The Best of Both Worlds:
A Hybrid Approach to Calculating Value at Risk

Jacob Boudoukh[1], Matthew Richardson, and Robert F. Whitelaw
Stern School of Business, NYU

The *hybrid approach* combines the two most popular approaches to VaR estimation: RiskMetrics and Historical Simulation.

It estimates the VaR of a portfolio by applying exponentially declining weights to past returns and then finding the appropriate percentile of this time-weighted empirical distribution.

This new approach is very simple to implement. Empirical tests show a significant improvement in the precision of VaR forecasts using the *hybrid approach* relative to these popular approaches.

[1] Contact person: The Salomon Center, Stern School of Business, 44 West 4th St., NYC, NY 10012, USA, or 34/19 Burla St., Tel Aviv 69364, Israel, email: *jboudouk@stern.nyu.edu*, phone: 00-972-52-817-244.

1. Introduction

Over the last few years Value at Risk (VaR) has gained recognition as the primary tool for market risk measurement in financial institutions. In fact, a large portion of these financial institutions report VaR measures to the public and/or to regulatory agencies. Regulatory agencies, in turn, allow the use of VaR for capital adequacy calculations. While there is widespread agreement on the use of VaR as a matter of principle, there is very little consensus on the preferred method for calculating VaR.

The difficulty in obtaining reliable VaR estimates stems from the fact that all extant methods involve some tradeoffs and simplifications. Determining what is the best methodology for VaR estimation becomes an empirical question and a question of implementation. It may depend on the number of underlying assets, their type, and the exact objective of the calculation. For instance, calculating the VaR of a derivative security whose value depends on a specific Brady bond index poses different difficulties than calculating the VaR of a well-diversified global equity portfolio. The Brady derivative involves a nonlinear instrument on a single, typically fat-tailed, return series. The equity portfolio involves multiple positions in multiple currencies, and hence numerous volatilities and correlations.

The challenge is to come up with a single, easily implementable approach that will provide an effective tradeoff. The hybrid approach presented here provides exactly that. It combines methodologies of two popular approaches to come up with an easy to implement methodology which inherits many of the advantages of each of these two approaches.

The two most widespread approaches to VaR estimation are exponential smoothing (EXP) (e.g., RiskMetrics) and historical simulation (HS). The EXP approach applies exponentially declining weights to past returns in order to calculate conditional volatilities. Using declining weights allows us to capture the cyclical behavior of return volatility. However, in order to calculate the VaR of a portfolio from its conditional volatility a conditional normality assumption must be invoked. Unfortunately, such an assumption seems to be at odds with financial data. Typically, financial data series seem to exhibit fat tails and skewness, properties which are very difficult to account for within the EXP approach.

The HS approach gets around the need to make distributional assumptions when inferring percentiles from volatility. This approach estimates percentiles directly, by using the empirical percentiles of the historical return distribution in order to obtain the VaR. Fat tails, skewness and any other property are accounted for directly. There are, however, two severe problems with the HS approach. First, extreme percentiles of the distribution (e.g., the 1% or 5% VaR) are notoriously hard to estimate with little data. Second, the HS approach essentially assumes that returns are independent and identically distributed (i.i.d.), and hence does not allow for time-varying volatility. One possible remedy to the first problem is to extend the observation period. Specifically, while the EXP approach effectively uses less than one year of daily return data, it is not uncommon to see the use of up to five years of historical data in applications of the HS approach. In doing so, however, the only remedy to the second problem (cyclical volatility) is lost. That is, the only way to put more weight on recent information within the HS approach is to use shorter historical windows. To sum up, with long histories the value of recent information diminishes, while with short histories we encounter estimation problems.

The hybrid approach combines the two approaches by estimating the percentiles of the return directly, using declining weights on past data. The approach starts with ordering the returns over the observation period just like the HS approach. While the HS approach attributes equal weights to each observation in building the conditional empirical distribution, the hybrid approach attributes exponentially declining weights to historical returns. Hence, while obtaining the 1% VaR using 250 daily returns involves identifying the third lowest observation in the HS approach, it may involve more or less observations in the hybrid approach. The exact number of observations will depend on whether the extreme low returns were observed recently or further in the past. The weighting scheme is similar to the one applied in the EXP approach.

Empirical results across different asset classes are very encouraging. Specifically, we study the proximity of the expected number of extreme returns to the actual number of extreme return across different methods. For example, in estimating the 1% VaR of the S&P500 the hybrid approach provides an absolute error which is 30% to 43% lower than the EXP approach and 14% to 28% lower than the HS approach. The results are similar or even stronger for a Brady bond index return series, for an oil price series, and an exchange rate series.[2]

We also devise a new benchmarking tool, which tests for the temporal unpredictability of extreme returns. While extreme returns will always occur, they should not “bunch up”. That is, a positive attribute of a VaR estimator is that once an increase in tail probability has occurred, the VaR forecast will increase in such a way that, given this new VaR estimate, the next large move has the same probability. This can be tested by examining the autocorrelation of the occurrence of tail events, which, under the null, should be zero. Our results indicate that the hybrid method provides the lowest level of autocorrelation across assets and parameters.

It is important to note that relative to extant approaches the improvements achieved by implementing the hybrid method come practically for free from a computational complexity, data intensity and programming difficulty standpoint.

2. Existing Approaches

For the purpose of our discussion we may start from the traditional flat-weight variance covariance approach. This approach calls for the calculation of variances and covariances across assets using a historical window of a certain length. The VaR of a portfolio of positions can be calculated by assuming multivariate normality, which, in turn, provides a full description of the

[2] This improvement is, in fact, an understatement of the true improvement, since an error in the 1% VaR should always occur even under the null.

portfolio return distribution, and hence the relevant percentiles. Given the return volatility we can calculate the 1% VaR by multiplying the volatility by 2.33, or the 5% VaR using a multiplier of 1.65.

The EXP approach, promoted by Morgan's RiskMetrics, makes the point that estimating conditional volatilities and correlations, which is precisely the task at hand, could benefit from giving greater weight to recent returns. This assertion comes from the recognition that financial data series exhibit time-varying and cyclical volatility. As such, placing more weight on recent observations while not ignoring distant observations represents a reasonable tradeoff between statistical precision and adaptiveness to recent news.

There are a number of disadvantages to the EXP approach. First, it is highly parametric, in that it assumes multivariate normality. Financial data series do not seem to fit this assumption, exhibiting fat-tails, skewness, unstable correlations, etc. As such, the method relies on a diversification effect, which exists only in large portfolios. The accuracy of this method becomes weaker when portfolios exhibit asset concentration. The second difficulty of the EXP method is in dealing with nonlinear positions (e.g., options and embedded options). Gamma and delta corrections are often incomplete solutions.

The portfolio aggregation (PA) approach (see JPMorgan's RiskMetrics Monitor for Q1 of 1997) is designed to get around the problem of nonlinear assets and unstable correlations, and is a simple variation on the EXP approach. The approach calls for the calculation of portfolio return volatility from simulated portfolio returns. Such returns are obtained by applying today's portfolio weights to historical returns. Declining weights are then applied to the simulated return information, similar to the EXP approach. This approach still suffers from the conditional normality assumption for the aggregate portfolio.

The HS approach avoids the parameterization problem entirely by letting the data dictate precisely the shape of the distribution. If we are interested in a particular percentile of a given return series, all we need to do is to sort the most recent K observations, and pick the corresponding observation as the VaR of the portfolio. If the VaR of a portfolio is desired, all we need to do is take the current portfolio composition, along with the last K returns, and generate K simulated portfolio returns (similar to the PA approach). These returns are the returns we would have observed if we had held today's portfolio over the past K trading days. This approach is completely nonparametric (aside from the choice of the observation window K). The idea is to let the data tell us about tail behavior, about conditional correlations, and about nonlinear asset behavior.

This method has a number of disadvantages as well. First, percentiles are notoriously difficult to estimate. This is especially true for extreme tails of the distribution. In estimating the 1% VaR with one year (250 days) of daily return data, we use only the third observation to determine the precise point of the tail (plus the fact that it is the third lowest of 250 observation). The EXP and the flat-weight approaches, in contrast, use every data point plus the normality assumption to come up with the VaR estimate. For this reason VaR estimates using the HS approach tend to be very "choppy" and volatile. The second shortcoming of this approach is that it uses flat weights over the most recent K observations. In doing so we lose the insight that the informativeness of historical data regarding the conditional distribution of current returns diminishes through time. Implicitly we assume that return series are i.i.d. For example, if three extreme negative returns are observed this week, the estimated 1% VaR may be large and *constant* for the entire next year (using K=250 days).

3. The Hybrid Approach

The hybrid approach is implemented in three steps:

Step 1: denote by $R(t)$ the realized return from t-1 to t. To each of the most recent K returns: $R(t)$, $R(t-1)$,...,$R(t-K+1)$, assign a weight $[(1-\lambda)/(1-\lambda^K)]$, $[(1-\lambda)/(1-\lambda^K)]\lambda$, ..., $[(1-\lambda)/(1-\lambda^K)]\lambda^{K-1}$, respectively. Note that the constant $[(1-\lambda)/(1-\lambda^K)]$ simply ensures that the weights sum to 1.

Step 2: order the returns in ascending order.

Step 3: in order to obtain the *x%* VaR of the portfolio, start from the lowest return and keep accumulating the weights until *x%* is reached. Linear interpolation is used between adjacent points to achieve exactly *x%* of the distribution.

Consider the following example, where we examine the VaR of a given series at a given point in time, and a month later, assuming that no extreme observations were realized during the month, and where the parameters are λ=0.98, K=100

Order	Return	Periods Ago	Weight	Cumul. Weight	Weight	Cumul. Weight
Initial Date:						
1	-3.30%	3	0.0221	0.0221	0.01	0.01
2	-2.90%	2	0.0226	0.0447	0.01	0.02
3	-2.70%	65	0.0063	0.0511	0.01	0.03
4	-2.50%	45	0.0095	0.0605	0.01	0.04
5	-2.40%	5	0.0213	0.0818	0.01	0.05
6	-2.30%	30	0.0128	0.0947	0.01	0.06
25 Days Later:						
1	-3.30%	28	0.0134	0.0134	0.01	0.01
2	-2.90%	27	0.0136	0.0270	0.01	0.02
3	-2.70%	90	0.0038	0.0308	0.01	0.03
4	-2.50%	70	0.0057	0.0365	0.01	0.04
5	-2.40%	30	0.0128	0.0494	0.01	0.05
6	-2.30%	55	0.0077	0.0571	0.01	0.06

The top half of the table shows the ordered returns at the initial date. Since we assume that over the course of a month no extreme returns are observed, the ordered returns 25 days later are the same. These returns are, however, further in the past. The HS approach, assuming an observation window of 100 days, estimates the 5% *VaR* to be 2.35% for both cases (note that VaR is the negative of the actual return). The hybrid approach initially estimates the 5% VaR as 2.63%. As time goes by and no large returns are observed, the VaR estimate smoothly declines to 2.34%.

There are a couple of technical issues that warrant some discussion. First, we assume that _ of a given return's weight is to the right and _ to the left of the actual observation. For example, the -2.40% return represents 1% of the distribution in the HS approach, and we assume that this weight is split evenly between the intervals from the actual observation to points halfway to the next highest and lowest observations. As a result, under the HS approach, -2.40% represents the 4.5th percentile, and the distribution of weight leads to the 2.35% VaR (halfway between 2.40% and 2.30%).

Second, some interpolation rule may be required in order to obtain a given percentile given two data points and their respective percentiles. For example, under the hybrid approach, the 5th percentile at the initial date lies somewhere between -2.70% and -2.60%. Using the allocation rule above, -2.70% represents the 4.78th percentile, and -2.60% represents the 5.11th percentile. We define the required VaR level as a linearly interpolated return, where the distance to the two adjacent cumulative weights determines the return. In this case, the 5% VaR is

$$2.70\%-(2.70\%-2.60\%)*[(0.05-0.0478)/(0.0511-0.0478))=2.63\%.$$

Similarly, the 5% VaR 25 days later is

$$2.35\%-(2.35\%-2.30\%)*[(0.05-0.0494)/(0.0533-0.0494)]=2.34\%.$$

4. An Empirical Comparison of Methods

In this section we first discuss the methodology and the relevant benchmarks for comparing the methods. We introduce a new simple test, and explain its relevance. We then perform an empirical analysis on a number of return series.

4.1 Comparison Methodology

The dynamic VaR estimation algorithm provides, for each of the three methods we examine here, an estimate of the *x%* VaR for the sample period. Therefore, the probability of observing a return lower than the calculated VaR should be *x%:*

$$Prob[R(t+1)<-VaR(t)]=x\%.$$

There are a few attributes which are desirable for *VaR(t).* We can think of an indicator variable *I(t)*, which is one if the VaR is exceeded, and zero otherwise. There is no direct way to observe whether our VaR estimate is precise; however, a number of different indirect measurements will, together, create a picture of its precision.

The first desirable attribute is *unbiasedness*. Specifically, we require that the VaR estimate be the *x%* tail. Put differently, we require that the average of the indicator variable *I(t)* should be *x%*:

$$avg[I(t)]=x\%$$

This attribute alone is an insufficient benchmark. To see this, consider the case of a VaR estimate which is constant through time, but is also highly precise unconditionally (i.e., achieves an average VaR probability which is close to *x%*). To the extent that tail probability is cyclical, the occurrences of violations of the VaR estimate will be "bunched up". This is a very undesirable property, since we require dynamic updating which is sensitive to market conditions.

Consequently, the second attribute which we require of a VaR estimate is that extreme events do not "bunch up". Put differently, a VaR estimate should increase as the tail of the distribution rises. If a large return is observed today, the VaR should rise to make the probability of another tail event exactly *x%* tomorrow. In terms of the indicator variable, *I(t)*, we essentially require that *I(t)* be i.i.d.. This requirement is similar to saying that the VaR estimate should provide a filter to transform a serially dependent return volatility and tail probability into a serially independent *I(t)* series.

The simplest way to assess the extent of independence here is to examine the empirical properties of the tail event occurrences, and compare them to the theoretical ones. Under the null that *I(t)* is independent over time,

$$corr[I(t-s)*I(t)]=0 \qquad s\,,$$

i.e., the indicator variable should not be autocorrelated at any lag. Since the tail probabilities that are of interest tend to be small, it is very difficult to make a distinction between pure luck and persistent error in the above test for any individual correlation. Consequently, we consider a joint test of whether the first 5 daily autocorrelations (approximately 1 week) are equal to zero.

Note that for both measurements the desire is essentially to put all data periods "on an equal footing" in terms of the tail probability. As such, when we examine a number of data series for a given method, we can aggregate across data series, and provide an average estimate of the unbiasedness and the independence of the tail event probabilities. While the data series may be correlated, such an aggregate improves our statistical power.

The third property which we examine is related to the first property -- the biasedness of the VaR series, and the second property -- the autocorrelation of tail events. We calculate a rolling measure of the absolute percentage error. Specifically, for any given period, we look forward in time 100 periods and ask how many tail events were realized? If the indicator variable is both unbiased and independent, this number is supposed to be the VaR's percentage level, namely *x*. We calculate the average absolute value of the difference between the actual number of tail events and the expected number across all 100-period windows within the sample. Smaller deviations from the expected value indicate better VaR measures.

4.2 Data and Results

The data we use includes a number of series, chosen as a representative set of "interesting" economic series. These series are interesting since we a priori believe that their high order moments (skewness and kurtosis) and, in particular, their tail behavior, pose different degrees of challenge in VaR estimation. The series span the period from 1/1/1991 to 5/12/1997, and include data on the following

- DEM: The dollar/DM exchange rate
- OIL: The spot price for brent crude oil
- S&P: The S&P 500 Index
- BRD: A general Brady bond index (JP Morgan Brady Broad Index)

We have 1663 daily continuously compounded returns for each series.

In the following tables, in addition to reporting summary statistics for the four series, we also report results for

- EQW: an equally weighted portfolio of the four return series
- AVG: statistics for tail events average across the four series

The EQW results will give us an idea of how the methods perform when tail events are somewhat diversified (via aggregation). The AVG portfolio simply helps us increase the effective length of our sample. That is, correlation aside, the AVG statistics may be viewed as using four times more data. Its statistics are therefore more reliable, and provide a more complete picture for general risk management purposes. Therefore, in what follows, we shall refer primarily to AVG statistics, which includes 6656 observations.

In the following tables we use a 250 trading day window throughout. This is, of course, an arbitrary choice, which we make in order to keep the tables short and informative. The statistics for each of the series include 1413 returns, since 250 observations are used as back data. The AVG statistics consist of 5652 data points, with 282 tail events expected in the 5% tail, and 56.5 in the 1% tail.

In Table 1 we document the percent of tail events for the 5% and the 1% VaR. There is no apparent strong bias for the 5% VaR. The realized average varies across methods, between 4.62% and 5.65%. A bias is observed, however, when examining the empirical performance for the 1% VaR across methods. Overall, there seems to be a bias in favor of the nonparametric methods, namely HS and Hybrid. These methods, by design, are better suited to addressing the well known tendency of financial return series to be fat-tailed. Since the estimation of the 1% tail requires a lot of data, there seems to be an expected advantage to high smoothers within the hybrid method.

Table 1. Percent in the VaR Tail

5% Tail	Historical	Historical	EXP		Hybrid	
	STD	Simulat.	0.97	0.99	0.97	0.99
DEM	5.18	5.32	5.74	5.18	5.25	5.04

OIL	5.18	4.96	5.60	5.39	5.18	5.18
S&P	4.26	5.46	4.68	4.18	6.17	5.46
BRD	4.11	5.32	4.47	4.40	5.96	5.46
EQW	4.40	4.96	5.04	4.26	5.67	5.39
AVG	**4.62**	**5.21**	**5.11**	**4.68**	**5.65**	**5.30**
1% Tail	**Historical**	**Historical**	**EXP**		**Hybrid**	
	STD	**Simulat.**	**0.97**	**0.99**	**0.97**	**0.99**
DEM	1.84	1.06	2.20	1.63	1.84	1.28
OIL	1.84	1.13	1.77	1.77	1.70	1.35
S&P	2.06	1.28	2.20	2.13	1.84	1.42
BRD	2.48	1.35	2.70	2.41	1.63	1.35
EQW	1.63	1.49	1.42	1.42	1.63	1.21
AVG	**1.97**	**1.26**	**2.06**	**1.87**	**1.73**	**1.32**

In Table 2 we document the mean absolute error (MAE) of the VaR series. The MAE is a conditional version of the previous statistic (percent in the tail from Table 1). The MAE uses a rolling 100-period window. There is, again, an advantage in favor of the nonparametric methods, HS and Hybrid, with the hybrid method performing best for high smoothers. Since a statistical error is inherent in this statistic, we can not possibly expect a mean absolute error of zero. As such, the 38% improvement of the hybrid method with lambda of 0.99 (with MAE of 0.90% for the AVG series' 1% tail) relative to the EXP method with the same smoother (with MAE of 1.45), is an understatement of the level of improvement. A more detailed simulation exercise would be needed in order to determine how large this improvement is. It is worthwhile to note that this improvement is achieved very persistently across the different data series.

Table 2. Rolling Mean Absolute Percentage Error of VaR

5% Tail	**Historical**	**Historical**	**EXP**		**Hybrid**	
	STD	**Simulat.**	**0.97**	**0.99**	**0.97**	**0.99**
DEM	2.42	2.42	1.58	2.11	1.08	1.77
OIL	2.84	2.62	2.36	2.67	1.93	2.44
S&P	1.95	1.91	1.52	1.85	1.72	1.68
BRD	3.41	3.53	3.01	3.34	2.54	2.97
EQW	2.43	2.36	2.48	2.33	1.50	2.20
AVG	**2.61**	**2.57**	**2.19**	**2.46**	**1.76**	**2.21**
1% Tail	**Historical**	**Historical**	**EXP**		**Hybrid**	
	STD	**Simulat.**	**0.97**	**0.99**	**0.97**	**0.99**
DEM	1.29	0.87	1.50	1.12	1.02	0.88
OIL	1.71	0.96	1.07	1.39	0.84	0.80
S&P	1.45	1.14	1.40	1.42	0.99	0.82
BRD	2.15	1.32	1.98	2.06	1.03	1.12
EQW	1.57	1.52	1.25	1.25	0.72	0.87
AVG	**1.63**	**1.16**	**1.44**	**1.45**	**0.92**	**0.90**

The adaptability of a VaR method is one of the most critical elements in determining the best way to measure VaR. When a large return is observed, the VaR level should increase. It should increase, however, in such a way which will make the next tail event's probability precisely *x%*. We can therefore expect these tail event realizations to be i.i.d. events with x% probability. This i.i.d.-ness can be examined using the autocorrelation of these tail events, with the null being that it is zero. As we see below, the hybrid method's autocorrelation for the AVG series is closest to zero, and so is also the case for the majority of the different series. Interestingly, this is especially true for the more fat-tailed series, BRD and OIL. As such, the hybrid method is very well suited for fat-tailed, possibly skewed series.

Table 3. First Order Autocorrelation of the Tail Events

5% Tail	Historical	Historical	EXP		Hybrid	
	STD	Simulat.	0.97	0.99	0.97	0.99
DEM	0.39	0.09	-2.11	-1.06	-2.63	-2.28
OIL	1.76	2.29	2.11	1.25	3.20	0.31
S&P	0.77	1.09	-0.15	0.94	0.77	2.46
BRD	11.89	12.69	13.60	12.27	10.12	12.08
EQW	5.52	2.29	3.59	4.26	-2.04	-0.14
AVG	**4.07**	**3.69**	**3.41**	**3.53**	**1.88**	**2.49**
1% Tail	**Historical**	**Historical**	**EXP**		**Hybrid**	
	STD	Simulat.	0.97	0.99	0.97	0.99
DEM	2.04	-1.08	1.05	2.76	-1.88	-1.29
OIL	-1.88	-1.15	2.27	2.27	-1.73	-1.37
S&P	4.94	9.96	7.65	8.04	2.04	8.70
BRD	15.03	9.30	10.75	12.60	-1.66	3.97
EQW	2.76	3.32	3.63	3.63	2.76	4.73
AVG	**4.58**	**4.07**	**5.07**	**5.86**	**-0.09**	**2.95**

In tables 4 and 5 we test the statistical significance of the autocorrelations we saw above. Specifically, we examine the first though fifth autocorrelations of the tail event series, with the null being that all of these autocorrelations should be zero. The test statistic is simply the sum of the squared autocorrelations, appropriately adjusted to the sample size. Under the null this statistic is distributed as $\chi^2(5)$. These test statistics are generally lower for the hybrid method relative to the EXP. For the specific series four rejections out of possible eight are obtained with the hybrid method, relative to seven out of eight for the EXP method.

Table 4. Test statistic for independence (autocorrelations 1-5)

5% Tail	Historical	Historical	EXP		Hybrid	
	STD	Simulat.	0.97	0.99	0.97	0.99
DEM	7.49	10.26	3.80	8.82	3.73	6.69
OIL	9.58	12.69	5.82	4.90	4.71	3.94
S&P	8.09	8.32	0.88	4.31	0.81	3.87
BRD	66.96	87.80	88.30	78.00	46.79	69.29
EQW	16.80	6.30	11.66	14.75	4.87	12.10
AVG	**21.78**	**25.07**	**22.09**	**22.16**	**12.18**	**19.18**
1% Tail	**Historical**	**Historical**	**EXP**		**Hybrid**	
	STD	Simulat.	0.97	0.99	0.97	0.99
DEM	3.34	5.33	4.56	4.39	7.58	3.83
OIL	33.98	8.29	3.82	18.89	8.53	3.54
S&P	14.67	36.15	22.68	25.18	3.26	24.10
BRD	88.09	29.37	41.60	82.77	11.26	11.36
EQW	41.55	14.69	16.85	16.85	5.08	13.05
AVG	**36.32**	**18.77**	**17.90**	**29.61**	**7.14**	**11.18**

Table 5. P-value for independence (autocorrelations 1-5)

5% Tail	Historical	Historical	EXP		Hybrid	
	STD	Simulat.	0.97	0.99	0.97	0.99
DEM	0.19	0.07	0.58	0.12	0.59	0.24
OIL	0.09	0.03	0.32	0.43	0.45	0.56
S&P	0.15	0.14	0.97	0.51	0.98	0.57
BRD	0.00	0.00	0.00	0.00	0.00	0.00
EQW	0.00	0.28	0.04	0.01	0.43	0.03
AVG	**0.09**	**0.10**	**0.38**	**0.21**	**0.49**	**0.28**

1% Tail	Historical	Historical	EXP		Hybrid	
	STD	Simulat.	0.97	0.99	0.97	0.99
DEM	0.65	0.38	0.47	0.49	0.18	0.57
OIL	0.00	0.14	0.58	0.00	0.13	0.62
S&P	0.01	0.00	0.00	0.00	0.66	0.00
BRD	0.00	0.00	0.00	0.00	0.05	0.04
EQW	0.00	0.01	0.00	0.00	0.41	0.02
AVG	**0.13**	**0.11**	**0.21**	**0.10**	**0.28**	**0.25**

5. Conclusion

In this paper we put forth a new method for estimating VaR. the method is easy and computationally cheap to implement. The hybrid method combines some of the benefits of extant methods for VaR estimation, such as giving less weight to distant observations (as in the EXP method) and accounting for the empirical tails skewness and other properties of the observed data series (as in the HS method). Empirical results show significant improvement in statistical performance over the two competing methods, EXP and HS. This improvement is most pronounced for fat-tailed data series, in our case the oil price and Brady index series.