Defence Research and Recherche et développement
Development Canada pour la défense Canada

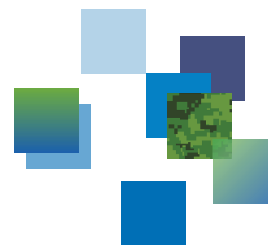# DRDC I RDDC

Foreign exchange value-at-risk with multiple currency exposure

A multivariate and copula generalized autoregressive conditional heteroskedasticity approach

David W. Maybury

Defence Research and Development Canada

Scientific Report
DRDC-RDDC-2014-R62
November 2014

Canada

Foreign exchange value-at-risk with multiple currency exposure

A multivariate and copula generalized autoregressive conditional heteroskedasticity approach

David W. Maybury

Defence Research and Development Canada

Scientific Report
DRDC-RDDC-2014-R62
November 2014

Abstract

Large DND projects and acquisitions are exposed to more than one foreign currency at the same time which complicates management's foreign exchange risk assessments. We extend the Centre for Operational Research and Analysis' (CORA) in-house Generalized Autoregressive Conditional Heteroskedasticity (GARCH) models to a full multivariate setting. Our extensions involve two models types: multivariate GARCH and copula-GARCH. We find that both models give qualitatively similar value-at-risk (VaR) estimates, and that both models provide a much improved risk assessment relative to the current practice – correcting VaR estimates on the order of 25% in cases in which multiple currency exposures are of similar size. Using the USDCAD, the EURCAD, and the GBPCAD, we demonstrate estimation techniques for each model. Finally, we show the strength of our improved models through a 100-day VaR calculation.

Significance for defence and security

The Department of National Defence (DND) expends approximately 10% of its annual budget in foreign currencies in support of equipment, acquisitions, and operations. The foreign currency transactions within DND's budget expose the department to financial market risk in the form of fluctuating exchange rates, placing additional risk mitigation pressures on internal management. Given the Government of Canada's planned military acquisitions in the form of the CF-188 replacement and the renewal of the Royal Canadian Navy – both of which will contain significant foreign content – exchange rate risk has quickly become an important risk element for DND management to understand.

In this study, we extend our in-house financial risk models to help DND understand the extent of risk mitigation that occurs naturally through exposure to more than one currency. DND's main foreign currency exposure involves the US dollar, the euro, and the British pound sterling. To gauge the effect of adverse currency fluctuations on DND's budgets and planning, we require an understanding of the strength of natural diversification benefits within value-at-risk (VaR) analyses. We examine two methods to attack this problem. Our first approach uses multivariate correlation methods which extract time dependence and persistence in correlations between exchange rate pairs. In our second approach, we build dependency models between exchange rate pairs using copulas. Both modelling techniques offer decision makers a deeper understanding of portfolio VaR – correcting the independence assumed VaR estimate by 25% in cases in which multiple currency exposures are of similar size.

Résumé

Dans ses grands projets et achats, le MDN doit souvent composer avec plusieurs devises à la fois, ce qui complique son évaluation des risques. Nous avons amplifié des modèles d'hétéroscédasticité conditionnelle autorégressive généralisée (GARCH) du Centre d'analyse et de recherche opérationnelle (CARO) de manière à y inclure toutes les variables nécessaires. Nous avons, de fait, modifié deux types de modèle : le modèle GARCH à variables multiples et le modèle GARCH à copules. Nous avons constaté que les deux modèles donnent des estimations de la valeur à risque (VaR) de qualités comparables et qu'ils permettent d'effectuer une meilleure évaluation du risque qu'avec la méthode courante. On obtient désormais des estimations de la VaR 25% plus exactes pour les cas où plusieurs devises entrent en jeu. Au moyen des combinaisons USD-CAD, EUR-CAD et GBP-CAD, nous avons démontré les techniques d'estimation des deux modèles. Enfin, nous avons démontré la fiabilité de nos modèles améliorés au moyen d'un calcul de la VaR sur 100 jours.

Importance pour la défense et la sécurité

Le ministère de la Défense nationale (MDN) dépense environ 10% de son budget annuel en devises étrangères pour l'entretien de son équipement, ses achats et ses opérations. Ces transactions en devises étrangères exposent le MDN à un risque financier lié aux variations du taux de change, et les responsables de la gestion interne se voient donc pressés de trouver des stratégies d'atténuation. Comme le gouvernement du Canada prévoit faire d'importantes acquisitions militaires, notamment pour remplacer le CF188 et renouveler la Marine royale canadienne (deux projets dans lesquels il fera beaucoup affaire avec l'étranger), le risque lié au taux de change est rapidement devenu un élément important que la direction du MDN se doit de maîtriser.

Dans notre étude, nous avons amplifié nos modèles maison d'évaluation du risque financier dans le but d'aider le MDN à comprendre l'importance de l'atténuation des risques lorsqu'on doit composer avec plus d'une devise. Les principales devises étrangères que le MDN rencontre le plus fréquemment sont le dollar américain, l'euro et la livre sterling britannique. Pour bien mesurer les effets négatifs que peuvent avoir les fluctuations monétaires sur le budget et la planification du MDN, il faut connaître le poids des bénéfices de la diversification naturelle dans les analyses de la valeur à risque (VaR). Pour aborder ce problème, nous avons examiné deux méthodes. Dans la première, nous avons utilisé des modèles de corrélation à multiples variables qui isolent la dépendance et la persistance temporelles dans les corrélations entre deux devises. Dans la seconde méthode, nous avons élaboré des modèles de dépendance à copules. Ces deux méthodes offrent aux décideurs une meilleure compréhension de la VaR d'un portefeuille, ce qui leur permet d'obtenir des estimations 25% plus exactes dans les cas où plusieurs devises entrent en jeu.

Table of contents

List of figures

This page intentionally left blank.

1 Introduction

1.1 Background

The Department of National Defence (DND) expends approximately 10% of its annual budget in foreign currencies in support of equipment, acquisitions, and operations. The foreign currency transactions within DND's budget expose the department to financial market risk in the form of fluctuating exchange rates, placing additional risk mitigation pressures on internal management. Given the Government of Canada's planned military acquisitions in the form of the CF-188 replacement and the renewal of the Royal Canadian Navy – both of which will contain significant foreign content – foreign exchange risk has quickly become an important risk element for DND management to understand.

Over the last seven years, the Centre for Operational Research and Analysis (CORA) has developed in-house financial risk management tools to help DND understand foreign exchange market risk within its budgets [1], [2]. CORA's approach uses two main lines of inquiry – models based on Generalized Autoregressive Conditional Heteroskedasticity (GARCH) [1], and models based on the internal consistency of foreign exchange derivative prices [2]. Both modelling methods yield value-at-risk (VaR) [3]: given a confidence level p, and a prescribed time horizon T, the VaR yields the foreign exchange loss level that is exceeded with probability $1 - p$ at time T. Decision makers can use VaR calculations matched to foreign payment schedules to assess the likelihood of project budget shortfalls.

In conjunction with risk modelling, CORA has helped DND prioritize corporate risk by using expert opinion with rank correlation. The analysis showed that "increased financial pressure due to external factors represents one of the highest risk factors among all corporate activities at DND" [4]. Specifically, ADM(Fin-CS) and ADM(Mat) singled out foreign exchange risk as a significant liability. Senior decision makers at DND continue to hold the view that foreign exchange represents a significant risk in DND planning activities.

In this study, we extend our in-house GARCH models to help DND understand the extent of risk mitigation that occurs naturally through exposure to more than one currency. DND's main foreign currency exposure involves the US dollar, the euro, and the British pound sterling. To gauge the effect of adverse currency fluctuations on DND's budgets and planning, we require an understanding of the strength of natural diversification benefits within VaR analyses. We examine two methods to attack this problem. Our first approach uses multivariate GARCH models which extract time dependence and persistence in correlations between exchange rate pairs. While this method offers a powerful modelling tool, the method implicitly contains restrictive assumptions about the nature of the dependency between the two exchange rates. We address these shortcomings in our second approach by building dependency models between exchange rate pairs using copulas. Both modelling techniques offer decision makers a deeper understanding of portfolio VaR in which more than one currency affects total project costs.

1.2 Scope

ADM(FinCS) requires an understanding of foreign exchange risk using VaR calculations for its budgets. In this paper, we propose a VaR analysis based two methods:

- multivariate GARCH models; and

- copula-GARCH models.

We apply these techniques to EURCAD-USDCAD, and GBPCAD-USDCAD exchange rate pairings using daily data from April 30, 2004 to May 1, 2014.

2 Generalized autoregressive conditional heteroskedasticity (GARCH) basics

Foreign exchange markets show no evidence of autocorrelation in the return data, implying that past observations of foreign exchange prices do not help forecast future returns. This observation – that past prices offer no predictive power and thus no advantage in trading – conforms with the efficient market hypothesis (EMH)[5]. However, even though unconditionally return variances are nearly stationary[1], for most assets, and especially for foreign exchange, they are not conditionally stationary. To understand the return sequence of foreign exchange data, we must account for time changing conditional return variances. The framework for modelling this phenomenon is called generalized autoregressive conditional heteroskedasticity (GARCH) [7] (see [8] for a complete review).

Under the simplest GARCH model, called GARCH(1,1), the variance process takes the form,

$$\sigma_{t+1}^2 = \omega + \alpha R_t^2 + \beta \sigma_t^2, \tag{1}$$

where $\alpha + \beta < 1$ and R_t denotes the asset return. In this paper, we use the arithmetic return,

$$R_t = \frac{S_{t+1} - S_t}{S_t}, \tag{2}$$

where S_t is the observed exchange rate at market close.[2] In GARCH modelling, we assume that the return sequence follows,

$$R_t = \sigma_t z_t, \tag{3}$$

where z_t follows an iid process.

In addition to modelling univariate time series with GARCH models, we can model the linear correlation between time series using similar methods. Note that we can re-write the univariate GARCH(1,1) model by recognizing that the unconditional variance has the relationship,

$$\sigma^2 = \mathbb{E}(\sigma_{t+1}^2) = \omega + \alpha \mathbb{E}(R_t^2) + \beta \mathbb{E}(\sigma_t^2) \tag{4}$$
$$= \omega + \alpha \sigma^2 + \beta \sigma^2, \text{ implying,}$$
$$\sigma^2 = \frac{\omega}{1 - \alpha - \beta},$$

where we assume that the daily return has a vanishing expectation. Using this approach, we can eliminate ω from the equations, allowing us to write the bivariate GARCH(1,1)

---

[1]For details on long term predictability in asset returns, see [6].

[2]In the literature, log differences are often taken as the return, but we conform to the approach taken in [9]. Note that for small changes, log and arithmetic returns are nearly identical.

process as,

$$Q_{t+1} = \mathbb{E}(\bar{z}_t \bar{z}_t^{tr})(1 - \alpha - \beta) + \alpha(\bar{z}_t \bar{z}_t^{tr}) + \beta Q_t$$

$$= \begin{pmatrix} q_{11,t+1} & q_{12,t+1} \\ q_{12,t+1} & q_{22,t+1} \end{pmatrix}$$

$$= \begin{pmatrix} 1 & \rho \\ \rho & 1 \end{pmatrix}(1 - \alpha - \beta)$$

$$+ \alpha \begin{pmatrix} z_{1,t}^2 & z_{1,t}z_{2,t} \\ z_{1,t}z_{2,t} & z_{2,t}^2 \end{pmatrix} + \beta \begin{pmatrix} q_{11,t} & q_{12,t} \\ q_{12,t} & q_{22,t} \end{pmatrix}, \tag{5}$$

where

$$\rho_{t+1} = \frac{q_{12,t+1}}{\sqrt{q_{11,t+1}q_{22,t+1}}} \tag{6}$$

and $z_{i,t}$ denotes the i-th univariate GARCH filtered return. The unconditional correlation coefficient is given by,

$$\rho = \frac{1}{T}\sum_{t=1}^{T} z_{1,t}z_{2,t}. \tag{7}$$

The multivariate GARCH model allows us to track the time dependence and the persistence of the correlation coefficient. Again, we use maximum likelihood methods for parameter estimation. The underlying structure of the multivariate GARCH model is a bivariate normal distribution with a time changing correlation coefficient.

3   A brief introduction to copulas

In statistical modelling, we often desire an understanding of the full joint distribution function of a collection of random variables. For example, the 2-dimensional case with random variables, X and Y, have the joint distribution function,

$$\mathbb{P}(X \leq x, Y \leq y) = H(x, y), \tag{8}$$

and in an optimal decision making context, we might require an inference of $H(x, y)$ from the empirical data. When confronted with the problem of estimating the dependence between random variables, we often rely on key statistical parameters, such as the covariance or the correlation coefficient. Unfortunately, dependence runs far deeper than simple covariances.

To gain an appreciation of the problem at hand, suppose that we make independent draws from the standard normal, keeping the result and its square – that is, we construct the bivariate process (X, X^2). Clearly X^2 depends perfectly on X, knowledge of the draw implies complete knowledge of its square. But notice that the covariance vanishes,

$$\operatorname{cov}(X, X^2) = \int_{-\infty}^{\infty} x^3 e^{\frac{x^2}{2}} \, dx - \left(\int_{-\infty}^{\infty} x e^{\frac{x^2}{2}} \, dx \right) \left(\int_{-\infty}^{\infty} x^2 e^{\frac{x^2}{2}} \, dx \right) = 0. \tag{9}$$

Thus, even though we know the second of our random variables depends explicitly on the first, the standard covariance relationship fails to help us understand dependence. While the correlation coefficient captures linear dependence between random variables, we miss all the non-linear information.

Copulas address this problem by specifying a parametric distribution from the marginals of the random variables.[3]

To begin our introduction to copulas, we require some definitions. We follow [10] and [11] throughout.

Definition 1 Let $S_1, \ldots, S_n$ be non-empty subsets in the extended real number line, $\overline{\mathbb{R}}$, $[-\infty, \infty]$. Let H be a real function of n variables such that $\operatorname{Dom} H = S_1 \times \cdots \times S_n$ and for n-tuples $\bar{a} \leq \bar{b}$ ($a_k \leq b_k$ for all k) let $B = [\bar{a}, \bar{b}] = [a_1, b_1] \times \cdots \times [a_n, b_n]$ be an n-box whose vertices are in $\operatorname{Dom} H$. Then the H-volume of B is given by,

$$V_H(B) = \sum_{\bar{c}} \operatorname{sgn}(\bar{c}) H(\bar{c}), \tag{10}$$

---

[3]Readers familiar with Kendall's tau and Spearman's rho for testing statistical dependence will find couplas a powerful tool. Both Kendall's tau and Spearman's rho can be re-expressed as integrals over copulas which generate extra insight into concordance relationships between random variables. For further details see [10].

where the sum over $\bar{c}$ is over all vertices of B and $\mathrm{sgn}(\bar{c})$ denotes,

$$\mathrm{sgn}(\bar{c}) = \begin{cases} 1 & \text{if } c_k = a_k \text{ for an even number of } k's, \\ -1 & \text{if } c_k = a_k \text{ for an odd number of } k's. \end{cases} \tag{11}$$

∎

For example, if B is the 3-box $[x_1, x_2] \times [y_1, y_2] \times [z_1, z_2]$, then the H-volume of B reads,

$$V_H(B) = H(x_2, y_2, z_2) - H(x_2, y_2, z_1) - H(x_2, y_1, z_2) - H(x_1, y_2, z_2)$$
$$+ H(x_2, y_1, z_1) + H(x_1, y_2, z_1) + H(x_1, y_1, z_2) - H(x_1, y_1, z_1). \tag{12}$$

In more compact notation, we write,

$$V_H(B) = \Delta_{\bar{a}}^{\bar{b}} H(\bar{t}) = \Delta_{a_n}^{b_n} \ldots \Delta_{a_1}^{b_1} H(\bar{t}) \tag{13}$$

where,

$$\Delta_{a_k}^{b_k} H(\bar{t}) = H(t_1, t_2, \ldots, b_k, t_{k+1}, \ldots, t_n) - H(t_1, t_2, \ldots, a_k, t_{k+1}, \ldots, t_n). \tag{14}$$

Definition 2 A real function H of n variables is n-increasing if $V_H(B) \geq 0$ for all n-boxes whose verticies lie in Dom H. ∎

We say that the function H is grounded if $H(\bar{t}) = 0$ for all $\bar{t}$ in Dom H such that $t_k = a_k$ for at least one k.

Definition 3 An n-dimensional distribution is a function H with domain in $\bar{\mathbb{R}}^n$ such that H is grounded, n-increasing and $H(\infty, \ldots, \infty) = 1$. ∎

If each S_k has an greatest element, b_k, then H has margins, given by

$$H_k(x) = H(b_1, \ldots, b_{k-1}, x, b_{k+1}, \ldots, b_n). \tag{15}$$

We obtain higher dimensional margins by fixing fewer places.

Definition 4 An n-dimensional copula is a function C with domain $[0, 1]^n$ such that:

- C is grounded and n-increasing;

- C has margins C_k, $k = 1, 2, \ldots, n$, which satisfy $C_k(u) = u$ for all u in $[0, 1]$.

∎

The above definitions lead us to Sklar's Theorem:

Theorem 1 (Sklar's Theorem) Let H be an n-dimensional distribution function with margins $F_1, \ldots, F_n$. Then there exists and n-copula C such that for all $\bar{x}$ in $\overline{\mathbb{R}^n}$,

$$H(x_1,\ldots,x_n) = C(F_1(x_1),\ldots,F_n(x_n)). \tag{16}$$

If $F_1,\ldots,F_n$ are all continuous, then C is unique, otherwise C is uniquely determined on $\operatorname{Ran} F_1 \times \cdots \times \operatorname{Ran} F_n$. Conversely, if C is an n-copula and $F_1,\ldots,F_n$ are distribution functions, then the function H defined aboves is an n-dimensional distribution function with margins $F_1,\ldots,F_n$. For proof, see [10]. ∎

Corollary 1 Let H be an n-dimensional distribution function with continuous margins $F_1,\ldots,F_n$ and copula C. Then for any $\bar{u}$ in $[0,1]^n$,

$$C(u_1,\ldots,u_n) = H(F_1^{-1}(u_1),\ldots,F_n^{-1}(u_n)). \tag{17}$$

Again, see [10] for details. ∎

From a model building persepective, copulas allow us to build joint distribution functions from the marginals of individual random variables. In this paper, we will restrict ourselves to 2-dimensional copulas, namely,

$$H(x,y) = C(u,v) = C(F^{-1}(x), G^{-1}(y)), \tag{18}$$

where $F(x)$ and $G(x)$ denote the respective 1-dimensional distribution functions of each random variable. We concentrate on two copula types:

Gaussian Copula:

$$
\begin{aligned}
C_N(u,v;\rho) &= \Phi_N(\Phi^{-1}(u),\Phi^{-1}(v);\rho) \\
&= \frac{1}{2\pi(1-\rho^2)^{1/2}} \int_{-\infty}^{\Phi^{-1}(u)} \int_{-\infty}^{\Phi^{-1}(v)} \exp\left(\frac{-(s^2 - 2\rho st + t^2)}{2(1-\rho^2)}\right) ds\,dt,
\end{aligned} \tag{19}
$$

where $\Phi(\cdot)$ is the standard normal distribution function.

Gumbel Copula and the Gumbel Survival Copula:

$$C_G(u,v;\alpha) = \exp\left[-\left\{(-\log(u))^{1/\alpha} + (-\log(v))^{1/\alpha}\right\}^{\alpha}\right] \tag{20}$$

$$C_{GS}(u,v;\alpha) = u + v - 1 + C_G(1-u,1-v). \tag{21}$$

We can easily show that the convex sum of copulas is also a copula, and thus we consider two model class for our foreign exchange analysis

$$
\begin{aligned}
C_1(u,v;\theta,\rho,\alpha) &= (1-\theta)C_N(u,v;\rho) + \theta C_G(u,v;\alpha) \\
C_2(u,v;\theta,\rho,\alpha) &= (1-\theta)C_N(u,v;\rho) + \theta C_{GS}(u,v;\alpha)
\end{aligned} \tag{22}
$$

where the convexity parameter satisfies $\theta \in [0,1]$. Using maximum likelihood estimations, we can determine which model best describes the foreign exchange data.

Our coupla selection offers a wide range of dependency fitting potential with parsimony [12]. The foreign exchange data exhibits features contained in both copula classes. The Gaussian copula belongs to a symmetric multivariate elliptical family – a feature that foreign exchange data exhibits at a gross level – and the Gumbel (survival) captures strong right (left) tail dependence, allowing us to model asymmetrical fat tails. In principle, we could create a five dimensional model by using a convex sum of the Gaussian, Gumbel, and Gumbel survival copula, but we find that the reduced parsimony does not give qualitative better results.

Once we have a copula fit to the data, we require a method for sampling to generate Monte Carlo simulations. Observe that,

$$\mathbb{P}(V \le v | U = u) = \lim_{\Delta u \to 0} \frac{C(u + \Delta u, v) - C(u, v)}{\Delta u} = C_u(u, v). \tag{23}$$

Thus we can calculate the conditional probability by taking the derivative of the copula with respect to one of its arguments. To sample from a copula, we generate two realizations from the uniform distribution, u, and t, and then solve

$$C_u(u, v) = t, \tag{24}$$

for v. The resulting pair, (u, v), represents the value of each marginal sampled from the copula. Inverting the marginals $u = F(x)$, and $v = G(y)$, we obtain the Monte Carlo realization of our random variables.

4 Foreign exchange data and analysis

We use ten years of foreign exchange data [13], 30 April 2004 to 1 May 2014, for USD-CAD, EURCAD, and GBPCAD, giving a total of 2609 observations in each time series. Figures 1, 2, 3 show the historical observations. DND's largest foreign exchange exposure is the USDCAD, followed by EURCAD, and GBPCAD [1] and we therefore focus on the pairwise dependency relationship between the USDCAD and DND's more minor exposures. From a risk analysis perspective, since the USDCAD dominates DND's foreign exchange obligations, any diversification benefit comes primarily from the USDCAD's relationship to the other two currencies.

In the QQ-plots of Figures 4, 5, and 6 we immediately see that the financial return data is not iid $N(0,1)$. We fit GARCH(1,1) models to our data using quasi-maximum likelihood methods [8]. Using the Kolmogorov-Smirnov test on the filtered returns, we find that we cannot reject the null hypothesis that the innovations for each times belong to iid $N(0,1)$. The QQ-plots of Figures 7, 8, and 9 show the quality of the GARCH filtered returns relative to the standard normal.

In Figures 10, 11, and 12 we see the time changing stochastic volatility for each time series expressed as an annualized percentage. Notice the sharp rise in volatility beginning in late September 2008, demarking the onset of the financial crisis. Using the GARCH filtered return data, we compute the dynamics of the linear correlation coefficient between the currency pairs EURCAD-USDCAD and GBPCAD-USDCAD; we display the results in Figures 13, and 14.

While the multivariate GARCH model captures time dependence and persistence, it does not allow asymmetrical heavy tailed relationships between pairs of time series. To reveal these features in the data, we use convex sums of Gaussian, and Gumbel (survival) copulas. We take each GARCH filtered time series pair and convert the data to their marginals. We do not parametrically model the marginals, instead relying on the empirical distribution functions themselves (recall that the Kolmogorov-Smirnov test does not reject $N(0,1)$ as the distribution function for each GARCH filtered return time series). In Figures 15, and 16 we see the scatter plots in the marginal space which immediately show a Southwest to Northeast direction – the time series pairs exhibit dependence, which we already discovered using the multivariate GARCH technique. Fitting the convex copula sum, we find that the Gaussian and Gumbel survival copula summed together provide the best fit to the data for both time series pairs. We show in Figures 17 and 18 the scatter plot data of Figures 15, and 16 with the iso-contours of the best fit copula density function. The effect of the asymmetric tail, captured by the Gumbel survival copula component, is readily apparent. To examine the robustness of our result, we perform a 500 sample bootstrap of the data to generate error estimates on each parameter of the copula model. We show the bootstrap results in Figure 19 and Figure 20.

We test each copula model for time dependence in each of its parameters in turn, holding the others fixed, by partitioning the data according to a 4×4 grid of the unit square. Each of the copula parameters $\Theta = (\theta, \rho, \alpha)$, belong to the grid through the sum,

$$\Theta_{i,t} = \sum_{j}^{16} d_j [(z_{1,t-1}, z_{2,t-1}) \in \mathscr{A}_j] \tag{25}$$

where $\mathscr{A}_j$ corresponds to the j-th element of the partitioned unit square and d_j is the 16-tuple, $(d_1, d_2, \ldots, d_{16})$, extension of the copula parameter considered. For example, $\mathscr{A}_1 = [0, p_1) \times [0, q_1)$; $\mathscr{A}_2 = [p_1, p_2) \times [0, q_1]$; etc. The choice of the partition is arbitrary and we follow [14], by setting $(p, q) \in (0.15, 0.50, 0.85)$ to ensure that we focus on large values within the partition elements. Thus the likelihood sum becomes,

$$\tilde{\Theta}_i = \arg\min \sum_{t=1}^{T} \ln \left(c(u_1, u_2); \Gamma(z_{1,t-1}, z_{2,t-1}; \Theta_i) \right), \tag{26}$$

where $c(u_1, u_2)$ is the copula density as a function of the empirical marginals and,

$$\Gamma(z_{1,t-1}, z_{2,t-1}; \Theta_i), \tag{27}$$

corresponds to the selection procedure of eq.(25). While this process can elicit time dependence, it cannot reveal persistence as in the multivariate GARCH estimation (the partition method looks back only one time step). The maximum likelihood method allows us to estimate each d_j (16-tuple) along with its standard error, which allows us to test for statistical significance in differences by using the $\chi^2(15)$ test [14]. We do not find evidence for time dependence in any of the copula parameters for each of the time series pairs.

To test the quality of the copula fit to the data, we construct a 7×7 contingency table for each pairing, comparing the theoretical occupation number in each cell, $A_{i,j}$, to the empirical occupation number $B_{i,j}$. The test statistic

$$M = \sum_{i}^{k} \sum_{j}^{k} \frac{(B_{i,j} - A_{i,j})^2}{A_{i,j}}, \tag{28}$$

is χ^2 distributed with $(k-1)^2$ degrees of freedom. We find that in both time series pairs, we cannot reject the null hypothesis that the empirical contingency table is generated by the theoretical distribution. We display the results in Tables 1, 2 and as a heatmap in Figures 21, 22, 23, and 24.

Table 1: Seven quantile contingency table test for the copula fit to EURCAD - USDCAD.

Exchange rate pairings	Quantile counts							$\chi^2(36)$
	1st	2nd	3rd	4th	5th	6th	7th	
EURCAD-USDCAD (empirical)								
Quantile counts								
7th	16	21	26	38	46	74	151	
6th	18	29	42	48	74	80	82	
5th	30	40	55	63	81	60	44	
4th	38	51	68	82	57	51	25	
3rd	48	77	57	52	54	52	33	
2nd	83	77	72	54	35	32	20	
1st	139	78	53	35	26	24	17	
EURCAD-USDCAD (predicted)								
Quantile counts								
7th	10	20	29	40	54	79	139	
6th	20	33	41	52	66	81	79	
5th	29	41	50	62	70	66	54	
4th	39	50	62	69	62	52	40	
3rd	52	65	73	62	50	41	29	
2nd	76	88	65	50	41	33	20	
1st	146	76	52	39	29	20	10	43.4 (c.v. 51.0)

Table 2: Seven quantile contingency table test for the copula fit to GBPCAD - USDCAD.

| Exchange rate pairings	Quantile counts							$\chi^2(36)$
	1st	2nd	3rd	4th	5th	6th	7th	
GBPCAD-USDCAD (empirical)								
Quantile counts								
7th	9	16	20	33	50	89	155	
6th	13	28	44	56	63	80	89	
5th	19	35	49	67	80	59	63	
4th	39	52	63	62	82	51	24	
3rd	48	71	65	69	54	40	26	
2nd	84	88	75	51	28	36	11	
1st	160	83	57	34	16	18	4	
GBPCAD-USDCAD (predicted)								
Quantile counts								
7th	5	14	24	37	55	85	154	
6th	14	28	41	54	69	83	85	
5th	23	40	53	63	70	69	55	
4th	35	53	64	67	63	54	37	
3rd	52	69	72	64	53	41	24	
2nd	82	87	69	53	40	28	14	
1st	163	82	52	35	23	14	5	36.1 (c.v. 51.0)
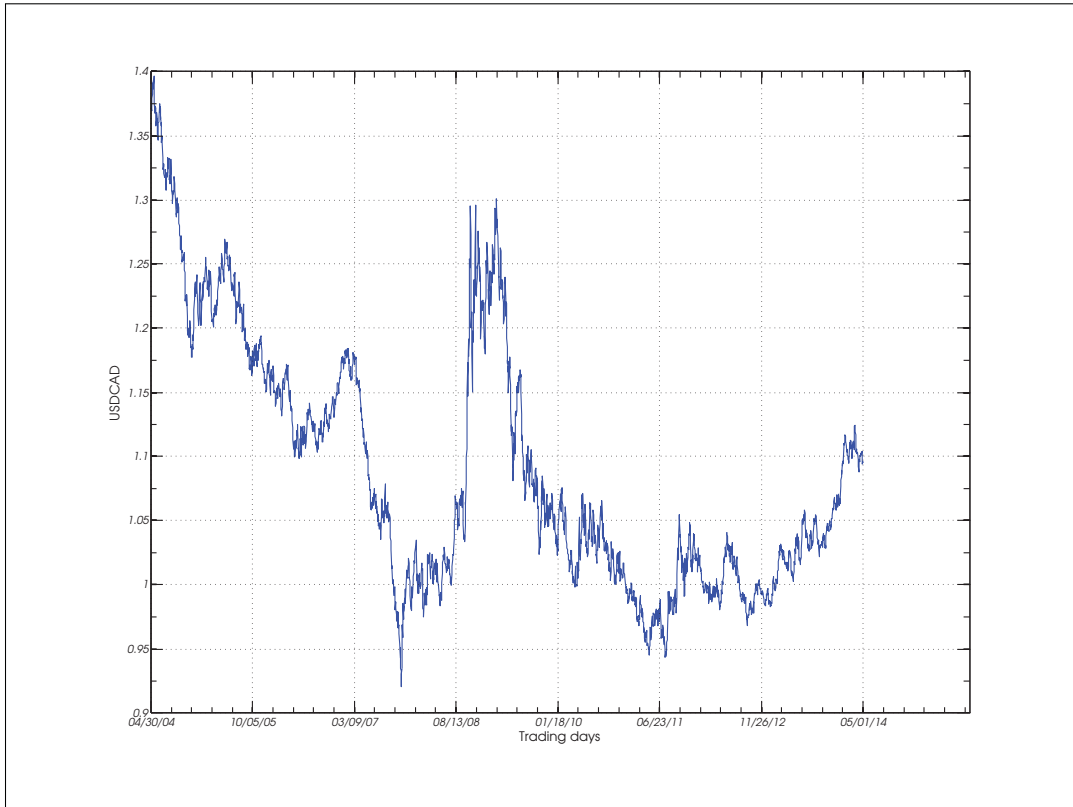
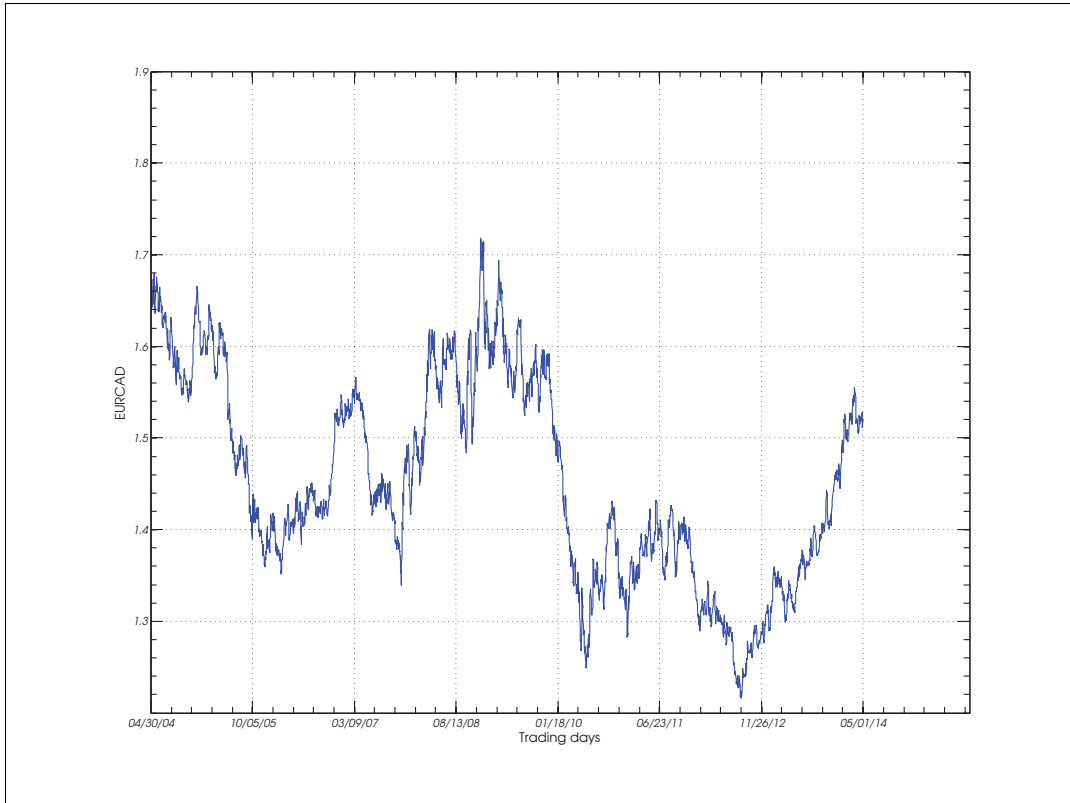Figure 1: USDCAD performance April 30, 2004 – May 1, 2014.

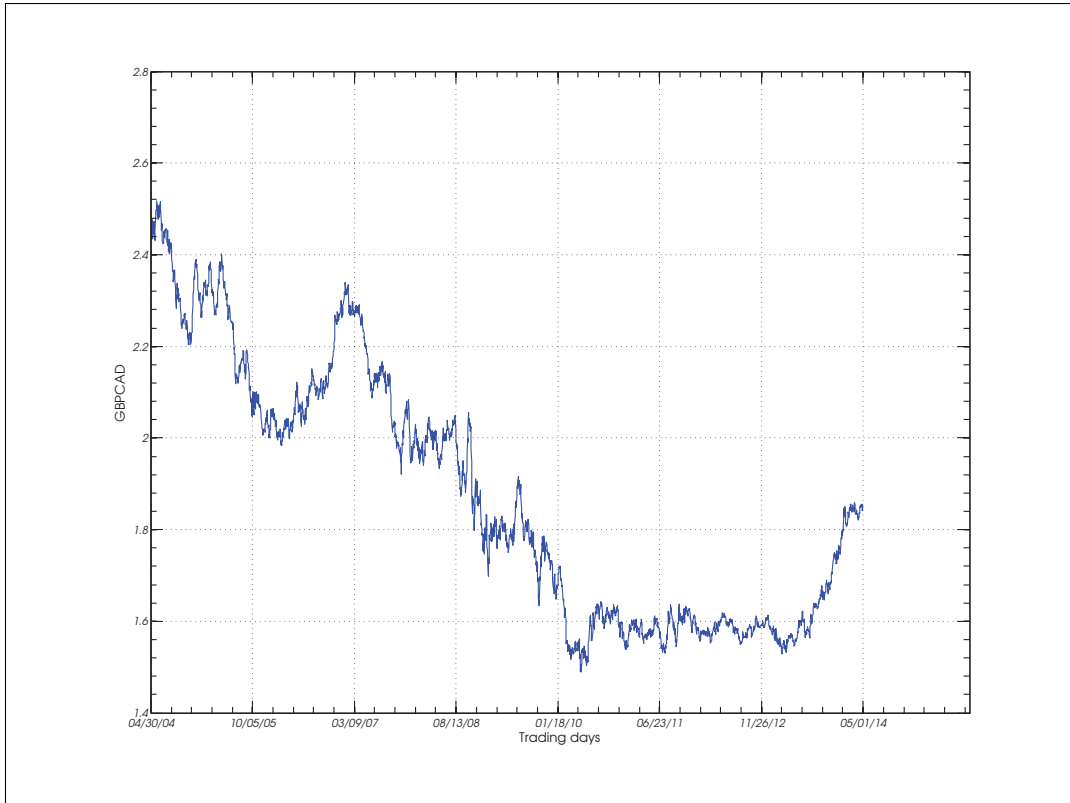Figure 2: EURCAD performance April 30, 2004 – May 1, 2014.

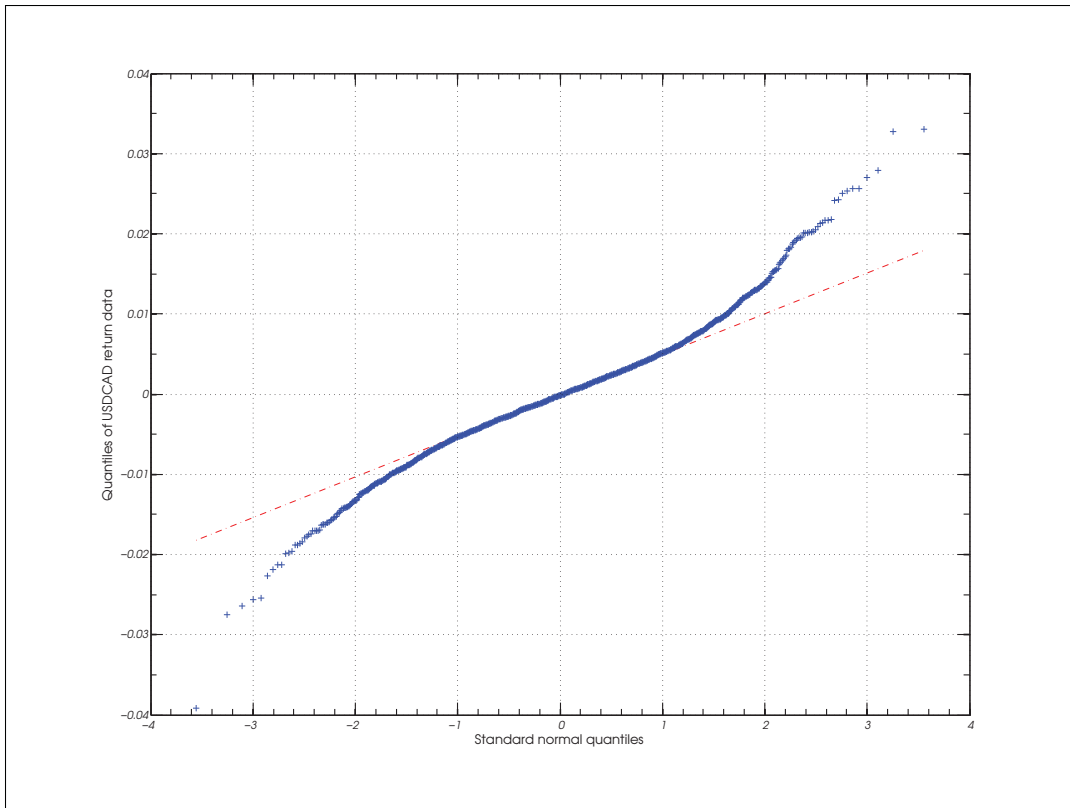Figure 3: GBPCAD performance April 30, 2004 – May 1, 2014.

Figure 4: USDCAD QQ-plot returns vs standard normal April 30, 2004 – May 1, 2014.
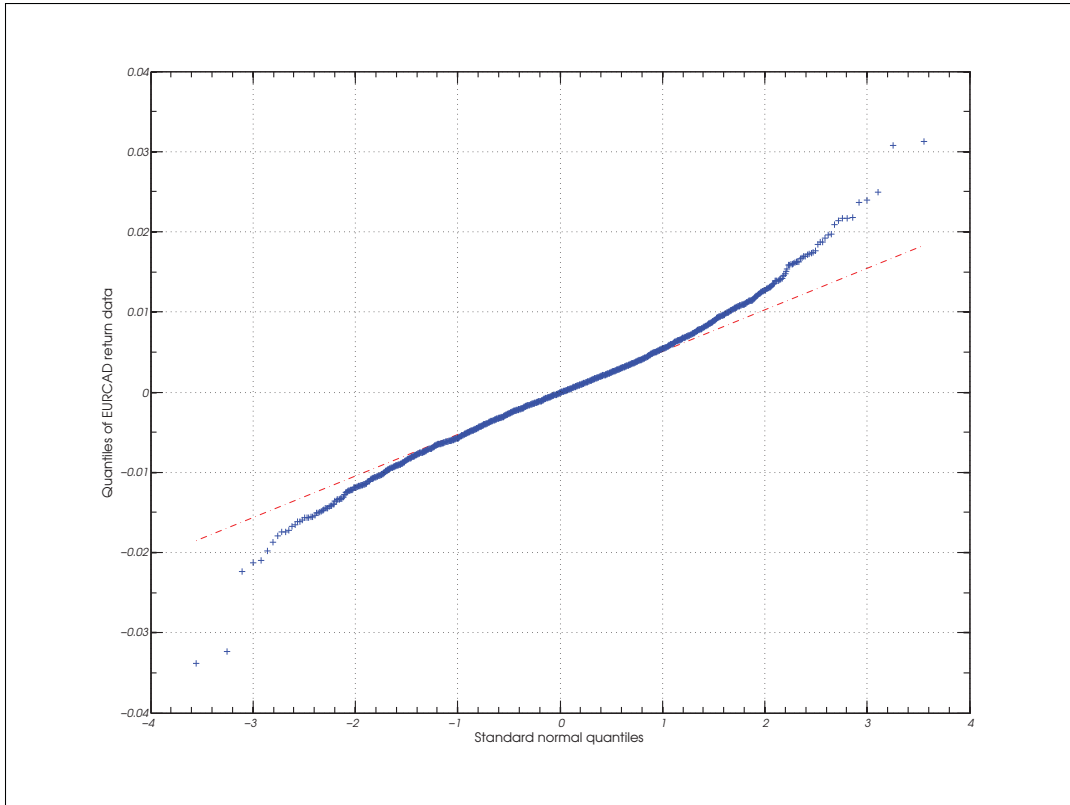
DRDC-RDDC-2014-R62

Figure 5: EURCAD QQ-plot returns vs standard normal April 30, 2004 – May 1, 2014.
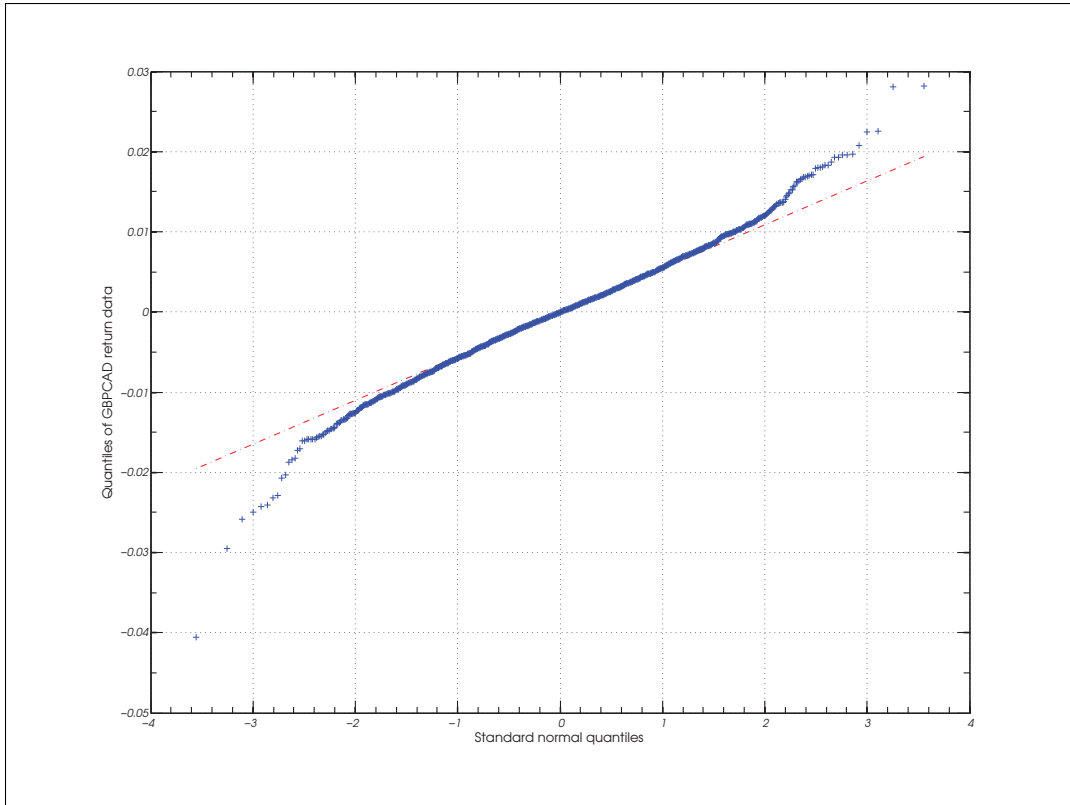
Figure 6: GBPCAD QQ-plot returns vs standard normal, April 30, 2004 – May 1, 2014.
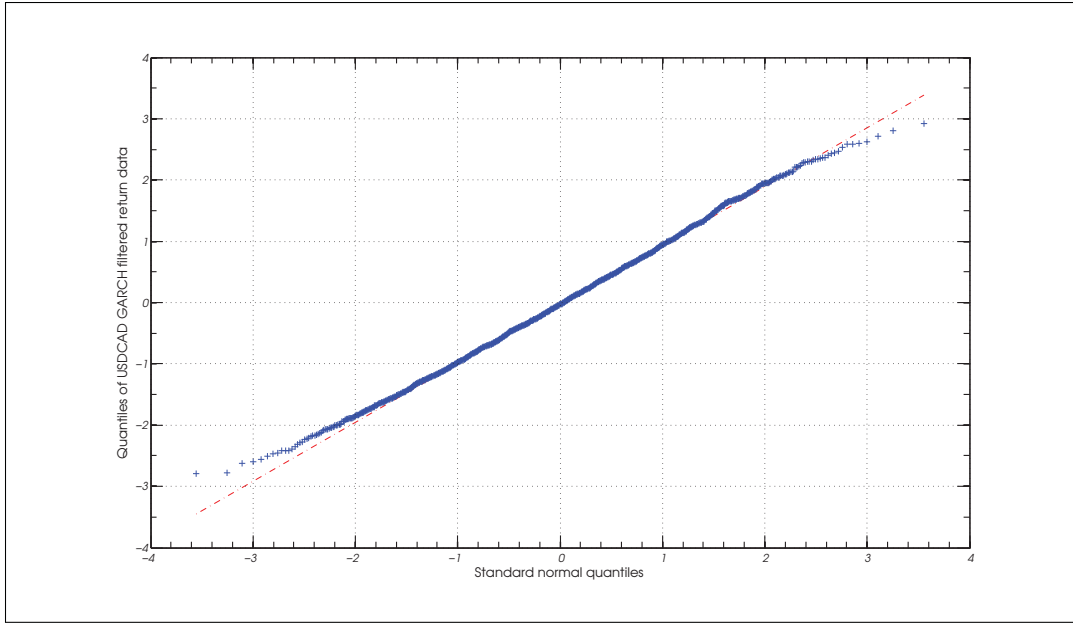
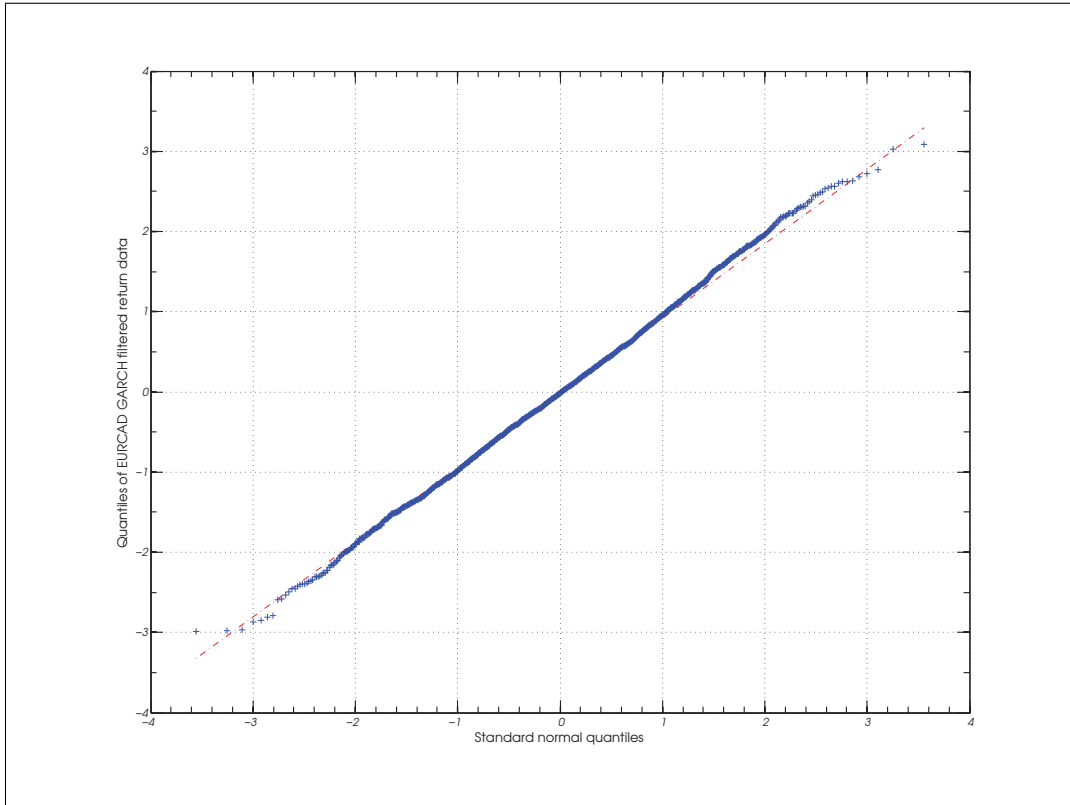Figure 7: USDCAD QQ-plot GARCH filtered returns vs standard normal April 30, 2004 – May 1, 2014.

Figure 8: EURCAD QQ-plot GARCH filtered returns vs standard normal April 30, 2004 – May 1, 2014.
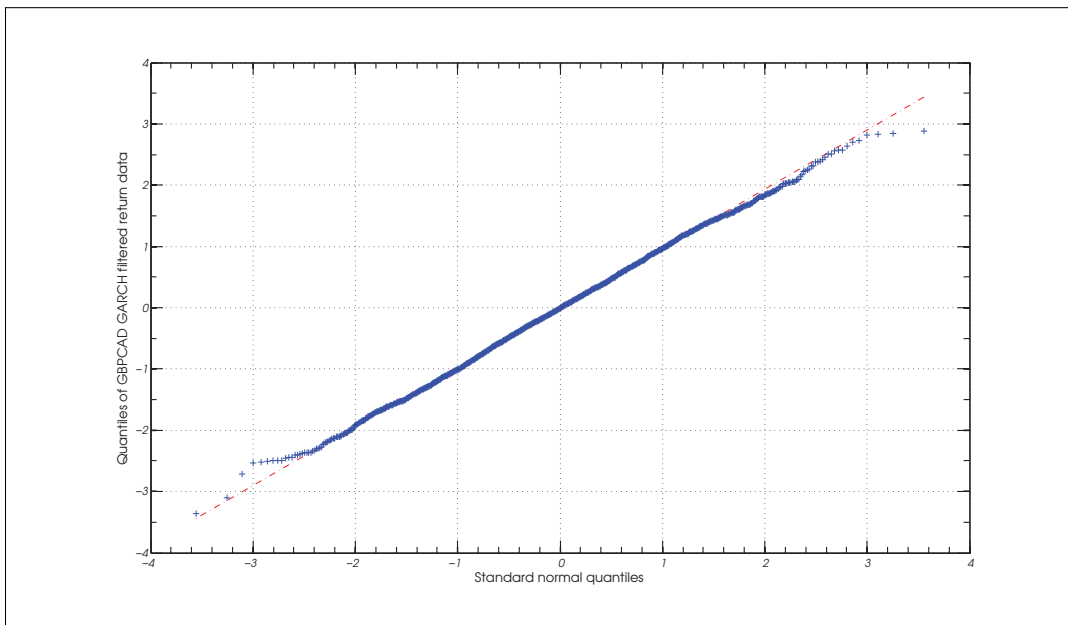
Figure 9: GBPCAD QQ-plot GARCH filtered returns vs standard normal, April 30, 2004 – May 1, 2014.

Figure 10: USDCAD GARCH modelled annualized volatility April 30, 2004 – May 1, 2014.
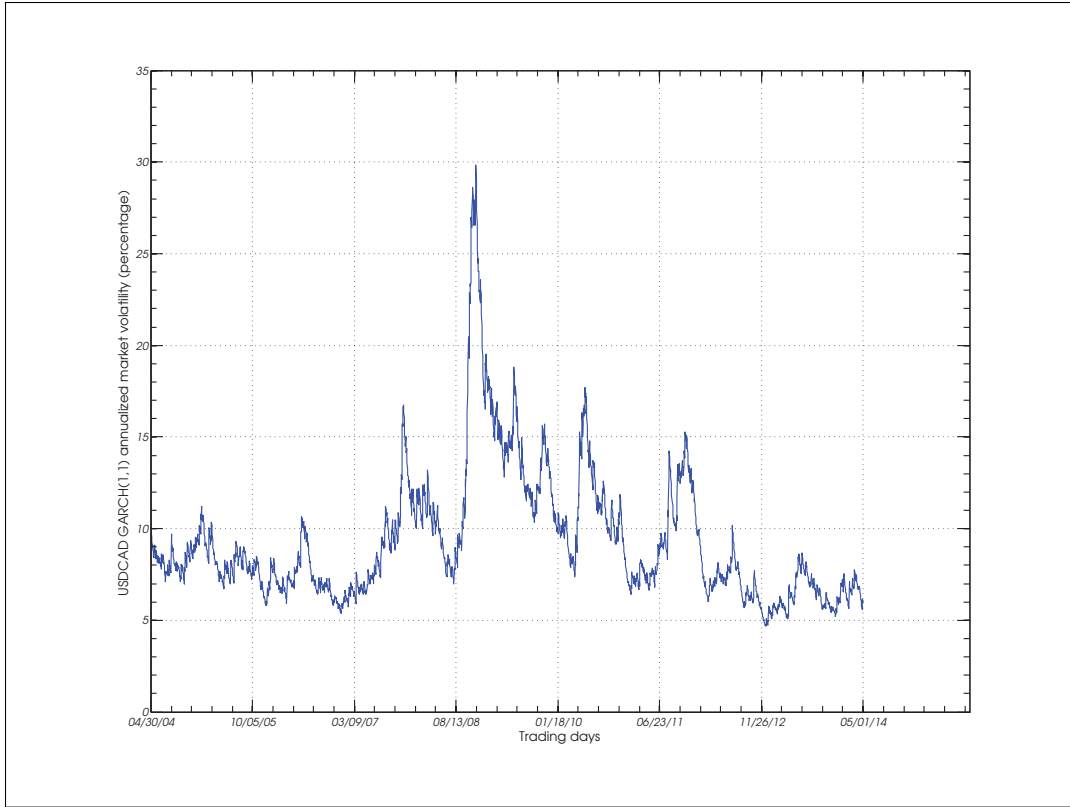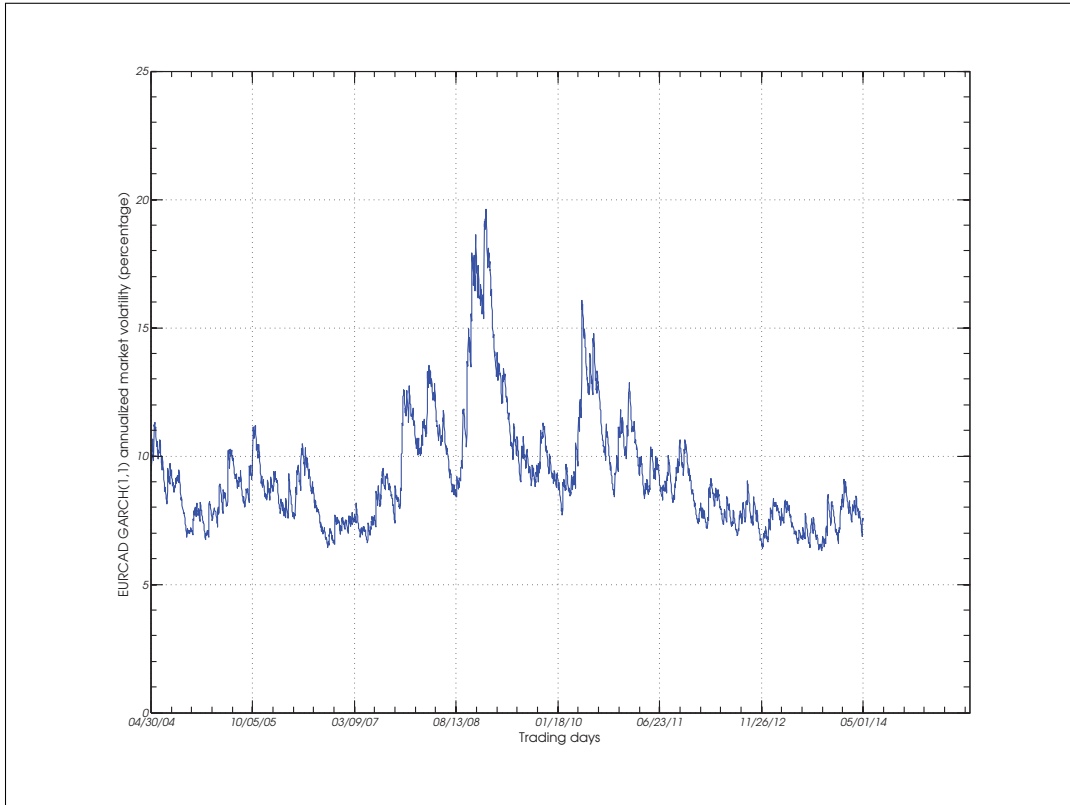
Figure 11: EURCAD GARCH modelled annualized volatility April 30, 2004 – May 1, 2014.
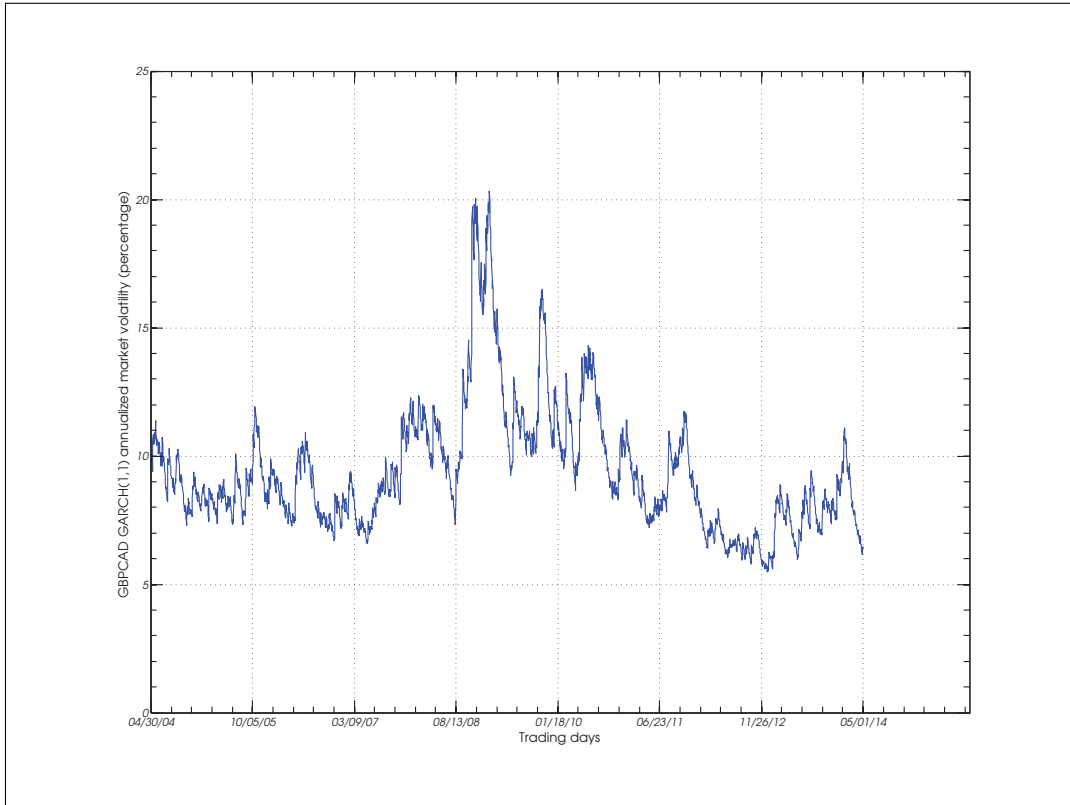
Figure 12: GBPCAD GARCH modelled annualized volatility, April 30, 2004 – May 1, 2014.
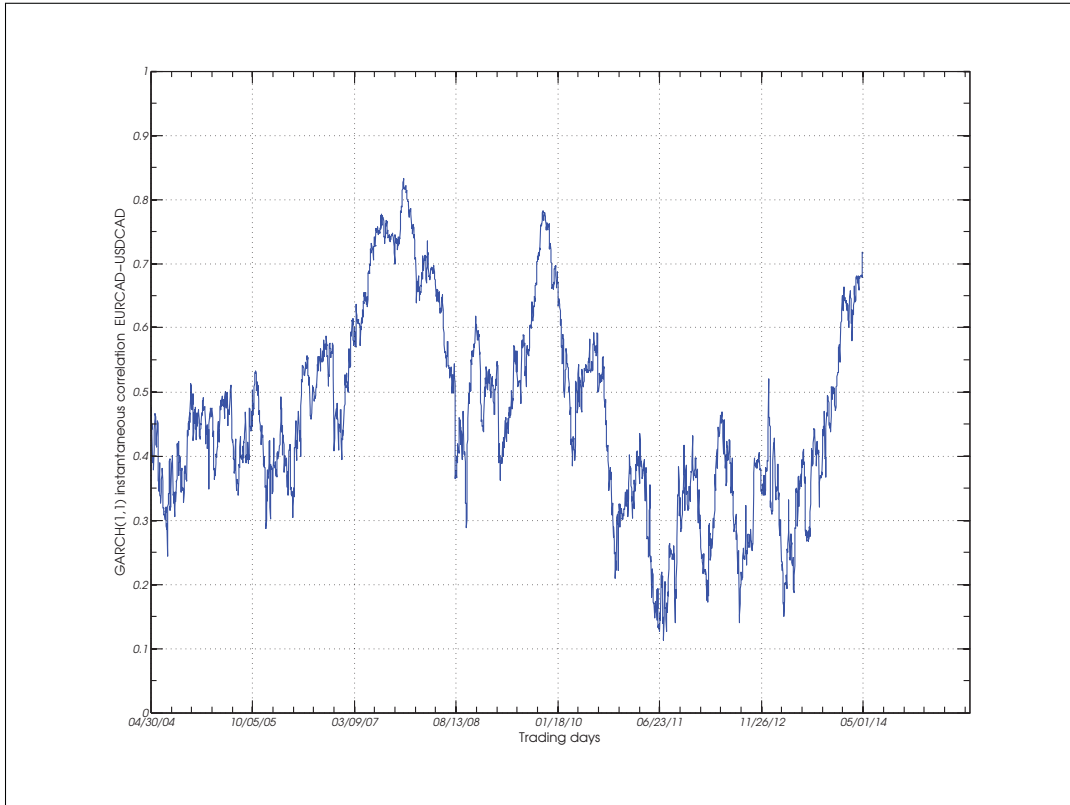
Figure 13: The correlation coefficient for the EURCAD-USDCAD pair April 30, 2004 –
May 1, 2014.
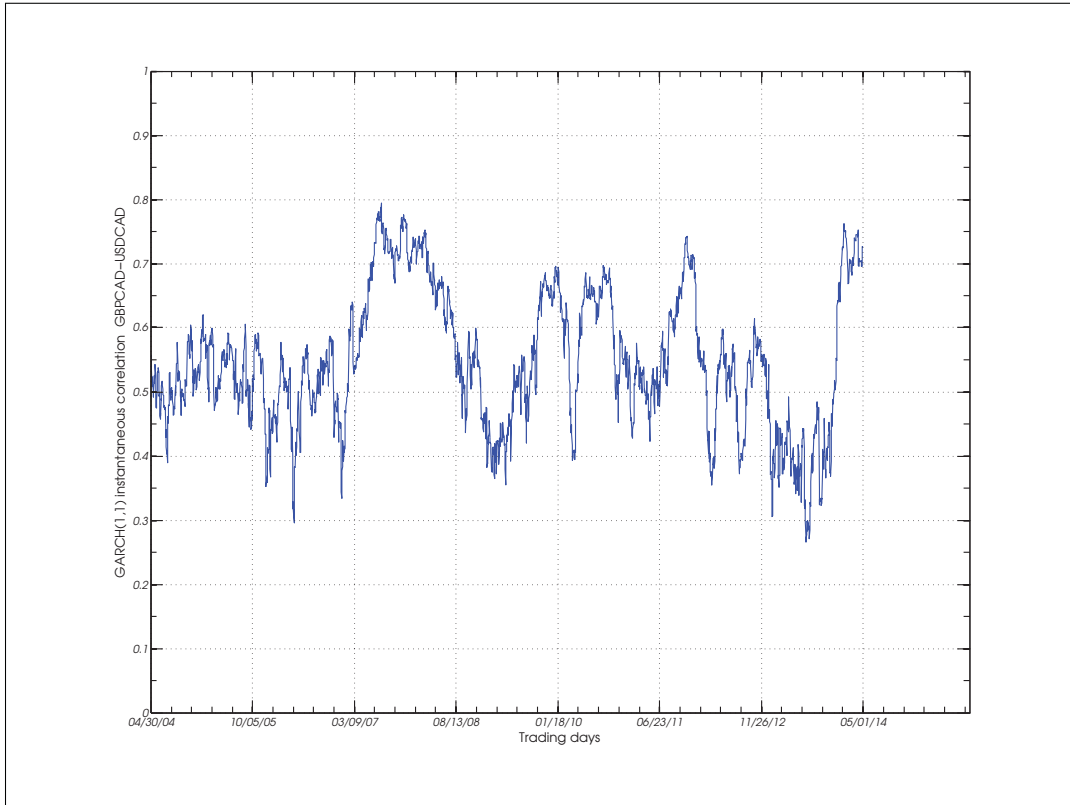
Figure 14: The correlation coefficient for the GBPCAD-USDCAD pair April 30, 2004 – May 1, 2014.
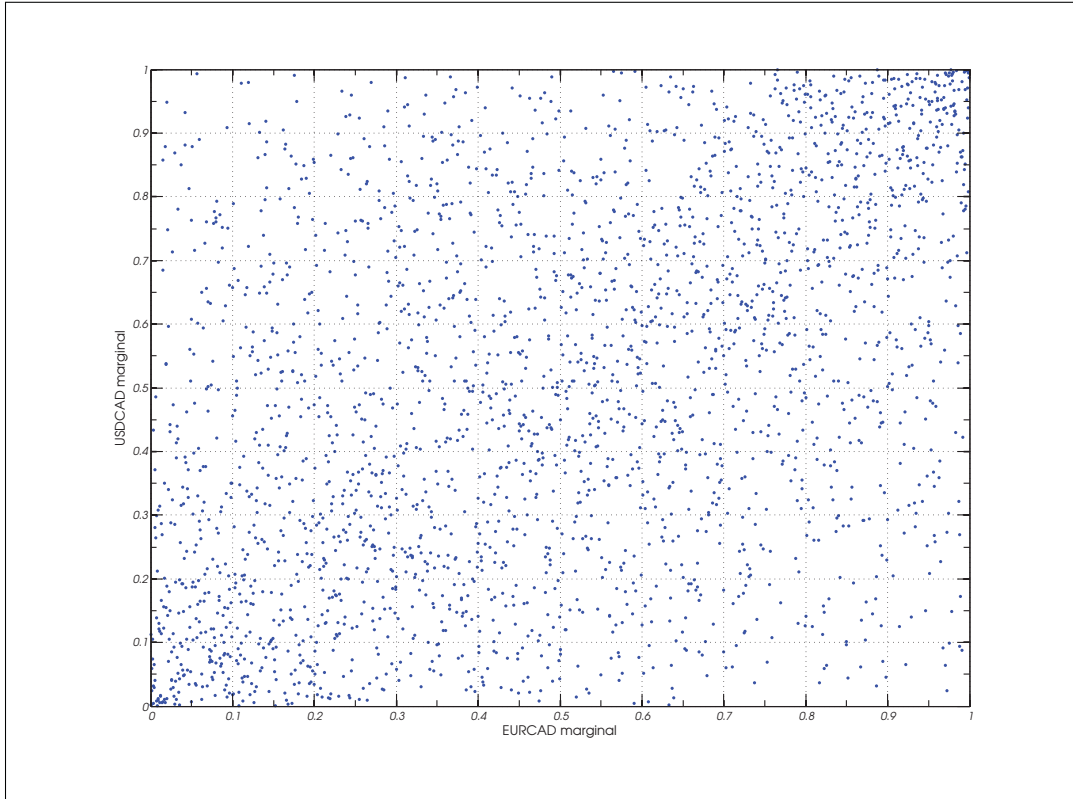
Figure 15: EURCAD-USDCAD marginals, April 30, 2004 – May 1, 2014.
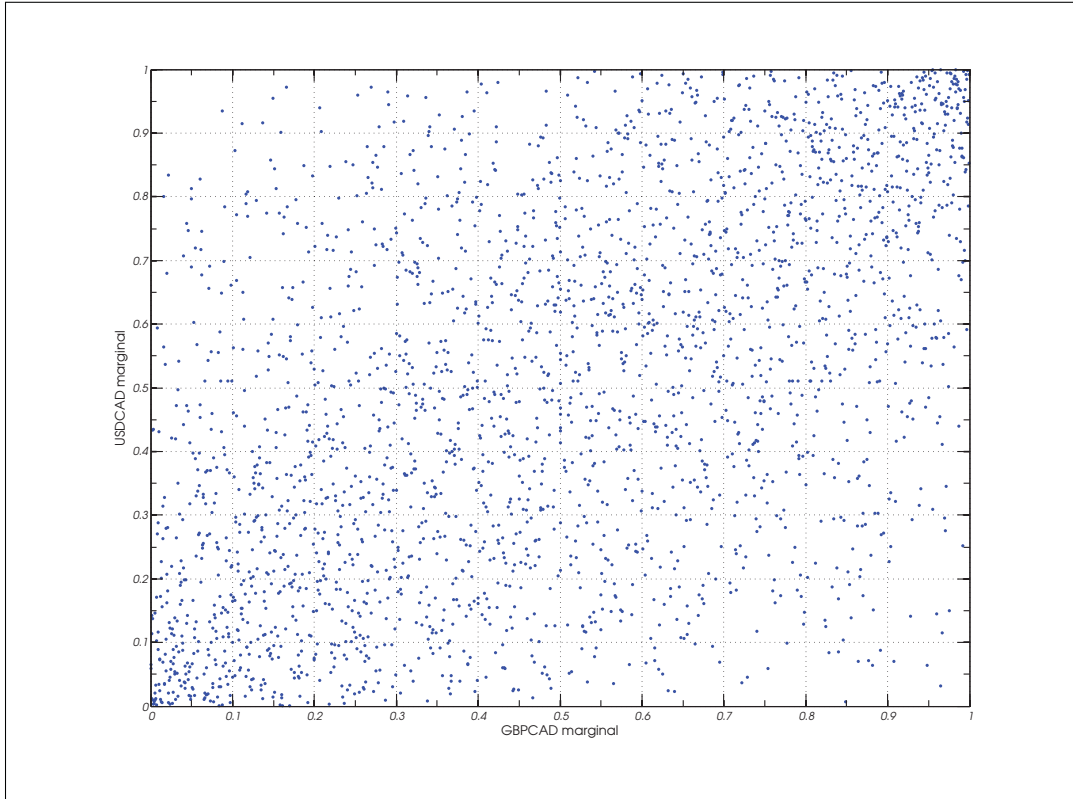
Figure 16: GBPCAD-USDCAD marginals, April 30, 2004 – May 1, 2014.
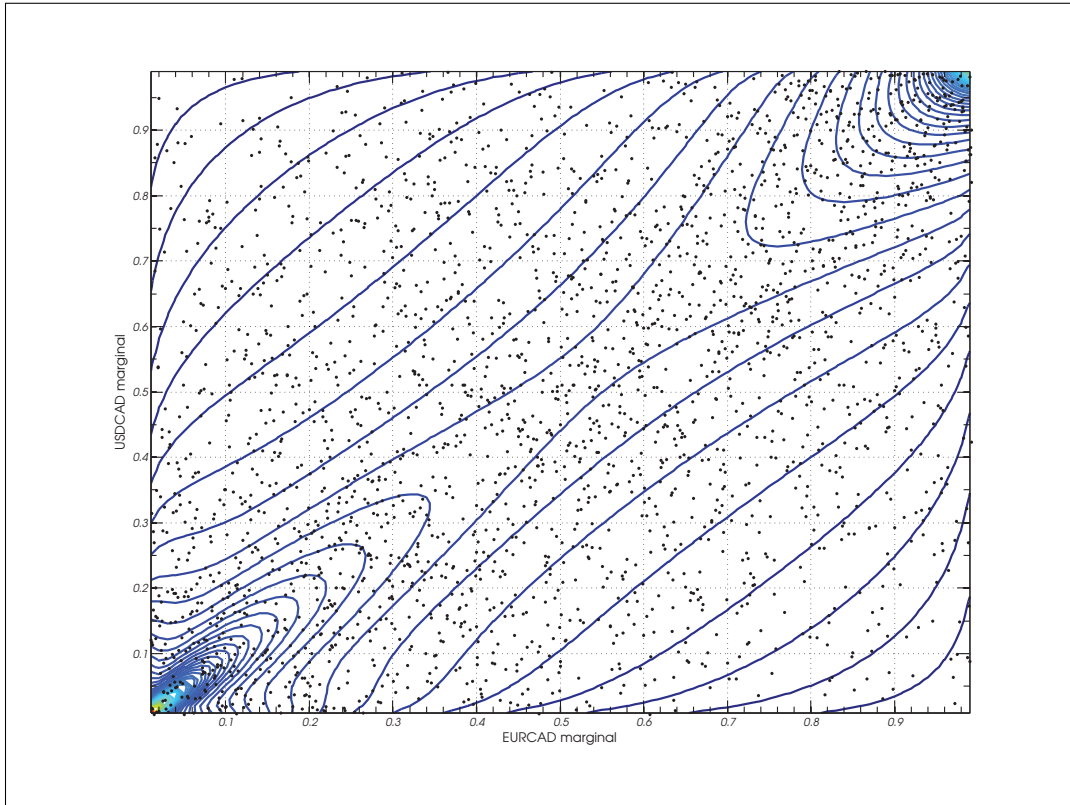
Figure 17: EURCAD-USDCAD Gaussian Gumbel survival copula density fit, April 30, 2004 – May 1, 2014.
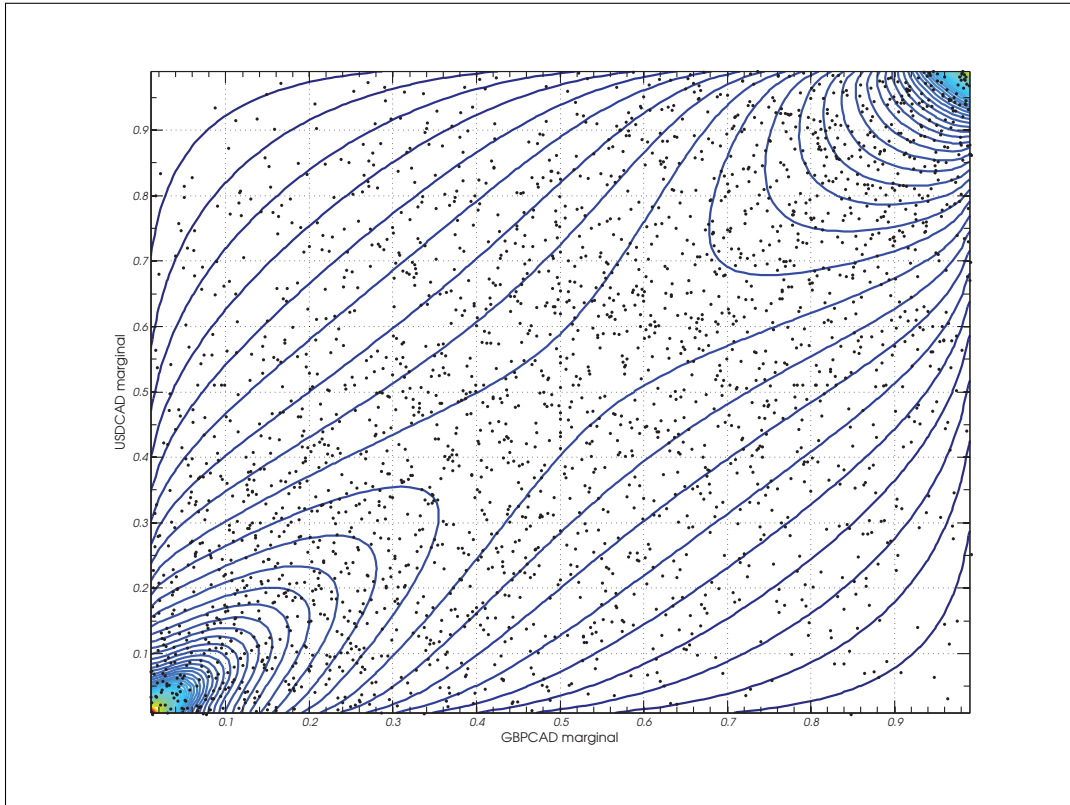
Figure 18: GBPCAD-USDCAD Gaussian Gumbel survival copula density fit, April 30, 2004 – May 1, 2014.

Figure 19: EURCAD-USDCAD bootstrapped (500 samples) copula parameter error distributions.
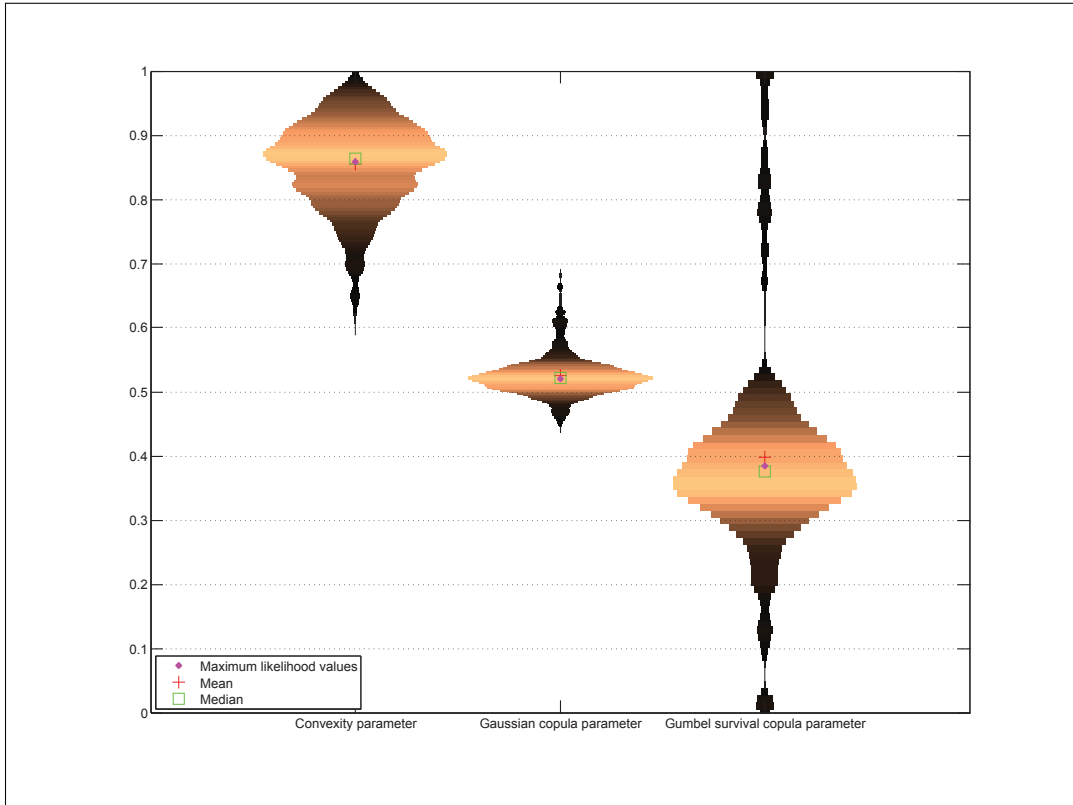
Figure 20: GBPCAD-USDCAD bootstrapped (500 samples) copula parameter error distributions.
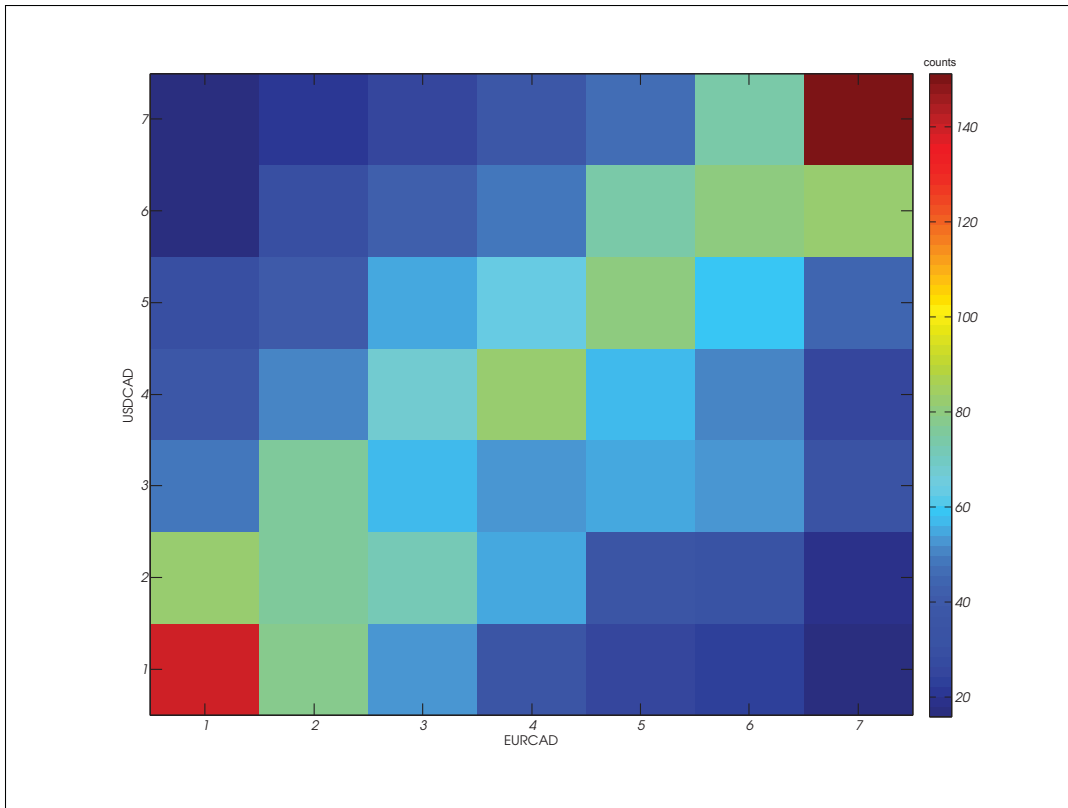
Figure 21: Heatmap depiction of the EURCAD-USDCAD empirical contingency table.
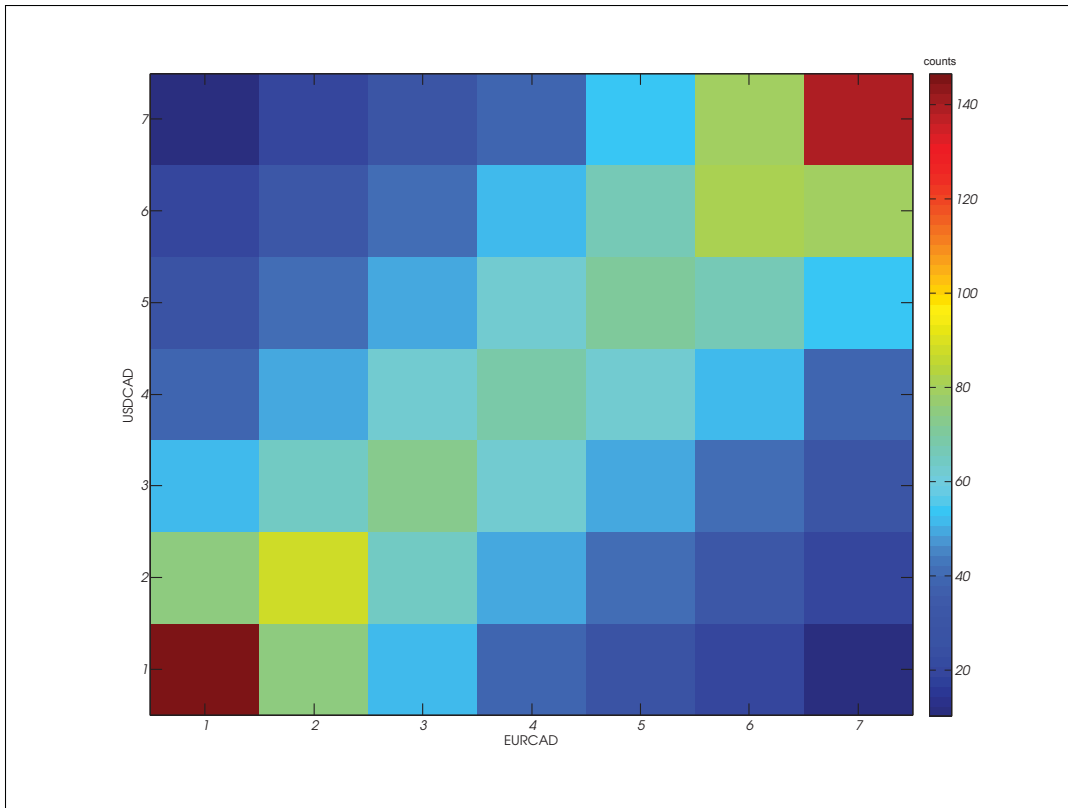
Figure 22: Heatmap depiction of the EURCAD-USDCAD copula predicted contingency table.
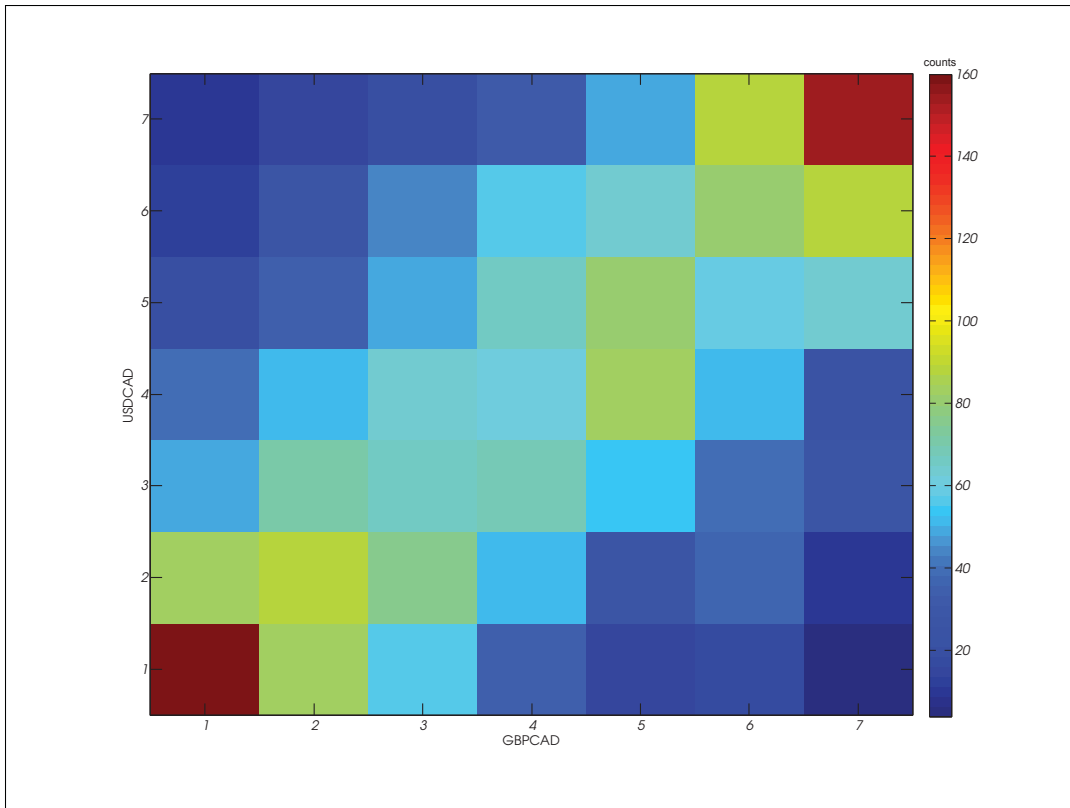
Figure 23: Heatmap depiction of the GBPCAD-USDCAD empirical contingency table.
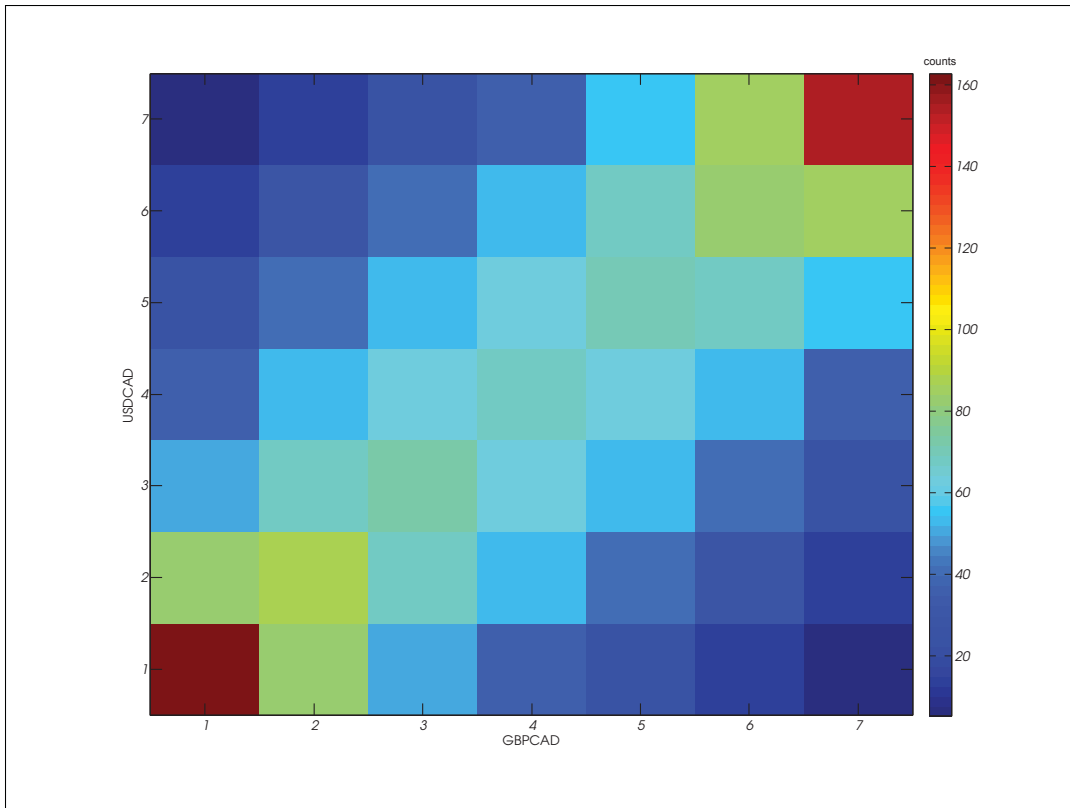
Figure 24: Heatmap depiction of the GBPCAD-USDCAD copula predicted contingency table.

Table 3: 100 trading day VaR levels from May 1, 2014 for a portfolio short 3 million USD and 1 million GBP. Results based on Monte Carlo simulation (sample size: 3×10^4).

Model	VaR$_{10\%}$ ($M CAD)	VaR$_{5\%}$ ($M CAD)	VaR$_{1\%}$ ($M CAD)
Copula-GARCH	0.24	0.32	0.48
Multivariate-GARCH	0.27	0.34	0.52
Independence	0.21	0.27	0.41

5   VaR analysis

We can use the copula and multivariate GARCH models to build VaR estimates through Monte Carlo techniques. The multivariate GARCH approach ensures that the resulting simulations contain the modelled correlation's time dependence and persistence between exchange rate pairs. The copula GARCH method models the joint distribution function as a fixed copula, ensuring VaR levels consistent with dependency relationships.

To demonstrate the applicability of the VaR estimate, consider a hypothetical project that is short 3 million USD and short 1 million GBP, both due 100 trading days after May 1, 2014. In table 3 we see the foreign exchange VaR levels relative to the May 1, 2014 spot purchase of 5.14 million CAD. We see that both dependency models give similar results while the independence assumption underestimates the risk in the portfolio: at the 1% VaR level, the independence assumption yields a result 25% lower than the dependency models. In figure 25 and 26 we show 1000 simulated joint exchange rate realizations for both dependency models.

Figure 25: GBPCAD-USDCAD simulation: copula scatter plot 100 trading day horizon from May 1, 2014 (1000 points).

Figure 26: GBPCAD-USDCAD simulation: multivariate GARCH scatter plot 100 trading day horizon from May 1, 2014 (1000 points).

6 Discussion

We present two methods – the multivariate GARCH, and the copula-GARCH – for estimating the foreign exchange return process. Using both methods, we can compute the VaR for a portfolio of foreign exchange obligations, properly including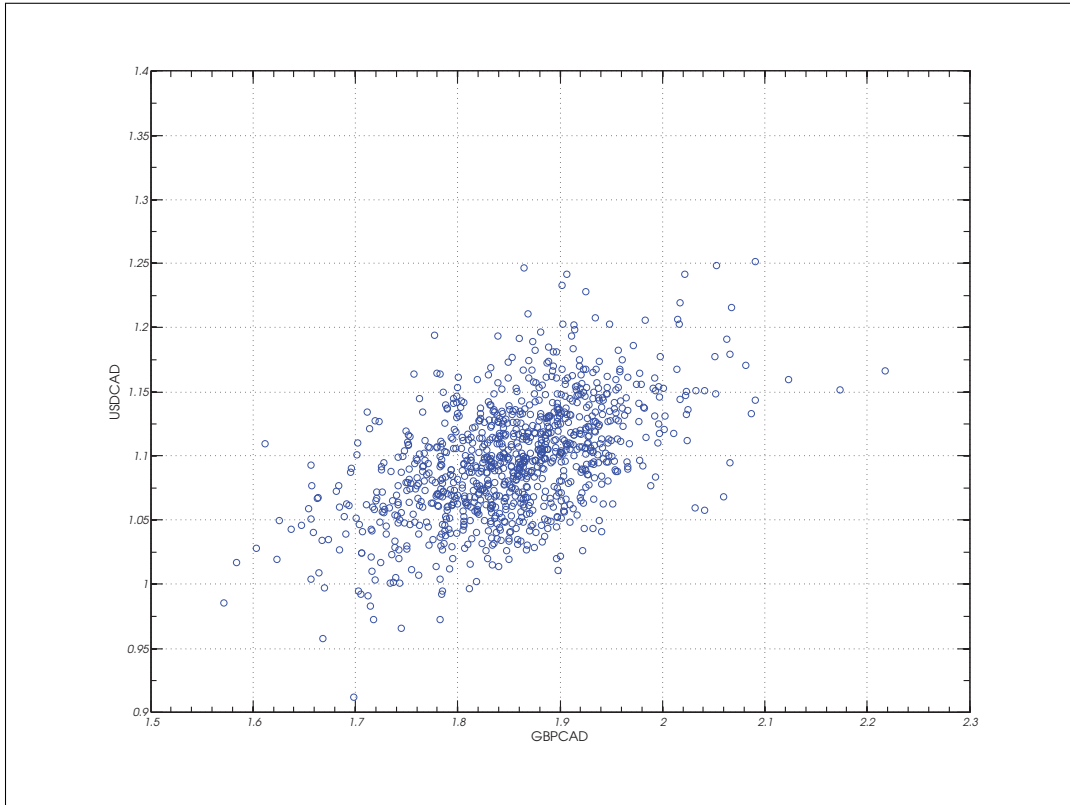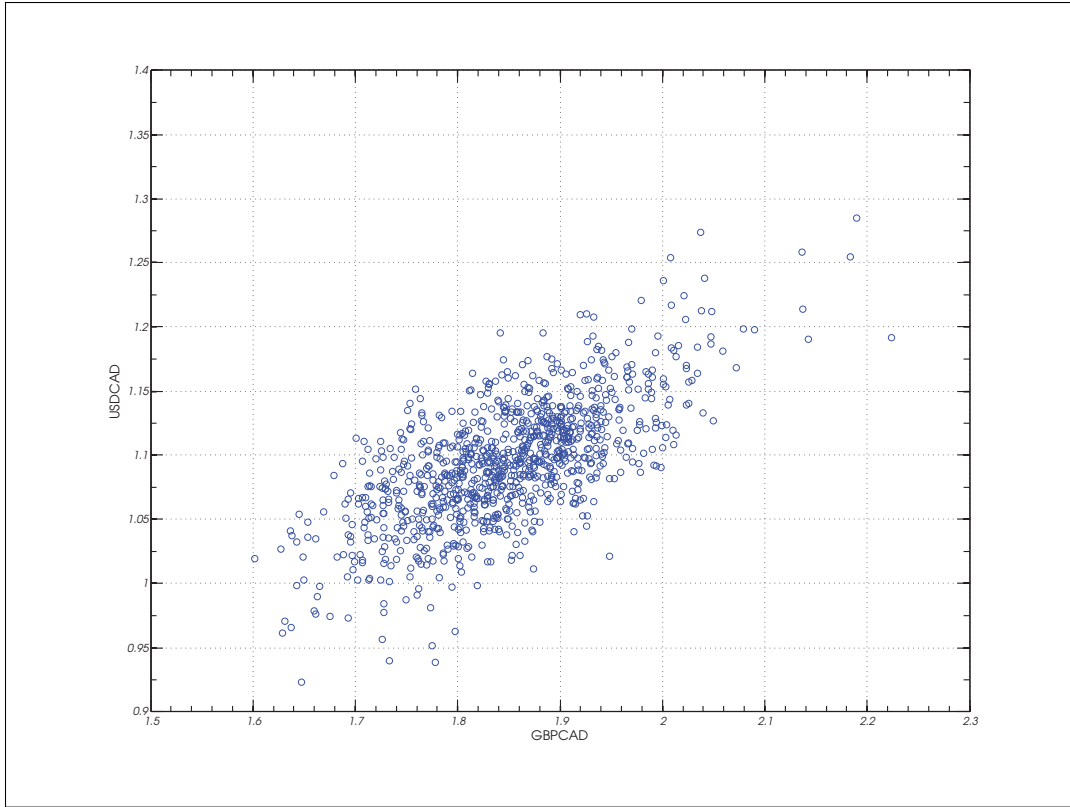 the dependence between currency pairs. We can construct Monte Carlo samples from each model and reverse the GARCH process to give simulated returns. Both models offer advantages – the multivariate GARCH model captures the persistence of shocks to the correlation coefficient, while the copula-GARCH model captures asymmetric tail dependence. Using both models will help DND managers understand multiple foreign currency exposures risk within project budgets.

We find that the pairings EURCAD-USDCAD and GBPCAD-USDCAD exhibit strong dependence in both modelling approaches. When the Canadian dollar does well or poor against one of the currencies, the Canadian dollar tends to follow the same performance against the other currency. Thus, DND's multiple currency exposures provides some, but limited, diversification benefit – especially during tail events.

Senior decision makers can use these models to gauge the foreign exchange risk within project budgets. We see that the exposure to multiple currencies does not eliminate foreign exchange risk and we cannot simply convolve distributions from univariate time series for each currency to estimate portfolio VaR. The dependency between the exchange rates plays an important role in determining the overall VaR level of the portfolio.

While VaR offers a risk monitoring mechanism for DND, it does not in and of itself offer risk mitigation. DND must decide on the risks it is comfortable holding. As a cautionary note, before embarking on any risk mitigation strategy in foreign exchange – contingency reserve or derivative based – DND must understand the source of possible gains. In financial firm evaluation, creating or buying insurance against unfavourable outcomes does not affect the value of a firm. Insurance repackages the firm's cash flows, and cash flow repackaging cannot increase firm value.[4] If insurance does provide value, it must come from advantages that the insurance separately creates. For example, if management can re-focus its efforts in areas where it has a comparative advantage by entering an insurance contract, then the repackaging can have a positive effect. In DND's case, financial risk mitigation of any kind must be tied to the real source of potential gains – superior project management.

CORA has afforded DND three powerful methods for modelling and understanding foreign exchange risk across department budgets: mulitvariate GARCH models, copula-GARCH models, and models based on the prices of vanilla foreign exchange options (in particular, the risk neutral probability distribution from the Heston model). Each modelling method

---

[4]This is a statement of the famous Modigliani-Miller Theorem – a no arbitrage argument for firm pricing. See [15] for details.

offers insight into the exchange rate market risk that DND assumes in foreign transactions. They capture realistic features of market data including heavy tails, asymmetric density functions, and leptokurtosis. The inclusion of regular VaR updates in risk reporting over time horizons sensitive to major project deliveries will help senior DND decision makers better understand their foreign currency liabilities in military operations.

References

[1]   Desmier, P. (2007), Estimating Foreign Exchange Exposure in the Department of National Defence, (DRDC-RDDC-2006-TR023) Defence Research and Development Canada – CORA.

[2]   Maybury, D. (2013), Modelling the US Dollar Trading Range: Bounds from the Risk Neutral Measure, (DRDC-RDDC-2013-TM086) Defence Research and Development Canada – Atlantic.

[3]   Duffie, D. and Pan, J. (1997), An Overview of Value at Risk, *The Journal of Derivatives*, 4(3), 7–49.

[4]   Yazbeck, T. (2010), Analytical Support to the Prioritization of Corporate Risks in the Department of National Defence, (DRDC-RDDC-2010-TM158) Defence Research and Development Canada – CORA.

[5]   Fama, E. (1970), Efficient Capital Markets: A Review of Theory and Empirical Work, *Journal of Finance*, 25(2), 383–417.

[6]   Cochrane, J. (2005), Asset Pricing, Princeton University Press.

[7]   Bollerslev, T. (1986), Generalized Autoregressive Conditional Heteroskedasticity, *Journal of Econometrics*, 31(3), 307–27.

[8]   Hamilton, J. (1994), Time Series Analysis, Princeton University Press.

[9]   Hull, J. (2007), Options, Futures, and Other Derivatives, 6th Ed., Prentice Hall.

[10]   Nelsen, R. (2006), An Introduction to Copulas, 2nd Ed., Springer.

[11]   Embrechts, P., Lindskog, F., and McNeil, A., Modelling Dependence with Copulas and Applications to Risk Management.

[12]   Hu, L. (2006), Dependence patterns across financial markets: a mixed copula approach, *Applied Financial Economics*, 16(10), 717–729.

[13]   Bloomberg, L. P., Foreign exchange data EURCAD, GBPCAD, USDCAD 04/30/04 to 05/01/14. Retrieved May 1, 2014.

[14]   Jondeau, E. and Rockinger, M. (2006), The Copula-GARCH model of conditional dependencies: An international stock market application, *Journal of International Money and Finance*, 25, 827–853.

[15]   Copeland, T., Weston, J., and Shastri, K. (2005), Financial Theory and Corporate Policy, 4th Ed., Addison Wesley.

	DOCUMENT CONTROL DATA			
	(Security markings for the title, abstract and indexing annotation must be entered when the document is Classified or Designated.)			

1.	ORIGINATOR (The name and address of the organization preparing the document. Organizations for whom the document was prepared, e.g. Centre sponsoring a contractor's report, or tasking agency, are entered in section 8.) DRDC – Centre for Operational Research and Analysis Dept. of National Defence, MGen G. R. Pearkes Bldg., 101 Colonel By Drive, 6CBS, Ottawa ON K1A 0K2, Canada	2a.	SECURITY MARKING (Overall security marking of the document, including supplemental markings if applicable.) UNCLASSIFIED
		2b.	CONTROLLED GOODS (NON-CONTROLLED GOODS) DMC A REVIEW: GCEC APRIL 2011

3.	TITLE (The complete document title as indicated on the title page. Its classification should be indicated by the appropriate abbreviation (S, C or U) in parentheses after the title.) Foreign exchange value-at-risk with multiple currency exposure

4.	AUTHORS (Last name, followed by initials – ranks, titles, etc. not to be used.) Maybury, D. W.

5.	DATE OF PUBLICATION (Month and year of publication of document.) November 2014	6a.	NO. OF PAGES (Total containing information. Include Annexes, Appendices, etc.) 54	6b.	NO. OF REFS (Total cited in document.) 15

7.	DESCRIPTIVE NOTES (The category of the document, e.g. technical report, technical note or memorandum. If appropriate, enter the type of report, e.g. interim, progress, summary, annual or final. Give the inclusive dates when a specific reporting period is covered.) Scientific Report

8.	SPONSORING ACTIVITY (The name of the department project office or laboratory sponsoring the research and development – include address.) DRDC – Centre for Operational Research and Analysis Dept. of National Defence, MGen G. R. Pearkes Bldg., 101 Colonel By Drive, 6CBS, Ottawa ON K1A 0K2, Canada

9a.	PROJECT OR GRANT NO. (If appropriate, the applicable research and development project or grant number under which the document was written. Please specify whether project or grant.)	9b.	CONTRACT NO. (If appropriate, the applicable number under which the document was written.)

10a.	ORIGINATOR'S DOCUMENT NUMBER (The official document number by which the document is identified by the originating activity. This number must be unique to this document.) DRDC-RDDC-2014-R62	10b.	OTHER DOCUMENT NO(s). (Any other numbers which may be assigned this document either by the originator or by the sponsor.)

11.	DOCUMENT AVAILABILITY (Any limitations on further dissemination of the document, other than those imposed by security classification.) (X) Unlimited distribution ( ) Defence departments and defence contractors; further distribution only as approved ( ) Defence departments and Canadian defence contractors; further distribution only as approved ( ) Government departments and agencies; further distribution only as approved ( ) Defence departments; further distribution only as approved ( ) Other (please specify):

12.	DOCUMENT ANNOUNCEMENT (Any limitation to the bibliographic announcement of this document. This will normally correspond to the Document Availability (11). However, where further distribution (beyond the audience specified in (11)) is possible, a wider announcement audience may be selected.)

13. ABSTRACT (A brief and factual summary of the document. It may also appear elsewhere in the body of the document itself. It is highly desirable that the abstract of classified documents be unclassified. Each paragraph of the abstract shall begin with an indication of the security classification of the information in the paragraph (unless the document itself is unclassified) represented as (S), (C), or (U). It is not necessary to include here abstracts in both official languages unless the text is bilingual.)

Large DND projects and acquisitions are exposed to more than one foreign currency at the same time which complicates management's foreign exchange risk assessments. We extend the Centre for Operational Research and Analysis' (CORA) in-house Generalized Autoregressive Conditional Heteroskedasticity (GARCH) models to a full multivariate setting. Our extensions involve two models types: multivariate GARCH and copula-GARCH. We find that both models give qualitatively similar value-at-risk (VaR) estimates, and that both models provide a much improved risk assessment relative to the current practice – correcting VaR estimates on the order of 25% in cases in which multiple currency exposures are of similar size. Using the USDCAD, the EURCAD, and the GBPCAD, we demonstrate estimation techniques for each model. Finally, we show the strength of our improved models through a 100-day VaR calculation.

14. KEYWORDS, DESCRIPTORS or IDENTIFIERS (Technically meaningful terms or short phrases that characterize a document and could be helpful in cataloguing the document. They should be selected so that no security classification is required. Identifiers, such as equipment model designation, trade name, military project code name, geographic location may also be included. If possible keywords should be selected from a published thesaurus. e.g. Thesaurus of Engineering and Scientific Terms (TEST) and that thesaurus identified. If it is not possible to select indexing terms which are Unclassified, the classification of each should be indicated as with the title.)

Copula; Financial Risk Management; Foreign Exchange; Heteroskedasticity; Value-at-Risk

DRDC | RDDC