Tungsten - Risk Forecasting

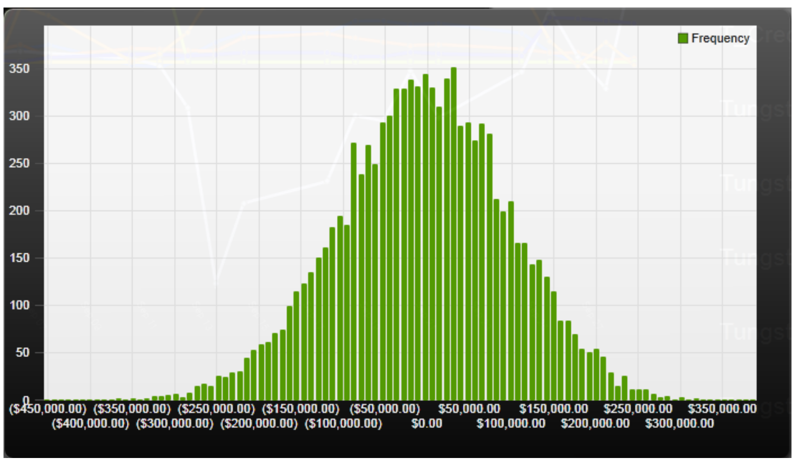

Model Overview

Hamilton, Bermuda

October 2012

Tungsten - Risk Forecasting - Model Overview

This document provides an overview of the Tungsten risk forecasting models. These models are available through the Tungsten GUI, Tungsten API /SaaS and the Eze Software Group Tradar PMS accounting platform.

Introduction

One of the main pillars of risk management is the ability to calculate and monitor firm wide Value At Risk. Value At Risk is a method to forecast risk, and it answers a simple question; what type of portfolio volatility (at 68% confidence) can I expect with my current set of positions for a given horizon and confidence. VaR does not tell us anything about maximum loss, it simply gives us an idea at a given confidence and horizon the expected loss. Let us take a simple example: If our VaR model tells us that our Value At Risk for a given date is $100,000 using model parameters such as 95% confidence and a one day horizon. This means we can expect to see a loss of $100,000 (or more) one day out of 20, or roughly one day per month.

VaR is useful as a firm wide risk metric, as it gives us one number that can be understood and compared across different types of strategies and asset classes on any level of the portfolio. Example one can look at a top level fund VaR and then drill down into each strategy to see how risk is distributed across the books. We can understand what parts of our portfolio are adding risk and what parts are acting as hedges or risk reductors.

Value At Risk works best on portfolios with liquid positions where our risk factors can be updated daily with clean time series data.

There are several models available to calculate VaR, with the most familiar ones being parametric, historical scenario and Monte Carlo simulation.

To get a fuller understanding of the VaR numbers and the dynamics of the portfolio it is common to also look at Marginal VaR and Expected Shortfall (tail risk, CVAR). The Tungsten platform can calculate all on any grouping level.

In Tungsten, Marginal VaR is defined as the change of VaR of the portfolio if a specific risk bucket were to be removed. This is calculated by removing the set of positions constituting the bucket (strategy) and then re-calculating VaR. The difference with and without the bucket is calculated and reported as the Marginal VaR. This is the same definition as implemented in Risk Metrics.

Tail risk or expected shortfall is calculated using the asset distribution result of a historical simulation or a Mont Carlo simulation (standard and hybrid, more about those models in the sections below). The tail risk is then the average loss in the tail at the specific percentile, e.g. 5%.

Tail VaR is best illustrated with the below graph (taken from this page):

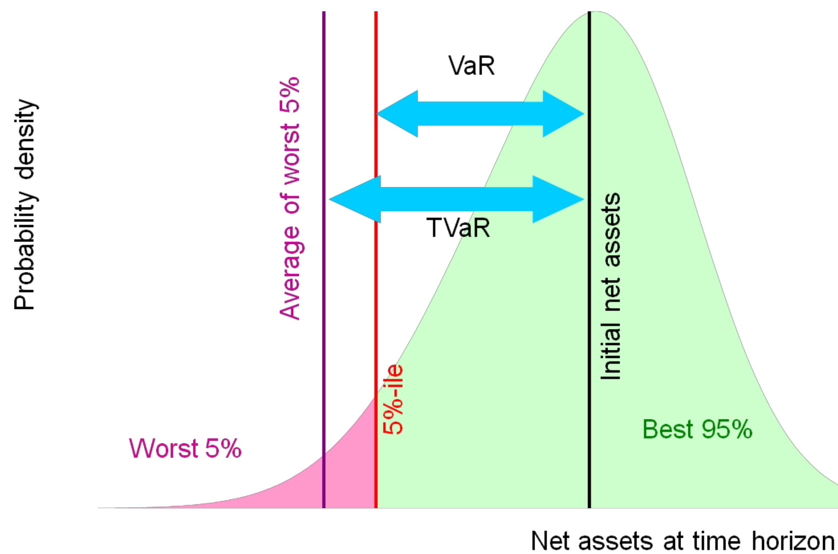

The Tungsten models

At the time of writing, Tungsten utilize four different models, an analytical (parametric) delta-gamma model, two historical scenario based models and Monte Carlo simulation.

Analytical model

The analytical delta-gamma model gives an accurate VaR for linear and simple derivatives portfolios. The model use the co-variances of the risk factors assuming a normal return distribution.

As the parametric model is a closed form solution it is straightforward and efficient to calculate a portfolio VaR on any given portfolio given the position weights (w) and the variance-covariance matrix.

The closed form formula for the *delta-normal* case we can use the following formula: $Var(Rp) = \alpha\sqrt{w^T V w}$ where w is the weights of the portfolio positions, and V is the variance co-variance matrix of the asset returns.

The Loading... is the constant for the specific confidence, e.g. 1.645 for a 95% one tailed normal distribution.

Tungsten is implemented using a *delta-gamma* approach which means we have added the second derivative (gamma), thus the name delta-gamma approximation. The second derivative takes into account the curvature of the non-linear assets in our portfolio, giving us a good estimate on our derivative positions.

Some of the key benefits of the parametric model is the speed of computation and simplicity of implementation.

Scenario based models

The Tungsten scenario based models are based on historical simulation and simulation via Monte Carlo. For the historical simulation models, we have implemented a simple model and a hybrid model based on the work by Boudoukh, Richardson and Whitelaw (1998) (BRW). [“The Best of Both Worlds: A Hybrid Approach to Calculating Value at Risk”]

The historical simulation approach has a few advantages to the analytical parametric model. It is intuitive and we do not assume a normal distribution of the asset returns. It also deals with the inherent nature of fat-tailed asset returns. The approach is relatively straightforward, we use a specific window of returns and re-value our portfolio given the *non-weighted asset* returns in this window.

The resulting distribution is sorted in ascending order and the VaR is then given by the specific percentile, e.g. 5% percentile for a 95% VaR. As the distribution is discrete and depends on the amount of historic data we use, the exact percentile value is calculated using linear interpolation. Given we have the distribution we can also estimate the expected shortfall (tail risk) as explained in the introduction.

The below image shows a sample of a historical simulation distribution. We note that the left tail is elongated compared to a normal distribution.

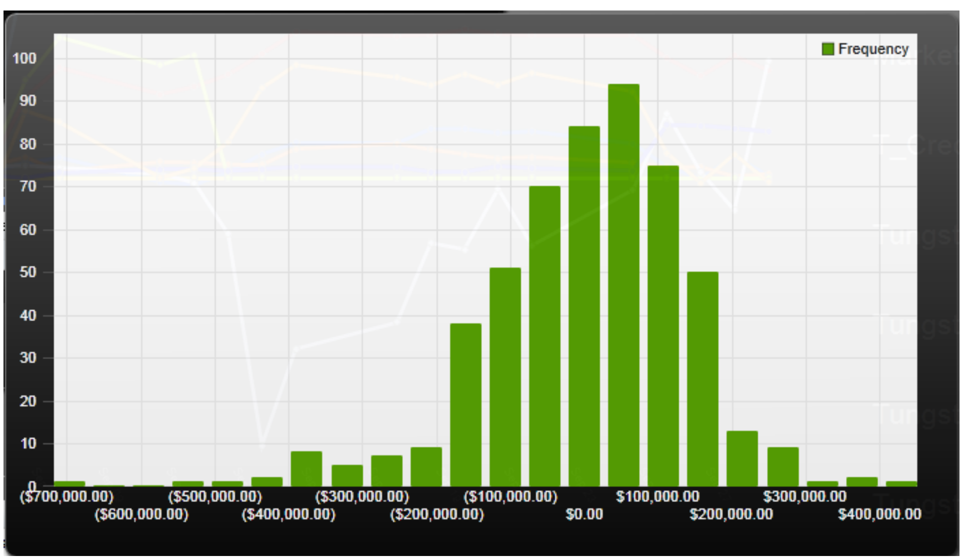

Hybrid model (BRW)

The basis for the hybrid model is the same as simple historical simulation, but instead of using non-weighted asset returns we apply exponentially declining weights to the asset return series. The effect is a more reactive model. i.e. we have no "ghosting effects" and the model generates VaR numbers that are more "in tune" with the current market environment. The weighting factor is given at the time of calculation. The number of days to include in the analysis is determined by the calculation type.

Boudoukh, Richardson and Whitelaw showed in their paper that the hybrid model resulted in a significant improvement in statistical performance over a parametric (analytical) and standard HS model. The improvement is most pronounced in series exhibiting fat-tails.

Monte Carlo Simulation

Another scenario based model is Monte Carlo simulation. Monte Carlo models can solve highly complex financial problems by simulating thousands of scenarios. The simulation is a stochastic process where the result is a distribution of possible outcomes. As in historical simulation we look at the percentiles of the distribution to find our Value At Risk. Monte Carlo models are especially good at measuring VaR on non-linear portfolio's.

The Tungsten's Monte Carlo engine is a utilizing pseudo random number generator to generate the random paths of each risk factor. The random paths are then fitted to the portfolio covariance matrix using a bridged Cholesky decomposition. The simulation model is using a standard Geometric Brownian motion (Black Scholes with zero drift) to generate continuously compounded returns.

The Monte Carlo engine is highly flexible, the user can select anything from the number of simulations (defaulted to 10,000), the random path generator, and distribution assumptions (Gaussian or Student-T) of the risk factors.

Sampling and Decay

Both the Monte Carlo engine and the analytical delta-gamma model can chose sampling and decay as parameter input that will affect the variances and covariances in our sampling universe - Tungsten is set to use daily, weekly and monthly sampling by default. However it is possible to use any custom sampling such as bi-monthly. This is normally calibrated at installation, but can be done at any time.

It is also possible to decay asset returns using an exponentially smoothed time series (EWMA). In general we would use a decay factor of 0.94 on daily returns and 0.97 on monthly returns - based on research by Risk Metrics™. This is fully configurable, depending on the users own requirements. The lower the decay (e.g. 0.94) the asset returns will be more reactive to market changes. A risk forecast model using EWMA is similar to GARCH (1,1), but more straightforward to use with no calibration required.

Variances, covariances and correlations are all decayed. **Note:** It is important that the users understand the implications of using decayed data as the changes in the risk estimates can be quite volatile the stronger decay we use (e.g. a lambda of 0.94 would give stronger weight to current returns than 0.97).

What model to use?

With four different models available at our disposal, which one is the best to use? This depends highly on the application the VaR data is to be used for and the portfolio. If your portfolio is highly linear, containing mostly equity based products and indexes and simple derivatives of those assets, the analytical model works perfectly fine.

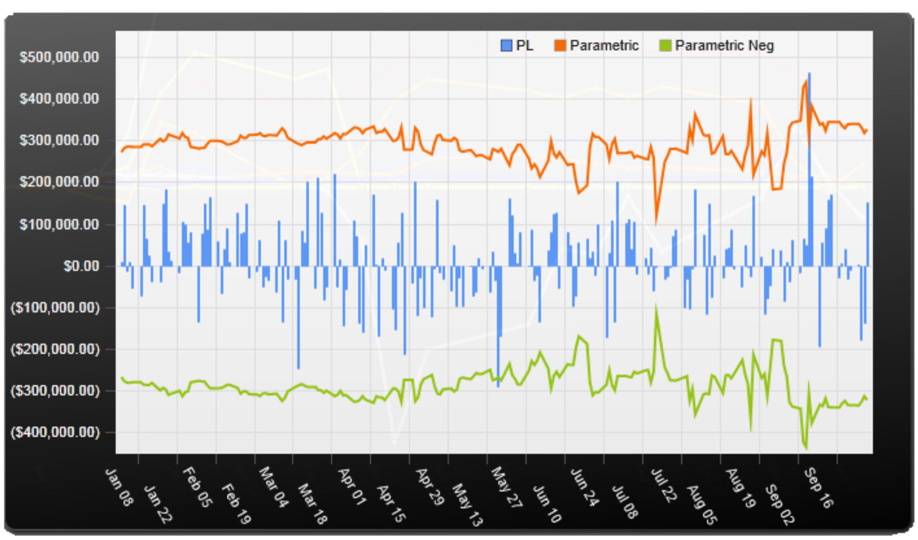

The analytical model give us quick results, and we can efficiently generate a P&L time series overlaid with the daily risk forecast as pictured below. This gives us an idea on how well the model is predicting our risk. In this case, we are looking at a year to date daily risk forecast using a 99% confidence.

We expect to see the model have a breach (downside or upside) roughly once every 100 trading days. In our example, we have had two breaches in 10 months of trading so our parametric model looks pretty good on this portfolio containing equities, currencies and derivatives of those instruments.

However if your portfolio contains assets that exhibit fat tails where the analytical model is poor at capturing these tails, then the simulation models would work better. The standard historical model will give a more constant VaR number that could be a better model to use for example a VaR limit, The hybrid model is highly reactive depending on the weight used, hence the VaR will change more than the standard model rendering it difficult to use for VaR limit management.
Having all four models at your fingertips and the ability to compare the data side by side, will give your better understanding of the dynamics of your portfolio.

riskname	Monte Carlo	Historical	Historical BRW	Parametric
biotech	26,344	36,852	30,274	25,224
brazil	17,646	24,105	15,193	16,870
commercial realestate	1,703	22,320	1,199	1,713
commodities	10,094	16,589	11,850	9,959
conglomerates	62,494	98,599	62,372	61,033
currency tactical	134,599	161,492	153,174	135,184
global macro	14,299	22,149	11,958	13,913
hedge	12,076	10,990	12,592	11,261
metals	9,034	10,347	12,245	8,569
multi strategy	15,727	31,876	18,965	15,445
residential/commercial	12,488	19,977	15,481	11,812
tactical market allocation	70,388	100,450	53,806	66,988
technology	77,014	85,020	71,441	73,898
value	28,906	10,787	15,047	27,116
vertical spread	0	0	0	0
vol_termstruct	110,146	149,507	87,346	101,884
16 rows 00:19 sec	248,020	381,966	229,615	234,965

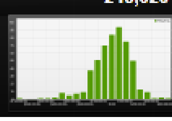

The image above shows all the models in action using the same sampling, confidence and horizon (two year daily returns with 0.94 decay factor, 99% confidence and one day horizon).

We observe that the Total portfolio VaR is relatively similar using the Parametric versus the Monte Carlo model (234,965 vs 248,020 respectively). Historical VAR is significantly higher which can be explained by it is not decaying the returns, and hence the weights from all two years of returns are factored in. This can be seen in the tail of the historical simulation distribution, where there is evidence of an elongated tail.

The Historical BRW model is using a smoothing factor of 0.97 and we note that current returns are given more weight, resulting in a risk reduction similar to the Parametric and Monte Carlo models.

riskname	Monte Carlo	MarginalMonte Carlo	Parametric	MarginalParametric
biotech	26,344	248	25,224	660
brazil	17,646	12,871	16,870	10,411
commercial realestate	1,703	348	1,713	-6
commodities	10,094	3,020	9,959	3,784
conglomerates	62,494	33,069	61,033	32,009
currency tactical	134,599	-64,262	135,184	-55,559
global macro	14,299	6,346	13,913	7,099
hedge	12,076	-8,820	11,261	-8,780
metals	9,034	2,661	8,569	1,471
multi strategy	15,727	4,543	15,445	4,645
residential/commercial	12,488	-2,430	11,812	-2,132
tactical market allocation	70,388	39,714	66,988	37,505
technology	77,014	59,895	73,898	53,862
value	28,906	17,310	27,116	16,841
vertical spread	0	0	0	0
vol_termstruct	110,146	74,147	101,884	66,110
16 rows 00:06 sec	248,020	248,020	234,965	0

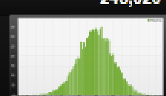

In the above example we are showing the Marginal VAR contribution of each risk bucket using both Parametric and Monte Carlo simulation.

Marginal VAR is defined in Tungsten as follows:

(Total portfolio VAR of all positions) - (Total portfolio VAR without position)

A negative Marginal VAR means the total VAR of the portfolio would increase if the position (risk bucket) was removed. Negative Marginal VAR are said to be risk reducing positions/strategies where the risk reduced is the amount of the marginal VAR. A positive Marginal VAR on the other hand is adding risk to the portfolio by that amount.

In our above example we note that Currency Tactical, Hedge and Residential/Commercial are all risk reducing strategies using the parametric and Monte Carlo models.