Original Article

Implementable tail risk management: An empirical analysis of CVaR-optimized carry trade portfolios

Received (in revised form): 24th May 2011

Hakan Kaya

is a vice president of Neuberger Berman and joined the firm in 2008. Dr. Kaya is a member of the Quantitative Investment Group (QIG) and is a portfolio manager for the commodity strategies and a member of the research team for asset allocation strategies. He focuses on research and asset allocation with an emphasis on portfolio risk management. Before joining the firm, he was a consultant with Mount Lucas Management Corporation where he developed statistical relative value and directional models for commodities investment, focusing mainly on the agricultural sector, as well as models for weather risk management. Dr. Kaya received a BS in Mathematics and Industrial Engineering from Koc University in Turkey and holds a PhD in Operations Research & Financial Engineering from Princeton University.

Wai Lee

is a Managing Director of Neuberger Berman and the Chief Investment Officer and Director of Research for the Quantitative Investment Group with overall responsibility for the quantitative investment function. Previously, he was the Head of the Quantitative Engineering group at Credit Suisse Asset Management (CSAM), responsible for the application of quantitative research across various products and strategies. He joined CSAM in 2000 from J.P. Morgan Investment Management, where he was in charge of quantitative research and risk management for the Global Balanced group. Previously, he was a postdoctoral research fellow at the Harvard Graduate School of Business. Lee is the author of the book *Theory and Methodology of Tactical Asset allocation*. He has been serving on the Advisory Board of the *Journal of Portfolio Management* since 1997. His research work has appeared in academic refereed journals and industry journals. Lee holds a BS (Hon) in Mechanical Engineering from the University of Hong Kong and MBA and PhD degrees in Finance from Drexel University.

Bobby Pornrojnangkool

is a Senior Vice President of Neuberger Berman and a member of the Quantitative Investment Group. His primary responsibility is the research, development and implementation of global macro strategies. Before joining the firm, he worked at the World Bank, where he developed a quantitative currency trading system for its pension plan. Previously, Dr Pornrojnangkool did consulting work for Citibank on an econometric modeling project and for an insurance consulting subsidiary of Seabury Group on several risk modeling projects. Dr Pornrojnangkool earned a BA (Hon) in Business from Chulalongkorn University, Thailand, an MS in Finance from University of Wisconsin at Madison, and holds a PhD in Finance from Columbia Business School, where he also spent 2 years in its MBA program. He is also a GARP Certified Financial Risk Manager.

Correspondence: Hakan Kaya, Quantitative Investment Group, Neuberger Berman, 605 3rd Avenue, 38th Floor, New York, NY 10158, USA

ABSTRACT Although it is relatively easy to identify limitations of the mean-variance framework in managing tail risk, offering a coherent, implementable alternative is more difficult. The goal of this article is to propose an implementable solution to the significant puzzle of portfolio construction in non-Gaussian markets, that is, in markets with more rare events than expected in a mean-variance framework. In this article, we explain how we improved on traditional risk management approaches with heavy-tailed distributions tailored to take into account extreme comovements. Our findings show that this new

quantitative asset allocation method, with non-Gaussian dynamic risk models, leads to enhanced downside protection without constraining upside potentials. Finally, we conduct out-of-sample tests to demonstrate these risk control capabilities on a currency carry portfolio allocation example.

Journal of Derivatives & Hedge Funds (2011) **17**, 341–356. doi:10.1057/jdhf.2011.15;
published online 30 June 2011

Keywords: copula; CVaR; GARCH; portfolio optimization; tail risk

INTRODUCTION

‘In a rare unplanned investor call, the bank revealed that a flagship global equity fund had lost over 30 percent of its value in a week because of problems with its trading strategies created by computer models. In particular, the computers had failed to foresee recent market movements to such a degree that they labeled them a *25-standard deviations event* – something that only happens once every 100,000 years or more’.¹

In our opinion, any viable investment process should comprise two important themes: return generation and risk management. In return generation, regardless of the investment approach, be it qualitative or quantitative, the attractiveness of each asset is implicitly or explicitly modeled to generate expected relative rankings or, in general, return forecasts. As long as the portfolio manager gets the relative attractiveness of assets right, positive returns may be generated. All else being equal, the success of return generation may easily be measured. Risk management, on the other hand, is not as well defined. Although risk is defined as quantifiable uncertainty by Knight,² appropriate measures of risk remain debatable. Perhaps because of the ambiguity of measuring risk, a typical risk report these days always includes a set of statistical risk measures. As a result, success of risk management is not as easily evaluated.

Markowitz’s³ approach, which is referred to as mean-variance analysis, is convenient and simple because it only requires estimation of expected returns, variances, the correlation matrix and nothing else. Putting aside expected returns, the remaining so-called second-moment risk measures can be estimated based on historical data and these estimates can be used as inputs into a quadratic programming solver to come up with optimal portfolio allocations. Following the seminal work of Markowitz, as well as the convenience of assuming a normal distribution, investment science started to measure and model risk as standard deviation of returns in an attempt to capture the dispersion or uncertainty of the investment outcome.

In practice, however, many investors appear to pay more attention to losses than to gains of the same magnitude and, as such, the goal of risk management is often interpreted as the mitigation of investment losses by measuring, forecasting and monitoring the uncertainty of asset returns in order to establish relative hedging positions. It was not until the early 1990s that more modern risk management practices emerged in addressing the apparent asymmetry in the perception of risks. Following the Securities and Exchange Commission and Bankers Trust in the late 1980s, in 1994, J.P. Morgan launched the RiskMetrics service, which offered the first widely followed

drawdown risk measure called value-at-risk (VaR). This measure gained further popularity in 1995 when the Basel Committee on Banking Supervision encouraged its use to determine capital requirements against market risks.

Briefly, VaR tells the investor about the probable magnitude of loss in the worth of a security or portfolio over a given period for a given confidence level. For example, a 99 per cent confidence 1-day VaR of 2 per cent (or simply 99 per cent 1-day VaR is 2 per cent) means there is less than a 1 per cent probability of losing more than 2 per cent within the next day. A 95 per cent 21-day VaR of US\$1 million means that there is less than a 5 per cent probability of losing more than \$1 million within the next month (21 trading days). A dollar VaR can be translated into percentage VaR simply by dividing the dollar VaR by the current value of the investment. Although it is intuitive and improves on variance by taking into account the negative tail of the return distribution, VaR does have some drawbacks that limit its use in portfolio optimization. Mainly, VaR can be shown as an incoherent risk measure⁴ as it does not encourage diversification. Furthermore, because of its non-convexity, efficient global optimization by numerical algorithms cannot be guaranteed.

As a better alternative to VaR, conditional value-at-risk (CVaR) solves the problem of incoherency. CVaR, in simple terms, measures the expected loss, during a given period at a given confidence level. For example, if the investor expects to lose 70 per cent in a month within the 5 per cent worst-case scenarios, the 70 per cent loss in this example is known as the monthly CVaR at 5 per cent confidence. Given that CVaR incorporates both the chance and expected magnitude of loss, it better

summarizes the extreme risks that can be realized, whereas volatility and VaR cannot directly infer potential tail events.⁵

Although we have seen some of the more sophisticated managers starting to report CVaR in their risk reports, its use remains largely as a statistic measuring the output of a portfolio rather than as part of the objective function of portfolio construction. The main reason is that optimizing CVaR is technically a lot more challenging than minimizing variance as in the Markowitz approach. Similar to VaR, CVaR is defined with respect to a quintile so that when asset returns do not follow Gaussian distributions,⁶ an analytical solution for the optimal CVaR portfolio does not exist. Instead, return distributions are discretized by employing Monte Carlo simulations and the portfolio is then optimized over a sufficiently large set of simulated scenarios.⁷ Unlike mean variance analysis in which only variances and correlations matter, a carefully created scenario matrix can capture stylized facts such as persistence in volatilities, generally known as heteroskedasticities in statistics, and rare events, including occurrences of *25 standard deviations* returns, as well as extreme dependencies when correlations move toward one or minus one.

Empirical analysis of many of the stylized facts in asset returns briefly discussed above has been well documented. As noted, although criticizing a particular framework such as mean-variance is easy, offering an alternative, coherent framework is difficult. The goal of this article is, more than just showing how traditional methods become suboptimal in non-Gaussian markets, to knit the pieces together to offer an implementable solution to the big puzzle of portfolio construction in such market environments.

To this end, in the next section, we discuss how we handle the challenge of volatilities that appear to be time varying and persistent. The subsequent section introduces fat-tail modeling of asset returns in an attempt to capture rare events that are deemed almost unlikely in accordance with the mean-variance approach. Next, the section following that introduces the uses of copula functions in joining individual fat-tailed return distributions to complete the multivariate modeling of all asset returns. Finally, after we describe the tail risk problem in the penultimate section, the final section compares the back-test results of CVaR optimization against the traditional mean variance optimization by using the currency carry trade as a case study.

VOLATILITY CLUSTERING

Different regimes of riskiness as measured by, for example, volatilities have been well documented. For instance, the mid-1990s can be classified as a less risky period when investors enjoyed a relatively stable growth in wealth with low anxiety. However, in 2008, we witnessed one of the most unusual and volatile market environments ever as risky assets swung up and down with no clear indication as to what range they would trade in, and how long the anxiety would last.

Numerous empirical analyses of financial data conclude that volatility is persistent so that large changes in asset returns tend to be followed by large changes, both positive and negative, and small changes tend to be followed by small changes, a phenomenon usually referred to as volatility clustering, autocorrelation or serial dependence in volatility. Technically, although returns themselves may be uncorrelated through

time, absolute returns (or squared returns) display positive, slowly decaying autocorrelations. Since the seminal work of Engle⁸ on autoregressive conditional heteroskedasticity (ARCH), conditional volatility research has resulted in many different versions of this model all mainly trying to capture volatility clustering with additional stylized facts. Among these, the generalized ARCH (GARCH)⁹ is probably the most widely employed owing to its parsimony. A lucid description and application of this class of models is in Lee and Yin.¹⁰ Nevertheless, we provide some details below for the sake of completeness.

Let r_t denote the demeaned returns of an asset and assume $r_t = \nu_t z_t$, where z_t 's are independent and identically distributed standard normal random variables. In this model, the variable ν_t captures the properties of volatilities of the asset. In detail, in a GARCH(p, q) system, the dynamics of these volatilities can be written as

$$\nu_t^2 = \alpha_0 + \sum_{i=1}^q \alpha_i r_{t-i}^2 + \sum_{j=1}^p \beta_j \nu_{t-j}^2$$

with constraints that are needed for stationarity

$$\alpha_0 > 0$$

$$\alpha_i \geq 0 \text{ for each } i \in \{1, 2, \dots, q\}$$

$$\beta_j \geq 0 \text{ for each } j \in \{1, 2, \dots, p\}$$

$$\sum_{i=1}^q \alpha_i + \sum_{j=1}^p \beta_j < 1 \quad (1)$$

For instance, GARCH(1, 1) forecasts the volatility today, ν_t , as a weighted average of (i) a constant, α_0 , which captures the long-term variance of the asset, (ii) yesterday's squared demeaned return, r_{t-1}^2 , and (iii) yesterday's forecast of variance, ν_{t-1}^2 . As α_1 is non-negative, a big jump in prices today implies a higher-than-normal volatility forecast for tomorrow.

Similarly, given today's high daily volatility forecast and a non-negative β_1 , this will again force the volatility forecast for tomorrow to be higher than average. Thus, a regime of high volatility is assured after jumps.

Although it may not be relevant for all asset classes, one of the weaknesses of the GARCH model is that it assumes that positive and negative shocks on returns have the same effects on volatility. In practice, it is well known that asset returns (especially equities) respond differently to positive and negative shocks. For example, volatility tends to be lower or at least stable as an asset price is going up. However, when an asset price is plummeting, the arrival of negative shocks tends to have a larger impact on volatility, commonly known as the leverage effect.

Inclusion of the leverage effect in asset returns is addressed in asymmetric GARCH models such as in EGARCH¹¹ or in GJR.¹² The latter simply extends (1) with the inclusion of indicator functions as follows:

$$v_t^2 = \alpha_0 + \sum_{i=1}^q \alpha_i r_{t-i}^2 + \sum_{j=1}^p \beta_j v_{t-j}^2 + \sum_{i=1}^q \gamma_i \xi_{t-i} r_{t-i}^2$$

where $\xi_{t-i} = 1$ if $r_{t-i} < 0$ and $\xi_{t-i} = 0$ otherwise with constraints that are needed for stationarity

$$\alpha_0 > 0$$

$$\alpha_i \geq 0 \text{ for each } i \in \{1, 2, \dots, q\}$$

$$\alpha_i + \gamma_i \geq 0 \text{ for each } i \in \{1, 2, \dots, q\}$$

$$\beta_j \geq 0 \text{ for each } j \in \{1, 2, \dots, p\} \text{ and}$$

$$\sum_{i=1}^q \alpha_i + \sum_{j=1}^p \beta_j + \frac{1}{2} \sum_{i=1}^q \gamma_i < 1 \quad (2)$$

To illustrate the effect of asymmetry in GJR-GARCH(1, 1), assume that we experience a

price decline. With GARCH(1, 1)'s forecast as the starting point, the additional term, $\gamma_1 r_0^2$, where r_0 denotes today's negative return, gives rise to a higher volatility forecast for the subsequent day.

As an example, we apply GJR-GARCH modeling to Canadian dollar excess returns. The graph on the upper left of Figure 1 shows the weekly excess returns. The larger scale in recent years indicates that the volatility of returns increased over time. The autocorrelations of the weekly absolute excess returns in the second plot on the top suggests that the current week's absolute return is correlated with the previous weeks' absolute returns, and the correlation is quite persistent. The lower left plot graphs the history of the annualized volatility forecasts, v_t 's. It confirms that volatility forecasts of the Canadian dollar excess return have jumped to much higher levels, consistent with the much higher realized returns magnitude. Finally, to demonstrate that the GJR-GARCH(1, 1) successfully captures the observed volatilities of the Canadian dollar, we feed the original time series of excess returns into the model, and then save the residuals for further diagnostics, a process known as filtering. Next, we plot the autocorrelations of the filtered return previously discussed as the fourth chart in Figure 1. The filtered returns no longer show autocorrelations, and therefore we conclude that asymmetric volatility clustering has been captured by the GJR-GARCH(1, 1).

FAT TAILS

Table 1 reports the summary statistics of the weekly excess returns of six major currencies against the US dollar from January 1990 to August 2009. These include the Japanese yen (JPY), Eurodollar (EUR), British pound (GBP),

Figure 1: Conditional volatility of weekly Canadian dollar (CAD) excess returns.

Note: Upper left plot shows the weekly excess return series of CAD. Upper right plot shows the autocorrelations of absolute excess returns of the CAD. Lower left plot graphs the conditional volatility of a GJR-GARCH(1, 1) model. Lower right chart shows the autocorrelations of absolute returns standardized by conditional volatilities. The red horizontal lines in the right plots define the intervals to reject 0 autocorrelations.

Source: Quantitative Investment Group (Period from 3 January 1990 through 12 August 2009).

Australian dollar (AUD), Canadian dollar (CAD) and Swiss franc (CHF). Consider the CAD, the least dramatic in this sample, as an example.

The worst week of CAD in this sample period is a loss of 5.81 per cent, an event with 0.000001 per cent probability, given its 1.02 per cent standard deviation according to the normal distribution assumption. Hence, in a normally distributed world, we expect this event to occur

once in every hundred million weeks.

Furthermore, its kurtosis of 8.05, a measure of tail thickness in the return distribution, is significantly higher than the kurtosis of 3 for a normally distributed random variable. Finally, the skewness of -0.17 suggests that the CAD excess return distribution is likely not symmetric around its mean, unlike the bell-shaped normally distributed random variable. In short, a normal

Table 1: Summary statistics of weekly currency excess returns

	<i>JPY</i>	<i>EUR</i>	<i>GBP</i>	<i>AUD</i>	<i>CAD</i>	<i>CHF</i>
Mean	−0.01%	0.03%	0.05%	0.05%	0.02%	0.02%
SD	1.49%	1.48%	1.36%	1.54%	1.02%	1.55%
Median	−0.12%	0.03%	0.11%	0.10%	0.01%	−0.02%
Min	−6.70%	−6.51%	−9.35%	−15.58%	−5.81%	−6.34%
Max	12.45%	10.74%	5.10%	7.30%	5.50%	11.79%
5th Percentile	−2.10%	−2.24%	−2.11%	−2.38%	−1.49%	−2.38%
95th Percentile	2.42%	2.37%	2.09%	2.34%	1.62%	2.60%
1st Percentile	−3.59%	−3.60%	−3.96%	−4.05%	−2.66%	−3.39%
99th Percentile	4.35%	3.71%	3.32%	3.33%	2.89%	3.88%
Skewness	1.00	0.22	−0.82	−1.21	−0.17	0.52
Kurtosis	9.20	6.51	7.15	14.36	8.05	6.70

Source: Quantitative Investment Group (Period from 3 January 1999 through 12 August 2009).

distribution appears to be a poor approximation of CAD in this sample period, and the same can be said about all other currencies during the same sample period.

To better visualize these anomalies within the observed empirical distributions, we first filter the CAD returns with GJR-GARCH(1, 1) to remove the volatility clusters as discussed in the previous section. The residuals, as a result, are then saved for further analysis, as depicted in Figure 2. The stair lines in the upper right and left plots depict the empirical cumulative distribution function (CDF) in the left (negative) and right (positive) tails of the distribution. On the same plots, we also include the CDF of a fitted corresponding normal distribution for the purpose of comparison. In each case, the normal distribution approaches 0 and 1 at both ends much faster than the empirical distribution. To observe this better, in the lower left chart, we plot the quintiles of the fitted normal distribution against the quintiles of the filtered CAD weekly excess return residuals. If, indeed, these two distributions were the same,

we would expect that the respective quintiles lie on a straight line. Although a good match is observed in the middle of the distribution, significant deviations in both tails are observed. Last but not least, when we employ a statistical goodness-of-fit test (Jarque–Bera normality test in this case), we reject normal distribution at the 5 per cent level. To conclude, the filtered CAD excess return distribution has much fatter tails than a normal distribution.

We note above that the CDF of a normal distribution decays quickly at both ends with exponential speed. Hence, a slower decaying function may be used to better approximate the tail probabilities. To this end, we first partition the observed returns into three disjoint ranges, covering the lower tail, middle and upper tail, respectively. Next, we fit a generalized Pareto distribution (GPD) to the tails, and a cubic spline to the middle.^{13–16}

Going back to Figure 2, we can see in the upper plots that the Pareto distribution (GPD) closely converges to the observed data in the

Figure 2: Filtered return distribution calibrations for CAD.

Source: Quantitative Investment Group.

tails, and hence is more reliable in extrapolating beyond an empirical distribution when compared with the normal distribution. The lower right plot shows that the quintiles of the empirical distribution of CAD weekly excess return residuals and generalized Pareto distribution quintiles match quite closely as we do not observe any consistent deviations from the straight line anymore, unlike the case with a normal distribution in the lower left plot. Finally, a Kolmogorov–Smirnov goodness-of-fit test confirms that the two distributions are statistically indistinguishable at level 5 per cent.

So far, we have demonstrated how to successfully approximate the observed distributions of individual asset returns through a combination of GARCH, splines and GPD. The next section shifts the focus on knitting them together in order to capture comovements.

EXTREME DEPENDENCY

An important feature of risky financial assets is that they exhibit strong positive or negative dependency during turbulent times owing to changes in the perception of crash risks.¹⁷ This

is especially the case in currencies as the unwinding of carry trades can lead to significant depreciations in 'investment currencies' and large appreciations in 'funding currencies'. In October 2008, for instance, we witnessed a similar episode when JPY, as a funding currency, assumed one of its highest returns, whereas other investment currencies realized quite dramatic losses simultaneously.

Although the covariance matrix may correctly identify the directions of the dependencies, it is not as informative when it comes down to comovements near the tails of the distributions. For example, conditional on the occurrence of a rare, four standard deviation event in one currency, a covariance matrix-driven multivariate normal distribution gives almost 0 per cent probability to a simultaneously rare, four standard deviation event in a different currency. In addition, even if one can successfully forecast extreme correlations among some assets in the universe to approach one or minus one, it is not straightforward to incorporate forecasts of just several elements into the correlation matrix, which has a very rigid structure that comes with its own integrity. Any changes introduced to some elements of a matrix can potentially destroy its semi-positive definiteness property, which is required in order to guarantee that the resulting portfolio standard deviation is non-negative. As a result, hedging away negative skewness and fat tails with a regular covariance matrix may not be as easy as it seems.

A potential solution to this challenge is the utilization of a technique called copula. First developed by Sklar,¹⁸ a copula is a function that joins univariate marginal distributions through their quintiles. It does so by taking into account interrelations of its constituents, as well as (depending on its type) extreme dependencies

in the tails.^{18–20} Its flexibility in linking general distributions has made copulas popular in a variety of fields. Li²¹ was among the first to introduce the copula into the financial industry by using it to model default correlations for credit default swap valuation purposes. The goal was to link exponentially distributed survival times until default so as to satisfy default dependency structures between companies. Yet in another application, Kaya²² studies copulas to generate weather scenarios over a region of agricultural districts in a way to preserve observed spatial and temporal climatological relationships. In this analysis, after calibrating heavy-tailed precipitation distributions for each district, a copula is used as a tool to simulate correlated rainfalls.

To illustrate how normal distribution fails to characterize some extraordinary comovements, we use the weekly excess returns of EUR and CHF from January 1990 to August 2009 as an example. We first plot in Figure 3 the 10 000 scenarios generated from a bivariate normal distribution with a historical covariance matrix estimate based on the data in the same sample period. As one can see in Figure 3, the simulated scenarios generated by a bivariate normal distribution are symmetric and bounded between four standard deviations from both tail ends.

Figure 4 presents the scatter plot of the observed standardized excess returns of EUR versus CHF, as well as the corresponding histograms. Although the bulk of the data is concentrated in the middle of the plot around the means, in sharp contrast to the simulated scenarios in Figure 3, there are a number of weeks when both currencies realize four standard deviations or more of losses. In particular, there is one week when EUR

Figure 3: Standardized weekly excess returns of EUR versus CHF using bivariate normal distribution simulation with a covariance matrix.

Note: Scatter plot of the simulated GJR-GARCH(1, 1) standardized weekly excess returns of the EUR CHF pair under bivariate normal distribution assumption.

Source: Quantitative Investment Group (Period from 3 January 1990 through 12 August 2009).

appreciates around five standard deviations and the CHF simultaneously gains more than six standard deviations. These outliers clearly show that a bivariate normal distribution can be a poor approximation for the joint behavior of this particular pair of currencies, especially when quantifying potential impacts of tail events that are considered critical in risk management.

To circumvent this, we first model marginal distributions by fitting fat-tailed distributions to the residuals of GJR-GARCH(1, 1) processes as described in the sections ‘Volatility clustering’ and ‘Fat tails’, and calibrate a t-copula to join these marginal distributions. While still requiring the estimation of a

correlation matrix, the t-copula additionally parameterizes the tail dependency. This so-called tail parameter allows us to extrapolate multivariate fat-tailed distributions in a way that is consistent with historical realizations. As such, with the t-copula, it is possible to allocate non-zero probabilities to joint outliers. To show this, we plot the simulated 10 000 scenarios from the t-copula-driven, bivariate, fat-tailed distribution in Figure 5. This simulated set of scenarios approximates the observed empirical distributions in Figure 4, with many more rare events and tail dependency, far better than the bivariate normal distribution in Figure 3.

Figure 4: Standardized weekly excess returns of EUR versus CHF using GARCH conditional volatilities.

Note: Scatter plot of the t-copula simulated GJR-GARCH(1, 1) standardized weekly excess returns of the EUR CHF pair.

Source: Quantitative Investment Group (Period from 3 January 1990 through 12 August 2009).

OPTIMIZATION WITH TAIL RISK: PUTTING IT ALL TOGETHER

As we discussed above, portfolio construction that takes tail risks into consideration is far more challenging, and therefore is not yet commonly adopted. The highly nonlinear nature and inclusion of unconventional risk measures such as CVaR introduce further non-convexities into the problem. The copula-driven multivariate distribution, for example, does not lead to tractable analytical solutions for optimal asset allocation. Needless to say, under these conditions, any optimization trial will end up resulting in suboptimal portfolios.

Instead, we can employ a scenario-based optimization approach where we first simulate

random samples with a set of distribution assumptions, then optimize the portfolio under these forward-looking scenarios. This approach has been implemented and tested in many neighboring domains. For instance, Mulvey *et al*²³ solve a large asset liability system on a dense scenario tree for pension plans. A similar study by Schwartz and Tokat²⁴ compares asset allocation under normal and fat-tailed scenarios based on a multi-period asset allocation model using different risk measures, concluding that scenarios generated from normal distributions significantly underestimate risks. On the active management space, Giacometti *et al*²⁵ show how the Black–Litterman model²⁶ can be improved

Figure 5: Standardized weekly excess returns of EUR versus CHF using t-copula simulation.
Note: Scatter plot of the GJR-GARCH(1, 1) standardized weekly excess returns of the EUR CHF pair.

Source: Quantitative Investment Group (Period from 3 January 1990 through 12 August 2009).

with realistic models of asset returns in a scenario-driven optimization model.

In our analysis, we use CVaR as a tail risk measure. Again, α per cent-VaR measures the amount of capital required to prevent a negative balance that occurs $1 - \alpha$ per cent of the time, and the α per cent-CVaR is the expected loss if a loss greater than or equal to the VaR materializes.²⁷

Under certain conditions, the CVaR function, denoted by $\theta_\alpha(\mathbf{w})$, can be discretized to a convex and piecewise linear function, $\tilde{\theta}_\alpha(\mathbf{w})$. The discretization and its following linearization are carried out by generating samples from the distribution of returns, $r \in \mathcal{R}^N$.²⁸ This approximation allows us to nest the copula-

driven, fat-tailed simulation scenarios in an optimization problem. Hence, given an expected return vector $\mu_t = E_t[r] \in \mathcal{R}^N$ at time t , we can solve

$$\begin{aligned} &\text{Maximize} && \mu_t w \\ &\text{Subject to} && \tilde{\theta}_\alpha(w) \leq \gamma \\ &&& w \in X \end{aligned} \quad (3)$$

to obtain the weights of a portfolio that maximizes the expected return while controlling CVaR at $\gamma \in \mathcal{R}$ at confidence level α . Here, $\mathbf{X} \subseteq \mathcal{R}^N$ can constrain market neutrality, dollar neutrality, short sales, transaction costs and the like. We call this optimization framework the Mean-CVaR problem.

IMPLEMENTATION: CARRY TRADE CASE STUDY

In this section, we compare the mean variance (MVO) and mean CVaR (MCO) optimization frameworks side by side using the out-of-sample performance of currency carry trade portfolios with the same universe of the six currencies in Table 1 for an example. Carry trade has been documented as a profitable strategy over the long term, but is often observed with crash risk having significant drawdowns.¹⁷ We also include a naive sorting strategy to determine whether these two optimization frameworks improve investment performance in terms of better risk management.

Controlling for expected return estimates allows us to focus merely on risk. To this end, we run monthly rebalanced strategies that all use the same expected returns as estimated by the interest rate differentials of the 1-month risk-free rate observed 1 day before the portfolio rebalancing dates. Hence, in the absence of currency and default risks, the currencies with the highest yields would be our optimal positions.

The naive portfolio (NAI) will go long in the three highest yielding currencies (investment currencies) and short in the three lowest yielding currencies (funding currencies). The weights assigned among long and short groups are all equal weighted, and hence the portfolio weights sum to zero. With this approach, no additional information is used to capture the dependence among currencies, time-varying volatilities and other usual statistics that are believed to be relevant in risk management.

The mean-variance portfolio requires a covariance matrix, which is estimated with weekly return data. In order to capture the dynamic behavior of changing correlations and

variances, we use a time-weighted covariance matrix with a half life of 52 weeks.²⁹ We further force the optimized weights to sum to 0 to have a dollar neutral portfolio, and constrain the annualized portfolio volatility to be less than 10 per cent.

The mean-CVaR portfolio requires the confidence level, α , and CVaR bound, γ . Without loss of generality, we set the confidence level α at 95 per cent and the monthly CVaR bound, γ , at 6 per cent to make it comparable to the observed risk characteristics of the resulting MVO portfolio so that we can compare these approaches at similar risk levels.

Figure 6 plots the log wealth paths of the optimal portfolios based on these three approaches. During the sample period of January 2000 to mid-2007, the carry trade enjoyed stable returns. However, a significant drawdown started in mid-2007 and lasted until the end of 2008 before the trade started to bounce back,

Figure 6: Out-of-sample test results.

Note: NAI: Naive approach, MVO: Mean variance optimization, MCO: Mean CVaR optimization.

Source: Quantitative Investment Group (Period from January 2000 to August 2009).

Table 2: Out-of-sample test results

	<i>Naive approach (NAI)(%)</i>	<i>Mean variance optimization (MVO)</i>	<i>Mean CVaR optimization (MCO)</i>
Annualized mean	3.32%	6.24%	8.36%
Annualized SD	7.08%	11.64%	11.92%
Monthly min	−8.86%	−10.42%	−8.25%
Monthly max	4.41%	6.96%	8.91%
Monthly VaR (5%)	3.20%	4.97%	4.44%
Monthly CVaR (95%)	5.13%	7.07%	6.12%
Max drawdown	25.29%	32.48%	27.72%
Sharpe ratio	0.47	0.54	0.70
Monthly mean/VaR	8.67	10.44	15.71
Monthly mean/CVaR	5.40	7.35	11.39
Monthly Mean/MaxDD	1.10	1.60	2.51

Source: Quantitative Investment Group (Period from January 2000 through August 2009).

making it a good candidate for evaluating tail risk management.

Table 2 tabulates the performance statistics. The out-of-sample realized risk of MVO turns out to be 11.64 per cent, which is considerably higher than its 10 per cent target. On the contrary, the MCO target of monthly 6 per cent CVaR is closely matched at the level of 6.12 per cent. This underscores the ability of MCO in consistent risk targeting.

In order to compare the relative performance of these approaches, we report some common reward-to-risk ratios, using different measures of risk including volatility, VaR, CVaR and maximum drawdown. For instance, per 1 per cent risk as measured by standard deviation, MVO achieves 0.54 per cent return, whereas MCO achieves 0.70 per cent, even though its objective is not to maximize the Sharpe ratio as in the case of MVO. Ratios such as mean-to-VaR, mean-to-CVaR and

mean-to-maximum drawdown indicate significant improvements when we switch from NAI to MVO, and MVO to MCO.

CONCLUSIONS

We certainly understand, and agree, that risk management goes far beyond the statistical perspectives. Nevertheless, we fail to imagine how a risk management system without statistical measures of risk, however risk is defined, can be sound. In our opinion, risk management must include at least the following three steps: (i) risk measurement, (ii) the incorporation of risk measures into portfolio construction or optimization and, finally (iii) risk monitoring.

Given the near impossibility of a 25-standard deviations event during an investor's lifetime, the standard deviation is almost surely wrong. The importance of tail risk measures has been

growing with more occurrences of crash events. As a result, deviations from the traditional risk measures have commenced and significantly gained momentum in various branches of finance. The incompetence of normal distribution in explaining real-world returns and consequently the underestimation of risks in the mean-variance framework lead to the analysis of fat-tailed extreme dependency distributions in risk modeling.

In this article, we employ modern statistical methods to better model the stylized facts about return distributions. Although some of these methods have been documented elsewhere and even implemented to some extent in the investment industry, we add to the literature by showing in a portfolio tail risk management case how a combination of these allows us to approximate the observed distributions of returns much more effectively than would a typical multivariate normal distribution. In summary, we show how persistence in volatilities can be filtered via GARCH models to achieve stationarity for fat tail and dependency modeling. We next fit fat-tailed distributions to the filtered data and calibrate copulas to model extreme dependencies.

Many of the statistical risk measures we discuss may be found in some of the more sophisticated managers' risk reports. However, probably because of the technical challenges, these measures are largely used for the purpose of risk monitoring after positions are taken, rather than as part of the objective function in the portfolio optimization stage, a step that we insist must be included in order to implement sound risk management. In the final part of this article, we demonstrate how to implement this important step by using the currency carry trade as an example. We simulate fat-tailed

scenarios with possible extreme dependencies, and optimize portfolio tail risk on these forward-looking samples by constraining CVaR. Results from out-of-sample tests show that, in general, managing risks even with a traditional method such as mean-variance may improve investment performance; however, only a non-normal model with a tail risk control capability appears to be effective in shrinking the size of drawdowns and rare substantial losses.

REFERENCES AND NOTES

- 1 Gangahar, A. and Tett, G. (2007) Limitations of computer models. *Financial Times*, 14 August, http://www.ft.com/cms/s/0/b54f3ea8-4a9-11dc-95b5-0000779fd2ac.html?ncklick_check=1.
- 2 Knight, F.H. (1921) *Risk, Uncertainty and Profit*. Boston, MA: Hart, Schaffner & Marx.
- 3 Markowitz, H. (1952) The utility of wealth. *Journal of Political Economy* 60(2): 151.
- 4 Artzner, P., Delbaen, F., Eber, J.M. and Heath, D. (1999) Coherent measures of risk. *Mathematical Finance* 9(3): 203–228.
- 5 Stemming from its convexity, another important feature of CVaR is its optimizability. Numerical optimization routines only guarantee global solutions under certain conditions. One of these necessary conditions is the convexity of the objective and constraint functions. If the risk measure is not convex, like VaR, then after optimization, it is always possible to find a (stochastically) dominating portfolio which is closer to or on the efficient frontier.
- 6 We use the words Gaussian and Normal interchangeably in describing a statistical distribution.
- 7 Rachev, S.T., Stoyanov, S.V. and Fabozzi, F.J. (2008) *Advanced Stochastic Models, Risk Assessment, and Portfolio Optimization: The Ideal Risk, Uncertainty, and Performance Measures*. Hoboken, NJ: Wiley.
- 8 Engle, R.F. (1982) Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation. *Econometrica* 50(4): 987–1007.
- 9 Bollerslev, T. (1986) Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics* 31(3): 307–327.
- 10 Lee, W. and Yin, J. (1997) Modeling and forecasting interest rate volatility with GARCH. In: F.J. Fabozzi (ed.) *Advances in Fixed Income Valuation Modeling and*

- Risk Management*. New Hope, PA: Frank J. Fabozzi Associates.
- 11 Nelson, D.B. (1991) Conditional heteroskedasticity in asset returns: A new approach. *Econometrica* 59(2): 347–370.
 - 12 Glosten, L.R., Jagannathan, R. and Runkle, D.E. (1993) On the relation between the expected value and the volatility of the nominal excess return on stocks. *The Journal of Finance* 48(5): 1779–1801.
 - 13 See Newman¹⁴ for details of Pareto distribution, its properties, estimation and other uses.
 - 14 Newman, M.E.J. (2005) Power laws, Pareto distributions and Zipf's Law. *Contemporary Physics* 46: 323–351.
 - 15 Cubic splines are functions that smoothly approximate the underlying function as closely as possible by interpolating the observed data. Splines are most powerful when there is ample data. Because we fit them to the center of the distribution where the bulk of the data resides, approximations are therefore less prone to errors. See Deboor¹⁶ for details.
 - 16 Deboor, C. (1978) *A Practical Guide to Splines*. Berlin: Springer-Verlag; Heidelberg: GmbH & Co.
 - 17 Brunnermeier, M.K., Nagel, S. and Pedersen, L.H. (2008) Carry Trades and Currency Crashes. National Bureau of Economic Research. Working Paper 14473, <http://www.nber.org/papers/w14473>.
 - 18 Sklar, A. (1973) Random variables, joint distribution functions and copulas. *Kybernetika* 9: 449–460.
 - 19 To be more precise, if R_i , $i=1, 2, \dots, N$ denotes a list of random variables (such as asset returns), with cumulative distribution functions $F_i(\cdot)$, $i=1, 2, \dots, N$, then $U_i := F_i(R_i)$, $i=1, 2, \dots, N$ are uniformly distributed on the unit interval and the joint distribution function of these U_i , $i=1, 2, \dots, N$ is called a copula. Note that U_i , $i=1, 2, \dots, N$ still preserve the dependency information among R_i 's. If $F: [0,1]^N \rightarrow [0,1]$ denotes the multivariate distribution of R_i , then Sklar¹⁸ proves that for this distribution one can find a copula function $C: [0,1]^N \rightarrow [0,1]$ such that $F(R_1, R_2, \dots, R_N) = C(F_1(R_1), F_2(R_2), \dots, F_N(R_N))$. In other words, the theorem makes it possible to separate the univariate modeling and dependency modeling where the dependence can be captured with a suitable copula. See Nelsen²⁰ for complete technical details.
 - 20 Nelsen, R.B. (1998) *An Introduction to Copulas*. New York: Springer, Lecture Notes in Statistics Series, Vol. 139.
 - 21 Li, D.X. (1999) On default correlation: A copula function approach. SSRN, <http://ssrn.com/abstract=18728>.
 - 22 Kaya, H. (2008) Applying statistical learning theory: Agricultural commodities and weather risks. PhD thesis, Princeton University, Princeton, NJ.
 - 23 Mulvey, J.M., Gould, G. and Morgan, C. (2000) An asset and liability management system for Towers Perrin-Tillinghast. *Interfaces* 30(1): 96.
 - 24 Schwartz, E.S. and Tokat, Y. (2002) The impact of fat tailed returns on asset allocation. *Mathematical Methods of Operations Research* 55: 165–185.
 - 25 Giacometti, R., Bertocchi, M., Rachev, S.T. and Fabozzi, F.J. (2007) Stable distributions in the Black-Litterman approach to asset allocation. *Quantitative Finance* 7(4): 423–433.
 - 26 Black, F. and Litterman, R. (1990) *Asset Allocation: Combining Investor Views with Market Equilibrium*. New York: Goldman, Sachs & Co. Fixed Income Research.
 - 27 In an N assets portfolio setting, CVaR is defined as follows. Let us denote by $f(\mathbf{w}, r)$ the loss of a portfolio $\mathbf{w}$ in a set $\mathbf{X}$ under random returns r . The space $\mathbf{X}$ here denotes the set of acceptable portfolios. Suppose that the random returns have a probability density $p(r)$. Then, the probability of $f(\mathbf{w}, r)$ being less than a threshold ξ can be calculated as $\Psi(\mathbf{w}, \xi) = \int_{f(\mathbf{w}, r) \leq \xi} p(r) dr$. Then, $\xi_\alpha(\mathbf{w}) = \min\{\xi \in \mathbb{R}: \Psi(\mathbf{w}, \xi) \geq \alpha\}$ and $\Theta_\alpha(\mathbf{w}) = (1 - \alpha)^{-1} \int_{f(\mathbf{w}, r) \geq \xi_\alpha} f(\mathbf{w}, r) p(r) dr$ are called VaR at level α and CVaR at level α , respectively.
 - 28 Rockafellar, R.T. and Uryasev, S. (2002) Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance* 26(7): 1443–1471.
 - 29 For example, if a weight of 0.2 is assigned to the latest observation, the observation 1 year ago will be assigned a weight of 0.1. The matrix with infinite half-life and with sum of all weights equal to 1 converges to the regular historical covariance matrix.