Non-parametric Forward Looking Value-at-Risk

Marcus Nossman*and Anders Vilhelmsson†

June 14, 2012

*Marcus Nossman (Ph.D) Kyos Energy Consulting Nieuwe Gracht 49, 2011 ND Haarlem, The Netherlands

†Anders Vilhelmsson, Senior Lectuer (Ph.D) School of Economics and Management, Lund University. Department of Business administration P.O. Box 7080 S-220 07 Lund, Sweden Email: anders.vilhelmsson@fek.lu.se. Phone +46 46 222 8586.

Executive summary

This paper proposes a new model for computing Value-at-Risk forecasts. The model is fully non-parametric and easy to implement. Further, it incorporates information about the market's perceived uncertainty about the future. The forward looking information is obtained from the option market via CBOEs implied volatility index VIX. Using SP500 data from 1990-2010 we find that the use of option implied volatility compares favorably to GARCH type models in terms of forecast performance. By comparing the model primarily used in the banking sector to our new model, we find that a financial institution using our model has on average a lower market induced capital requirement (MCR). However, during the time period leading up to the financial crises our model gives a 40% higher MCR.

Keywords: Risk management, Forecasting, Value-at-Risk, Implied volatility

1 Introduction

One of the perceived culprits behind the financial crisis is the risk management tools used by financial institutions, particularly the Value-at-Risk (VaR) measure. While some critique is well justified such as VaR not being sub additive for non-elliptical return distributions (Szegö, 2002) there is some confusion in media's critique. One such example is Joe Nocera's column (Nocera, 2009) in The New York Times where drawbacks regarding a particular implementation of the VaR measure are presented as a failure inherent to the VaR method itself. The two most prevalent examples are the normality assumption of portfolio returns and the fact that VaR relies on history repeating itself. Both of these critiques actually only apply to certain implementations of VaR.

It is the purpose of this paper to suggest a VaR model that uses future looking information and further does not rely on any specific assumption regarding the return distribution. We achieve this by utilizing the information content about market participants' perceived uncertainty of the future coupled with the filtered historical simulation method of Hull and White (1998). The necessary views about the future are available from today's option prices which importantly have become easily and freely available since CBOE's construction of the implied volatility index VIX, based on SP500 options. The major benefits of the model is that it is completely non-parametric, forward looking and utilizing what should be the best variance forecast that exists. Also it is robust, under some assumptions, to the difference in implied volatility under the risk-neutral and physical probability measure which is otherwise one of the major problems when using option implied variance forecasts.

The rest of this paper is organized as follows, section 2 gives an overview of VaR and presents the new model, section 3 displays the data and explains the estimation, section 4 introduces the VaR evaluation tests, section 5 gives the results and section 6 concludes.

2 Historical simulation of Value at Risk

The Value at Risk litterature is vast and rapidly growing. Because of this the description here will by necessity be narrow and only focus on what is most often called historical simulation[1] (HS) since the model proposed in this paper belongs to this class of models. For a general overview of the field the reader is referred to the textbooks by Dowd (2005) and Jorion (2006).

Consider a δ day return $y_{t,t-\delta} = \ln(p_t/p_{t-\delta})$ with p_t being the price of a financial asset or a portfolio of such assets with time index t ranging from 1 (the beginning of the sample) to $T - \delta$ denoting the end of the sample period. Value at Risk is the maximum loss expected to incur over a certain time period

[1] There is yet a lack of a common terminology in the VaR area, this article follows the terminology of Dowd (2005).

(h) with a given probability q. Equivalently, it can be stated that the loss will be less than $VaR(q,h)$ dollars, $(1-q)*100\%$ of the time. Statistically

$$VaR_t\left(q,h\right)=F_{t+h}^{-1}\left(q\right)|\Omega_t \tag{1}$$

where F_{t+h}^{-1} is the h-step conditional forecast of the inverse cumulative distribution function (cdf) of $y_{t,t-\delta}$ and Ω_t is the information set up to and including time t information. From the definition it follows that VaR can simply be calculated as a percentile from the historical return distribution, a method known as historical simulation (HS). Historical simulation is the most used (Perignon and Smith, 2010) and perhaps also the easiest method to understand and to implement for calculating VaR forecasts. HS may also appear very flexible since it does not assume a particular data generating process, while it in fact relies on the very strong assumption that returns are independently and identically distributed (*iid*).

2.1 Filtered historical simulation

It is now, and has been for some time, very well documented that financial returns exhibit volatility clustering. Early references include Mandelbrot (1963) and Fama (1965). To formalize, volatility clustering means that the variance parameter that describes the return distribution varies over time. Several methods exist that try to augment heteroskedasticity to HS, most well known are probably Hull and White (1998) and Boudoukh et al (1998). Both of these methods (or at least the particular implementation of the method in the Hull and White case) require some parametric assumptions, thereby somewhat lessening the appeal for a user that wants a fully non-parametric method.

As previously stated historical simulation requires that the distribution of $y_{t,t-\delta}$ is independent and identical over time. Hull and White (1998) suggest to transform $y_{t,t-\delta}$ as $y^*_{t,t-\delta}=\frac{y_{t,t-\delta}}{\hat{\sigma}_{t,t-\delta}}\hat{\sigma}_{T,T-\delta}$ with $\hat{\sigma}_{t,t-\delta}$ the filtered volatility from time $t-\delta$ to t given $T-\delta$ information and $\hat{\sigma}_{T,T-\delta}$ being a volatility forecast, again using $T-\delta$ information. This approach changes the *iid* assumption from applying to $y_{t,t-\delta}$ to the much more plausible assumption of $\frac{y_{t,t-\delta}}{\hat{\sigma}_{t,t-\delta}}$ being *iid*. However Hull and White suggest to get the $\hat{\sigma}_{t,t-\delta}$ estimates from the exponentially weighted moving average model (EWMA), making their approach only semi-parametric.[2] Further the EWMA model, as any parametric model, will according to traditional market efficiency arguments only use a subset of the information available to the market participants in forming the variance forecasts.

2.2 Option implied variance forecasts

The variance forecast literature has a tradition of trying to use the information about future variance contained in the price of options. The early literature, such as Canina and Figlewski (1993), find that option implied variance predict

[2] Barone-Adesi and Giannopoulos (2001) instead use a GARCH model.

variance poorly. There are at least three reasons for this i) the forecast is of the risk-neutral and not the physical variance ii) the forecast rely on a specific option pricing model iii) the forecast was compared to a noisy proxy of the ex post volatility.

Since the work of Britten-Jones and Neuberger (2000) it is now possible to extract a variance forecast assuming only absence of arbitrage and not any particular option pricing model. Further with the realized volatility literature (Andersen and Bollerslev, 1998) it is now possible to construct ex post measures of volatility that allows for proper evaluation of volatility forecast models. The first reason for the poor performance remains however, since, at least for index options, the risk neutral volatility is often documented to be more than 50% higher than the realized volatility, see eg Bollerslev et al. (2011). The difference between the risk-neutral and physical variance can be thought of as a variance risk premium. That is, if you could own volatility[3] as an asset, you would on average earn a large negative return as documented in eg Mixon (2007) as well as Carr and Wu (2009). A negative volatility risk premium corresponds to the risk-neutral variance being higher than the variance under the real world probability measure. In the next section we will propose a method to calculate VaR that incorporates the forward looking variance and that is also robust to the presence of a variance risk premium.

2.3 The Historical simulation VIX model

The theoretically best solution for estimating the filtered volatility $\hat{\sigma}_{t,t-\delta}$ would be to use high frequency data and calculate the so called realized variance (Andersen and Bollerslev, 1998), since the realized variance converges to the true latent variance $\sigma_{t,t-1}$ as the sample frequency tends to infinity. This method was very recently suggested in Andersen et al. (2012). In practice however, data availability will often mean that only one observation is available for each variance that is to be estimated, in this case the realized variance is still unbiased but very noisy. Further, some kind of parametric model would have to be used to forecast $\hat{\sigma}_{T,T-\delta}$.

In this paper we suggest to use the option implied volatility that is freely available from CBOE's volatility index VIX. We call this new model the Historical simulation VIX model (HS-VIX). The "new" VIX introduced in 2003 is based on SP 500 options and uses the model free method of Britten-Jones and Neuberger (2000) to calculate the implied volatility. That is, $\hat{\sigma}_{t,t-\delta}$ is "estimated" by using the close of the VIX index on day $t-\delta$ denoted by $VIX_{t+22,t}$ since the VIX is a 22 trading day forecast[4] and $\hat{\sigma}_{T,T-\delta}$ is forecasted by using the close of the VIX on day $T-\delta$ (the last day in the sample) so that we get

$$y^*_{t,t-\delta} = y_{t,t-\delta}\frac{VIX_{T+22,T}}{VIX_{t+22,t}} \tag{2}$$

[3] Indeed you can, futures and options exist with variance as the underlying asset.

[4] The subscript t in used instead of $t-\delta$ since we use the close value and not the open value of the VIX on day $t-\delta$.

If the return series $y_{t,t-\delta}$ has a non-zero mean, rescaling with the volatility implies a dependence between the mean and the volatility similar to a GARCH in mean effect. For an asset with primarily systematic risk such as the SP500 this is probably desirable. However this effect can be removed by subtracting the mean from $y_{t,t-\delta}$ before the scaling and then it can be added back.

There are several advantages with the choice of using the VIX index, but also some drawbacks. The major benefits is that the method is completely non-parametric, forward looking and utilizing what should be the best variance forecast that exists. The fact that the model is non-parametric should make it much more robust to varying economic conditions where a parametric model may have to deal with structural breaks (eg a GARCH model should not be used when there is a shift in the unconditional variance). That is, a non-parametric model is basically immune to the Lucas (1976) critique.

It should also be mentioned that it is very easy to calculate expected shortfall (ES) (conditional VaR) by simply computing the average of the observations that fall below the VaR threshold for all methods based on historical simulation. The reason ES is not computed in this paper is that the methodology for evaluating ES forecasts is still not very developed.

The drawbacks with the model is that the variance forecast includes a variance risk premium, has a horizon of 22 days and is only valid for the SP500 which may differ substantially from the portfolio the user wants to compute his VaR. We will now look at each of the three drawbacks in greater detail.

2.3.1 The variance risk premium

If the variance risk premium is proportional over time to the VIX level, the real world volatility forecast can be written as $c \cdot VIX$ with c being a constant of unknown magnitude. Reassuringly we get

$$y^*_{t,t-\delta} \equiv y_{t,t-\delta} \frac{VIX_{T+22,T}}{VIX_{t+22,t}} = y_{t,t-\delta} \frac{c \cdot VIX_{T+22,T}}{c \cdot VIX_{t+22,t}} \tag{3}$$

so the method is actually immune to a variance risk premium of unknown size as long as it is proportional to the VIX level. Bollerslev et al. (2011) find evidence of time variation in the variance risk premium which makes our assumption of a proportional variance risk premium more plausible than assuming a constant variance risk premium.

2.3.2 The forecast horizon

As can be seen in equation (2) a δ day return is rescaled by a 22 day volatility which introduces a potential problem for all $\delta \neq 22$. If eg a one day volatility can be written as a constant times the 22 day volatility our method will not be affected by using the "wrong" forecast horizon. This will be the case if the VIX index follows a random walk. There is however very strong empirical evidence (the whole GARCH and Stochastic volatility literature) that variance is mean reverting. If mean reverting, the 1 day variance forecast will be more than

1/22 of the 22 days variance forecast when the variance is higher than its mean reversion level and similarly lower when the variance is below its mean reversion level. So the deviations in forecast horizon will probably have some effect that we investigate empirically by making VaR forecasts from 1 day, 10 days and 22 days returns.

There are at least two possible solutions to the mismatch between the VIX forecast horizon and the desired VaR forecast horizon. If the user has access to option data an implied variance forecast of the desired length can of course be constructed, this would completely eliminate the problem and keep the method non-parametric. The second solution is to use some parametric assumption about the volatility process and adjust the forecast accordingly. Both of these approaches are left for future research.

2.3.3 Other assets than SP500

Finally if we are interested in the VaR of some other asset than the SP500, the method once again works if the standard deviation of the other asset can be written as a constant times the level of the VIX. If one is not willing to make this assumption other volatility indexes exists, both domestic and international, which may be more suitable for a specific portfolio. It exists today implied volatility indices for the major equity indices not only in the US but also in eg Germany, the UK, France, Switzerland, Hong Kong and India. In addition to this implied volatility indices for commodities such as oil, gold and silver also exist. For an exhaustive listing of available volatility indices the reader is referred to table 1 in Andersen et al. (2011). Once again, if the user of the VaR forecast has access to relevant option data, the volatility forecast can be constructed from that data.

In addition to this, Birtoiu and Dragu (2011) use the method proposed in this paper on real profit and loss data from one American and three European banks during January 2007 through December 2010 and find that the HS-VIX model performs well. This result is encouraging although the sample period in Birtoiu and Dragu (2011) is admittedly short.

3 Data and estimation

The VIX index is available from 2 January, 1990 which is also the start of our sample period and the end date is August 30, 2010. The VIX index, which is constructed by CBOE, measures the square-root of the risk neutral expected variance of the S&P 500 index by using S&P 500 index options with a maturity of one month (22 trading days)[5]. The VIX data is combined with returns from the SP500 index.

[Insert Table 1 about here.]

[5] Details about the calculations are found at http://www.cboe.com/micro/vix/vixwhite.pdf

As can be seen from table 1 the annual standard deviation of the SP500 index is lower than the average level of the VIX index. This is consistent with a negative variance risk premium, although the difference is smaller than in Bollerslev et al. (2011) who uses the same data series. Their data however ends in 2004. Further there is persistence in the squared SP500 returns which is consistent with volatility clustering (ARCH effects).

Our aim is not to conduct a horse race between many existing models. Instead we estimate three different models on the SP500 data with three different forecast horizons. The three models are a standard historical simulation model (HS), a GARCH historical simulation model (HS-GARCH), and the model proposed in this paper called HS-VIX. The three models are selected so that we first of all can gauge if there is an improvement to be had from introducing volatility updating into historical simulation. This is done by comparing the HS model and the HS-GARCH model. Further we investigate if there is an additional improvement by using the variance forecast from the VIX index by comparing the HS-GARCH model and the HS-VIX model. The three forecasts horizons are 1 day, 10 days and 22 days.

The full sample, including 5,207 daily price observations, is divided into an estimation sample and a forecast sample. The estimation sample is set to the first 500 daily observations for the 1 day forecast horizon, 1,000 daily observations for the 10 days forecast horizon and 2,500 daily observations for the 22 days forecast horizon. The increase in the number of daily observations for the longer forecast horizons is necessary since the empirical return distribution that is used to make the VaR forecasts are constructed from returns equal to the forecast horizon, ie a 22 days VaR forecast is calculated as a one step ahead forecast from the empirical distribution of 22 days returns. The forecasts are calculated using a rolling scheme so that the size of the estimation sample is held constant and the start point and end point is moved forward with a number of days equal to the forecast horizon after each forecast is made.

The GARCH model is the standard model of Bollerslev (1986) given by

$$\begin{aligned} y_{t+\delta,t} &= \mu + z_{t+\delta,t}\sigma_{t+\delta,t} \\ \sigma^2_{t+\delta,t} &= \omega + \alpha\varepsilon^2_{t,t-\delta} + \beta\sigma^2_{t,t-\delta} \\ \text{with } \varepsilon_t &= z_t\sigma_t \text{ and } z_t \sim IID\ N(0,1). \end{aligned} \tag{4}$$

The GARCH model is estimated on the same rolling windows as described above with the three different frequencies $\delta = \{1, 10, 22\}$.

This method gives 4,707 daily forecasts, 421 ten days forecasts and 124 twenty-two days forecasts. Each forecast is made for the VaR levels 1%, 2%, 3%, 4% and 5%.

4 Backtesting VaR models

Our selection of methods for evaluating VaR models is based in part on the popularity of the tests and in part on their performance in the simulation study conducted in Berkowitz et al. (2009). We briefly summarize the tests below and refer to Berkowitz et al. (2009) for a more detailed description. Define an indicator series, I_t, according to:

$$I_t = \begin{cases} 1 \; if \; y_{t+h,t} < VaR_t(q, \delta) \\ 0 \; Otherwise \end{cases} \tag{5}$$

meaning that VaR exceptions are coded as 1 and non-exceptions as zero. As shown in Berkowitz et al. (2009) VaR evaluation tests can be presented in a unified framework by noticing that a correctly specified VaR model implies

$$E[I_t - q|\Omega_{t-1}] = 0 \tag{6}$$

with Ω_{t-1} being the information set available at time $t-1$, note that the information set is not limited to information in just the indicator series. The three popular tests proposed in Christoffersen (1998) that check for a correct number of exceptions (LR_{UC}), independence of exceptions (LR_{ind}) and jointly tests for a correct number of independent exceptions (LR_{cc}) fit into the framework above when first order Markov dependence is used as the alternative hypothesis. From (6) it follows that $E[(I_t - q)(I_{t-k})] = 0$ for all $k > 0$. This means that all autocorrelations of the mean adjusted indicator series should be equal to zero. Berkowitz et al. (2009) suggest to test this with a Ljung-Box test which we label LB1 and LB5 for one and five lags respectively. Berkowitz et al. (2009) also suggest a test based on the Caviar model of Engle and Manganelli (2004) that includes the VaR estimates from the model being evaluated. The test consists of estimating the unrestricted log likelihood from the logit model

$$I_t = \alpha + \beta_1 I_{t-1} + \beta_2 VaR_t + u_t$$

and comparing it to the restricted likelihood by setting $\beta_1 = \beta_2 = 0$ and $e^{\alpha}/(1+e^{\alpha}) = q$. This test, which we will call the Caviar test, was generally found to have the best power properties in the simulation study in Berkowitz et al. (2009). Because of the size distortions in the tests documented in Berkowitz et al. (2009) we report p-values from the tests based on the simulated finite sample distribution. We follow the method of Dufour (2006) and use 5,000 simulations.

5 Results

This section will look at the economic consequences the different models will have in terms of capital requirements according to the Basel II regulations. Also it presents the results from the statistical backtesting of the models.

5.1 Economic results

From visual inspection of the top panel of figure 1 it is apparent that HS is very slow to respond to changes in returns and because of this produces VaR estimates that are constant for prolonged periods of time. For example, on the 12 September, 2008, the Friday just before Lehman brothers filed for Chapter 11 the following Monday, the one day ahead VaR forecast from the HS-VIX and HS-GARCH were 3.50% and 4.07%. Five days later they had increased to 4.98% and 7.60%. In contrast to this, the VaR from HS is 3.02% the Friday before Lehman Brothers and almost unchanged at 3.20% five days later.

Both the VIX and GARCH filtered HS respond in a similar fashion to changes in volatility but generally the most extreme movements are created by the HS-GARCH model. The average VaR for the 1 day forecast horizon is -2.79 for the HS model, -2.59 for HS-GARCH and -2.58 for HS-VIX.

[Insert figure 1 about here.]

The market induced capital requirements (MCR) for a bank using the internal models approach under Basel II is based on the 10 days ahead 99% VaR forecast (or the average of the last 60 days if this is higher) times a scaling factor that depends on the number of exceptions (e) the bank's VaR model had when backtested on 250 days one step ahead 99% VaR forecasts (BIS, 1996). The scaling factor ranges from 3 for a maximum of four exceptions (a correct model will on average have 2.5 exceptions) to 4 for a model that has produced 10 or more exceptions. Specifically the scaling factor S_t it is determined by

$$S_t = \begin{array}{c} 3 \; if \; e \leq 4 \\ 3 + 0.2(e-4) \; if \; 5 \leq e \leq 9 \\ 4 \; if \; e \; \geq 10 \end{array} \tag{7}$$

and the capital requirement by

$$MCR_t = S_t \max(VaR_t(0.01, 10), \frac{1}{60} \sum_{i=0}^{60} VaR_{t-i}(0.01, 10)) \tag{8}$$

which is plotted in figure 2 for the different models using 10 days-ahead non-overlapping forecasts.

[Insert figure 2 about here.]

Figure 2 displays the MCR over the sample period for a constant position of 100$ in the SP500 index using the three different models. For the whole sample period the MCRs is 28.61$ for the HS, 24.07$ for the HS-VIX and 25.50$ for the HS-GARCH. From figure 1 we saw that the HS-VIX and HS-GARCH models reacted much faster in increasing the VaR after the Lehman Brothers filed for chapter 11. However the real purpose of the HS-VIX model is the use of forward looking information, so can we see an increase in the MCR already before Lehman files for chapter 11?

In figure 2, observation 373 is the first observation after Lehman Brothers's bankruptcy but we can see that already from observation 336 (April 10, 2007) the MCR from the HS-VIX model starts to be consistently higher than that of the HS and HS-GARCH models.[6] The average MCR for the HS-VIX model is 21.99$ from April 10, 2007 until September 11, 2008 (from the observed divergence until just before the Lehman bankruptcy), 15.67$ for the HS model and for the HS-GARCH model the MCR is 16.09$. This shows that a financial institution that would have based their market induced capital requirements on the HS-VIX model would have been shielded by 36.7% more capital than if they had used the HS-GARCH model and 40.4% more than if they had used the HS model. Given the results in Perignon and Smith (2010), that the HS model is the most commonly used model by the banking sector, it is interesting to ponder what would have happened during the financial crisis if banks had instead used a model that would have required about 40% more risk capital[7].

During the financial crisis we can see that the MCR is actually higher then the maximum loss of the position which because of limited liability is equal to 100$. This inconsistency in the Basel II rules is the result of the scaling factor that does not take into account what the maximum loss of the position is.

5.2 Statistical results

All results from the statistical backtesting are presented in table 2. The test labeled unconditional coverage (LR_{UC}) tries to answer the question if the models provide a correct number of exceptions. All three models somewhat underestimate the VaR at the 1 day horizon, particularly the HS-GARCH model. The test for unconditional coverage can reject the HS-GARCH model as providing correct coverage at the 5% significance level for all five VaR level, the HS model can be rejected at three VaR levels and the HS-VIX model at two VaR levels. The 10 and 22 days forecast horizons seem much more difficult and both the HS and HS-GARCH models severely underestimate the risk with the number of exceptions being two to three times too high. Both the HS and HS–GARCH models can be rejected[8] as providing a correct number of exceptions on all VaR levels for both the 10 and 22 day forecast horizon. The HS-VIX model is better and can be rejected on three VaR levels at the 10 days horizon and at two VaR levels for the 22 days horizon. Even though the HS-VIX model overall achieved the best unconditional coverage we would expect to see the largest benefits when it comes to avoiding clustering of exceptions and this is indeed also the case.

[Insert table 2 about here.]

[6] This is one week after the subprime mortage lender New Century Financial Corporation files for Chapter 11.

[7] Most of the banks' risk capital is based on credit risk, the results presented here only apply to market risk.

[8] All inference reported in the rest of the result discussion is based on a 5% significance level.

The three tests LR_{ind}, LB1 and LB5 all look at clustering of exceptions without taking any account of the actual number of exceptions. Looking at these tests there are dramatic differences between HS and HS-VIX. Taking all three tests, five VaR levels and three forecasts horizons together the HS model is rejected as providing independent exceptions on 27 of a maximum of 45 occasions with most of the violations occurring at the 1 day forecast horizon. The HS-VIX model can be rejected on 11 occasions and the HS-GARCH model fall in between with 18 rejections. Looking at the combined results for the unconditional coverage and independence test we see that the HS-GARCH model is better at capturing the dynamics in the data than at capturing the level. This indicates that the α and β parameters are estimated with better precision than the ω parameter in equation (4)

The LR_{cc} and the Caviar tests jointly if a model can produce a correct number of exceptions that are independently distributed over time. Both the HS and HS-GARCH models can be rejected by both tests on all VaR levels and forecast horizons with the single exception of the HS model at the 1 day forecast horizon for 4% VaR. The HS-VIX model is rejected 14 times out of 30 (2 tests, 3 forecast horizons and 5 VaR levels).

6 Conclusion

The performance of the HS-GARCH model is markedly worse than in Hull and White (1998). There are several reasons for this: i) the joint tests for a correct number of independent exceptions were not available and hence not included in the Hull and White (1998) study, ii) our data set is challenging since it contains the financial crisis iii) we include 10 and 22 days forecast horizons, whereas Hull and White (1998) only used 1 day forecasts. In view of this, the results that the HS-VIX model is rejected 14 times out of 30 on the joint tests, whereas the HS-GARCH and HS models are rejected 30 and 29 times out of 30 respectively, must be considered a success for the new model.

The results presented in this paper have implications both for users of VaR models, such as banks, and also for regulators. A financial institutions using the HS-VIX model proposed in this paper would have approximately 40% higher market induced risk capital than an otherwise identical institution using the HS model during the time period leading up to the financial crises. The large difference between the HS-VIX and HS model is especially important in light of the finding in Perignon and Smith (2010) that the HS model is by far the most used among commercial banks. A forward looking model, such as the HS-VIX, could hence play an important role in crises mitigation.

References

Andersen, T., Bollerslev, T., Christoffersen, P. & Diebold, F. (2012), 'Financial risk measurement for financial risk management', *NBER working paper 18084* .

Andersen, T., Bondarenko, O. & Genzalez-Perez, M. (2011), 'Coherent model-free implied volatility: A corridor fix for high-frequency vix', *Working Paper , SSRN* .

Andersen, T. G. & Bollerslev, T. (1998), 'Answering the skeptics: yes, standard volatility models do provide accurate forecasts', *International Economic Review* **39**, 885–905.

Barone-Adesi, G. & Giannopoulos, K. (2001), 'Non parametric var techniques. myths and realities', *Economic Notes* **30**, 167–181.

BASEL (1996), 'Supervisory framework for the use of "backtesting" in conjuction with the internal models approach to market risk capital requirements', *Basel Committee publication* **22**.

Berkowitz, J., Christoffersen, P. & Pelletier, D. (2009), 'Evaluating value-at-risk models with desk-level data', *Management Science* **Articles in Advance**.

Birtoiu, A. & Dragu, F. (2011), 'Market risk management: The applicability and accuracy of value-at-risk models in financial institutions', Master Thesis, Lund University.

Bollerslev, T. (1986), 'Generalised autoregressive conditional heteroskedasticity', *Journal of Econometrics* **51**, 307–327.

Bollerslev, T., Gibson, M. & Zhou, H. (2011), 'Dynamic estimation of volatility risk premia and investor risk aversion from option-implied and realized volatilities', *Journal of Econometrics* **160**, 235–245.

Boudoukh, J., Richardson, M. & Whitelaw, R. (1998), 'The best of both worlds: A hybrid approach to calculating value at risk', *Risk* **11**, 64–67.

Britten-Jones, M. & Neuberger, A. (2000), 'Option prices, implied price processes, and stochastic volatility', *Journal of Finance* **55**, 839–866.

Canina, L. & Figlewski, S. (1993), 'The informational content of implied volatility', *Review of Financial Studies* **6**, 659–681.

Carr, P. & Wu, L. (2009), 'Variance risk premiums', *Review of Financial Studies* **22**, 1311–1341.

Christoffersen, P. (1998), 'Evaluating interval forecasts', *International Economic Review* **39**, 841–862.

Dowd, K. (2005), *Measuring Market Risk 2nd edition*, Wiley, Chichester, West Sussex, England.

Dufour, J.-M. (2006), 'Monte carlo tests with nuisance parameters : A general approach to finite-sample inference and nonstandard asymptotics in econometrics', *Journal of Econometrics* **133**, 443–477.

Engle, R. & Manganelli, S. (2004), 'CAViaR: Conditional autoregressive value at risk by regression quantiles', *Journal of Business and Economic Statistics* **22**, 367–381.

Fama, E. F. (1965), 'The behaviour of stock market prices', *Journal of Business* **38**, 34–105.

Hull, J. & White, A. (1998), 'Incorporating volatility updating with the historical simulation method for value at risk', *Journal of Risk* **Fall**, 115–119.

Jorion, P. (2006), *Value at Risk. The new benchmark for managing financial risk, 3rd edition*, McGraw-Hill, New York.

Lucas, Jr., R. E. (1976), Econometric policy evaluation: A critique, *in* 'The Phillips Curve and Labor Markets', Vol. 1 of *Carnegie Rochester Conference Series on Public Policy*, North-Holland, Amsterdam, pp. 19–46.

Mandelbrot, B. (1963), 'The variation of certain speculative prices', *Journal of Business* **36**, 394–419.

Mixon, S. (2007), 'The implied volatility term structure of stock index options', *Journal of Empirical Finance* **14**, 333–354.

Nocera, J. (2009), 'Risk mismanagement', `http://www.nytimes.com/2009/01/04/magazine/04risk-t.html?_r=1&ref=joe_n%ocera`.

Perignon, C. & Smith, D. (2010), 'The level and quality of value-at-risk disclosure by commercial banks', *Journal of Banking and Finance* **34**, 362–377.

Szegö, G. (2002), 'Measures of risk', *European jounal of operational research* **63**, 5–19.

Figures and Tables

Table 1 - Descriptive statistics		
	SP500	VIX
Mean	5.14%	20.4
St dev	18.59%	8.31
Skewness	-0.20	2.01
Kurtosis	11.84	10.11
$\rho(1)$	-0.055	0.983
$\rho(5)$	-0.031	0.942
$\rho^2(1)$	0.205	
$\rho^2(5)$	0.326	

This table presents descriptive statistics for the CBOE implied volatility index (VIX) and for the percentage returns of the SP500. The time period is January 2, 1990 until August 30, 2010. Mean and St dev is the annual percentage mean and standard deviation, ρ is the autocorrelation and ρ^2 is the autocorrelation of the squared series (unreported for VIX).

Table 2 - VaR Forecast performance

Model	Forecast horizon 1 day					Forecast horizon 10 days					Forecast horizon 22 days				
	VaR Level														
	1%	2%	3%	4%	5%	1%	2%	3%	4%	5%	1%	2%	3%	4%	5%
	VaR Exceptions														
HS	1.46%	2.55%	3.41%	4.52%	5.41%	2.61%	4.04%	6.41%	9.03%	10.21%	4.03%	6.45%	6.45%	8.87%	11.29%
HS-VIX	1.33%	2.25%	3.61%	4.42%	5.34%	1.90%	4.28%	5.94%	7.13%	7.13%	3.23%	4.84%	6.45%	8.06%	8.87%
HS-GARCH	1.63%	2.67%	3.68%	4.94%	5.86%	3.56%	5.23%	8.08%	9.50%	10.93%	4.03%	6.45%	10.48%	12.10%	13.71%
	LR UNC														
HS	7.51	5.65	2.24	2.76	1.41	7.66	6.91	12.8	20.66	18.76	6.54	7.95	3.85	5.75	7.73
	0.01	0.02	0.13	0.09	0.24	0.02	0.01	<0.01	<0.01	<0.01	<0.01	<0.01	0.05	0.01	<0.01
HS-VIX	4.14	1.23	4.83	1.83	0.95	2.73	8.42	9.78	8.76	3.56	3.91	3.66	3.85	4.16	3.21
	0.05	0.28	0.03	0.18	0.32	0.12	<0.01	<0.01	0.01	0.06	0.02	0.11	0.05	0.06	0.08
HS-GARCH	13.66	8.36	6.04	8.72	5.93	16.82	15.55	25.74	24.24	23.61	6.54	7.95	14.71	13.99	13.72
	<0.01	<0.01	0.01	<0.01	0.02	<0.01	<0.01	<0.01	<0.01	<0.01	0.01	<0.01	<0.01	<0.01	<0.01
	LR IND														
HS	6.48	5.48	6.8	2.61	3.19	1.16	1.87	4.96	5.71	4.88	0.42	1.11	1.11	2.16	3.61
	<0.01	0.01	0.02	0.11	0.08	0.04	0.06	0.01	<0.01	0.01	0.02	0.04	0.08	0.05	0.02
HS-VIX	0.1	3.3	3.75	4.32	3.59	0.31	1.54	1.39	1.54	1.54	0.27	0.62	0.41	0.05	0.00
	0.99	0.09	0.06	0.04	0.06	0.08	0.08	0.14	0.22	0.31	0.03	0.06	0.33	0.95	0.99
HS-GARCH	2.42	4.52	1.13	<0.01	0.06	6.32	2.39	3.59	6.74	5.11	0.42	1.11	0.14	0.56	1.25
	0.05	0.02	0.30	0.96	0.81	0.01	0.04	0.02	<0.01	0.01	0.02	0.05	0.72	0.38	0.15
	LR CC														
HS	13.99	11.13	9.04	5.37	4.59	8.82	8.77	17.76	26.37	23.64	6.96	9.07	4.96	7.92	11.34
	<0.01	<0.01	0.01	0.06	0.10	<0.01	<0.01	<0.01	<0.01	<0.01	0.01	0.01	0.03	0.01	<0.01
HS-VIX	4.23	4.52	8.58	6.14	4.54	3.04	9.95	11.17	10.3	5.10	4.18	4.28	4.26	4.21	3.21
	0.09	0.11	0.02	0.05	0.11	0.14	0.01	<0.01	<0.01	0.06	0.03	0.03	0.04	0.09	0.15
HS-GARCH	16.08	12.88	7.17	8.72	5.99	23.13	17.94	29.32	30.98	28.72	6.96	9.07	14.84	14.55	14.97
	<0.01	<0.01	0.05	0.01	0.05	<0.01	<0.01	<0.01	<0.01	<0.01	0.01	<0.01	<0.01	<0.01	<0.01

table continued

	Forecast horizon 1 day					Forecast horizon 10 days					Forecast horizon 22 days				
	VaR Level														
Model	1%	2%	3%	4%	5%	1%	2%	3%	4%	5%	1%	2%	3%	4%	5%
	LB1														
HS	11.92	7.79	9.13	3.02	3.63	1.87	2.74	7.08	7.38	6.01	0.23	0.61	0.61	1.22	2.09
	0.01	0.01	<0.01	0.08	0.05	0.04	0.08	0.02	0.01	0.02	0.02	0.04	0.09	0.11	0.08
HS-VIX	0.11	4.53	4.65	5.19	4.12	0.16	2.16	1.75	1.88	1.88	0.14	0.33	0.52	0.05	<0.01
	0.99	0.03	0.03	0.02	0.04	0.08	0.10	0.11	0.09	0.10	0.04	0.06	0.11	0.87	0.99
HS-GARCH	3.55	6.21	1.28	<0.01	0.06	12.31	3.33	4.58	8.72	6.22	0.23	0.61	0.13	0.50	1.07
	0.05	0.02	0.26	0.96	0.81	0.02	0.07	0.04	0.01	0.02	0.02	0.05	0.52	0.22	0.12
	LB 5														
HS	23.93	31.44	38.04	62.03	64.34	4.69	6.56	8.93	8.42	9.78	4.5	12.36	12.36	4.53	13.89
	0.01	<0.01	<0.01	<0.01	<0.01	0.17	0.2	0.12	0.12	0.07	0.05	0.08	0.08	0.26	0.03
HS-VIX	1.82	15.26	11.35	13.47	13.52	5.6	4.04	5.6	5.94	5.94	0.75	3.63	2.83	2.76	5.39
	0.91	0.02	0.05	0.02	0.02	0.17	0.36	0.24	0.23	0.24	0.07	0.17	0.36	0.51	0.22
HS-GARCH	9.05	9.90	10.64	9.58	7.42	18.28	5.41	6.73	12.52	7.08	4.51	8.19	4.48	4.58	4.19
	0.13	0.08	0.06	0.08	0.18	0.06	0.28	0.19	0.04	0.17	0.06	0.09	0.25	0.27	0.34
	Caviar														
HS	16.92	13.85	10.82	8.09	8.71	21.12	13.46	21.93	28.47	25.12	7.43	11.65	7.12	18.85	25.36
	<0.01	<0.01	0.02	0.04	0.03	<0.01	<0.01	<0.01	<0.01	<0.01	0.02	0.01	0.05	<0.01	<0.01
HS-VIX	5.87	6.21	10.4	7.21	6.08	7.00	17.56	18.62	14.78	9.71	4.75	4.94	5.43	6.15	4.78
	0.10	0.10	0.02	0.07	0.11	0.03	<0.01	<0.01	<0.01	0.02	0.07	0.10	0.08	0.08	0.15
HS-GARCH	22.09	22.61	12.6	14.38	12.63	29.12	25.88	34.77	35.12	31.9	10.15	11.93	20.57	18.76	15.62
	<0.01	<0.01	0.01	<0.01	0.01	<0.01	<0.01	<0.01	<0.01	<0.01	0.01	0.01	<0.01	<0.01	<0.01

This table presents 1, 10 and 22 days ahead VaR forecast results for VaR levels ranging from 1% to 5%. VaR exceptions give the percentage of VaR exceptions for each model and VaR level. Test statistics for each test are given with simulated finite sample p-values below based on 5,000 simulations. The six different tests are described in section 4.

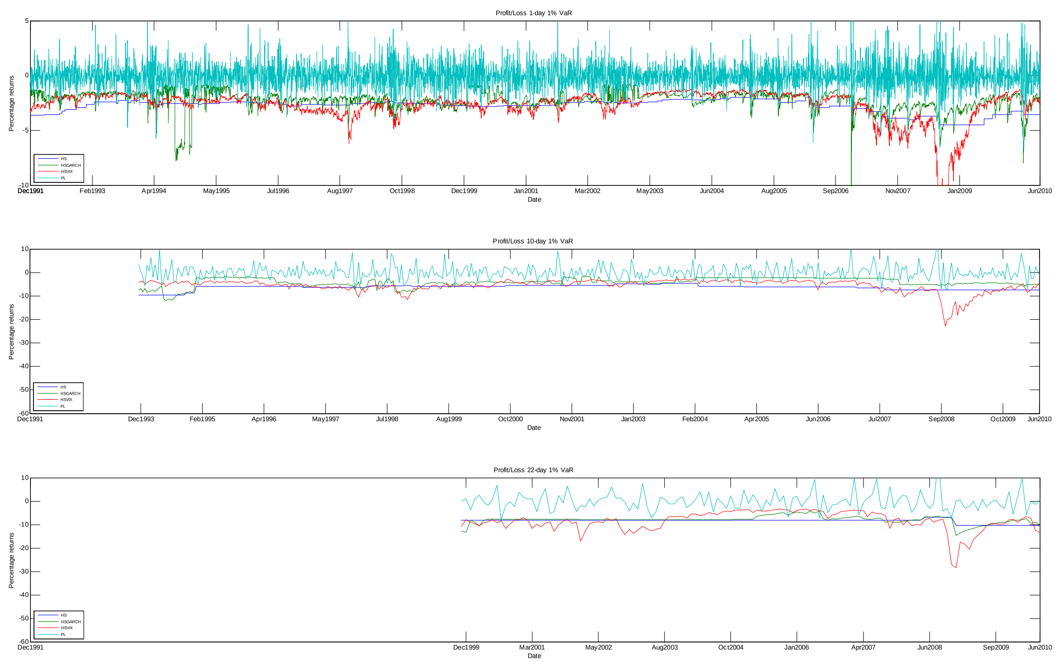

Figure 1. This figure shows the returns for the SP500 index (P/L) and the VaR forecasts for the three models, HS, HS-GARCH and HS-VIX.

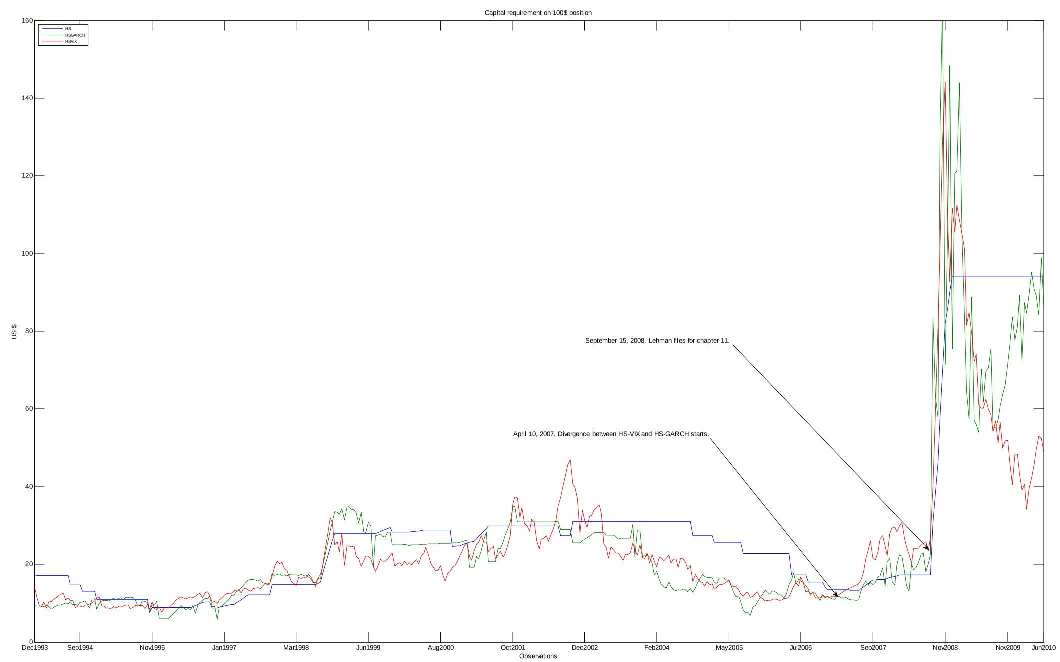

Figure 2. This figure shows the market induced capital requirements for the three models, HS, HS-GARCH and HS-VIX calculated according to equation (7).