Mongolia's TOP-20 Index Risk Analysis

Federico M. Massari

January 12, 2017

This work targets investors interested in adding Mongolian equities to their well-diversified portfolios through mimicking the TOP-20 Index, the country's main stock market indicator. Including the index, as opposed to picking single stocks, is a superior choice thanks to the availability of a consistent stream of daily price data, which facilitates risk-based analysis.

Overview

This work targets investors interested in portfolios mimicking Mongolia's TOP-20 Index, the main indicator of the Mongolian stock market.

The index is calculated based on the market capitalisation and average daily trade of the top 20 securities listed on the Mongolian Stock Exchange[a] (MSE). As of December 2016, the index constituents are Apu Company (ticker: APU), Mongolia's oldest enterprise and the country's premier brewer and beverage producer[b]; Baganuur (BAN), one of the big coal mining companies in the country (and the closest to the capital city)[c]; Bayangol Hotel (BNG), the capital's oldest and trustable hotel complex[d]; Ulaanbaatar BUK (BUK), a supplier of concrete products for builders and homeowners[e]; Arig Gal (EER), a manufacturer of wool and cashmere products; Gobi Cashmere (GOV), a leading designer and manufacturer of cashmere, wool, and yak down goods[f]; Genco Tour Bureau (JTB), a leading tour company[g]; Telecom Mongolia (MCH), a provider of internet and telecommunication services[h]; Material Impex (MIE), a supplier of building materials[i]; MIK Holdings (MIK), a purchaser of mortgage loan pools, and provider of affordable housing for Mongolian citizens[j]; Makhimpex Shareholding Company (MMX), a manufacturer of meat and meat products; Mongol Post (MNP), the postal system of Mongolia[k]; Merex (MRX), a concrete plant[l]; Darkhan Nekhii (NEH), a manufacturer and distributor of sheepskin coats, leather goods, and shoes[m]; Remicon (RMC), a concrete producer[n]; Sharin Gol (SHG), a coal mining company[o]; Suu (SUU), Mongolia's dairy products pioneer[p]; Talkh Chikher (TCK), a bakery products maker[q]; Tavan Tolgoi (TTL), one of the largest coal miners in the country[r]; and Ulsiin Ikh Delguur (UID), the State Department Store, Ulaanbaatar's landmark shopping center[s].

Investing in Mongolian equities can be a great way to add diversification benefit to a portfolio. Yet, several listed companies have sparse time series of prices, since they trade very thinly: this lack of data generally implies poor forecasting abilities and suboptimal risk management. On the contrary, the TOP-20 Index has a very consistent and publicly available flow of daily trading data starting from January 3, 2005, and it is therefore suitable for fairly reliable risk-based analyses. Creating a portfolio of TOP-20 constituents, with weights equal to those implied by the index, is currently the only feasible way to invest in it, as it is not possible to purchase the index and, to our knowledge, index derivatives are not currently present on the market. The purchase procedure shoud be relatively inexpensive, since only twenty securities are involved, and these are usually among the most actively traded on the MSE. This work aims to determine whether buying Mongolia's TOP-20 Index is beneficial, in terms of risk and reward involved. It uses econometric techniques in an attempt to forecast the distribution of index returns and to understand the likelihood, and entity of potential future losses [1]. Due to the uncertainty related to investing in emerging markets equities, we advise to include the asset in an already well-diversified portfolio.

Data sample

We use TOP-20 Index daily close prices over nine and a half years, from August 10, 2007 to November 30, 2016* (Figure 1). This time frame is considered to be the most adequate to give a faithful representation of Mongolia's current stock market conditions. In addition, we keep the most recent month of data, December 1-29, for out-of-sample (OOS) testing. We calculate daily log returns, the conventional choice for univariate time series analysis, from close prices (Figure 2), and provide summary statistics (Table 1). On average, daily returns stay in the ±5.00% band, with most of them, in absolute value, close to zero. Inspection of their distribution (Figure 3) provides a clearer picture of this fact. The sum of the two tallest histogram bars, representing the frequency of returns between ±0.625%, is above 50%, and more than 95% of the observations lie in ±2.50%. In addition, as it is usually the case for daily data,

Figure 1: TOP-20 Daily Close Prices

Figure 2: TOP-20 Daily Log Returns

Table 1: TOP-20 sample statistics	
Beginning	10/08/07
Ending	30/11/16
Sample Size	2320
Frequency	*Daily*
Mean	0.0050%
Variance	0.0174%
Standard Deviation	1.3195%
Skewness	0.2637
Excess Kurtosis	11.7881

Figure 3: Sample Distribution of Returns

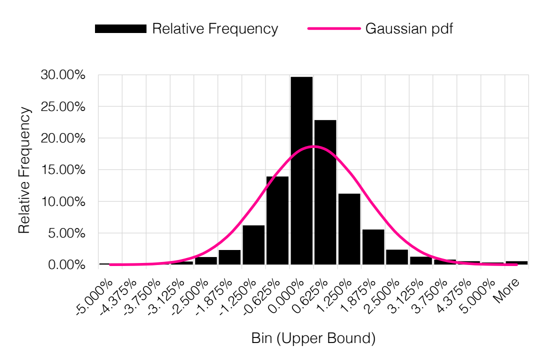

the mean sample return is indistinguishable from zero (0.0050%).

Evidence of non-normality

Having independent and identically distributed (i.i.d.) Normal observations is desirable in a forecasting exercise. The Normal distribution is completely specified by its first two moments, so if the data at hand come from such a density, precise estimation of the population mean and variance is ensured by increasing sample size.

However, TOP-20 Index evidences non-normality features, as signalled by (in contrast to the Gaussian distribution, for which both parameters are equal to zero): positive skewness (there are more large positive than large negative returns in the sample) and positive excess kurtosis (the distribution has fat tails, high peak, and thin shoulders). Both features are apparent in the histogram above, as well as in the plot of the empirical against the Gaussian cumulative distribution functions (Figure 4).

Another visual tool to detect non-normality is the autocorrelation function (ACF). Autocorrelation (AC) measures the linear dependence between today's value of a time series variable and the past value of the same. By definition, i.i.d. Normal returns, as well as their transformations (such as the squares) have no autocorrelation at all lags. Autocorrelation hinders the ability to forecast. We test AC for both TOP-20 Index log returns and squared returns. Generally, at daily intervals, log returns display close to zero AC at all lags, meaning no predictability of tomorrow's value based on today's and previous days' ones. Yet, at lag orders up to 17, TOP-20 Index exhibits statistically significant autocorrelation (Figure 5). Statistical significance occurs if the value of the AC coefficient lies outside Bartlett's bounds,

Figure 4: TOP-20 Index Cumulative Distribution Function Against the Gaussian

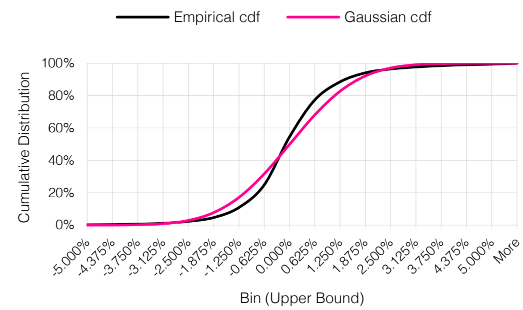

Figure 5: Autocorrelation of TOP-20 Index Log Returns

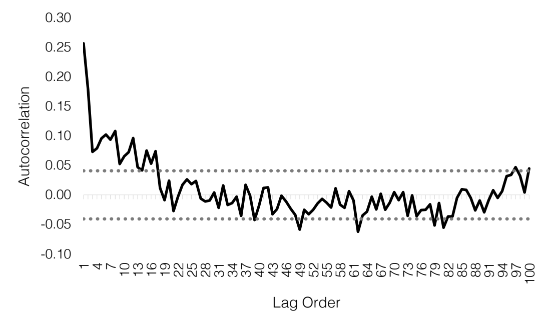

Figure 6: Autocorrelation of TOP-20 Index Squared Returns

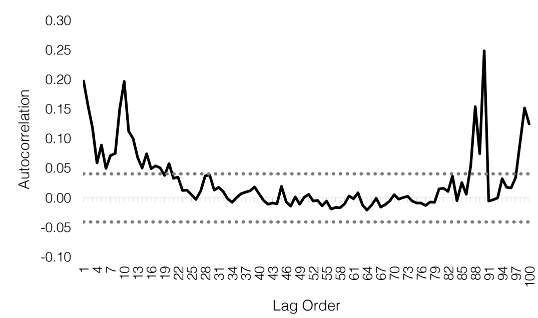

provided at the 5% significance level (2.5% probability per tail) and, for sample size 2320, equal to ±4.07%. One possible explanation of the anomaly could be the low liquidity of the index constituents. Squared returns, instead, should naturally show positive and statistically significant AC at several lags. Squared returns are a valid proxy for conditional variance, which empirically clusters: if today's variance is high (low), it is likely that tomorrow's variance will be high (low) as well. Although it is an additional symptom of the non-Gaussian nature of the data, positive AC of squared returns is by itself beneficial, as volatility plays an essential role in the forecasting process since it helps making returns i.i.d. Normal. As expected, the AC of TOP-20 Index squared returns is strongly persistent, spiking at 20% at lag 10, and gradually fading only after 20 lags (Figure 6). The AC in both log and squared returns is to be reduced using variance forecasting techniques.

Variance forecasting techniques

To make TOP-20 Index returns close to i.i.d. Normal, we divide them by conditional standard deviations able to capture the non-Gaussian features of the data. This way, via the usual fraction simplification, such features are as much as possible eliminated. The procedure is equivalent to data standardisation, since for daily returns the standard deviation (1.3195%) completely dominates the mean (0.0050%), so it is safe to assume it is equal to zero. The result of the standardisation process is a time series of shocks (returns divided by standard deviation), which are more or less i.i.d. Normal depending on the quality of the volatility proxy, and whose distribution should more closely resemble the Gaussian than the density of shocks standardised by unconditional volatility (homoscedastic) (Table 2), mostly in terms of higher moments, skewness and excess kurtosis.

There exist several conditional volatility models. We focus on two, RiskMetrics and GARCH(1,1), as well as on a less computational-intensive variation of the latter, Variance Targeting (VT).

Table 2: Homoscedastic shocks statistics	
Mean	0.0038
Variance	1.0000
Standard Deviation	1.0000
Skewness	0.2637
Excess Kurtosis	11.7881

Table 3: RiskMetrics parameter estimation	
λ	0.9172
Maximised log-likelihood	9101.89
Intercept (S.E.)	5.33E-05 (1.63E-05)
Slope (S.E.)	0.6914 (0.0564)
R-squared	6.09%

Figure 7: RiskMetrics Fit of TOP-20 Squared Returns

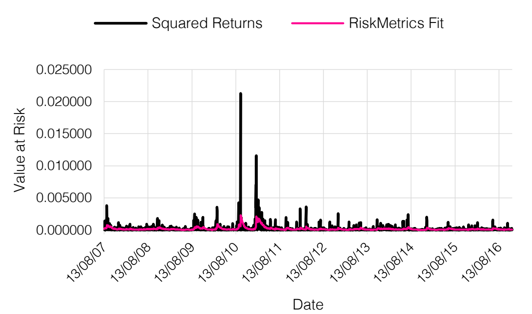

RiskMetrics

This conditional volatility model is completely specified by parameter λ. For daily data, λ is usually set to 0.94, although we here estimate it via maximum likelihood (MLE). MLE maximises the probability that the data sample is generated by the model used. In other words, it optimises the agreement of the selected model with the observed data. $\lambda = 0.9172$, with a goodness of fit, obtained by regressing the time series of conditional variances (X) on squared returns (Y), of 6.09% (Table 3). R-squared is usually poor (5-10%), so the model is not necessarily bad. In fact, RiskMetrics fit of the time series of squared returns turns out to be quite satisfactory (Figure 7). The model successfully captures some non-normality features of the data: both skewness and excess kurtosis for

RiskMetrics shocks are reduced, down respectively 37% and 67%, compared to the same parameters for homoscedastic (constant variance) shocks (Table 4).

GARCH(1,1)

This widely popular model is completely specified by the three parameters ω, α, and β, whose values are necessarily estimated via MLE (Table 5). ω = 1.13E-05, α = 0.1579, β = 0.7766, with a goodness of fit of 6.65%, almost 10% higher than with RiskMetrics. The likelihood ratio (LR) test (183.85) is positive and significant, showing that choosing GARCH over RiskMetrics leads to a better fit of the squared returns (Figure 8). The improvement is particularly visible in the peaks occurring on September 23, 2010 and February 1, 2011, following abnormal squared returns. GARCH also bests RiskMetrics in terms of higher moments for the shocks (Table 6). Skewness is reduced by a staggering 96% (homoscedastic) or 60% (RiskMetrics), excess kurtosis by 70% (homoscedastic) or 2% (RiskMetrics). In particular, variance and skewness values are very close to Normality.

GARCH(1,1)-VT

Variance Targeting (VT), a slight variation of the GARCH model, conveniently sets ω equal to the unconditional (long-run) variance, thus reducing the number of parameters to estimate by one. Apart from computational reasons, VT is useful whenever the sample variance (a proxy for the unconditional one) estimated from the data is thought to be representative of current stock market conditions. This occurs since GARCH itself usually leads to small deviations from this value in the long-run volatility forecast (Table 7). GARCH-VT calibrates the parameters (Table 8) in order to obtain a predicted level of variance in line with current sample variance (Table 9). GARCH-VT does not seem to add much to the fit, with very little variation in the value of ω, α, and β, and almost no improvement in terms of R-squared (+0.05%). Moreover, the LR test is negative, although not statistically signifi-

Table 4: RiskMetrics shocks statistics	
Mean	–0.0220
Variance	1.1893
Standard Deviation	1.0905
Skewness	0.1672
Excess Kurtosis	3.8669

Table 5: GARCH(1,1) parameter estimation	
ω	1.13E-05
α	0.1579
β	0.7766
Persistence ($\alpha + \beta$)	0.9345
Maximised log-likelihood	9193.81
Intercept (S.E.)	4.35E-05 (1.65E-05)
Slope (S.E.)	0.7525 (0.0586)
R-squared	6.65%
Likelihood Ratio test	183.85

Figure 8: GARCH(1,1) Fit of TOP-20 Squared Returns

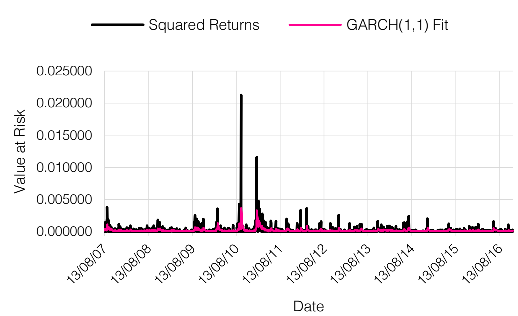

Table 6: GARCH(1,1) shocks statistics	
Mean	–0.0279
Variance	1.0016
Standard Deviation	1.0008
Skewness	0.0094
Excess Kurtosis	3.6535

Table 7: GARCH(1,1) long-run variance forecast	
Unconditional variance	0.000174
Predicted variance	0.000172
Error	–2.36E-06

cant. VT shock statistics are also very close to those of GARCH (Table 10), with slightly worse odd moments (mean and skewness) and slightly better even moments (variance, standard deviation, and excess kurtosis). Nevertheless, we consider the estimated sample variance to be adequate to explain current Mongolian stock market conditions: in a 21-day out-of-sample experiment, between December 1 and December 29, 2016, conditional variance is found to oscillate about the long-run, unconditional value, so it is safe to assume it will eventually converge to this figure (Figure 9).

Effectiveness of variance forecasting

To test the effectiveness of the variance forecasting techniques, we rely on two visual tools. The first one is the QQ-plot, which maps out the input sample quantiles of shocks against those of a standard normal distribution; the second one is the autocorrelation function (ACF) of squared shocks.
In a QQ-plot, data conforming to the Normal density lie on a straight, 45-degree line. Homoscedastic shocks (log returns divided by sample variance) exhibit visible departure from Normality, especially in the tails (Figure 10). In risk management, particularly relevant are deviations occurring in the left tail, because they represent failure of the model to account for extreme negative events, usually linked to severe monetary losses. Division by constant volatility does not adequately reduce the magnitude of the largest negative shock, –11.0539, corresponding to a single-day return of –14.59% on September 22, 2010. RiskMetrics offers an improvement, which is nevertheless only partial. Although the shock on the same day becomes

Table 8: GARCH-VT parameter estimation	
ω	1.11E-05
α	0.1581
β	0.7779
Persistence (α + β)	0.9360
Maximised log-likelihood	9193.81
Intercept (S.E.)	4.36E-05 (1.65E-05)
Slope (S.E.)	0.7488 (0.0583)
R-squared	6.65%
Likelihood Ratio test	–0.0097

Table 9: GARCH-VT long-run variance forecast	
Unconditional variance	0.000174
Predicted variance	0.000174
Error	0.00E+00

Table 10: GARCH-VT shocks statistics	
Mean	–0.0278
Variance	0.9988
Standard Deviation	0.9994
Skewness	0.0115
Excess Kurtosis	3.6379

Figure 9: GARCH-VT Variance Forecast

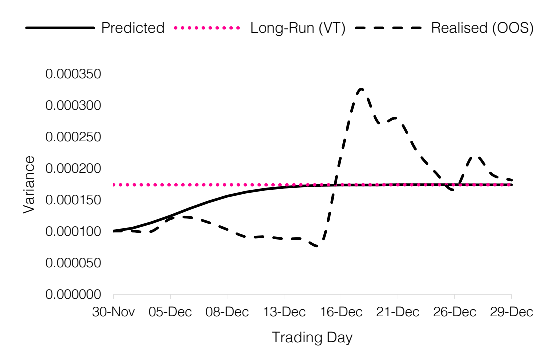

Figure 10: QQ-plot with Homoscedastic Shocks

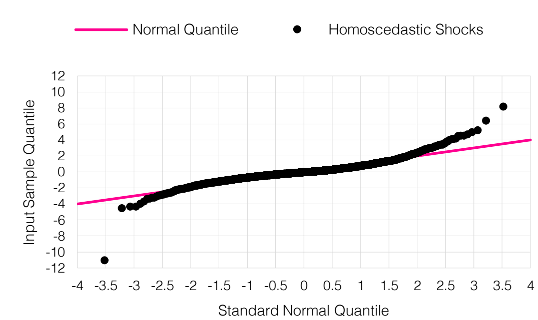

now –6.3992, several other extreme values are farther from the 45-degree line than they were before (Figure 11), because variance is now 20% higher (1.1893, compared to 1.000). In addition, the right tail is still clearly misspecified. GARCH-VT provides the best fit, globally (Figure 12). Both tails are well-modelled, with extreme shocks much closer to the straight line than they were before. Of particular interest is the analysis of the 100 worst outcomes, which shows that, apart from the most extreme value, whose magnitude (–7.6606) is slightly higher than with the RiskMetrics fit (–6.3992), all shocks are fairly close to the 45-degree line (Figure 13).

Inspection of the ACF of squared shocks leads to similar conclusions. While both conditional volatility models are effective at reducing predictability, RiskMetrics represents a worse fit, with statistically significant AC at lag 1, and with past variables still able to explain about 10% of the values of those the following day (Figure 14). GARCH-VT, is a much better fit, with only mildly significant AC at lag 1: previous day variables are only able to explain about 4.36% of the values of those the following day, a statistic very close to that for Bartlett's upper bound (4.07%) (Figure 15).

Overall, the GARCH-VT model remarkably captures most of the non-normality embedded in the sample minimising, in particular, skewness. We keep GARCH-VT shocks and look for a superior distribution fit, able to account for asymmetry and, most importantly, fat tails.

Distribution fitting

By treating data as Gaussian, it is unlikely to obtain a better fit even by further refining the conditional volatility model. The Normal den-

Figure 11: QQ-plot with RiskMetrics Shocks

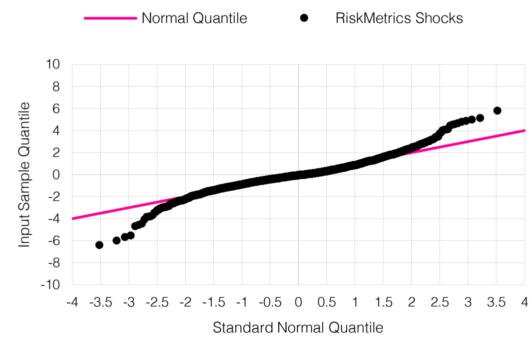

Figure 12: QQ-plot with GARCH-VT Shocks

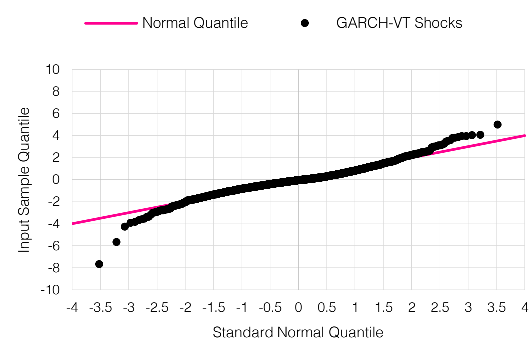

Figure 13: Left Tail Goodness of Fit

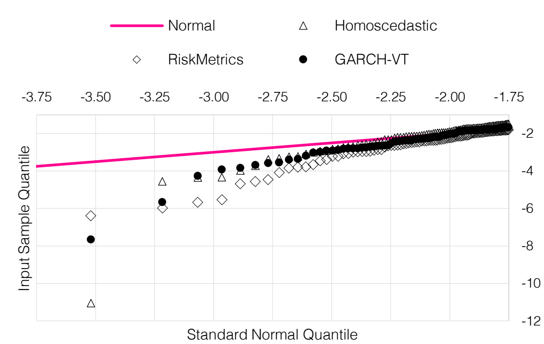

sity does not account for higher order moments, which are still positive in the distribution of GARCH-VT shocks: skewness = 0.0115, excess kurtosis = 3.6379. We first attempt to reduce the latter, due to its magnitude and importance in risk management, using the symmetric, standardised t distribution. We then consider skewness, by introducing an asymmetric version of t, obtained by merging two densities with different degrees of freedom, and we test whether factoring this parameter in the model brings any improvement to the fit. Finally, we further refine the left tail, using Extreme Value Theory (EVT).

GARCH-VT standardised t distribution

The Student's t is able to capture the fatter tails and the thinner shoulders in the distribution of data. We use the standardised version of t, because we are interested in the distribution of shocks, with unit variance. The t is fully specified by one parameter, the number of degrees of freedom, which dictates the shape of the density. The smaller this number, the fatter the tails and the less pronounced the peak of the resulting distribution. Variance, the second central moment, is well-defined only if the number of degrees of freedom, $d > 2$. Kurtosis, the fourth central moment, exists only if $d > 4$. We estimate the optimal value for d via quasi-maximum likelihood (QMLE), a two-step process: first, find the values of the conditional volatilities; then, calculate the value of the degrees of freedom, maximising the log-likelihood function (MLE). For step one, we rely on the GARCH-VT model. For step two, we input 5.65, an initial approximation of d given by 6/(excess kurtosis) + 4, and we calibrate it to the data set to obtain $d = 4.13$ (Table 11).
The standardised t offers a much nicer fit to the distribution of GARCH-VT shocks than the Normal density does. The left tail, of particular interest to risk management, is nicely modelled, with only mild deviations for few extreme observations (Figure 16). The right tail, however, is still partly misspecified, a consequence of the symmetry of the distribution, which is at odds with the positive skewness of the sample. Although we are not much interested in perfectly modelling the right tail, we now test whether including skewness in the model brings any significant

Figure 14: Autocorrelation of RiskMetrics Squared Shocks

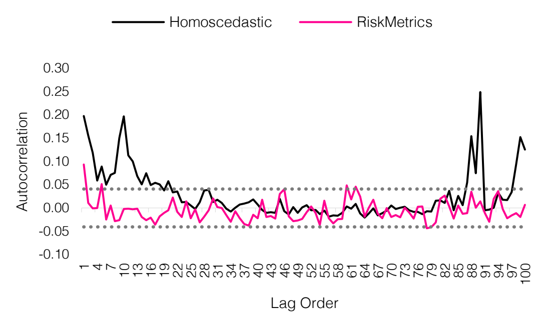

Figure 15: Autocorrelation of GARCH-VT Squared Shocks

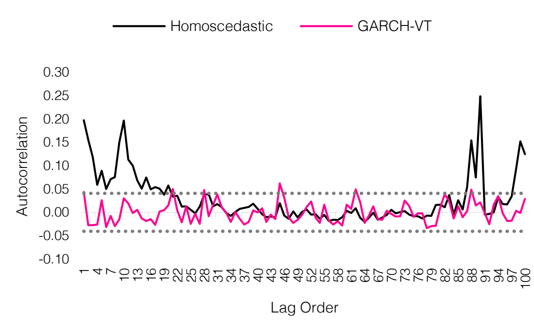

Table 11: Standardised t parameter estimation	
Degrees of freedom (d)	4.13
Maximised log-likelihood	–3159.68
d approximation	5.65

Figure 16: QQ-plot of GARCH-VT Shocks against the Standardised t Distribution

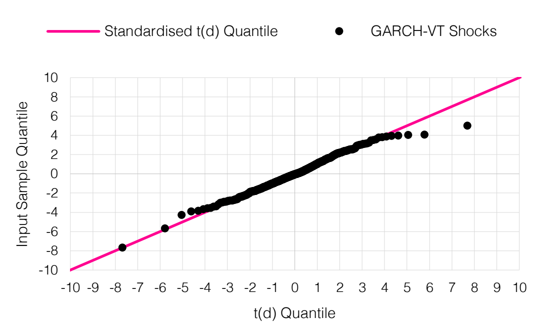

improvement to the fit.

GARCH-VT asymmetric t distribution

The asymmetric Student's t is a generalised version of the standardised t introduced above. It is obtained by merging two distributions at a point $-A/B$ on the horizontal axis. The density is completely specified by two degrees of freedom, $d_1 > 2$ and $-1 < d_2 < 1$, dictating the shape of the distribution. The standardised t is a special case, with $d_1 = 4.13$ and $d_2 = 0$, so that $A = 0$ and $B = 1$ and the distribution is symmetric and centered at the origin. We find optimal values for d_1 and d_2 via QMLE, using the available GARCH-VT shocks as input. Parameter calibration generates $d_1 = 4.09$ and $d_2 = 0.09$, so that $-A/B = -0.1322$ (Table 12). A positive value for d_2 indicates right skewness, as expected from the data (Figure 17). Including an additional parameter brings some enhancement to the fit, as shown by the statistically significant likelihood ratio test, 18.17 (99% confidence level, Chi-squared critical value: 6.63). The improvement is mostly visible in the shoulders and in the parts of the tails closer to the center of the distribution, with the effect of bringing shocks in those regions nearer to the 45-degree line (Figure 18). However, accounting for asymmetry in the model comes at the expense of the fit in the farthest regions of the tails: extreme shocks, especially negative ones, are now slightly offset, more distant from the straight line than they were before. The reason is that incorporating right skewness in the model makes the left tail thinner, understating the probability of negative extremes and overstating that of positive ones. Because precise left tail specification is crucial for risk control, we now turn to an approach that only makes use of negative extreme data.

Table 12: Asymmetric t parameter estimation

d_1	4.09
d_2	0.09
C	0.5252
A	0.1328
B	1.0043
$-A/B$	-0.1322
Maximised log-likelihood	-3150.60
Likelihood ratio test	18.17

Figure 17: Density Comparison between Standardised and Asymmetric Student's t

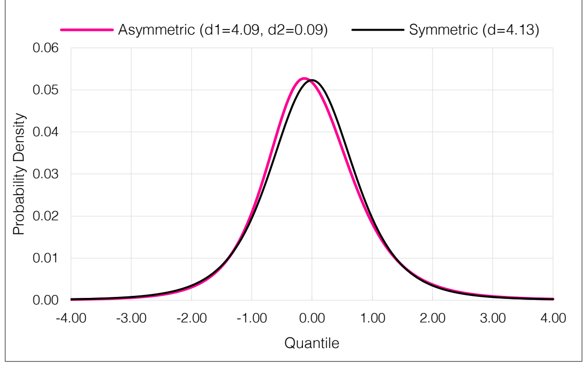

Figure 18: Comparison between Asymmetric and Standardised t Distribution Fit

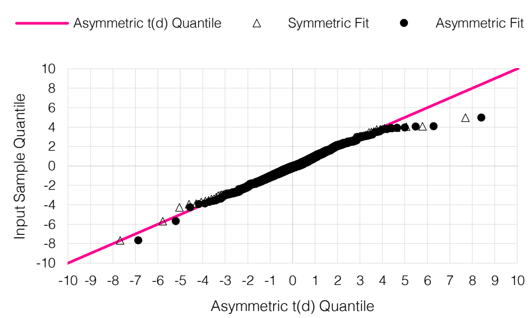

Extreme value theory

Extreme value theory (EVT) asserts that the extreme tail of a broad range of distributions can be approximated by the Generalised Pareto Distribution (GPD). The tail of the GPD has shape dictated by parameter ξ, with $\xi > 0$ (power tails) in presence of positive excess kurtosis. The two cases $\xi = 0$ (Gaussian density, exponential tails) and $\xi <$

0 (platykurtic densities) are usually ignored in risk management. For ξ strictly positive, we can estimate the GPD distribution in closed form, using Hill's estimator. The latter is completely specified by u, the threshold beyond which lie the extremes included in the analysis, T_u, the total number of observations larger than the threshold, and c, a constant. We resort to a rule of thumb and set T_u = 50, to obtain optimal value for ξ = 0.2796 (Table 13).

The EVT fit appears to be slightly superior to that of the Standardised t distribution (Figure 19). In particular, the third, fourth and fifth most extreme shocks are now quite in line with their predicted values. Overall, EVT provides a satisfactory approximation to the tail of the empirical density.

We now introduce a widely used portfolio risk measure, Value at Risk.

Value at Risk

Value at Risk (VaR) is a threshold such that the probability that a given loss will be greater than this threshold in a predefined time window is equal to the selected confidence coefficient. We compute VaR at the 99% confidence level for the most significant models discussed above, plus two: historical simulation (HS), a popular, albeit very poor model, and filtered historical simulation (FHS). Moreover, we provide an approximation to VaR that accounts for the estimated values of skewness and excess kurtosis, known as the Cornish-Fisher (CF) quantile.

We start by analysing RiskMetrics (Figure 20) and GARCH-VT (Figure 21) VaR, under the Normal density assumption. The former does not seem to adequately react to negative innovations, and tends to underestimate current market conditions in the computation

Table 13: Hill's estimator	
u	–2.1024
T_u	50
ξ	0.2796
c	0.3074

Figure 19: QQ-plot of GARCH-VT Tail Shocks against the EVT Distribution

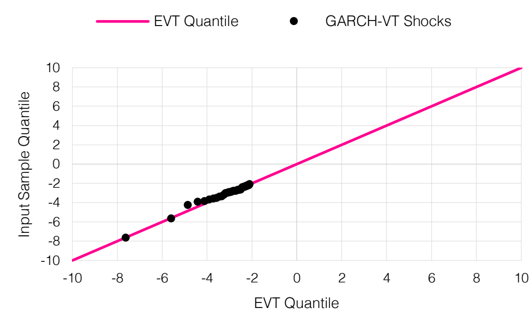

Figure 20: RiskMetrics Gaussian VaR

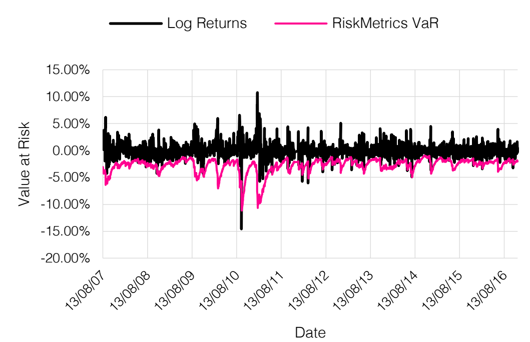

Figure 21: GARCH-VT Gaussian VaR

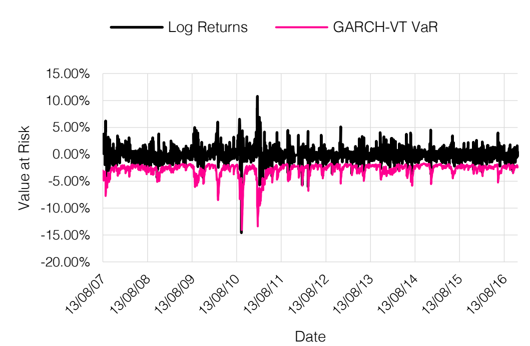

of VaR for the subsequent trading days. Predictions for September 23, 2010, following a return of –14.59% the day before, was only

80% of this figure (Table 14). The latter is more adaptive, and better describes the volatility clustering phenomenon, although it still slightly misrepresents VaR when market crashes occur. On the same day, model predicted VaR was only 3.6% lower.

We turn to the GARCH-VT model under the assumption of a Standardised t distribution (Figure 22). VaR appears to be extremely reactive in all situations, and fully covers potential future losses following negative events, although it is probably too conservative at times, forcing to set apart a large amount of assets after a loss occurs. For the day following the crash, predicted VaR was overestimated by 11.5%.

GARCH-VT EVT (Figure 23) produces similar results, as it was already apparent from the analysis of the QQ-plots. It fully incorporates losses, and it is still expensive, although less than before. Following the crash, predicted VaR was overvalued by only 8%.

The models considered up to this point are parametric in nature, because they make inference on the shape of the probability density of sample data, whether Gaussian, Student's t, or GPD. We now turn to two widely used models, historical simulation (non-parametric) and filtered historical simulation (semi-parametric). HS lets the data convey information of the distribution of returns: an arbitrary number of past returns (here, the most recent 250) is sorted in ascending order, and VaR is a percentile defined by a cut-off point. HS is easy to compute but misleading, as it is unresponsive to market crashes and frequently gives rise to box-shaped patterns (Figure 24). Forecast HS VaR for September 23, 2010 missed the target by almost 80%, as the model failed to update to new market conditions.

Table 14: Sep 23, 2010 predicted VaR, % loss	
September 22, 2010 return	−14.59
Gaussian, RiskMetrics	11.17
Gaussian, GARCH-VT	14.07
Standardised t, GARCH-VT	16.27
EVT, GARCH-VT	15.76
Historical Simulation	3.39
FHS, GARCH-VT	16.56
Cornish-Fisher quantile	19.16

Figure 22: GARCH-VT Standardised t VaR

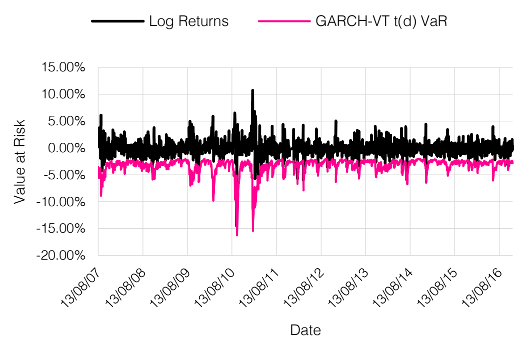

Figure 23: GARCH-VT EVT VaR

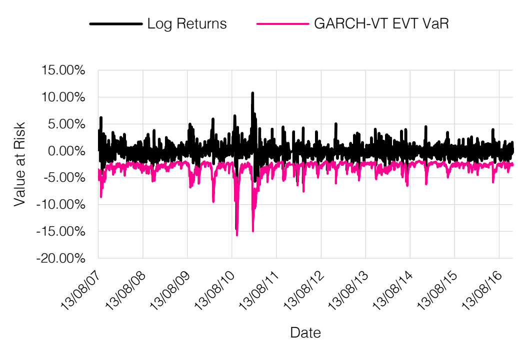

Figure 24: Historical Simulation VaR

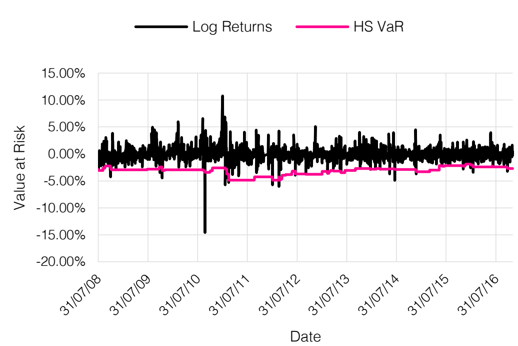

FHS combines the best of both parametric and non-parametric worlds: it relies on a conditional volatility model to calculate shocks, but it avoids making inference on their distribution, so it uses HS to find the percentile corresponding to VaR. FHS with GARCH-VT shocks (Figure 25) is a fairly accurate, albeit expensive model: it promptly reacts to losses, but it requires to set apart a large amount of assets. After the crash, VaR was overestimated by 13.5%.

We finally provide the Cornish-Fisher quantile, a Taylor expansion around the Gaussian density that accounts for skewness and excess kurtosis and works as an approximation of VaR. We input moment statistics for GARCH-VT shocks to obtain a CF quantile equal to –3.1683 (Table 15).

We compare the performances of the main VaR models for the most recent year (Figure 26). VaR is here defined in terms of maximum loss, hence it is positive. On one hand, Gaussian GARCH-VT gives the lowest predicted figure for VaR, often inadequate given current market conditions. On the other, CF GARCH-VT generates the highest value, also inadequate because too expensive. Both models are not recommended. Standardised t and EVT lead to similar and well-balanced results, and these models we favour.

We now report cumulative maximum potential daily profit or loss on TOP-20 Index for the whole data period (August 10, 2007 to November 30, 2016) given a trader's daily VaR limit of MNT 1,000,000 (EUR 384.89) (Figure 27). The Gaussian model imposes the lowest risk-taking limits on the trader, maximising both the outstanding position and the profit or loss for the day. On the contrary, the CF model is very conservative, as it sets the highest risk-taking burden, minimising both potential risks and rewards. Standardised t and EVT are in-between. Overall, an individual or firm investing in the

Figure 25: Filtered Historical Simulation VaR

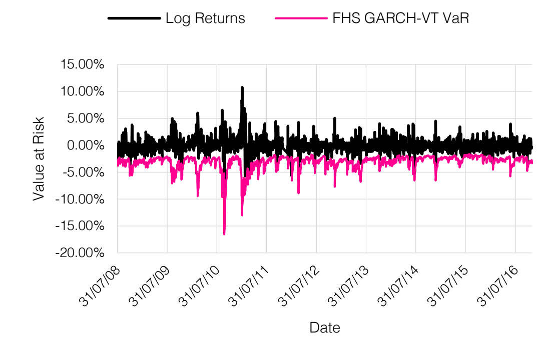

Table 15: Cornish-Fisher quantile statistics	
Φ^{-1}_p	–2.3263
Skewness (ζ_1)	0.0115
Excess kurtosis (ζ_2)	3.6379
CF quantile (CF^{-1}_p)	–3.1683

Figure 26: GARCH-VT VaR Model Comparison

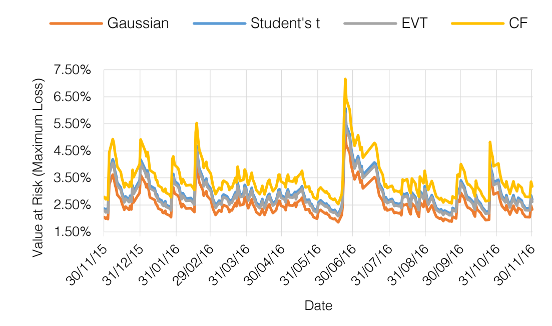

Figure 27: Cumulative Daily Maximum P&L

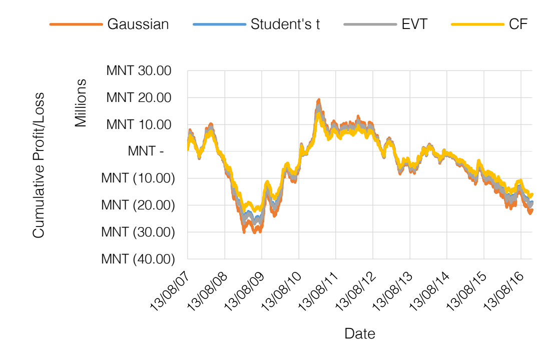

Mongolian stock market during the analysed period would have been better off adopting a conservative model, such as CF, during the turbulent period from mid 2008 to mid 2010 and an in-between one before and after. In the near future, we suggest to keep going on with one of the following: standardised t, EVT or FHS, and with GARCH-VT or an analogous model for conditional volatility. We advise to switch back to conservative if market conditions deteriorate.

Backtesting Value at Risk

Backtesting allows to check whether the model used to compute VaR is adequate, given the data. The procedure counts the number of exceptions (days in which the return was lower than VaR), checks if this number is consistent with the theoretical value (the significance level) and, if the actual number of exceptions is higher than the predicted one, uses a Chi-squared test to decide if the model is still adequate given the data.

We backtest three VaR models, Gaussian, Standardised t, and EVT, with GARCH-VT conditional volatility, at 1% significance level. We perform three tests: unconditional coverage, independence, and conditional coverage. The first, unconditional coverage (UC), checks if the fractions of violations coming from a particular VaR model (π) is seriously different from the theoretical fraction (p). The second, independence (I), tests if the violations are clustered, i.e., the likelihood that a violation will occur tomorrow, given that today it did as well. If risk managers can predict an exception based on today's outcome, the VaR model is not satisfactory and should be rejected. The clustering hypothesis is tested by counting the number of 00 (no violation, followed by no violation), 01 (no violation-violation) and 10 (violation-no violation) items in the sample.

Table 16: Gaussian GARCH-VT VaR Backtesting	
Hit = 0	2283
Hit = 1	37
00 sequence	2246
01 sequence	36
10 sequence	36
11 sequence	1
Unconditional Coverage	
π	1.59%
L(π)	3.63E-83
p	1.00%
L(p)	1.08E-84
Unconditional Coverage	7.0237
Reject VaR model?	REJECT
Independence	
π_{01}	1.58%
π_{11}	2.70%
L(π_1)	4.17E-83
Independence	0.2763
Reject VaR model?	DO NOT REJECT
Conditional Coverage	
Conditional Coverage	7.3000
Reject VaR model?	REJECT

Table 17: Std. $t_{4.13}$ GARCH-VT VaR Backtesting	
Hit = 0	2296
Hit = 1	24
00 sequence	2271
01 sequence	24
10 sequence	24
11 sequence	0
Unconditional Coverage	
π	1.03%
L(π)	9.65E-59
p	1.00%
L(p)	9.51E-59
Unconditional Coverage	0.0276
Reject VaR model?	DO NOT REJECT
Independence	
π_{01}	1.05%
π_{11}	0.00%
L(π_1)	1.25E-58
Independence	0.5228
Reject VaR model?	DO NOT REJECT
Conditional Coverage	
Conditional Coverage	0.5503
Reject VaR model?	DO NOT REJECT

The last, conditional coverage (CC), jointly tests for the above: if the number of exceptions is consistent with predictions, and if the exceptions are independent. All tests are distributed according to a Chi-squared with one degree of freedom. We set significance level to 5% (3.84). If at least one statistic is above the critical value we reject the VaR model; else, we do not.

As expected, the Gaussian VaR model performs poorly (Table 16). The number of violations is statistically significant (1.59%), although they do not cluster. Overall, we reject the Gaussian VaR model as too inaccurate. The standardised t and EVT VaR models, instead, perform similarly and nicely. The former is more precise, with a smaller number of exceptions in the period (1.03%) almost in line with the predicted value (Table 17), even though the latter is still completely satisfactory (Table 18).

We also include graphs of VaR exceedances. Gaussian VaR (Figure 28) often underestimates the amount of daily resources to set apart to cover potential losses. The result is a series of short, but frequent (37) bars in the graph. Standardised t VaR (Figure 29) and EVT VaR (Figure 30) promptly react to current market conditions, generally predicting the right amount of capital to keep and giving rise to a much lower number of violations (respectively, 24 and 28).

Out-of-sample forecast

We now incorporate the GARCH-VT conditional volatility model in simulation techniques to predict the level of VaR for the out-of-sample trading month December 1-29, 2016. We run FHS, and introduce Monte Carlo, a computational algorithm that uses repeated random sampling to generate

Table 18: EVT GARCH-VT VaR Backtesting	
Hit = 0	2292
Hit = 1	28
00 sequence	2263
01 sequence	28
10 sequence	28
11 sequence	0
Unconditional Coverage	
π	1.21%
$L(\pi)$	1.59E-66
p	1.00%
L(p)	9.91E-67
Unconditional Coverage	0.9410
Reject VaR model?	DO NOT REJECT
Independence	
π_{01}	1.22%
π_{11}	0.00%
$L(\pi_1)$	2.26E-66
Independence	0.7087
Reject VaR model?	DO NOT REJECT
Conditional Coverage	
Conditional Coverage	1.6497
Reject VaR model?	DO NOT REJECT

Figure 28: Gaussian GARCH-VT VaR Exceedances

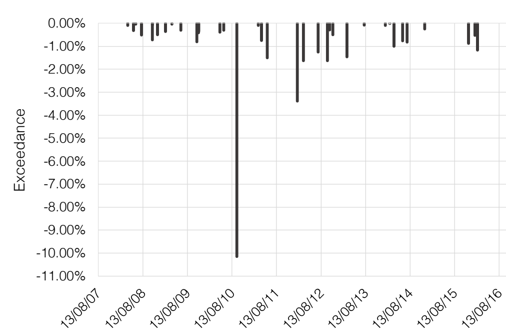

Figure 29: Standardised t GARCH-VT VaR Exceedances

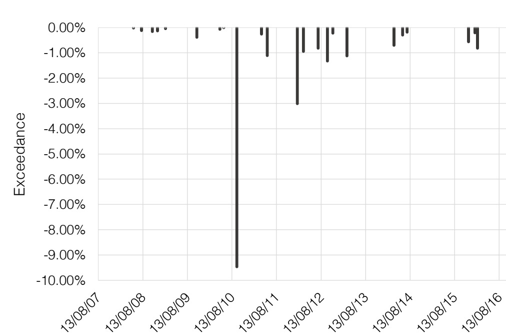

fresh scenarios. MC requires that a distribution be specified in advance, so we opt for two: the Gaussian and the standardised t. We run 10,000 simulations for each trading day. FHS assumes that the historical distribution of TOP-20 Index returns is a valid proxy for the future, and randomly extracts, with replacement, observations from the data available, using an index based on discrete random variables generated beforehand. The steps are the following:

1. We create 21 series of 10,000 discrete RV in the closed interval [1, sample size]. Each series represents a day, each RV the index of a particular scenario (one of the TOP-20 daily returns in the sample).
2. From the sample of historical GARCH-VT shocks, we extract the ones whose indices correspond to the discrete numbers included in the first series.
3. We multiply each shock by the level of volatility predicted for the day, to obtain a series of simulated returns.
4. We calculate the vector of GARCH-VT conditional variances using the squared simulated returns and the predicted level of volatility (with ω, α, and β already estimated).
5. We repeat each step 21 times, then aggregate the figures to obtain a series of 21-day cumulative returns.
6. We provide the 21-day VaR, as well as the term-structure of VaR for the period.

Continuous sampling from the limited amount of data available produces oscillating statistics for the resulting distributions of returns, even for as many as 10,000 scenarios (Table 19). Mean and variance are gener-

Figure 30: EVT GARCH-VT
VaR Exceedances

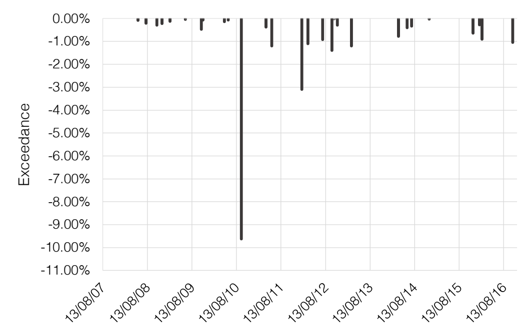

Table 19: FHS GARCH-VT statistics	
No. simulations	10,000
Mean	–0.6225%
Variance	0.2793%
Standard Deviation	5.2846%
Skewness	–0.0546
Excess Kurtosis	6.0962
Simulated 21-day VaR	13.99%

Figure 31: FHS GARCH-VT
Distribution of 21-day Returns

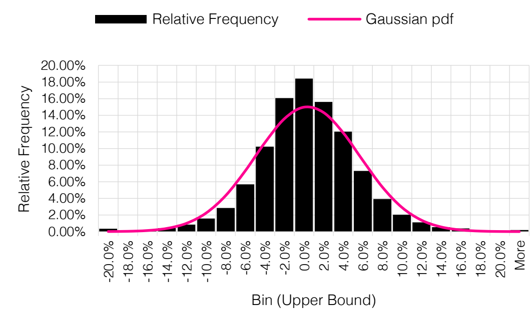

Figure 32: FHS Term Structure
of VaR and Loss Coverage

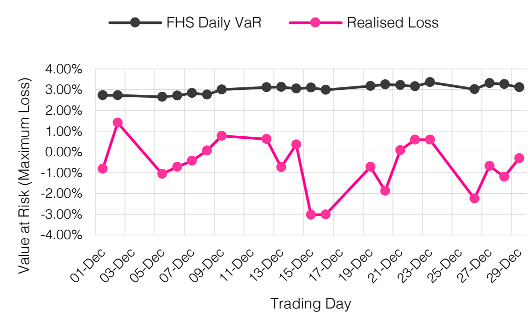

ally stable, but skewness and excess kurtosis greatly vary: the former swings from negative to mildly positive, the latter is in the range [1,6], with rare peaks above 10. The reason is the overrepresentation of extreme events in some simulations, as in the case presented. This issue is visible in the histogram, on which we imposed the Gaussian density (Figure 31).

We predicted maximum aggregated loss, for the 21-day period, to be equal to 13.99%, on average, at 99% confidence level. Market conditions were good in the period, with the highest one-day loss being 1.41% on December 2. Overall, FHS VaR was adequate at covering potential losses even if, at times, a bit expensive (Figure 32 – a negative value represents a gain, or negative loss).

As opposed to FHS, which draws from historical data, Monte Carlo generates new, hypothetical scenarios from a predetermined density. We run two simulations under the assumptions that returns are, respectively, normally and t distributed. The steps for the former are the following:

1. We create 21 series of 10,000 standard Normal RV. The latter are assumed to be random shocks (standardised returns).
2. We calculate the vector of returns one day ahead by multiplying the first set of variables with the most recent GARCH-VT figure.
3. We build the vector of GARCH-VT conditional variances using the squared simulated returns and the predicted level of volatility.
4. We repeat each step 21 times, then aggregate the figures to obtain a series of 21-day cumulative returns.
5. We provide the 21-day VaR, as well as the term-structure of VaR for the period.

Table 20: MCS Gaussian GARCH-VT statistics	
No. simulations	10,000
Mean	–0.0300%
Variance	0.2843%
Standard Deviation	5.3320%
Skewness	–0.0012
Excess Kurtosis	1.3445
Simulated 21-day VaR	13.49%

Figure 33: Gaussian GARCH-VT Monte Carlo Distribution of 21-day Returns

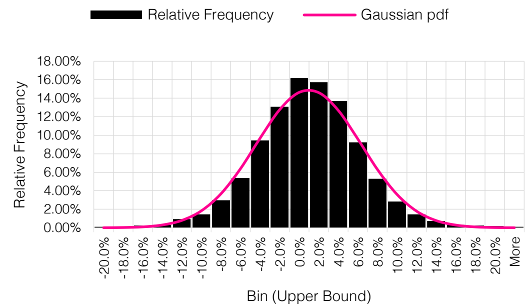

Figure 34: MCS Gaussian Term Structure of VaR and Loss Coverage

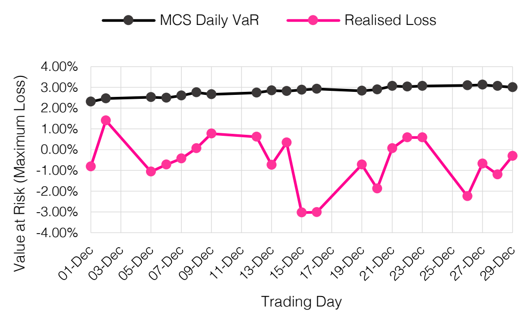

With 10,000 scenarios, MC generated data is very close to being normally distributed. Simulated 21-day return variance is in line with its theoretical value, one-day variance multiplied by 21 (equivalently, standard deviation multiplied by the square root of 21); skewness and excess kurtosis are very close to zero. All figures are stable in repeated experiments. The mean, instead, is found to oscillate between negative and positive values, but the statistic is not too meaningful, as it is still overwhelmed by the standard deviation (Table 20). Data behaviour is apparent in the histogram, with the Gaussian density superimposed (Figure 33).

We found maximum aggregated loss for the 21-day period to be equal to 13.49%, on average, at 99% confidence level. The statistic is inferior to the one predicted for FHS, as expected: the Gaussian density underestimates the probability of extreme negative events. Overall, given normal market conditions in the period, MC simulated Gaussian VaR was enough to face actual losses (Figure 34). MCS under the assumption of Standardised t shocks follows a similar process, apart from step 1 which we modify as follows:

1. Generate 21 sets of 10,000 standardised $t_{4.13}$ RV. The variables are still assumed to be random shocks.

The resulting sample of aggregated returns has values of skewness and excess kurtosis higher, in magnitude, than in the previous case (Table 21). These statistics oscillate, although they are, on average, in line with expectations. Assuming that shocks are distributed according to t gives a better representation of extreme events, as it is visible in the histogram of returns (Figure 35). We superimposed both the standardised $t_{4.13}$ and the Normal densities, to show that the former captures fat tail events, the latter does not. As a consequence, the predicted level of aggregated VaR for the period was higher at 13.84%. Overall, MC simulated $t_{4.13}$ VaR was adequate at capturing potential losses throughout the trading month (Figure 36). At times, we found it to be quite expensive, although the total amount of resources to set apart in the period was reasonable.

Expected Shortfall

VaR has become the industry standard for risk reporting because it is intuitive and it can be applied to all types of securities, including complex portfolios [2]. However, this measure is only concerned with the percentage of losses that exceed its value, and not the magnitude of these losses. Expected shortfall (ES) reveals the value of tomorrow's loss, conditional on it being worse than the VaR. It is an average of all the sorted losses above a predefined confidence level. We present 21-day ES for the three simulation methodologies above, at 99% confidence level (Table 22). The different treatment of

Table 21: MCS Std. $t_{4.13}$ GARCH-VT statistics	
No. simulations	10,000
Degrees of freedom	4.13
Mean	0.0071%
Variance	0.2842%
Standard Deviation	5.3312%
Skewness	–0.1707
Excess Kurtosis	5.9422
Simulated 21-day VaR	13.84%

Figure 35: Standardised $t_{4.13}$ GARCH-VT Monte Carlo Distribution of 21-day Returns

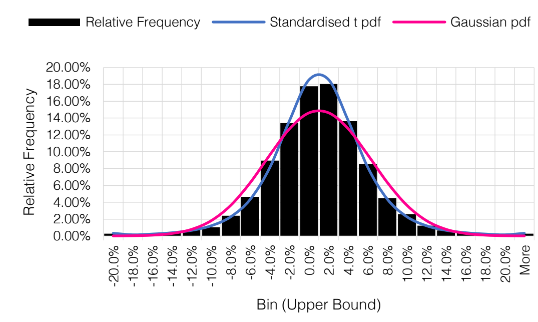

Figure 36: MCS $t_{4.13}$ Term Structure of VaR and Loss Coverage

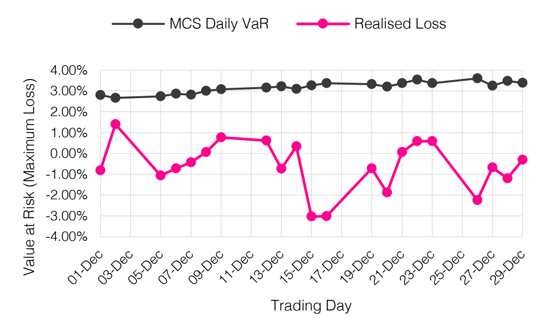

the tails by the three distributions, not so clear-cut when we considered VaR, is apparent with ES. The Gaussian severely underestimates the probability of extreme events, with the likely consequence of not having enough resources to face the actual losses throughout the period. The $t_{4.13}$ distribution, by contrast, is correct in giving a higher weight to the likelihood of crashes. FHS is in-between.
We also compare the term structures of ES for the three (Figure 37).

Is TOP-20 Index a good investment?

We further inspect the distributions of simulated 21-day returns to test whether adding TOP-20 Index to a portfolio of securities is beneficial, in terms of risk and reward involved. We first check which percentage of simulated returns falls within the positive and negative regions. This step is particularly meaningful for FHS, whose distribution is highly asymmetric due to continuous resampling from a limited amount of data; for the other methods, we expect returns to be equally divided between the areas, a consequence of both the symmetric nature of the densities and given the large number of scenarios generated (10,000), the central limit theorem. We then calculate the frequency of observations falling beyond a large return band, ±10%, the so-called extreme events, and find the two expectations of returns conditional on being outside the band. We finally locate the realised, 21-day OOS return, in each distribution, and check how it ranks, compared to others.
Each simulation model requires to set a beginning level of conditional variance for the computation of shocks on day one. We bet on a future rise in volatility, and apply multiplier 2.0 to the most recent figure for conditional variance (estimated on November 30,

Table 22: 21-day simulated Expected Shortfall, %	
FHS GARCH-VT	18.91
MC Gaussian GARCH-VT	16.44
MC $t_{4.13}$ GARCH-VT	19.48

Figure 37: Term Structures of ES

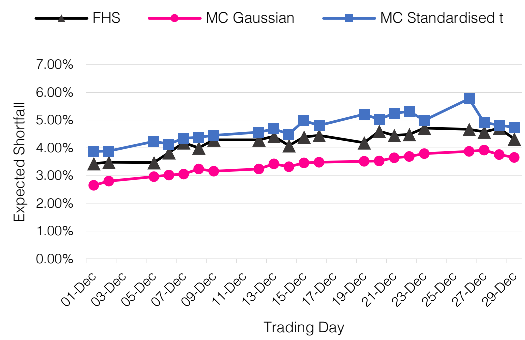

Table 23: FHS distribution inspection, %	
Neg./pos. return frequency	56.48; 43.52
Freq. returns beyond ±10	5.64; 4.04
Conditional expectations	–14.39; 14.56

Table 24: MCS Gaussian distribution inspection, %	
Neg./pos. return frequency	50.23; 49.77
Freq. returns beyond ±10	5.41; 5.09
Conditional expectations	–13.69; 13.77

Table 25: MCS Std. $t_{4.13}$ distribution inspection, %	
Neg./pos. return frequency	49.97; 50.03
Freq. returns beyond ±10	4.30; 4.53
Conditional expectations	–15.06; 14.64

2016) to obtain input value 0.0201% (1.42% conditional volatility). This increased variability also accounts for fatter tails, and it is a conservative estimate of the likelihood of extreme events.
About 56.5% of FHS-simulated data are negative, the remainder positive (Table 23). Also, approximately 5.6% of the observations fall below the ±10% band, 4% above. However, the expectation of returns conditional on being below –10% is lower than that of returns conditional on being above 10%. All these features are a likely consequence of the right skewness of the empirical distribution, preserved by FHS.

MC simulated data are almost equally distributed between negative and positive, as expected. Interestingly, in this particular experiment the percentage of returns falling beyond the ±10% band is higher under the assumption of Gaussian density (Table 24) than under that of Standardised t (Table 25). However, the expectations of returns conditional on being extreme are much higher, in magnitude, for the latter case.

We found the OOS, 21-day return, to be equal to 12.41%, an exceptionally high value which reflects a period of intense trading on the Mongolian stock market, with fast appreciation of TOP-20 Index.

In the distribution of historical aggregate returns, this figure lies on the 93rd percentile, with 160 out of 2300 observations above it (Table 26). The rank is high, although most likely downward biased, since the empirical distribution does not adequately reflect the strong decline in volatility that occurred after mid 2011 (Figure 38). In the other cases, the return lies above the 97th percentile, a datum we think reasonable.

Overall, we believe the frequency of negative returns to be slightly higher than that of positive ones. However, the majority of them is very small in magnitude. Moreover, given the predicted upsurge in volatility, we deem the likelihood of extreme events (both positive and negative) to be approximately 10%, with an expected gain or loss, in case an extreme event takes place, of ±14.5% (conservative estimate). Within the context of a well-diversified portfolio, we think TOP-20 index makes a good investment opportunity.

Derivatives

We conclude this study with a section on derivatives. There exists no traded contract on TOP-20 Index, so it is not possible to make a detailed analysis on the subject. We just provide theoretical prices for European

Table 26: OOS return distribution ranks, %	
OOS 21-day return	12.41
Empirical percentile	93.04
FHS percentile	97.81
MCS Gaussian percentile	97.39
MCS $t_{4.13}$ percentile	97.53

Figure 38: TOP-20 Index Trailing 21-day Returns

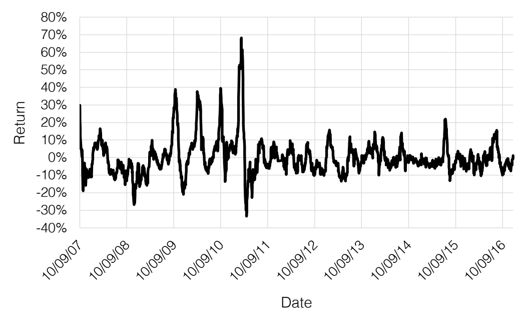

Table 27: TOP-20 EU options parameters	
S_0	MNT 12,456.06
K	∈ [8,000; 16,000]
τ	∈ [3m, 6m, 9m, 12m]
σ (annualised)	25.74%
Current exchange rate	1.00 EUR : 2,601.53 MNT

options using the Black-Scholes-Merton (BSM) formula. We price the contracts in euro, using the most recent exchange rate (December 30, 2016). Theoretical call and put prices are fully specified, in absence of dividends on the index, by the following parameters: current underlying price (S_0), strike (K), time to maturity (τ), volatility (σ), interest rate (r). We include values for all but the last one in a chart (Table 27), and provide the term structure of interest rates, updated with figures from the latest bond issue (December 23) separately (Figure 39). Annualised volatility is calculated as the square root of the product between the most recent level of GARCH-VT conditional variance (0.0182%, as of December 29) and 365, the number of calendar days in a year for option valuation.

We compute call prices using the BSM formula, and derive put prices via put-call parity. We plot the results in terms of moneyness, the ratio between the most recent index level and the strike, for all strikes in the range [8,000; 16,000], with steps of 50 MNT. Fairly typical option patterns emerge. Calls (Figure 40) become more valuable not only when strike price is lower (equivalently, moneyness is higher), but also when maturity increases. The latter is the joint effect of a high level of predicted volatility, which allows for greater swings in the underlying price when more time is given (but with limited downside risk) and of a generally upward sloping interest rate curve, which applies a greater discount to the strike to pay for longer maturities. Puts (Figure 41), instead, become more expensive both when the strike is higher (equivalently, moneyness decreases) and when maturity either lengthens (for near-the-money contracts) or shortens (for far-in-the-money ones). NTM contracts profit from a higher τ because of the greater likelihood that the option will become in-the-money in the future; far ITM instruments profit from a lower τ because the passing of time might make the contract less attractive than it is at the moment. In addition, a generally upward sloping interest rate curve makes shorter maturity instruments more valuable, due to the lower discount applied to the strike price to be received.

We valued call contracts in the range 38.14 MNT (K = 16,000; τ = 3m) to 5,732.58 MNT (K = 8,000; τ = 12m), about € 0.01 to € 2.20, and put contract in the range 0.02 MNT (K= 8,000; τ = 3m) to 2,925.26 MNT (K = 16,000; τ = 3m), about € 0.00 to € 1.12).

We also provide theoretical prices for at-the-money options (moneyness = 1.00), in euro (Table 28).

Figure 39: Term Structure of Mongolian Annualised Risk-Free Rate

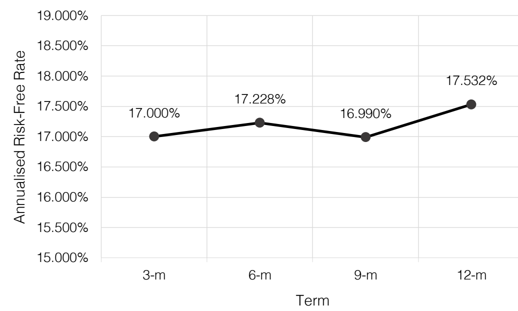

Figure 40: TOP-20 BSM EU Call Prices

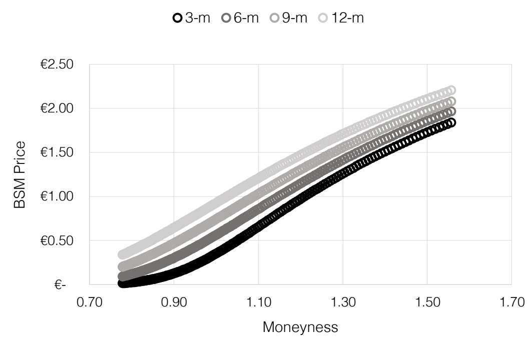

Figure 41: TOP-20 BSM EU Put Prices

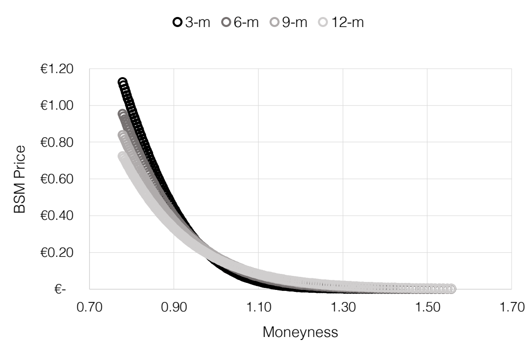

Table 28: TOP-20 ATM option prices (EUR)

	EU Call	EU Put
3-m	0.35	0.15
6-m	0.57	0.17
9-m	0.75	0.17
12-m	0.93	0.17

Conclusion

We presented a detailed risk assessment of TOP-20 Index, the main indicator of the Mongolian Stock Market. We showed that the index exhibits strong non-normality features, in terms of right skewness (higher frequency of large positive, rather than large negative, returns) and visible excess kurtosis (fat tails, thin shoulders, high peak), which hinder the ability to forecast.

To improve predictabily, we found GARCH-VT to be the most accurate to model conditional volatility, Standardised t and EVT to offer the best fit for the empirical distribution of returns.

To monitor daily risk, we suggested Standardised t, EVT (both passed backtesting) or FHS VaR/ES, and advised to switch to CF (a very conservative measure) if market conditions deteriorate. All methods are generally useful to gauge the right amount of resources to set apart to face potential losses.

To model the distribution of returns one month ahead (21 trading days), we think the best course of action is to combine FHS, anchored to the past, with Monte Carlo Standardised t, forward looking, both paired with GARCH-VT. It is always good to have more points of view.

We believe FHS gives a clearer picture of the frequency of both negative and positive 21-day returns in the distribution (about 55% to 45%), and estimate the likelihood of events beyond the ±10% return band at 10%, equally allocated, with an expected gain or loss, conditional on the occurrence of such events, of ±14.5%. The results might be conservative, and reflect our view of a sudden increase in volatility, which translated into a conditional variance multiplier of 2.0 at the time the simulation took place. In case there are no views on the future direction of volatility, we suggest to simply use the most recent level of conditional variance available, with no multiplier.

Although no derivative contract on TOP-20 Index exists yet, we provided theoretical call and put option prices. We found calls would sell in the range € 0.01 to € 2.30 (MNT 40 to MNT 6,000), with at-the-money options between € 0.30 and € 1.00, and put would sell in the range € 0.00 to € 1.15 (MNT 0 to MNT 3,000), with at-the-money contracts around € 0.15 - € 0.20.

Overall, given the recent upsurge in volatility, the higher likelihood of extreme positive, as opposed to extreme negative, events, and the availability of efficient risk monitoring techniques, we consider TOP-20 Index a good investment opportunity, in the context of a well-diversified portfolio.

The report is made for Standard Investment LLC by Federico M. Massari, a long distant volunteer risk analyst, using the sources provided.

---

* All data from mse.mn. We modified the value of the close price recorded on August 13, 2010 from MNT 11,145.50 to MNT 10,145.50; the previous datum was most likely an outlier resulting from transcription error.

Bibliography

[1] Christoffersen, P.F.: *Elements of Financial Risk Management*, 2nd Ed., 2012. Academic Press, Elsevier. Book and companion material.

[2] Ruppert, D.: *Statistics and Data Analysis for Financial Engineering*, 2011, Springer Science+Business Media LLC.

Additional Sources

[a] en.wikipedia.org/wiki/Mongolian_Stock_Exchange
[b] apu.mn/about-us
[c] baganuurmine.mn
[d] bayangolhotel.mn
[e] ubbuk.mn
[f] company.gobi.mn/en/about-us/
[g] genco-tour.mn
[h] mtcone.net
[i] materialimpex.mn
[j] mik.mn/en/m25
[k] mongolpost.mn/more/71
[l] merex.mn/en/page/intro
[m] nekhii.mn/en/about-us
[n] remicon.mn
[o] sharyngol.mn/our-company
[p] mongolmilk.mn/eng/pages/introduction
[q] talkh-chikher.mn/v/4
[r] tavantolgoi.mn/p/4
[s] nomin.mn/en/c/48

Contacts

Federico M. Massari
Long Distant Volunteer Risk Analyst
federico.massari@libero.it
Tel. +39 340 1011568

Standard Investment, LLC
Jigjidjaw St. 5/3, 1st khoroo, Chingeltei district
Ulaanbaatar, Mongolia

Postal Address:
PO Box 1487, Central Post Office
Ulaanbaatar 15160

Tel. +976 7011 4433
info@standardinvestment.mn

Disclaimer

Investors act on their own risk. This is not an offer to buy or sell or the solicitation of an offer to buy or sell any security/instrument or to participate in any particular trading strategy.

All information in this report is for general information only. The information is derived from sources which the company believes to be reliable and prepared in good faith. Standard Investment LLC makes no guarantee of accuracy, timeliness and completeness of the information. Neither Standard Investment nor its affiliates shall be liable for any damages arising out of any person's reliance upon this report.