

# Investment Analytics

## F1000 Volatility Opportunity Strategy

---

### Strategy Analysis

---

**PRIVATE AND CONFIDENTIAL**

Prepared for:

Copy #

**NOTICE:**

This document is intended only for the use of the individual or entity to which it is addressed, and contains private and confidential information. If the reader of this document is not the intended recipient you are hereby notified that any review, retransmission, dissemination, distribution, copying or other use of this document is strictly prohibited. All material contained herein is copyright © 2004 Proteom Capital Management Ltd.

---

# Contents

Executive Summary .....	2
Data Set .....	3
Portfolio Composition .....	3
Performance Summary .....	4
Outperformance and Consistency in Strategy Returns.....	4
Stock- and Sector-Specific Effects in Strategy Returns.....	5
Seasonal Effects in Strategy Returns .....	7
Correlation Effects in Strategy Returns.....	7
Strategy Returns – Negative Findings.....	8
Portfolio Risk.....	10
Stock-Specific Risk.....	11
Tail Risk .....	11
The Model.....	15
Optimizing The Model .....	15
Impact on Recovery Probabilities .....	15
Summary of Forecasting Performance.....	17
Analysis of Forecasting Performance .....	17
Implications for Option Pricing.....	19
Appendix 1 – Statistical Measures.....	20
Appendix 2 – Simulation of Returns Process .....	22
Appendix 3 – Test Statistic Distributions.....	23

# 1 Executive Summary

## Executive Summary

This study analyses in detail the Volatility Opportunity Strategy returns and its risk characteristics. The following are the key findings:

- The average tenor of the portfolio varies between 17 and 35 days, depending on which point we are at in the expiration cycle.
- The Strategy has outperformed the S&P 500 index by an average of just over 100bp per day.
- Winning days substantially exceeds losing days in proportion and Sharpe ratio.
- There are no significant seasonal effects in Strategy returns.
- The evidence suggests that the strategy is making money by volatility arbitrage, rather than by selling “terrorism insurance” or other insurance to the market.
- Correlation between Strategy returns and SP500 index returns is a negligible 0.09.
- Correlation between the Strategy and the VIXX index is 0.00.
- Upside volatility is 1.6 times larger than downside volatility.
- Strategy returns process has a large positive skewness and kurtosis.
- Maximum tail risk in any individual stock is limited to no more than 2% of total risk capital.
- Strategy is highly robust and could likely withstand, and even fully recover from, a tail event of a catastrophic magnitude.
- Worst-case scenario of a 1-day drawdown of -50% (around 12 standard deviations) is likely to arise once in every 12,000 years.
- Risk-adjusted returns can be optimized by reducing strategy volatility to around 25%, perhaps by operating at around one half of the current levels of leverage.
- Strategy models account for almost 32% of variation in the volatility processes, approximately ten times greater than is typically achieved with similar models applied to asset returns processes.
- Strategy models are almost 71% accurate in predicting the future direction of volatility.
- When the models are wrong in predicting the future direction of volatility, the percentage error is typically of the same order of magnitude as during periods when the models correctly predict the change in sign.

# 2

# Strategy Performance

## Data Set and Analysis Summary

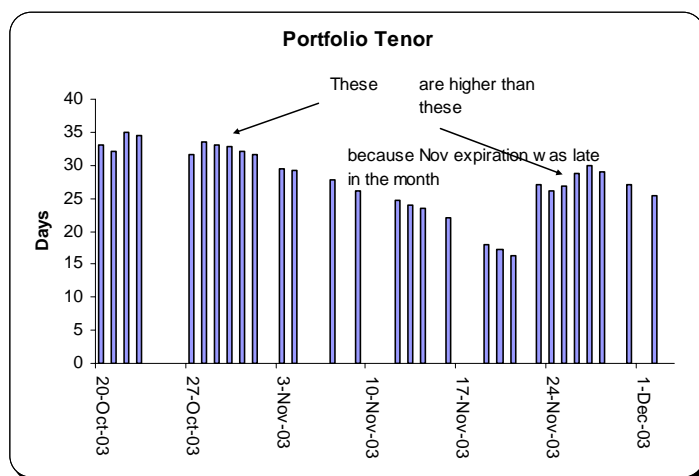
### Data Set

The data consists of fund daily gross returns for the period 1<sup>st</sup> October 2002 to 25<sup>th</sup> November 2003. The daily returns on the S&P 500 Index have also been included for comparison purposes.

### Portfolio Composition

The strategy invests in equity and equity index options in a universe comprising the S&P500 and QQQ indices and 100 of the S&P 500 index constituent stocks. Option portfolios consist of long and short positions in the nearest two option expirations, with a maximum tenor of 45-50 days. The average tenor of the investment portfolio will fluctuate over time. After expiration we start to add options with around 45 days to expiration, gradually increasing the tenor of the portfolio as we do so. Opposing this, the passage of time ages existing option positions causing the tenor to decline. Around a week to ten days after expiration the decrease in tenor due to time passing overcomes the increase in tenor produced by adding new option positions. At expiration of one series of options the tenor exactly matches the time to the next expiration.

This pattern is illustrated in the following chart.



---

# Analysis of Strategy Returns

## Performance Summary

A summary analysis for the return series is shown in the table below

	SPX	VoF	Up Mkt	Dwn Mkt	Win Days	Lose Days	W/L Ratio
avg	0.089%	1.092%	1.560%	0.591%	2.941%	-1.673%	1.76
compound	22.5%	1152.5%	3785.6%	272.2%	123124.4%	-98.7%	
st. dev	1.267%	4.244%	4.675%	3.684%	4.192%	2.574%	
ann.vol	20.11%	67.36%	74.22%	58.47%	66.54%	40.86%	1.63
#	290	290	152	136	172	113	1.52
t-stat	1.19	<b>4.38</b>	<b>4.11</b>	1.87	<b>9.20</b>	<b>-6.91</b>	
win%	52%	59%	61%	58%	100%	0%	
best	4.6%	36.3%	36.3%	14.3%	36.3%	0.0%	
worst	-3.59%	-20.05%	-11.23%	-20.05%	0.00%	-20.05%	
skewness	0.3	2.6	3.8	-0.5	4.5	-4.2	
kurtosis	0.7	22.0	24.1	9.2	27.7	24.5	
correl		0.09	-0.08	0.14	0.05	0.19	
beta		0.29	-0.21	0.22	0.14	0.42	
Sharpe	1.02	17.08	50.98	4.62	1850.33	-2.46	
autocorrel	-0.15	0.02	-0.09	0.38	0.27	0.74	

## Outperformance and Consistency in Strategy Returns

The Volatility Opportunity Strategy has outperformed the S&P 500 index by an average of just over 100bp per day over the sample period. Annually compounded, the differential return over the sample period amounts to 1130% per annum. Correspondingly, the Sharpe ratio for the strategy is around 17 times greater than that for the S&P 500 index.

The results also indicate other elements of the robustness of the strategy and its ability to outperform consistently under most market conditions:

- The winning % on down-market days (when the S&P 500 index closes lower) is very close to the winning % on up-market days (58% vs. 61%)
- The proportion of winning days substantially exceeds the proportion of losing days (60% to 40%)
- On winning days the average daily return is 1.76 the size of the average daily loss of days when the strategy loses money (+2.94% vs. -1.67%)
- The Sharpe ratio for the strategy on winning days exceeds by several order of magnitude the ratio for losing days (1850.3 vs. -2.5)

- On up-market days (when the SP500 index closes higher on the day) the strategy outperforms the average daily return by approximately 47bp. On down-market days the strategy under-performs the average daily return by approximately 50bp. However the strategy still returns a creditable daily average of 59bp, even on days when the overall market is down.

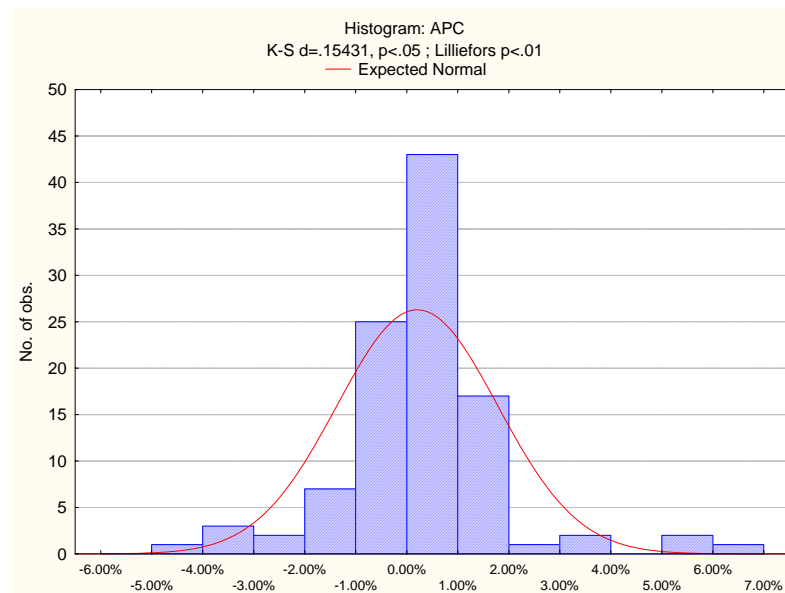
## Stock- and Sector-Specific Effects in Strategy Returns

In examining the source of strategy returns, one problem we encounter is the choice of a suitable metric. The obvious measure, percentage of total return, suffers from the problem that gains and losses cancel each other out, resulting in a denominator which is small relative to the profit or loss on each individual stock. Faced with a similar problem in assessing forecasting accuracy, one solution is to use absolute percentage errors. Analogously, the Absolute Percentage Contribution (APC) is defined as follows:

$$APC_i = P_i / \sum_{i=1}^n |P_i|$$

where  $P_i$  is the profit on stock  $i$  (including realized and unrealized gains and losses on all stock and option positions),

The chart below shows the distribution of APC across the investment universe.



The great majority (88%) of the individual values lie inside the range +/-2%. The largest winners and losers are as follows:

Losers		Winners	
Stock	APC	Stock	APC
HDI	-4.47%	AIG	2.54%
WYE	-3.96%	MMM	3.51%
BA	-3.21%	FNM	3.93%
THC	-3.08%	PG	5.10%
QQQ	-2.96%	SPX	5.97%
TXN	-2.81%	EDS	6.38%

Some of the stocks in the tails of the APC distribution can be approximately “paired off” as follows:

- QQQ vs SPX (indices)
- TXN vs EDS (technology)
- BA vs FNM (financial)
- WYE vs PG (pharmaceuticals)

A breakdown by GICS sector is shown in the table below. The standouts are the indices

GICS Sector	# Stocks	Avg.APC
Consumer Discretionary	10	0.11%
Consumer Staples	9	-0.14%
Energy	7	0.29%
Financials	13	0.59%
Healthcare	19	-0.02%
Index	2	1.51%
Industrials	8	0.11%
Information Technology	25	0.27%
Materials	4	0.39%
Telecommunication Services	7	0.01%
<b>Total</b>	<b>104</b>	<b>0.21%</b>

(+1.51%) , Financials (+0.59%), and Materials (+0.39% ). The very low contribution from Airlines, Travel and other industries in the consumer discretionary sector argues against the theory that the strategy is making money by selling “terrorism insurance” to the market.

The relative concentration of the Absolute Percentage Contribution from the indices is not as concerning as it might initially appear:

- The liquidity in the indices is orders of magnitude greater than for any individual stock
- The tail risk is smaller than for any individual stock
- Arbitrage in index options represents more of a pure volatility play

## Seasonal Effects in Strategy Returns

An analysis of the pattern of returns indicates that there are no significant weekday effects. None of the t-statistics comparing the average weekday return vs. the overall daily average is significant at the 95% level (see table below).

<b>Weekday</b>	<b>Average of VoF</b>	<b>StdDev of VoF</b>	<b>t-Stat</b>	<b>t-Stat vs Avg</b>
1	0.301%	4.16%	0.541	-1.402
2	1.234%	4.48%	<b>2.150</b>	0.267
3	0.938%	2.91%	<b>2.451</b>	-0.374
4	1.393%	3.69%	<b>2.923</b>	0.655
5	1.493%	5.48%	<b>2.077</b>	0.574
<b>Grand Total</b>	<b>1.081%</b>	<b>4.22%</b>	<b>4.381</b>	

An analysis of day-of-month effects indicate that the strategy has produced abnormally low returns on the 12<sup>th</sup> and 26<sup>th</sup> of each month, and abnormally large returns on the 18<sup>th</sup> of the month, compared to the overall daily average return (table not shown). The above-average return on the 18<sup>th</sup> of the month appears to tie in with option expiration, but there is no ready explanation for the other effects. There is a danger in reading too much into these results – not only are the sample sizes small, but one would expect to see at least one such “significant” result (at the 95% level) in a sample of 31 t-tests, simply by random chance.

## Correlation Effects in Strategy Returns

The question arises as to whether strategy performance is tied to the performance of the overall market. For example, does the strategy tend to make more money on days when the overall market performs well? More importantly, since this is a volatility strategy, what is the relationship, if any, between strategy performance and the overall level of market volatility?

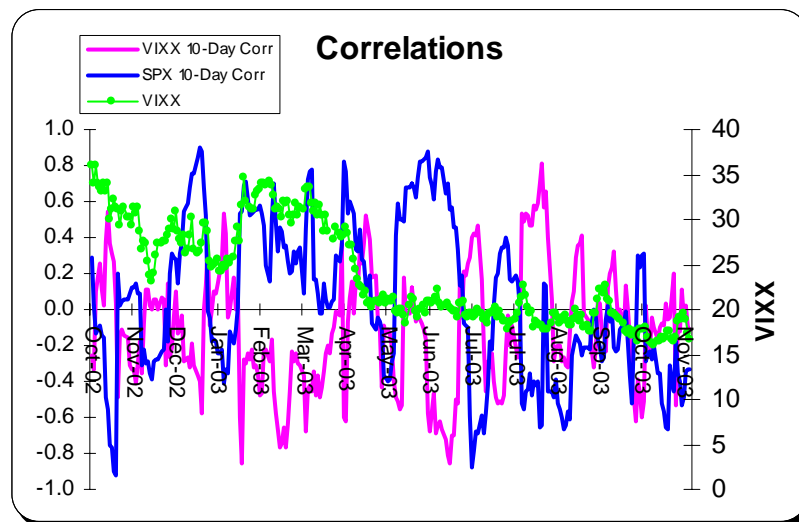
To delve into these questions we have analyzed the performance of the strategy in relation to both the S&P500 index and the VIXX volatility index over the sample period. The correlation between strategy returns and SP500 index returns is a negligible 0.09, while that between the strategy and the VIXX index is 0.00.

A chart of the 10-day correlation between the strategy returns and the two indices is shown below. The correlation between the strategy and the index varies widely between -0.9 to +0.9, and there is no evidence of any pattern in the process.

The same holds true for the correlation between the strategy and the VIXX index.

Furthermore, statistical tests reject the conjecture that the degree of correlation between strategy performance and the VIXX is in some way related to the level of the index itself.

In other words, the strategy is correlated neither with the market nor with the level of market volatility. It has demonstrated the ability to perform well in periods when the market is strong or weak, and when market volatility is at historically high or low levels.



## Strategy Returns – Negative Findings

There are some aspects of the strategy's performance that fall short of the ideal:

- Although the correlation between the strategy and the benchmark returns is insignificant overall (0.09), the correlation is slightly positive on days when the market is down. (The correlation for up-market days is not statistically significant).
- The strategy tends to have a small negative beta on days when the market is up and a small positive beta when the market is down. However, given the insignificant overall correlation between the strategy and the index, the conception of portfolio beta (and, for that matter, alpha) is quite meaningless.
- The autocorrelation in the strategy returns process is larger for losing days than it is for winning days. In other words, a losing day typically tends to be followed by more

losing days, and this pattern is stronger than the corresponding behavior on winning days. Put another way, the strategy tends to make small losses over several days, while its wins are often single large gains. Of course, the overall returns from the gains far outweigh the losses from losing runs.

---

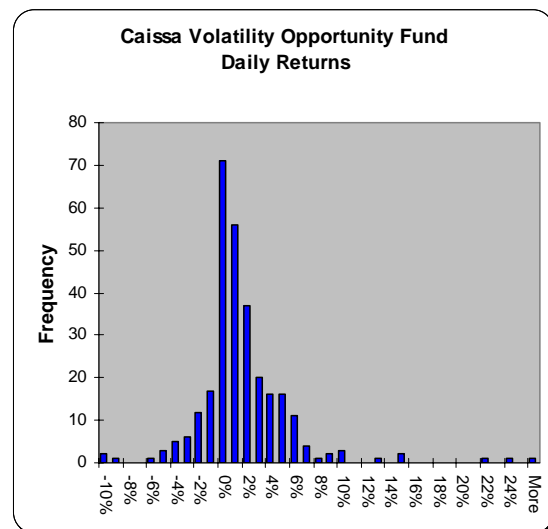
# Analysis of Strategy Risk

## Portfolio Risk

A breakdown of performance into winning and losing days highlights a very important aspect of the risk inherent in the strategy: the majority of the variation in the strategy returns process results from upside rather than downside volatility. This is implied by the fact that:

- Upside volatility, at 66.5%, is 1.6 times larger than downside volatility, at 40.9%.
- Overall volatility is weighted more towards upside volatility rather than downside volatility, since 60% of the returns were positive.

Taken together, these findings suggest that the appropriate quantity to consider for risk management purposes is the downside volatility estimate of 40.9% and consequently that the strategy is much less risky (in the sense of the probability of loss) than might otherwise be inferred. Furthermore, an analysis of Value-at-Risk based on the standard Delta-Normal model is likely to seriously overstate the downside risk (and underestimate the upside potential) if the higher, overall strategy standard deviation is used. Analysis of higher moments of the returns distribution add support to the theory that upside “risk” far outweighs the risk to the downside. Overall, the strategy returns process has a large positive skewness and kurtosis. This is reflected in both the number and magnitude of the extreme positive returns (9 daily returns of 10% or more vs. only 2 negative returns of -10% or greater).



In other words, while the strategy returns distribution is definitely non-normal, it is so in a manner which is favorable both for risk and return: the right tail is significantly “fatter” than the left tail and the overall distribution is positively skewed.

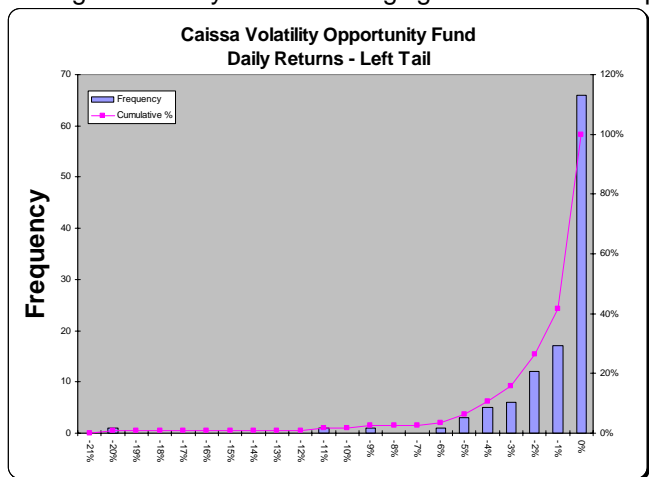
## Stock-Specific Risk

A number of risk management procedures are applied to limit the stock specific risk:

- Any individual stock and option position which accounts for more than 15% of the total portfolio value-at-risk is hedged using the underlying stock and/or options on the underlying stock as required to reduce the VaR contribution to below 10%.
- The maximum tail risk in any individual stock is limited to no more than 2% of total risk capital, with the stock subject to a maximum move of +/- 20%

## Tail Risk

With only thirteen months of daily observations, it is impossible to be certain of the true shape of the left tail of the distribution. So far we have seen only one case of a loss exceeding 20% in a sample of 290 observations [the chance of a greater loss than this arising is limited by the crash hedging structured into the portfolio]<sup>1</sup>. However, there is no

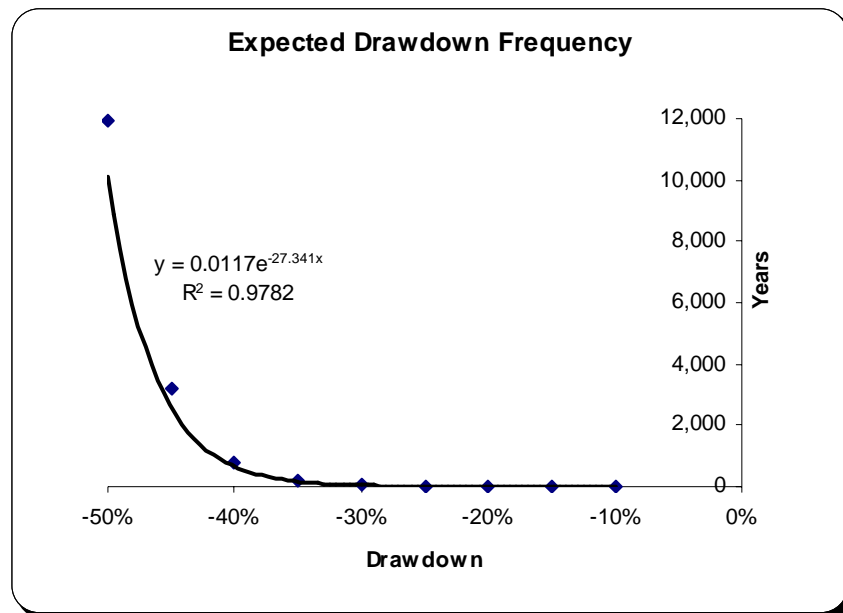


indication of how typical the sample may be, and because of the significant amount of discretion employed in the strategy, reliable back-testing is problematic. In this regard, it is reassuring that back-tests of the more automated Strategic Volatility Fund strategy have produced results

which are almost identical in every respect with the live performance of the fund over the same sample period.

<sup>1</sup> The single observation we have seen of this magnitude was probably the result of mark-to-market error in the prime brokerage reporting system rather than an actual loss.

To get a handle on the tail risk in the strategy and its likely impact on performance we have constructed the following simulation scenario. First, we assume that, during the course of a year, a single tail event occurs with a one-day drawdown in the range -10% to -50%, which is between 2 and 12 standard deviations. This range is up to 2 ½ times the size of the maximum daily drawdown experienced during the sample period and the probability of a tail event of this magnitude is extremely small indeed. To illustrate, we have used a loglogistic distribution function with “fattened” tails to give an idea of how frequently the simulated tail events are likely to occur. The results are shown in the chart below. Based on our analysis we would expect a -10% 1-day drawdown about once every six months, and a -20% drawdown about once a year (in line with the empirical evidence). **However, the worst-case scenario, a 1-day drawdown of -50%, is likely to arise only once in every 12,000 years!**



Having established the point that we are making extremely conservative assumptions about the tail risk in the strategy we continue with the analysis as follows. We make the further simplifying assumption that the tail event may occur on the first day of any month of the year and, up until that point the strategy has been generating returns at the average daily rate produced over the sample period ( $R_A = 1.092\%$  per day). We then compute the daily rate of return that the strategy would require to generate over the remainder of the year in order to completely recover 100% of the initial capital. This “recovery return” (RR) is computed as:

$$RR = -1 + \left( \frac{1}{C_N} \right)^{\frac{1}{(252-N)}}$$

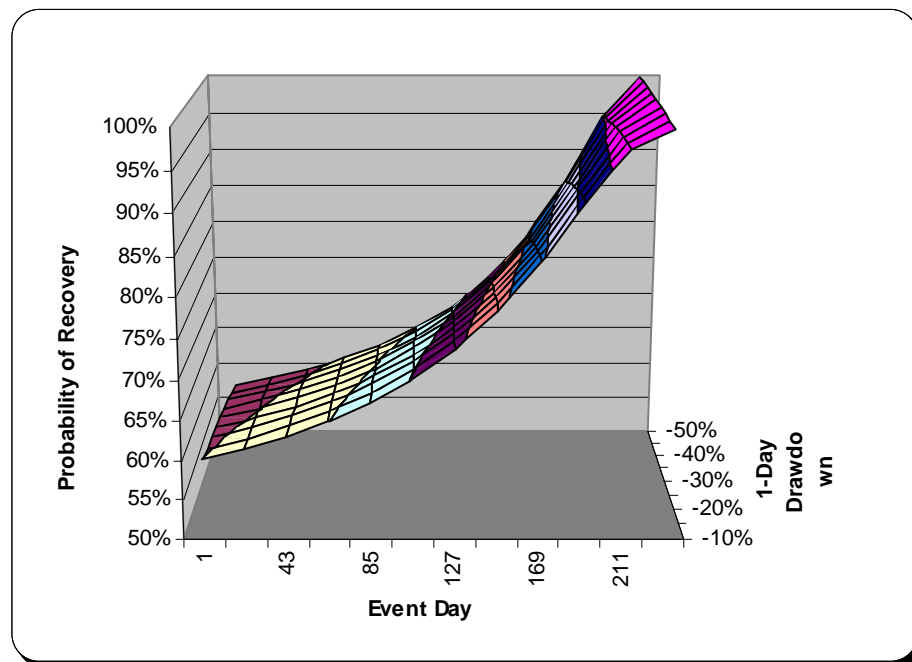
Where  $C_N$  is the capital remaining after the drawdown on day  $N$ , given by:

$$C_N = (1 + R_A)^{(N-1)} (1 + D)$$

Where  $D$  is the drawdown percentage:  $-10\% \geq D \geq -50\%$

Finally, we can estimate the probability of achieving the recovery return,  $RR$ , using the empirical returns distribution over the sample period. We approximate this by generating probabilities from a loglogistic distribution, which provides a best fit to the empirical distribution (especially in the tails) – see Appendix 2 for details.

The results are summarized in the chart below.



The simulation indicates that the probability of recovering from a -25% drawdown on the first day of trading, a tail event that may be expected to arise once in every 6 ½ years, is around 59%. If the strategy experiences such a drawdown six months into the year, the recovery probability will have risen to over 71%. And by the final quarter of the year the chances of recovering from a drawdown of this magnitude rise to over 90%.

Even in the worst case, where a drawdown of -50% occurs on day 1 (i.e. half the initial capital is lost immediately), the probability of making a full recovery is better than 56%. By the half-year stage the recovery probability will have risen to over 65%. If the tail event is postponed until the final quarter of the year the recovery probability is at least 85%.

This simple model suggests that the strategy is highly robust and could likely withstand and even fully recover from a tail event of a magnitude far greater than has been so far experienced, or even one that is unlikely to be seen in the next several millennia.

#### **Extreme Value Modeling**

We use extreme value theory to provide a more realistic assessment of the likelihood of a tail event.

Our tail-augmented log-logistic distribution has a left-tail area approximately 10 times the magnitude of the equivalent Normal distribution.

#### **Model Limitations**

The model is less than fully realistic in several respects:

- No account is taken of the potential for a sequence of drawdowns, rather than a single large event.
- A tail event shock might alter the strategy returns process and lead to lower average returns after the event. This would reduce the likelihood of recovery.
- It is not clear how well the strategy will scale and the capacity that can be accommodated without resulting in significant deterioration in performance.

---

# Strategy Optimization

## The Model

This section is intended to give some broad guidelines as to how to optimize the strategy. We need to begin with some basic assumptions. We continue to assume that the returns distribution is approximately distributed as loglogistic, with location parameter gamma, and shape parameters alpha and beta. While the fit to the empirical distribution is less than perfect, its behavior in the all-important tails is the closest to what the empirical data suggests we should expect.

Our initial model is one which matches the empirical distribution as closely as possible, with parameters as follows:

<b>Gamma</b>	<b>Beta</b>	<b>Alpha</b>
-0.382	0.389	22.576

With this distribution the simulated returns process yields shows an annual return of 867%, an annual volatility of 50% and a Sharpe ratio of just over 17.

## Optimizing The Model

If we optimize the shape parameters to maximize the Sharpe ratio we find the following results:

<b>Gamma</b>	<b>Beta</b>	<b>Alpha</b>	<b>Mean</b>	<b>SD</b>	<b>Ann Vol</b>	<b>Ann Return</b>	<b>Sharpe</b>
-0.382	0.390	44.707	0.861%	1.58%	25.2%	867.1%	34.39

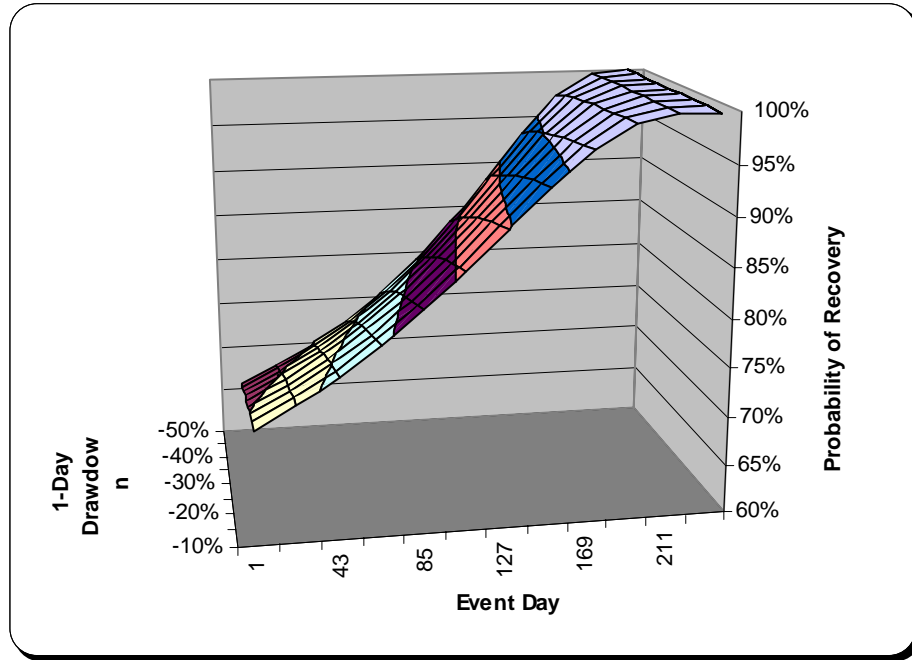
In other words, our theoretical model suggests that we can maximize risk-adjusted returns by reducing strategy volatility to around 25%, perhaps by operating at around one half of the current levels of leverage.

Of course, this analysis does not take into account the impact that other investments or strategies may have in determining the optimal configuration of any specific strategy or the optimal mix of strategies – that goes to the issue of portfolio management theory, which is not the subject of this analysis.

## Impact on Recovery Probabilities

How does portfolio optimization impact recovery probabilities? We can assess this using the same procedure described in the previous section on tail risk, but this time we use the

parameters of the optimized returns distribution to estimate the recovery probabilities. Naturally, as we have notionally reduced the strategy volatility, value-at-risk will be reduced, while the probability of recovering from a drawdown of any given size will be greater. In fact the impact on recovery probabilities is substantial, as illustrated in the chart below.



Now we find that the probability of recovering from a -25% drawdown on day 1 is as high as 70% for the optimized strategy, compared with 59% for the current version for the strategy. In fact, for the optimized strategy the recovery probability after a drawdown of any size occurring after the first two months of the year are at least 70%, rising to 80% or more in the case of a drawdown occurring in the second half of the year.

A table of recovery probabilities for the optimized portfolio is shown below.

	-10%	-15%	-20%	-25%	-30%	-35%	-40%	-45%	-50%
<b>1</b>	71.4%	70.9%	70.3%	69.7%	69.0%	68.2%	67.4%	66.5%	65.5%
<b>22</b>	73.7%	73.1%	72.5%	71.8%	71.1%	70.3%	69.5%	68.5%	67.5%
<b>43</b>	76.2%	75.6%	75.0%	74.3%	73.5%	72.7%	71.8%	70.8%	69.7%
<b>64</b>	79.1%	78.5%	77.8%	77.1%	76.3%	75.5%	74.6%	73.5%	72.3%
<b>85</b>	82.3%	81.7%	81.1%	80.3%	79.6%	78.7%	77.7%	76.7%	75.4%
<b>106</b>	85.9%	85.3%	84.7%	84.0%	83.2%	82.4%	81.4%	80.3%	79.1%
<b>127</b>	89.7%	89.2%	88.6%	88.0%	87.3%	86.5%	85.6%	84.5%	83.3%
<b>148</b>	93.5%	93.1%	92.7%	92.2%	91.6%	90.9%	90.1%	89.2%	88.1%
<b>169</b>	96.9%	96.6%	96.4%	96.0%	95.6%	95.1%	94.6%	93.9%	93.0%
<b>190</b>	99.2%	99.0%	98.9%	98.8%	98.6%	98.4%	98.1%	97.8%	97.3%
<b>211</b>	99.9%	99.9%	99.9%	99.9%	99.9%	99.8%	99.8%	99.7%	99.7%
<b>232</b>	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%	100.0%

# 3

# Forecasting

## Analysis of Forecasting Performance

### Summary of Forecasting Performance

In this section of the analysis we examine the performance of one class of forecasting model applied to data series comprising daily volatility processes in our current universe of 100 stocks. The selected class of model is a single factor model which combines short-term and long-term memory effects and interactions. The performance results are summarized in the table below.

	R2	MSE	MAD	MAPE	MAPE+1	MAPE-1	DP	Theil's U
Avg	31.8%	19.9%	35.1%	10.2%	8.7%	13.7%	70.9%	0.77

### Analysis of Forecasting Performance

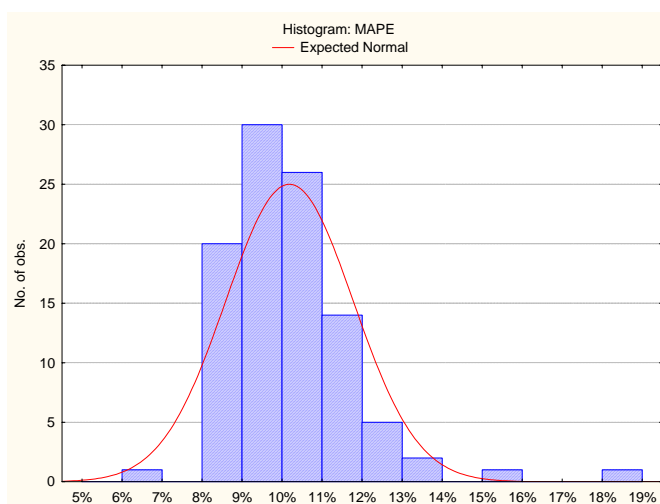
In practice we assess the performance of each model on over twenty different criteria, but for the purpose of this assessment we focus on a limited subset which have immediate relevance to assessing the impact of forecasting model performance on strategy performance. The performance metrics we use are described in Appendix 1. Histogram plots of the test statistic results are given in Appendix 2.

The average R-square of the models indicates that they are accounting for almost 32% of the variation in the volatility processes, over the sample period. While this leaves of great deal of unexplained variation ("noise"), the signal to noise ratio is approximately ten times greater than is typically achieved with similar models applied to asset returns processes. In other words, compared to stock returns, stock volatility is inherently much more predictable given an appropriate model framework.

Theil's-U compares the performance of the forecasting models against a naïve predictor, which uses the volatility in the previous period as the forecast of volatility in the next period. A value of 1 indicates no special forecasting ability compared to the naïve predictor. The average value of 0.77 achieved by the models indicates a highly statistically significant ability to outperform the naïve forecast.

The Direction Prediction indicator looks at how good the models are at predicting the direction, rather than the magnitude, of the 1-period change in the process. In other words it measures how successfully the models predict whether volatility is likely to expand or contract in the next period. A model without predictive power would return an average value of around 50%. A talented market timer might have an edge of around 58% correct direction predictions. Clearly these models far exceed the ability of the great majority of analysts to forecast the direction of market volatility.

The MAPE is the most intuitively obvious of the performance measures, indicating by how much (in percentage terms) the models tend to under- or over-estimate future volatility.



The average value of just over 10% is very low by comparison with results for, say, asset return forecasting models. More importantly, this low error threshold leaves plenty of scope for identifying significant mispricings in the related options series (see section 4.2).

The next two columns of the analysis break down the MAPE results and provide further insight into forecast accuracy. The MAPE+1 statistic indicates that on average the models under- or over-forecast by 8.7% during periods when they correctly anticipate the sign change in volatility. In periods when the direction forecast is incorrect the MAPE is, unsurprisingly, much larger at 13.7%. The overall MAPE, being a weighted average of the two directional MAPE measures is nearer the lower MAPE+1 result because the models correctly predict the sign change around 71% of the time. These findings show that, even when the models are wrong in predicting the future direction of volatility, the percentage error is typically of the same

order of magnitude as during periods when the models correctly predict the change in sign.

## Implications for Option Pricing

Consider a simple Black-Scholes world in which we permit some degree of uncertainty as to the future level of actual volatility over the life of an option. For at-the-money options, the Black-Scholes price is approximately linear in the volatility parameter,  $\sigma$ . Hence an ATM option priced at  $\sigma \pm 10\%$  will be over (under) priced by approximately 10%. The Caissa volatility arbitrage strategy identifies options which are mis-priced by a **minimum** of 30% (and mis-priced by around 60%, on average). This means that the models are finding differentials between implied and forecast actual volatility of at least 30%. Hence, with a MAPE of around 10%, there is scope for **at least** a minimum 20% mis-pricing (and an average closer to 50%), even allowing for forecast error.

# 4

# Appendices

## Appendix 1 – Statistical Measures

### R-Squared

The regression R-squared found by regressing the actual values  $y_t$  against the forecast values  $f_t$ .

### Mean Square Error (MSE)

The variance of the error process  $e_t = (f_t - y_t)$ .

$$MSE = \frac{1}{n} \sum_{i=1}^n e_i^2$$

### Mean Absolute Deviation (MAD)

The average of the absolute values of the forecast errors.

$$MAD = \frac{1}{n} \sum_{i=1}^n |e_i|$$

### Mean Absolute Percentage Error (MAPE)

The average of the absolute values of the forecast errors, expressed as a percentage of the actual value.

$$MAPE = \frac{1}{n} \sum_{i=1}^n \left| \frac{e_i}{y_i} \right|$$

### MAPE + 1

The average of the absolute values of the forecast errors, expressed as a percentage of the actual value, for periods when the sign change is correctly forecast (1 step ahead).

### MAPE -1

The average of the absolute values of the forecast errors, expressed as a percentage of the actual value, for periods when the sign change is incorrectly forecast (1 step ahead).

### Direction Prediction (DP)

% of correct changes in sign in the forecast series (one step ahead).

$$DP = \frac{1}{n} \sum_{i=1}^n z_i$$

$$z_i = \begin{cases} 1 & \text{if } (f_{i+1} - y_i)(y_{i+1} - y_i) > 0 \\ 0 & \text{otherwise} \end{cases}$$

Theil's U

Compares the accuracy of the model forecasts against a naïve predictor ( $f_{t+1} = y_t$ ). Values under 1 indicate superior forecasting performance by the model.

$$U = \sqrt{\frac{\sum_{t=1}^n (FPE_{t+1} - APE_{t+1})^2}{\sum_{t=1}^n APE_{t+1}^2}} \bigg/ \sqrt{\frac{\sum_{t=1}^n \left(\frac{f_{t+1} - y_{t+1}}{y_t}\right)^2}{\sum_{t=1}^n \left(\frac{y_{t+1} - y_t}{y_t}\right)^2}}$$

## Appendix 2 – Simulation of Returns Process

The distribution that provides a best-fit<sup>2</sup> to the empirical returns distribution of the strategy is the loglogistic distribution with parameters as follows:

Gamma = -0.3816, Beta = 0.3889, Alpha = 22.576

The cumulative density function is given by:

$$F(x) = \frac{1}{1 + \left(\frac{1}{t}\right)^\alpha} \quad \text{where } t = \frac{x - \gamma}{\beta}$$

The mean of the distribution is given by

$$\beta\theta \csc(\theta) + \gamma \quad \text{where } \theta = \frac{\pi}{\alpha}$$

The variance of the distribution is given by

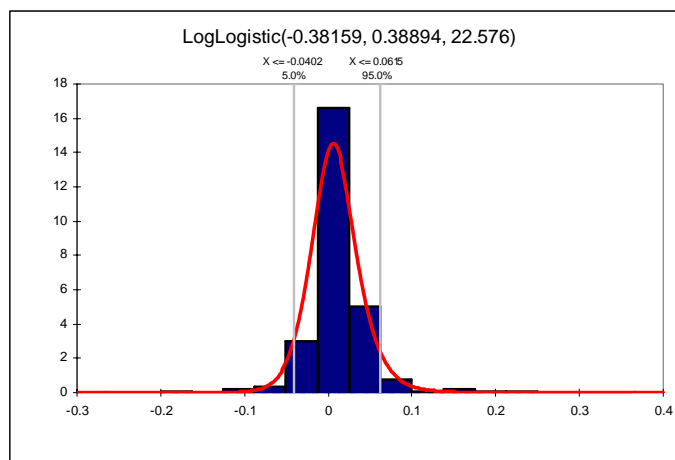
$$\beta^2\theta [2 \csc(2\theta) - \theta \csc^2(\theta)]$$

We use the Maclaurin expansion of the cosecant function:

$$\csc x = \frac{1}{x} + \frac{1}{6}x + \frac{7}{360}x^3 + \frac{31}{15120}x^5 + \dots + \frac{(-1)^{n+1} 2(2^{2n-1} - 1)B_{2n}}{(2n)!}x^{2n-1} + \dots,$$

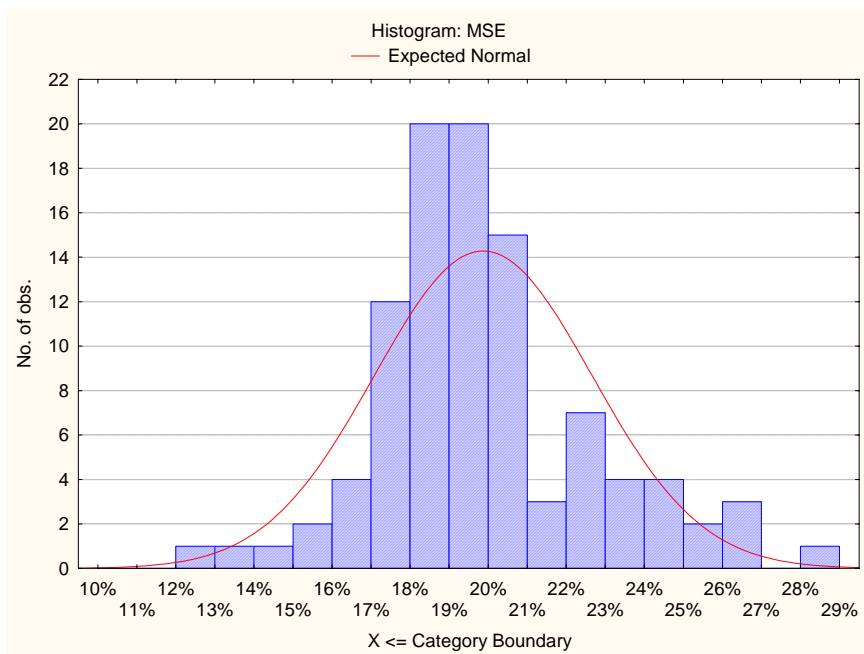
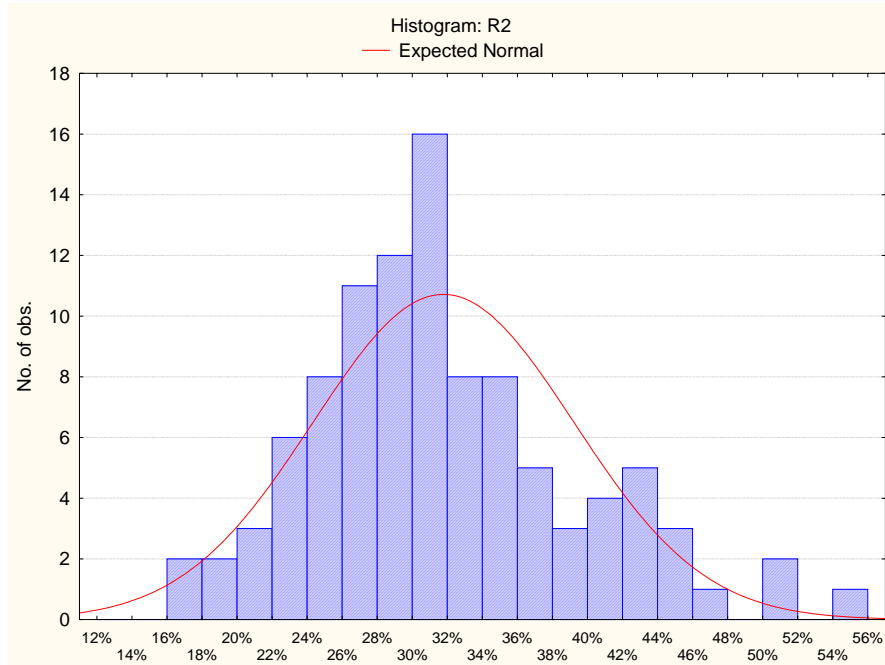
where  $B_{2n}$  is a Bernoulli number.

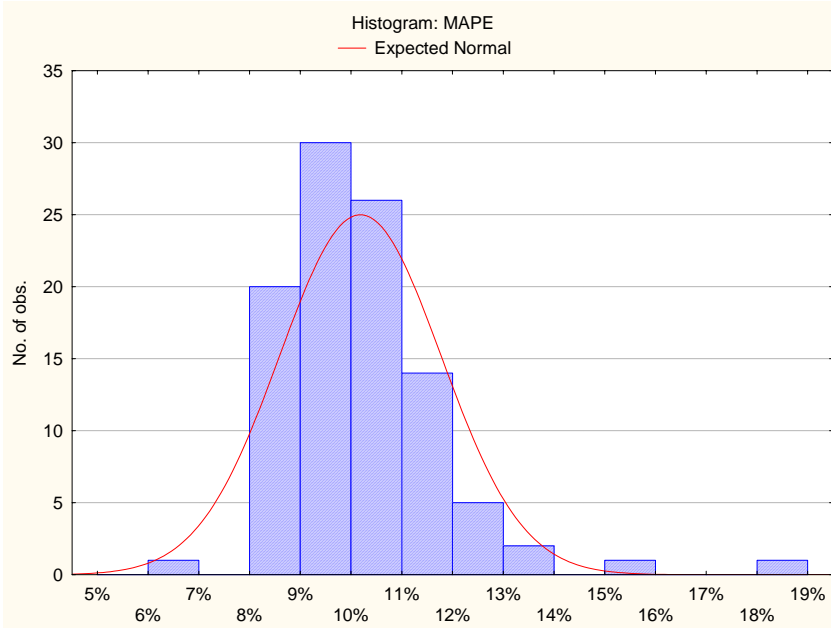
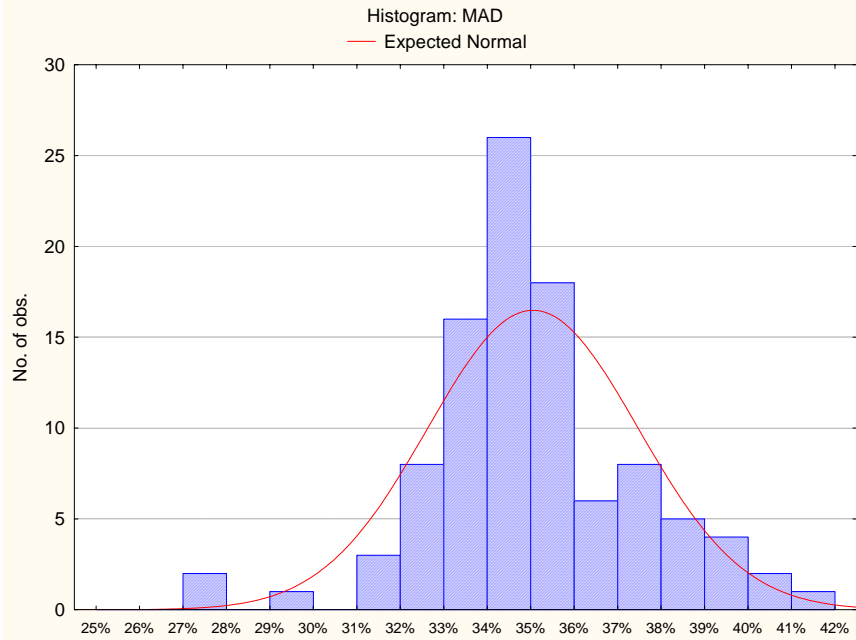
A histogram of the strategy returns is shown in the chart below, together with an overlaid plot of the loglogistic function.

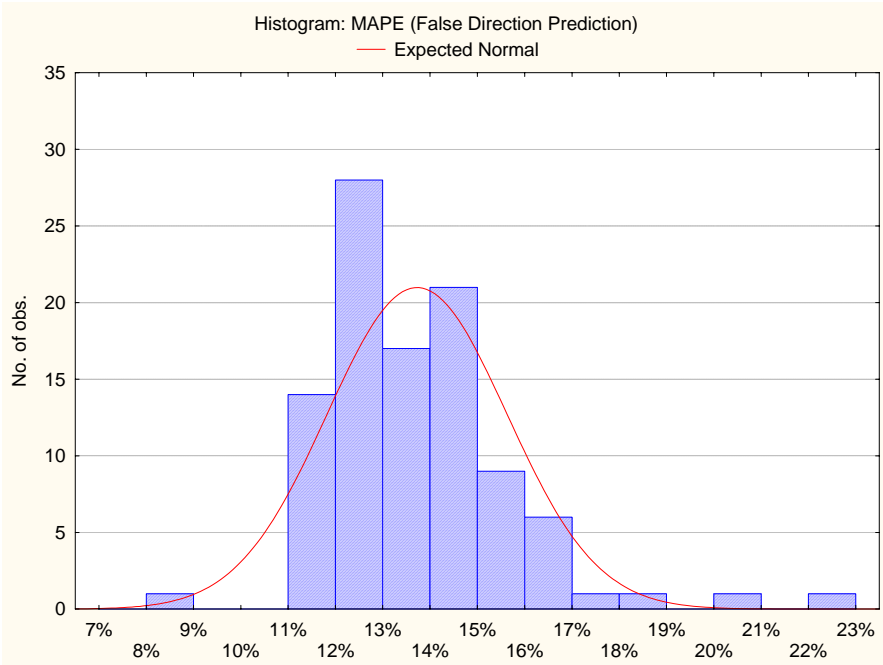
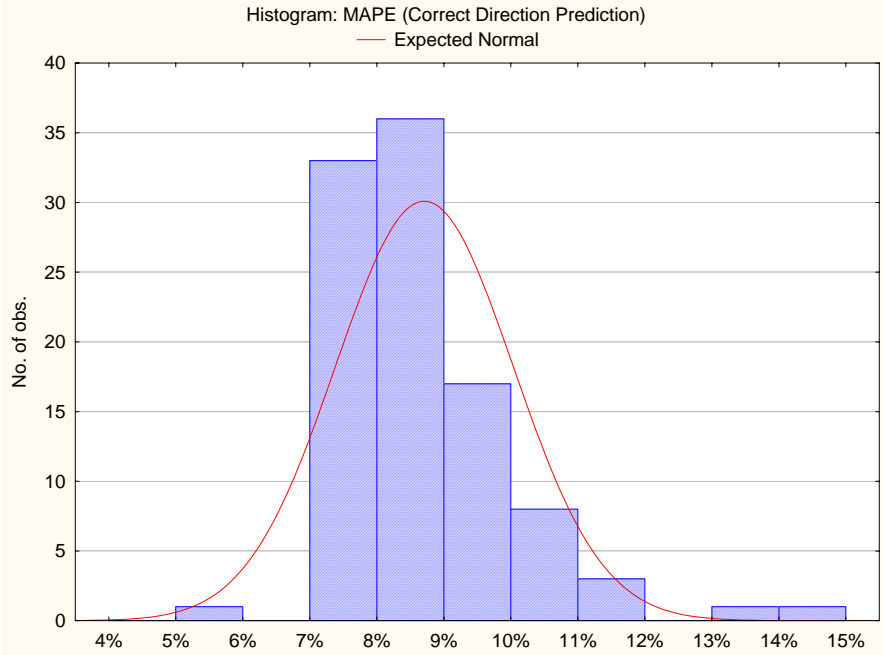


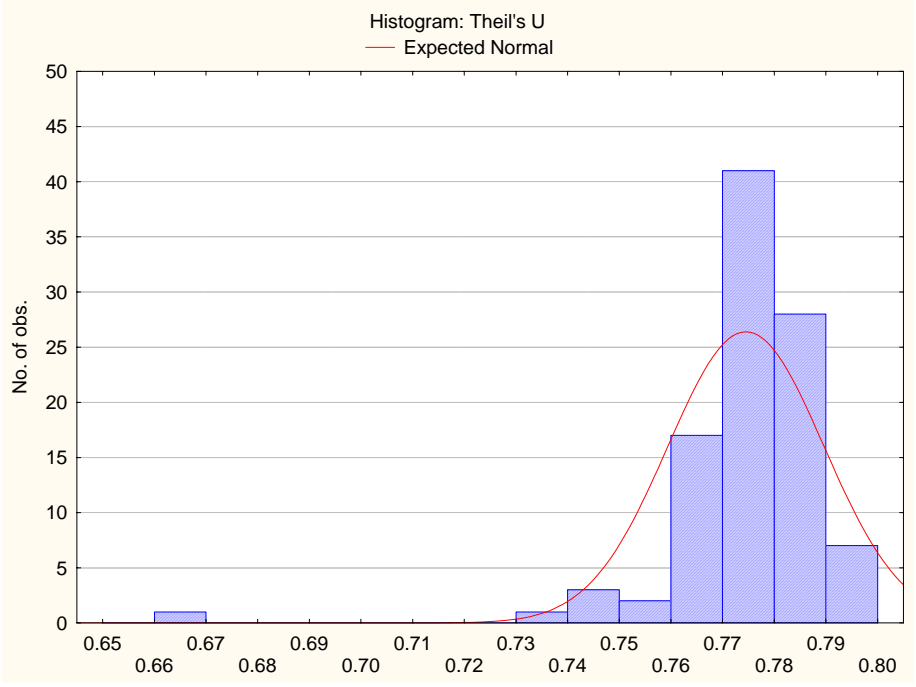
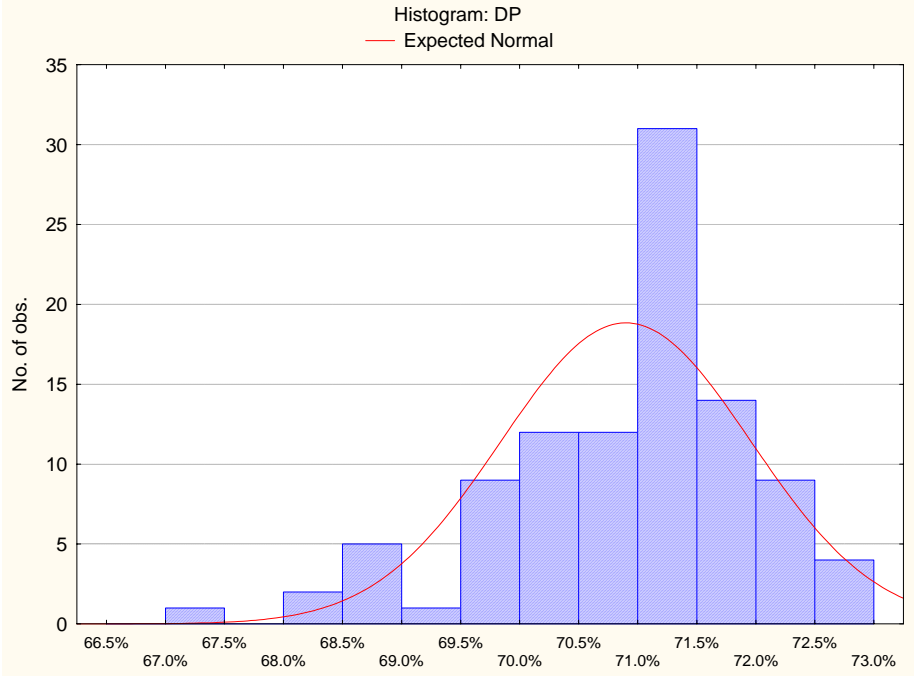
<sup>2</sup> Goodness of fit is measured in terms of Chi-Squared, Anderson-Darling and Kolmogorov-Smirnov test statistics

# Appendix 3 – Test Statistic Distributions









Scatterplot: MAPE vs. DP  
DP = .74326 - .3365 \* MAPE  
Correlation: r = -.5072

