avoid the possible impact of the recession on the exchange market. The exchange rates are retrieved from the PACIFIC Exchange Rate Service, provided by Antweiler (2015). Daily returns are obtained by taking the log-difference of the official noon spot rates.

Table 1. Descriptive Statistics

| | EUR/USD | GBP/USD |
|---|---|---|
| Mean | 0.013346 | -0.002342 |
| Median | 0.004056 | -0.000766 |
| Maximum | 2.677503 | 3.468275 |
| Minimum | -4.606501 | -4.301831 |
| Std. Dev. | 0.658011 | 0.584191 |
| Observations | 1970 | 1970 |

*Note:* All statistics are in %.

Table 1 provides the descriptive statistics of the data. The mean and the median of the return series are essentially zero. Comparing the two exchange returns, the Euro returns are more variable than the

British pound's returns, having a larger range and standard deviation. This may suggest that trading the Euro dollar is riskier than trading the Great British Pound. Thus, we might anticipate that larger VaR and ES values may be associated with the EUR/USD returns than with the GBP/USD returns. We return to this pint later. There are 1,970 observations in total for each time-series.

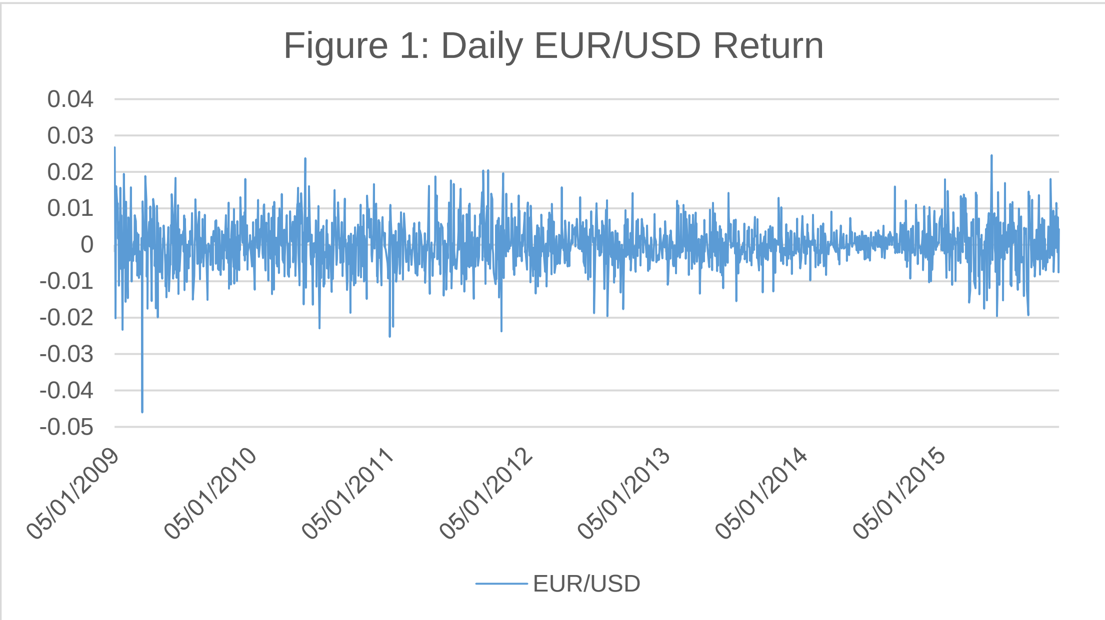


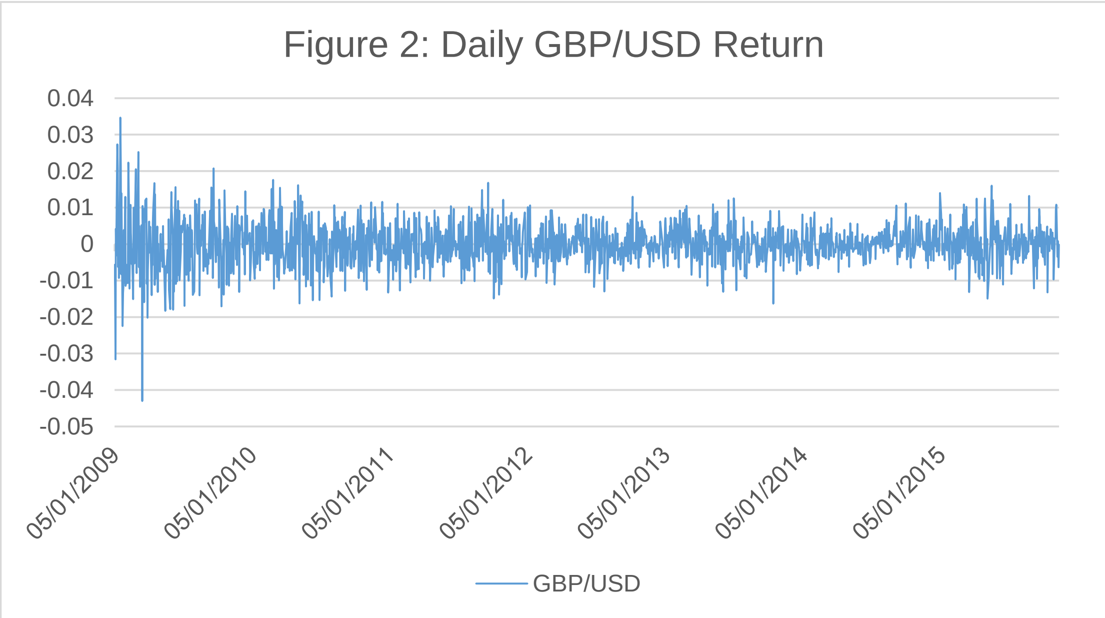


Figures 1 and 2 depict the daily returns for the EUR/USD and GBP/USD returns respectively. From these graphs, the relatively higher volatility at the beginning of 2009 may suggest the possible presence of conditional heteroskedasticity within the samples. However, the results that we obtain after the GARCH filtering, discussed by McNeil and Frey (2000), are almost identical to those based on the original data series. This indicates that the GARCH effect is not significant in our samples. Nevertheless, we use the data after the GARCH filtering for the rest of this paper.

## 4.2 Preliminary tests

The normality and the stationarity of the sample data are examined through the Jarque–Bera (J-B) test and the Augmented Dickey-Fuller (ADF) test, with no drift and no trend. The test results are shown in Table 2. The (essentially) zero p-value of the J-B test clearly leads us to reject the null hypothesis that the data series are normally distributed. Further, the positive excess kurtosis confirms our previous expectation about the fat tail property of the financial time series. The two daily returns series are stationary in their levels based on the ADF test results.

Table 2. Preliminary Test Results

| | EUR/USD Returns | GBP/USD Returns |
|---|---|---|
| **Normality Test** | | |
| Skewness | -0.131233 | -0.05131 |
| Excess Kurtosis | 2.129536 | 3.218265 |
| J-B Test | 377.8962 | 851.0206 |
| (p-value) | (0.0000) | (0.0000) |
| **Stationarity Test** | | |
| ADF | -44.52135 | -43.8718 |
| (p-value) | (0.0001) | (0.0001) |

Due to the asymmetry and non-negative properties of the GPD, we separated the data into two parts: the positive returns and the absolute value of the negative returns. For the daily returns of the EUR/USD, there are 969 positive returns and 1,001 negative returns. On the other hand, the GBP/USD sample contains 976 positive returns and 994 negative returns. For more intuitive results, the daily returns are scaled by a factor of 100.

Quantile-quantile (Q-Q) plots were used to examine whether the actual data distribution is consistent with the theoretical distribution which we selected. In particular, the first part of the data should follow the Normal distribution but truncated at the threshold. The R package, *truncgof*, developed by Wolter (2015), facilitates the construction of these plots for truncated data. However, the package is available only for left truncated data, which is the opposite of our right truncated samples. Thus, we modified the data sets by recording the negative value first, then fitting against the theoretical distributions. As a result, the left end of the Q-Q plot actually relates to the right tail of the empirical data distribution, and *vice versa*.

Figure 3 compares each daily returns series with the theoretical truncated normal distribution (with estimated mean and variance). From this we can conclude that all of the samples exhibit the same property. The empirical data are mostly consistent with the theoretical distributions from the left end (right side) of the graph (actual data). However, the plots deviate away from the 45° line and curve up to the right. This might suggest that the actual daily returns come from a different distribution family than the truncated normal, which is inconsistent with our assumption.

Figure 3: Quantile-Quantile Plot (Truncated Normal)

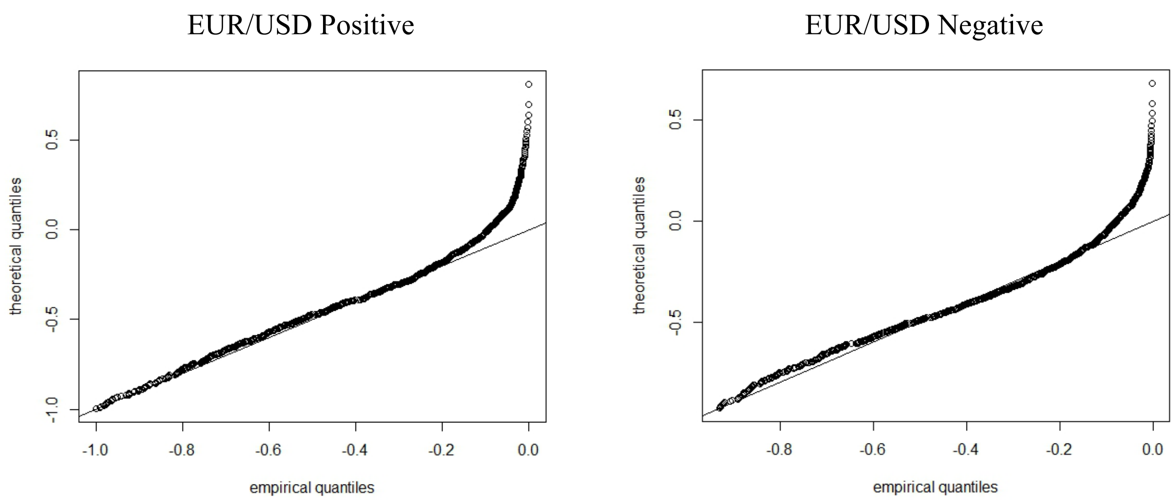

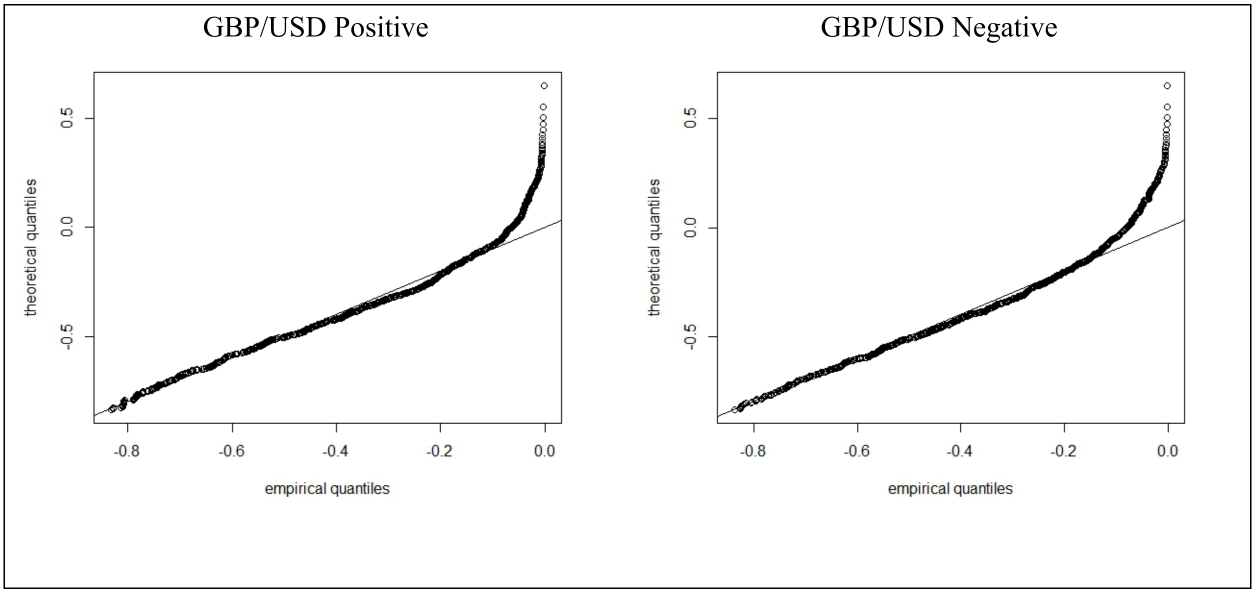


However, in addition to analyzing the preceding distribution graphically, formal hypothesis tests are also performed. Chernobai *et al.* (2005) have developed the proper method to construct the appropriate test statistics for traditional goodness-to-fit tests, such as the Kuiper test and the Cramér-von Mises test in the context of truncated data sets. Based on their work we apply both tests to the exchange returns. Table 3 presents the test statistics and the corresponding p-values of the two goodness-of-fit tests. At the 5% significance level, we cannot reject the null hypothesis that the empirical exchange returns are consistent with the theoretical truncated normal distribution, using at least one of the tests. In other words, the test results provide reasonable support for the assumption of a (truncated) normal distribution for the data below the threshold.

Table 3. Goodness-of-Fit Tests

| | EUR/USD Returns | | GBP/USD Returns | |
|---|---|---|---|---|
| | *Positive* | *Negative* | *Positive* | *Negative* |
| Kuiper | 5.5519 | 5.5895 | 5.8893 | 5.6384 |
| (p-value) | (0.06) | (0.02) | (0.06) | (0.06) |
| Cramér-von Mises | 3.1513 | 2.028 | 1.8116 | 2.0571 |
| (p-value) | (0.08) | (0.08) | (0.02) | (0.02) |

## 5. Results

### 5.1 Threshold and Parameters Estimation

As we have seen already, an accurate threshold is a prerequisite for constructing the risk measurements. The ME plot and the parameter plot are adopted to find the most fitted threshold in the traditional method. As mentioned in the previous section, the mean excess function of the GPD is linear with a positive slope if the data follow the GPD. Ideally, the threshold is chosen where the ME plot is positive sloping while the parameter plot begins to become unstable. From Figure 4, the ME plot of the EUR/USD negative returns is downward sloping at the beginning, but starts to fluctuate and become positive sloping. However, it is difficult to locate the turning point. This exposes the major drawback of the traditional method. Further, the parameter plot only provides an approximate range for the threshold values. A truly reliable point "estimate" of the threshold cannot be obtained directly from these plots.

Figure 4: ME Plot and Parameter Plot of EUR/USD Negative Returns

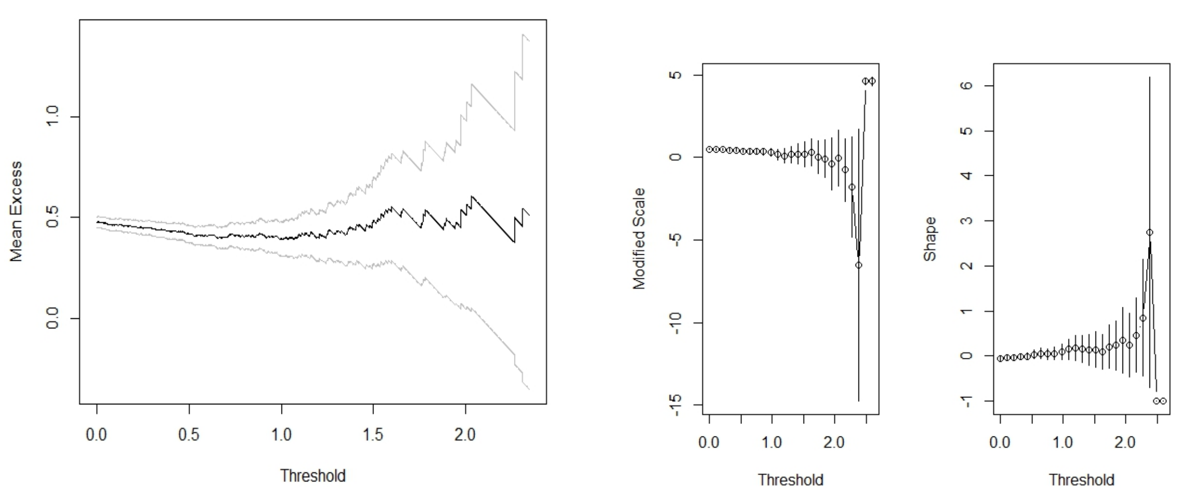


In contrast, the Bayesian approach that we use can overcome such problems. The threshold is treated as an unknown parameter in the model, and is estimated directly, together with the other parameters. In particular, equation (3.5) is applied to describe the posterior distribution. Inferences are drawn *via* MCMC with Metropolis-Hastings sampling.

Figure 5: Marginal Posterior Densities (GBP/USD Negative Returns)

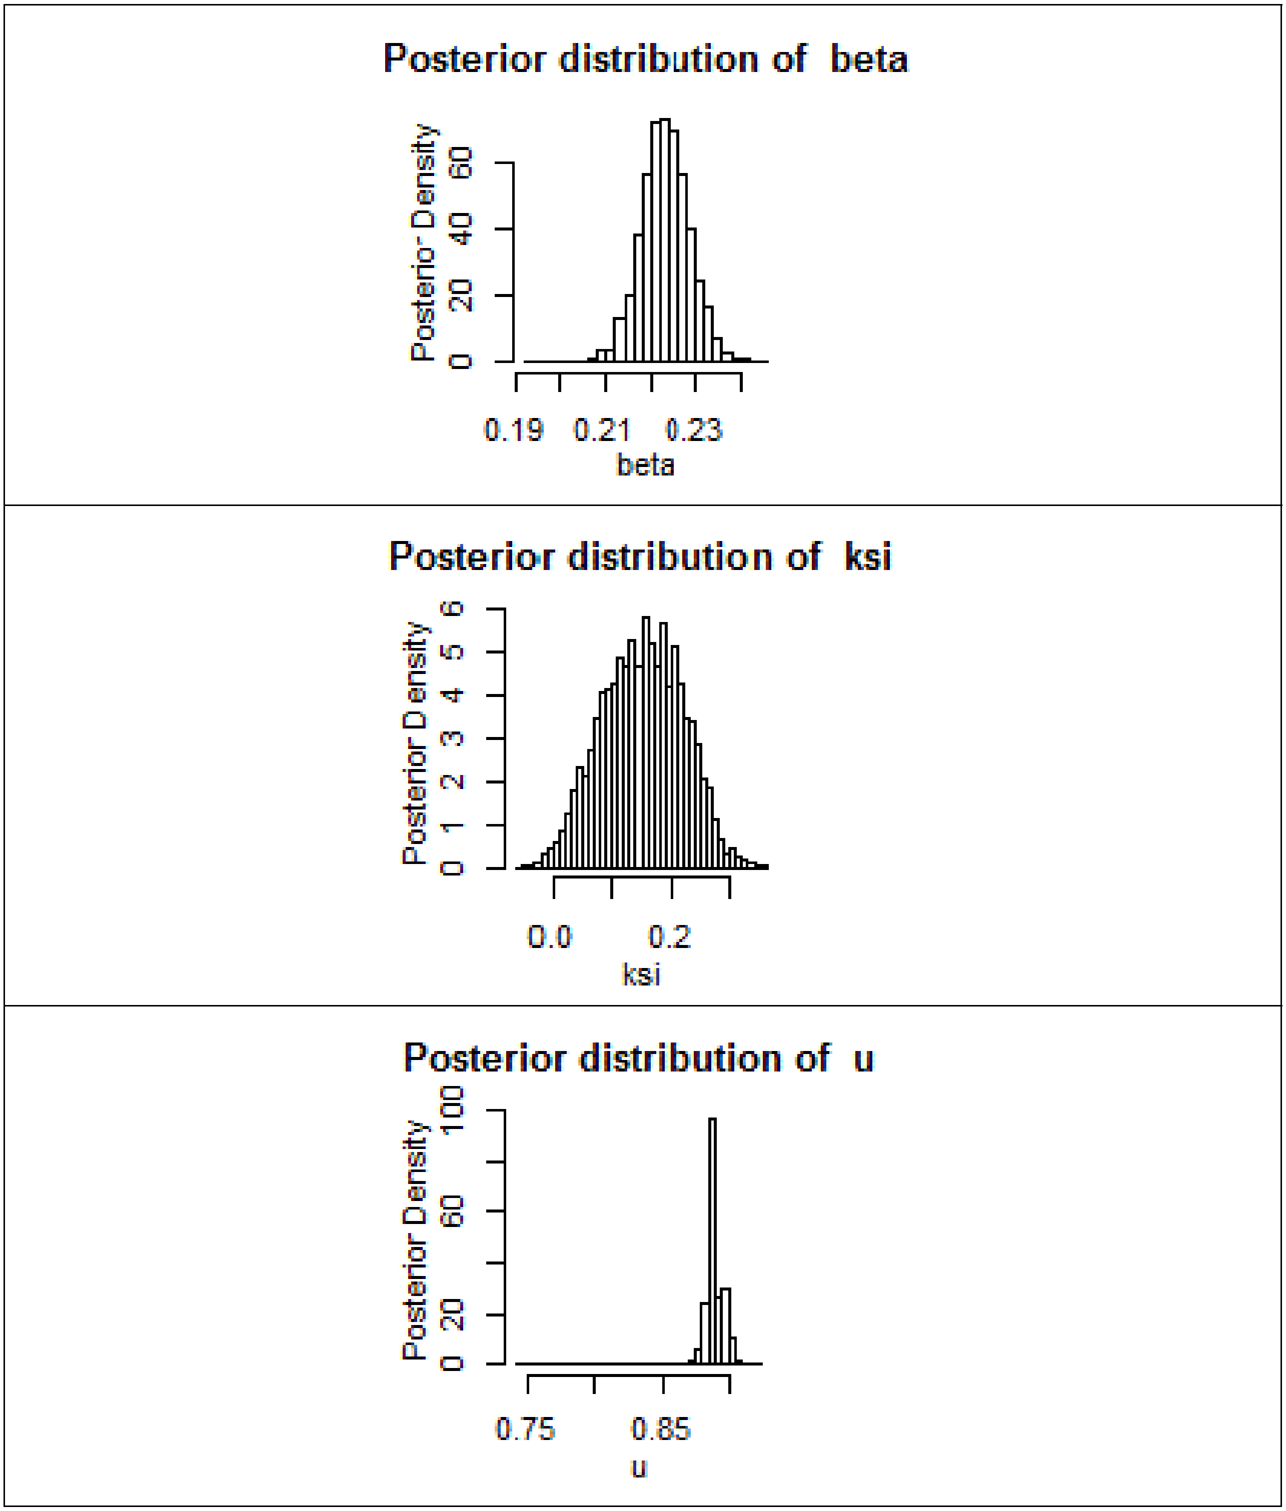

The marginal posterior densities for the case of the GBP/USD negative returns are shown in Figure 5 by way of illustration, and these are fully representative of their counterparts for the other data-sets. Table 4 summarizes the Bayes estimates for the GPD parameters and the corresponding 95% credible intervals. The point estimates are the means of the marginal posterior densities, so we are using a quadratic loss function.

Table 4. Bayes Estimates

| | EUR/USD Returns | | GBP/USD Returns | |
|---|---|---|---|---|
| | *Positive* | *Negative* | *Positive* | *Negative* |
| Threshold ($u$) | 0.9949 | 0.9039 | 0.8412 | 0.8878 |
| (C.I) | (0.9129 1.0516) | (0.8234 0.9891) | (0.7785 0.9212) | (0.8705 0.9022) |
| No. above | 131 | 142 | 132 | 125 |
| Scale ($\beta$) | 0.4286 | 0.3722 | 0.2875 | 0.2253 |
| (C.I) | (0.313 0.5553) | (0.2933 0.4577) | (0.2271 0.3644) | (0.1090 0.3264) |
| Shape ($\xi$) | -0.1301 | 0.0962 | 0.1555 | 0.1473 |
| (C.I) | (-0.2736 0.0551) | (-0.0350 0.2685) | (-0.0137 0.3029) | (0.0201 0.2699) |

*Notes:* C.I stands for the 95% credible interval; all results are simulated using a burn-in of 1,000 iterations and 50,000 further iterations.

The thresholds of the EUR/USD exchange rate are around 0.9 for both positive and negative returns. The GBP/USD returns have the lower thresholds which are slightly lower. Recall that from the fundamental theory of the POT method, the threshold should be sufficiently large enough to ensure the asymptotic distribution of the extreme values is GPD. Meanwhile, care must be taken to select a threshold value which provides enough observations above it to keep the variance of the parameter estimators from becoming too high. From Table 4, the number of observations above the thresholds is between 125 and 145, which is around 12.5% to 14% of the total sample size. While this may raise concerns about estimation bias when applying the POT method, Ren and Giles (2010) suggest that treating no more than 20% of the population as exceedances is reasonable in practice. All of the shape parameter estimates are positive for the two currencies, except for the EUR/USD positive returns. As is discussed in section 2.2, the underlying distribution has a fat tail if the shape parameter is greater than zero, otherwise the distribution has the thinner tail than the normal distribution. The negative shape parameter of EUR/USD positive returns suggests that the underlying distribution does not exhibit fat tails. This conflicts with the usual fat tail characteristic of financial time series data, and may require further investigation.

Figure 6 shows the trace plots for the Markov chains associated with the marginal posterior densities for the various GPD parameters, and exhibit good “mixing”. For both EUR/USD positive and negative returns, the estimates of the thresholds as well as the shape and scale parameters are centered around the posterior mean and are stable within the 95% credible interval. For the GBP/USD positive returns, the traces of the estimated parameters deviate at the beginning of the simulations. However, the estimates tend to be stable around the mean after about 1,000 iterations. The trace of the estimated threshold for the GBP/USD negative returns suggests that the estimates of the thresholds seem to contain more outliers than other samples. But for the most part of estimates are stable within the 95% credible interval.

Figure 6: Trace Plots for Markov Chains

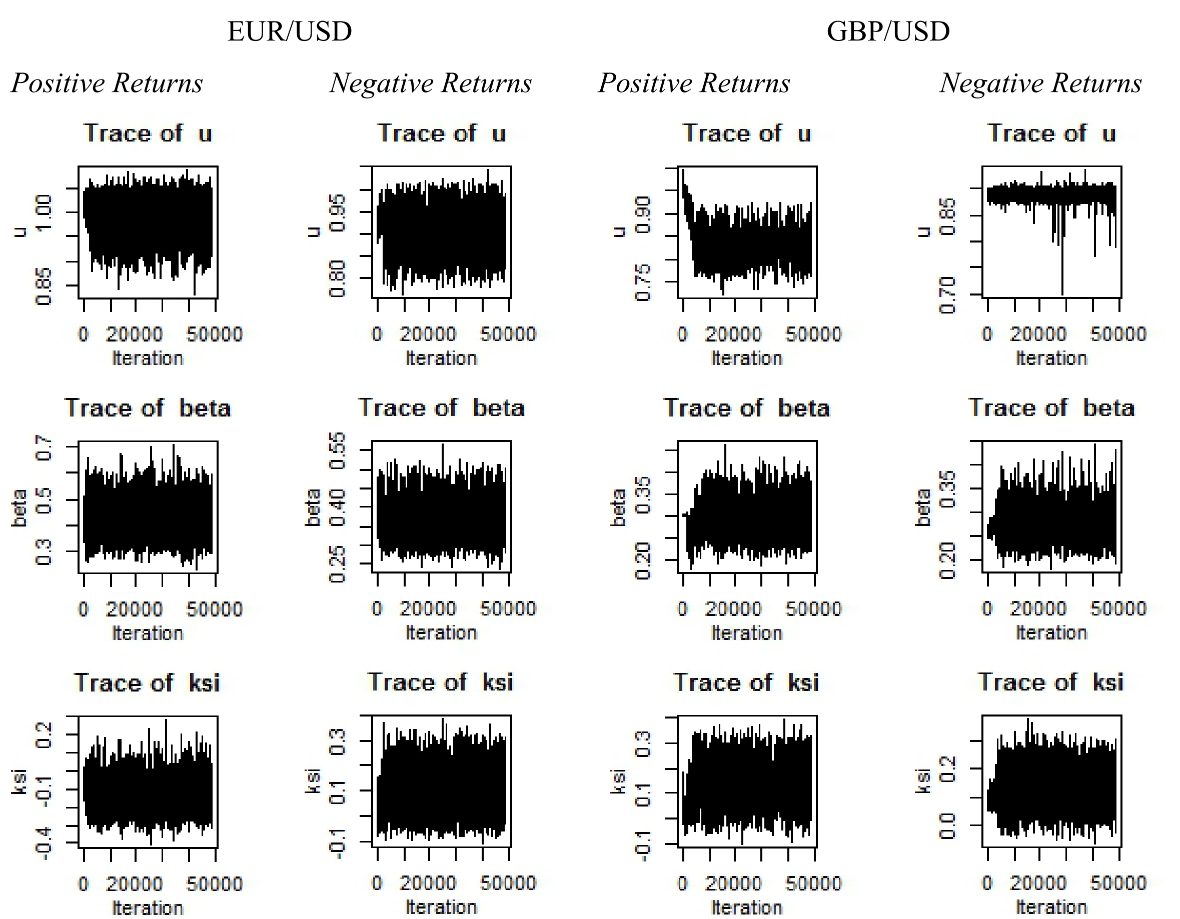

Figure 7: Quantile-Quantile Plots (GPD)

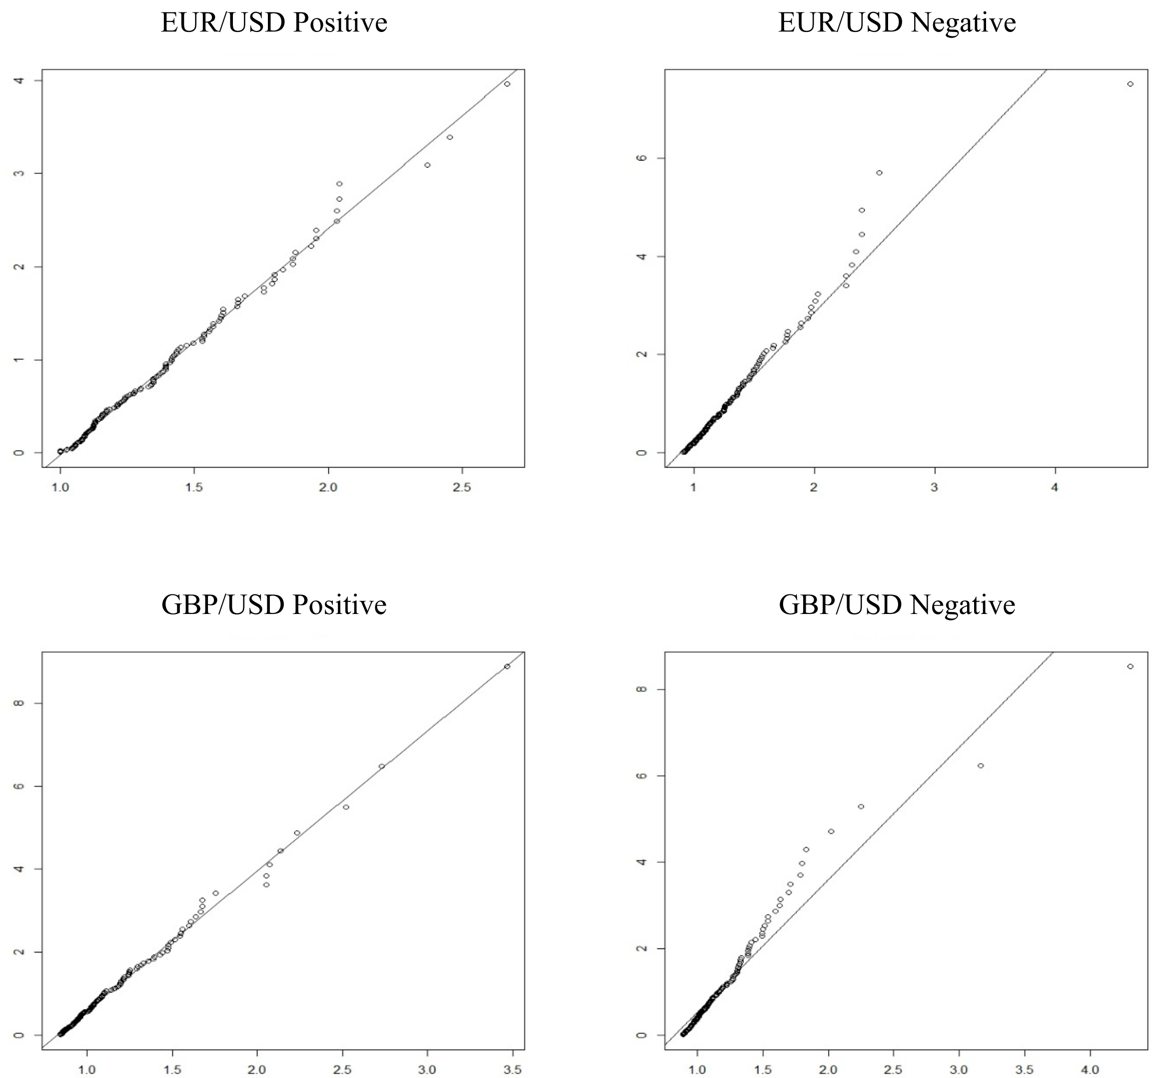


Figure 7 presents the Q-Q plots for the tail distributions (horizontal) of each returns series against the theoretical GPD (vertical). The theoretical distribution is determined by the threshold and the estimated shape parameters through the R package ‘evir’ (Pfaff, McNeil and Stephenson, 2015). Although there are departures from the 45° line at the right ends of each returns distribution, the major part of the data fits the theoretical GPD well. The right ends of the plots are where the outliers are located, and it is not unusual that the outliers do not fit the corresponding distribution. Recalling the negative shape parameter of the EUR/USD positive returns, it is inconsistent with the fat tail assumption of the financial data and brings the concern about the reliability of the POT method under such a situation. However, the Q-Q plot of these returns indicates that the GPD models the exceedances quite well. This, lends credence to the estimation results associated with the EUR/USD positive returns data. Compared with other three exchange returns, the GBP/USD negative returns exhibit the most variation from the theoretical GPD. This explains why the trace of the threshold estimates of the GBP/USD negative returns are less stable than the others to some extent. Overall, the GPD is supported well by these results as a description of the tail distributions of our daily returns.

For comparison, we also used the graphical method to locate the threshold and the two-step EVT method, suggested by McNeil and Frey (2000), to estimate the GPD parameters *via* maximum likelihood (ML). Chen and Giles (2016) point out that, for the graphical method, the determinants of the thresholds should not be based only on the ME and the parameter plots, but one should also consider the significance of ML estimates as well as the Akaike's Information Criterion (AIC) value. Thus, the thresholds are chosen to minimize the AIC value and the asymptotic standard errors of the corresponding ML estimates. The results are shown in Table 5. Conditional on the selected thresholds, these estimates can be viewed a Bayes estimates based on a diffuse (non-informative) prior for the other parameters in the model.

Table 5. Maximum Likelihood Estimates (two-step EVT)

| | EUR/USD Returns | | GBP/USD Returns | |
|---|---|---|---|---|
| | *Positive* | *Negative* | *Positive* | *Negative* |
| Threshold ($u$) | 1.30 | 1.20 | 1.02 | 1.25 |
| AIC | -11.8085 | 13.5972 | 3.2327 | -11.3623 |
| No. above | 63 | 70 | 72 | 41 |
| Scale ($\beta$) | 0.3761 | 0.3284 | 0.3005 | 0.2033 |
| (a.s.e.) | (0.0642) | (0.0604) | (0.0585) | (0.0516) |
| Shape ($\xi$) | -0.1731 | 0.1821 | 0.1971 | 0.4059 |
| (a.s.e.) | (0.0875) | (0.1416) | (0.1569) | (0.2102) |

*Note*: "a.s.e." denotes "asymptotic standard error".

The graphical method generates higher thresholds and hence smaller effective sample sizes than does the Bayesian methodology. Both the scale and shape parameter estimates are different in Tables 4 and 5.

The ML estimate of the shape parameter for the EUR/USD positive returns is still negative, again providing evidence that the tail distribution for EUR/USD positive returns has a thinner tail than the normal distribution. Although Giles, Feng, and Godwin (2016) find that the maximum likelihood estimation is still valid with a negative shape parameter, the unexpected negativity suggests that further research and analysis is required.

## 5.2 Risk Measures Estimation

The point estimates of our two risk measures are obtained using the functions in (2.3) and (2.8). VaR and ES are calculated based on the Bayes estimates of the thresholds as well as the GPD parameters, summarized in Table 4. VaR and ES are both calculated under a 1% probability that the event will happen. These estimates are presented in Table 6.

Table 6. Risk Measures (Using Bayes Estimates)

| | EUR/USD Returns | | GBP/USD Returns | |
|---|---|---|---|---|
| | *Positive* | *Negative* | *Positive* | *Negative* |
| Threshold | 0.9949 | 0.9039 | 0.8412 | 0.8878 |
| VaR | 1.9415 | 2.0285 | 1.7645 | 1.7327 |
| ES | 2.2118 | 2.5601 | 2.2750 | 2.2016 |

*Note:* For VaR and ES, the numbers are in percentages.

For the EUR/USD positive returns, with 1% probability, the daily gain from trade will exceed 1.9415% and if this gain is exceeded the expected value of this gain is 2.2118%. That is, there is a 1% probability that a currency holder will make $19,415 profit or more in one day with a $1million investment. *If this happens*, the expected value for such a gain is $22,118. For the EUR/USD negative returns, there exists a 1% probability that the daily loss will exceed 2.0385%. The expected value of such loss (of this amount or greater) is 2.56601%. This implies that with $1million investment, there is a 1% chance that a trader will suffer a loss of $20,385 or more in one day. *If this happens*, then the expected value of such loss is $25,601.

Compared with the EUR/USD returns, the GBP/USD returns exhibit extreme risk, with smaller VaR as ES estimates, for the most part. For positive returns, there is a 1% chance that the Pound can yield gain of 1.7645% or more overnight. If this occurs, the expected gain is 2.2750%, which is better than the extreme positive returns for the Euro. On the other hand, there is a 1% chance of a potential loss when trading the Great British Pound of 1.7327% or more for one trading day. If this occurs, the expected loss is 2.2016%.

Again, for comparison purposes, we also apply the two-step EVT methods to compute the risk measures. The estimators of the VaR and ES are the same, but they are based on the ML estimates of the GPD parameters. We used the R package "POT" (Ribatet, 2006) to obtain the ME and parameter charts, as well as the maximum likelihood estimates. The results appear in Table 7. Interestingly, the risk measures are similar to the previous results, despite the relatively large differences between the threshold values that are used here, and the corresponding Bayes estimates, and the differences in the estimates of the GPD parameters under the two different approaches.

Table 7. Risk Measures (Using Maximum Likelihood Estimates)

| | EUR/USD Returns | | GBP/USD Returns | |
|---|---|---|---|---|
| | *Positive* | *Negative* | *Positive* | *Negative* |
| Threshold | 1.30 | 1.20 | 1.02 | 1.25 |
| VaR | 1.9152 | 1.9664 | 1.7560 | 1.6394 |
| ES | 2.1640 | 2.5383 | 2.3109 | 2.2476 |

*Note:* For VaR and ES, the numbers are in percentages, and calculated at 1% probability.

Both risk measures, obtained with the different estimation methods, confirm that the Great British Pound appears to show less variation than does the Euro under both positive and negative extremes. From studying the thick tail of the distribution, the results show that the Great British pound has the smallest risk relative to all the foreign currencies studied. So, we can confirm that the Great British pound is less risky than the Euro. This is also consistent with the previous analysis of the descriptive statistics of two returns series.

## 6. Sensitivity Tests

Other than the truncated normal prior, we also consider a diffuse prior for $u$ to study the possible influence on the estimates of the threshold of different choices of the prior distributions. These sensitivity test results are presented in the Table 8. Compared with the previous estimates in Table 4, there are minor changes in the estimated values of each parameter. Further, the number of exceedances above the thresholds is identical in both tables. Therefore, we can conclude that the different choice of the prior distributions for the threshold has insignificant impact on the estimation of the threshold as well as the GPD parameters. The risk measures are also similar to the previous results. On the other hand, the trace plots of the estimated threshold of the GBP/USD negative returns still exhibit the more volatility than those for other three data sets. Further, under the different specification of the posterior distribution, the GBP/USD negative return data do not fit with the theoretical GPD particularly well, as was the case previously. This might suggest that the unfavourable estimation results come from the data's own characteristics, rather than from the choice of prior density.

Table 8. Sensitivity Tests Results

| | EUR/USD Returns | | GBP/USD Returns | |
|---|---|---|---|---|
| | *Positive* | *Negative* | *Positive* | *Negative* |
| Threshold ($u$) | 0.9989 | 0.9052 | 0.8414 | 0.8873 |
| (C.I) | (0.9194 , 1.0518) | (0.8259 , 0.9885) | (0.7778 , 0.9341) | (0.8359 , 0.9031) |
| No. above | 131 | 142 | 132 | 125 |
| Scale ($\beta$) | 0.4243 | 0.3715 | 0.2916 | 0.2758 |
| (C.I) | (0.3140 , 0.5478) | (0.2956 , 0.4577) | (0.2290 , 0.3685) | (0.2214 , 0.3482) |
| Shape ($\xi$) | -0.1253 | 0.0976 | 0.1498 | 0.1505 |
| (C.I) | (-0.2723 , 0.0662) | (-0.0343 , 0.2735) | (-0.0259 , 0.2985) | (0.0217 , 0.2754) |
| VaR | 1.9418 | 2.0298 | 1.7704 | 1.7373 |
| ES | 2.2139 | 2.5631 | 2.2771 | 2.2126 |

*Notes:* C.I stands for the 95% credible interval; all results are simulated using a burn-in of 1,0000 iterations and 50,000 further iterations. for VaR and ES, the numbers are in percentages, and calculated at 1% probability.

Table 9. Sensitivity Tests for EUR/USD Positive Returns

| Starting Values | | | | | Estimates | | | | |
|---|---|---|---|---|---|---|---|---|---|
| $\mu_u$ | $s_u^2$ | $u$ | $\beta$ | $\xi$ | $u$ | $\beta$ | $\xi$ | VaR | ES |
| 0.36 | 0.26 | 0.99 | 0.3 | -0.13 | 0.9917 | 0.4296 | -0.1320 | | |
| | | | | | (0.9159 , 1.0505) | (0.3219 , 0.5479) | (-0.2696 , 0.0555) | 1.9385 | 2.2076 |
| 0.40 | 0.30 | 1.00 | 0.43 | -0.10 | 0.9949 | 0.4286 | -0.1301 | | |
| | | | | | (0.9129 , 1.0516) | (0.3131 , 0.5553) | (-0.2736 , 0.0550) | 1.9415 | 2.2118 |
| 0.42 | 0.32 | 1.30 | 0.38 | -0.15 | 1.0177 | 0.4034 | -0.1135 | | |
| | | | | | (0.9139 , 1.2924) | (0.2177 , 0.5457) | (0.1761 , -0.2704) | 1.9202 | 2.1905 |

*Notes:* 95% credible intervals are reported in parentheses; all results are simulated using a burn-in of 1,000 iterations and 50,000 further iterations; for VaR and ES, the numbers are in percentages, and calculated at 1% probability.

The MCMC simulations require a starting value for each parameter of the posterior distribution. Possibly, changing in the starting values might change the final estimates as well. To study this possibility, we conduct the simulations under various combinations of the initial parameters' values using the criterion discussed in section 3.3. We find that the different staring values have very little impact on the estimation results, as long as the hyper-threshold values are chosen as a high quantile of the data. To conserve space, we present only the sensitivity test results for the EUR/USD positive returns, which appear in Table 9. Other results are also available upon request.

## 7. Conclusions

In order to account for systematic risk, large financial institutions are interested in the tail behavior—extreme events—for portfolios. Extreme value theory (EVT) is commonly applied to study the extreme distributions of financial time series. In this paper, we apply the peak over threshold (POT) method from the EVT to investigate the volatility of the exchange returns for two major reserve currencies: the Euro and the Great British Pound. The POT method is based on the results that data under the extreme situation (as defined by a predetermined threshold) follow a Generalized Pareto Distribution (GPD). Motivated by the analysis of Behrens *et al*. (2004), we conduct a modified Bayesian approach to estimate the threshold as well as the shape and the scale parameters of the GPD. Markov Chain Monte Carlo simulation with the Metropolis-Hastings sampling method is used to draw the posterior inferences. Then, two risk measures - Value at Risk and Expected Shortfall are estimated based on the corresponding

GPD. From the estimation results, we conclude that the daily returns for the Euro are more volatile than those for the Pound. In other words, the Euro is riskier than the Pound (in each case priced in U.S. dollars). This implies that for those who want to make a (relatively) safe investment, trading the Great British Pound is preferable to trading the Euro.

Further, compared with the traditional two-step peaks-over-threshold method, introduced by McNeil and Frey (1999), we find that the Bayesian approach is superior. The two-step procedure locates the threshold that determines extreme data graphically. However, the associated plots are often hard to interpret, and the resulting inferences can be unreliable. In contrast, the Bayesian approach treats the threshold as an unknown parameter and it is estimated simultaneously with the other parameters of the underlying model. Despite of the difference in the threshold choices and the estimates of the shape and scale parameters for the GPD, our approach generates the risk measures that are very similar to those obtained *via* the two-step EVT/POT approach. Some limited sensitivity testing indicates that the different choice of the prior distribution for the threshold has little impact on the Bayesian estimation results.

Although the results of our Bayesian analysis are extremely encouraging, further work remains to be done in relation to our empirical application. For example, alternative distributions for the data below the threshold values need to be considered. The treatment of the parameters associated with these distributions should to be generalized through the introduction of a prior distribution with additional hyper-parameters. Finally, there is the potential to perform a Bayesian analysis of the extreme behavior of both currencies *jointly*, using the bivariate generalized Pareto distribution. This last extension of our work is the most challenging as Bayesian analyses of the GPD to date have been limited to the univariate case. However, it would seem to be a fruitful line of research, given the success of non-Bayesian bivariate GPD studies using financial data (*e.g*., Chen *et al*., 2012).

## References


Antweiler. W., (2015). PACIFIC Exchange Rate Service. Sauder School of Business, University of British Columbia. Retrieved from: http://fx.sauder.ubc.ca/

Bank for International Settlements (2010). Basel III: A global regulatory framework for more resilient banks and banking systems. ISBN web: 92-9197-859-0

Behrens, C. N., Lopes, H. F., and Gamerman, D. (2004). Bayesian analysis of extreme events with threshold estimation. *Statistical Modeling,* 4, 227-244.

Chen, Q., Giles, D. E., and Feng, H. (2012). The extreme dependence between Chinese and other international stock markets". *Applied Financial Economics*, 22, 1147-1160.

Chen, Q.L., and Giles, D. E. (2016). Risk analysis for three precious metals: An application of extreme value theory. Mimeo., Department of Economics, University of Victoria.

Chernobai A., Rachev, S.T., and Fabozzi, F.J. (2005). Composites goodness-of-fit tests for left-truncated loss samples. *Handbook of Financial Econometrics and Statistics*, 575-596. Springer, New York.

Chivers, C. (2015). Package 'MHadaptive' (Version 1.1-8). Retrieved from: https://cran.r-project.org/web/packages/MHadaptive/MHadaptive.pdf

de Dieu, N. J., Mwita, P. N., and Mung'atu, J. K. (2014). Estimation of extreme value at risk in Rwanda exchange rate. *European Journal Statistics and Probability*, 2, 14-22.

de Jesús, R., Ortiz, E., and Cabello, A. (2013). Long run peso/dollar exchange rates and extreme value behavior: Value at risk modeling. *North American Journal of Economics and Finance*, 24, 139-152

de Zea Bermudez, P., Amaral Turkman, M. A., and Turkman, K. F. (2001). A predictive approach to tail probability estimation. *Extremes,* 4, 295-314.

Eugenia Castellanos, M., and Cabras, S. (2007). A default Bayesian procedure for the generalized Pareto distribution. *Journal of Statistical Planning and Inference,* 137, 473-483.

Fisher, R. A., and Tippett, L. H. C. (1928) Limiting forms of the frequency distribution of the largest or smallest member of a sample, *Proceedings of the Cambridge Philosophical Society*, 24, 180-190

Fretheim, T., and Kristiansen, G. (2015). Commodity market risk from 1995 to 2013: An extreme value theory approach. *Applied Economics*, 47, 2768-2782.

Giles, D.E., Feng, H., and Godwin, R.T. (2016). Bias-corrected maximum likelihood estimation of the parameters of the generalized Pareto distribution. *Communications in Statistics – Theory and Methods*, 45, 2465-2483.

Gilli, M., and Këllezi, E. (2006). An application of extreme value theory for measuring financial risk. *Computational Economics,* 27, 207-228.

Gnedenko, B. V. (1943). Sur la distribution limite du terme d'une se′ rie ale′ atoire, *Annals of Mathematics,* 44*,* 423–453.

Haile, F. D., and Pozo, S. (2006). Exchange rate regimes and currency crises: An evaluation using extreme value theory. *Review of International Economics*, 14, 554-570.

McNeil, A. J., and Frey, R. (2000). Estimation of tail-related risk measures for heteroscedastic financial time series: An extreme value approach. *Journal of Empirical Finance,*7, 271-300.

Oehlert, G. W. (1992). A note on the delta method. *American Statistician*, 46, 27-29.

Pfaff, B., McNeil, A., and Stephenson, A. (2015) Package 'evir' (Version 1.7-3). Retrieved from: https://cran.r-project.org/web/packages/evir/evir.pdf

Pickands, J. (1975) Statistical inference using extreme order statistics, *Annals of Statistics,* 3*,* 119–131.

Ren, F., and Giles, D. E. (2010). Extreme value analysis of daily Canadian crude oil prices. *Applied Financial Economics,* 20, 941-954.

Ribatet, M. (2006). A user's guide to the POT package (Version 1.4). Department of Hydrological Statistics, INRS.

Santamaría, G. S., Méndez, O. F. J., Velandia, L. F. M., and Moreno, C. A. Q. (2016). Comparison of methods for estimating the uncertainty of value at risk. *Borradores de Economía, Núm.* 927.

Wolter, T. (2015). Package 'truncgof' (Version 0.6-0). Retrieved from: https://cran.r-project.org/web/packages/truncgof/truncgof.pdf