Extending Value at Risk to a Corporate Setting

An Application to Fonterra Cooperative

A thesis presented in partial fulfilment of the requirements for the degree of
Master of Business Studies in Finance

at Massey University

Palmerston North, New Zealand

Jared M Cahan

13 December, 2005

1

Abstract

This paper demonstrates the development and application of a corporate Value-at-Risk model. Using the RiskMetrics Group's CorporateMetrics as a starting point we show how the framework can be modified to meet the specific needs of Fonterra Cooperative, a major New Zealand dairy exporter. We develop a Monte-Carlo simulation model that uses univariate ARIMA and multivariate Vector Error Correction (VECM) forecast models to estimate the Value-at-Risk on Fonterra Group Treasury's interest rate and FX hedge portfolio over a 15-month period.

2

Acknowledgements

Confidentiality

The commercial sensitivity of certain Fonterra specific information necessitates its removal from the publicly available version of this thesis.

Appendix D, containing information regarding the Group Treasury's FX and interest rate hedging strategies and graphs of proprietary basic commodity price (BCP) data series, has been removed from this report. Requests to access this information should be directed to the author.

It is our intention that the removal of these details does not detract from the academic rigor of the analyses. Rather than focus on the specific model developed for Fonterra, we emphasise the development of a general long-term corporate VaR model.

Contents

1 Introduction

In today's environment of volatile financial markets and rapid globalisation, risk management has become a crucial element in the success of any organisation. Over the last three decades, high volatility as a result of floating exchange rates, interest rate instability and fluctuating commodity prices, coupled with the subsequent explosion in financial derivatives has led to the need for a simple, easy to understand risk measure (Dowd, 1998). The well publicised and oft quoted collapse of such financial giants as Barings Bank, Metallgesellschaft and California's Orange County only serve to hammer home the need for some measure to quantify firm-wide risk that is accessible both to management and shareholders.

One such measure is value at risk (VaR). Since its inception in the late eighties and early nineties, value at risk has become the de facto standard for risk measurement in financial institutions (Jorion, 1996). Value at risk allows firm-wide market risk to be summarised as a single number, expressed as the maximum loss expected over a fixed period at a given probability level. Since VaR was formalised in a practical sense by JP Morgan's RiskMetrics in 1994, an expansive body of academic literature has developed. However, the majority of such research has focused on VaR applied to financial institutions as a measure of the short-term potential loss on a portfolio of financial instruments (Lee, 1999). Recently there has been increased interest in applying a VaR type measure within a corporate environment. In such a setting VaR is complicated by the inclusion of business risk and the need to work to a long time horizon. Because of these added complications, applying a VaR type measure within a corporate setting is a much more firm-specific exercise than traditional portfolio VaR.

Although the recent creation of the CorporateMetrics framework by the RiskMetrics group seeks to simplify this task, the onus remains on the company to identify the unique risks facing it and to link these to financial performance measures. In this paper we demonstrate how the CorporateMetrics approach can be tailored to fit the needs of a specific company. We find that, given limited data and a long-term forecast horizon, the forecast framework utilised by CorporateMetrics is inappropriate for our project. Instead, we focus on one particular long-run, multivariate forecast model -

namely the Vector Error Correction Model (VECM) - described in LongRun, CorporateMetric's companion forecasting technical document (Kim, Malz and Mina, 1999). We extend the model by allowing the user to choose to use normally distributed or bootstrapped residuals as the stochastic error term within the Monte-Carlo simulation. In addition, we also fit univariate ARIMA models to each of the data series. These provide a base case against which we can determine the extent to which the inclusion of additional variables and use of cointegrating relationships in the forecast model improve out-of-sample forecast accuracy. The flexibility in model choice allows end users to gain at least a qualitative feel for the level of model risk inherent in the simulation.

The company used as an example throughout this project is Fonterra Cooperative. Fonterra is the world's second largest exporter of dairy products (after Nestlé) with an annual turnover of USD 6.8 billion, single-handedly generating a staggering 20% of New Zealand's export receipts and 7% of the GDP. As a result, understanding the impact of changes in market rates on Fonterra's profitability is vital - not just for the 13,000 farmer shareholders who collectively own the company - but also for the New Zealand economy as a whole. Fonterra has recognised the need to quantify its exposure to changes in the NZDUSD exchange rate, NZ interest rates and world commodity prices in a single firm-wide measure appropriate for reporting at all levels of management. VaR[1] is the natural choice.

The aim of this project is two-fold. Firstly, we seek to provide a comprehensive literature review of the current state of VaR. Over the past decade VaR has become one of the most researched topics in risk management literature and as such a vast number of alternative techniques have been developed. Given the plethora of VaR techniques we think that it is important to outline the numerous methods available, with particular emphasis on the assumptions and drawbacks of each. Secondly, we develop a VaR model for Fonterra Cooperative and implement it in VBA/Excel. This model serves to illustrate the problems of developing a VaR model within the

---

[1] We use the term VaR loosely throughout this project to refer to any VaR type risk measure such as cash-flow at risk (CFaR) and earnings per share at risk (EPSaR).

constraints of a real-world corporate setting. Many of the techniques described in the first section become unrealistic within a practical setting. Our work serves to demonstrate the assumptions and simplifications that must be made – and, more importantly, the implications of them - when implementing a corporate VaR model.

The primary complication that arises in applying VaR within in corporate setting is the addition of business risk to the risk management process. In a financial institution the goal is to measure the uncertainty surrounding the market value of a portfolio of financial instruments, so market risk is the primary concern. In a corporate, however, market risk and business risk are inextricably linked.

For example, consider an appreciation of the NZD against the USD. A trading institution would be primarily concerned with the impact of the changing FX rate on the value of the financial instruments in its various trading portfolios. In contrast, an exporter such as Fonterra would be concerned not just with the effect such an appreciation would have on its FX hedges (market risk), but also with the decreased competitiveness and subsequent potential reduction in volume of its exports to the US (business risk).

However, the inclusion of business risk, which requires the development of firm-wide models, is an extension for future research. This project analyses the market risks facing Fonterra. This risk is in the form of a gain or loss on the Group Treasury hedge portfolio (consisting of foreign exchange and interest rate hedge instruments) over a long-term[2] horizon. As such the project provides an in-depth look at the impact of market rate changes on Fonterra's profitability, while also providing a natural platform from which to launch later investigations into the impact of business risk.

We also place little emphasis on the actual VaR numbers obtained. Given the fluid nature of Fonterra's hedge book we are content to outline the method behind the calculations without attempting to model all the intricacies of the hedge book. The

---

[2] Throughout this report we take long-term as referring to periods greater than one year. Given that the usual VaR horizon of less than 3 month, this seems appropriate terminology.

primary aim of the second half of the project is to develop a working model for use by Fonterra Cooperative. As such, we are limited in the amount of historical backtesting and the extent to which we can verify the concordance of simulated and historical data. An in-depth analysis of the robustness of the model is left for further research. The the contribution of this project is to illustrate the process by which a VaR model can be developed within a corporate setting and provide Fonterra with a viable model for their risk management needs.

2 Literature Review

A wide body of academic VaR literature exists and indeed VaR is one of the most researched topics in modern financial risk management. The bulk of this research however is focused on traditional VaR applied to the trading portfolios of financial institutions. It is only more recently that VaR has began to be extended to other risk management applications including, credit risk and corporate VaR. Nonetheless, because the concepts behind VaR are so similar across different applications, particularly with respect to the mathematics and performance testing, the VaR literature provides a sound platform for beginning an analysis of corporate VaR.

2.1 VaR concepts

For a broad, albeit somewhat outdated, overview of the basic concept of VaR see Dowd (1998), Jorion (1996) or Pan (1997). In simple terms VaR is the maximum amount we would expect to lose over a given time horizon at a given confidence level[3].

VaR has a number of oft-cited advantages. The first is that VaR provides a common yardstick across different risk factors (Dowd, 1998). A VaR number on an equity position is directly comparable to a VaR figure on a derivatives postion, thus VaR provides a consistent and comparable risk measure. In addition VaR summarises complex risk exposures in a single concrete number that is accessible and understandable to all levels of the organisation. With increasingly esoteric derivative products, that are often far beyond the understanding of senior management, the importance of a clear and concise measure of risk should not be underestimated.

Of course, VaR also has its limitations. Dowd (1998) points out several primary concerns. Foremost of these is that any VaR system is inherently backward looking; that is, past losses are used to forecast future losses. The assumption that past

---

[3] For example, a 95% VaR of 18 million over a time horizon of one week can be interpreted as follows: the maximum amount we would expect to lose 95% of the time on the given portfolio over a one week period is 18 million.

behaviour is a good predictor of the future is particularly dubious where VaR is concerned since by definition we are concerned with abnormal events. Extreme events, such as a total market crash, may not even appear within the historical period used to estimate the VaR model, and yet it is these events with which we should be most concerned. One needs only to look at the fall of Long Term Capital Management to see vividly illustrated the consequences of overlooking abnormal events. Secondly any VaR number is entirely contigent upon the models used and the assumptions made. This should come as no surprise, these are fundamental limitations of any measurement framework. The problem with VaR is that in presenting a single, seemingly indisputable number, it becomes easy to forget that the apparently concrete number is the result a multitude of assumptions and approximations. Dowd (1998) suggests that the answer is not to discard VaR, but to be clearly aware of the assumptions upon which the particular VaR model is built.

Essentially VaR methodology can be broken down into three main categories, namely the variance-covariance method, the historical simulation method and the Monte-Carlo simulation method. We will look at each method in turn.

2.2 Variance-covariance

The variance-covariance method is arguably the most simplistic approach to computing VaR. This approach is based on the assumption that asset returns are normally distributed. Such an approach has two key advantages. First, calculating a VaR number from a normal distribution is a trivial exercise (see, for example, Dowd (1998) or Jorion (1996)) and second, it is very informative[4]. However, the assumption of normality is dubious at best. Much research exists indicating that this is not the case, for example, often asset returns demonstrate excess kurtosis or "fat tails". Similarly the method assumes returns are a linear function of normal risk variables which is clearly not the case for instruments such as options.

Current research focuses primarily on addressing the dual problems of non-normality and non-linearity. One path of research looks at accommodating non-linearity through

---

[4] That is, obtaining a VaR figure for a given confidence level and holding period tells us the VaR for all combinations of holding periods and confidence level (Dowd, 1998, p. 64).

higher-order approximations such as the Delta-Gamma (second order) methods. Wilson (1994) proposes finding VaR as the solution to a corresponding optimisation problem using second order approximations. However, Wilson's approach can be very inaccurate. As a result Jamshidian and Zhu (1996) suggest using market delta and gamma to simplify calculations while Fallon (1996) applies a non-linear econometric procedure to estimate non-linear relationships. Several papers also examine the use of higher moments in VaR calculation such as Zangari who adjusts the confidence level for skewness (1996d) and investigates estimating the first four moments of the returns distribution and then fitting these to a known distribution (1996c).

Other research focuses on the problem of fat-tails. Zangari (1996a) and Venkataraman (1997) apply a normal mixture approach while Zangari (1996b) proposes the use of a generalised error distribution.

2.3 Historical simulation

The historical simulation (HS) methodology makes the implicit assumption that the historical distribution of asset returns are a good proxy for the distribution of asset returns in the next period (Dowd, 1998). Under the broad banner of HS a plethora of methods have been developed to improve the basic methodology.

Hendricks (1996) provides a comprehensive assessment of basic historical models. More recent innovations include Hull and White (1998), who propose the use of using GARCH models to scale the historical returns used in forecasting to better reflect current market volatility. Other examples of more recent research include the use of kernel quantile estimators in HS (Butler and Schachter, 1998), and the use of a generalised error distribution model for weighting historical returns (Lin and Chang-Cheng, 2003).

Contemporary research in the area revolves around computationally intensive algorithms for computing HS VaR for large portfolios of assets. See, for example, Audreno and Barone-Adesi (2002) who use a functional gradient descent algorithm to compute VaR for a large portfolio of stocks. Of course it is vital to understand the

shortcomings of each method. Pritsker (2001) carries out a comprehensive review of some of the currently favoured HS techniques and finds a number of shortcomings.

2.4 Monte-Carlo simulation

A Monte Carlo (MC) simulation approach to VaR involves estimating VaR from simulation results obtained from mathematical forecast models. The techniques tend to follow the same general methodology, namely: select a model to describe the future behaviour of the variable (e.g. GARCH, random walk, Markov chain, etc.), estimate the parameters for the model from historical data, construct price paths using random numbers to obtain a distribution at each point in time and use this to estimate a VaR number.

The most well known, and certainly the most influential, Monte-Carlo implementation is that of RiskMetrics (see Morgan Guaranty Trust Company (1996) and Mina and Xiao (2001)). The RiskMetrics approach forms the cornerstone of VaR calculations within many financial institutions[5] and it is worth briefly reviewing it here. The RiskMetrics approach uses a random walk model for the log-price of each financial instrument, that is the log-price at time t is given by

$$p_t = \mu + p_{t-1} + \sigma_t \varepsilon_t$$

where $p_t = \ln(P_t)$. Since $r_t = \ln\left(\dfrac{P_t}{P_{t-1}}\right) = \ln(P_t) - \ln(P_{t-1}) = p_t - p_{t-1}$ it follows that $r_t = \mu + \sigma_t \varepsilon_t$. RiskMetrics makes the further assumption that the mean return, μ, is zero. An implication of the model is that the variance is a linear function of time, thus the daily standard deviation can be scaled up by the square root of time to obtain the standard deviation over the longer time period. Thus the price for asset i at time t is given by

$$P_t^{(i)} = P_0^{(i)} e^{\sigma_i \sqrt{t} \varepsilon_{i,t}}, \qquad \varepsilon_{i,t} \sim N(0,1).$$

---

[5] Indeed, the introduction of Basel Committee on Banking Supervision requirements in 1996 requires standardised VaR calculations.

14

At each time horizon the error terms for each asset are drawn from a multivariate normal distribution (MVN)[6] taking into account the historical correlations between the variables. The variance is allowed to change, at each point in time the variance and correlations are forecast using exponentially weighted moving averages[7]. The fundamental assumption of the RiskMetrics model is that returns are conditionally normally distributed, that is the returns are normally distributed conditional on the variance. Appendix A and B of the *RiskMetrics Technical Document* (Morgan Guaranty Trust Company (1996, pp. 227-242)) explores the assumptions behind the RiskMetrics model and how they might be relaxed. Mina and Xiao (2001) revisit the RiskMetrics model from the standpoint of the several additional years of development and discussion.

Since RiskMetrics much of VaR research has focused on the Monte-Carlo approach. In general such simulations are computationally intensive and almost all research in the area focuses on reducing computational complexity. Pritsker (1996) proposes a "grid MC" whereby the portfolio is mapped onto a grid of factors for which price paths are computed reducing the number of simulations. Clewlow and Caverhill (1994) and Boyle, Broadie and Glasserman (1997) achieve success in reducing computation time through the use of variance reduction techniques.

Quasi Monte Carlo (QMC) simulation is also the subject of much research. Instead of using pseudo-random numbers, which are often clustered close together, QMC takes numbers that are evenly and uniformly distributed in the domain[8]. Dowd (1998) lists a multitude of papers making a contribution in this area; including Boyle et al. (1997), Owen and Tavella (1996), Brotherton (1994), Joy, Boyle and Tan (1995), Paskov and Traub (1995) and Papageorgiou and Traub (1996) among others. Further advances in

---

[6] That is $\varepsilon_t = MVN(0, \Sigma_t)$ where $\varepsilon_t = [\varepsilon_{1,t} \quad \varepsilon_{2,t} \quad ... \quad \varepsilon_{N,t}]$ and Σ_t is the variance-covariance matrix and time t.

[7] The *RiskMetrics Technical Document* also explores alternative methods of forecasting variances and covariances such as GARCH.

[8] This avoids the inefficency that results when two random numbers very close together are selected leading to almost identical price paths. Such a situation effectively means that little information is added for an additional computation step.

efficiency are made by Moskowitz and Caflisch (1995), Neiderreiter and Xing (1995) and Tezuka (1994).

2.5   Validation and backtesting

Any VaR model is subject to both random error and possible systematic bias. As a result the backtesting and validation of a VaR model against historical data is essential. A number of VaR evaluation methods have developed over time in an attempt to improve on earlier methods, all of which have specific drawbacks. Early methods developed by Kupiec (1995) are based on a binomial assumption and look at whether VaR estimates exhibited correct unconditional (average) coverage. However such methods fail to test whether a model is accurate at any given point in time, i.e. they do not test for correct conditional coverage.

Christoffersen (1998) introduces an interval forecast method which tests for correct unconditional and conditional coverage. However, Crnkovic and Drachman (1996) and more recently Berkowitz (2001) point out that interval forecasts, such as that of Christoffersen (1998) do not utilise all available information (since they look only at exceptions and do not backcheck the model at every percentile) and thus require large samples to obtain sufficient data for backtesting. Crnkovic and Drachman (1996) and Berkowitz (2001) develop density evaluation techniques in an attempt to overcome this problem.

All the aforementioned techniques are based on a hypothesis testing framework and generally have relatively low statistical power. An alternative is presented by Lopez (1996), who proposes a method based on standard probability forecast evaluation techniques instead.

2.6   Corporate VaR

There is a small, but growing, body of literature on VaR techniques applied in a corporate setting. Techniques such as CFaR applied in a corporate setting are increasingly being discussed in practitioner publications such as *Financial Engineering News* and *International Finance & Treasury*. See, for example, Ripp and Cohn (2001), Pitt (2001) and Stein, LaGattuta and Youngen (2000). The CFaR

concept is originally attributed to the National Economic Research Associates (NERA), a consultancy firm who have produced a number of publications on its use[9].

By far the most comprehensive treatment thus far of the concept is the RiskMetrics Group's CorporateMetrics (Lee, 1999). CorporateMetrics is a comprehensive framework for extending the VaR concept to corporate cash flows or earnings. Along with LongRun (Kim, Malz and Mina, 1999), the companion long-term forecasting framework, CorporateMetrics represents a foundation for extending VaR to corporates. It is particularly relevant here because it highlights many of the problems inherent in translating VaR to a corporate setting. For example, in contrast to a financial institution, where market risk (exposure to foreign exchange, interest rate and commodity price changes) is the primary concern, corporates are exposed to both market risk and the business risk[10] associated with their specific operations. In addition, the level of market risk is a function of business risk, making risk management a far more complex task (Lee, 1999). For example, in Fonterra's case a change in the NZDUSD exchange rate will not only lead to a gain/loss due to its foreign exchange hedges (market risk), but could also effect sales volume in the US as the relative competitiveness of its products in the US changes (business risk).

CorporateMetrics does not attempt to model business risk, rather it seeks to provide a consistent framework for quantifying the impact of market risk on the operation of a corporate. The process is divided into five distinct steps (as summarised in Lee (1999, p. 28-29)). The first step is to decide on the metric specification, that is, the risk measure to calculate (CFaR, EaR, VaR, etc.), the time horizon and the confidence level. The next step involves the development of "exposure maps" that link the changes in market rates through to the financial result of interest. These exposure

---

[9] For example Guth (2002).

[10] Lee (1999, p. 5) defines business risk as the "uncertainty of future financial returns related to the business decisions that companies make and to the business environment in which companies operate". In contrast market risk is defined as "the uncertainty of future financial results that arises from market rate changes".

maps might take the form of a simple formula or a complex set of forecast financial statements[11].

Once the exposure maps are in place, long-run forecasting methodologies are applied to forecast the distribution of each market rate at each time horizon. CorporateMetrics makes use of two forecast methods. The first method uses current market information such as forward and future prices and known relationships between them to forecast the future prices. Implicit in the method is the assumption that market expectations of future prices are embedded in current prices. In cases where there is a lack of relevant market data, such as for forecasts greater than two years, CorporateMetrics uses econometric time series models in an attempt to capture the long-run relationship between variables.

Price paths are then obtained by sampling from the forecast distributions at each time horizon, taking into account the historical correlation between variables. In the fourth step each scenario is iteratively input into the exposure maps to generate a distribution of the financial results. Finally, in step five the simulated distribution of the financial result is used to calculate the desired VaR figure.

2.7 Long-run forecasting

Given the pivotal role of forecasting in economics and finance, it is not surprising that the long-run forecasting of economic and financial time-series is a heavily researched area. In this section we focus on alternative models that have been proposed for the forecasting of foreign exchange rates, interest rates and commodity prices.

2.7.1 FX rate forecasting

Since introduction of flexible exchange rates in the mid-1980s the forecasting of foreign exchange rates has become an increasingly important topic. Much of this research has centred around the question of whether floating exchange rates are non-stationary or whether they exhibit mean reversion in the long-run, an issue of vital

---

[11] For example, the NZDUSD exchange rate may be linked to the final Fonterra milk payout figure via a complex series of forecast financial statements. Simulated NZDUSD rates fed into the statements as inputs would then feed through to impact on the milk payout figure.

importance when forecasting. The theory of Purchasing Power Parity (PPP) suggests that, in the long-run at least, flexible exchange rates will tend to be mean reverting (see Rogoff (1996) and Rose (1996)). Similarly, the argument that countries undertake monetary policy with the aim of maintaining a stable exchange rate (Sweeney, 2001) can be raised as a reason for expecting mean reversion. Yet despite these economic arguments many authors report a lack of mean reversion. Roll (1979), Adler and Lehmann (1983) and Pigott and Sweeney (1985) all conclude that exchange rates under the managed float regime are essentially a random walk (Jorion and Sweeney, 1996). This is somewhat surprising since, as Jorion and Sweeney explain, "the random walk hypothesis has the disturbing implication that shocks to the real exchange rate are never reversed, which implies that there is no tendency for Purchasing Power Parity (PPP), a fundamental building block of exchange rate models" (1996, p. 536).

Given the lack of concordance of earlier studies with economic theory, a number of more recent studies look further into the issue. Huizinga (1987) finds that real exchange rates slowly revert to a wandering mean while Abuaf and Jorion (1990) reject the null of a unit root, albeit only over periods of many decades that span multiple exchange rate regimes. Jorion and Sweeney (1996) find strong evidence that exchange rates are mean reverting over the 1973-93 period using the multivariate framework developed in Abuaf and Jorion (1990). Rather than using the usual OLS regression test for stationarity[12] on each country individually, Jorion and Sweeney estimate the SUR system

$$R_{i,t} = \alpha_i + \beta R_{i,t-1} + u_{i,t}$$

---

[12] That is, for each country estimate $R_{i,t} = \alpha_i + \beta_i R_{i,t-1} + \mu_{i,t}$ where $R_{i,t}$ is the real exchange rate. Under the null of a random walk $\beta_i = 1$ while under the alternative of mean reversion $\beta_i < 1$. Because of the downward bias in β_i under the null, Dickey-Fuller critical values are used rather than the normal t-statistic (Jorion and Sweeney, 1996).

where $R_{i,t}$ is the real exchange rate and β is constrained to be equal across all regressions[13]. Abuaf and Jorion (1990) argue that the estimation of a system of equations should lead to more powerful tests and under the Jorion and Sweeney framework fail to reject the null of no unit root. They go on to examine the forecast ability of four alternative models and find that while in the short term a random walk model is about as good as the mean reversion model, over longer horizons the mean reversion model is substantially more accurate.

Siddique and Sweeney (1998) provide a more in-depth look at the forecast performance of the Jorion and Sweeney (1996) mean reversion model relative to a random walk model. In a 12-month ahead forecast with a 10-year estimation period they find that the random walk model never outperforms the mean reversion model. More recently Sweeney (2001) looks at mean reversion in G-10 nominal FX rates and finds further evidence of mean reversion. Again, he goes on to show that in out-of-sample forecasts mean reverting models outperform random walks on average. It is worth noting at this stage that for our purposes whether or not exchange rates are stationary is not important. What is important is how well the alternative models perform in out-of-sample tests. Campbell and Perron (1991) argue that the choice between treating a series as stationary or not should not necessarily depend on the actual stationarity of the series, but rather on the application in mind.

2.7.2 Interest rate forecasting

Having examined various long-run forecasting methodologies for exchange rates we now turn our attention to similar models for interest rates. Given the regulatory environment in which interest rates are determined we would expect such rates to exhibit mean reversion in the long-run.

Extensive literature exists in the area of interest rate models, for example Chan, Karolyi, Longstaff and Sanders (1992a) point to the following papers: Merton (1973), Brennan and Schwartz (1977, 1979, 1980), Vasicek (1977), Dothan (1978), Cox,

---

[13] In this model $R_i^* = \dfrac{a_i}{1-\beta}$ is the long-run equilibrium rate and (1-β) is the common speed of adjustment to equilibrium across the exchange rates.

Ingersoll and Ross (1980, 1985), Constantinides and Ingersoll (1984), Schaefer and Schwartz (1984), Sundaresan (1984), Feldman (1989), Longstaff (1989a), Hull and White (1990), Black and Karasinski (1991) and Longstaff and Schwartz (1992).

Chan et al. (1992a) bring together much of the earlier work in interest rate modelling. They consider eight alternative continuous time short-term interest rate models, all of which are shown to be special cases of a discretised approximation to a single stochastic differential equation

$$dr = (\alpha + \beta r)dt + \sigma r^{\gamma} dZ$$

where r is the interest rate, t is time and dZ is a Weiner process. Eight models well established in literature are nested within this differential equation:

1. Merton (1973)	$dr = \alpha dt + \sigma dZ$	
2. Vasicek (1977)	$dr = (\alpha + \beta r)dt + \sigma dZ$	
3. Cox, Ingersoll, Ross (1985)	$dr = (\alpha + \beta r)dt + \sigma r^{1/2} dZ$	
4. Dothan (1978)	$dr = \sigma r dZ$	
5. Geometric Brownian Motion	$dr = \beta r dt + \sigma r dZ$	
6. Brennan, Schwartz (1979)	$dr = (\alpha + \beta r)dt + \sigma r dZ$	
7. Cox, Ingersoll, Ross (1980)	$dr = \sigma r^{3/2} dZ$	
8. Constant Elasticity of Variance	$dr = \beta r dt + \sigma r^{\gamma} dZ$	

The differential equation is discretised as

$$r_{t+1} - r_t = \alpha + \beta r_t + \varepsilon_t$$

and the parameters are estimated by the Generalised Method of Moments (GMM) of Hansen (1982). They then test the explanatory power of the various models as well as the forecast performance on US interest rate data. Overall they find that the models that best capture the dynamics of short-term interest rates are those that allow the volatility of interest rate changes to be sensitive to the level of the risk free rate. Tse

21

(1995) extends the methodology of Chan et al. (1992a) to 11 different countries and shows that no single model satisfactorily describes the stochastic structure of interest rates for all countries.

2.7.3 Commodity price forecasting

The literature on long-run modelling of commodity prices is much less developed. Tomek (2000) provides an excellent overview of the state of commodity price modelling, illustrating and commenting on problems – some of which are unique to commodity prices – that arise. A rather haphazard array of different models has been applied to commodity prices, necessitated to a degree by the non-uniformity of the various commodities. Some look at time series approaches such ARIMA and GARCH models while others apply general equilibrium models built on an economic framework. For example, Yang and Brorsen (1992) fit Mixed Jump Diffusion, GARCH and Deterministic Chaos models to the daily prices of seven agricultural commodities and show that the GARCH model is preferred for explaining the dynamics of daily prices. Myers (1994) and Myers and Tomek (1993) look at commodity prices from an econometric standpoint. In contrast Deaton and Laroque (1991) apply a rational expectations competitive storage model to 13 commodities in an attempt to explain the high level of autocorrelation and occasional violent price explosions often observed in commodity price series. Jesse (2002) proposes a systematic, user-friendly process by which future milk prices can be forecast. The method has a practical focus and is laid out as a step-by-step in which milk prices are derived from forecasts of milk production levels. The papers mentioned illustrate the widely different viewpoints from which commodity price research is approached, one of the main reasons for the lack of standard forecasting methodologies.

2.7.4 Time series models

We have looked at models developed specifically for a particular economic time series. Before closing we briefly touch on some more generalised statistical time series modelling methods. An advantage or pure time series models is that they can be applied consistently across a wide range of time series without the need to specifically consider the dynamics or underlying economic rationale behind each individual series. Box and Jenkins (1976) describe the traditional ARIMA methodology for fitting univariate autoregressive and moving average models to time series.

A primary concern with respect to VaR is how to model the co-movement between relevant variables. Two models that do so are the Vector Autoregressive Model (VARM), in which each variable depends on past values of other variables in the system as well as its own past values, and the Vector Error Correction Model (VECM), which models long-run equilibrium relationships between variables (Kim, Malz and Mina, 1999). We explore statistical modelling techniques in more depth later within this report. Examples of VECM modelling can be found in Fisher, Fackler and Orden (1995), who model the interaction between money, prices and output in New Zealand using a VECM model, and Fisher (1996), who examines dynamic interactions between the exchange rate, domestic price level and terms of trade in Australia and New Zealand.

3 Methodology

3.1 Overview

In this section we outline the methodology behind the Fonterra Value-at-Risk model. Particular emphasis is placed on why we select the model that we do. With this in mind we begin by briefly looking at the problems associated with alternative VaR models.

Since its inception at JP Morgan in 1994, the RiskMetrics framework (see Section 2.4) has become the industry standard for VaR calculations within trading institutions. However, we have already noted that the RiskMetric model has significant drawbacks when we move from a financial institution to a corporate, primarily relating to the complex interaction between business and market risks. In addition to this over-riding concern, which here is somewhat mitigated by the fact that we limit the scope of the investigation to market risk alone, the RiskMetrics approach suffers from more specific problems in our context.

Firstly, the RiskMetrics model assumes that the log prices of financial instruments follow a random walk; that is, log daily price changes (returns) are independent. This is generally a reasonable assumption for daily prices over the short time frames typically associated with VaR (usually less than three months), since such changes are effectively random. However, we are concerned with monthly prices over a long-term timeframe (greater than one year). Given the long-term trends inherent in foreign exchange and interest rate series, monthly returns are likely to exhibit significant autocorrelation. Indeed, we find evidence that this is the case. The random walk model does not provide the best explanation for the long-term dynamics of financial returns over long horizons (Kim, Malz and Mina, 1999).

Secondly, the RiskMetrics model assumes that the mean return is zero. This assumption has little impact over short-term horizons but is untenable over longer horizons. Finally, in the RiskMetrics model the variance is a linear function of time with the implication that the daily standard deviation can be multiplied by the square

root of time to obtain the standard deviation for a longer period (Kim, Malz and Mina, 1999). Again, this "scaling up" approximation is inappropriate when dealing with longer horizons.

The RiskMetrics group's CorporateMetrics (Lee, 1999) seeks to address the issues inherent in extending VaR to a corporate setting. While CorporateMetrics offers a comprehensive framework that goes a long way towards solving the problems of corporate VaR, a number of issues render it inappropriate for our purposes. As we have seen, CorporateMetrics utilizes long-term forecast models to forecast the distribution of returns at each time horizon. Price paths are then generated by sampling from these forecast distributions taking into account the historical correlation between variables. The model used to forecast the mean price and variance at each time horizon varies. At shorter time horizons extensive use is made of current market information[14], based on the rationale that market expectations of future prices are embedded in current spot and derivative prices (Kim, Malz and Mina, 1999). In this regard problems arise with respect to the Fonterra VaR model. Such a forecasting methodology requires extensive market data such as option prices and forward rates, data which Fonterra does not have readily available. Even if procedures were put in place to obtain necessary market data[15] the forecast methodology effectively applies only to traded financial instruments. A lack of derivatives or other instruments on Fonterra's commodity prices means that this method is not applicable. On top of these data difficulties, the plethora of different forecast models required is impractical for this project. The CorporateMetrics approach necessitates not only the use of different models for different instruments, but also different models for different forecast horizons[16]. Thus the CorporateMetrics approach is both too data intensive and too model intensive for out purposes.

---

[14] For example, the volatility forecast for a particular instrument might be obtained from the implied volatility of options on that instrument.

[15] For example, import from Reuters or purchase from RiskMetrics.

[16] While market data may be used to estimate the mean return in one year's time, the mean return in two year's time might be estimated using a VECM model or an economic equilibrium model. Loosely speaking there is a trade-off between the forecast accuracy obtained by using a greater number of models and the additional model risk that increasing the number of models entails.

What we require is a modelling methodology that takes into account the autocorrelation in monthly returns, requires minimal data (ideally we would like to use only the time series that we wish to forecast themselves) and takes into account the relationship between the series. One method might be to estimate multivariate regression models using lagged values of economic time series as predictors in an attempt to determine the long-term relationships between variables. However, as we will show, the variables in question exhibit non-stationarity and therefore relationships estimated by OLS are likely to be spurious (Granger and Newbold, 1974).

Univariate time series models, such as ARIMA models, fit the first two criteria. Such models are easy to estimate[17] and require only past observations of the variable to be forecast. However, ARIMA models do not explicitly account for the correlation between variables, although they do so implicitly (since past realisations of the variable being modelled, from which the ARIMA model is estimated, depend on the value of other variables).

An alternative to univariate time series models is the vector error correction model (VECM). In cases where variables are cointegrated it is possible to estimate a VECM model that models the long-run relationship between non-stationary variables. The VECM model meets all our criteria: it requires data only for the series being forecast, it models the autocorrelation and its forecasts take into account the long-run relationships between variables. Of course, the existence of a VECM model is conditional on the presence of a cointegrating relationship between variables. For this reason we estimate ARIMA models for all variables and VECM models where possible. We then assess the forecast accuracy over a 12-month hold-out period for both models. The ARIMA models provide a base case against which the forecast accuracy of the VECM models can be assessed[18]. In the final VaR model we give the

---

[17] Most statistical packages, for example SAS and Minitab, allow for the easy estimation of ARIMA models.

[18] In a sense, the difference in forecast accuracy between the ARIMA and VECM models is the extent to which the relationship between variables aids in forecasting. If the VECM models outperform the

user the option of choosing the forecast model, thus providing a rough means of assessing the model risk. A large difference in the VaR figure generated using alternate models is indicative of a high degree of model risk (since the result is not robust to model choice).

One final comment is worth making here. Pure time series models may seem naïve in that they do not specifically seek to model the dynamics of market rates based on theoretical economic relationships. Instead, they rely entirely on estimating a statistical relationship from historical data. While this is certainly a valid criticism, there are good reasons for our choice. We have already looked at different forecast models developed specifically for foreign exchange and interest rates, for example, the mean-reverting Siddique and Sweeney (1998) model for exchange rates[19] and the Vasicek (1977) and Cox-Ross-Ingersoll (1985) models for interest rates.

One problem with using these models is simply the number of different models that must be estimated. Not only are the models different, but the method of estimation is also different. For example, one might be estimated via the generalised method of moments (GMM) while another could be estimated using maximum likelihood. As well, no similar models based on economic fundamentals have been developed for commodity prices.

In addition, while these models have been applied extensively to exchange and interest rates individually, it is more difficult to determine how to incorporate correlations between rates into the separate models. Like the univariate ARIMA models, no account is made for relationships between variables. Having said this, however, it should be possible to incorporate such correlation information via the stochastic error term in a Monte-Carlo simulation by drawing correlated random variables. Developing such an approach is an avenue for future research.

ARIMA models then the long-term relationship between variables is important in forecasting and vice-versus.

[19] In fact, the Siddique and Sweeney model is merely an AR(1) model, a special case of the general ARIMA model.

3.2 Data analysis

Having outlined our general approach for developing the Monte-Carlo VaR model we now detail the specifics of the data analyses required to develop appropriate ARIMA and VECM forecast models.

3.2.1 Basic statistics

We begin by computing basic statistics for the data series to determine the unconditional normality, autocorrelation properties and correlation of return series. We employ three separate normality tests. The Anderson-Darling and Kolmogorov-Smirov tests are ECDF (empirical cumulative distribution function) and chi-squared based test respectively, while the Ryan-Joiner test is correlation based. We compute the Durbin-Watson statistic to test for first order autocorrelation in the returns of each series. The presence of autocorrelation among monthly returns is of particular importance since non-independence of returns means the random walk model is inappropriate. Finally, we calculate the correlation coefficient between each of the return series to give an initial indication of the relationships between series.

3.2.2 Stationarity

The stationarity of a time series is vital in determining an appropriate model. Strictly speaking a time series is said to be stationary if the joint and conditional probability distributions of the process are independent of time; however, in most instances we are interested primarily in weak stationarity in which the mean and variance are constant over time and the covariance between two points depends only on the lag between those points and not on the point in time (Ellwood, 2000). More specifically the conditions for weak stationarity can be expressed as

1) $E[y_t] = \mu$ for all t, i.e. the mean is independent of time,

2) $\text{var}(y_t) = \sigma_y^2$ for all t, i.e. the variance is a finite constant, and

3) $\text{cov}(y_t, y_{t+k}) = \gamma_k$, i.e. the covariance between two points depends only on the lag between them and not on t.

Time series that are non-stationary present difficulties with regards to econometric analysis since the normal properties of the least squares regression break down when

non-stationary time series are involved. Granger and Newbold (1974) describe and demonstrate the problems that can occur when one non-stationary time series is regressed on another. In such regressions the least squares estimator is not consistent and the conventional t and F test statistics are not distributed as t and F distributions when the null hypothesis is true, leading to apparent significant relationships where none exist (commonly referred to as "spurious regressions"). Since the usual properties of the least squares estimator break down when the time series involved are non-stationary, such series must be differenced to achieve stationarity before estimating regression models.

To test for stationarity we apply the Augmented Dickey-Fuller (ADF) unit root test. A time series is said to have a unit root if it can be made stationary by differencing. The ADF test tests the null hypothesis that the series is difference stationary[20], i.e. has a unit root, versus the alternative that the series is trend stationary[21]. We apply the Dickey-Fuller test based on three different models as described in Ellwood (2000):

(i)   a zero mean model $\Delta z_t = \beta_0 z_{t-1} + \sum_{j=1}^{n} \beta_j \Delta z_{t-j} + e_t$

(ii)   a single mean model $\Delta z_t = \alpha_0 + \beta_0 z_{t-1} + \sum_{j=1}^{n} \beta_j \Delta z_{t-j} + e_t$, and

(iii)   a time trend model $\Delta z_t = \alpha_0 + \alpha_1 t + \beta_0 z_{t-1} + \sum_{j=1}^{n} \rho_j \Delta z_{t-j} + e_t$

The test applied is the augmented Dickey-Fuller test[22]. The relevant parameter in the above regressions is β_0, the coefficient on z_{t-1}. If $\beta_0=0$ then z_t has unit root; that is, z_t

---

[20] A difference stationary series is one that can be made stationary by taking the difference $z_t - z_{t-1}$.

[21] A series is trend stationary if it can be expressed in the form $z_t=f(t) + e_t$ where $e_t \sim (0, \sigma^2)$ so that the residuals will form a detrended stationary series.

[22] Since the Dickey-Fuller test relies on the OLS estimation of model equation, the standard OLS assumption that the residuals e_t are white noise applies. Any deviation from this will impact on the results of the test. For this reason the augmented Dickey-Fuller test adds additional lags of the difference with the aim of bringing the residuals closer to white noise.

can be made stationary by differencing. Estimating one of the equations above via OLS regression yields an estimate of β_0 along with its standard error. We can test the hypothesis that $\beta_0=0$ by computing an appropriate t-statistic and comparing to the appropriate critical value in the Dickey-Fuller tables as outlined in Ellwood (2000).

Ellwood (2000) points out that the number of lags to be included in the ADF regression needs to be selected carefully. Too few lag terms yields residuals that are not white noise while too many lag terms impacts on the power of the unit root test due to few degrees of freedom (Enders, 1995, pp. 226-227). One means of determining the appropriate lag is to set $p=p_{max}$, where p_{max} is large enough to include the true p. The method is to test the significance of the coefficient on successively smaller lag terms from p_{max} downwards until the coefficient is found to be significantly different from zero (see Ellwood, 2000, pp. 13-14). The first lag that has a coefficient significantly different from zero is selected as the lag length to use in the test. Alternatively an information criterion, such as the Akaike Information Criterion (AIC), can be used to distinguish between models of different lag. Such criteria attempt to measure the explanatory power of alternative models whilst punishing for over-parameterisation.

Ellwood (2000) also outlines a process for selecting between models (i), (ii) and (iii) since once we have selected an appropriate lag length we need to choose the appropriate model. According to Bannerjee, Dolado, Galbraight and Hendry (1993, pp. 100) cited in Ellwood (2000), incorrect choice of model can invalidate standard inferences due to the presence of nuisance parameters in the data generating process. The procedure is to begin with the most general model (iii) and test the significance of the additional constant and trend coefficient to determine whether the model is appropriate, and hence whether the stationarity conclusion drawn from that model is meaningful (see Ellwood, 2000, pp. 16-17).

3.2.3 Cointegration

The concept of cointegration (Granger, 1981) is important in econometrics since it is linked closely to the idea of long-term equilibrium. In general, as we have discussed briefly, if two time series, x_t and y_t, are non-stationary then we must take differences

before attempting to model the relationship between them because of the dangers of spurious regressions. However, there exists a special case when a long-run equilibrium exists between x_t and y_t that allows us to model the relationship between the un-differenced (level) series. Because differencing removes all information contained in the level series, any long-run relationship between level series is lost. Cointegration means that we do not need to difference the data, thus the existence of cointegration means that we can model the long-run relationship between variables.

If two time series y_t and x_t are non-stationary, but their first differences y_t-y_{t-1} and x_t-x_{t-1} are stationary, then y_t and x_t are said to be integrated processes of order 1, denoted I(1). In general when y_t and x_t are I(1), a linear combination $y_t - \alpha - \beta x_t = e_t$ of y_t and x_t will also be I(1). It is possible, however, that in some cases e_t may be stationary, or I(0), if the trends in both series tend to cancel out. In such a case the variables are said to be cointegrated. In the two variable case, cointegration implies the presence of a unique long term relationship between y_t and x_t; however, when the number of variables is greater than two more than one such cointegrating relationship can exist. Thus in general cointegration implies the existence of *at least one* long-term equilibrium relationship between the variables.

We test for cointegration using the maximum likelihood Trace Test proposed by Johansen (1988). The test is conducted by estimating a vector autoregressive (VAR) model by maximum likelihood under different assumptions about the number of cointegrating vectors[23] r and then conducting likelihood ratios tests. The maximum likelihood $L_{max}(r)$ is a function of the cointegration rank r. Johansen's test is based on the log-likelihood ratio

$$\ln\left(\frac{L_{max}(r)}{L_{max}(k)}\right) \text{ where } r = k\text{-}1, ..., 1, 0$$

---

[23] If a linear combination of variables is stationary then the vector of coefficients is called a *cointegrating vector*. For example if $e_t = z_t - \beta_0 - \beta_1 w_t - \beta_2 x_t - \beta_3 y_t$ is I(0) then the vector of coefficients $(1, -\beta_0, -\beta_1, -\beta_2, -\beta_3)$ is the cointegrating vector (see Thomas, 1997, p. 438). When more than two variables are cointegrated more than one such vector can exist.

and tests the null hypothesis that the cointegration rank is r against the alternative that the cointegration rank is k.

The VAR(k) model used in the test is of the form

$$X_t = \mu + A_1 X_{t-1} + A_2 X_{t-2} + ... + A_k X_{t-k} + \varepsilon_t$$

where X_t is an n × 1 vector of I(1) variables (Ellwood, 2000). We can rewrite this using the difference operator $\nabla = 1 - L$, where L is the lag operator, so that

$$\nabla X_t = \mu + \Gamma_1 X_{t-1} + ... + \Gamma_{k-1} X_{t-k-1} + \Pi_k X_{t-k} + \varepsilon_t$$

where

$$\Gamma_i = -(I - A_1 - ... - A_i) \text{ for } i = 1, 2, ..., k\text{-}1$$
$$\Pi = -(I - A_1 - ... - A_k).$$

The matrix Π contains all long-run information in X_t since all the other terms are differenced and hence lose all long-run information. If we factor Π such that $\Pi = \alpha \beta'$, where α and β are n × r matrices, then the columns of β are the cointegrating vectors while the i^{th} row of α indicates the importance of each of these cointegrating relationship to the i^{th} series (see Ellwood, 2000, p. 21). The number of cointegrating vectors is determined by the rank of Π. If $r=n$, then Π is full rank and each series is stationary, i.e. any linear combination of the variables is stationary. If $r=0$ then Π has rank zero and there are no linear combinations of variables that are stationary, i.e. there are no cointegrating vectors. Finally, if $0 < r < $ n then Π has rank r and there are r linear combinations of the variables that are stationary.

It is important that the error terms for each series are white noise. Thus it is often necessary to change the number of lags included in the model to ensure that the

residuals are normally distributed and uncorrelated. We select the simplest model that satisfies these conditions at the 5% level.

Johansen's approach has the advantage that it allows us to determine not only whether cointegration exists, but also the number of cointegrating vectors (Ellwood, 2000). Kasa (1992) points out that if there are r cointegrating vectors relating a set of n variables then there must exist $n-r$ stochastic trends. A stochastic trend is a common trend among the variables. In general, the more cointegrating vectors that exist, the more inter-related are a set of variables. Likewise, the lower the number of stochastic trends, the more inter-related are the variables. For example, if five variables are explained by a single common trend they are highly inter-related. The factor loadings on each variable within the common trend are an indication of the importance of the trend on that variable. If a common trend is relatively unimportant to a particular variable then that variable is more independent compared to the other variables.

3.2.4 Forecast models

We now describe the methodology employed to estimate and select appropriate forecast models. In this section we apply two different time series modelling techniques, namely ARIMA and VECM, to the data series and evaluate the forecast accuracy over a 12-month period. Model parameters are estimated using the period May-1995 to Jan-03, leaving the 12-month period Feb-03 to Jan-04 against which the accuracy of a 12-month forward forecast can be assessed.

3.2.4.1 ARIMA models

We first use the ARMA (or Box-Jenkins) approach to model each time series individually. The Box-Jenkins approach is a generalized method for fitting models to time series in which the observed autocorrelation properties of the series are compared to the autocorrelation properties that would be generated by various theoretical series to assist in assigning a model (Box and Jenkins, 1976). The ARMA process applies only to stationary time series. ARIMA modeling is an extension of the ARMA process, the only difference being that the method is applied to a time series that has been differenced to achieve stationarity.

ARIMA models have the advantage that we require only the past values of the variable in question to predict future values. Of course, this also excludes any potential improvement that including past values of other variables may make to the forecast. Nonetheless, ARIMA models have the advantage of simplicity and so provide a useful starting point for evaluating the alternative forecast models.

If the series are first difference stationary, as we would expect, then the models will be of the form ARIMA(p,1,q), where 1 is the level of differencing required to make the series stationary and p and q are the orders of the autoregressive and moving average components respectively.

In general we can write the ARIMA(p,1,q) model as

$$\nabla z_t = \phi_1 \nabla z_{t-1} + \phi_2 \nabla z_{t-2} + ... + \phi_p \nabla z_{t-p} + a_t - \theta_1 a_{t-1} - \theta_2 a_{t-2} - ... - \theta_q a_{t-q}$$

where $\nabla z_t = z_t - z_{t-1}$ are the first differences which form a stationary series $\{\nabla z_t\}$, a_t are white noise disturbances and ϕ_1, ϕ_2,..., ϕ_p and θ_1, θ_2,..., θ_q are the coefficients of the autoregressive and moving average components respectively. More compactly in backshift notation

$$(1 - B)(1 - \phi_1 B - \phi_2 B^2 - ... - \phi_p B^p)z_t = (1 - \theta_1 B - \theta_2 B^2 - ... - \theta_q B^q)a_t .$$

In line with the standard Box-Jenkins, methodology the ACF (autocorrelation function) and PACF (partial autocorrelation function) are used to determine an initial model for each of the differenced series. Once a potential model is assigned the residuals are checked for autocorrelation using the Box-Pierce chi-squared test statistic at various lags. If the model successfully removes all significant autocorrelation from the residuals it is tentatively selected. We then over-fit using an ARIMA(p+1,1,q) and an ARIMA(p,1,q+1) model. If the additional coefficients, ϕ_{p+1} and θ_{q+1}, are not significant the selected ARIMA(p,1,q) model is deemed to be appropriate. In cases where higher order models also seem appropriate the most parsimonious model is selected.

In practice it is often difficult to determine an appropriate model from the behaviour of the ACF and PACF alone. In cases where more than one model may be appropriate we estimate the possible models and then compare the AIC statistic[24], the autocorrelation of the residuals, the significance of the estimated parameters and the MSE of each model to determine a preferred model. The assignment of ARIMA models is often a qualitative exercise at best so it is important that the selected models are tested for robustness. One means of doing so is to estimate separate ARIMA models in different sub-periods[25]. If the models and coefficient estimates in each sub-period are similar then we can be more confident that the model is robust. Here, however, we are constrained by a lack of data so we must be content with re-estimating the models as more data becomes available. If the coefficient estimates do not remain stable over time then we are likely exposed to significant model risk.

Two measures of forecast accuracy are computed, the mean average percentage error (MAPE), which provides a measure of the average percentage error of the 12 forecast values, and the point forecast percentage error which is simply the percentage error between the 12 month forward forecast value and the realised 12 month forward value. Specifically the accuracy measures are computed as

$$MAPE = \frac{1}{12} \sum_{t=T+1}^{12} \frac{|y_t - f_t|}{|y_t|} \times 100$$

$$Point\ Forecast\ \%\ Error = \frac{|y_{T+12} - f_{T+12}|}{|y_{T+12}|} \times 100$$

where y_t is the actual price, f_t is the forecast price and T is the number of observations.

---

[24] The AIC (Akaike Information Criteria) value is an indication of the fit of the model taking into account over-parameterisation, i.e. it selects the model with the minimum SSE while penalising for over-parameterisation. A lower AIC value indicates a better fit.

[25] Abauf and Jorion (1990), in an effort to model exchange rates, describe parameter instability as one of the primary drawbacks of autoregressive models.

3.2.4.2 VECM models

Having developed forecast models for each variable in isolation we move on to look at multivariate models. One such model is the vector error correction model (VECM). Granger and Engle (1987), in the Granger representation theorem, state that if a set of variables are cointegrated then there exists a valid VECM representation. Such a representation means that future changes in variables can be predicted by the deviatation from long-run equilibrium. In simple terms a VECM for two I(1) variables, x_t and y_t, can be expressed as

$$\Delta y_t = \delta + \lambda e_{t-1} + \omega \Delta x_t + \varepsilon_t$$

where $e_t = y_{t-1} - \alpha - \beta x_{t-1}$, the deviation from the long-run equilibrium $y_t = \alpha + \beta x_t$, is I(0). Thus y_t depends on both past values of x_t and the equilibrium error term e_t. The presence of the equilibrium error term in the model acts to bring the model back towards the long-run equilibrium. For example, if e_t is positive then y_{t-1} is above its equilibrium level of $\alpha - \beta x_{t-1}$. If λ is negative then the next value Δy_t will be smaller so when y is out of equilibrium it will tend to move back towards equilibrium in the next period (see Koop, 2000). Because VECM incorporate both the long-run and short-run properties of the relationship between variables it is a popular model.

We first look at whether utilizing past values of other basic commodity price series (BCP) aids in the forecasting each individual BCP. Then, in cases where cointegrating relationships between BCPs and market rates where found to exist, we develop models in which exchange and interest rates are used to aid in the forecast of individual BCPs. Model fit is assessed by looking at the normality and autocorrelation of the residuals for each BCP price series, as well as the cross correlation of residuals between series.

Once we have estimated appropriate models we compare the forecasting performance, using the measures outlined previously, over the same 12-month period as the ARIMA models.

3.3 VaR model

3.3.1 Group treasury hedge portfolio

Fonterra Group Treasury is responsible for hedging foreign currency inflows and debt book interest rate sensitivity. The hedge book consists of four main instruments, namely: FX forwards, NZD call options, interest rate swaps and cross currency interest rate swaps.

Under traditional VaR the composition of the portfolio of financial instruments is assumed to be constant over the VaR time horizon. This is a reasonable assumption when the time horizon is short, as is the case in a trading institution. A complication of calculating VaR over a long-term horizon is the fact that the composition of the portfolio of interest changes over the period. Thus it is necessary to model the evolution of Fonterra's hedge portfolio through time before a VaR calculation is possible. This is done on a monthly basis and is subject to a number of simplifications.

The FX and interest rate portions are modelled separately based on the Fonterra hedging policy for each. Once the composition of the portfolio is known at each month in the time horizon we can calculate the gain/loss in each month due the difference between the realised market rate in that month and the hedged rate. The total hedging gain/loss over the time horizon is then simply the sum of these monthly gains and losses over the time horizon. The portfolio models are linked to a Monte-Carlo simulation model (described in the next section) which generates price paths for each market rate using the estimated ARIMA and VECM models. Each simulated price path is fed into the portfolio model to calculate the total hedging gain/loss under that particular scenario. Repeated iterations generate a distribution of gains/losses on the hedge portfolio from which a value at risk number can be calculated. Thus our VaR measures the *maximum loss over the specified time horizon that would be expected 95% of the time as a result of Fonterra's hedging activities.*

Details of Fonterra's hedging policy are confidential so we can offer only a broad overview here[26].

3.3.1.1 FX hedge portfolio

Fonterra's cash inflows are almost exclusively denominated in USD and all other foreign currency inflows are hedged back to USD dollars. The Board approved Treasury policy, adopted in March 2003, dictates that all forecast USD inflows out to 15 months must be fully hedged against any change in the value of the NZD against the USD[27]. That is, at the end of any given month all forecast cash inflows for the month occurring in 15 months time must be 100% hedged. This is achieved by purchasing forwards and call options to buy NZD dollars (sell USD) in 15 months time. Appendix A provides an example.

The type and proportion of instruments used is also constrained by policy. The implication being that for time horizons up to 15 months the composition of the FX hedge portfolio is known with certainty (all instruments are held to expiration). For time horizons greater than 15 months the composition of the hedge portfolio, given a particular simulated NZDUSD price path, needs to be forecast using the hedging policy. Forecasting the composition of the portfolio introduces additional assumptions[28]. The risk is that the forecast portfolio under each scenario will not match the portfolio that would actually be created under that scenario (dealers have some discretion within the broad constraints of the policy). For example, given a

---

[26] Additional information, for those who have signed a confidentiality agreement with Fonterra Cooperative only, is contained in Appendix D (restricted access).

[27] Although this removes FX risk, additional risks arise due to the fact that forecast cash inflow in 15 months may not exactly match the realised cash inflow and thus the inflow may not be fully hedged. Also non-USD inflows need to be converted back to USD and then hedged with the USD inflows.

[28] Forecasting the FX portfolio composition requires that the forward rate be forecast in addition to the spot rate since we must know at what rate each month's forwards are purchased. Similarly we need to know the strike price and premium at which options are purchased. We can make the assumption that the future forward rate is related to the simulated spot rate by some fixed forward premium/discount and can likewise assume that options are purchased with strike price some fixed number of points above this forward rate. Once the strike price is known we can estimate the option premium either by using the Black-Scholes option pricing model or using a lookup table of volatilities and prices.

scenario in which significant hedging losses are made it is unlikely that Fonterra's hedge policy would remain unchanged. For these reasons we limit our VaR calculations within this report to time horizons less than 15 months, so that the composition of the FX portfolio is already known at each month.

3.3.1.2 Interest rate hedge portfolio

All Fonterra debt is raised on a floating basis; however, Treasury policy dictates that a certain proportion of each year's debt must be fixed for a specified number of years forward via interest rate swaps at the relevant swap rate. The remaining debt is floating and interest is paid at a fixed margin above the NZ 30 day bank bill rate. Thus at any point in time Fonterra has various amounts of fixed debt, each at a potentially different rate, with the remainder of debt floating. The gain/loss on the interest rate hedge portfolio is the amount of interest paid when hedged, less the amount of interest that would have been paid had Fonterra remained unhedged[29].

Given the number of assumptions that must be made, forecasting debt composition is a tenuous exercise at best. One concern is how to forecast the 30 day bank bill rate and the relevant swap rate. To forecast both rates together effectively requires a forecast of the yield curve at each month. Forecasting the time dynamics of the yield curve is not an easy task and we make the simplifying assumption that the relevant swap rate sits at some fixed margin above the 30 day bank bill rate and simulate the 30 day bank bill rate only. Obviously the shape of the yield curve is not constant so this assumption has major implications. With all assumptions we give the user the ability to input different values so that sensitivity analysis on the impact of the particular assumption can be carried out.

3.3.2 Monte-Carlo simulation

The Excel/VBA Monte-Carlo simulation model generates multiple monthly price paths for each market rate series[30] over a selected time horizon. Each price path is

---

[29] Further details of the interest rate hedging policy have been removed due to commercial sensitivity. Authorized persons can find additional detail in Appendix D (restricted access).

[30] Note that the hedge portfolio VaR model does not explicitly require the basic commodity price (BCP) series to be forecast. Nonetheless we include these throughout our data analysis since an

then iteratively input into the hedge portfolio model to calculate the gain/loss on the hedge portfolio under each given scenario. Repeated simulation yields a distribution of such gains/losses from which the final VaR number can be obtained. Figure 1 illustrates the process behind the VaR calculation.

The model allows the user to select either the ARIMA or VECM forecast models. The stochastic error term for both models is obtained in one of two ways depending on the user choice. The first method is to draw the error terms for each forecast path from a normal distribution with mean and variance equal to the historical mean and variance of residuals from the fitted models. Alternatively, the error terms are obtained by bootstrapping with replacement the residuals from the fitted models[31]. The first method makes the implicit assumption that the models fully explain all pattern in the data and thus the residuals are i.i.d. In contrast, the bootstrapping method does not make any such assumption about the distribution of residuals (Brock, Lakonishok and LeBaron, 1992, p. 1745). Both methods make the implicit assumption that the historical errors provide the best estimate of the future stochastic error series.

---

analysis of these series is vital for future work to incorporate the model into milk payout models. Ultimately the gain/loss from the hedge portfolio could be offset against revenues to calculate a VaR around the yearly milk payout.

[31] See Efron (1982), Freedman (1984), Freedman and Peters (1984a, 1984b), Efron and Tibshirani (1986). Brock, Lakanishok and LeBaron (1992) provide a good overview of the bootstrap methodology.
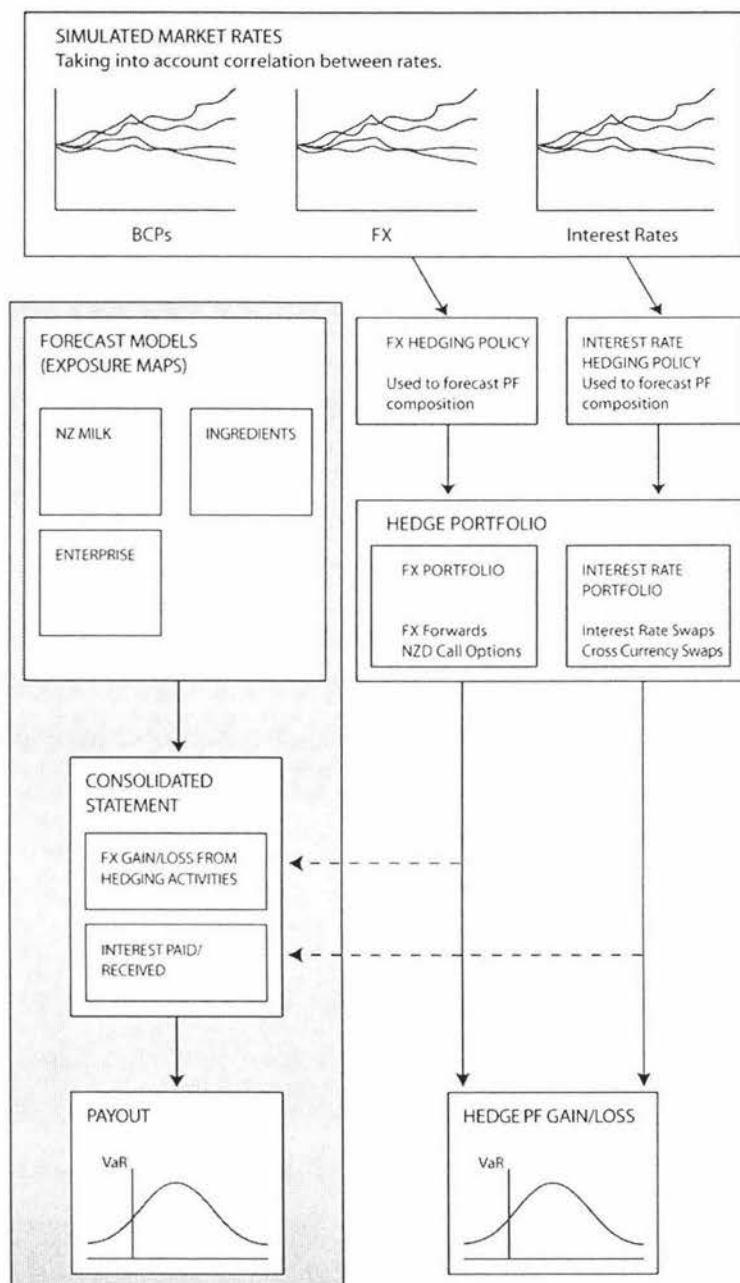
Figure 1: Fonterra VaR model. The current project looks only at the part of the project not enclosed within the grey box. This corresponds to a VaR on the Treasury hedge portfolio. The grey box shows how the project might be extended to incorporate business risk and provide a VaR figure around the yearly milk payout.

4 Data

4.1 Foreign exchange and interest rates

For the NZDUSD exchange rate we use the Reserve Bank of New Zealand (RBNZ) daily series[32] consisting of the mid-rate as at 11 am on each day from Jan-1985 to Jan-2004. We create a monthly series by taking the daily rate on the last trading day of each month. This gives us 229 monthly observations. Given that the NZD was floated on 4 March 1985 data earlier than the 1985 is unlikely to reflect the current market determined exchange rate so we begin the series in 1985.

We use two interest rate series, both sourced from the RBNZ. The NZ 30 Day Bank Bill Rate is used as the short-term rate while the 10 Year T-Bond rate is used to proxy longer term interest rates. Both rates are the rate as at 11 am on the last trading day of each month from Jan-1985 to Jan-2004 and consist of 229 observations. We choose to model yields, rather than prices, for fixed income instruments since a well documented shortcoming of modeling price returns is that the method ignores a bond's price pull to par phenomenon (Zangari, 1996).

Graphs of the level and return[33] series for the exchange and interest rates are displayed in Appendix A.

4.2 Commodity prices

Ten basic commodity product (BCP) price series are used in the analysis as shown in Table 1. Each BCP price series consists of the monthly volume weighted average price (USD per Metric Ton) of commodity contracts in the month. We use data from the period May-1995 to January-2004. This gives 105 monthly observations for each of the ten BCP price series. Graphs of each of the commodity price series can be found in Appendix D (restricted access).

---

[32] RBNZ historical exchange and interest rate series are available from the RBNZ website at http://www.rbnz.govt.nz.

[33] Where returns are calculated we use log returns. The return at time t is given by $r_t = ln\left(\dfrac{P_t}{P_{t-1}}\right)$.

Variable Name	Commodity
WMP	Whole Milk Powder
SMP	Skim Milk Powder
DS Cheese	Dry Salted Cheese
BS Cheese	Brine Salted Cheese
Casein	Casein
Butter	Butter
BMP	Butter Milk Powder
AMF	Anhydrous Milk Fat
WPC	Whey Protein Concentrate
Lactose	Lactose

Table 1: BCP variables.

5 Results

5.1 Data analysis

We examine the data series with the aim of specifying ARIMA and VECM forecast models and identifying problems with respect to a VaR model. We first look at each data series in isolation, examining properties such as excess kurtosis that are likely to complicate the VaR model, and then move on to look at the extent to which the data series are interrelated. The second aspect, as we have previously noted, is particularly relevant from a Monte-Carlo simulation perspective. Any co-movement between the key variables must be taken into account when simulating price paths into the future, thus identifying cointegration between variables is a critical step of the analysis.

5.1.1 Summary statistics

Summary statistics for the NZDUSD exchange rate and interest rates, along with their respective return series, are presented in Table 2. Table 3 displays the same statistics for the BCP price series. Due to confidentiality constraints this table is only available in the restricted appendix (see Appendix D).

	Levels			Returns		
	NZDUSD	30 Day BB	10Y T-Bond	NZDUSD	30 Day BB	10Y T-Bond
Mean	0.563292	10.36105	9.565022	0.001543	-0.00509	-0.00457
Standard Error	0.005075	0.386719	0.25167	0.002111	0.00634	0.002723
Median	0.5658	8.24	7.82	0.002158	-0.00374	-0.00667
Mode	0.5542	5.28	12.26	#N/A	0	0
Standard Deviation	0.076802	5.852118	3.808458	0.031879	0.095735	0.04111
Sample Variance	0.005899	34.24729	14.50435	0.001016	0.009165	0.00169
Kurtosis	-0.56021	0.939997	-0.4927	4.44355	6.944146	0.265737
Skewness	-0.18716	1.277914	0.919489	0.331034	0.979702	0.025236
Range	0.30795	26.36	13.54	0.281337	0.877818	0.227482
Minimum	0.40115	3.64	5.36	-0.10594	-0.32733	-0.11778
Maximum	0.7091	30	18.9	0.175398	0.550491	0.109699
Sum	128.994	2372.682	2190.39	0.351713	-1.16061	-1.04233
Count	229	229	229	228	228	228

Table 2: Selected summary statistics for exchange and interest rate series.

5.1.2 Normality

Of particular interest with regards to VaR modelling is whether returns are normally distributed since, as we have seen, normality of returns significantly simplifies VaR modelling. The NZDUSD returns have kurtosis 4.44. In contrast a normal distribution has kurtosis of 3. Excess kurtosis, or "fat-tails", is a well-documented feature of financial returns (see Fama (1965) and Mandelbrot (1963)) and complicates a VaR analysis since extreme values are more likely to occur than is the case under an assumption of normality. As well as exhibiting excess kurtosis, the return series demonstrate positive skewness. Like excess kurtosis, positive skewness is a common characteristic of financial time series.

Normality tests on the exchange and interest rate returns are displayed in Table 4. For the NZDUSD and 30 Day bank Bill rate returns the null hypothesis of normality is strongly rejected in all three tests. In contrast, we fail to reject normality for the 10Y T-Bond return series.

	Returns		
	NZDUSD	30 Day BB	10Y T-Bond
Anderson-Darling	3.098	9.241	0.58
	(0.0000)***	(0.0000)***	(0.1300)
Ryan-Joiner	0.967	0.929	0.9964
	(<0.0100)***	(<0.0100)***	(>0.1000)
Kolmogorov-Smirnov	0.092	0.152	0.048
	(<0.0100)***	(<0.0100)***	(0.1500)

Table 4: Return normality tests.

It is important to note that these are tests of *unconditional* normality. A lack of unconditional normality does not necessarily preclude *conditional* normality. Even if returns are not normally distributed overall they can still be normally distributed conditional on their variance[34]. Indeed, this is the basic assumption of the RiskMetrics approach. Conditional normality allows for fat tails in the unconditional return distribution, however, the distributions of many observed financial return series have tails that are "fatter" than those implied by conditional normality (Morgan Guaranty

---

[34] That is, while the returns r_t may not be normally distributed, r_t/σ_t may be normally distributed.

Trust Company, 1996) so the assumption of conditional normality is still often insufficient to adequately model financial returns.

It is also worth remembering at this point that it is the normality of the *portfolio* returns that is important here. If a portfolio is reasonably diversified and if the individual returns are sufficiently independent it is possible that the portfolio return can be reasonably approximated by a normal distribution, even if the returns of the individual instruments are themselves non-normal (Dowd, 1998)[35].

In the case of Fonterra's hedge portfolio, however, these conditions are unlikely to be met. Fonterra's foreign exchange exposure is almost exclusively due to the NZDUSD exchange rate through NZD call options and forwards, thus it is unlikely that sufficient diversification will exist to eliminate the non-normality of the NZDUSD returns. In addition, a significant portion of Fonterra's hedge portfolio (the NZD call options) is non-linear in foreign exchange risk. As we have seen, non-linearity in risk factors is an immediate obstacle to the normal VaR approach. Although methods have been devised for overcoming this difficulty, such as the Delta-Normal approach, most tend to have the drawback of making further simplifying assumptions. Finally, it is doubtful that the individual returns are wholly independent. For example, economic theory would suggest that exchange rate movements are, at least in the long run, linked to interest rate differentials. Determining the extent to which the variables move together is the subject of much of this analysis.

5.1.3 Autocorrelation

Table 5 shows the first order autocorrelation and Durbin-Watson statistic for each return series. The Durbin-Watson statistic is often very different from two, indicating that significant autocorrelation is present in many of the series. This suggests that in many cases a random walk model will be insufficient to model the data[36].

---

[35] Wilson (1994) and Frain and Meegan (1996) examine the so-called "portfolio-normal" approach in which portfolio returns are assumed to be normal without the additional assumption that individual asset returns are normally distributed.

[36] In addition, the presence of persistence in regressors can lead to spurious regression problems in much the same way as non-stationarity (see Ferson, Sarkissian and Simin, 2003). Powell, Shi and

Variable	1st order autocorrelation coefficient	Durbin-Watson statistic
WMP	0.23	1.48
SMP	0.28	1.35
DS_Cheese	-0.05	2.08
BS_Cheese	-0.08	2.15
Casein	0.41	1.15
Butter	0.07	1.83
BMP	0.05	1.88
AMF	-0.03	1.98
WPC	0.61	0.77
Lactose	0.45	1.08
BB_RATE	-0.12	2.24
NZDUSD	0.10	1.80
TBOND	0.14	1.71

Table 5: First order autocorrelation and Durbin-Watson statistic for return series.

5.1.4 Correlations

Table 6 displays the matrix of correlations between the return series. The correlation between commodity prices and the exchange rate is of particular importance as is demonstrated vividly by the story of Western Mining Corporation Holdings, an Australian exporter facing similar risks to Fonterra (see Maloney, 1990). There appears to be little evidence of correlation between the NZDUSD exchange rate returns and the commodity price returns.

5.1.5 Stationarity

The results of the Augmented Dickey-Fuller unit root test are displayed in Table 7 and 8. The results for each three of the alternative models are presented; in almost all cases the model used for the hypothesis test is the time trend model.

Smith (2003) provide a good discussion of the problem in relation to the use of persistent dividend yields as a predictor of future returns.

Variable	WMP	SMP	DS_Cheese	BS_Cheese	Casein	Butter	BMP	AMF	WPC	Lactose	BB_RATE	NZDUSD	TBOND
WMP	1.00	0.66	0.23	0.11	0.01	0.34	0.16	0.47	0.26	0.01	0.05	0.03	-0.14
SMP	0.66	1.00	0.34	0.12	0.18	0.34	0.25	0.41	0.21	0.07	0.10	0.01	-0.17
DS_Cheese	0.23	0.34	1.00	0.24	0.29	0.20	0.07	0.29	0.14	0.23	0.08	0.06	0.00
BS_Cheese	0.11	0.12	0.24	1.00	0.21	0.06	-0.10	0.04	0.11	0.12	-0.09	-0.07	-0.12
Casein	0.01	0.18	0.29	0.21	1.00	0.02	0.08	0.08	0.22	0.33	0.03	-0.11	0.07
Butter	0.34	0.34	0.20	0.06	0.02	1.00	0.20	0.47	0.05	-0.12	-0.13	-0.05	-0.16
BMP	0.16	0.25	0.07	-0.10	0.08	0.20	1.00	0.30	0.18	0.03	0.03	0.06	0.06
AMF	0.47	0.41	0.29	0.04	0.08	0.47	0.30	1.00	0.09	-0.02	-0.14	0.12	-0.13
WPC	0.26	0.21	0.14	0.11	0.22	0.05	0.18	0.09	1.00	0.29	0.03	0.01	-0.07
Lactose	0.01	0.07	0.23	0.12	0.33	-0.12	0.03	-0.02	0.29	1.00	-0.07	0.09	0.18
BB_RATE	0.05	0.10	0.08	-0.09	0.03	-0.13	0.03	-0.14	0.03	-0.07	1.00	-0.12	0.21
NZDUSD	0.03	0.01	0.06	-0.07	-0.11	-0.05	0.06	0.12	0.01	0.09	-0.12	1.00	-0.19
TBOND	-0.14	-0.17	0.00	-0.12	0.07	-0.16	0.06	-0.13	-0.07	0.18	0.21	-0.19	1.00

Table 6: Return correlation matrix.

	NZDUSD	30 Day BB	10Y T-Bond
Zero mean model	0.51	-2.07	-2.06
	(0.8256)	(0.0375)**	(0.03774)**
Single mean model	-1.99	-1.99	-1.76
	(0.2930)	(0.2898)	(0.4000)
Time trend model	-2.09	-2.22	-1.82
	(0.5483)	(0.4737)	(0.6928)

Table 7: Dickey-Fuller unit root tests on FX and interest rate series.

For all series we fail to reject the null hypothesis of a unit root at any significant confidence level. This suggests that we need to difference the series to obtain stationarity.

To confirm that the level series are first difference stationary we repeat the analysis on the first difference of each series as displayed in Table 9 and Table 10. In every first difference series we reject the null of a unit root at the 1% level, confirming that the first difference series are indeed stationary. We also test the NZDUSD and interest rate return series and, like the first difference series, these also exhibit stationarity as would be expected.

	WMP	SMP	DS Cheese	BS Cheese	Casein
zero mean model	-0.72	-0.90	-0.03	-0.26	-0.60
	(0.4013)	(0.3250)	(0.6727)	(0.5887)	(0.4529)
single mean model	-2.27	-2.22	-1.24	-1.25	-2.35
	(0.1846)	(0.1990)	(0.6543)	(0.6517)	(0.1592)
time trend model	-1.89	-1.87	-0.57	-1.71	-2.22
	(0.6525)	(0.6650)	(0.9784)	(0.7422)	(0.4717)
	Butter	BMP	AMF	WPC	Lactose
zero mean model	-0.82	-0.92	-0.96	-0.65	0.29
	(0.3581)	(0.3177)	(0.2984)	(0.4338)	(0.7671)
single mean model	-1.84	-2.49	-2.03	-2.94	-2.39
	(0.3606)	(0.1199)	(0.2751)	(0.0448)**	(0.1464)
time trend model	-1.67	-2.33	-1.17	-2.91	-2.42
	(0.7568)	(0.4146)	(0.9114)	(0.1626)	(0.3664)

Table 8: Dickey-Fuller unit root tests on BCP price series.

	NZDUSD	30 Day BB	10Y T-Bond
zero mean model	-4.92	-6.70	-5.37
	(0.0000)***	(0.0000)***	(0.0000)***
single mean model	-4.96	-6.86	-5.67
	(0.0001)***	(0.0001)***	(0.0001)***
time trend model	-4.94	-7.00	-5.87
	(0.0004)***	(0.0001)***	(0.0000)***

Table 9: Dickey-Fuller unit root tests on differenced FX and interest rate series.

	WMP	SMP	DS Cheese	BS Cheese	Casein
zero mean model	-5.33	-4.58	-5.94	-6.90	-4.99
	(0.0000)***	(0.0000)***	(0.0000)***	(0.0000)***	(0.0000)***
single mean model	-5.32	-4.58	-5.92	-6.86	-4.96
	(0.0001)***	(0.0003)***	(0.0001)***	(0.0001)***	(0.0001)***
time trend model	-5.43	-4.70	-6.01	-6.83	-4.99
	(0.0001)***	(0.0012)***	(0.0000)***	(0.0001)***	(0.0005)***
	Butter	BMP	AMF	WPC	Lactose
zero mean model	-5.58	-7.88	-5.62	-5.75	-4.49
	(0.0000)***	(0.0000)***	(0.0000)***	(0.0000)***	(0.0000)***
single mean model	-5.57	-7.85	-5.62	-5.72	-4.49
	(0.0001)***	(0.0001)***	(0.0001)***	(0.0001)***	(0.0004)***
time trend model	-5.62	-7.88	-5.83	-5.70	-4.46
	(0.0001)***	(0.0001)***	(0.0000)***	(0.0001)***	(0.0028)***

Table 10: Dickey-Fuller unit root tests on differenced BCP price series.

The non-stationarity of the NZDUSD exchange rate is worth discussing further. Conventional wisdom is that floating exchange rates are non-stationary (Sweeney, 2001); however, studies by Jorion and Sweeney (1996) and later Siddique and Sweeney (1998) find evidence that exchange rates are in fact mean reverting. These studies make use of tests with more power than the ADF test employed here. However, given that we will be forecasting the exchange rates over a one to two year period, even if exchange rates are mean reverting we would expect such behaviour to be exhibited over a longer period than this[37].

As we have discussed, any regression analysis between series must be conducted on the differenced series to avoid the spurious regression problem outlined by Granger and Newbold (1974). However, differencing a time series has the effect of removing all information contained in the levels of that series. Thus information relating to the long-run relationship between the levels of two variables is lost. In the next section we look for cointegrating relationships which can allow us to model such long-run relationships when the series involved are non-stationary.

5.1.6 Cointegration

We first check for cointegration between the NZDUSD exchange rate and each of the two interest rates using Johansen's (1988, 1991) Eigenvalue Trace Test (Table 11). The first row of the table tests the null hypothesis that $r = 0$ against the alternative that $r = 1$ and the second row tests $r = 1$ against $r > 1$. The hypotheses are tested using two different VAR models, one with a linear drift and one with a constant drift[38]. If the trace statistic is larger than the 5% critical value we reject the null hypothesis that r cointegrating vectors are present. The Johansen test is sensitive to the number of lags p included in the VAR model. Verbeek (2000, p. 301) points out that choosing too small a p will invalidate the test while choosing too large a p will result in a loss

---

[37] Campbell and Perron (1991) argue that the choice between treating a time series as stationary or non-stationary should depend on the intended application, and not necessarily on the actual stationarity. For example, while a time series may found to be stationary, a non-stationary model may produce better forecasts for the purpose at hand.

[38] In all cases we fail to reject the restriction that the drift is constant, hence, we present the results for the constant drift model only.

of power. We conduct the test using lags of $p = 6$ and $p = 12$ to explore the sensitivity of the results to the choice of p.

			30 Day Bank Bill Rate			
		$p = 6$		$p = 12$		
H0	H1	Eigen Value	Trace Statistic	Eigen Value	Trace Statistic	5% Critical Value
$r = 0$	$r > 0$	0.0343	11.52	0.0373	15.11	19.99
$r = 1$	$r > 1$	0.0166	3.73	0.0311	6.86	9.13
			10Y T-Bond Rate			
		$p = 6$		$p = 12$		
H0	H1	Eigen Value	Trace Statistic	Eigen Value	Trace Statistic	5% Critical Value
$r = 0$	$r > 0$	0.0389	13.77	0.0412	16.32	19.99
$r = 1$	$r > 1$	0.0218	4.92	0.0326	7.2	9.13

Table 11: Johansen cointegration test between NZDUSD and interest rates.

We fail to reject the null hypothesis that no cointegrating vectors exists at the 5% level for both interest rate series. Thus there is no evidence of cointegration between the NZDUSD exchange rate and the short or long term interest rates in NZ, at least not over the time period that we examine. The results for lags of 6 and 12 are similar, indicating the robustness of this result[39]. Lack of cointegration between the exchange rate and interest rates is advantageous from a simulation perspective since it suggests that the NZDUSD rate and interest rate can be simulated independently, removing the need for multivariate models that the presence co-movement between these rates would necessitate.

Theoretically in the long-term we would expect the NZDUSD exchange rate to equal the ratio of the US and NZ price levels (absolute purchasing power parity). Following an example of Verbeek (2000, p. 287) we can express absolute PPP in logarithms as

$$s_t = p_t - p_t{}^*$$

where s_t is the log of the spot exchange rate, p_t the log of domestic prices and $p_t{}^*$ the log of foreign prices. The presence of a cointegrating vector would suggest a long-

---

[39] Given this result we report the results using a lag length of 12 only for the remaining analysis.

term equilibrium relationship between the exchange rate and foreign and domestic price levels. Thus, while we do not find any relationship between the exchange rate and interest rates that could be used to assist in long-term forecasting, methods based on economic fundamentals such as PPP or interest rate parity could provide a fruitful avenue for future research.

We also investigate whether the BCP price series are cointegrated with the exchange or interest rates. Identifying relationships (or lack thereof) between commodity prices and the market rates is vital for quantifying Fonterra's risk exposure. For a given exchange and interest rate scenario we must be able to state how commodity prices will be impacted and vice versus. Table 12 presents the Johansen Trace Test results for each individual BCP price series with each of the market rates. Here the tests are conducted on two series at a time, later we look for cointegrating relationships between all the BCP price series.

Variable	H0	H1	NZDUSD		30 Day BB		10Y T-Bond	
			Eigenvalue	Trace	Eigenvalue	Trace	Eigenvalue	Trace
WMP	$r = 0$	$r > 0$	0.08	13.19	0.13	19.62	0.14	23.19**
	$r = 1$	$r > 1$	0.05	5.19	0.07	6.65	0.10	9.64**
SMP	$r = 0$	$r > 0$	0.09	13.30	0.10	17.48	0.11	19.15
	$r = 1$	$r > 1$	0.05	4.63	0.08	7.30	0.08	8.11
DS_Cheese	$r = 0$	$r > 0$	0.15	21.45**	0.22	29.21**	0.14	22.61**
	$r = 1$	$r > 1$	0.06	5.96	0.07	6.60	0.09	8.34
BS_Cheese	$r = 0$	$r > 0$	0.17	22.52**	0.14	20.76**	0.15	23.71**
	$r = 1$	$r > 1$	0.05	5.11	0.07	6.32	0.08	8.06
Casein	$r = 0$	$r > 0$	0.15	19.27	0.16	23.83**	0.19	28.29**
	$r = 1$	$r > 1$	0.04	4.04	0.08	8.07	0.09	9.12
Butter	$r = 0$	$r > 0$	0.23	27.94**	0.12	16.60	0.14	21.46**
	$r = 1$	$r > 1$	0.04	3.61	0.05	4.95	0.08	7.94
BMP	$r = 0$	$r > 0$	0.11	15.01	0.14	20.96**	0.15	24.65**
	$r = 1$	$r > 1$	0.04	4.12	0.07	6.97	0.09	9.13**
AMF	$r = 0$	$r > 0$	0.24	28.94**	0.14	20.34**	0.13	23.09**
	$r = 1$	$r > 1$	0.04	3.95	0.07	6.47	0.10	9.89**
WPC	$r = 0$	$r > 0$	0.08	11.99	0.09	16.95	0.11	17.44
	$r = 1$	$r > 1$	0.05	4.44	0.08	7.94	0.07	6.81
Lactose	$r = 0$	$r > 0$	0.12	15.81	0.21	33.90**	0.13	22.71**
	$r = 1$	$r > 1$	0.05	4.41	0.12	11.91**	0.10	10.09**

Table 12: Johansen cointegration test between individual BCPs and each market rate.

The 5% critical value for H_0 is 19.99 and 9.13 for H_1. In cases denoted by ** we reject the null hypothesis of r cointegrating vectors. In most cases we reject $r = 0$ and subsequently cannot reject the null hypothesis that $r = 1$ implying the existence of a stable long-term relationship between the pair of variables. In some cases, for example WMP with 10Y T-Bond, we also reject the hypothesis that there is one cointegrating vector. This implies that the cointegration rank is two, which is equal to the number of variables. Such a result means that both series are themselves stationary. This conflicts with our earlier findings using the Dickey-Fuller unit root test. We note that in these cases the null of one cointegrating vector is rejected only marginally, so one possibility is that only one cointegrating vector exists. Alternatively it has been noted that the Dickey-Fuller unit root test lacks power (Jorion and Sweeney, 1996) so it is possible that test may have failed to reject the presence of a unit root in the these time series when the series where in fact stationary in levels. It should be stressed that discretion needs be exercised when selecting the cointegration rank, particularly given the Johansen test's sensitivity to the number of lags included. Although the Johansen test may indicate cointegration, if the long-term relationship obtained from the estimated cointegration vector makes little sense economically, caution should be exercised.

Table 13 summarises the results, indicating those series which are cointegrated at the 5% level.

	NZDUSD	30 Day BB	10Y T-Bond
WMP	No	No	No
SMP	No	No	No
DS_Cheese	Yes	Yes	Yes
BS_Cheese	Yes	Yes	Yes
Casein	No	Yes	Yes
Butter	Yes	No	Yes
BMP	No	Yes	No
AMF	Yes	Yes	No
WPC	No	No	No
Lactose	No	No	No

Table 13: BCP cointegration with market rates.

Finally we look for cointegrating relationships between all the BCP price series. The time series graph presented in Appendix D (restricted access) suggests that some of

the BCP price series are highly correlated so we expect cointegrating relationships to exist. The results are presented in Table 14.

H0	H1	Eigenvalue	Trace	5% Critical Value
r = 0	r > 0	0.77	533.72**	244.56
r = 1	r > 1	0.60	387.47**	203.34
r = 2	r > 2	0.53	297.74**	165.73
r = 3	r > 3	0.44	222.49**	132.00
r = 4	r > 4	0.42	164.86**	101.84
r = 5	r > 5	0.32	110.3**	75.74
r = 6	r > 6	0.30	71.88**	53.42
r = 7	r > 7	0.16	36.22**	34.80
r = 8	r > 8	0.11	19.26	19.99
r = 9	r > 9	0.08	8.03	9.13

Table 14: Johansen Trace Test for BCP cointegration.

At the 5% level the null that there are seven cointegrating relations is rejected while we cannot reject the null that there are eight cointegrating vectors. Thus we conclude that there are eight cointegrating vectors linking the ten BCP variables. Since there are eight cointegrating vectors among ten variables we conclude that there are two common trends among the ten variables. This result highlights one of the problematic issues with cointegrating relationships, namely if there are eight different linear combinations of BCPs that yield a stationary error series which one do we regard as the "true" long run relationship?

We use results from this analysis to develop vector error correction (VECM) models to take advantage of the long-term relationships identified in a later section.

5.2   Forecast Models

Now that we have examined the characteristics of the data series we develop forecast models for use in the Monte-Carlo simulation. In this section we apply two different time series modelling techniques, namely univariate ARIMA models and multivariate VECM models, to the data series and evaluate the forecast accuracy over a 12-month period. Model parameters are estimated using the period May-1995 to Jan-03 leaving the 12-month period Feb-03 to Jan-04 against which the accuracy of a 12-month forward forecast can be assessed.

5.2.1 ARIMA Models

As outlined in the methodology section we apply the standard Box-Jenkins approach (Box and Jenkins, 1976) to identify appropriate ARIMA models. Table 15 presents the selected model and estimated coefficients for each BCP price series. The p-value for the significance or each estimated coefficient is given in brackets and the significance is denoted in the usual manner.

BCP	Model	Parameter Estimates				Forecast Accuracy	
		ϕ_1	ϕ_2	θ_1	θ_2	MAPE	Point Forecast % Error
WMP	ARIMA(1,1,1)	0.7818		0.5693		12.75	6.93
		(0.0002)***		(0.0247)**			
SMP	ARIMA(1,1,1)	0.8120		0.5646		17.01	14.21
		(<0.0001)***		(0.0072)***			
DS_Cheese	ARIMA(0,1,0)					13.33	28.36
BS_Cheese	ARIMA(0,1,0)					16.52	31.42
Casein	ARIMA(1,1,0)	0.3775				20.92	34.14
		(0.0002)***					
Butter	ARIMA(2,1,2)	1.2286	-0.6310	1.2063	-0.8353	9.89	14.25
		(<0.0001)***	(0.0002)***	(<0.0001)***	(<0.0001)***		
BMP	ARIMA(2,1,2)	-0.5365	-0.9319	-0.6730	-0.9835	8.27	8.20
		(<0.0001)***	(<0.0001)***	(<0.0001)***	(<0.0001)***		
AMF	ARIMA(1,1,2)	0.6159		0.6800	-0.3715	5.64	10.78
		(0.0015)***		(0.0004)***	(0.0005)***		
WPC	ARIMA(2,1,0)	0.8164	-0.2733			9.54	5.66
		(<0.0001)***	(0.0085)***				
Lactose	ARIMA(1,1,0)	0.4728				9.49	21.95
		(<0.0001)***					
					Average	12.34	17.59

Table 15: BCP ARIMA Models.

ARMA models are also estimated for the NZDUSD, 30 Day Bank Bill and 10Y T-Bond return series (Table 16). Because of the extreme levels of interest rates in the late 80s a log-transformation is found to be necessary before we can adequately model the interest rate series. The ARIMA model for the 30 Day Bank Bill Rate and 10Y T-Bond rate are thus both estimated using the natural log of the series. The NZDUSD model applies to the original series.

		Parameter Estimates				Forecast Accuracy	
Series	Model	ϕ_1	ϕ_2	θ_1	θ_2	MAPE	Point Forecast % Error
ln 30 Day BB	ARIMA(1,1,1)					5.20	4.62
NZDUSD	ARIMA(1,1)	-0.7048		-0.7886		7.98	18.33
		(<0.0085)***		(0.0008)**			
ln 10Y T-Bond	ARIMA(0,1,1)			-0.1838		2.35	0.06
				(0.0067)***			
					Average	5.18	7.67

Table 16: FX and interest rate ARIMA models. Interest rate models are estimated using log-transformed series.

We initially estimated each model with a constant term; however, in all cases this term was insignificant at the 10% level so each model was re-estimated without the constant.

Two measures of forecast accuracy are computed, the mean average percentage error (MAPE), which provides a measure of the average percentage error of the 12 forecast values, and the point forecast percentage error which is simply the percentage error between the 12 month forward forecast value and the realised 12 month forward value. Specifically the accuracy measures are computed as

$$MAPE = \frac{1}{12} \sum_{t=T+1}^{12} \frac{|y_t - f_t|}{|y_t|} \times 100$$

$$Point\ Forecast\ \%\ Error = \frac{|y_{T+12} - f_{T+12}|}{|y_{T+12}|} \times 100$$

where y_t is the actual price, f_t is the forecast price and T is the number of observations.

Appendix C presents graphs of the actual and forecast series for the NZDUSD exchange rate and the two interest rate series. Because of their confidentiality, graphs of the BCP forecasts are only shown in Appendix D (restricted access).

The ARIMA models specified have the advantage that they utilize only the past observations of the variable in question. From the perspective of the VaR project these models offer a simple means of forecasting market rates over a 12-month horizon without the need to utilize additional data series or develop complex multivariate models that in the end might have only a limited effect on the accuracy of the forecast.

However, commodity prices tend to move together, so it may be advantageous to utilize the past values of other commodity price series when making forecasts. This is the next step of the modeling process. We develop models that use the cointegrating relationships between variables in addition to lagged values of the variable in question.

One point that should be noted is the size of the 95% confidence intervals around the forecast. In all but two cases[40] the realized price path falls within these bounds, however, these bounds are very wide. This illustrates the danger of relying too heavily on a single forecast path. Rather, in the Monte-Carlo model we simulate a large number of price paths (roughly 95% of which would be expected to fall within these bands) using bootstrapped residuals from the model estimation phase to provide the stochastic error term, thus better capturing this large future uncertainty. At the same time it is important that these bounds make economic sense. For example, it may be impossible for a commodity price to reach a certain level due to economic or market forces such as the US milk support price. An important part of model testing is to validate the simulation results against historical data to ensure that any boundary constraints are met.

Also it is important to note that the measures of forecast accuracy do not tell us a great deal in isolation. Once we specify VECM models we can use the forecast accuracy to assist in selecting the best modeling method.

---

[40] The realised DS Cheese and BS Cheese price paths fall outside the confidence interval, however, it should be noted that for both these series the first difference series is already effectively white noise so no model is applied. Thus the confidence intervals may not be as meaningful as for the other models.

5.2.2 VECM Models

Having developed forecast models for each variable in isolation we now move on to look at multivariate models.

We first look at whether utilizing past values of different BCPs aids in forecasting each individual BCP. Then, in cases where cointegrating relationships between BCPs and market rates were found to exist, we develop models in which exchange and interest rates are used to aid in the forecast of individual BCPs.

| | VECM p = 1 | | VECM p = 2 | | VECM p = 6 | |
| | Point Forecast % | | Point Forecast % | | Point Forecast % | |
BCP	MAPE	Error	MAPE	Error	MAPE	Error
WMP	18.16	16.11	15.98	5.34	33.81	16.44
SMP	32.90	39.30	27.87	19.37	59.98	57.57
DS_Cheese	8.07	3.66	8.52	9.09	37.60	26.11
BS_Cheese	7.21	6.52	8.46	13.21	22.39	38.41
Casein	10.44	10.09	14.35	13.40	38.50	61.30
Butter	8.63	5.04	8.76	17.49	49.81	61.34
BMP	32.43	32.27	37.06	16.35	83.28	73.55
AMF	4.73	7.60	6.62	13.81	48.66	75.10
WPC	41.24	46.23	42.17	45.27	131.51	163.18
Lactose	4.57	1.21	7.82	3.59	39.43	63.59
Average	16.84	16.80	17.76	15.69	54.50	63.66

Table 17: Forecast error for BCP VECM model.

We have noted the existence of significant cointegrating relationships exist between the ten BCPs series (Table 13). Table 17 shows the forecast error for three different VECM models. Each model uses all ten BCP price series to generate the forecast for each individual BCP. The only difference between models is the number of lag terms included. Model fit is assessed by looking at the normality and autocorrelation of the residuals for each BCP price series, as well as the cross correlation of residuals between series. In general, the individual residual series are white noise and each individual BCP model is significant as assessed by the F-statistic[41]. However, in most

---

[41] There are, of course, exceptions. In some cases the residuals for an individual BCP price series fail the chi-squared normality test or exhibit significant ARCH effects. However, such cases are infrequent

cases significant cross correlation exists between residual series suggesting that the models fail to totally explain the relationship between variables.

In terms of overall average forecasting error, the VECM models under-perform simple ARIMA models. Appendix D (restricted access) displays graphs of the VECM forecasts for each BCP price series using the lag 1 model which gives the lowest forecast error. It is important to remember here that even if VECM models do not perform as well in forecast accuracy terms, from a Monte-Carlo perspective they still may offer an advantage over the univariate ARIMA models in that they incorporate the co-movement between individual series[42]. For example, under a single forecast scenario using ARIMA models the price paths for two closely correlated BCP price series could be widely divergent since no account is made of the relationship between them. Using a VECM model, on the other hand, ensures that the forecast price paths are bound by the long-run cointegrating relationships between variables. Thus under a single scenario the price paths of correlated BCP price series will be similar. Of course, if forecast error is significantly worse under a VECM model then the models are unlikely to be modelling the relationships between series correctly anyway, nullifying this perceived advantage.

We now turn our attention to using the exchange and interest rates to assist in forecasting the BCP price series. Table 12 in section 5.1.6 displayed the cointegrating relationships between each BCP price series and the exchange and interest rates. Not all BCPs exhibit such cointegration, but in cases where they do we estimate a VECM model with the aim of improving the forecast error for these BCP price series. The model and forecast error is displayed in Table 18.

---

and overall the VECM models do a reasonable job of explaining the variation in individual BCP price series.

[42] Although the prevailing view is that the use of cointegrating relationships where possible will produce superior long-term forecasts (see, for example Stock (1995, p. 1)), recent work by Christoffersen and Diebold (1997) suggests that nothing is lost when forecasting at long horizons by ignoring cointegrating relationships (in fact, univariate Box-Jenkins forecasts are just as accurate), at least when standard measures of forecast accuracy are used.

		Forecast Error	
BCP	Additional Explanatory Variables	MAPE	Point Forecast % Error
DS_Cheese	NZDUSD, BB_RATE, TBOND	18.62	2.06
BS_Cheese	NZDUSD, BB_RATE, TBOND	7.74	9.43
Cascin	BB_RATE, TBOND	9.86	22.06
Butter	NZDUSD, TBOND	10.30	10.38
BMP	BB_RATE	22.59	0.67
AMF	BB_RATE, NZDUSD	5.75	14.55
	Average	12.48	9.86

Table 18: VECM models and forecast errors using exchange and interest rates.

The average MAPE over the six BCP price series when using lagged exchange and interest rates as explanatory variables (12.48%) is slightly higher than that of the same six BCPs when using all the lagged BCPs as explanatory variables (11.92%). However, the point forecast accuracy is slightly better (9.86% compared to 10.87%). Given these results it is difficult to conclude that using lagged values of the exchange and interest rates aids in forecasting performance. Of course, here we look only at a single period. Ideally we should look at forecast performance over a range of sub-periods, but - as is so often the case in time series modelling - insufficient data is a major hindrance.

	Forecast Error	
BCP	MAPE	Point Forecast % Error
WMP	26.13	26.26
SMP	29.37	31.61
DS_Cheese	9.50	1.19
BS_Cheese	6.46	4.14
Cascin	20.22	28.83
Butter	10.83	8.42
BMP	49.76	48.85
AMF	12.70	9.16
WPC	38.03	49.73
Lactose	2.54	2.47
Average	20.55	21.07

Table 19: Forecast errors for VECM model using all BCPs and exchange and interest rates.

Finally we estimate a VECM model using all BCPs, the exchange rate and both the short and long term interest rates. Given that using the exchange and interest rates

alone as explanatory variables does not significantly improve forecast performance, we would not expect using these variables in conjuction with all BCPs to provide significantly better forecasts than using the BCPs alone. The results (Table 19) confirm this.

The average MAPE and point forecast error are slightly worse than when the VECM model is estimated without using the exchange and interest rates.

As an aside it is interesting to note the forecasts of the exchange and interest rates under this VECM model. Intuitively we would not expect lagged values of the BCP price series to aid in forecasting NZ market rates. Although dairy exports represent a large portion of New Zealand's exports, it seems unlikely world dairy commodity prices are a significant driver of the exchange rate or interest rate. Table 20 tends to confirm this. The forecasts of the exchange and interest rates using BCPs as explanatory variables are not better than the ARIMA models.

BCP	Forecast Error	
	MAPE	Point Forecast % Error
NZDUSD	5.76	17.12
BB_RATE	8.15	8.99
TBOND	15.30	18.79
Average	9.74	14.96

Table 20: Forecast error for exchange and interest rates using VECM model incorporating all BCPs and FX and interest rates.

5.3   Monte-Carlo simulation

We now illustrate the application of the model to Fonterra's FX hedge portfolio[43]. We use the Monte-Carlo model to calculate the maximum foreign exchange loss over the 15-month horizon from Feb-04 to Apr-05 that would be expected 95% (and 99%) of the time, i.e. the greatest loss expected with 95% (99%) confidence as a result of Fonterra's FX hedging activities.

---

[43] Monthly forecast debt levels were unavailable at the time of writing so we do not apply the model to the interest rate portion of the hedge portfolio.

The relevant forecast variable is the NZDUSD exchange rate. Figure 2 illustrates 20 price paths generated by the VECM model with bootstrapped residuals.
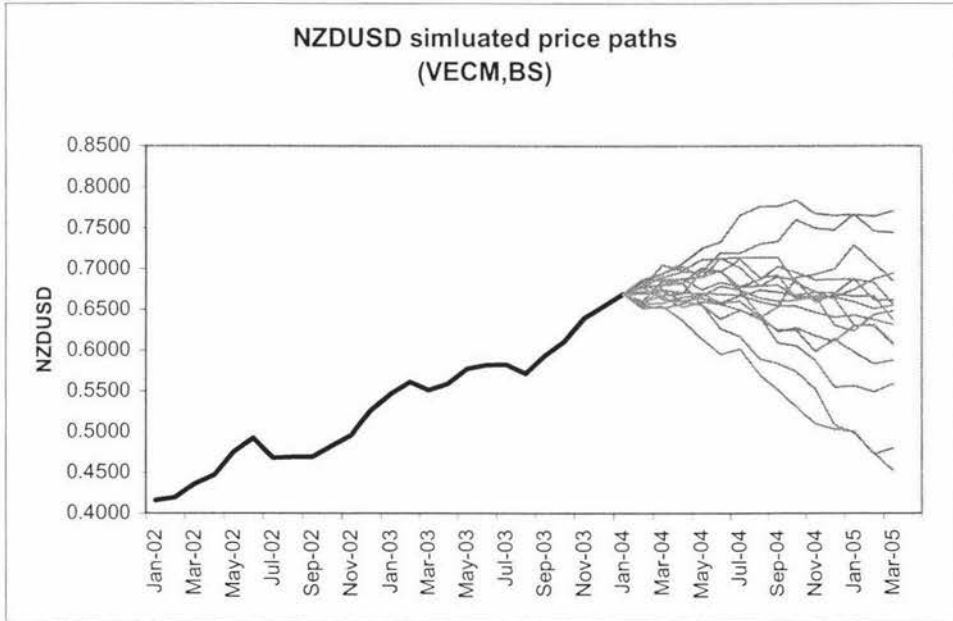

Figure 2: 20 simulated NZDUSD price paths using VECM,BS model.

Table 21 shows selected summary statistics for the simulated distribution of returns from each forecast model and the historical distribution.

	Historical	ARIMA,BS	ARIMA,N	VECM,BS	VECM,N
N	103	280	280	280	280
Mean	-0.0002	0.0011	0.0010	0.0014	0.0010
Std. Dev.	0.0305	0.0311	0.0269	0.0223	0.0235
Skewness	-0.2246	0.0916	-0.0011	-0.3536	-0.0110
Kurtosis	0.8388	2.1717	0.6358	0.9800	-0.4070
t-stat (mean difference)		-0.37	-0.36	-0.49	-0.37
		(0.7090)	(0.7220)	(0.6250)	(0.7150)

Table 21: Selected summary statistics for historical and simulated return distributions. Simulation returns are obtained from 20 price paths.

The forecast models are ARIMA with bootstrapped residuals (ARIMA,BS), ARIMA with normal residuals (ARIMA,N), VECM with bootstrapped residuals (VECM,BS) and VECM with normal residuals (VECM,N). We t-test the difference between mean

of the historical return distribution and the mean of the forecast distribution for each model and find the difference to be not significantly different from zero in each case.

Figure 3 shows the forecast cumulative return distribution from each model compared to the historical cumulative return distributions. The cumulative distribution of simulated returns appears comparable to the historical distribution. Jones (2004) proposes a means of testing the concordance of the unconditional distribution of simulated returns with the observed unconditional distribution. His method involves applying the bootstrap methodology to generate many return series and then simply counting the number of times each of the first four moments is greater than the original return series. A good forecast model should generate return distributions with each moment greater than the original series close to half the time. Although we do not pursue this approach here, his method suggests a simple means by which we might determine how closely the unconditional distribution of the simulated returns matches the observed distribution.
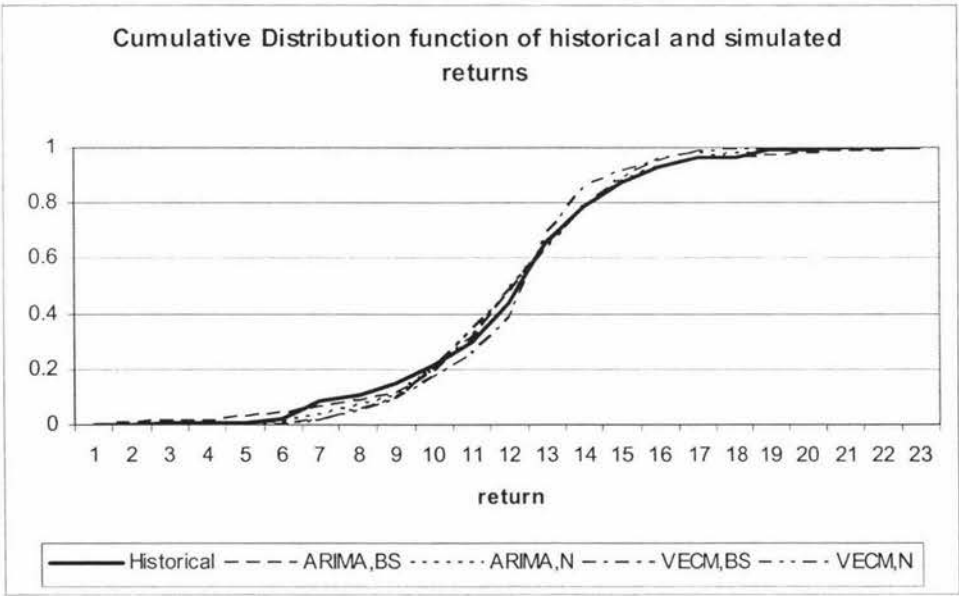
Figure 3: Cumulative distribution function of simulated returns (from 20 simulation paths) versus the observed cumulative distribution function.

To obtain our final VaR number we simulate 10,000 price paths for the NZDUSD exchange rate for 15 months from Feb-04. Each price path is iteratively fed into the

FX hedge portfolio model to generate the gain/loss on the portfolio under that specific exchange rate scenario. Figure 4 displays the forecast distribution of FX hedge portfolio NZD gains/losses using the VECM,BS model with 1000 simulation paths.
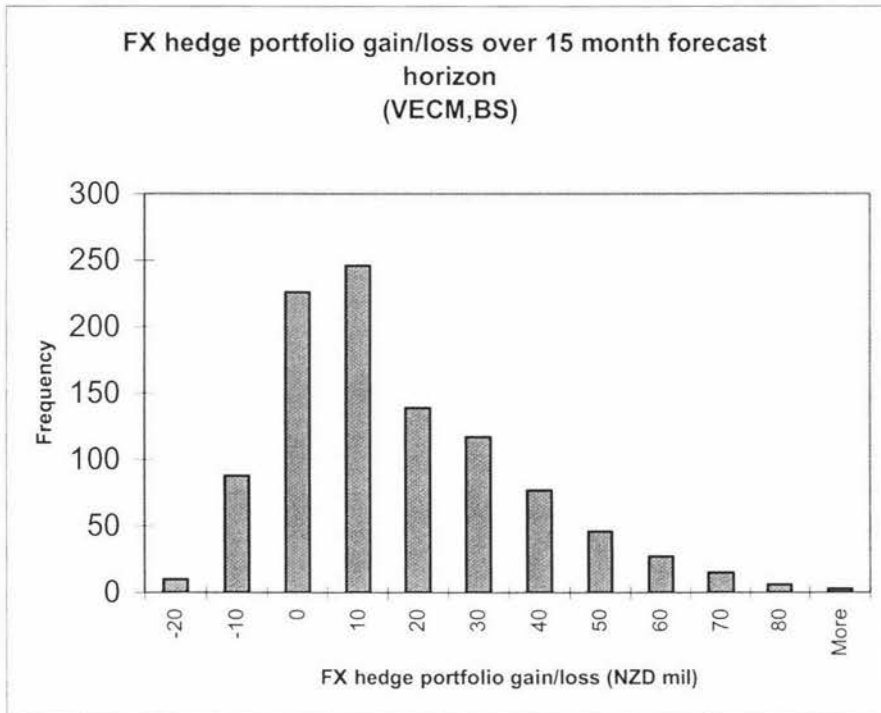

Figure 4: Distribution of simulated FX hedge portfolio gain/loss over 15-month forecast horizon using the VECM,BS model with 1000 simulation paths.

From the distribution we can calculate the 95% and 99% VaR by ordering the gain/loss figures and selecting the 950[th] and 990[th] worst result respectively. Table 22 displays the 95% and 99% VaR for each of the four forecast models. The 95% VaR of −13.31 million using the VECM,BS model can be interpreted as the greatest amount that Fonterra would expect to lose over the next 15 months as a result of current FX hedges 95% of the time.

It is apparent that while the BS and N models produce similar results, a significant difference exists between the ARIMA and VECM models. The difference lies in the fact that the price paths generated by the VECM models exhibit more than twice the

variation of the ARIMA price paths. Ongoing out-of-sample testing is required to determine which model provides the best description of future variability.

		ARIMA, BS	ARIMA,N	VECM,BS	VECM,N
FX gain/loss (NZD mil)	Minimum	-12.81	-11.66	-25.79	-28.77
	Maximum	37.11	33.73	96.21	74.67
	Average	4.63	4.90	11.36	10.63
	Std. Dev.	6.95	7.02	19.57	17.57
	95% VaR	**-4.79**	**-4.61**	**-13.31**	**-13.25**
	99% VaR	**-7.82**	**-8.01**	**-20.67**	**-23.89**
15-month NZDUSD forecast	Minimum	0.5578	0.5653	0.4727	0.4532
	Maximum	0.7793	0.7682	0.9730	0.9024
	Average	0.6638	0.6650	0.6808	0.6790
	Std. Dev.	0.0321	0.0324	0.0802	0.0750

Table 22: 95% and 99% VaR on Fonterra FX hedge portfolio using VECM,BS model with 1000 simulation paths.

As Table 22 illustrates, the VaR number obtained via Monte-Carlo simulation is conditional upon the forecast model used and the assumptions made. The degree of difference between the results from the four alternative models provides a rough indication of the level of model risk inherent in the result; that is, the sensitivity of the result to the forecast model employed. Similar sensitivity analysis should be conducted around the different assumptions made to get an idea of the impact of each on the result.

6 Conclusion

This project has illustrated the creation of a VaR type model for use within the treasury department of a major corporate. Academic research in the area of VaR is increasingly technical and esoteric, often at the expense of practicality. Our model is developed under the real-world constraints of cost and time pressure, limited data resources and standard software platforms. Thus the project not only provides a framework for the development of a corporate VaR model, but also illustrates the extent to which academic theory is applicable when bounded by real-world constraints.

Throughout the project we have sought to emphasise the *process* by which a corporate VaR model might be developed. With that goal in mind we begin by looking at the problems inherent in the widely used RiskMetrics and CorporateMetrics approaches and propose an alternative Monte-Carlo method using ARIMA and VECM time series forecast models in an effort to improve on these limitations. Such models have the advantage of minimal data requirements and ease of estimation, and are therefore well suited to the corporate environment. More importantly, VECM models take into account the autocorrelation inherent in monthly return series and model the long-term relationship between variables, thus accounting for co-movement between variables. We then show how the model can be applied to Fonterra's Treasury FX hedge portfolio to find the maximum loss over a 15 month period that would be expected 95% and 99% of the time as a result of Fonterra's FX hedging policies.

While, in adopting the approach that we do, we illustrate the method behind the model we stop short of an exhaustive analysis of the robustness of the VaR result. Given the changing composition of the hedge book and fluidity of hedging policies, adequately backtesting the model is difficult. In addition, Fonterra has only existed in its current state since 2001 so a lack of historical data compounds the problems. We are content to look briefly at sensitivity to different forecast model choices and leave a more in-depth quantitative assessment of model accuracy and robustness as an avenue for further research. It is our opinion that quantitative tests of the "accuracy" of a model of this nature are dubious at best and worse, can often be misleading. The importance

of the human element in assessing the model cannot be over-emphasised. Here we return again to the theme of practicality that runs through this report. The first test of the accuracy of the model should be whether the results make "sense" to the users of that model. Likewise, the assumptions behind any model should be rigorously tested on an ongoing basis by those using the model. Like any model, the Fonterra VaR model is only as good as the assumptions on which it is built.

7 References

Abuaf, N., & Jorion, P. (1990). Purchasing power parity in the long run. *Journal of Finance, 45,* 157-174.

Adler, M., & Lehman, B. (1983). Deviations from purchasing power parity in the long run. *Journal of Finance, 38,* 1471-1487.

Audrino, F., & Barone-Adesi, G. (2002). A multivariate technique to improve VaR computation in equity markets. Working paper, University of Southern Switzerland.

Bannerjee, A., Dolado, J., Galbraight, J., & Hendry, D. (1993). *Co-integration error-correction and econometric analysis of non-stationary data.* Oxford: Oxford University Press.

Berkowitz, J. (2001). Testing density forecasts with applications to risk management. Working paper, University of California, Irvine.

Black, F., & Karasinski, P. (1991). Bond and option pricing when short rates are lognormal. *Financial Analysts Journal, July/August,* 52-59.

Boyle, P., Broadie, M., & Glasserman, P. (1997). Monte Carlo methods for security pricing. *Journal of Economic Dynamics and Control, 21,* 1267-1321.

Box, G., & Jenkins, G. (1976). *Time series analysis: Forecasting and control.* Rev. ed. San Francisco: Holden Day.

Brennan, M.J.M., & Schwartz, E.S. (1977). Savings bonds, retractable bonds, and callable bonds. *Journal of Financial Economics, 3,* 133-155.

Brennan, M.J.M., & Schwartz, E.S. (1979). A continuous time approach to the pricing of bonds. *Journal of Banking and Finance, 3*, 133-155.

Brennan, W., Lakonishok, J., & LeBaron, B. (1992). Simple technical trading rules and the stochastic properties of stock returns. *Journal of Finance, 47*, 1731-1763.

Brock, M.J.M., & Schwartz, E.S. (1980). Analyzing convertible bonds. *Journal of Financial and Quantitative Analysis, 17*, 75-100.

Brotherton, R. (1994). Monte Carlo monitoring. *Risk, 7* (December), 53-57.

Butler, J., & Schachter, B. (1998). Estimating value at risk with a precision measure by combining kernel estimation with historical simulation. *Review of Derivatives Research, 1*, 371-390.

Campbell, J.Y., & Perron, P. (1991) Pitfalls and opportunities: What macroeconomists should know about unit roots. In: Blanchard, O., & Fischer, B. Eds. *NBER Macroeconomics Annual 6*. Cambridge MA: MIT Press.

Chan, K. C., Karovlt, G.A., Longstaff, F.A., & Sanders, A.B. (1992). The volatility of Japanese interest rates: A comparison of alternative term structure models. *Pacific-Basin Capital Markets Research, Vol III*, 119-136.

Christoffersen, P., & Diebold, F. (1997). Cointegration and long-horizon forecasting. Working paper, IMF.

Christoffersen, P. (1998). Evaluating interval forecasts. *International Economic Review, 39*, 841-862.

Clewlow, L., & Carverhill, A. (1994). Quicker on the curves. *Risk, 7* (May), 63-65.

Constantinides, G.M. (1984) A theory of the nominal term structure of interest rates. *The Review of Financial Studies, 5*, 531-552.

Cox, J.C., Ingersoll, J.E., & Ross, S.A. (1980) An analysis of variable rate loan contracts. *Journal of Finance, 35* (May), 389-403.

Cox, J.C., Ingersoll, J.E., & Ross, S.A, (1985). A theory of the term structure of interest rates. *Econometrica, 53* (March), 385-407.

Covec Limited. (2004). Econometric modelling of dairy prices. Internal Report, Fonterra Cooperative.

Crnkovic, C., & Drachman, J. (1996). Quality control. *Risk, 9* (September), 138-142.

Cryer, J. (1986). *Time series analysis.* Boston: Duxbury Press.

Deaton, A., & Laroque, G. (1992). On the Behavior of Commodity Prices. *Review of Economic Studies, 59*, 1-23.

Dothan, U.L. (1978). On the term structure of interest rates. *Journal of Financial Economics, 6* (March), 59-69.

Dowd, K. (1998). *Beyond value at risk: the new science of risk management.* New York: John Wiley and Sons.

Ellwood, N. (2000). Cointegrating relationships in international stock markets. Unpublished Masters Thesis, University of Otago.

Enders, W. (1995). *Applied econometric time series.* New York: John Wiley.

Fallon, W. (1996). *Calculating value at risk.* Working paper, Wharton School, Financial Institutions Center.

Fama, E. (1965). The behavior of stock market prices. *Journal of Business, 38*, 34–105.

Feldman, D. (1989). The term structure of interest rates in a partially observable economy. *Journal of Finance, 4,* 789-812.

Ferson, W., Simin, T., & Sarkissian, S. (2003). Spurious regressions in financial economics? *Journal of Finance, 58*, 1393-1414.

Fischer, L., Fackler, P., & Orden, D. (1995). Long-run identifying restrictions for an error-correction model of New Zealand money, prices and output. *Journal of International Money and Finance, 14*(1), 127-147.

Fisher, L. (1996). Sources of exchange rate and price level fluctuation in two commodity exporting countries: Australia and New Zealand. *The Economic Record, 73*(21), 345-58.

Frain, J., & Meegan, C. (1996). *Market risk: An introduction to the concept and analytics of value-at-risk.* Mimeo. Economic Analysis Research and Publication Department, Central Bank of Ireland.

Granger, C., & Newbold, P. (1974). Spurious regressions in econometrics. *Journal of Econometrics, 2*, 111-20.

Granger, C. (1981). Some properties of time series data and their use in econometric model specification. *Journal of Econometrics, 16,* 121-30.

Griffiths, W., Hill, R., & Judge, G. (1993). *Learning and practicing econometrics.* New York: John Wiley & Sons, Inc.

Guth, L. (2002). *Twelve things I've learned about enterprise-wide risk management of non-financial firms.* Promotional brochure, NERA Economic Consulting.

Hamiltom, J. (1994). *Time series analysis.* Chichester: Princeton University Press.

Hansen, P. (1982). Large sample properties of generalized method of moments estimators. *Econometrica, 50,* 1029-1054.

Hendricks, D. (1996). Evaluation of value-at-risk models using historical data. *Federal Reserve Bank of New York Economic Policy Review, 2* (April), 39-70.

Hull, J., & White, A. (1990). Pricing interest-rate derivative securities. *Review of Financial Studies, 3,* 573-592.

Hull, J., & White, A. (1998). Incorporating volatility updating into the historical simulation method for value-at-risk. *Journal of Risk, 1* (1), 5-19.

Huizinga, J. (1987). *An empirical investigation of the long run behavior of real exchange rates.* Carnegie-Rochester Conference Series on Public Policy, 27, 149-210.

Jamshidian, F., & Zhu, Y. (1996). Scenario simulation model for risk management. *Capital Market Strategies, 12,* 26-30.

Jesse, E., & Schuelke, J. (2002). *Forecasting class III and class IV milk prices.* Staff paper No. 453, University of Wisconsin-Madison.

Johansen, S. (1988). Statistical analysis of cointegrating vectors. *Journal of Economic Dynamics and Control, 12,* 231-254.

Johansen, S. (1991). Estimation and hypothesis testing of cointegrating vectors in Gaussian vector autoregressive models. *Econometrica, 59,* 1551-1580.

Jones, B. (2004). Is arch useful in high frequency foreign exchange applications? Working paper, Macquarie University.

Jorion, P. (1996). *Value at risk: the new benchmark for controlling market risk.* Chicago: Irwin.

Jorion, P., & Sweeney, R.J. (1996). Mean reversion in real exchange rates: evidence and implications for forecasting. *Journal of International Money and Finance 15*, 535-550.

Joy, C., Boyle, P., & Tan, K. (1995). *Quasi-Monte Carlo methods in numerical finance.* Mimeo, University of Waterloo.

Kim, J., Malz, A., & Mina, J. (1999). *LongRun.* Technical document, RiskMetrics Group.

Kupiec, P. (1995). Techniques for verifying the accuracy of risk management models. *Journal of Derivatives, 3,* 73-84.

Lee, A. (1999). CorporateMetrics. Technical document. RiskMetrics Group.

Lin, C., & Chang-Cheng, C. (2003). Incorporating a GED model into the historical simulation method for value-at-risk. Working paper.

Longstaff, F. (1989). A nonlinear general equilibrium model of the term structure of interest rates. *Journal of Financial Economics, 23,* 195-224.

Longstaff, F., & Schwartz, E. (1992). Interest-rate volatility and the term structure: A two-factor general equilibrium model. *Journal of Finance, 47*(4), 1259-82.

Lopez, J. (1996). Regulatory evaluation of value-at-risk models. Working paper, Research and Market Analysis Group, Federal Reserve Bank of New York.

Mandlebrot, B. (1963). The variations of certain speculative prices. *Journal of Business, 36,* 394–419.

Merton, R.C. (1973) Theory of rational option pricing. *Bell Journal of Economics and Management Science, 4*, 141-183.

Mina, J., & Xiao, J. (2001). *Return to RiskMetrics: The evolution of a standard.* Technical Document. Risk Metrics Group.

Morgan Guaranty Trust Company (1996). *RiskMetrics Technical Document* (4th ed.). Technical Document. Risk Metrics Group.

Moskowitz, B., & Caisch, P. (1995). Smoothness and dimension reduction in quasi-Monte Carlo methods. Working paper, UCLA.

Myers, R.J. (1994). A time series econometrics and commodity price analysis: A review. *Review of Marketing and Agricultural Economics, 6,* 167-181.

Neiderreiter, H., & Xing, C. (1995). *Low-discrepency sequences and global function fields with many rational places.* Mimeo, Austrian Academy of Sciences.

Owen, A., & Tavella, D. (1996). *Scrambled nets for value at risk calculations.* Mimeo, Stanford University.

Pan, J. & Duffie, D. (1997). An overview of value at risk. *Journal of Derivatives, 4*(3), 7-49.

Papageorgiou, A., & Traub, J. (1996). Faster valuation of financial derivatives. *Risk, 9* (June), 63-65.

Paskov, S., & Traub, J. (1995). Faster valuation of financial derivatives. *Journal of Portfolio Managment, 22*, 113-120.

Peter J.M. (1990). Managing currency exposure: The case of western mining. *Journal of Applied Corporate Finance, Winter*, 29-34.

Pigott, C., & Sweeney, R., (1985). Purchasing power parity and exchange rate dynamics: some empirical. results. In: Arndt, S., Sweeney, R., Willett, T., Eds. *Exchange Rates, Trade and the US Economy*. Cambridge: Ballinger.

Pitt, W. (2001). Future shock: Measuring cash flow volatility. *MCC Viewpoint, Spring.*

Powell, J., Shi, J., & Smith, T. (2004). Dividend persistence and return predictability. Working paper, Massey University.

Pritsker, M. (1996). *Evaluating value at risk methodologies: accuracy versus computational time*. Mimeo, Board of Govenors of the Federal Reserve System.

Pritsker, M. (2001). The hidden dangers of historical simulation. Working paper, University of California, Berkeley.

Ripp, A., Moore, L., & Cohn, P. (2001). Cash flow-at-risk. *International Finance & Treasury, 26* (March), 11-13.

Rogoff, K.S., (1996). The purchasing power parity puzzle. *Journal of Economic Literature, 34,* 647-668.

Roll, R., (1979). Violations of purchasing power parity and their implications for efficient international commodity markets. In: Sarnat, M., & Szego, G. Eds. *International Finance and Trade*. Cambridge: Ballinger.

Rose, A.K., (1996). International finance and macroeconomics. *NBER Reporter, Winter,* 1-6.

S&OP Group. (2004). *Price forecasting update*. Internal Report, Fonterra Cooperative.

The SAS Institute. (2004). *http://v8doc.sas.com/sashtml/*. Online.

Schaefer, S., & Schwartz, E. (1984). A two factor model of the term structure: An approximate analytical solution. *Journal of Finance and Quantitative Analysis, 19,* 413-424.

Siddique, A., & Sweeney, R.J. (1998). Forecasting real exchange rates. *Journal of International Money and Finance, 17,* 63-70.

Stein, J., Usher, S., LaGattuta, D., & Youngen, J. (1997). A comparables approach to cashflow-at-risk for non-financial firms. Working paper #39, NERA Economic Consulting.

Stock, J.H. (1995). *Point forecasts and prediction intervals for long horizon forecasts,* Manuscript, J.F.K. School of Government, Harvard University.

Sundaresan, S. (1984). Consumption and equilibrium interest rates in stochastic production economies. *Journal of Finance, 3,* 77-92.

Sweeney, R. (2001). Mean reversion in G-10 nominal exchange rates. Working paper, Georgetown University, Washington, D.C.

Tezuka, S. (1994). *A generalisation of Faure sequences and its efficient implementation.* Research report rto105, IBM Research, Tokyo.

Tomek, W. (2000). Commodity prices revisited. *Agriculture and Resource Economics Review, 29*(2), 125-37.

Tomek, W., & Myers. R.J. (1993). Empirical analysis of agricultural commodity prices: A viewpoint. *North Central Journal of Agricultural Economics, 15,* 181-202.

Tse, Y. (1995). Some international evidence on the stochastic behaviour of interest rates. *Journal of International Money and Finance, 14*(5), 721-738.

Vasicek, O. (1977). An equilibrium characterization of the term structure. *Journal of Financial Economics, 5* (November), 177-188.

Venkataraman, S. (1997). Value at risk for a mixture of normal distributions: the use of quasi-Bayesian estimation techniques. *Economic Perspectives, March/April*, 2-13.

Verbeek, M. (2000). *A guide to modern econometrics.* Chichester: Wiley.

Wilson, T. (1994). Plugging the gap. *Risk, 7* (October), 74-80.

Yang, S.R., & Brorsen, W.B. (1992). Nonlinear dynamics of daily cash prices. *American Journal of Agricultural Economics, 74*(3), 706-715.

Zangari, P. (1996a). A value-at-risk analysis of currency exposures. *RiskMetrics Monitor, Second Quarter*, 26-33.

Zangari, P. (1996b). When is non-normality a problem? The case of 15 time series from emerging markets. *RiskMetrics Monitor, Fourth Quarter*, 20-32.

Zangari, P. (1996c). An improved methodology for measuring VaR. *RiskMetrics Monitor, Second Quarter*, 7-25.

Zangari, P. (1996d). A VaR methodology for portfolios that include options. *RiskMetrics Monitor, First Quarter*, 4-12.

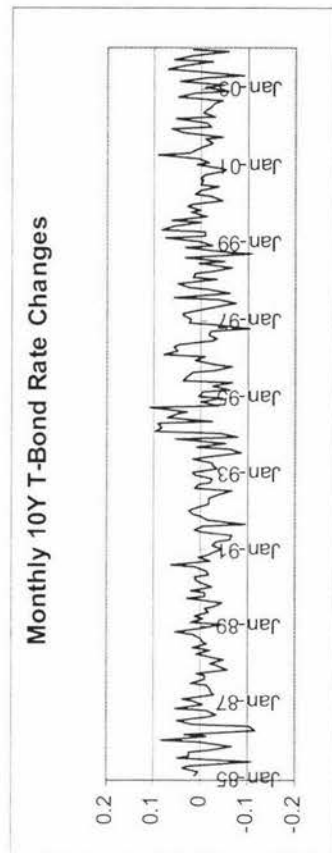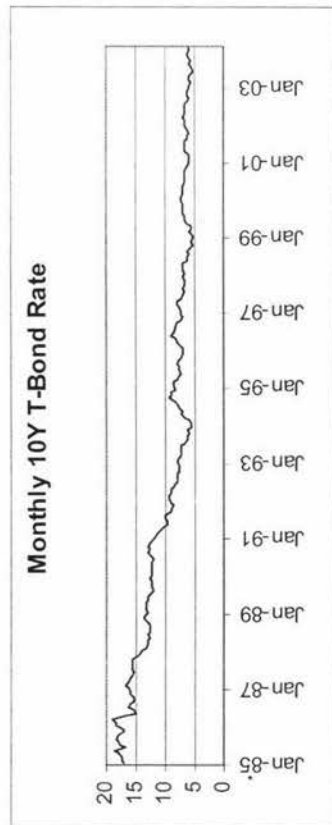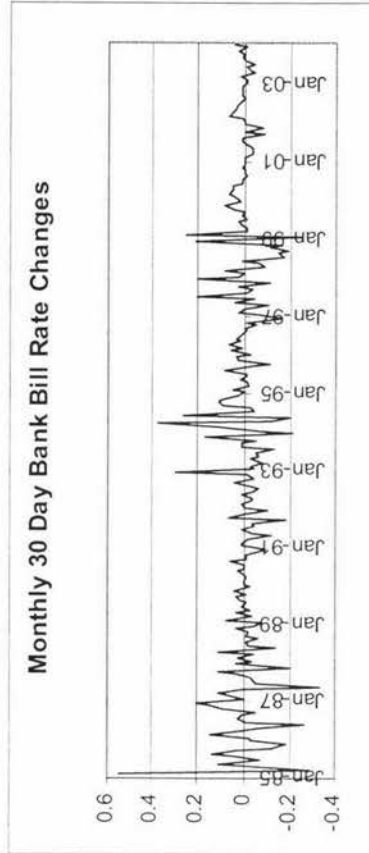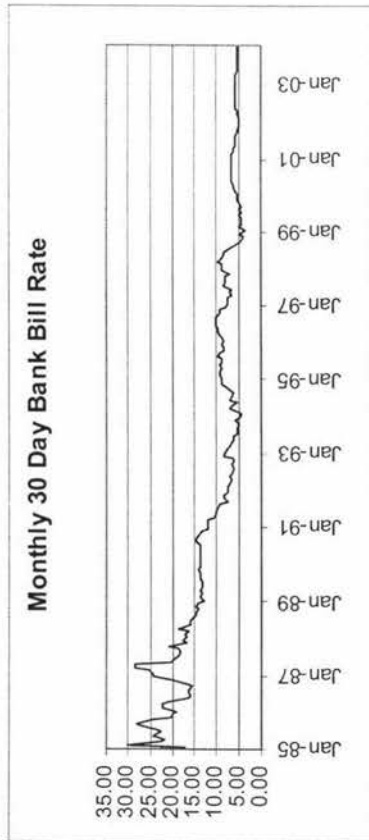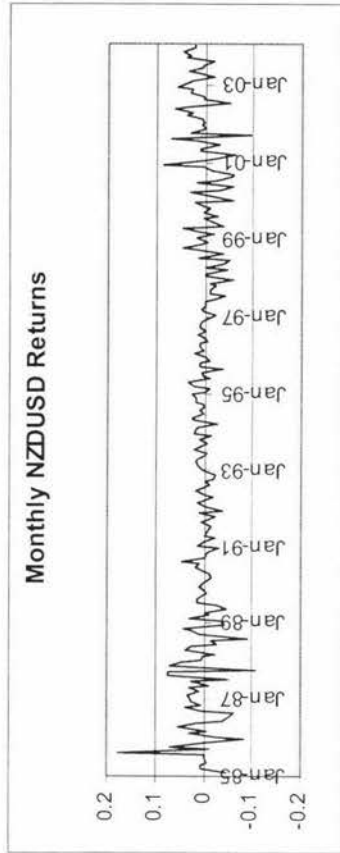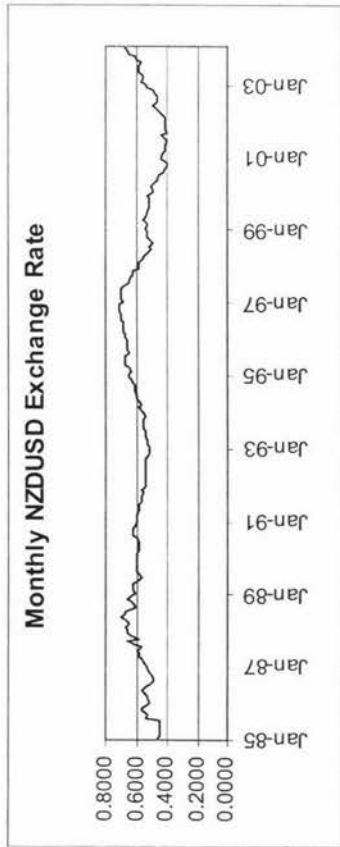# Appendix A: Treasury hedge portfolio

A1. FX hedge portfolio
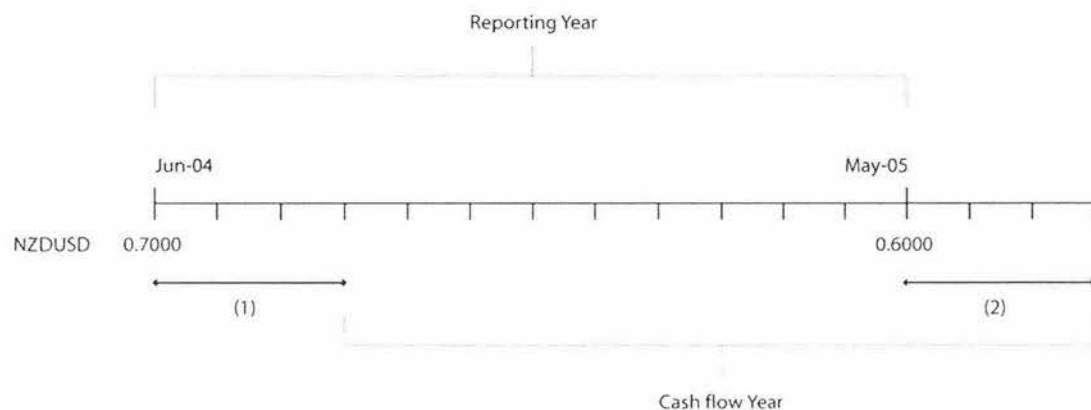
The FX gain/loss in month t is the sum of the gain/loss on forwards and options expiring in month t, i.e.

$$FX\ gain/loss_t = NZD\ amount\ forwards_t \times (Actual\ rate_t - Forward\ rate_t)$$
$$+ NZD\ amount\ options_t \times \max(0, Actual\ rate_t - Strike\ rate_t).$$

The cash flow year is offset by approximately three months from the financial year so Fonterra fully hedges 15 months. As a result, three months of outstanding forwards and options relating to the current financial year are revalued at the end of year exchange rate. For example, the options and forwards expiring in the three months labeled (2) on the diagram are revalued using a strike and forward rate of 0.6000 and are included in the 2004 reporting period. This also means that the forwards and options in the first three months of each reporting year are valued at the exchange rate as at the end of the previous reporting period. For example, the gain/loss on options and forwards in the period labeled are (1) is calculated using a strike and forward rate of 0.7000. The revaluation is for accounting purposes only and does not affect the total gain/loss from any given position.
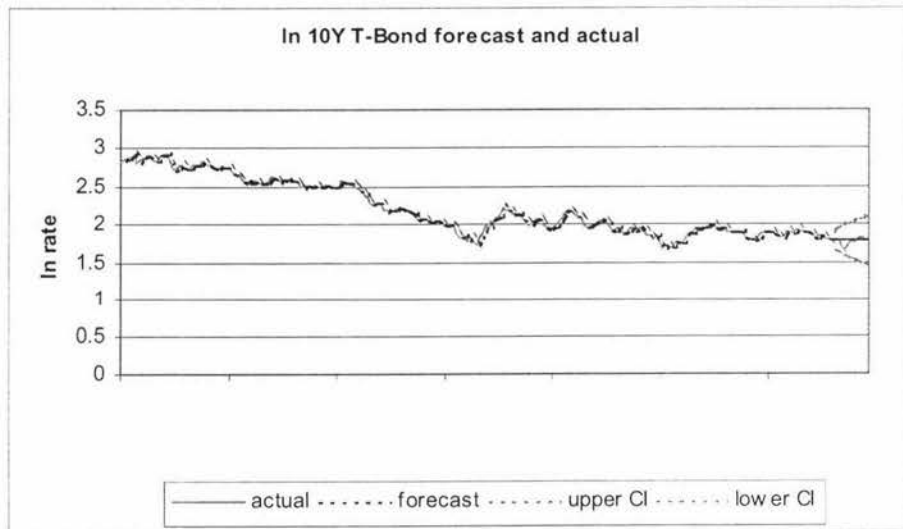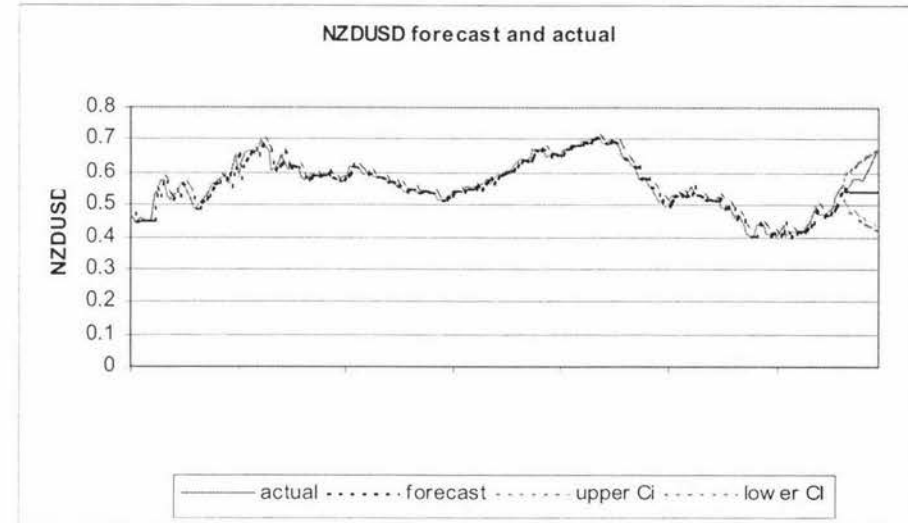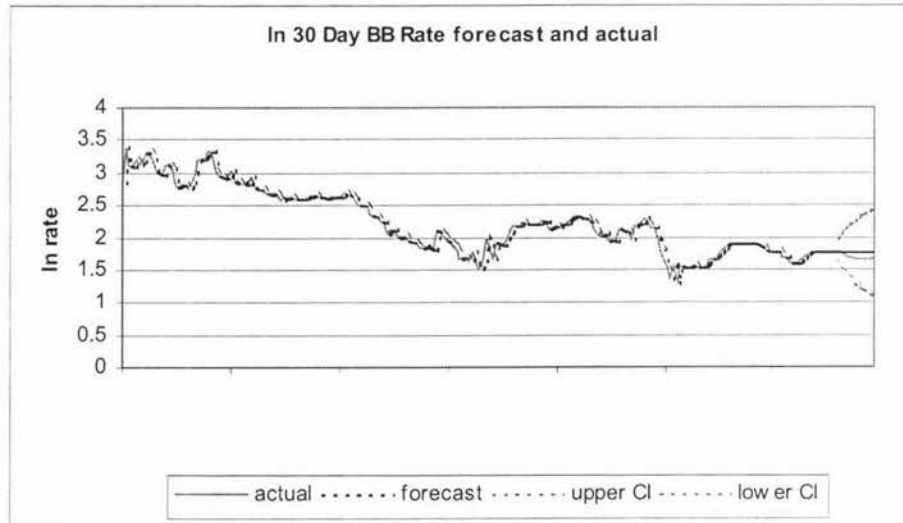
Further description of the FX hedge portfolio and an overview of the interest rate hedging policy are commercially sensitive and are included in Appendix D only.

Appendix B: Time series graphs

Monthly NZDUSD Exchange Rate

Monthly NZDUSD Returns

Monthly 30 Day Bank Bill Rate

Monthly 30 Day Bank Bill Rate Changes

Monthly 10Y T-Bond Rate

Monthly 10Y T-Bond Rate Changes

Appendix C: ARIMA model forecast graphs

In 30 Day BB Rate forecast and actual

actual ······ forecast ······ upper CI ······ lower CI

NZDUSD forecast and actual

actual ······ forecast ······ upper CI ······ lower CI

In 10Y T-Bond forecast and actual

actual ······ forecast ······ upper CI ······ lower CI

Appendix D: Confidential information

This section is confidential and is only to be accessed by persons who have signed a confidentiality agreement with Fonterra Cooperative. Please contact the author for further information.