Return or Position-based Value at Risk?

Jérôme Huillard d'Aignaux, Ph.D.*

Benjamin Baranne, M.Sc.

Jürgen Fuchs, Ph.D.

Stavros Siokos, Ph.D.

Sciens Capital Limited

4th Floor

25 Berkeley Square

London W1J 6HN

Phone: +44 (0) 20 7078 0600

Fax: +44 (0) 20 7078 0611

email: jhuillard@sciensam.com

* Corresponding author

SUBMITTED FOR PUBLICATION

Abstract

The emergence of managed account platforms for investing in hedge funds has given the investor access to a range of new risk measures. This study analyses the accuracy of various Value at Risk (VaR) methodologies in the context of hedge fund investing. We found that in our sample of data, position-based VaR provides the best risk assessment measure for short horizons whereas for longer term horizons, VaR can sometimes be inaccurate. Our interpretation is that despite being based on actual underlying positions, position-based VaR is not designed to capture the dynamic of hedge fund portfolios. Over longer term horizons, VaR calculated using historic returns still provides a rough but reliable risk indicator for the investor. Our study also highlights the importance of regular monitoring of any risk model's accuracy in its specific context.

1 Introduction

Value at Risk (VaR) is the most commonly used measure in risk assessment. It is meant to be the answer to a simple question: how much can an investment depreciate with a given probability over a pre-set horizon? What makes VaR attractive is its ability to aggregate several risk dimensions into one number that is easily interpreted and understood. In the context of hedge fund investing, the use of VaR poses a certain number of issues. There is obviously the debate about technicalities and assumptions made for the statistical modelling of the data (as with any other assets). These debates sometimes leave aside a major specificity of hedge funds, namely that a hedge fund is not a static asset class [1]. Hedge funds have evolved significantly over the years, but their main feature is their dynamic nature. The rate of change of hedge fund's underlying investment depends on the strategy employed and the market environment. In the early years of the alternative industry, transparency was very limited. Typically, the only available data were the historic returns of the manager/strategy. The VaR methodologies were, therefore, limited to return-based analysis. With the emergence of managed accounts where investors can gain full transparency on the underlying investments, the risk manager is given more options. Position-based VaR seems *a-priori* a more accurate way of measuring the overall risk of a portfolio [2]. However, given the potential dynamic nature of hedge funds, what is the benefit of using position data knowing that the portfolio that is being analysed may change completely within the time horizon that is being considered? In this paper, we examine the accuracy of various VaR methodologies when compared with a benchmark VaR computed using daily position data. The daily position and historic return data are taken from a number of hedge funds available on the Sciens Capital managed account platform.

2 Data and methodology

2.1 What is a managed account?

Most hedge fund managers are reluctant to disclose their positions and the information they provide to investor tends to be opaque. Transparency has been a long term issue for investment managers and, particularly, for risk managers assessing the risk of hedge fund investing. One answer to improving the framework of hedge fund investment has been the development of managed accounts. A managed account is an investment vehicle that can be used to replicate any alternative strategy. The hedge fund manager becomes the account trading advisor and replicates his investment book on the account. However, he no longer has entire control. The investor himself or a third party becomes the legal fund manager and, in this way, has full transparency of the current portfolio exposure. For many investors, though, setting up a trading account can be time-consuming and economically non viable. The answer to this has been the emergence of managed account platforms which undertake to set up the structure and manage the day-to-day operations. A platform can also offer the investor a range of investment strategies and additional services such as monitoring specific investment restrictions, assuring compliance with particular regulatory constraints and initiating bespoke reporting requirements. In the area of risk management, this transparency feature is of paramount interest only if the platform provider is able to translate the complex trades into consolidated reports. It is the collation of the aggregated data on a daily basis into portfolio exposures that is truly valuable, as few investors have the resources to administer hedge fund investments themselves. From the platform provider, the investor can obtain information from both a return-based and position-based perspective. The objective of this paper is to determine to what extent both measures are useful.

2.2 Data sources

The data used in this study come from different sources. We have obtained daily position-based VaR on fourteen accounts over a period of sixteen months between October 2010 and February 2012. This data is calculated daily by RiskMetrics and is based on the end-of-day position-level data available to the platform operator. Position-based VaR is calculated using RiskMetrics' 2006 methodology (RM 2006). It is based on a Monte Carlo simulation of all underlying risk factors in the portfolio with full revaluation of all securities. It is assumed that returns follow a student's T distribution with five degrees of freedom and that the ex ante volatility of the portfolio can be modelled by a long memory ARCH process. The position data are obtained from various brokerage counterparties contracted to the platform. Daily positions are thoroughly verified and checked before being processed by RiskMetrics. The accounts' Net Asset Values (NAVs) are obtained from the platform operator and are calculated daily by a third party valuation agent. For each account on the Sciens platform, the historic monthly track record of the manager (benchmark fund) is obtained directly from the trading advisor. Funds for which the account settings implied a significant drift from the trading advisor's original strategy (e.g. as a result of investment restrictions or change of mandate during the life of the account) were not included in the analysis.

2.3 VaR methodologies

A number of models exist for estimating VaR and each model has its own set of assumptions. When historical returns are the only data available, the choice is limited to the assumption made about the shape of the probability density function of the return distribution. The most common assumption is that historical returns follow a normal distribution with a constant mean μ and variance σ^2 estimated using the historic return from the sample. This leads to the common Gaussian VaR which is simply calculated as the α-quantile of the normal distribution with mean $\hat{\mu}$ and variance $\hat{\sigma}^2$. However, the assumption of normality has been contradicted by a number of empirical studies [3]. The return distribution of hedge funds typically exhibits non-normal patterns, which invalidates the Gaussian approximation. Several alternatives have been proposed. In this study, we used the

Cornish-Fisher expansion of the VaR because of its relative simplicity over more advanced and complex methods [4]. The Cornish-Fisher methodology takes into account two particular characteristics of hedge fund returns: i) the asymmetry: more frequent returns above the mode than below or vice versa and, ii) the occurrence of extreme return values significantly more or less prevalent than expected in a normal distribution. Statistically, these two characteristics are measured by the skewness and kurtosis of the distribution. The calculation is not radically different from the Gaussian method. However, the quantile used in this calculation is adjusted by the third and fourth moments of the distribution. The calculation is based on mathematical properties of quantiles of any random variable which can be approximated by its first moments. The quantile is calculated as $y_\alpha = \mu + \sigma x_\alpha$ where μ and σ are the mean and standard deviation of the return distribution (α is the quantile being considered - typically 1% or 5% for a 99% or 95% VaR). The term x_α is calculated as

$x_\alpha = z_\alpha + \frac{1}{6}(z_\alpha^2 - 1)\gamma_1 + \frac{1}{24}(z_\alpha^3 - 3z_\alpha)\gamma_2 - \frac{1}{36}(2z_\alpha^3 - 5z_\alpha)\gamma_1^2$

where γ_1 and γ_2 are the skewness and kurtosis of the return distribution and z_α is the α-quantile of the standardised normal distribution. The time horizon of the calculated VaR is determined by the original frequency of the data used in the calculation. When using the mean, volatility, skewness and kurtosis of a monthly return time series, the resulting VaR is for a one month horizon. However, the value can be transformed into a VaR for an any time horizon by multiplying the original VaR by the factor $\sqrt{\frac{T_{dest}}{T_{source}}}$, where T_{source} is the original frequency of the VaR and T_{dest} is the frequency of the desired time horizon. For instance, to convert a one month VaR into a daily VaR, the factor to use is $\sqrt{\frac{1}{20}}$ (we assume one month is equivalent to 20 trading days).

2.4 VaR backtesting

The evaluation of VaR accuracy is far less controversial. Given a α-VaR value, one needs to evaluate the frequency of VaR breaches (how many times the actual return was below the predicted VaR level). By definition, the expected value of the frequency of α-VaR breach events is α: on average, we expect the 95% daily VaR to be breached once in twenty trading days. If there are significantly less than the expected VaR breaches, the VaR is overly conservative. If there are more, the proposed VaR model underestimates the portfolio's actual level of risk. This analysis is referred to as the unconditional coverage property: the probability of realising a loss in excess of the reported α VaR must be precisely α. Several statistical tests can be used to formally examine if this condition is verified [5] [6]. In this study, we used Kupiec's test [7] to assess the statistical significance of the estimated unconditional coverage of the different VaR methods. The test derives from the assumption that the number of α-VaR breaches follows a binomial distribution with parameter α. It is a straightforward log-likelihood ratio test that is expressed as:

$$POF\left(N, N_{\alpha}\right) = -2\log\left(\left(\frac{1-\hat{\alpha}}{1-\alpha}\right)^{N-N_{\alpha}}\left(\frac{\hat{\alpha}}{\alpha}\right)^{N_{\alpha}}\right)$$

where N is the sample size, N_{α} is the number of exceptions, $\hat{\alpha}$ is the sample frequency of α-VaR breaches ($\hat{\alpha} = N_{\alpha}/N$) and α is the VaR level being considered (e.g. 1%). Under the null hypothesis (alpha $= \hat{\alpha}$), this statistic follows a χ^2 distribution with one degree of freedom. One of the commonly cited limitations of the above test is that it is undefined when $\hat{\alpha} = 0$. We suggest that the test can be extended to the situation where no VaR breaches are observed in the sample (e.g.

$\hat{\alpha} = 0$). [1] In this case, the log-likelihood ratio test expression is reduced to:

$$POF_0\left(N,0\right) = -2N\log\left(1-\alpha\right)$$

The second expected property of an accurate VaR model is the statistical independence: the occurrence of an event should not be related to the occurrence of any other event in the sample. For instance, exceptions have to be randomly spread over time. Clustering of exceptions would indicate some form of failure to capture or adapt to a rapid change in a market environment and some form of time dependency that is not desirable. Independence is a more difficult subject for a statistician. The difficulty lies with the limitation of the hypothesis testing framework which requires distributional assumption of the data under the alternative hypothesis. While there are a number of ways statistical independence can be violated, there are also as many (e.g. infinite) possible statistical tests for this independence. For practical purposes, we are mostly interested in time clustering and several statistical tests have been suggested [8]. The Christoffersen contribution is aimed at reducing the problem of determining the accuracy of a VaR model in an analysis of whether or not hit sequences (VaR violations) satisfy the unconditional coverage and the independence property [5]. Christoffersen's Markov test of independence property examines whether the likelihood of a VaR violation depends on whether or not a VaR violation has happened the day before. According to the independence property, a violation that has happened in the past may not give an indication of whether an additional VaR violation may occur in the future.

[1] When no failure is being observed in the sample, the log-likelihood of the observation is simply proportional to $N\log\left(1-p\right)$ which is maximum for $p = \hat{\alpha} = 0$ (i.e our best estimate of the actual failure rate is 0). Under the null hypothesis that $p = \alpha$, the likelihood is $N\log\left(1-\alpha\right)$. The log-likelihood ratio test therefore simplifies into:
$2N\log\left(1-\hat{\alpha}\right) - 2N\log\left(1-\alpha\right) = 2N\log\left(1\right) - 2N\log\left(1-\alpha\right) = -2N\log\left(1-\alpha\right)$.

The test supposes that the hit sequence follows a first-order Markov sequence with a switching probability matrix. Consider a binary first-order Markov chain with a transition probability matrix:

$$\Pi = \begin{pmatrix} 1-\pi_{01} & \pi_{01} \\ 1-\pi_{11} & \pi_{11} \end{pmatrix}$$

where π_{ij} is the probability of an i on day t - 1 being followed by a j on day t. The test of independence (ind) is then $H_{0,ind} : \pi_{01} = \pi_{11}$. The approximate likelihood function under H_0 is:

$$\pi_2 = \frac{(n_{01}+n_{11})}{n}$$
$$LRH_0 = (1-\pi_2)^{(n_{00}+n_{10})}\pi_2^{(n_{01}+n_{11})}$$

where n_{ij} is the number of observations with value i followed by j. The approximate likelihood function under H_1 is:

$$\pi_{01} = \frac{n_{01}}{n_{00}+n_{01}} \; \pi_{11} = \frac{n_{11}}{n_{10}+n_{11}}$$
$$LRH_1 = (1-\pi_{01})^{n_{00}}\pi_{01}^{n_{01}}(1-\pi_{11})^{n_{10}}\pi_{11}^{n_{11}}$$

The likelihood ratio test of independence is asymptotically distributed as a χ^2 with one degree of freedom:

$$LR_{ind} = -2(log(LRH_0) - log(LRH_1))$$

The idea is that clustered violations point out a risk model inaccuracy. In our study, the problem is compounded by the fact that we have considered rolling periods for VaR horizons of five and twenty days. For these horizons, we would expect pairs of VaR violations to be naturally clustered when using overlapping rolling windows - even if the data are independent. The proposed test of independence would not be able to determine to what extent observed clustering is purely random or the result of an issue with the VaR model. To address this, we have calculated the distribution of the frequency of pairs of VaR violations in consecutive overlapping time windows under the null hypothesis by simulation. For each simulation, we have used an independent random sample of a normal deviate and calculated the number of consecutive VaR violations in periods of five and twenty days. Each simulation was repeated ten thousand times in order to obtain the probability distribution of consecutive VaR breaches. These values were used in our test of independence. A violation of the independence property will highlight the fact that VaR violations happen in pairs and that the VaR measure is slow to adjust to market changes. Both the unconditional coverage and independence property analyse a dimension of an accurate VaR model.

The self evaluation of VaR models is a routine part of regulatory compliance for banks and financial institutions. The Basel Committee (1996) [9] set up a regulatory backtesting framework consisting of comparing the last 250 daily 99% VaR estimates with corresponding daily trading outcomes. In essence, this is a practical implementation of Kupiec's test for unconditional coverage. The committee classifies backtesting outcomes into three categories: green, yellow and red which roughly correspond to critical values of Kupiec's test (0.05 and 0.0001). VaR models are classified into colour categories according to the observed number of exceptions.

3 Results

An example of the different VaR models is illustrated on Fig 1. It shows three time series corresponding to the three VaR methods that have been used in this study for one randomly selected account. The position VaR varies significantly in the period with occasional peaks. By contrast, historic return-based VaRs are very smooth, with occasional jumps. This is to be expected since the distribution of historic returns is stable, relative to the underlying position in the account which can change daily. The monthly return VaR appears to be lower (in absolute value) than both the position and historic daily return VaRs. This is because, in this example, the historic monthly track record had a lower historic volatility than that observed during the study period. This obviously varies from one account to the other.

Unconditional coverage analysis was carried out on rolling windows by comparing the following n-day actual return (n ranging from one to twenty trading days) to the VaR available on the first day of the period. VaR values were adjusted for the corresponding time horizon (see Methodology). Results are presented in Table 1. For a one day horizon, the accuracy of the position-based VaR was perfectly satisfactory with a 1.1% failure rate. This compares with 1.9% and 1.4% failure rates for the monthly return and daily return-based VaR respectively. As explained above, if the VaR model is adequate, we expect the failure rate to be 1% plus some sample variation. Results by strategy give more insight into how the different VaR models perform. In the diversified managers group (two funds), we found that the daily position VaR significantly underestimated the risk, with a near 3% failure rate, but was quite satisfactory for other strategies. For the one day horizon, the

Figure 1: Example of VaR Time Series

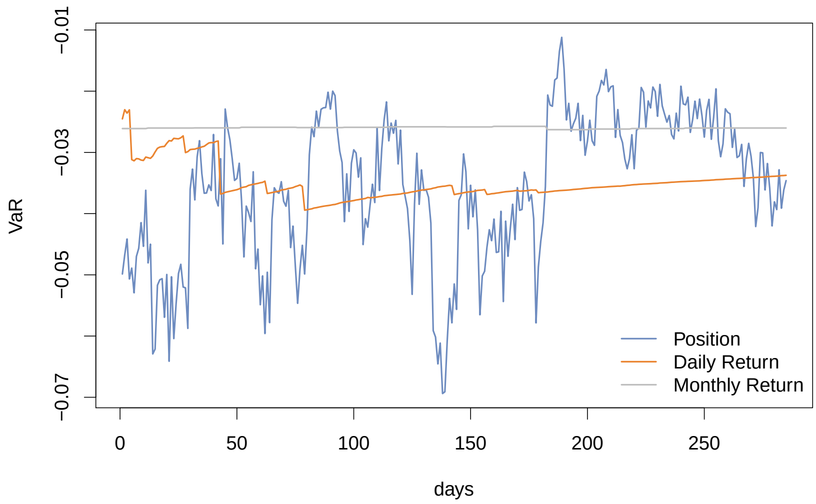

accuracy of return-based VaR is mixed, ranging from a too conservative model (failure rate of 0% for relative value) to a very serious underestimation of the risk (3% failure rate for managed futures).

Table 1: VaR Models Unconditional Properties by Strategy and Time Horizon

Strategy	Number of Managers	Value at Risk Model: 99% Position VaR (Daily)			99% Historic VaR (Daily)			99% Historic VaR (Monthly)		
		1 Day	7 Days	30 Days	1 Day	5 Days	20 Days	1 Day	5 Days	20 Days
Diversified	1	3.6%	3.1%	0.7%	2.4%	3.1%	0.0%	1.8%	0.6%	0.0%
Equity-Variable Bias	2	0.7%	2.0%	3.2%	1.6%	3.2%	4.7%	1.0%	1.6%	3.2%
Macro	4	1.3%	2.1%	3.7%	1.1%	1.0%	0.5%	1.3%	1.0%	0.1%
Managed Futures	6	0.7%	1.5%	0.9%	1.3%	1.2%	0.2%	2.9%	4.0%	2.4%
Relative Value	1	1.4%	1.1%	0.0%	2.4%	3.2%	1.5%	0.0%	0.0%	0.0%
Grand Total	14	1.1%	1.8%	2.0%	1.4%	1.7%	1.1%	1.9%	2.3%	1.6%

Table 2: Test of Independence, p-values by Strategy and Time Horizon

	Value at Risk Model								
Strategy	99% Position VaR (Daily)			99% Historic VaR (Daily)			99% Historic VaR (Monthly)		
	1 Day	5 Days	20 Days	1 Day	5 Days	20 Days	1 Day	5 Days	20 Days
Diversified	0.093	<0.001	0.873	0.093	<0.001	0.873	0.093	0.888	0.873
Equity-Variable Bias	0.093	0.008	<0.001	0.093	<0.001	<0.001	0.093	0.008	<0.001
Macro	0.093	0.173	0.002	0.093	0.463	0.669	0.093	0.463	0.873
Managed Futures	0.093	0.341	0.571	0.093	0.463	0.873	<0.001	<0.001	0.022
Relative Value	0.093	0.888	0.873	<0.001	0.003	0.571	0.093	0.888	0.873
All	0.093	0.173	0.071	0.093	0.116	0.274	0.005	0.004	0.091

Table 3: Independence Test Results by Time Horizon

VaR methodology	1 Day	5 Days	20 Days
Position	0.093	0.173	0.071
Historic Daily	0.093	0.116	0.274
Historic Monthly	0.005	0.004	0.091

Approximate p-values of the test of independence are shown on Tables 2 and 3. Position and daily return VaR violations appear to be independent from each other on a one day horizon, as indicated by p-value greater than 0.01. When using a monthly return VaR, there is a significantly high number of pairs of successive VaR violations. On the five and twenty day horizons, there is again no evidence of particular clustering when using both the position and daily return VaRs. Monthly return VaR is the worst affected, particularly on a short term horizon, but appears acceptable on a twenty day horizon. As with the coverage test, there are large discrepancies in model adequacy between strategies. For equity-variable managers, we found relatively strong evidence of time dependency in the VaR violation process in all three methods. Results for longer time horizons lead to more surprising results. Overall VaR violations for up to twenty days horizon (e.g. one month) are plotted on Fig 2. The accuracy of the position VaR appears to deteriorate as the horizon increases (with the failure rate stabilising around 2%). But, interestingly, return-based VaR accuracy appears to improve as the horizon increases. The daily return VaR appears to be very well calibrated (with the failure rate stable around 1% from ten to twenty days horizon). The accuracy of the monthly return VaR improves markedly and performs as well as, if not better, than both

Figure 2: Frequency or VaR Breaches by Time Horizon

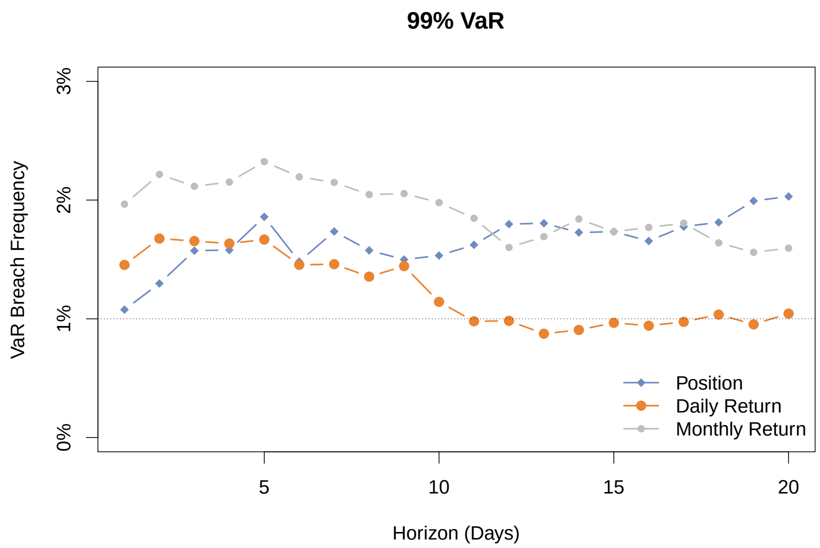

the position-based VaR for horizons of ten days and above. The analysis by time horizon and strategy is difficult to interpret. First, the number of observations - and therefore the statistical power - is obviously reduced. A failure rate of 0% - when 1% is expected - is not completely surprising when the total number of observations is one hundred (in fact, the p-value of the coverage test is 0.15 and the null hypothesis would not be rejected in this case). There are large variations between strategies regarding the change of coverage with a time horizon. The position-based VaR appeared to be over conservative for diversified and relative value and to significantly underestimate the risk in equity variable bias and macro managers. Return-based VaR showed the same trend except in macro where there is a divergence in the VaR model's performance over a longer time horizon. However, given the low number of observations, these results should be considered with caution.

4 Discussion

We are conscious that our data sample has a number of biases. First, it is not representative of the alternative industry. Most managers on managed account platforms are on more liquid edge of the whole alternative investment universe. Our sample of managers is no different: the largest group falling into either CTA or macro groups. Liquidity is known to lead to an underestimation of risk in general. One other limitation of our study is the use of net daily performance when evaluating the accuracy of the position-based VaR when the latter is applicable to gross returns. Net returns would typically have a slightly lower average and lower volatility. Because each effect would have an opposite impact on a perfectly calibrated VaR, it is difficult to evaluate the combined effect on the results. Return-based VaR is calculated using net of fee returns. The period of analysis could also be influential. The European sovereign debt crisis reached its culmination in the middle of the study period. These events, and the simultaneous market reactions, could be considered as tail events for which most VaR models are not very well prepared. This could explain why some VaR models appear to significantly underestimate the portfolio risk.

5 Conclusion

Reassuringly, position-based VaR appears to be the best indicator in our sample for short term risk management, at least from the hedge fund manager's point of view. However, the investment horizon of the average hedge fund investor would typically be several months, if not years. Investors are also generally restricted by the liquidity terms of a fund and the number of hedge funds offering daily liquidity is very limited. The accuracy of the position VaR over a typical investment horizon or over the dealing frequency is more debatable. In strategies that involve a good degree of turnover, position-based VaR could be misleading when extrapolated to longer time horizons. Interestingly, we found that for a longer term inference, a more traditional, simple historic return-based analysis could also provide very accurate results. This could be because an historic return would naturally incorporate the dynamic of the strategy. Nonetheless, a one day horizon risk indicator remains an extremely valuable tool for investors. It can be used to evaluate how the trading advisor manages the risk of the portfolio. For instance, one can look at the timing of VaR changes adjusted for the current market conditions as an indication of how the trading advisor is controlling the risk of the portfolio.

The answer to the initial question of whether return-based or position data is best is therefore not beyond dispute. From the fund manager's perspective, there is no doubt that position-based VaR provides the best instantaneous picture of the risk of a portfolio. From the investor's risk management perspective, however, an instantaneous risk picture could be of less use for typical weekly or monthly dealing periods. Its use for purely risk management purposes could therefore be limited. Conversely, return-based analysis appears very inaccurate over very short term horizons but can provide accurate information over the longer term. Our interpretation is that return-based analysis incorporates the dynamic nature of hedge funds. The study also highlights the importance of the regular monitoring of the accuracy of any chosen risk model.

References

[1] W. Fung and David A. Hsieh. Asset-based hedge fund styles and portfolio diversification. *Financial Analyst Journal, Forthcoming*, page 6, 2001.

[2] P. Jorion. Risk management for hedge funds with position information. *Journal of Portfolio Management*, 34:127–134, 2007.

[3] F.Goltz and D. Schroeder. Hedge fund reporting survey. *EDHEC Risk and Asset Management Research Centre Publication*, 2008.

[4] P. Zangari. A var methodology for portfolios that include options. *RiskMetrics Monitor*, 1st Quarter:4–12, 1996.

[5] P. Christoffersen. Evaluating interval forecasts. *International Economic Review*, 39:842–862, 1998.

[6] P. Christoffersen and D. Pelletier. Backtesting value-at-risk: A duration-based approach. *Journal of Empirical Finance*, 2:84–108, 2004.

[7] P. Kupiec. Techniques for verifying the accuracy of risk measurement models. *Journal of Derivatives*, 3:73–84, 1995.

[8] S. Campbell. A review of backtesting and backtesting procedures. Finance and Economics Discussion Series, Divisions of Research and Statistics and Monetary Affairs, Federal Reserve Board, Washington, D.C., 2005.

[9] Basel Committee on Banking Supervision. Supervisory framework for the use of 'backtesting' in conjunction with the internal models approach to market risk capital requirements. 1996.