Tungsten - Risk Forecasting

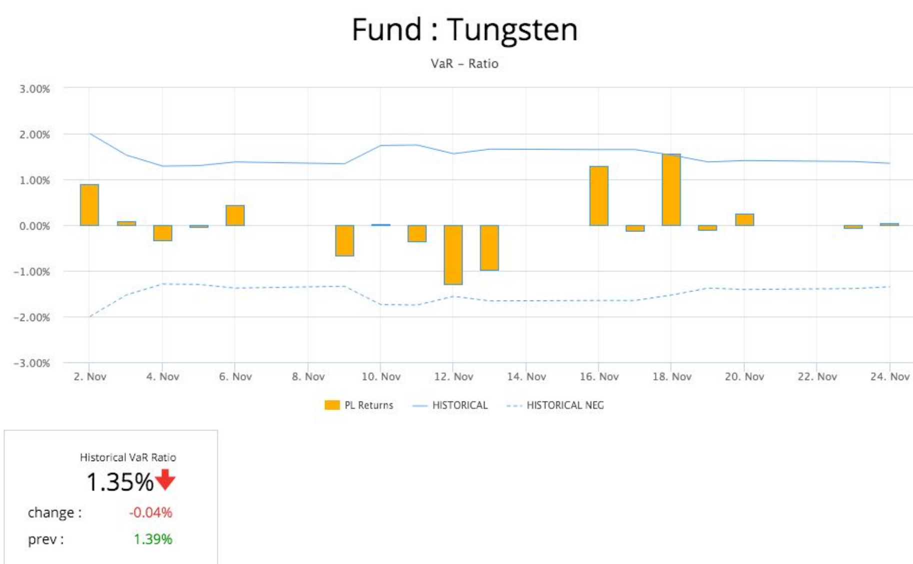

Model Overview

Hamilton, Bermuda

December 2019

Tungsten - Risk Forecasting

This document provides an overview of the Tungsten risk forecasting models. These models are available through the Tungsten GUI, Tungsten API /SaaS and the SS&C/Eze Software Group PMA 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 of 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.

Risk Attribution

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

Incremental VaR

Incremental 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/grouping) and then re-calculating VaR. The difference with and without the bucket is calculated and reported as the Incremental VaR **(Total portfolio VAR of all positions) - (Total portfolio VAR without position).**

With the Incremental VaR we can see what risk buckets (strategies) are adding to the total VAR or reducing (such as portfolio hedges).

Component VaR

Component VaR is similar to Incremental VaR in that it gives us an idea on what positions.risk buckets are risk reducing vs risk adding. The difference with component VaR vs Incremental VaR (other than the way it is calculated) is that the aggregate is additive and equals the total VaR. This allows us to calculate the Component VaR Ratio (Component VaR / Total VaR). The fact that component VaR is additive is one of the main benefits of this calculation method. Note: Component VaR is estimated using a kernel density estimator function which works well on most linear portfolio's. Component VaR on portfolios with a large exposure to positions with optionality are less accurate and we advise using Incremental VaR instead.

Marginal VaR

The Marginal VaR estimates the change in VaR given a small change in position - in this case 1% change. The Marginal VaR gives us an idea on what parts of the portfolio are more sensitive to position changes. A positive change means the position/risk bucket is adding to the overall risk of the portfolio.

Expected Shortfall

Tail risk or expected shortfall is calculated using the asset distribution result of a historical simulation or a Monte 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%.

Expected shortfall (tail risk) is best illustrated with the below graph:

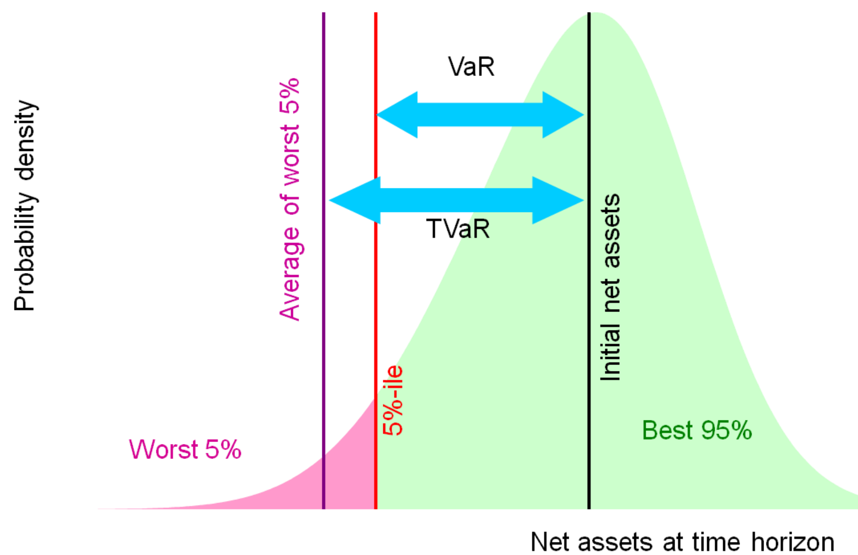

Forecast Beta

Another useful metric available in the Tungsten VaR view is the forecast beta. The forecast beta can be calculated versus a benchmark index of any of the default indexes available in Tungsten. The indexes can also be augmented with time series data imported from PMA.

Forecast beta is estimated simply by taking the covariance of the current portfolio vs the benchmark index and dividing it by the variance of the index using. This is different from the **realized beta** that is available in the P&L View and it can be useful to compare the two. Obviously the result will be highly dependent on how frequently the portfolio is turned over as the realized beta measures the actual P&L time series vs the benchmark index through time, so adjustments to positions will affect the realized beta. The forecast beta assumes a static portfolio without any adjustments to positions.

The forecast beta can also give us an idea on what parts of the portfolio are risk reducing (negative beta) vs risk adding (positive beta).

Consider the below portfolio, as we can see the aforementioned calculations are performed on a portfolio using a risk category attribute as defined in our test portfolio. As we can see the first column is the stand alone Monte Carlo VaR by risk bucket. As an example, the Hedge bucket shows 0.18% stand alone VaR. Stand alone being VaR of positions in just that group / NAV of total portfolio. The next column is showing the Incremental VaR and if we follow along with our Hedge bucket we can see that it is reducing the risk of the total portfolio by -0.13%, or in other words if we were to remove the Hedge bucket, the VaR of the total portfolio would increase by 0.13%.

The component VaR reflects this fact as well - we can see that most of the Incremental VaR are close to the component VaR's as we would expect, there are slight variations here and there as the calculations are different.

RiskSubCatNew	Monte Carlo VaR Ratio	Monte Carlo Incremental Ratio	Monte Carlo Component Ratio over NAV	Monte Carlo Component Ratio	Beta Forecast Ratio
Biotech	0.01%	-0.01%	0.00%	-0.25%	-0.01%
Discretionary	0.50%	0.23%	0.33%	28.00%	3.27%
EV Infrastructure	0.11%	0.06%	0.02%	2.09%	3.88%
FX Intraday	0.00%	0.00%	0.00%	0.02%	0.01%
Hedge	0.18%	-0.13%	-0.14%	-11.65%	-15.34%
Hedge Fund	0.23%	0.14%	0.19%	16.15%	17.97%
Internet Value Add	0.09%	0.06%	0.07%	5.96%	3.43%
Macro	0.21%	0.09%	0.14%	11.72%	7.53%
Momentum	0.34%	0.28%	0.33%	27.57%	24.74%
Renewables	0.17%	0.14%	0.12%	9.78%	7.38%
Risky Long Term	0.03%	0.01%	-0.01%	-0.44%	-0.05%
Sustainable Foods	0.06%	0.00%	0.00%	-0.25%	0.73%
Technology	0.19%	0.15%	0.12%	10.21%	11.83%
TrendSpider	0.24%	0.04%	0.00%	-0.05%	3.44%
Utilities	0.10%	0.01%	0.01%	1.15%	4.23%
Total	1.19%	1.19%	1.19%	100.00%	73.02%

The component ratio is simply the component VaR / total VaR, where the sum of the component VaR ratio's equals 100%. Lastly we can see the forecast beta is showing -15.34% for the Hedge bucket, which means that the Hedge positions are negatively correlated to the S&P 500 index (the index we picked as our benchmark). Here the total portfolio forecast beta is showing 73.02%, which simply means that the portfolio is slightly less risky than the benchmark, or in other words if the S&P 500 would increase by 1%, our portfolio is expected to increase by 0.73%.

The Tungsten VaR Models

At the time of writing, Tungsten utilizes 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 uses the covariances and 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 covariance matrix of the asset returns.

The α 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 and for simpler portfolio's this method works great.

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.

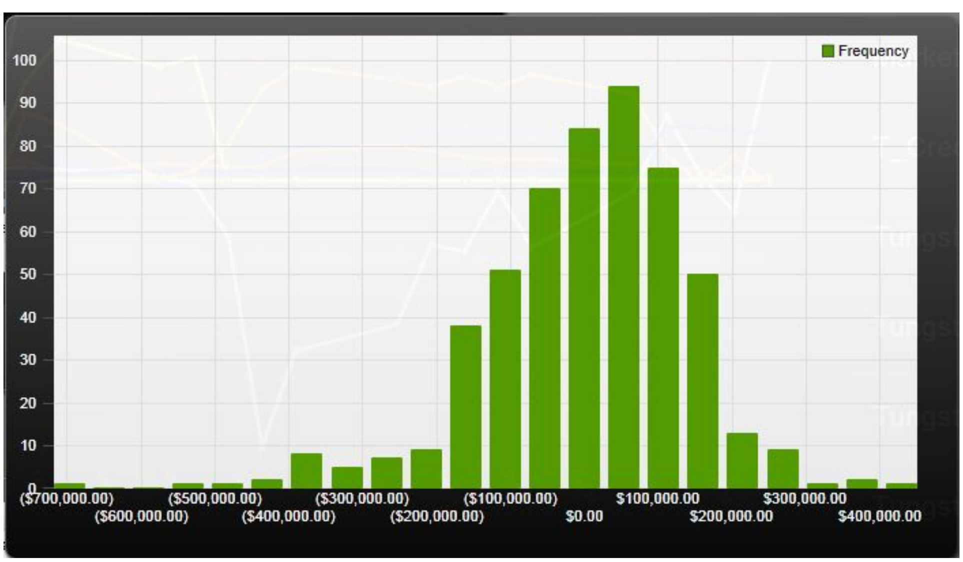

Hybrid model (BRW)

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 nonlinear portfolios.

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 choose 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) method. 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.

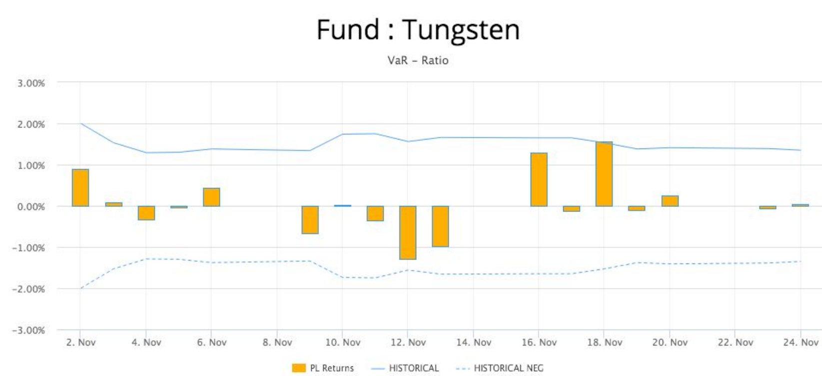

The analytical model gives us quick results, and we can efficiently generate a P&L time series overlaid with the daily risk forecast as pictured above. This gives us an idea on how well the model is predicting our risk. In this case, we are looking at a year to date daily risk forecast using a 99% confidence.

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

However if your portfolio contains assets that exhibit fat tails where the analytical model is poor at capturing these tails, then the simulation models could 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.

This type of VaR back testing is crucial to understanding the dynamics of the various VaR models and is often a requirement to report to financial regulators. VaR back testing can be done in Tungsten by using the date range settings, please note that as the model needs to fully re-calculate the VaR on each trading date in the date range and retrieve the actual P&L for each date this process can take some time depending on the size of your portfolio. **Top tip:** Setup a VaR backtest report by saving your favorite back test as a view and then create a report of that view. This way you can let Tungsten crunch the numbers in the background and deliver the ready report to your email as a PDF and/or HTML report.

Another powerful way of getting to grips with the dynamics of your portfolio is to have all four models display the data side by side, and this is straightforward in Tungsten by simply picking the models you wish to see in one view/report.

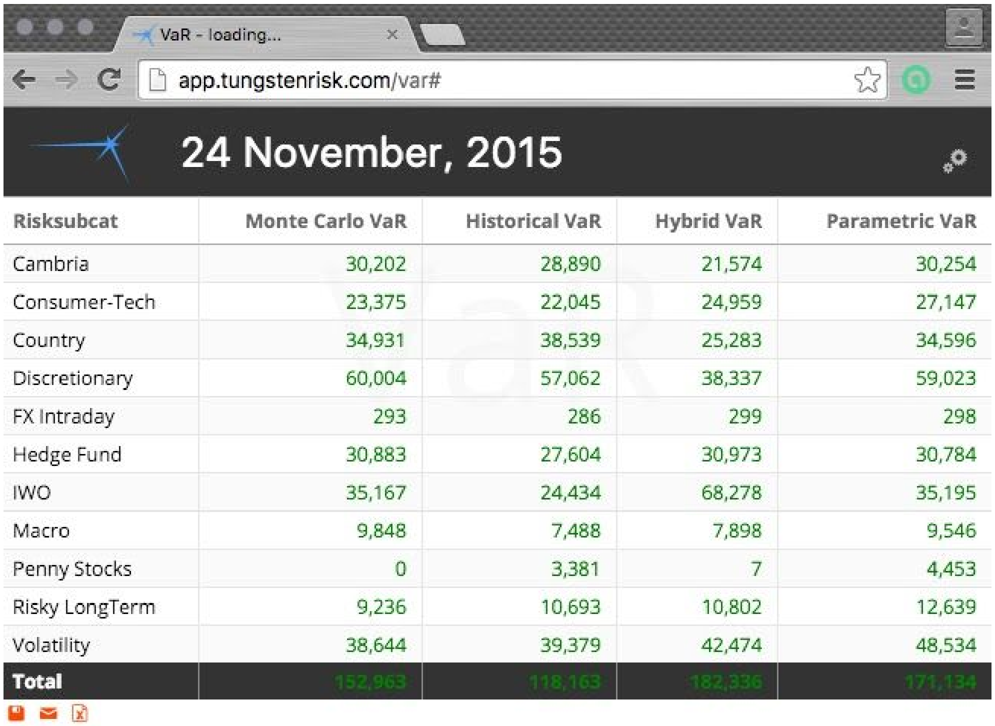

VaR - loading...

app.tungstenrisk.com/var#

24 November, 2015

Risksubcat	Monte Carlo VaR	Historical VaR	Hybrid VaR	Parametric VaR
Cambria	30,202	28,890	21,574	30,254
Consumer-Tech	23,375	22,045	24,959	27,147
Country	34,931	38,539	25,283	34,596
Discretionary	60,004	57,062	38,337	59,023
FX Intraday	293	286	299	298
Hedge Fund	30,883	27,604	30,973	30,784
IWO	35,167	24,434	68,278	35,195
Macro	9,848	7,488	7,898	9,546
Penny Stocks	0	3,381	7	4,453
Risky LongTerm	9,236	10,693	10,802	12,639
Volatility	38,644	39,379	42,474	48,534
Total	152,963	118,163	182,336	171,134

The image above shows all the models in action using the same sampling, confidence and horizon (four year daily returns no decay factor, 95% confidence and one day horizon).

We observe that the Total portfolio VaR is relatively similar using the Parametric versus the Monte Carlo model (152,963 vs 171,134 respectively). The Hybrid model is showing the highest VAR which can be explained by the weighting scheme on the asset returns, and at the time of writing (November 2015) markets have been relatively volatile.

CalcType, Horizon, Sampling

The calcType is one of the most important inputs to the various VaR models. It will tell the VaR models how much data to use, what decay to use (if any), what time period the data is extracted from, e.g. GFC 2008/2009 - The resulting variance and covariance matrix of the risk factors is highly dependent on the amount of data and the time period this data is coming from.

The horizon tells us the time into the future the VaR forecast is estimating - if we use daily sampling (730-Daily for example), the portfolio is not expected to have **daily** loss exceeding the VaR number. If we use 730-Weekly, the VaR number is now estimating the **weekly** loss instead, and finally 730-Monthly estimates a **monthly** horizon. It is also possible to use say a weekly sampling (e.g. 1460-weekly) and convert the weekly number to a daily estimate. The VaR models will then run the VaR calculations as usual and then convert the final result to daily. This can be useful in case you want less volatile market data input such as weekly returns, but you want to be able to back test this VaR estimate using the daily P&L returns. If the user wants to see a longer term horizon such as a year you can use any of the samplings (daily/weekly/monthly), however the horizon needs to be set according to the sampling type. If daily is used, a 250 day horizon needs to be used for a one year horizon. If weekly is used, a 52 week horizon should be used, and finally with monthly sampling a 12 month horizon should be set. The VaR result is then converted to yearly by taking the VAR amount * sqrt(250) when using the daily sampling.

To boost the number of data points it is possible to do sampling overlapping. Example the 730-Monthly = 24 data points (12 months * 2 years) which makes the VaR result difficult to prove statistically significant. This means instead if just taking the month end price changes, we take each day's monthly price change. This will give us a much more significant number of data points to work with. There can of course be issues with auto-correlation using overlapping returns that one has to be aware of. However it can be very challenging to find many years of monthly returns to ensure the results are statistically significant so the auto-correlation issue may be worth it. Again this is up to the risk analyst to decide on the best approach with the data at his disposal.