

A News-Based Approach for Computing Historical Value-at-Risk

Frederik Hogenboom, Michael de Winter, Flavius Frasinca, and Alexander Hogenboom

Abstract Within the field of finance, Value-at-Risk (VaR) is a widely adopted tool to assess portfolio risk. When calculating VaR based on historical stock return data, the data could be sensitive to outliers caused by seldom occurring news events in the sampled period. Using a data set of news events, of which the irregular events are identified using a Poisson distribution, we research whether the VaR accuracy can be improved by considering news events as additional input in the calculation. Our experiments show that when a rare event occurs, removing the event-generated noise from the stock prices for a small, optimized time window can improve VaR predictions.

1 Introduction

Despite its limitations in terms of interpretability and mathematical properties [2, 16], Value-at-Risk (VaR) is a widely adopted risk measure used by practitioners in the field of finance, quantifying the risk of loss on a portfolio of financial equities. It is defined as a threshold value and confidence level such that the probability that the loss on the portfolio over a given time horizon does not exceed a certain value at a given confidence level. It is generally assumed that there are no unexpected trend breaks. However, in reality we are faced with deviations from trends, mainly caused by emerging events. These events are usually reported in news and can greatly impact today's financial markets. For example, when Google announced a 29% increase in its 2011 Q3 net-income, within hours its shares went up by 7%.

According to the weak form of the efficient market hypothesis, news that contains information on an equity is not perfectly incorporated in the price when it is

Frederik Hogenboom (e-mail: fhogenboom@ese.eur.nl) · Michael de Winter (e-mail: m.r.dewinter88@live.nl) · Flavius Frasinca (e-mail: frasinca@ese.eur.nl) · Alexander Hogenboom (e-mail: hogenboom@ese.eur.nl)

Erasmus University Rotterdam, P.O. Box 1738, NL-3000 DR Rotterdam, The Netherlands

published. Studies have reported on the existence of such a delay [8], caused by initial over- or under-reactions to the news. Additionally, news events have an effect on the volatility of equities [15]. Hence, taking into account news events for VaR calculations (which are based on returns distributions) could be beneficial, as the volatility is the standard deviation of the distribution of returns.

As the usage of information extracted from text in a financial context has proven to be a vital strategy in many financial applications [6, 12], we hypothesize that we can improve VaR computations by introducing financial news events [3, 10] as an additional input. In our research, we employ the ViewerPro [18] software for the extraction of ticker data and news events. By using a Poisson distribution, we identify the irregular (and hence noisy) events. Subsequently, we cleanse the ticker data from event-generated noise, and aim to obtain a data set which is a more accurate representation of the expected returns distribution. In our experiments, we aim to optimize the time window for which the noise is removed by evaluating for different configurations the accuracies of the calculated VaR.

This paper is organized as follows. First, we describe related approaches to this research in Sect. 2. Then we introduce our framework in Sect. 3. Section 4 presents our implementation, our data set, and an evaluation of the framework on this data set. Last, in Sect. 5 we draw our conclusions and provide directions for future work.

2 Related Work

The existence of a relationship between the stock market and news events has been acknowledged by many previous studies [5, 7, 9]. Additionally, the number of news events and trading activity have proven to be correlated [15]. Even though the efficient market hypothesis supports that news information is fully and immediately processed into the value of shares, in practice this is not always the case [17]. Hence, for traders, timely and accurately reacting on news and estimating the VaR of portfolios correctly, is of utmost importance.

The three most widely used implementations for VaR calculations are the parametric method (assuming a specific distribution of equity returns), a Monte Carlo simulation-based method that predicts future returns by fitting a distribution based on historical data, and the historical method, which assumes that historical changes in the price accurately predict changes in the future. Common distributions for the parametric method are the normal and log-normal distributions, as they offer simplicity and robustness. However, in practice, equity returns are almost never normally distributed [1]. Assuming a specific distribution could therefore lead to a bias in the risk measure. Even though the Monte Carlo simulation overcomes this problem by randomly sampling the historical data multiple times to approximate its distribution, this method is rather slow as it is computationally intensive. As we aim for an application that is able to run real-time, Monte Carlo simulation-based methods are not suitable for our research. Similarly, the historical method also analyzes a set of historical returns instead of an assumed distribution. An advantage of the

historical method over the Monte Carlo simulation-based methods is its simplicity, which fosters real-time computation. Therefore, in this paper we utilize the historical method for VaR prediction in which we implement event-based improvements.

Hull and White [11] improve the VaR calculation by updating the volatility in the historical method by means of GARCH/EWMA models in order to reflect the difference between the volatility at the time of the observation and the current volatility. While Hull and White analyze multiple equity portfolios, in our work we only observe single equity portfolios in order to prevent heteroscedasticity (i.e., interdependencies between variances, which is often the case with different financial equities in a portfolio). The authors propose a method to update the volatility in the appropriate time interval so that the volatility becomes a more dynamic factor in VaR calculation. Based on mean absolute percentage error (MAPE), their work is compared to another method, involving the assignment of weights to observations that are more recent [4]. The authors find that their method outperforms both the traditional historical method and second method for exchange rates, yet for stock indices, results are mixed.

Other work that aims to improve technical indicators with news was performed by Zhai et al. [19]. The authors make use of a simple text classification algorithm with a supervising learning method. Instead of only using company specific news, they are also integrating general market news in combination with technical indicators. It is concluded that technical indicators and news events alone are inaccurate as estimators, but that the combination of both could lead to better results. Based on a real-life market simulation, the authors show that by using their approach it is possible to make profit.

3 Framework

In order to be able to assess whether the incorporation of news into the calculation of the VaR of a specific equity improves the overall quality of the outcomes, we propose a framework that is based on two inputs, i.e., a list of stock prices and a list of financial events, which are extracted from several feeds such as Reuters using the ViewerPro application.

In a pre-processing phase, we cleanse the collected equity prices as follows. As stock markets are only open on specific dates and times, we filter the prices and keep those within market opening times. Also, in order to decrease computational complexity, the time intervals between individual prices are defined per hour instead of per second.

Subsequently, we read in news events, stemming from news items processed by ViewerPro using computational linguistics, semantic analysis, and formal logic. ViewerPro determines the positive and negative impacts of the information described in the news on the equities that are relevant to the user. Large amounts of news messages are filtered for equity-specific news, and the semantic component of ViewerPro analyzes each individual news message for economic impact. This

yields a list of relevant annotated news events. Some general types of news events that are covered by the ViewerPro annotations are hiring and resignation of CEOs, acquisitions, profit announcements, etcetera.


An additional step is performed by identifying irregularly occurring event types from our event set, as these events are not likely to occur again and thus cause a significant noise in stock rates. As Poisson distributions are used in many fields to model the number of occurrences of events in a certain time interval (if the average rate of occurrences is known and we assume that events occur independently from each other), we apply a Poisson distribution F to a test set $test$, which is a function of the measured and expected number of occurrences in the test set, i.e., x and λ , respectively:

$$F(x; \lambda) = \frac{\lambda^x e^{-\lambda}}{x!} . \quad (1)$$

As depicted in Fig. 1, when using a threshold α of 0.05, for $x = 0$ (which means no event occurrences), $F(x; \lambda) < \alpha$ for $\lambda \geq 3$. For a training set $train$ the expected number of occurrences λ' is obtained by scaling λ by the proportion of the set cardinalities, i.e., $\lambda' = \lambda \times \vartheta$, with $\vartheta = |train|/|test|$. Hence, we consider event types that occur $\geq 3 \times \vartheta$ as regular events, and events occurring $< 3 \times \vartheta$ as rare events.

The identified (rare) events are subsequently associated with times in which they occurred and also with the recorded stock rates. We adjust the collected prices for a time window to account for the generated noise by updating their values to the previously measured value, which is illustrated by Algorithm 1 that processes a list of chronologically ordered (hourly) recorded prices. For each stock price $price$ in price list $prices$, we compare the stock price time with the time of each event $event$ stored in event list $events$ in order to check for event occurrences. If an event occurrence is identified, $impact$ is set to the window size $window$ (for which the optimization is given in Sect. 4), causing the value of the subsequent $price$ items to be set to the current value. The value of $impact$ is decreased with 1 every next $price$ in price list $prices$, so that subsequent price values are updated up until the window size has been reached. In case of overlapping events, the $impact$ counter is reset to the window size $window$. After processing all original prices stored in $prices_{hist}$, we obtain a new list of event-corrected prices, i.e., $prices_{event}$.

Fig. 1 Poisson distributions for various measured and expected occurrences, i.e., x and λ , respectively



Algorithm 1 News event processing (per equity)

Require: *prices* = array of stock prices and associated times

Require: *events* = array of events and associated times

Require: *window* = integer representing time window

```
1: previousprice.value = prices[1].value
2: for all price in prices do
3:   for all event in events do
4:     if impact > 0 then
5:       impact = impact - 1
6:       price.value = previousprice.value
7:     end if
8:     if price.time = event.time then
9:       impact = window
10:    end if
11:  end for
12:  previousprice.value = price.value
13: end for
```

Both sets of original (“*hist*”) and denoised (“*event*”) prices are converted to sets with hourly returns. We compute the return set *returns* of a price set *prices* as the relative change between the price at time $t + 1$ and the previous price at time t , i.e.,

$$returns = \frac{prices_{t+1} - prices_t}{prices_t} \quad \forall t = 1, \dots, N - 1. \quad (2)$$

where N represents the number of items in the list. A specific return $returns_t$ equals the profit that can be obtained if a share is bought at time t and sold at time $t + 1$.

We make use of the historical returns (both original and adapted) to estimate the future returns. The time horizon used for computing returns is 1 day. After sorting the return list *returns*, we calculate the Value-at-Risk, *VaR*, as

$$VaR = returns' [[\alpha \cdot \text{length}(returns)]] . \quad (3)$$

Here, $returns'$ represents the ordered (sorted) list of returns and where the confidence level is denoted by α . Thus, in a data set with 20 historical returns – with the first element being located on position 1, and the last on position 20 – we select the first worst return (i.e., position 19) for a confidence level of 0.95. With (3), we calculate VaR_{event} and VaR_{hist} using our adjusted method and the traditional method (i.e., the historical method without the improvements proposed in [4, 11]), respectively.

4 Evaluation

In order to evaluate the performance of the proposed historical VaR calculation, our framework is implemented as a Java-based application that calculates the VaR of a single equity based on a data set containing news events and stock prices.

The data set used in our experiments stems from the ViewerPro software, and – after filtering – covers news events and stock data collected on an hourly basis for 363 heterogeneous equities on weekdays during the year 2010, and contains approximately 2,000 stock data points, 119 different event types, and 50 up to 75 associated events per equity. In order to evaluate the performance of the calculation, we predict the VaR_{event} and VaR_{hist} for 75% of our data set. The remaining 25% is used as a test set for comparing the predicted VaR with the actual VaR.

Even though many VaR analyses are currently performed using the Kupiec test [14], we employ a different set of measures. As explained by Kupiec in his original work, the test is statistically weak with sample sizes of one year. As our data set covers only the 2010, we need different measures that provide insight into the effectiveness of our proposed event-based approach.

In order to analyze for how many equities our adjusted event-based historical method provides better quality predictions in comparison to the traditional historical method, we measure each method's squared error. The squared error SE for equity e is defined as the squared difference between the equity's actual VaR ($VaR_{e,actual}$) measured in our test set and the predicted VaR ($VaR_{e,predicted}$) that has been predicted based on our training set, i.e.,

$$SE_e = (VaR_{e,actual} - VaR_{e,predicted})^2, \quad (4)$$

where $VaR_{e,predicted}$ is one of VaR_{event} or VaR_{hist} .

Subsequently, the squared errors are combined into the mean squared error (MSE), yielding an MSE_{hist} and MSE_{event} . The MSE is calculated as the summation of the squared errors (SE) of all equities $e \in E$ divided by the number of equities, i.e.,

$$MSE = \frac{\sum_{e \in E} SE_e}{|E|}, \quad (5)$$

where $|E|$ denotes the total number of equities in set E , in our case 363.

Additionally, we evaluate the number of times both methods outperform one another, i.e., OPT (OutPerformed Total), by comparing the squared errors $SE_{e,hist}$ and $SE_{e,event}$ for each equity $e \in E$, yielding

$$OPT_{hist,event} = \sum_{e \in E} O(SE_{e,hist}, SE_{e,event}), \quad (6)$$

$$OPT_{event,hist} = \sum_{e \in E} O(SE_{e,event}, SE_{e,hist}), \quad (7)$$

$$O(X, Y) = \begin{cases} 1 & \text{if } X < Y \\ 0 & \text{else} \end{cases}. \quad (8)$$

In our experiments, we compare the MSE and OPT for the traditional and event-based VaR calculation methods, both on the full event data set, as well as on a data set containing only the rare events, using an arbitrary time window of 8 hours (determined based on initial estimates). Subsequently, we determine the optimal

time window size by observing plots of MSE and OPT values for the event-based VaR calculation method. Also, we take into account the number of overconfident predictions ($CONF$) of all equities $e \in E$, which is calculated as

$$CONF = \sum_{e \in E} C(VaR_{e,predicted}, VaR_{e,actual}), \quad (9)$$

$$C(X, Y) = \begin{cases} 1 & \text{if } X > Y \\ 0 & \text{else} \end{cases}, \quad (10)$$

where $VaR_{e,predicted}$ represents the predicted VaR_{event} for equity e based on our adjusted data set (only containing the rare events).

Last, we perform a two-sample one-tailed t -test on the sets of individual squared errors SE_{hist} and SE_{event} (containing $SE_{e,hist}$ and $SE_{e,event} \forall e \in E$, respectively) for our optimal configuration, in order to assess the significance of the measured difference between MSE_{hist} and MSE_{event} . For this, we use a significance level of 0.05 to reject the null hypothesis that there is no difference between the measured MSE values.

When comparing the results from both VaR calculation methods using our data set containing all (i.e., regular and rare) events and stock rates on an hourly basis, and when using a time window of 8 hours, we obtain the results depicted in Table 1, which shows the MSE and OPT values for both the traditional and the event-based historical VaR calculation methods (columns *hist* and *event*, respectively). We observe an improvement of 21.66% in terms of MSE when accounting for event-generated noise in our stock data. Additionally, the event-based VaR calculation method outperforms the traditional historical method (in terms of squared errors of predicted VaR values with respect to the actual VaR values) 232 times which is a share of 63.91% of all predictions, compared to 131 observations in which the traditional method outperforms our adapted method.

As presented in Table 1, repeating the same experiment on a filtered data set which only contains non-rare events (i.e., 345 in total) yields an additional improvement over the previous results. Now, in 71.88% of the cases (i.e., 248 out of 345), event-based historical VaR calculation outperforms the traditional method. Also, we see a performance gain in terms of MSE . The improvement in MSE values has increased from 21.66% to 26.29%. Both scores underline the added value of only considering the rare events.

Subsequently, we optimize the size of the time window by evaluating MSE and OPT values on the one hand, and the number of overconfident predictions ($CONF$) on the other hand. For this, we observe VaR prediction models with time windows

Table 1 Experimental results of the performance of traditional and event-based historical VaR calculation (columns *hist* and *event*, respectively), while employing a cleansing window of 8 hours

Measure	All events		Non-rare events	
	<i>hist</i>	<i>event</i>	<i>hist</i>	<i>event</i>
MSE	1.0590E-05	8.2965E-06	1.1220E-05	8.2700E-06
OPT	131	232	97	248

ranging from 1 to 24 (i.e., 3 working days of 8 hours, which is the maximum effect of a news event [13]). As depicted by the graphs in Fig. 2, cleansing the data with a window of 10 hours yields the highest score for *OPT* (i.e., 249). However, the lowest *MSE* value is observed for a window of 14 hours. The number of overconfident predictions increases for each increase in window size, and hence we opt for a time window of 10 hours, as this maximizes the number of outperforming predictions while minimizing the number of overconfident predictions.

As shown in Table 2, utilizing a window of 10 instead of 8 hours on a data set with rare events yields an improvement both in terms of *MSE* and *OPT*. The *MSE* of our event-based historical VaR prediction models improves with 31.73% over the traditional historical VaR prediction method's *MSE*. This improvement is a lot higher than the measured improvement of 21.66% when using a cleansing window of 8 hours. Alternatively, we can also determine an optimal cleansing window for each event type separately by evaluating the percentile differences from the mean stock rate per equity in order to determine the impact of an event type. Large differences (e.g., > 50.00%) after an event occurrence indicate noise that should be cleansed, while a small difference (i.e., the smallest difference after an event) indicates that the market has returned to normal, hence not requiring any cleansing. This strategy

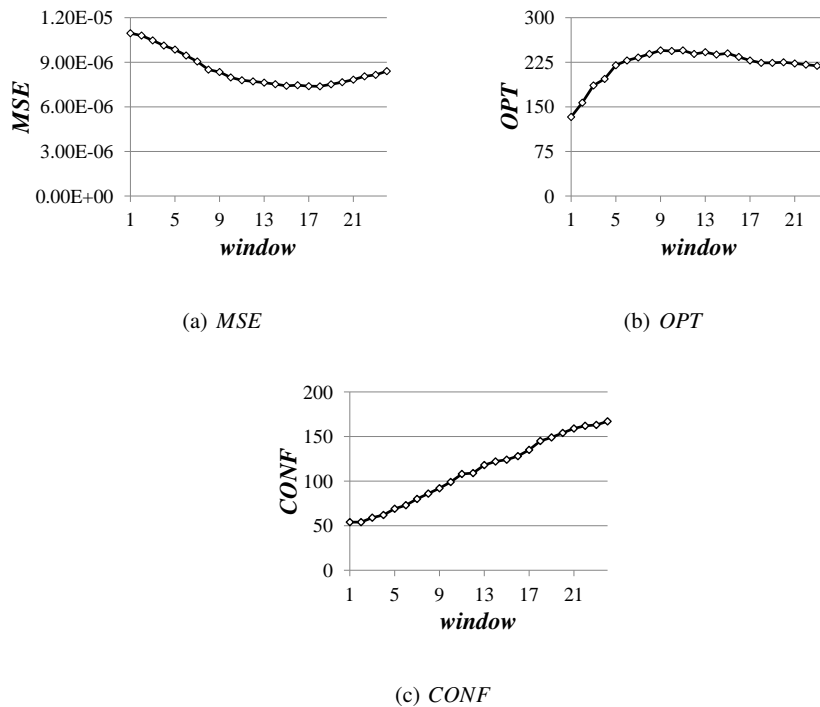


Fig. 2 Performances of event-based VaR prediction models with various time windows

Table 2 Experimental results of the performance of traditional and event-based historical VaR calculation (columns *hist* and *event*, respectively) on non-rare events

Measure	<i>window</i> = 10		<i>window</i> = event-based	
	<i>hist</i>	<i>event</i>	<i>hist</i>	<i>event</i>
<i>MSE</i>	1.1220E-05	7.6600E-06	1.1220E-05	7.2400E-06
<i>OPT</i>	96	249	100	245

for determining (individual) window sizes yields even higher improvements. For the measured *MSE* values we obtain a decrease of 35.47%, whereas 71.01% of the event-based VaR predictions outperform the traditional ones.

In order to assess the significance of the measured *MSE* improvement of 35.47%, we perform a paired two-sample one-tailed *t*-test based on SE_{hist} and SE_{event} , containing squared errors for all equities. We obtain a *p*-value of 0.0027, hereby rejecting the null hypothesis that there is no difference between the measured *MSE* values when applying a significance level of 0.05. Hence, the proposed event-based historical VaR calculation method (using non-rare events and event-based window sizes) produces more reliable VaR predictions when compared to the traditional method.

5 Conclusions

When calculating Value-at-Risk (VaR) – a widely adopted tool in the field of finance to assess portfolio risk – based on historical stock returns data, the data could be sensitive to outliers caused by seldom occurring news events in the sampled period. In order to address this shortcoming, we have proposed a way to enhance the calculation and prediction of VaR based on historical data, by removing the event-induced noise. This would enable practitioners to make better predictions of risk in terms of distributions of expected future returns.

Using a substantial data set of stock rates and news events of 2010 stemming from the proprietary ViewerPro software, we have identified rare (and hence noisy) events using a Poisson distribution. Subsequently, the event-generated noise was removed from the stock rates. From our experiments, in which we evaluated various cleansing window sizes, we can conclude that the calculation and prediction of VaR can be improved with news (i.e., extracted events and stock rates) as an additional input. Our event-based method demonstrates a significant *MSE* improvement of 35.47% compared to the traditional historical method, and outperforms the latter in 71.01% of the cases.

For future work, we suggest to investigate accounting for the type of news events, which could affect the influence of an event on equity prices (e.g., mergers could generate a larger noise than quarterly profit announcements). Another direction is related to additionally accounting for general stock market events such as financial crises, instead of only just the company specific news. Last, we would also like to build a real-life market simulation for our improved historical VaR method.

Acknowledgment

The authors are partially sponsored by the NWO Physical Sciences Free Competition project 612.001.009: Financial Events Recognition in News for Algorithmic Trading (FERNAT) and the Dutch national program COMMIT. We would like to thank Mark Vreijling and Wijand Nuij from Semlab for fruitful discussions and technical support during this research.

References

1. Andersen, T.G., Bollerslev, T., Diebold, F.X., Ebens, H.: The Distribution of Stock Return Volatility. *Journal of Financial Economics* **61**(1), 43–76 (2001)
2. Artzner, P., Delbaen, F., Eber, J.M., Heath, D.: Coherent Measures of Risk. *Mathematical Finance* **9**(3), 203–228 (1999)
3. Borsje, J., Hogenboom, F., Frasincar, F.: Semi-Automatic Financial Events Discovery Based on Lexico-Semantic Patterns. *International Journal of Web Engineering and Technology* **6**(2), 115–140 (2010)
4. Boudoukh, J., Richardson, M., Whitelaw, R.F.: The Best of Both Worlds: A Hybrid Approach to Calculating Value at Risk. *Risk* **11**(5), 64–67 (1998)
5. Byström, H.: News Aggregators, Volatility and the Stock Market. *Economics Bulletin* **29**(4), 2673–2682 (2009)
6. Chan, W.S.: Stock Price Reaction to News and No-News: Drift and Reversal After Headlines. *Journal of Financial Economics* **70**(2), 223–260 (2003)
7. Engelberg, J.E., Parsons, C.A.: The Causal Impact of Media in Financial Markets. *Journal of Finance* **66**(1), 67–97 (2009)
8. Fama, E.F.: The Behavior of Stock-Market Prices. *Journal of Business* **38**(1), 34–105 (1965)
9. Goonatilake, R., Herath, S.: The Volatility of the Stock Market and News. *International Research Journal of Finance and Economics* **3**(11), 53–65 (2007)
10. Hogenboom, A., Hogenboom, F., Frasincar, F., Kaymak, U., van der Meer, O., Schouten, K.: Detecting Economic Events Using a Semantics-Based Pipeline. In: Twenty-Second International Conference on Database and Expert Systems Applications (DEXA 2011), *Lecture Notes in Computer Science*, vol. 6860, pp. 440–447. Springer (2011)
11. Hull, J., White, A.: Incorporating Volatility Updating into the Historical Simulation Method for Value-at-Risk. *Journal of Risk* **1**(1), 5–19 (1998)
12. Ikenberry, D.L., Ramnath, S.: Underreaction to Self-Selected News Events: The Case of Stock Splits. *Review of Financial Studies* **15**(2), 489–526 (2002)
13. Kalev, P.S., Liu, W.M., Pham, P.K., Jarnecic, E.: Public Information Arrival and Volatility of Intraday Stock Returns. *Journal of Banking & Finance* **28**(6), 1441–1467 (2004)
14. Kupiec, P.H.: Techniques for Verifying the Accuracy of Risk Measurement Models. *Journal of Derivatives* **3**(2), 73–84 (1995)
15. Mitchell, M.L., Mulherin, J.H.: The Impact of Public Information on the Stock Market. *Journal of Finance* **49**(3), 923–950 (1994)
16. Rockafellar, R.T., Uryasev, S.: Conditional Value-at-Risk for General Loss Distributions. *Journal of Banking & Finance* **26**(7), 1443–1471 (2002)
17. Rosenberg, B., Reid, K., Lanstein, R.: Persuasive Evidence of Market Inefficiency. *Journal of Portfolio Management* **11**(3), 9–16 (1985)
18. Semlab: ViewerPro. <http://viewerpro.semlab.nl/> (2011)
19. Zhai, Y., Hsu, A., Halgamuge, S.K.: Combining News and Technical Indicators in Daily Stock Price Trends Prediction. In: 4th International Symposium on Neural Networks (ISSN 2007), pp. 1087–1096. Springer-Verlag Berlin, Heidelberg (2007)