Quantifying and Managing Foreign Exchange Risk in the Canadian Department of National Defence

P.E. Desmier* and D.W. Maybury

October 28, 2011

ABSTRACT

Quantifying foreign exchange risk within military acquisition presents the Department of National Defence (DND) with a nontrivial budget problem. Industry has accepted standard financial reporting methods that center on value-at-risk (VaR), which provide a quantitative measure on the downside risk associated with foreign currency transactions. We construct a model designed to forecast DND's foreign exchange expenditures within departmental financial accounts and to capture the effect of time-varying volatility on DND's foreign exchange transaction risk. Our resulting VaR model yields the maximum expected loss from adverse currency fluctuations over a budget horizon. Finally, we show that DND can mitigate its foreign exchange risk by applying simple hedging techniques.

*Corresponding author contact: Dr. Paul E. Desmier, Defence Research & Development Canada (paul.desmier@forces.gc.ca)

Contents

List of Tables

List of Figures

1.0 INTRODUCTION

To support military operations at home and abroad, the Canadian Department of National Defence (DND) often acquires new equipment and supplies from foreign sources. Between fiscal years (FY) 99/00 and 10/11, DND expended $12.0 billion Canadian (CAD) in foreign currencies of which 86% was spent in U.S. dollars (USD), 5% in U.K. sterling, and 6% in euros. Currently, foreign currency transactions account for approximately 10% of overall DND spending.

Canada, like most industrialized nations, uses a floating exchange rate, allowing its currency to fluctuate freely according to the demands of the foreign exchange market. As floating exchange rates automatically adjust, they help nations absorb economic shocks by responding to each nation's balance of trade with the rest of the world. Over the long term, the value of a nation's currency reflects economic fundamentals, such as global competitiveness and internal monetary policies. In Canada, the central bank and the government believe that the market should determine the value of the Canadian dollar and thus neither body has an exchange rate target. As the global economy realigns after the financial shock of 2008, the Canadian dollar may see persistent periods of high volatility, presenting DND with additional risk during military procurement activities.

Foreign exchange exposure refers to the sensitivity of an organization's cash flows to changes in the exchange rates. Firms who conduct global business while reducing their foreign exchange exposure have access to both operational and financial hedging strategies [1, 2]. If not properly quantified, understood, and managed, foreign exchange risk can expose firms to significant budgetary consequences.[3]. Today, the standard method for reporting financial risk focuses on the value-at-risk (VaR) methodology, which defines the predicted worst-case loss at a specific confidence level over a given period of time [4]. Typically, financial institutions report the daily VaR 95% loss confidence level. VaR provides a quantitative measure of the downside risk associated with the uncertain payoff. Measuring foreign exchange risk represents the first step any organization must take before considering appropriate risk mitigation strategies.

As a matter of policy, DND does not use foreign currency hedging techniques during military acquisition, instead choosing to internalize all foreign exchange risk. DND's policy leads to high uncertainty in final project costs. In this paper, we quantify DND's foreign exchange risk through VaR modelling which allows decision-makers to see more clearly the actual level of risk retained inside departmental spending activities. Our VaR modelling techniques are based on earlier research found in [5–7]. We use linear models for currency pairs with conditional variances to capture the characteristics of each return series. By combining our results, we present DND with a department wide VaR characterization of its total foreign exchange risk. Finally, we demonstrate risk mitigation techniques using vanilla currency options through a counterfactual study [8, 9] on six capital projects with a foreign exchange exposure of $849M USD.

We divide the paper into two parts. In part one we detail the VaR construction and show results over a four month time window. In the second part we describe a counterfactual hedge and present results from the hypothetical hedging activity on six DND capital projects.

2.0 QUANTIFYING THE RISK

While VaR methods apply generally, we use DND's USD capital equipment account as a specific example to demonstrate an application that quantifies DND's foreign exchange risk. Specifically, VaR measures seek to answer the question: "What is the loss such that it will only be exceeded $p \times 100\%$ of the time in the next K trading days?", where $Pr(Loss > VaR) = p$. Figure 1 shows a simplified illustration[1] of the VaR calculation.

[1]Although the return distribution in Figure 1 is normal, in general the return distribution will have additional structure. The VaR interpretation remains the same.

We see that the VaR analysis yields the area under the loss distribution associated with losses that exceed the VaR level, in this case 5%.

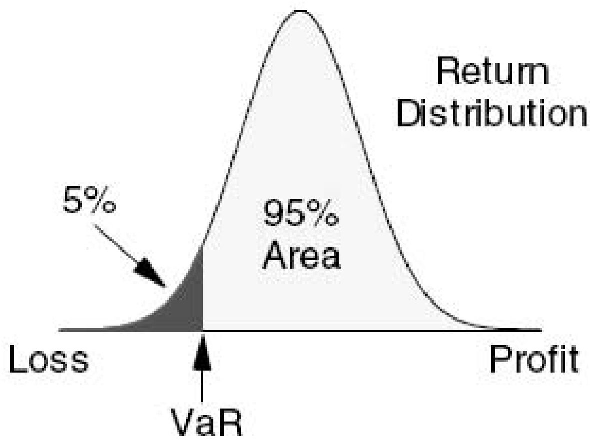

Figure 1: Value-at-Risk (VaR) Example

DND's annual budget planning process begins in March-June of the preceding year, in which budgetary requirements are estimated in Canadian dollars by using forecasted rates for specific currencies. In the overall process, the vast majority of foreign exchange exposure arises from the variance (difference) between the exchange rates observed at the time of the budget obligation (b) as compared to the prevailing exchange rates at the time of liquidation (p). The observed transaction differences, when multiplied by the expenditure, (E), are absorbed within the local DND budgets used to procure the foreign service or equipment. We see that an internal capability that affords a reasonably accurate variance prediction ($b-p$) helps ensure proper management of public funds. The monthly realized budget variance (V) is calculated as the difference between the budget rate (b) and the liquidated rate (p) multiplied by the expenditure (E), i.e.,

$$V = E \times (b - p) \; . \tag{1}$$

Equation 1, in its simplified form, is the basic relationship that defines all VaR calculations for this study. If the liquidated exchange rate is greater than the budget rate, a negative variance (loss) is forecasted and a shortfall is presented to the local budget for which funds must be acquired from other sources.

2.1 A Model for the USD Capital Account

Capital (equipment) transactions represent the largest spending factor in DND's transaction data, easily exceeding other spending amounts by over an order of magnitude. We model the USD Capital Account as a monthly discrete time series in which we assume all transactions accumulate at the end of each month [2].

In figure 2(a), we show the monthly time series data for the USD Capital account. The series contains 156 data point with sample mean and standard deviation \$26.4M and \$$3.05 \times 10^7$ respectively. All values are positive and none are zero.

In modelling the expenditures, we use the expert system *Autobox* – a time series software package that can forecast both univariate and multivariate time series based on Box-Jenkins methods. The software takes a user

[2] In the following, we assume that the reader has some knowledge of time series analysis

specified input series which Autobox corrects for omitted variables that have had historical effects, e.g., pulses, seasonal pulses, level shifts and local time trends. Autobox enhances the forecast model through dummy variables and/or autoregressive memory schemes. Autobox evaluates numerous possible models and reports the model which satisfies all necessity tests to guarantee statistically significant coefficients, and all sufficiency tests to ensure that the residuals are a linear combination of zero-mean, uncorrelated random variables or a zero-mean Gaussian white noise process [10].

In figure 2(a) we see the raw time series data, while in figure 2(b) we see Autobox has defined the structure of the series and has adjusted the values to account for one-time events, highlighted as a single pulse ("P"). The unadjusted series is indicated by the data points without the "P" label. All points viewed by Autobox as pulses must be modelled as increments or reductions on the final series. For example, Autobox found 17 single, non-repeatable pulses in the time series data. Point 48 (December 2002) is the largest pulse found with a magnitude of 203.6×10^6 or, as specified by Autobox, an increment of 182.8×10^6 over the final series. There are no points over which the final series would have to be reduced. The model that best describes the series is an Autoregressive lag 12, i.e.,

$$y_t = 1.18 \times 10^7 + 0.400 y_{t-12} + \sum_{i=1}^{17} c_i x_{t,i} + \varepsilon_t \ , \tag{2}$$

where the $x_{t,i}$ are the 17 single pulses with coefficients c_i, and ε_t represents the error distribution.

Figure 2(b) shows how well the model fits the actual data by superimposing the fit (red) on the actual observations. The coefficient of multiple determination, R^2, for the model has a value of 0.834 which implies that 83.4% of the variance in USD Capital expenditures can be explained by eqn. (2). Since the model is only predictive after lag 12, the first fitted value starts at lag 13. The number of residuals is 144 and the mean squared error is 1.85×10^{14}.

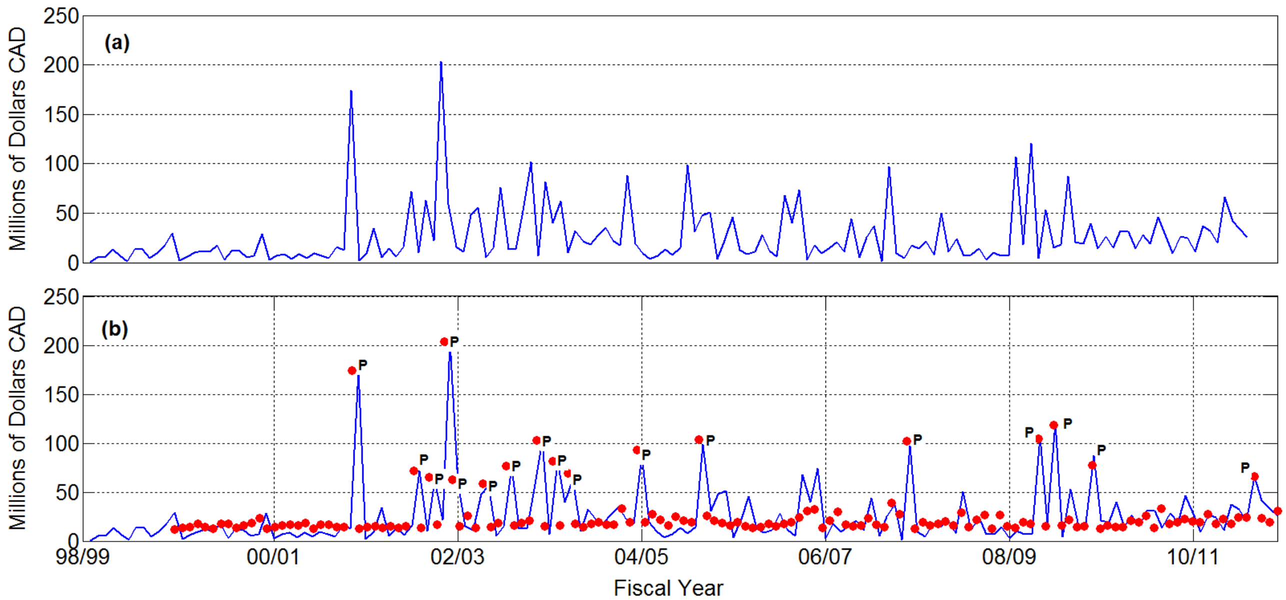

Figure 2: (a) USD Capital expenditures; (b) USD Capital expenditures with model fit

2.2 The Currency Model

Financial returns are known to exhibit certain stylized properties that are common across a wide range of markets and time periods. Examples of these properties include volatility clustering, the leptokurtic[3] distribution of returns, high autocorrelation of squared returns and no autocorrelation of raw returns [11, 12]. The extreme values in the tails of the distribution require a distribution with fatter tails than the normal [13]. Practitioners often model leptokurtic behaviour of the observed returns using Generalized Error or Student's *t* distributions. The degrees of freedom parameter in these models control the fatness of the tails in the fitted model.

Figure 3 shows the time series plots of (a) the daily closing CAD/USD exchange rate, (b) the percent return, and (c) the modelled daily volatility, with significant events in time highlighted by vertical lines. In figure 3 we see the evidence for volatility clustering in the exchange rate data. The highest volatility and price changes occurred at the start of the 2008 recession and the failure of major U.S. financial institutions.

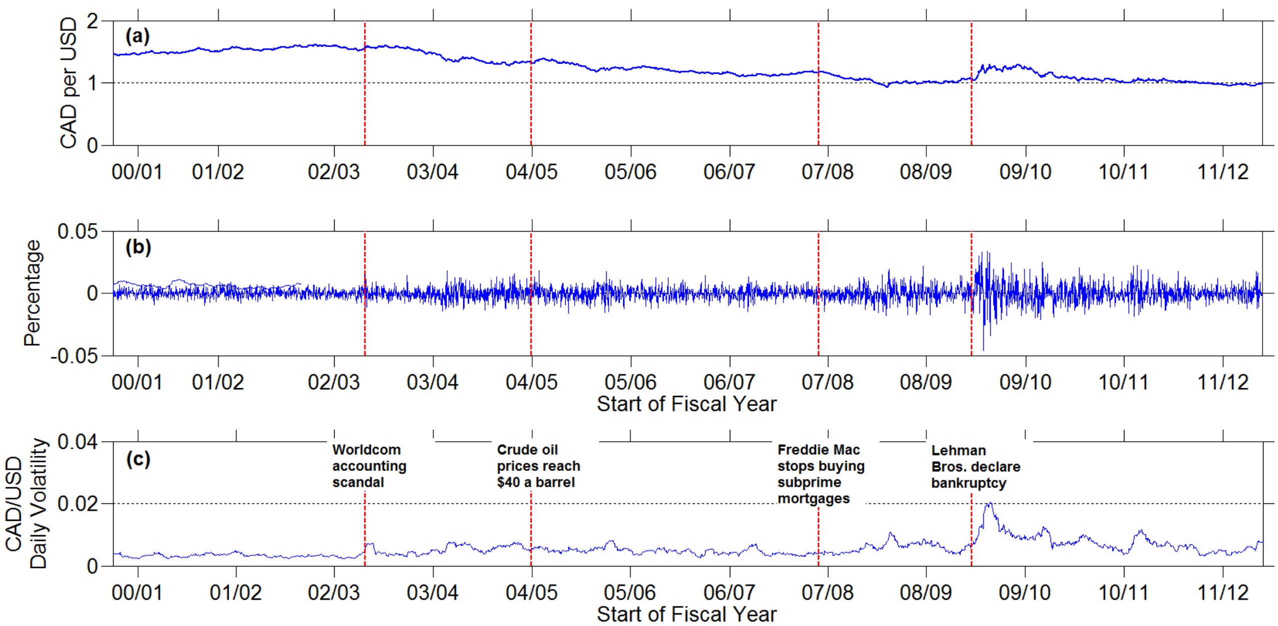

Figure 3: CAD/USD Rates (a), Returns (b) and Volatility (c) from 01 Jan 2000 to 30 Aug 2011

2.2.1 The Univariate GARCH(1,1,) Model

There are two aspects to the problem of calculating a VaR and determining DND's foreign exchange risk. First, we need to model the expenditures for each DND fund, and secondly, we need to develop models for the financial returns series that capture the characteristics of each currency's time-varying volatility and clustering along with its non-normal distributional returns.

GARCH, *Generalized Autoregressive Conditional Heteroskedasticity*[4], models [14, 15] have become important tools in financial time series analysis since their introduction in the 1980s. GARCH modelling recognizes that unconditional stationarity of the return variance does not imply stationarity of the conditional vari-

[3]The condition for a probability density curve to have fatter tails and a higher peak at the mean than the normal distribution.

[4]Autoregressive describes a feedback mechanism that incorporates past observations into the present; Conditional implies a dependence on observation of the immediate past; and, Heteroskedastic refers to time-varying variance or volatility.

ance. Given the lack of autocorrelation in the exchange rate return data, the standard GARCH(p,q) model has the conditional variance, σ_t, depending on p lags of the conditional variance and on q lags of the squared return through the relation,

$$\sigma_t^2 = \omega + \sum_{i=1}^{q} \alpha_i r_{t-i}^2 + \sum_{j=1}^{p} \beta_j \sigma_{t-j}^2 . \tag{3}$$

We model the return distribution as

$$r_t = \sigma_t z_t \quad \text{with } z_t \sim \tilde{t}(d) \, , \tag{4}$$

where z_t is the error term defined by the standardized $t(d)$ distribution with mean zero and unit standard deviation, and the conditional distribution of r_t coincides with the distribution of z_t. We adopt the simplest model, $p = q = 1$, GARCH$(1,1)$, namely,

$$\boxed{\sigma_t^2 = \omega + \alpha\, r_{t-1}^2 + \beta\, \sigma_{t-1}^2 \, .} \tag{5}$$

In eqn. (5), the parameters ω, α, and β denote unknown constants that satisfy the constraints $\omega > 0$, $\alpha \geq 0$, and $\beta \geq 0$, thereby ensuring positivity of the conditional variance, and the additional constraint, $\alpha + \beta < 1$, which is a necessary and sufficient condition to ensure covariance stationarity. Following Christoffersen [16], we estimate the GARCH(1,1) parameters by maximizing the full log-likelihood function using the $\tilde{t}(d)$ distribution,

$$f_{\tilde{t}(d)}(z;d) = \frac{\Gamma((d+1)/2)}{\Gamma(d/2)\sqrt{\pi(d-2)}} (1 + z^2/(d-2))^{-(1+d)/2} \, , \tag{6}$$

where d denotes the degrees of freedom ($d > 2$) and $\Gamma(*)$ is the usual gamma function, and thus the log-likelihood function becomes,

$$\begin{aligned} \ln L &= \sum_{t=1}^{T} \ln(f(r_t;d)) - \sum_{t=1}^{T} \ln(\sigma_t^2)/2 \\ &= T\{\ln(\Gamma((d+1/2)) - \ln(\Gamma(d/2)) - \ln(\pi)/2 - \ln(d-2)/2\} \\ &\quad -\frac{1}{2}\sum_{t=1}^{T}(1+d)\ln(1+(r_t/\sigma_t)^2/(d-2)) - \sum_{t=1}^{T}\ln(\sigma_t^2)/2 \, . \end{aligned} \tag{7}$$

2.3 The Departmental VaR Model

Using GARCH modelling techniques, we develop a VaR tool for departmental financial analysts, allowing them to gauge the implications of conditional exchange rate volatility on local budgets. Our VaR model helps the local budgeting process apportion funding for future contract invoices. Furthermore, our VaR model forecasts the conditional variance out to a maximum horizon of three months (one quarter). Our model of volatility effectively relaxes to the unconditional forecast within our horizon restriction[5]. In this section, we assemble all the times series expenditure models with the GARCH(1,1) model of the exchange rate volatility to produce a complete VaR model for the department.

[5] As stated in [5]:“[The forecast] is not very accurate if the horizon of interest is more than 20 days, since volatility is effectively not accurately forecastable beyond that limit” [17]. Therefore, forecasts up to one quarter should be treated with varying degrees of confidence.

2.3.1 Building the VaR Model

In section 2.1, we highlighted our USD Capital account model, built through time series methods as linear combinations of past expenditures, intervention variables and current values of white noise disturbance terms. We proceed by using Monte Carlo methods to sample from the time series model. Changing the notation slightly to fit eqn. (1) and the Monte Carlo simulation, the forecast expenditures become

$$E^k_{c,a,t+22n} = f_{c,a}(\varepsilon_{t+22n},\, \phi_j y_{t-j}) \,, \tag{8}$$

where the subscripts c,a denote the currency and account (or fund) respectively; $k = 1, \ldots, 10,000$, the number of iterations in the simulation process; $n = 1, \ldots, 12$, the number of months; $j = 1, \ldots, p$, the number of autoregressive terms respectively, with some ϕ taking on zero values. Similarly, based on the results of section 2.2, we can write the forecasted exchange rates as

$$p^k_{c,t+22n} = f_c(\hat{z}_{t+22n},\, \sigma_{t+22n},\, r_{t+22n}) \,, \tag{9}$$

where c, k and n were previously defined.

Given that the budget rates arise through monthly forecasts, but are fixed by external sources, i.e., b_{t+22n}, we can write the relationship that defines the fund variance using eqn. (1) as

$$\left\{ V^k_{c,a,n} = E^k_{c,a,t+22n} \times (b_{c,t+22n} - p^k_{c,t+22n}) \,, k = 1, \ldots, 10,000 \right\}^{12}_{n=1} \,, \tag{10}$$

where $V^k_{c,a,n}$ is the variance for currency c, account a, iteration k and month n, and b, the budget rate, is fixed for each n. The VaR, defined at the 5th percentile of eqn. (10), yields for any n month

$$VaR^{0.05}_{c,a,n} = \left\{ V^k_{c,a,n} \,, k = 1, \ldots, 10,000 \right\}^{0.05} \,. \tag{11}$$

2.3.2 Filtered Historical Simulation

In [5], we use a Filtered Historical Simulation (FHS) for sampling in our Monte Carlo routines as FHS captures all possible values of the historical distribution in price returns over the return history. In particular, observed tail events critical to VaR calculations automatically become part of the Monte Carlo simulation. FHS is a non-parametric method in the sense that the simulation imposes no structure on the distribution of returns [18, 19].

Following [16], we start the process by considering the set of past returns $\{r_{t+1-\tau} : \tau = 1,2,\ldots,T\}$ where T defines the size of the historical CAD/USD dataset. From eqn. (4), we can write the one-day ahead return as the product of the estimated standard deviation and the error term,

$$r_{t+1} = \sigma_{t+1} z_{t+1} \,, \tag{12}$$

where σ_{t+1} is defined through the GARCH variance eqn. (5), and calibrated using 21 years of historical data through the relation

$$\sigma_{t+1} = \left[\omega + \alpha r_t^2 + \beta \sigma_t^2\right]^{1/2} . \tag{13}$$

The parameters (ω, α, β) are specified by maximizing the log-likelihood of eqn. (7). Using the data set $\{r_{t+1-\tau} : \tau = 1,2,\ldots,T\}$ we can now estimate the model parameters and calculate the set of realized standardized returns, $\{\hat{z}_{t+1-\tau} : \tau = 1,2,\ldots,T\}$, defined by

$$\hat{z}_{t+1-\tau} = r_{t+1-\tau}/\sigma_{t+1-\tau} \,, \qquad \text{for } \tau = 1,2,\ldots,T \tag{14}$$

Therefore, given return data up to time t, we can immediately evaluate the GARCH variance and eqn. (13) for time $t+1$. To compute hypothetical returns for tomorrow we draw with replacement from the set of past standardized residuals, $\{\hat{z}_{t+1-\tau} : \tau = 1, 2, \ldots, T\}$, through sampling a discrete uniform distribution of elements consisting of the $\tau = 1, 2, \ldots, T$ standardized returns defined by eqn. (14). The estimated exchange rate, P_{t+1}, becomes[6]

$$P_{t+1} = e^{r_{t+1}} P_t , \tag{15}$$

where P_t is defined as the exchange rate on day t.

2.4 Forecasting Variance and Value-at-Risk

Figure 4 illustrates the VaR Monte Carlo output for CAD/USD forecasted capital account transactions for May 2011 – August 2011 inclusive. The shaded tails correspond to the lower and upper 5% of the distribution respectively. We report the 5th percentile VaR in the upper portion of each chart. Note that in figure 4 a zero budget variance corresponds to DND's internal forecasted budget exchange rate given by b, in eqn. 1. Closer inspection of each chart shows a biased distribution (most notable for the July and August results) as our model identifies a statistically significant non-zero expected variance. Thus, our result shows that internally generated DND budget exchange rates contain subjective opinions about the future strength of the Canadian dollar.

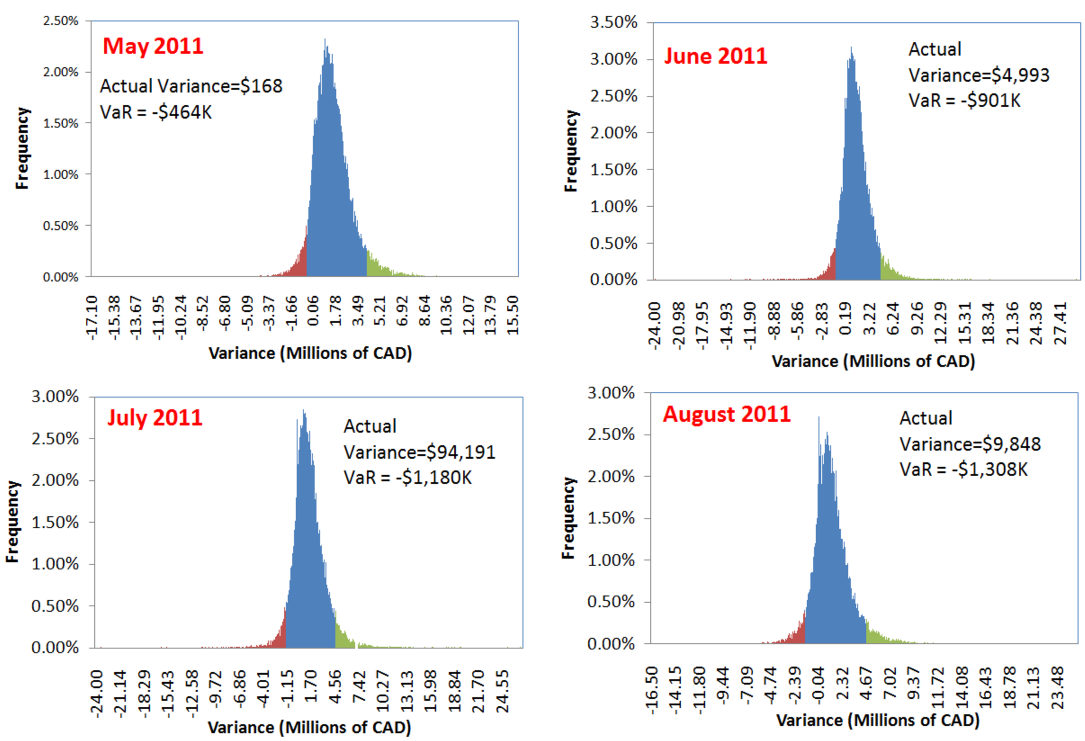

Figure 4: Variance forecasted distributions for CAD/USD capital account from May 2011 through August 2011. Shaded areas to left and right of average correspond to the lower and upper 5% of results respectively.

[6] Recall that we are interesting in calculating a VaR level for financial analysts within DND. Our model does not provide a prediction of the exchange rate over a given horizon. We ensure that the simulated return sequence behaves as a martingale over the horizon of interest.

3.0 MANAGING THE RISK

While our VaR model quantifies the foreign exchange risk in DND budgets, it does not offer solutions or risk mitigation strategies. In an operational sense, quantification represents only half of DND's foreign exchange risk picture. Based on the risk aversion of DND decision makers, we can use our VaR model in conjunction with simple hedging techniques to provide DND with insurance against adverse currency movements.

In 2005, a Canadian study [20] examined the hedging benefits associated with using a derivative known as a forward contract with DND's foreign currency expenditures. Forwards in foreign currency markets oblige the contract holder to buy currency on a predetermined date at a predetermined price. Since the payoff of a forward contract is fixed, the delivery price is readily computed from the riskless interest rates in each nation's market. In a formal sense[7], the delivery price of the currency in the contract represents the expected value of the currency on the delivery date and thus the forward fixes the exchange rate for the contract holder. The study found that DND stands to benefit from a hedging strategy, especially under periods of high volatility.

Forward and futures contracts fill an important role in hedging activity in the financial world. Viewed as an insurance contract, firms obtain a fixed price for a future foreign currency obligation, which can help with internal budgeting processes. Some firms view the rigid nature of the forward and futures contract as inappropriate for their business. Fortunately, the financial world has developed more sophisticated products to fill today's risk management needs.

To facilitate the construction of a hedging strategy for DND, we explore the application of a derivative called an option contract. In the foreign exchange context, an option contract gives the holder the right, but not the obligation, to buy or sell currency at a future date. Since an option contract gives the holder a privilege, the holder must pay the seller of the contract a premium. From the hedger's perspective, the premium is interpreted as the cost of the insurance required to mitigate an unwanted risk. The call option contract contains a strike price which sets the price at which the contract holder can buy the currency from the contract writer on the future date[8]. Figure 5 illustrates the call option's payoff at maturity where K and P denote the strike price and premium respectively.

Consider a simple call option strategy executed by DND. By acquiring a call option on the foreign currency once the foreign obligation becomes known, DND can protect itself from a strengthening foreign currency over the life of the contract. Specifically, if the prevailing exchange rate in the spot market at the time of the currency transaction lies above the strike price then DND will exercise its option, whereas if the prevailing exchange rate proves favourable then DND will let its contact expire worthless and buy the currency in the spot market.

Using a counterfactual approach, we examine the effect of a hypothetical hedge based on call options on six DND projects with significant USD exposure. We make the following assumptions about the transaction process and the foreign exchange markets:

1. We assume project milestones (equipment imports and/or invoices) associated with DND projects listed in Table 1, occur on the last day of the specified month with payment occurring exactly 30 calendar days later.
2. To give an accurate estimate of the price of the call options, we use the Black-Scholes-Merton (BSM) framework. Standard assumptions for applying BSM are:
    a. Markets are complete and efficient (past prices do not contain information on future prices);
    b. There are no transaction costs;

[7] i.e., under the risk neutral measure.
[8] In this study, we only consider vanilla European call options.

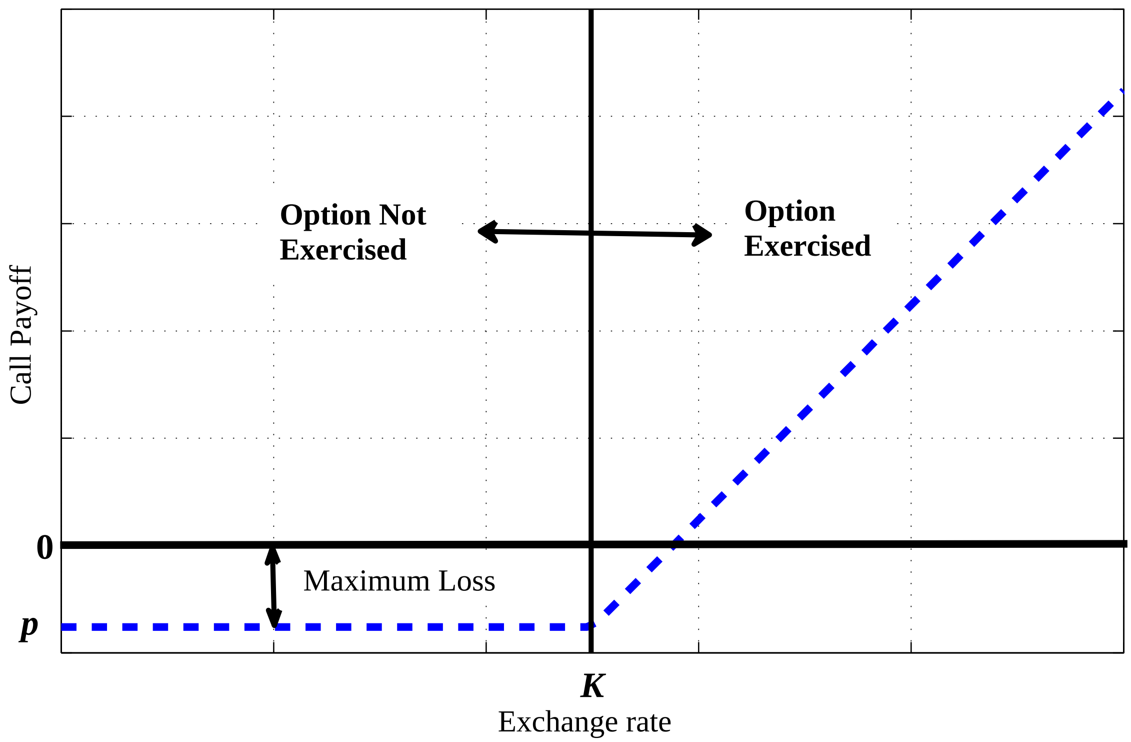

Figure 5: Payoff Scheme for the Call Option

c. A continuum of asset positions are possible;

d. All profits are taxed in a similar way, making tax considerations irrelevant;

e. Market participants can lend and borrow at the same riskless (LIBOR) interest rate;

f. Interest compounds continuously and the rates are constant over the time periods we consider;

g. Market participants use all arbitrage possibilities; and

h. Short selling with full profit is allowed.

3. The exchange rate follows a geometric random walk without drift over the time periods we consider.

4. The volatility of the exchange rate is assumed to be constant at 16.6% for USD over the time periods we consider.

5. We assume that foreign exchange markets never close.

6. There are no liquidity issues/problems.

7. We take all input data on interest rates and exchange rates as of October 30, 2009.

Table 1 specifies the invoiced amounts per month for the six projects: Airlift Capability Project – Strategic (ACP-S) and Tactical (ACP-T), Medium Heavy Lift Helicopter (MHLH), Maritime Helicopter Project (MHP), Medium Support Vehicle System (MSVS), and Strategic Air-To-Air Refuelling (SAAR) [21].

Table 1: Department Projects, Milestones and Payments (in thousands)

Project	Currency	Nov-09	Dec-09	Jan-10	Feb-10	Mar-10	Apr-10	May-10	Jun-10
ACP-S	USD	–	–	–	–	31,694	–	–	–
ACP-T	USD	44,162	63,628	27,989	23,943	17,380	46,821	51,472	116,609
MHLH	USD	43,767	61,804	6,808	–	60,435	6,808	48,389	13,554
MHP	USD	2,500	2,500	2,500	2,500	2,500	2,500	2,500	2,500
MSVS	USD	20,325	18,453	20,873	19,539	22,646	27,998	20,231	13,500

3.1 Analysis and Results

We initiate the USD hedging analysis with the hypothetical purchase of eight call options on October 30, 2009. Each option has a maturity date set on the USD payment date (30 days from the invoiced date in Table 1) and each option's strike price is set at the forward price of the USD obligation at maturity[9]. We model the exchange rate as a geometric random walk and produce 15,000 sample paths over eight simulated months. In each simulated run, each call option is exercised only if the Canadian dollar cost of the USD obligation on the option's maturity date sits above the strike price; otherwise the USD is bought in the spot market. The loss (interpreted as the insurance premium for the hedge) to DND is limited to the purchase price of the call options.

Figure 6 shows the USD hedge results over the entire eight month period for (a) hedging with call options, and (b) no hedge applied. The total USD obligations for all projects over the eight month period is $849M USD and the total sum of the call options amounts to $40 million CAD. We see in Figure 6 that the call options minimize the downside loss and prevent the total cost of the USD obligation from ever exceeding $955 million CAD (cost of the premiums included). On the other hand, in figure 6 panel (b) we see that that there is a 30% chance that the USD obligation will cost more than the maximum of the hedging scenario. In a worst case scenarios, the $849M USD obligation could exceed $1.2B CAD in the spot market on the obligation date, whereas the hedge would produce a savings of more than $245M dollars CAD. The call options behave like an insurance contract on the exchange rate.

[9] Setting the strike at the forward price is called an at-the-money-forward call (ATMF) option

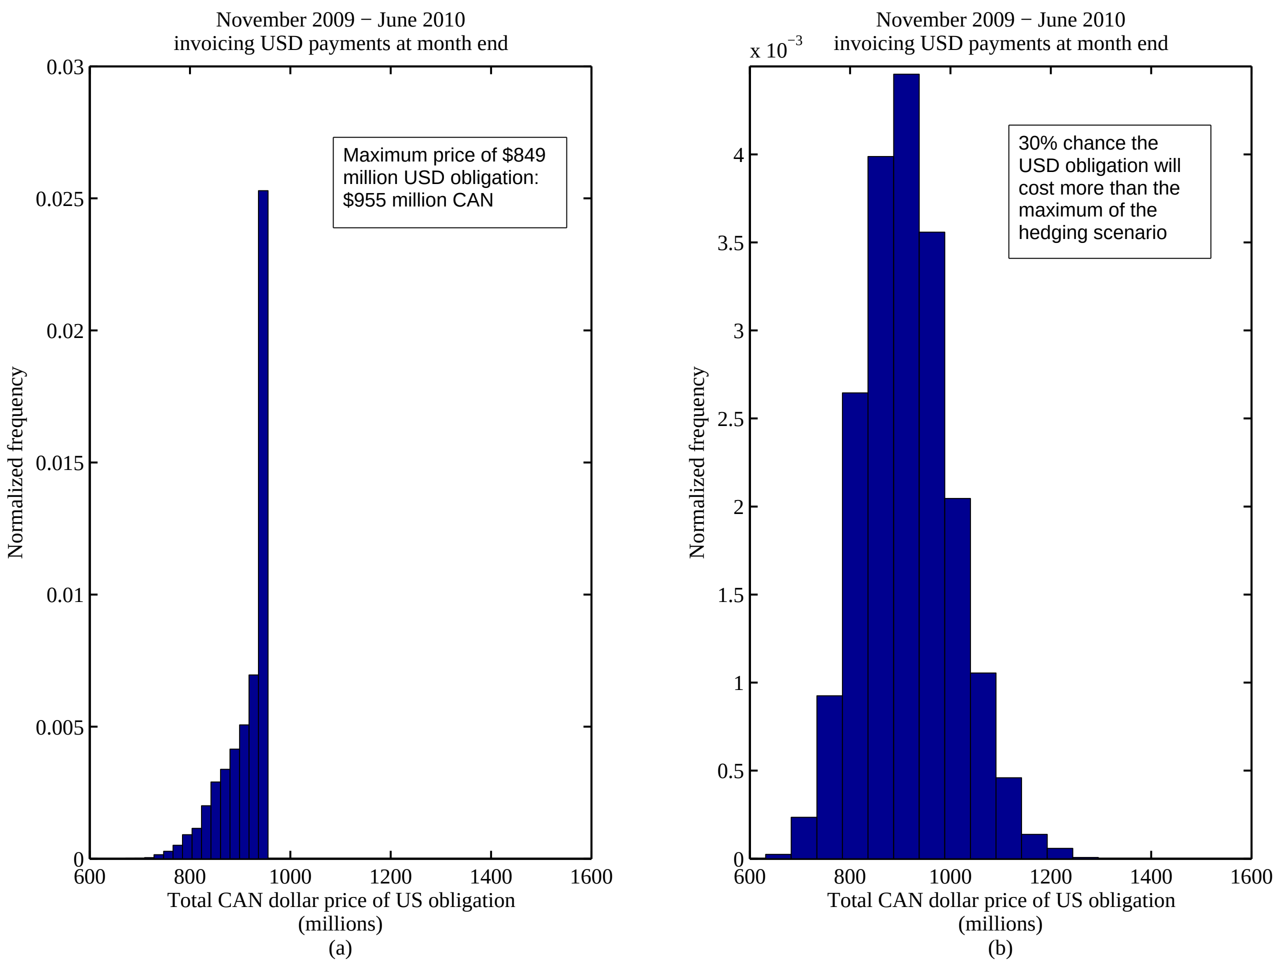

Figure 6: USD vanilla call (a) and no hedge (b) distributions

4.0 CONCLUSIONS

Quantifying foreign exchange risk for DND, given its wide variety of military procurement activities, represents a challenging problem. The VaR models and methodology we have presented addresses the quantification issue and has since been institutionalized in DND as a software application called *FOREX*. DND financial analysts and project managers now have access to a powerful tool that informs them on the likelihood that they will experience a significant shortfall due to adverse currency movements of quarterly planning horizons. As an additional benefit, we have also shown that the VaR model can help inform senior decision makers on the utility of using a simple hedging technique to mitigate unwanted foreign exchange risk. While our simple counterfactual hedge only represents an illustration of derived benefits, DND can use this analysis as a springboard into tailored hedging strategies that would involve other government departments and financial intermediaries.

Given the current volatility of foreign currency markets, and the complex nature of DND's military procurement efforts, understanding foreign exchange risk represents an important step toward successful operations. Using our VaR model in conjunction with risk mitigating strategies can help DND construct a consistent department wide view of risk, which can immeasurably help in planning and budgeting processes. Risk mitigation in areas beyond the control of DND – such as foreign exchange – can help DND deliver timely and effective results and allow DND to concentrate on its primary task of defending Canada and her interests.

REFERENCES

[1] Dominguez, K. M. and Tesar, L. L., “A Reexamination of Exchange-Rate Exposure,” *Papers and Proceedings of the Hundred Thirteenth Annual Meeting of the American Economic Association*, Vol. 91, The American Economic Review, May 2001, pp. 396–399.

[2] Dominguez, K. M. and Tesar, L. L., “Exchange Rate Exposure,” *Journal of International Economics*, Vol. 68, No. 1, February 2006, pp. 188–218.

[3] Groshek, G. M., “Foreign Currency Exposure in the Department of Defense,” *Public Budgeting and Finance*, Vol. 20, No. 4, 2000, pp. 15–35.

[4] RiskMetrics Group, *Risk Management - A Practical Guide*, RiskMetrics Group, 1st ed., August 1999.

[5] Desmier, P., “Estimating Foreign Exchange Exposure in the Department of National Defence,” Technical Report 2006-23, Defence R&D Canada, Centre for Operational Research and Analysis, Jan. 2007.

[6] Desmier, P. E., “Estimating Foreign Exchange Exposure in the Canadian Department of National Defence,” *Journal of Risk*, Vol. 10, No. 4, Summer 2008, pp. 31–68.

[7] Desmier, P., “The Foreign Exchange Exposure Model (FOREX) Expansion,” Technical Memorandum 2009-04, Defence R&D Canada, Centre for Operational Research and Analysis, Feb. 2009.

[8] D.W. Maybury and P.E. Desmier, “Hedging Foreign Exchange Risk in ADM(Mat),” Letter Report 2009-50, Defence R&D Canada, Centre for Operational Research and Analysis, Nov. 2009.

[9] David W. Maybury, “Hedging Foreign Exchange Risk in the Department of National Defence,” Technical Report 2011-09, Defence R&D Canada, Centre for Operational Research and Analysis, Feb. 2011.

[10] Automatic Forecasting Systems (http://www.Autobox.com), 2007.

[11] Cont, R., “Empirical properties of asset returns: stylized facts and statistical issues,” *Quantitative Finance*, Vol. 1, Institute of Physics Publishing, 2001, pp. 223–236.

[12] Taylor, S. J., *Asset Price Dynamics, Volatility, and Prediction*, Princeton University Press, 2005, Chapt. 4, pp. 51–96.

[13] Rachev, S. T., Fabozzi, F. J., and Menn, C., *Fat-Tailed and Skewed Asset Return Distributions : Implications for Risk Management, Portfolio Selection, and Option Pricing*, Wiley, 2005.

[14] Bollerslev, T., “Generalized Autoregressive Conditional Heteroskedasticity,” *Journal of Econometrics*, Vol. 31, No. 3, 1986, pp. 307–327.

[15] Engle, R., “Autoregressive Conditional Heteroskedasticity with Estimates of the Variance of United Kingdom Inflation,” *Econometrica*, Vol. 50, No. 4, July 1982, pp. 987–1007.

[16] Christoffersen, P. F., *Elements of Financial Risk Management*, Academic Press, 2003, Chapter 4.

[17] Christoffersen, P. F. and Diebold, F. X., “How Relevant is Volatility Forecasting for Financial Risk Management?” NBER Working Papers 6844, National Bureau of Economic Research, Inc, Dec. 1998.

[18] Barone-Adesi, G., Giannopoulous, K., and Vosper, L., "VaR without Correlations for nonlinear Portfolios," *Journal of Futures Markets*, Vol. 19, April 1999, pp. 583–602.

[19] Barone-Adesi, G., Giannopoulous, K., and Vosper, L., "Filtering Historical Simulation. Backtest Analysis," *Manuscript*, March 2000.

[20] N. Essaddam, C.H. Bucar, and R.A. Groves, "The Hidden Cost of Operations: Foreign Exchange Risk Management in the Canadian Department of National Defence," *Canadian Military Journal*, Vol. 6, No. 4, Winter 2005, pp. 61–66.

[21] A. Cholette, ADM(Mat)/DMPS 4-2, "Private Communication," 04 November 2009.