

Risk management at the trade level for foreign exchange strategies

Issam S. STRUB* and Christopher A. UDY†

The Cambridge Strategy (Asset Management) Ltd
7th Floor, Berger House, 36–38 Berkeley Square, London W1J 5AE, United Kingdom

March 12, 2010

1 Introduction

Currency markets tend to be highly speculative and volatile in nature; sudden changes in currency pair values can occur quickly due to a number of fundamental and technical factors, therefore efficient risk management is crucial for a successful trading strategy. The main idea consists of controlling risk by capping losses for each trade at a certain percentage of assets under management. This is done by allocating a certain amount to each trade called the *Risk Adjusted Trade Size* or RATS. Traditional risk management techniques such as Value at Risk and Monte Carlo simulations rely on the assumption that portfolio returns are normally distributed, which grossly underestimates tail risk, in particular for high beta FX strategies which typically generate so called fat tails. To address this issue, the Cambridge Strategy developed a number of proprietary algorithms which compute the risk adjusted trade size using advanced mathematical and statistical techniques which take into account the entire return distribution. These various formulæ will be presented and compared using daily returns data from the Asian strategy.

2 Sharpe ratio based RATS

The Sharpe ratio based RATS multiplies a Model RATS by a coefficient representing the past performance of the trading strategy, as modelled by its Sharpe ratio and volatility σ . We begin by setting a given level of Maximal Acceptable Loss (or *ruin*). Then, the Sharpe ratio over a certain period of time (for example the previous 12 months) is computed as well as the strategy's annualised volatility. Using this data, Monte Carlo simulations are run and yield a matrix of probabilities of returns (see WILLIAMS [2005]). A maximum probability (MaxP) of ruin is chosen and the loss corresponding to this probability is computed using the aforementioned matrix; more precisely, we look for a Loss Percentage such that $P(\text{Returns} \leq \text{Loss Percentage}) = \text{MaxP}$. Now that this value has been computed, it is divided by the volatility σ , yielding the *Unacceptable Loss Measure* (ULM):

$$\text{ULM} = \frac{\text{Loss Percentage such that } \{P(\text{Returns} \leq \text{Loss Percentage}) = \text{MaxP}\}}{\text{Annualised Daily Volatility}} \quad (1)$$

A trade size which corresponds to a probability MaxP of Maximal Acceptable Loss is now selected; this is given by the quotient of the maximum Percentage Acceptable Loss over the Loss Percentage reached with probability MaxP

*Research Scientist, issam.strub@thecambridgestrategy.com, 7th Floor, Berger House, 36–38 Berkeley Square, London W1J 5AE, United Kingdom.

†Director of Research, chris.udy@thecambridgestrategy.com, Level 40, Chifley Tower, Chifley Square, Sydney, NSW 2000, Australia.

then multiplied by the Model RATS:

$$\text{RATS} = \frac{\text{Percentage Acceptable Loss} \times \text{Model RATS}}{\text{Annualised Daily Volatility} \times \text{Unacceptable Loss Measure}} \quad (2)$$

In the practical case mentioned in THOMPSON [2007], the maximum Percentage Acceptable Loss is 10%, the chosen probability of reaching this ruin level is 5% and the volatility is $\sigma = 12.8\%$. Using Monte Carlo simulations it is found that for the current Sharpe ratio of 2.1, a probability of 5% corresponds to a loss of $7.4\% = 0.58 \times \sigma$. In other words

$$P(\text{Returns} \leq -7.4\%) = 5\%$$

This leads to the RATS:

$$\text{RATS} = \frac{10 \times \text{Model RATS}}{12.8 \times 0.58} \quad (3)$$

One possible extension would be to use the CS ratio (CASCON and SHADWICK [2007]) rather than the Sharpe ratio to compute the RATS as this would better estimate the downside risk. This is examined below.

3 CS Ratio based RATS

Indeed, the Sharpe ratio based RATS previously described relies on the Sharpe ratio and volatility of a given strategy to assess its probability of ruin using Monte Carlo simulations. However, return distributions with the same Sharpe ratio and standard deviation can exhibit widely differing tails and thus downside risk may not be properly estimated. Secondly, Monte Carlo simulations assume a known return distribution, typically normal, which significantly underestimates tail risk. To address this, the CS RATS was developed using the CS ratio introduced in CASCON and SHADWICK [2007] and the *filtered historical simulation* (FHS) method presented in BARONE-ADESI et al. [1999]. The CS ratio is defined by:

$$\text{CS ratio} = \frac{\mu - r_f}{\omega} \quad (4)$$

where $\omega = \frac{1}{2}\mathbb{E}(|x - \mu|)$ is the standard dispersion. Like the Sharpe ratio, the CS ratio has the property of being scale invariant and the added advantage of taking into account the full distribution including tails without making any assumption of normality.

The filtered historical simulation (FHS) method is an attempt to combine the advantages of a historical sampling technique and a parametric estimation model; the latter attempts to capture conditional heteroskedasticity but assumes a normal distribution while the former does not assume a specific distribution but does not capture conditional heteroskedasticity. The FHS method relies on a model based approach for the volatility, typically using a GARCH type model, while remaining model free in terms of the distribution. In particular, this method has the notable advantage of being able to simulate extreme losses even if they are not present in the sample historical returns used for the simulation, thus taking tail risk into consideration more accurately.

3.1 Filtered Historical Simulation

In the following we will apply this method to compute a CS Ratio based RATS. We use daily returns from a given strategy and denote R_t as the daily log return at time t and σ_t as the standard deviation of daily returns. As mentioned above, the historical simulation method assumes that daily returns are independent and identically distributed (i.i.d.) through time; however, significant autocorrelation is usually found for the daily squared returns. To produce a sequence of i.i.d. observations, we fit an AR(1) first order autoregressive model to the daily returns:

$$R_{t+1} = c + aR_t + \varepsilon_t \text{ where } \varepsilon_t = \sigma_t z_t \quad (5)$$

In our case we will choose a Student t distribution for the standardised returns $\{z_t\}$ to account for increased tail risk as the t distribution has fatter tails than the corresponding normal distribution.

To model the variance of returns we use an exponential GARCH (EGARCH) model developed in NELSON [1991], NELSON and CAO [1992] to capture the asymmetry in volatility induced by large positive and negative returns. Indeed, volatility usually increases more after a large negative return than after a large positive return due to the leverage effect (see BLACK [1976] for more on the subject). The following model will be used:

$$\ln \sigma_{t+1}^2 = \omega + \alpha(\phi R_t + \gamma(|R_t| - \mathbb{E}|R_t|)) + \beta \ln \sigma_t^2 \quad (6)$$

After fitting an AR(1)/EGARCH(1,1) model to the daily returns, the standardised residuals (which are now i.i.d.) are used in a historical sampling technique. The i.i.d. property is important for bootstrapping, as it allows the sampling procedure to safely avoid the pitfalls of sampling from a population in which successive observations are serially dependent. We now simulate a number N_1 of independent random trials over a given time horizon of N_2 days; unlike Monte Carlo simulations we do not make a specific distributional assumption regarding the standardised returns $\{z_t\}$ and instead use the past returns data. Eventually, we end up with N_1 daily return series each covering N_2 trading days.

We begin by computing the CS Ratio of the current strategy on a rolling one year window and select the simulated return series which have a CS Ratio within a certain specified range of the aforementioned CS Ratio, and we then compute the loss corresponding to a 5% probability for each simulated return series previously selected. We take the largest of these losses and use it to compute the RATS corresponding to a 10% maximum acceptable loss for the strategy as was done for the Sharpe Ratio based RATS. This in turn yields the CS Ratio based RATS (CS RATS) which is adjusted weekly.

4 Risk Adjusted Trade Size with Expected Shortfall and Extreme Value Theory

The latest risk management tool developed by the Cambridge Strategy called *Extreme RATS* (ERATS) uses the most advanced quantitative measures of extreme risk, Extreme Value Theory and Expected Shortfall.

4.1 Expected Shortfall

The RATS computation from the previous sections is based on a maximum acceptable loss at a certain confidence level, or, equivalently, a 95 % Value at Risk (VaR) of 10 %. However, Value at Risk is concerned only with the number of losses that exceed the VaR and not the size of these losses. Many investment professionals would like to consider a risk measure that accounts for both the magnitude of large losses as well as their probability of occurring. The most complete measure of large losses is the shape of the tail of the return distribution beyond the VaR. This led to the introduction of *Expected Shortfall* (ES) (see CHRISTOFFERSEN [2003], MCNEIL et al. [2005]) (also referred to as *Conditional Value at Risk* (CVaR) or TailVaR). The expected shortfall for a daily return distribution F at a given confidence level α is given by:

$$ES_\alpha = -\mathbb{E}\{X|X \leq -\text{VaR}_\alpha\} \quad (7)$$

where the VaR is:

$$\text{VaR}_\alpha = -F^{-1}(1 - \alpha) \quad (8)$$

While expected shortfall presents a number of additional interesting properties, such as being a coherent risk measure and a convex function of the portfolio weights (unlike VaR) which makes it suitable for portfolio optimisation, its computation requires an explicit expression of the return distribution function F which is usually unknown in practice. To address this problem, one possible method consists in fitting a known distribution to the left tail of historical daily returns and then making use of this distribution to compute the expected shortfall. One such distribution is the *Generalised Pareto Distribution* (GPD) originating from *Extreme Value Theory* (EVT), a branch of statistics dedicated to modelling extreme events.




Figure 1: Left: Comparative plot of the left tail of standardised residuals obtained after filtering daily returns, GPD distribution fitted to the standardised residuals through MLE, and normal standard distribution. The empirical data present a noticeably fatter tail than the normal distribution while the GPD provides a good approximation. Right: Comparison of daily cumulative returns obtained using either Sharpe RATS, CS RATS or ERATS computed weekly between January 2008 and April 2009.

4.2 Extreme Value Theory

The central result in extreme value theory states that the extreme tail of a wide range of distributions can approximately be described by the GPD (BALKEMA and DE HAAN [1974], PICKANDS [1975]):

$$G_{\zeta, \beta}(y) = \begin{cases} 1 - (1 + \frac{\zeta y}{\beta})^{-\frac{1}{\zeta}}, & \text{for } \zeta \neq 0; \\ 1 - \exp^{-\frac{y}{\beta}}, & \text{for } \zeta = 0. \end{cases} \quad (9)$$

The shape and scale parameters ζ and β can be estimated using *Maximum Likelihood Estimation* (MLE) by fitting a GPD distribution to the tail of the return distribution beyond a given threshold u . Once this is done, the expected shortfall can be computed:

$$ES_{\alpha} = \frac{VaR_{\alpha} + \beta - \zeta u}{1 - \zeta} \quad (10)$$

where the VaR for a GPD can be estimated by:

$$VaR_{\alpha} = u + \frac{\beta}{\zeta} \left(\left(\frac{\alpha N}{N_u} \right)^{-\zeta} - 1 \right) \quad (11)$$

with N the total number of observations and N_u the number of observations exceeding the threshold u .

4.3 Practical Implementation

The preceding result requires that observations be i.i.d., therefore, we will start by filtering daily returns using the AR(1)/EGARCH(1,1) model with a t-distribution as described earlier and apply EVT to the standardised residuals (see MCNEIL and FREY [2000], NYSTRÖM and SKOGLUND [2002]), with a GPD being fitted to the tails through MLE. An example of this technique is presented in Figure 1; the normal distribution is clearly inadequate to model the left tail and the GPD provides a much more suitable approximation. Once this is done, we obtain the shape and scale parameters corresponding to the GPD fitted to the left tail and replace these values in Equation (10) to compute the expected shortfall at the level of confidence desired (95 % in this case). This value is then used to obtain the RATS level using the following formula:

$$ERATS = \frac{\text{Model RATS} \times \text{Maximum Acceptable ES}}{\text{Computed ES}} \quad (12)$$

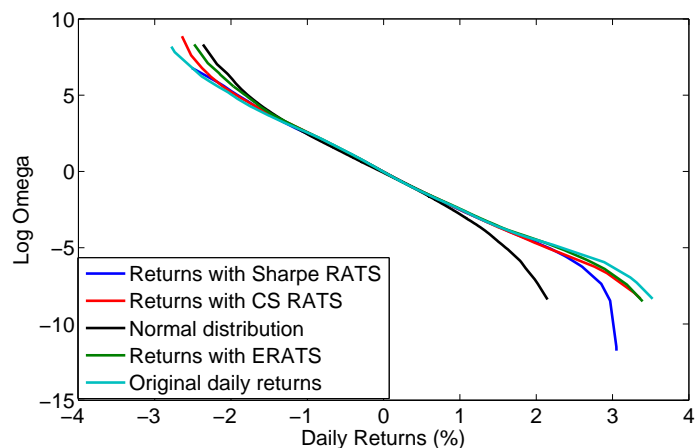


Figure 2: Logarithm of the Omega function for each strategy and for a normal standard distribution as a function of standardised daily returns. The Omega curve corresponding to the ERATS is substantially higher than the two others both for the left (downside) and right (upside) part.

where ERATS designates the RATS level computed using EVT. The maximum acceptable ES replaces the maximum VaR used in the previous RATS formulæ presented in this article. The Max ES can be computed through the rule of thumb that one would not want an ES higher than the one corresponding to the Max VaR for a normal distribution; this value, in turn, can be easily computed from the VaR, as for normal distributions, ES and VaR are related by the formula: $ES \simeq 1.26 \times VaR$. This yields the ERATS level, and as was done with the CS RATS, the daily returns for the following week are adjusted using the ERATS as the leverage indicator. Note that there is no longer a simulation step as in the CS RATS. This is not necessary as the ES coupled with EVT takes into account the entire left tail of the return distribution, thereby modelling downside risk accurately. This, in turn, means that the computational time is dramatically reduced and the ERATS algorithm can be readily integrated in trading platforms and provide real time extreme risk management. We now turn to the application of the previous approaches to compute RATS levels from daily returns as recorded by the Asian strategy from January 2008 to April 2009, using the Sharpe ratio based RATS, the CS RATS and the ERATS.

5 Results

	Original Returns	Sharpe RATS	CS RATS	ERATS
Cumulative return	12.6 %	11.4%	13.6 %	14.6 %
Annualised volatility	11.1 %	9.7 %	11 %	10.4 %
Average Leverage	1	1.08	1.22	1.10
Max Daily Gain	2.6 %	1.9 %	2.5 %	2.6 %
Max Daily Drawdown	-2.3 %	-1.9 %	-1.8 %	-1.4 %
CS Ratio	2.2	2.3	2.6	2.8
Sharpe Ratio	0.9	0.9	1	1.2

Table 1: Comparison of returns obtained using Sharpe RATS, CS RATS or ERATS computed weekly.

The numerical results are summarised in Table 1 as well as Figure 1. The ERATS yields higher cumulative returns than the Sharpe RATS or CS RATS with similar leverage overall; in particular, the daily returns with ERATS have

a higher CS ratio (2.8) compared to returns with Sharpe RATS (2.3). Figure 2 shows the logarithm of the Omega function (a novel risk measure introduced in KEATING and SHADWICK [2002a,b] which takes into account the entire shape of the return distribution) of standardised returns; the fact that the Omega curve corresponding to the ERATS strategy is noticeably higher than those related to the Sharpe RATS or CS RATS both on the left side (which represents downside risk) and on the right side (corresponding to upside potential) illustrates the ability of the ERATS to reduce risk while improving returns; in particular, the left part of the Omega curve for the ERATS is very close to the curve corresponding to the normal distribution, a direct consequence of the criterion used for the expected shortfall value. The ability of the ERATS to reduce downside risk (originating from the traditionally fatter left tails encountered in many trading strategies) while preserving upside potential is confirmed in Table 1 as the ERATS strategy presents the lowest Max Daily Drawdown and the highest Max Daily Gain.

References

- A. BALKEMA and L. DE HAAN. Residual life time at great age. *Annals of Probability*, 2:792–804, 1974.
- G. BARONE-ADESI, K. GIANNOPOULOS, and L. VOSPER. VaR without correlations for nonlinear portfolios. *Journal of Futures Markets*, 19(April):583–602, 1999.
- F. BLACK. Studies in stock price volatility changes. *Proceedings of the Meeting of the Business and Economic Statistics Section. American Statistical Association*, pages 177–181, 1976.
- A. CASCON and W. F. SHADWICK. The standard dispersion and its application to risk analysis for portfolio management. *The Journal of Investment Consulting*, 8(2):38–54, 2007.
- P. F. CHRISTOFFERSEN. *Elements of Financial Risk Management*. Academic Press, 2003.
- C. KEATING and W. F. SHADWICK. A universal performance measure. Technical report, The Finance Development Centre, London, 2002a.
- C. KEATING and W. F. SHADWICK. An introduction to Omega. Technical report, The Finance Development Centre, London, 2002b.
- A. J. MCNEIL and R. FREY. Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. *Journal of Empirical Finance*, 7:271–300, 2000.
- A. J. MCNEIL, R. FREY, and P. EMBRECHTS. *Quantitative Risk Management : Concepts, Techniques, and Tools*. Princeton University Press, 2005.
- D. B. NELSON. Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, 59:347–370, 1991.
- D. B. NELSON and C. Q. CAO. Inequality constraints in the univariate GARCH model. *Journal of Business and Economic Statistics*, 10:229–235, 1992.
- K. NYSTRÖM and J. SKOGLUND. Univariate extreme value theory, GARCH and measures of risk. Technical report, Preprint, Swedbank, 2002.
- J. PICKANDS. Statistical inference using extreme order statistics. *Annals of Statistics*, 3:119–131, 1975.
- R. THOMPSON. File note: Risk management – Ruin strategy. Technical report, The Cambridge Strategy (Asset Management) Limited, 2007.
- S. WILLIAMS. How risky is your trading strategy ? Technical report, HSBC Global Research, 2005.

About the authors:

DR ISSAM STRUB: Dr Strub joined the Cambridge Strategy in 2009 as a senior member of the firm research group. Prior to joining the firm, Dr Strub was a graduate student at the University of California, Berkeley, in the Mathematics and Civil & Environmental Engineering departments conducting research in an array of fields ranging from Partial Differential Equations, Fluid Mechanics and Scientific Computing to Control Theory and Optimisation. He obtained his PhD in Engineering in May 2009 and has authored a number of research articles in leading scientific journals and conferences. He received the Best Systems Paper Award at the 45th IEEE Conference on Decision and Control in 2006 and the Outstanding Graduate Student Instructor Award from the University of California in 2005.

CHRIS UDY: Chris joined the Cambridge Strategy in 2009 and is responsible for the firm ongoing research effort. His current research areas include Bayesian portfolio optimisation, heuristic performance measures, regime-switching models and Markov Chain Monte Carlo sampling techniques. Prior to joining the firm, Chris worked for Mount Row Capital, a systematic hedge fund, from 2004 where he was Chief Technology Officer responsible for developing their proprietary trading programmes across various asset classes. Before this, Chris worked for RadioScape Ltd (a London based mathematical modelling company) as a Research Engineer on their 3G Telecommunication System Design and Traffic modelling team where he was awarded a patent for his work in the digital IF arena. Subsequently he rose to hold various research positions in their Digital Audio Broadcast (DAB) stack design and product development teams. Chris has 1st class Honours degree in Engineering from the University of Auckland.

Disclaimer: Some services are not available to private inexperienced investors. Services may also not be available to certain investors due to regulatory or other constraints either in the UK or elsewhere. The information contained herein is not targeted at the residents of any particular country and is not intended for distribution to, or use by, any person in any country or jurisdiction where such distribution or use would be contrary to local law or regulatory requirements. Information contained herein may be subject to change without notice. Investments in securities for financial instruments (which include contracts for differences, futures, options spot and forward foreign exchange and off-exchange contracts) can fluctuate in value and you should be aware that you may not realise the initial amount invested and may incur additional liabilities. As investments in securities or financial instruments may entail above average risks you should carefully consider whether your financial circumstances permit you to invest and if necessary seek the advice of an independent financial adviser. Foreign exchange denominated securities and financial instruments are subject to fluctuations in exchange rates that may have a positive or negative effect on the value, price or income derived from the securities or financial instrument concerned. Past performance is not a reliable indicator of future performance. You are advised that The Cambridge Strategy (Asset Management) Limited is unable to provide advice as to tax consequences of a particular investment or investment strategy and you are advised to seek professional advice in this respect. This document is issued by the Cambridge Strategy (Asset Management) Limited and is not a solicitation or instruction to invest. It is provided for information purposes only. The Cambridge Strategy (Asset Management) Limited is authorised and regulated by the Financial Services Authority and registered with the SEC (US) and the SFA (HK). The Cambridge Strategy (Asset Management) Limited is exempt from the requirement to hold an Australian financial services licence under the Corporations Act 2001 (Ch) (Class Order 03/1099) in respect of the provision of financial services. It is regulated by the Financial Services Authority (FSA) under UK laws, which differ from Australian laws. Australian Investors: These materials are provided by a representative of The Cambridge Strategy (Asset Management) Limited ('Cambridge') and is intended for wholesale clients as defined in the Corporations Act 2001 (Ch). The information in this presentation is current unless stated otherwise and Cambridge is not under any obligation to update the information to the extent that it is or becomes out of date or incorrect. It is confidential and has been prepared by Cambridge solely for use in connection with its Programmes. This information must not be made available, published or distributed to any third party without the prior consent of Cambridge. This information has been prepared without taking into account anyone's objectives, financial situation or needs so before acting on it each person should consider its appropriateness to their circumstances before making any investment decision. In particular, a person should consider the Programme's investment objectives, risks, fees and other charges. Each person should carefully read and consider any offer documentation before making an investment decision. The information in this presentation is indicative and may change with market fluctuations. It does not purport to be a comprehensive statement or description of any markets or securities referred to within. Cambridge assumes no fiduciary responsibility or liability for any consequences financial or otherwise arising from any reliance on this information. Each person should make their own appraisal of the risks and should consult to the extent necessary, their own legal, financial, tax, accounting and other professional advisors in this respect to any investment in the Programmes. This document is not a solicitation or instruction to invest. It is provided for information purposes only.