

Value at Risk Incorporating Dynamic Portfolio Management

Stephen Lawrence

Carroll School of Management

Boston College

314 Fulton Hall

Chestnut Hill, MA 02467

Tel: (617) 552-2050

Email: stephen.lawrence@bc.edu

Web: <http://www2.bc.edu/~lawrenst>

Final Draft : June 26, 2000

I would like to thank the following people for their valuable comments and insight : George O. Aragon, Jushan Bai, Pierluigi Balduzzi, David Belsley, Stefan Bokor, Eric Jacquier, Arthur Lewbel, Ludan Liu and Serena Ng.

Abstract

Value at Risk has established itself as one of the key measures of risk in both academic literature and the business world. Recent academic research has highlighted some serious shortcomings of Value at Risk, in both the inconsistency of results across models and the failure to account for feedback effects from risk management measures on trading behavior. This paper uses a simulated model of exchange rates to investigate Value at Risk measurements, the dynamic behavior of foreign exchange traders and the effects of risk measures on trading behavior. We extend the calculation of Value at Risk to explicitly account for the dynamic nature of the portfolio holdings and account for the restrictions placed by risk management benchmarks on the value at risk of a portfolio.

Keywords: Value at Risk, Risk Management, Dynamic Portfolio, ARCH, Exchange Rates, Risk Budgeting.

1. Introduction

Value at Risk is becoming increasingly popular as a measure of the potential losses of a portfolio or business. Government and banking regulations set up in the past decade have required a Value at Risk measure to be reported or used in determining capital requirements. The introduction of JP Morgan's RiskMetrics system in 1995 increased the profile of Value at Risk substantially, and as the importance of Value at Risk has increased, so has the volume of academic literature developing, supporting or criticizing this risk measure. In his book on Value at Risk, Jorion (1997) defines Value at Risk as¹:

“the expected maximum loss (or worst loss) over a target horizon within a given confidence interval”

One of the major concerns about Value at Risk calculations is the lack of consistency between different parameterizations and implementations of the risk measure. In their overview of available methods, Duffie and Pan (1997) note that:

“Value at risk is merely a benchmark for relative judgements, such as the risk of one desk relative to another, the risk of one portfolio relative to another, the relative impact on risk of a given trade, the modeled risk relative to the historical experience of marks to market, the risk of one volatility environment relative to another, and so on. Even if accurate, comparisons

¹ For the sake of this paper, we will assume that the confidence level is 5%, and the target horizon is assumed to be 20 periods. Hence when discussing Value at Risk we will be implicitly referring to the 5th percentile of the distribution of the difference between today's portfolio value and the estimated portfolio value 20 periods hence.

The 5% level is the level used in JP Morgan's RiskMetrics and is used here as a tradeoff between measuring particularly bad outcomes and having enough data to infer results from. As the confidence level decreases towards zero, the amount of data available to characterize the tail to the left of the quantile decreases and the MSE of the risk measure increases.

The data used to calibrate the model is daily data; hence the horizon can be considered to be one month. We will assume that the investor's horizon is 20 periods or one month. However the simulation does not rely on the frequency of the data and hence the horizons can be thought of purely in terms of periods. The number of periods is important for the calculations in this paper as it changes the expected number of portfolio adjustments and hence the disparity between a dynamic and a fixed model.

such as these are specific to the time horizon and the confidence level associated with the value-at-risk standard chosen.”

Hence within a particular model, the results produced should be used as a guideline rather than a tool to make exact predictions. This becomes even more apparent when identical portfolios are evaluated using different models. Although some discrepancies between results can be attributed to the various simplifying assumptions made, the magnitude of the difference between models can be startling, as Beder (1995) highlights. She compares the Value at Risk in three portfolios using eight different VaR calculations and finds that results can vary by as much as a multiple of 6x to 14x between the smallest and the largest estimate.

Basak and Shapiro (1998) attack another weakness of the current Value at Risk based models, namely that the models fail to account for the effect of the risk measure itself on trader behavior. Their paper explains that Value at Risk places all of the attention on the percentile in question and ignores the question of what happens below that level. As an example consider two portfolios, identical except for the losses incurred in the left-hand tail (see Figure 1). The Value at Risk in the two portfolios is identical, however under normal circumstances it would be preferable to choose the portfolio with the smaller expected loss.

Basak and Shapiro detail how a portfolio manager limited by a VaR based risk measure² has an incentive to ignore the returns below a certain quantile and concentrate on maximizing the return in good states of the world. In extremely severe states of the world, where losses are expected, the risk measure has the unfortunate side effect of actually increasing the losses – exactly what the measure is designed to prevent.

The model proposed by Basak and Shapiro (1998) is, however, based on a continuous time model with one risky and one riskless asset. In this setting, if the penalty is severe enough, the portfolio manager has both the incentive and the means to remain within the Value at Risk limit at all times with probability one. This is achieved by adjusting the portfolio position infinitesimally and continually as the state of the world worsens. Although the model is important from a theoretical point of view, the continual adjustment of portfolio positions is unrealistic in practice.

² I.e. their performance is rewarded as long as the Value at Risk of their portfolio remains below a certain level.

Transaction costs and information-processing constraints would not allow the portfolio manager to always adhere to the constraints; the question of what happens to the portfolio if the constraint is exceeded is never addressed.

Wang (1998) works in discrete time to develop new models that take into account dynamic aspects of trading behavior. He establishes a theoretical basis upon which to build a dynamic risk measure. Wang's likelihood based measure encompasses traditional measures, including Value at Risk, but possesses qualities that perform better when handling multi-period risk. Cvitanic and Karatzas (1998) also look at the theory behind dynamic risk measures and investigate the properties of risk measures in a Bayesian setting.

The main focus of current dynamic trading and value at risk literature is on historical back-testing and mathematical models. Consequently, past studies have not investigated fully the potential differences between the risk measures set and the observed risk or potential risk resulting from a dynamically managed portfolio. There is also potential to investigate the effects on Value at Risk models when the IID normal returns assumption is relaxed.

This paper investigates the ex-post performance of Value at Risk based risk management techniques by Monte Carlo simulation. Sample paths from two foreign exchange rates are generated using IID normal data, a bivariate GARCH model (with t distributed innovations) and bootstrapped historical data. These three models are used to test the sensitivity of the risk measures to the simulation technique.

We consider two models of portfolio manager behavior and estimate the Value at Risk of the portfolios under different risk budgeting constraints and different policies for handling excessive risk.

We find that the Value at Risk of a dynamic portfolio is strongly affected by the investment style of the manager and the risk measure used. We find that when the portfolio is regularly adjusted, the ex-post Value at Risk is often vastly different from the Value at Risk specified in the risk measure. We also find that the method used to estimate the Value at Risk for the constraint has a large effect on the resulting distribution.

The layout of the remainder of the paper is as follows: In section 2, we explain the data used to parameterize the model. In section 3 we outline how the exchange rates are simulated and in section 4 we introduce the proposed models of trader behavior. In section 5 we report the results of the simulation and in section 6 we conclude.

2. Data

For this paper we will consider the exchange rates of the US Dollar against both the British Pound and the Japanese Yen. As an initial data set for the model, we use daily levels taken from DataStream for the period Jan 1990 – Jan 1994 (1000 observations). Any gaps in the data have been removed and the bivariate series is treated as continuous data. To transform the series into stationary data, first differences of the log rates are taken. The level as of January 6, 1994 is used as a reference level for time 0 in the model whenever the transformation needs to be reversed.

Time series plots of the data are shown in Figure 2 and Figure 3. Partial cross-correlograms of the are shown for the returns (Figure 4) and squared returns (Figure 5). Statistical properties of the series are given in Table 1.

From Table 1 we see that the data for the returns of US\$/GB£ is leptokurtic. This is further illustrated in Figure 6 and Figure 7, which show the distribution of the returns and the Q-Q plots of the returns against a normal distribution. The partial correlograms show that AARGH! It is clear from the Q-Q plots that the data cannot be approximated appropriately by a normal distribution. This is an issue that is addressed when we simulation the data.

3. Simulation of Data

To test the performance of the various models, we need to produce suitable simulated time series to represent the two exchange rates used in the model. An ideal simulation would capture the general properties of the distribution of returns as well as the time dependent properties, such as conditional heteroskedasticity.

Bootstrapping from the sample data is one method of generating simulated paths possessing similar characteristic to the original data. Random draws are taken from the vectors (y_{1t}, y_{2t}) , to

form a sample, $\{(y_1^{(1)}, y_2^{(1)}), \dots, (y_1^{(\tau_{\max})}, y_2^{(\tau_{\max})})\}$ of the required path length (τ_{\max}). Exchange rates can then be recovered by applying the following transformation:

$$\tilde{x}_{it} = \exp\left\{-\ln(x_{i,init}) - \sum_{\tau=1}^{\tau_{\max}} y_i^{(\tau)}\right\} \quad i = 1, 2 \quad \text{and } x_{i,init} \text{ is the starting level for the path.}^3$$

Bootstrapping produces a model whose distribution of returns asymptotically matches the sample returns without requiring any assumptions about the underlying distributions. However, as the returns are taken randomly, any time specific characteristic such as correlation and conditional heteroskedasticity are lost. We will use bootstrapping as a yardstick by which we will judge the performance of our parametric models.

One possible solution to the lack of correlation in returns of a bootstrap model is to take block bootstraps. A block bootstrap of length p involves taking chains of length p from the sample and pasting them together. This will help incorporate some of the correlations (up to lag p), however the joins between blocks break any time dependent patterns and conditional heteroskedasticity is not properly captured. Bootstrap methods capable of circumventing these problems have been investigated⁴.

For our purposes, it will be easier to use an appropriate parametric time series model that captures the necessary features of the data. The GARCH(p, q) model, as proposed by Bollerslev (1986) is suitable for describing exchange rates⁵ as it is capable of capturing time variation of the conditional variance of the exchange rates.

The BEKK(p, q) multivariate extension of the GARCH(p, q) model is parameterized as follows for $p = q = 1$:⁶

³ When the model was estimated, the exchange rates used were USD/GBP and USD/JPY. However, inverting the rates means that x_{1t} is the price of £1 in US Dollars at time t . This is far more intuitive for application to portfolio theory (as pounds and yen can be thought of as two risky assets). The minus signs inside the parentheses account for this, and these new rates will be denoted by $\tilde{x}_{it} = 1/x_{it}$.

⁴ Berkowitz (1996) gives a lengthy overview of bootstrapping methods that can be used with time series, including the block bootstrap.

⁵ Stochastic Volatility models could also be considered here.

⁶ Taken from the S+GARCH User's Manual (1996)

$$y_t = c + \varepsilon_t, \quad \text{var}(\varepsilon_t | \varepsilon_{t-1}, \dots) = V_t$$

$$V_t = AA^T + A_1(\varepsilon_{t-1}\varepsilon_{t-1}^T)A_1^T + B_1V_{t-1}B_1^T$$

The parameterization (c, A, A_1, B_1) ensures that the variance-covariance matrix, V is always positive definite. Typically the conditional distribution of the errors, ε_t is assumed to be multivariate normal. However in this situation it is better to assume that the residuals follow a multivariate t distribution (where the degrees of freedom becomes a parameter to be estimated). This increases the weight in the tails of the distribution and is more appropriate for Value at Risk calculations.

Estimation of the parameters in a BEKK(1, 1) model was done using S+GARCH. Details of the estimated model are given below in Table 2. Plots of the model residuals and (given in

Figure 8) show that the model has accounted for most of the conditional variance in the model. The Q-Q plots show that a multivariate t distribution with 5.23 degrees of freedom explains the error terms well.

With these parameters we can begin to simulate data from a BEKK(1,1) model by generating sequences of vectors taken from a multivariate t distribution⁷ and substituting these errors back into the above sequence of equations. The model requires an initial covariance matrix to begin the simulation, however S-Plus circumvents this by allowing for a burn-in period before taking the sample path. A method that might be more desirable is to use the estimated covariance matrix from the original data as a starting point for the simulation. However this option was not available with the software used, and the benefits from tying in the model with the empirical data are less than if we were using the model to predict the exact Value at Risk for a period.

The timeline of the samples is shown below in Figure 9. For each run, we generate a historical path (periods -100 to 0) which is used to calibrate the model of investors. Within each Monte Carlo run, several continuation paths are generated (periods 1 to 20) to provide different realizations for the calculation of Value at Risk. Each Monte Carlo run contains one historical path and 1000 future paths. The historical path is chosen randomly so that, by repeating the simulations, we can test the robustness of the model to initial conditions. The level at period 0 is fixed in all cases at (0.6748, 112.7000) to provide a reference level to compare the final results.

By calculating all risk measures inside the same loop, we can compare the individual results for each model against the other models using the same data to ascertain the source of any heterogeneity in the results. Running the simulation several times will allow us to calculate the Value at Risk of each portfolio model and average the results to check for systematic differences between models.

⁷ Unfortunately, this option is not implemented in S+GARCH, so we wrote our own C++ routines and replaced the normal samples with samples taken from the standardized multivariate t distribution. To generate vectors from a multivariate t distribution, we used the following transformation :

$$\epsilon_t = \frac{(X_{1t}, \dots, X_{nt})}{\sqrt{Y_t / \nu}}$$

Where X_{it} are independent drawings from a standard normal and Y_t is an independent chi-squared distribution with ν degrees of freedom.

4. Models of Investor Behavior

The simulated data is used as an underlying engine to drive the dynamics of the behavioral models, which are described below. As explained in the introduction, the aim of the paper is to assume certain properties of the portfolio manager and test the effects of these on the overall Value at Risk of the portfolio. We will start with a very basic model of the investor and work towards more complicated models. To make comparisons between models however we need some further restrictions and assumptions.

The following trading model will be assumed in all models. Each investor starts with \$1000 split equally between pounds and yen (i.e. £337.40 and ¥56,350). The investor may at the end of each period change his holdings by trading costlessly between the two currencies. The investor's portfolio value is measured in US dollars, however the investor is not allowed to hold any dollars directly in his portfolio. This model was chosen to provide interesting behavior when two exchange rates change in their relative properties – a risk-averse investor should be persuaded in certain situations to change his holdings if one currency becomes particularly volatile.

The model is designed to be an abstract model of the decisions facing the average portfolio manager: the manager is unlikely to find himself in a situation where they are required to transfer funds from pounds to yen on a weekly basis. However, if the model is viewed as representing the flexible part of a firm's foreign market exposure⁸, then the value at risk calculations give an estimate of the extra risk involved for the firm in the timing of its cash-flow transfers.

Our approach to calculating value at risk will follow the Monte-Carlo simulation method. At each period the value of the portfolio will be adjusted according to the specified trading model. The value at risk of the portfolio is the 5th percentile of the losses incurred across the individual simulations.

Investor 0 - The Benchmark Model

As a benchmark we consider an agent who maintains the same cash flow throughout. The only risk involved here is that inherent to the foreign exchange rates; the value at risk calculated here is exactly the same as the standard calculations mentioned in the introduction.

Investor 1 – Utility Maximizing Agent

We assume that the investor is rewarded by a convex function of his portfolio value in the next period⁹. This is a myopic strategy, chosen to keep the resulting optimizations simple. It can be justified however, if the investor's perception of future decisions on the final value is limited when compared with the current period's decision¹⁰. A limit on the maximum and minimum exposure to each currency is set by restricting θ , the decision variable, to remain between θ_0 and θ_1 .¹¹ The investor in this model chooses his holdings each week so as to maximize his expected utility given by:

$$EU_1 = \mathbf{E}_t(Z_{t+1}^2)$$

$$Z_{t+1} = \theta \frac{Z_t}{x_{1,t}} \tilde{x}_{1,t+1} + (1-\theta) \frac{Z_t}{x_{2,t}} \tilde{x}_{2,t+1}$$

Here Z_{t+1} is a random variable representing the portfolio at time $t + 1$, $\tilde{x}_{1,t}$ (likewise $\tilde{x}_{2,t}$) is the price of one pound (yen) in dollars at time t and θ is the fraction of the portfolio invested in pounds.

Hence the investors decision is to choose $\hat{\theta}$:

$$\hat{\theta} = \arg \max_{\theta} \mathbf{E} \left[\left(\theta \frac{x_{1,t+1}}{x_{1,t}} + (1-\theta) \frac{x_{2,t+1}}{x_{2,t}} \right)^2 \right]$$

⁸ Perhaps as a result of it's day to day cash flows or current assets.

⁹ The convex function basically treats the portfolio manager's payoff as option based on the value of his portfolio. If the portfolio loses value, the worst that can happen to the manager is that he will be fired and receive nothing. If the portfolio increases in value, the manager is rewarded with bonuses. As a result, his payoff is convex.

¹⁰ A specific example of this is the case where returns are distributed IID and the investor has CRRA. In this setting, the investor's decision can be split into individual periods where the investor maximizes his immediate return.

¹¹ $\theta_0 < \theta_1$. The investor is allowed to sell short if $\theta_0 < 0$ or $\theta_1 > 1$

The constant $V_t^{1-\alpha} / (1-\alpha)$ does not affect the optimization and is removed from the above equation. We now consider the ratio $p_{i,t+1} / p_{i,t}$. In our model, $p_{i,t+1} = p_{i,t} \times \exp\{-c_i - \varepsilon_{i,t+1}\}$, hence the optimization simplifies to

$$\hat{\theta} = \arg \max_{\theta} \mathbf{E} \left[\left(\theta \exp\{-c_1 - \varepsilon_{1,t+1}\} + (1-\theta) \exp\{-c_2 - \varepsilon_{2,t+1}\} \right)^2 \right]$$

$(\varepsilon_{1t}, \varepsilon_{2t})$ is a bivariate t distribution whose variance-covariance matrix is given by a GARCH process. Hence, the above expectation contains a non-trivial distribution function. Exact calculation of the optimal solution will be avoided here, and analysis of the expression will be done in a graphical and intuitive manner.

Given that $(\varepsilon_{1t}, \varepsilon_{2t})$ is generated by dividing a bivariate normal distribution by a function of a chi-squared distribution, the returns should be centered on zero with an elliptical density, the shape of which is determined by the variance-covariance matrix V .

Hence $(\exp\{-c_1 - \varepsilon_{1,t}\}, \exp\{-c_2 - \varepsilon_{2,t}\})$ is distributed around (1, 1), although the shape of the density no longer elliptical. The investor's convex payoff function implies that he will simply choose the portfolio with the highest variance. Without constrictions, the variance increases as $\theta \rightarrow \pm\infty$. We can approximate the expectation in the optimization decision by taking 1000 sample ε_t 's and taking an ensemble average:

$$\hat{\theta} = \arg \max_{\theta} \frac{1}{1000} \sum_{i=1}^{1000} \left[\left(\theta \exp\{-c_1 - \varepsilon_{1,t+1}^{(i)}\} + (1-\theta) \exp\{-c_2 - \varepsilon_{2,t+1}^{(i)}\} \right)^2 \right]$$

The optimal decision is to invest as heavily as possible in pounds iff:

$$f(\theta_0) < f(\theta_1) \text{ where } f(\theta) = \sum_{i=1}^{1000} \left[\left(\theta \exp\{-c_1 - \varepsilon_{1,t+1}^{(i)}\} + (1-\theta) \exp\{-c_2 - \varepsilon_{2,t+1}^{(i)}\} \right)^2 \right]$$

The investor's strategy will be simulated by comparing the two expressions at each point along the sample path and updating the holdings accordingly. At each point, the previous conditional variance matrix V_t and previous return ε_t must be used for the calculations. To decrease the

number of calculations involved at each step, the same realizations from a standard bivariate t distribution are used. They are then pre-multiplied by the square root of the calculated variance-covariance matrix to generate the required sample.

Investor 2 – Simple Day-trader (filter rule) strategy

The next two investors considered in this paper are aimed to look at the management of traders (whether they are within the same company as the risk manager or clients using the risk manager’s company as a brokerage) that use simple day-trading techniques to “extract profits from the market”.

The strategy we will consider is simple filter rule¹², where the investor buys further into a position if the price has risen by more than 1% in the last period, and sells if the price has fallen over 1%. The hope of such a strategy is to catch upward and downward trends before they level off and realize the difference.

Such a strategy is statistically unlikely to make money in the long run, however such strategies are widely discussed, especially on the Internet. In terms of the parameters we have already introduced for the paper, the investor’s decisions are:

Increase holdings in pounds if $\Delta \ln \left(\frac{\tilde{x}_{1t}}{\tilde{x}_{2t}} \right) = \varepsilon_{2t} - \varepsilon_{1t} \in [0.1, 0.2]$.

Decrease holdings in pounds if $\Delta \ln \left(\frac{\tilde{x}_{1t}}{\tilde{x}_{2t}} \right) = \varepsilon_{2t} - \varepsilon_{1t} \in [-0.2, -0.1]$.

The restrictions on exposure to a single currency will be the same as for investor 1, i.e. between θ_0 and θ_1 .

¹² The performance of filter rules is discussed in Campbell, Lo, and MacKinlay (1997).

Risk Limit A – Fixed Limits on exposure

The fixed limit on exposure simply requires that θ_0 and θ_1 be set to fixed limits in all cases. We will assume to begin with that the limit on exposure is $\theta \in (-1, 2)$, i.e. the investor can go short no more than 100% of his net worth in either direction. For investors 1 and 2, this simply means that the relative merits of setting $\theta = -1$ and setting $\theta = 2$ are compared and the best limit taken.

Risk Limit B – Root-t Value at Risk as a Restriction

This investor must now keep the absolute Value at Risk¹³ in their portfolio below a certain level¹⁴.

The challenge in the implementation of this model is to calculate the Value at Risk of the portfolio at each step in a computationally non-intensive manner. Clearly, using a Monte Carlo historical simulation model would not be appropriate as the simulation is already required to carry out two levels of Monte Carlo simulation – adding a third level would make the problem computationally infeasible.

We will use a Value at Risk model that follows the RiskMetrics framework loosely. Let

$$A \sim N_2(0, I_2), B \sim \chi_v^2.$$

$$\text{Then } \varepsilon = \frac{V^{1/2}A}{\sqrt{B/v}} \sim MVT(v, V).$$

The variance-covariance matrix, V , when used for a multivariate t distribution, is a little misleading as:

$$\text{var}(\varepsilon) = \mathbf{E}(\varepsilon - \mathbf{E}\varepsilon)^2 = \mathbf{E}\varepsilon^2 = \mathbf{E}\left[V \frac{A^2}{B/v}\right] \sim V \times MVF_{1,v}$$

¹³ Absolute Value at Risk is defined as the difference between the future value and the **current value** at the 5th percentile. Mean Value at risk is defined as the difference between the future value and the **mean of possible future values**. For simplicity, we assume that the mean is zero, hence the definitions coincide.

¹⁴ We assume that Value at Risk is calculated with respect to a 20 period window and that for the sake of these value-at-risk calculations, the portfolio remains fixed. Of course the aim of the paper is to see how the restrictions affect the actual value at risk, which is what the overall simulation will approximate.

Here $MVF_{1,v}$ is a multivariate F-distribution with degrees of freedom 1 and v respectively. The expected value of a $MVF_{1,v}$ is $v/(v-2)$, hence $\text{var}(\boldsymbol{\varepsilon}) = \frac{v}{v-2}V$. So the actual variance-covariance matrix of a MVT distribution is $v/(v-2)$ times V .

To simplify the calculations, we will initially assume that V is proportional to the square root of the horizon, despite the fact that we know that V is described by a bekk(1, 1) model. The reason we make this assumption is that in period 0, the Value at Risk requires knowledge of the distribution of $\boldsymbol{\varepsilon}_{20}$. We know that this is distributed $MVT(v, V_{20})$, however V_{20} depends on the values of $\boldsymbol{\varepsilon}_0, \dots, \boldsymbol{\varepsilon}_{19}, V_1, \dots, V_{19}$ - the resulting equation is messy and although it can be predicted by iteration, it is easier to simplify the model and just treat the variance as constant. This is the approach used typically in commercial VaR systems. We can use the current value of V_t to parameterize the VaR calculation¹⁵:

$$(\tilde{\boldsymbol{\varepsilon}}_{20} | t = \tau) \sim MVT(v, \tilde{V}) \quad \tilde{V} = V_\tau \sqrt{20 - \tau}$$

For consistency we assume that the horizon for this Value at Risk calculation is the same as for the overall calculations, i.e. VaR at time $t = 120$ (20 periods in the future).

If the returns were normally distributed, then the Value at Risk of a particular portfolio would be calculated as follows¹⁶:

$$VaR = \Phi^{-1}(.95)\boldsymbol{\sigma}_p \sqrt{\tau} = 1.65\sqrt{X^T \tilde{V} X}$$

$\Phi^{-1}(.95)$ is the 95th percentile of the normal distribution and X is a vector of dollar holdings in the pound and yen respectively. Given that the returns are distributed $MVT(v, \tilde{V})$, we first notice

¹⁵ We use $\tilde{\boldsymbol{\varepsilon}}$ to denote returns under the simplifying assumption that $V_t = V_\tau$ for $t = \tau, \dots, 20$.

¹⁶ This equation is taken from Jorion (1997) and www.riskmetrics.com

that¹⁷ $a_1 \tilde{\epsilon}_1 + a_2 \tilde{\epsilon}_2 \sim t(v, a^\top \tilde{V} a)$, where $t(v, a^\top \tilde{V} a)$ is a t distribution with zero mean and variance given by $\frac{v}{v-2} a^\top \tilde{V} a$. The corresponding VaR restriction for t-distributed returns is:

$$VaR = t^{-1}(.95) \sigma_p \sqrt{\tau} = 1.995 \sqrt{X^\top \tilde{V} X}$$

The restriction placed by Value at Risk is implemented as follows. A limit on the acceptable losses in the final period is set. At each step, the losses incurred up to this point are calculated and subtracted from the limit (so as to make the restriction tighter on the remaining value). If the remaining loss allowable is VaR_{remain} , then the restriction becomes¹⁸:

$$1.995 \sqrt{X^\top \hat{V} X} \leq VaR_{remain} \quad \text{or} \quad X^\top \hat{V} X \leq \frac{VaR_{remain}^2}{1.995^2}.$$

The budget constraint requires that $x_1 p_{1t} + x_2 p_{2t} = Z_t$. So $x_1 = \theta Z_t$ and $x_2 = (1 - \theta) Z_t$. To find the effect of the restriction on the behavior of the investor we substitute θ into the constraint and solve:

$$\left[\theta^2 \hat{v}_{11} + 2\theta(1 - \theta) \hat{v}_{12} + (1 - \theta)^2 \hat{v}_{22} \right] \leq \frac{VaR_{remain}^2}{1.995^2 Z_t^2}$$

$$\Rightarrow \theta^2 (\hat{v}_{11} - 2\hat{v}_{12} + \hat{v}_{22}) + \theta(2\hat{v}_{12} - 2\hat{v}_{22}) + \left(\hat{v}_{22} - \frac{VaR_{remain}^2}{1.995^2 Z_t^2} \right) \leq 0$$

$$\text{Hence the roots are } \theta_L, \theta_H = \frac{(\hat{v}_{22} - \hat{v}_{12}) \pm \sqrt{(\hat{v}_{12} - \hat{v}_{22})^2 - (\hat{v}_{11} - 2\hat{v}_{12} + \hat{v}_{22}) \left(\hat{v}_{22} - \frac{VaR_{remain}^2}{1.995^2 Z_t^2} \right)}}{(\hat{v}_{11} - 2\hat{v}_{12} + \hat{v}_{22})}$$

The investor is restricted to investing between θ_L and θ_H .

¹⁷ in the same way that $a_1 x_1 + a_2 x_2 \sim N(a^\top \mu, a^\top \Sigma a)$ if $x \sim N_2(\mu, \Sigma)$

¹⁸ Provided that $VaR_{remain} \geq 0$. If $VaR_{remain} < 0$ then we assume that no roots exist.

The behavior is the same as in risk limit A, except that at each stage θ is restricted to lying between $\theta_0 = \theta_L$ and $\theta_1 = \theta_H$, which vary with the portfolio, rather than between fixed θ_0 and θ_1 .

Risk Limit C – Projected GARCH Value at Risk Constraint

The expected variance matrix of the 20-period-ahead returns can be calculated by the following recursion:

$$\mathbf{E}(V_{t+1} | V_t) = AA^\top + \frac{v}{v-2} A_1 V_t A_1^\top + B_1 V_t B_1^\top$$

This variance-covariance matrix takes slightly longer to calculate than the root-t approach, however the effect on the performance of the simulation is minimal. We will compare the results of this method with the simpler approach to see how different the results are.

Handling of excessive positions

The only remaining question, which needs to be addressed, is what to do if the Value at Risk constraint does not have any roots, i.e. the constraint cannot be satisfied for any holdings.

Initially, we will assume that the portfolio is reverted to the minimum variance portfolio, which is given by setting

$$\theta_{\min} = \frac{\hat{v}_{12} - \hat{v}_{22}}{\hat{v}_{11} - 2\hat{v}_{12} + \hat{v}_{22}}$$

This would be the best solution if leaving the markets is not allowed.

We will also consider the case where the portfolio is frozen until the restriction can be satisfied. This typically leaves the portfolio in an heavily exposed position and should increase the Value at Risk dramatically.

Typically a business would close the positions completely upon discovery of an extremely large risk exposure. This would involve moving from holdings in pounds and yen to holding dollars.

5. Simulated Results

Comparison of Bootstrapped Data with the GARCH Model

As an initial test, we compare the distributional properties for investor 0 under both the GARCH model and the model where returns are bootstrapped from the initial returns. Histograms of final portfolio values are shown in Figure 10.

From the histograms it is clear that the GARCH process produces similar results to results using bootstrapped returns. The important difference, although slight, is in the tails of the distribution. The conditional heteroskedasticity in the GARCH process produces the fat tails that must be captured to produce accurate estimates of risk in exchange rates. We will hence use the GARCH model as parameterized above to compare the performance of the remaining investors.

Value at Risk of Investors

1000 simulated paths, each containing 20 periods were generated using the GARCH model described previously. These sample paths were each passed in turn through the investor model simulations to calculate the changes in holdings. Initially, we use the constraints $\theta \in (-1, 2)$ for risk measure A and $\text{VaR} < \$25$ for risk measures B and C. Value at Risk was calculated by taking the 50th (ordered) result, then subtracting it from the initial portfolio value (\$1000). Estimated Value at Risk for each of the investors is given Table 3. The largest simulated loss (i.e. the 0.1 percentile) as well as the expected future value of the investor's portfolio is also given. The number of times in which the VaR limit is given in column 3.

Table 3 shows that the Value at Risk of a fixed portfolio is \$60.62. When the portfolio is allowed to vary, the Value at Risk increases, naturally. First we notice that Risk Measure B produces results closest to the fixed portfolio and risk measure C produces results with the highest Value at Risk for both investors. It appears that on the whole, investor 1 takes on more risk than investor 2.

Comparisons between the severity of risk measures A, B and C are difficult, as there is no one figure that describes how restrictive a limit is. A few measures of the holdings of the investors under the various limits are given in Table 4. Column 3 shows that the average magnitude of the exposure is largest under risk limit C and with investor 1. The Value at Risk constraint of \$25

was intentionally chosen so that for investor 1, under risk limit B (root-t value at risk) the average exposure was similar to that of the fixed limit A (around 200% of net worth held in one currency).

The maximum and minimum positions taken are significantly larger for the Value at Risk constraint; however this only occurs once the investor has amassed substantial gains. It is apparent from column 3 in Table 4 and column 2 in Table 3 that the two methods of estimating the variance-covariance matrix are not equivalent. Using the GARCH estimation (Risk limit C) gives a significantly more lenient risk measure. This is reflected in the Value at Risk of the various portfolios: for risk measure C the Value at Risk is roughly 40% larger.

Column 3 of Table 3 shows that the Value at Risk constraint is exceeded more often than not. This is not surprising as the value at risk of a fixed portfolio is \$60 and we have required that the manager maintain a value at risk below \$25. We now investigate the performance of less restrictive risk limits.

The Effect of Relaxing the Risk Limits

In Table 6 we investigate the effect of relaxing the constraints. To make comparisons between risk measures, we increase limits A and B so that the average exposure for investor 1 is approximately equal¹⁹.

First, comparing the value at risk and expected returns for risk measures A and B, we notice that the performance of the Value at Risk measure is superior in both. Hence for the same average restrictions on exposure, Value at Risk's flexibility allows the portfolio manager to take risks with the accumulated gains.

The Value at Risk incurred by investors under risk measure B is significantly larger the prescribed value at risk limit (on average to the tune of \$40). Proportionally, however, the two measures do get closer as the limit is relaxed.

¹⁹ For now we will ignore risk limit C. Tests (not reported) show that because both Value at Risk based limits use the same philosophy the results obtained are similar and differ consistently by a scale factor of around 40%.

Hence, it is important to understand that in a high risk dynamic trading setting, the Value at Risk limit set, is not necessarily the value at risk obtained. This is caused in part by the constraint, which requires only that the value at risk of a *fixed* portfolio with a *fixed 20-day horizon* remain below a certain level.

Each new portfolio is chosen so that if the portfolio were to remain fixed, the risk limit would be exceeded 1 out of 20 times at the conclusion of 20 days. Such a risk limit does not require that the risk remain below the limit at *all* points during the 20 days.

Importance of Closing Positions

As recent collapses in the financial world have shown, it is important to close positions as soon as possible on the discovery of excessive risk taking²⁰. This can be demonstrated by comparing the results of Table 3 (where the positions are immediately returned to a minimum-variance portfolio) with those of Table 5 (where positions are left open).

The tables show that the Value at Risk is larger when the portfolio is left exposed and little additional return is gained from such an oversight. Many of the simulated runs exceed the VaR limit early on in the path. This limits the opportunities for the investors to exploit the volatility of the market and create excessive positions – those that could cause the most damage.

To understand the relative severity of the two different restrictions (fixed limits and the VaR based limits), a statistical summary of the average positions held by all four investors is given in Table 4.

Column 3 shows the mean absolute value of the holdings in pounds of each investor. By design, this is 2 for investors 1 and 3. For investors 2 and 4 the value is between 3 and 3.5 times larger. Roughly speaking this means that on average, investors 2 and 4 are allowed to take 3 times as large a holding in any one currency. This is a natural reason as to why there is such a difference between the VaR for investors 1 and 2 in Table 5. To level the field, consider tightening the VaR restriction to 11.29 for investors 2 and 4.²¹ Results for investors 2 and 4 are given in Table 6.

²⁰ Rawnsley (1995) is quick to highlight in her description of the aftermath of Barings' initial collapse, how some of the additional losses could have been avoided had positions been closed upon discovery.

²¹ I.e. final value must not fall below \$270 with 95% certainty at each step.

Distortion of the Tail Probabilities

Figure 11 shows the distribution of profits/losses for the 3 different portfolio managers (0, 1 and 2) and the two different risk limits (with θ and $\text{VaR}_{\text{limit}}$ chosen so that the ex-post Value at Risk is \$100 in all cases).

Investor 1B (Expected utility maximizing investor with VaR constraint) has a similar distribution to investor 1A, except for the extreme positive and negative tails, which are heavier for investor 1B. This distortion is caused by a similar effect to that outlined by Basak and Shapiro (1998). The investor is able to align his portfolio so that the 5th percentile falls at -\$100 but maximizes his return by infrequently taking large risks. This does not affect the main body of the distribution, but has an undesirable effect on the tails.

This is the “hidden” risk of using Value at Risk as a risk measure – whilst it captures the overall risk of the portfolio, it is possible for an investor to construct his portfolio in such a way that the risk is increased without affecting the constraint ex-ante.

6. Conclusion

This paper has shown the importance of taking into account dynamic aspects of portfolio holdings when calculating Value at Risk. Although it is unrealistic to assume that a risk manager can describe the behavior of his traders by as simple a model as given above, using such models alongside more traditional calculations can serve as an extra measure by which risk can be gauged. The use of a model that captures the inter-temporal characteristics of both the underlying financial instruments as well as the behavioral dynamics is vital if the results are to be used as an additional reporting tool.

Simple Value at Risk based measures can be used successfully to reduce risk whilst still allowing the portfolio manager to hold on average the same extremes in holdings.

The method of simulating conditional expectations by taking ensemble averages of sample paths proved to be a valid trade-off between computational complexity and closeness to the theoretical model. Reuse of the random innovations increases the bias of the results, however the

computational advantage obtained as a result far outweighs the small increase in bias produced. The algorithm worked fast enough on the specified models that multi-period dynamic programming models of investor behavior could be considered in future research without overstepping sensible computational limits on processing time²²

²² The simulations took about 2 minutes to run on a SUN Ultra 10 440. Increasing the dynamic programming problem to a 2 period look-ahead rule could increase this time to several hours.

7. References


- Basak, S. and A. Shapiro (1998) "Value-at-Risk Based Risk Management Optimal Policies and Asset Prices," Wharton working paper.
- Beder, T. (1995) "VAR: Seductive but Dangerous," Financial Analysts Journal, Capital Market Risk Advisors, Inc.
- Berkowitz, J., (1996) "Recent Developments in Bootstrapping Time Series," Board of Governors of the Federal Reserve System.
- Bollerslev, T. (1986) "Generalized Autoregressive Conditional Heteroskedasticity," Journal of Econometrics.
- Campbell, J., A. Lo, and A. MacKinlay (1997) "The Econometrics of Financial Markets," Princeton University Press.
- Cvitanic, J. and Karatzas, I. (1998) "On Dynamic Measures of Risk," Working Paper, Columbia University.
- Davison, A. and D. Hinkley (1997) "Bootstrap Methods and their Application," Cambridge University Press.
- Devroye, L. (1986) "Non-uniform Random Variate Generation," Springer-Verlag.
- Duffie, D. and J. Pan (1997) "An Overview of value-at-risk," Journal of Derivatives, spring.
- Engle, R. and S. Manganelli (1999) "CAViaR: Conditional Autoregressive Value at Risk by Regression Quantiles," NBER Working Paper.
- Hendrick, D. (1996) "Evaluation of Value-at-Risk Models Using Historical Data," FRBNY Economic Policy Review.
- Jorion, P. (1997) "Value-at-Risk: The New Benchmark for Controlling Market Risk," McGraw-Hill.
- Krause, A. and M. Olson (1997) "The Basics of S and S-PLUS," Springer.
- Li, D. (1999) "Value at Risk Based on the Volatility, Skewness and Kurtosis," RiskMetrics Group Working Paper.
- Martin, R., H. Gao, Y. Zhan and Z. Ding (1996) "S+GARCH User's Manual" MathSoft, Inc.

Rawnsley, J. (1995) "Total Risk: Nick Leeson and the Fall of Barings Bank," Harper Business.

RiskMetrics (1995)


Wang, T. (1998) "A Characterization of Dynamic Risk Measures," working paper, University of British Columbia.

Figure 1




The 5% quantiles for both distributions are identical, however the dotted curve places more weight on an extreme loss.

Figure 2



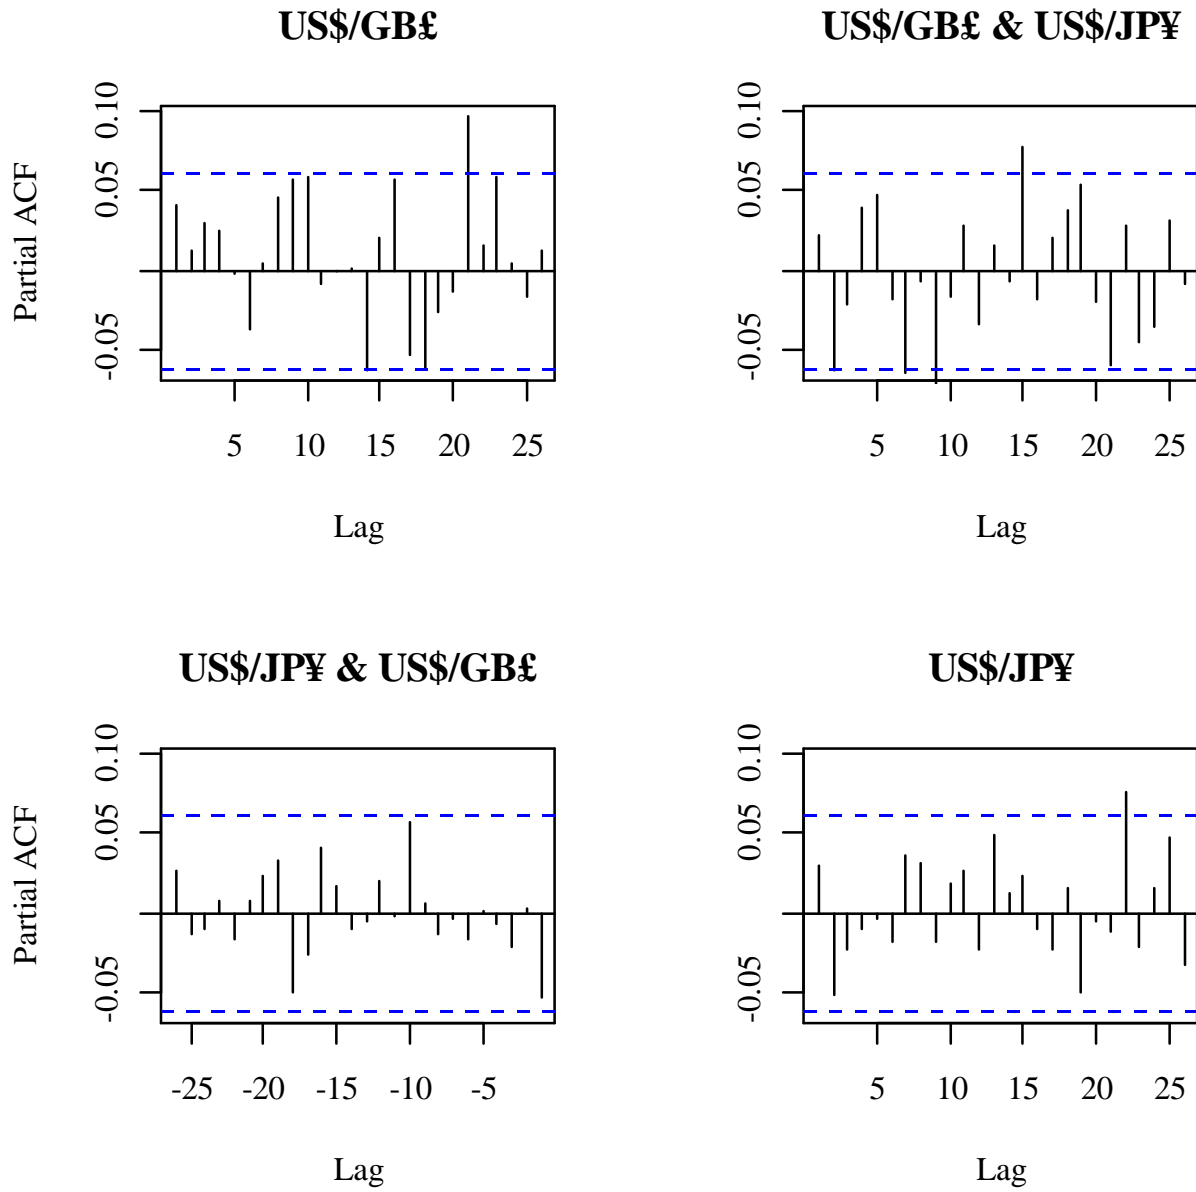
The raw time-series used as input data for the simulation

Figure 3



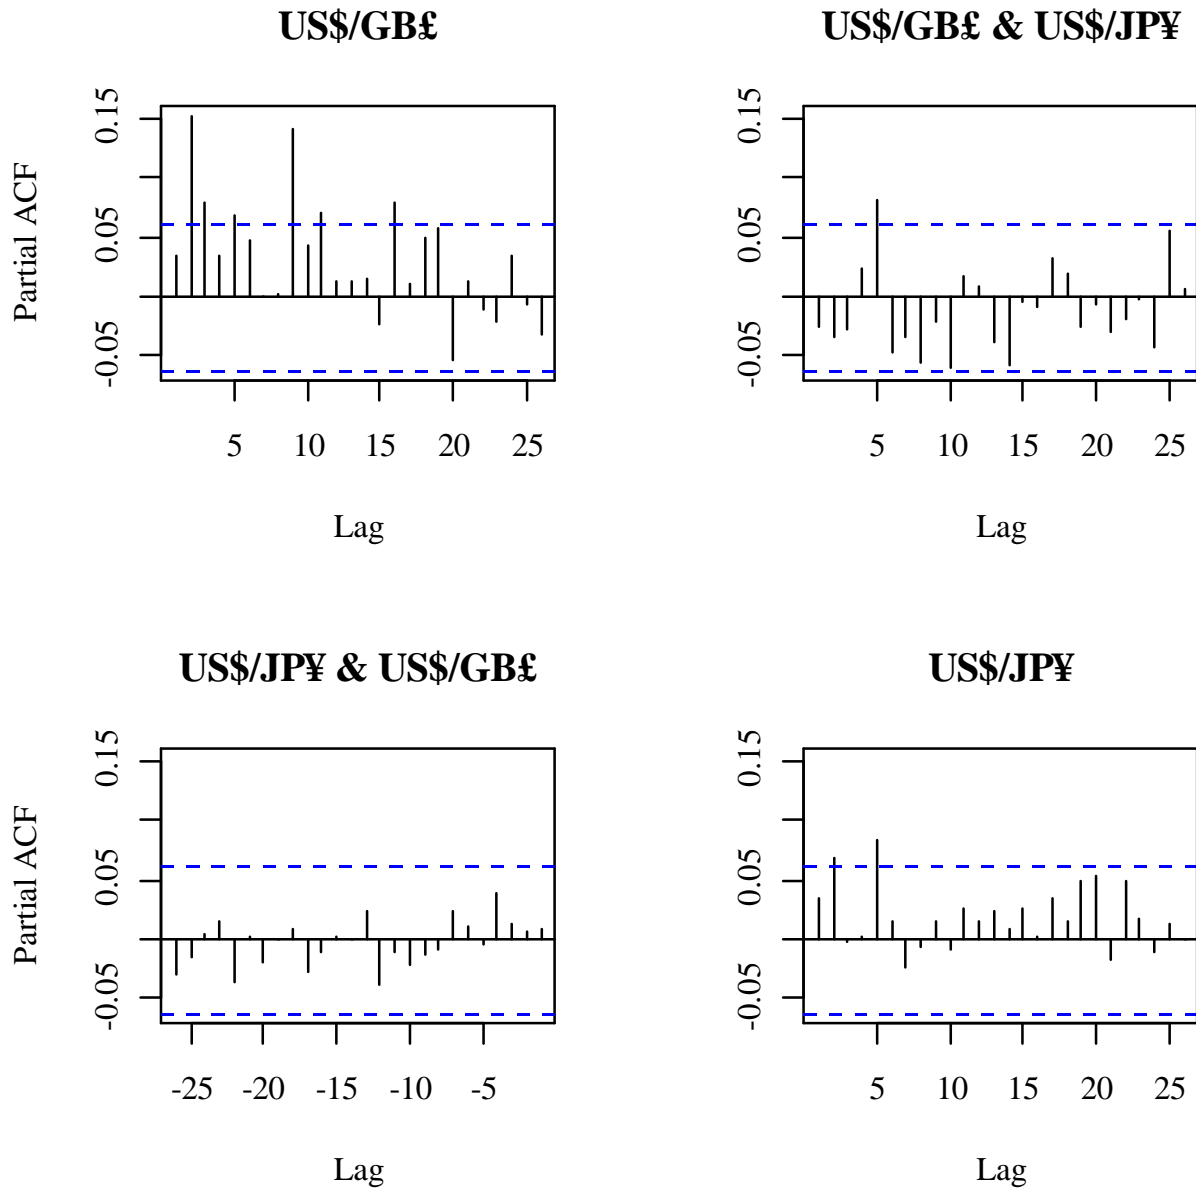
The time-series of returns, $y_t = \ln(x_t) - \ln(x_{t-1})$

Figure 4





Partial Autocorrelations and Cross-correlations of the Returns

Figure 5





Partial Autocorrelations and Cross-correlations of the *squared* Returns

Figure 6



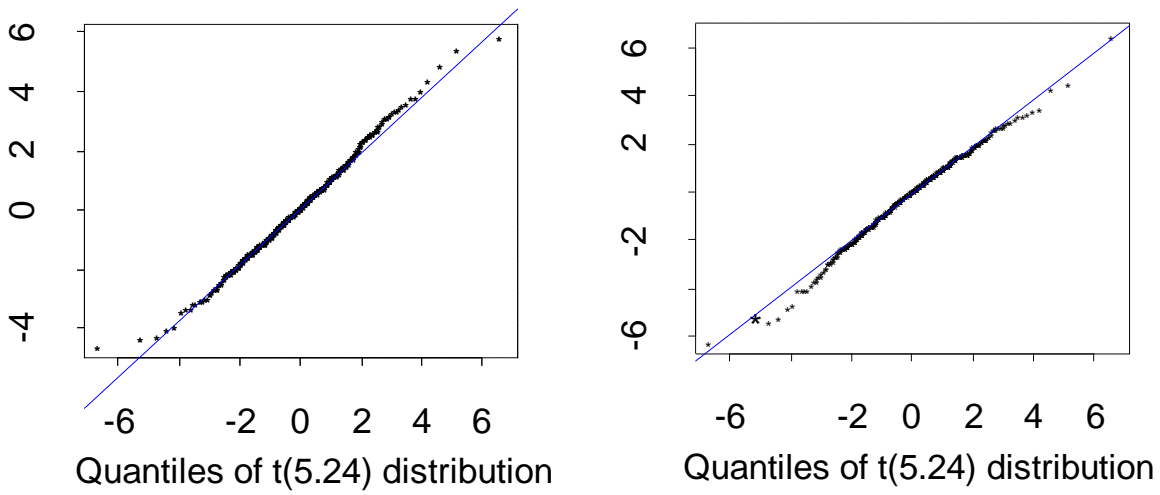
Distributions of the returns of the exchange rates

Figure 7






Q-Q Plots of the returns against a normal distribution

Figure 8




Q-Q plots of residuals against a t distribution with 5.23 degrees of freedom.



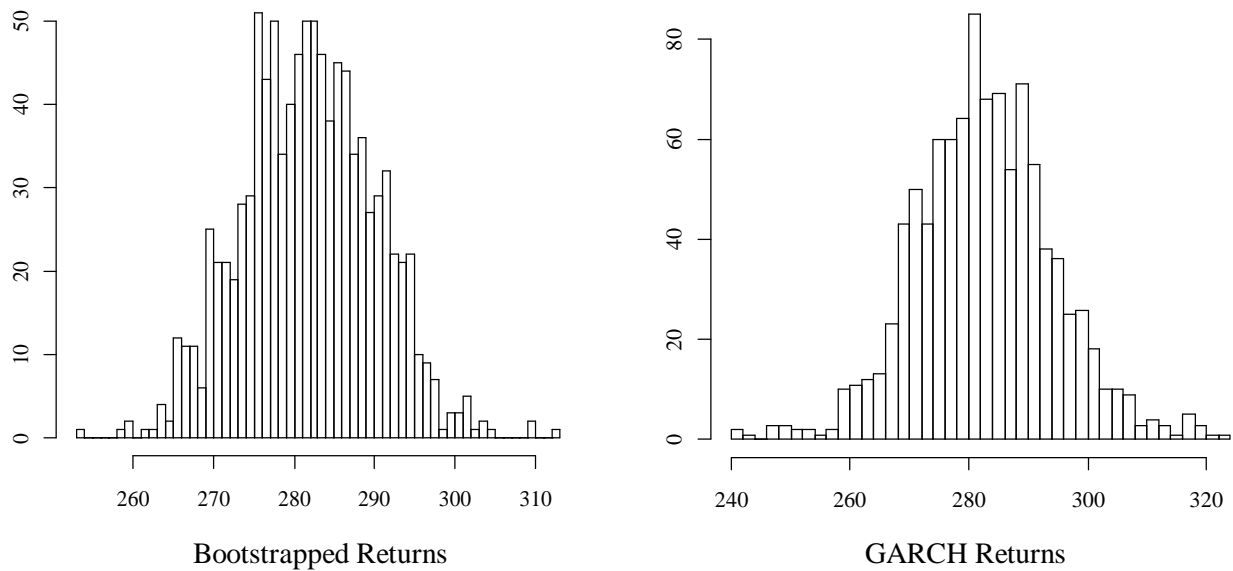
Autocorrelation plots of the cross products of the standardized model residuals.

Figure 9




One Monte Carlo Run: Periods 1 through 100 are used as a common path to develop investor behavior that is comparable across different realizations in periods 101 to 120. Period 100 is taken to be the present rate and is fixed for all Monte Carlo Runs.

Figure 10



Distribution of final portfolio value for investor 0 under bootstrap and GARCH model of returns

Figure 11



Distribution of final portfolio value for investors 0, 1A (EU/Fixed), 1B (EU/VaR), and 2B (Filter/VaR) under GARCH model of returns

Table 1

Mean of US\$/GB£	0.5916	Mean of US\$/JP¥	129.125
Mean of returns on US\$/GB£	7.777×10^{-5}	Mean of returns on US\$/JP¥	-2.509×10^{-4}
Variance of return on US\$/GB£	6.491×10^{-5}	Variance of return on US\$/JP¥	4.598×10^{-5}
Skewness of returns on US\$/GB£	0.2221	Skewness of returns on US\$/JP¥	-0.1414
Kurtosis of returns on US\$/GB£	2.321	Kurtosis of returns on US\$/JP¥	2.968
Covariance of returns	2.521×10^{-5}	Correlation of returns	0.4615

Statistical Properties of the time series $y_t = \ln(x_t) - \ln(x_{t-1})$

Table 2

Parameter	Est. value	s.e.	t value	Parameter	Est. value	s.e.	t value
c_1	-0.000369	0.000213	-1.73211	A_1^{12}	-0.045055	0.018348	-2.45550
c_2	-0.000137	0.000183	-0.74740	A_1^{22}	0.177092	0.025361	6.98295
A^{11}	0.000946	0.000244	3.88133	B_1^{11}	0.960745	0.010249	93.73683
A^{21}	0.000258	0.000322	0.80139	B_1^{21}	0.004733	0.012163	0.38914
A^{22}	0.000448	0.000226	1.97873	B_1^{12}	0.011480	0.008284	1.38574
A_1^{11}	0.185518	0.024556	7.55482	B_1^{22}	0.968794	0.008683	111.5754
A_1^{21}	-0.001234	0.031080	-0.03969	ν	5.236174	0.7083102	

AIC(14) = -14510.44

BIC(14) = -14441.73

S+GARCH estimation of the exchange rate model and estimated parameters.

Table 3

Investor	VaR (\$)	Fraction of runs in which VaR limit is exceeded at least once	Largest simulated loss (\$)	Mean return (\$)	Variance of final value
0	60.62	-	147.67	4.87	1769.4
1A	83.53	-	177.78	1.92	2893.8
1B	68.74	96.5%	337.95	14.13	77998.5 ^{23a}
1C	103.00	91.7%	1022.71	112.3	11327896 ^{23b}
2A	85.81	-	224.45	7.82	3825.6
2B	61.64	95.7%	153.75	3.27	3392.0
2C	98.20	90.1%	352.92	2.80	35482.3

Estimated Value at Risk for Investors and other key statistics

(Fixed limits (A) are $\theta \in [-1, 2]$.)

VaR limits (B and C) requires VaR < \$25)²⁴

Value at Risk is estimated using a root t approach

Table 4

Investor	Mean holding of pounds ($\bar{\theta}$)	Largest percentage held in £ or ¥	Minimum £ holding (θ_{\min})	Maximum £ holding (θ_{\max})
1A	0.05	200%	-1	2
1B	0.79	206%	-22.93	73.00
1C	0.95	339%	-87.34	95.93
2A	1.42	200%	-1	2
2B	0.80	154%	-19.05	21.92
2C	2.16	317%	-55.18	78.93

Summary statistics of positions taken in Pounds/Yen.

$\theta = 1$ represents 100% of wealth invested in Pounds.

Values less than 0 indicate selling short in Pounds.

Values greater than 1 indicate selling short in the Yen.

²³ The figures reflect a large skew caused by the large returns of the 10 best realizations.

²⁴ I.e. final value must not fall below \$250 with 95% certainty.

Table 5

Investor	VaR (\$)	Fraction of runs in which VaR limit is exceeded at least once	Largest simulated loss (\$)	Mean return (\$)	Variance of final value
1B	83.29	96.5%	721.46	16.91	89840.6
2B	71.01	95.7%	249.95	4.81	3888.2

Value at risk for investors 1 and 2 when overexposed portfolios are frozen (left exposed).

Table 6

Investor	Constraints on θ or VaR	VaR (\$)	Fraction of runs in which VaR limit is exceeded at least once	Largest simulated loss (\$)	Mean return (\$)	Average Exposure
1A	$\theta \in (-2\frac{1}{2}, 3\frac{1}{2})$	135.60	-	279.48	-1.11	3.5
	$\theta \in (-5, 6)$	187.37	-	377.37	-4.24	5
	$\theta \in (-8\frac{1}{4}, 9\frac{1}{4})$	340.01	-	627.61	-13.61	9.25
1B	VaR < 50	91.10	73.7%	476.65	24.20	3.54
	VaR < 75	116.31	60.3%	765.14	33.00	5.07
	VaR < 150	201.75	43.0%	1778.37	64.06	9.25
2A	$\theta \in (-2\frac{1}{2}, 3\frac{1}{2})$	138.89	-	334.86	10.93	3.5
	$\theta \in (-5, 6)$	198.82	-	433.33	14.17	5
	$\theta \in (-8\frac{1}{4}, 9\frac{1}{4})$	362.43	-	654.56	24.07	9.25
2B	VaR < 50	78.66	69.4%	159.08	7.14	2.73
	VaR < 75	105.56	54.8%	202.46	11.53	4.05
	VaR < 150	178.27	34.3%	630.63	23.75	7.72

Value at Risk and average positions for investors 1 and 2 under various levels on risk limits