Composite Value-at-Risk

Mathieu Mercadier [a,*], Jean-Pierre Lardy [b]

a: DCU Business School, Dublin City University, DUBLIN 9, IRELAND
b: JPLC SASU, 54 avenue de la Révolution, 87000 LIMOGES, FRANCE

This draft: 12 September 2023

Preliminary version. Please do not quote without the permission of the authors.

Abstract

Historical Value-at-Risk (VaR) may yield persistently high values that do not accurately represent the underlying risk, followed by dramatic drops. This phenomenon can be circumvented with some weighted historical methods like those proposed by Boudoukh et al. (BRW 1998) and Hull and White (HW 1998). In this article, we develop a simple, transparent, and quick-to-compute weighted historical VaR estimate, called composite VaR, based on the Normal and Laplace distributions. This approach provides probabilistic information on the nature of tail risk. In line with the literature, a dedicated example leads us to support the importance of the Laplace distribution for cryptocurrencies. An analysis based on common bias and clustering benchmarks is performed on this new method and its closest parents, the BRW and HW VaRs and the classical historical VaR. Although the composite VaR can be chosen for its rapidity, none of the examined methods can be unequivocally regarded as the most efficient VaR estimate. One additional key benefit of our method is its low bias.

JEL classification: C46; G32

Keywords: Risk management; Value-at-Risk; Normal distribution; Laplace distribution

Statements and Declarations: We have no conflicts of interests to disclose.

[*] Corresponding author.
E-mail addresses: mathieu.mercadier@dcu.ie (0000-0002-3531-956X), jp.lardy@jplc.fr (0009-0002-7549-9072).

1) Introduction

Utilizing classical percentile computations to assess Value-at-Risks (VaR)[1] during the COVID-19 pandemic has led to risk measurement variations that are both excessively slow and overly sharp. As explained by Engle & Manganelli (2004), this issue may mislead market finance professionals and academics with respect to timing and reactivity. Indeed, the rolling historical quantile method assumes that within a certain window, any return is equally likely, while a return outside this window has a zero probability of occurring. Consequently, the VaR of a portfolio drops dramatically as soon as an extreme return falls outside the window. To address this problem, academics devised simple weighted historical methods (Boudoukh et al., 1998; Hull & White, 1998).

In this paper, we propose a straightforward method, which we call the 'composite VaR', that provides comparable results and is both simpler and quicker to compute than the Hull & White (HW) VaR (Hull & White, 1998) and the Boudoukh, Richardson & Whitelaw (BRW) VaR (Boudoukh et al., 1998). As summarized by Abad and Benito (2013), VaRs are calculated either with historical simulation, Monte Carlo simulation, parametric methods, or relying on the extreme value theory; each with its own advantages and disadvantages. In addition to the time-weighted approaches discussed in this paper, filtered historical techniques combining historical simulation and conditional volatility have been employed (Sentana, 1995; Glosten et al., 1993; Engle & Ng, 1993; Nelson, 1991; Engle et al., 1987; Bollerslev, 1986; Engle, 1982)[2]. However, these approaches require extensive historical profits and losses, which can be costly for financial institutions to provide for parameter estimation.

The simple historical VaR is available to a large number of financial institutions and is in compliance with the Basel Committee on Banking Supervision (2019, 2009). Despite the fact that returns are generally not Normally distributed (Koundouri et al., 2016), many widely used financial tools, such as the VaR, still heavily rely on the Normal distribution due to its convenience. Therefore, in this study, we propose a parametric VaR method that assumes that returns can also arise from a more conservative distribution. The Laplace distribution was chosen over the Logistic distribution which itself can be considered as a mixture of the Normal and the Laplace distributions. Recent research has highlighted the role of the Laplace distribution for highly volatile assets, such as cryptocurrencies (Punzo & Bagnato, 2021; Hrytsiuk & Babych, 2020; Chan et al., 2017). The approaches proposed by Boudoukh et al. (1998) and Hull and White (1998) are respectively known as age-weighted and volatility-

---

[1] For some recent contributions, see e.g., Bernard et al. (2022); Hoga & Demetrescu (2022); Leung et al. (2021); Mercadier & Strobel (2021); Meng & Taylor (2020); Zhu et al. (2019); Leippold & Vasiljevic (2019); Taylor (2019); Babat et al. (2018); Chun et al. (2012); Berkowitz et al. (2011).

[2] More details are available in Abad et al. (2014).

weighted historical simulations. Our method can be defined as an absolute-return and volatility weighted historical parametric method.

The composite VaR is a risk measurement technique that possesses several advantages, such as transparency and ease of interpretation. One notable aspect of this method is the assessment of the likelihood of favoring the Normal over the Laplace distribution. This aspect can be utilized by both researchers and practitioners to extrapolate risk trends and detect periods with larger tails in the distributions. We provide an illustration that supports the usage of the Laplace distribution for cryptocurrencies. Subsequently, we provide a graphical comparison of three forgetting factor VaR estimates, namely BRW, HW, and composite VaRs, with the common historical one. Our findings indicate that none of the weighted historical methods can be definitely designated as the most efficient VaR estimate. A deeper statistical analysis conducted over a large sample emphasizes the low bias produced by our method. The composite VaR is primarily addressed to financial practitioners who wish to employ a transparent and easily understandable method to assess and manage risk.

The composite VaR is developed in section 2. Then, section 3 provides two illustrations. The first illustration focuses on the Normal versus Laplace likelihoods. The second illustration compares the performances of different VaR methods. Subsequently, a more detailed analysis is conducted to assess how bias and clustering properties are handled with each VaR methods.

2) Composite VaR combining the Normal and Laplace distributions

The location parameters μ are assumed to be zero and the scale parameters are respectively noted σ for the Normal distribution and b for the Laplace distribution. Assuming independent and identically distributed returns, noted r_t, and time weights w_t such that their sum equals 1, the maximum likelihood estimator of the Laplace scale parameter is the mean absolute deviation $\hat{b} = \sum_t w_t |r_t|$. For the Normal distribution scale parameter, the maximum likelihood estimator is its variance $\hat{\sigma}^2 = \sum_t w_t r_t^2$. Thus, to obtain a reactive VaR, we decrease the importance of the past such that weights decline by a forgetting factor $\lambda < 1$, where $w_{t-1} = \lambda w_t$, and rewrite AR(1) estimators as follows:

$$\hat{\sigma}_t = \sqrt{(1-\lambda) r_t^2 + \lambda \hat{\sigma}_{t-1}^2} \qquad \hat{b}_t = (1-\lambda)|r_t| + \lambda \hat{b}_{t-1}$$

A risk manager could consider using either the Normal or the Laplace distribution when most likely, but this would lead to discontinuities. The main idea here is to use both in proportion of their likelihoods. From the Normal and Laplace distributions, we derive the corresponding likelihoods, $\mathbb{P}_{Norm} = \left(\sqrt{2\pi e}\hat{\sigma}\right)^{-1}$ and $\mathbb{P}_{Lapl} = \left(2e\hat{b}\right)^{-1}$, with Euler's number 'e'. Thus, we write for likelihoods $\mathbb{P}$ and quantile functions $\mathbb{Q}$:

$$VaR\_cp(q) = \frac{\mathbb{P}_{Norm}\mathbb{Q}_{Norm} + \mathbb{P}_{Lapl}\mathbb{Q}_{Lapl}}{\mathbb{P}_{Norm} + \mathbb{P}_{Lapl}}$$

Which is rewritten:

$$VaR\_cp(q) = \alpha \mathbb{Q}_{Norm} + (1-\alpha)\mathbb{Q}_{Lapl}$$

Where at a given date t:

$$\alpha_t = \left(1 + \sqrt{\frac{\pi}{2e}}\frac{\hat{\sigma}_t}{\hat{b}_t}\right)^{-1}$$

The formula is finally expressed as follows:

$$VaR\_cp_t(q) = \alpha_t \mathcal{N}^{-1}(q)\hat{\sigma}_t + (1-\alpha_t)\mathcal{L}^{-1}(q)\hat{b}_t \quad (1)$$

Where, $\mathcal{N}^{-1}(q)$ and $\mathcal{L}^{-1}(q)$ respectively correspond to the Normal and the Laplace quantile functions.

3) Empirical illustrations and comparison of approaches

This section begins with two illustrations based on Bloomberg L.P. data. The first illustration in figure 1 shows the behavior of α_t for both the S&P 500 index and bitcoin, which is typically modelled using the Laplace distribution, consistent with prior studies by Punzo and Bagnato (2021); Hrytsiuk and Babych (2020); Chan et al. (2017). The second illustration in figure 2 and 3 depicts the composite VaR relative to the historical, HW and BRW VaRs during two crisis periods, namely the global financial crisis (GFC) and the COVID-19 pandemic. The two sectors studied have been severely impacted by these crises, namely the banking sector and the travel & leisure sector. Furthermore, we provide a comprehensive statistical analysis based on conventional VaR comparison metrics (Abad & Benito, 2013; Boudoukh et al., 1998; Hull & White, 1998). The chosen panel data sample is a blend of 18 equity indices, 12 currencies and 11 commodities for 16 developed and emerging countries[3]. The data set of the returns spans over twenty-four years and a half (6392 working days), from 1997-01-01 until 2021-07-01, and is provided entirely by Refinitiv Datastream.

3.1) The Normal vs. Laplace likelihood illustration

The figure displayed on the left-hand side of figure 1 highlights the fatter negative tails of bitcoin (vs. USD) returns against those of the S&P 500, apart from the onset of the COVID-19 pandemic and the volatile period in early February 2018. These fatter negative tails are emphasized in the figure on the right-hand side, where α_t is significantly below 50%, even reaching around 40% for bitcoin It suggests that the data is more likely to fit a Laplace distribution, which is consistent with previous studies by Punzo & Bagnato (2021); Hrytsiuk & Babych (2020); Chan et al. (2017).

Insert Figure 1

This transparent α_t parameter can prove beneficial for both academics and practitioners in measuring risk trends, primarily in identifying periods of fatter tails in the distributions.

[3] The United States, Japan, the Eurozone, Germany, France Spain, the United Kingdom, Canada, Australia, Hong Kong, China, India, Brazil, Mexico, South Africa & Poland.

3.2) Illustrations of four VaR approaches

The study examines the behavior of four VaR methods for two industrial sectors that were severely impacted by the GFC and the COVID-19 pandemic. For rolling windows of 260 days, $\lambda = 0.98$[4] and a 1% confidence level, the daily VaR is computed using the historical (percentile) VaR and three weighted historical methods: the composite VaR, the HW VaR and the BRW VaR. Figure 2 highlights the case of the Euro Stoxx Bank index during the last twenty years (rhs) and focuses on the GFC (lhs), which caused immense pressure on the European banking sector. Similarly, the COVID-19 pandemic is illustrated using the Euro Stoxx Travel & Leisure index as shown in figure 3 (lhs).

Insert Figure 2

Insert Figure 3

Figures 2 and 3 illustrate that the three time-weighted measures exhibit faster responsiveness compared to the historical method.

Nevertheless, all four methods can be local maxima and minima of the VaR, creating difficulty in selecting a single measure among them. Despite this challenge, our analysis indicates that when the risk escalates, the composite method exhibits greater reactivity compared to the historical VaR. Furthermore, the composite VaR avoids undesirable peaks that are likely to be encountered using the HW method and, to a lesser degree, the BRW method.

3.3) Comparison of four VaR approaches

Following a statistical presentation of the returns from the three asset classes under examination, we conduct a more comprehensive analysis of our approach by utilizing three benchmarks. These benchmarks comprise two widely employed metrics (Boudoukh et al., 1998; Hull & White, 1998), as well as a simple conditional probability measure. Each of these benchmarks evaluates relevant properties for risk measures.

The data set for this empirical analysis, composed of three common asset classes (i.e., equities, currencies and commodities), is summarized in the descriptive statistical table 1.

Insert Table 1

As expected, the means of the returns are close to zero and the standard deviations are on an annual basis respectively equal to 22% (Equity), 11% (Currency) and 31% (Commodity), in line with the ranges between the maximum and minimum values. It underlines that commodities are more volatile than the other classes. In addition, the Normality hypothesis was tested with the Jarque-Bera test and rejected for all individual assets.

Typically, the probability of observing a return lower than the computed VaR should be equal to the confidence level (Kupiec, 1995), otherwise a bias occurs. In line with Abad &

[4] Hull & White (1998) also set λ at 0.98 while Boudoukh et al. (1998) tested both $\lambda = 0.99$ and $\lambda = 0.97$.

Benito (2013); Boudoukh et al. (1998); Hull & White (1998), we begin by dividing the number of tail excessions – that is, the frequency with which tomorrow's return exceeds today's VaR - by the number of observations. We refer to this metric as the Excession Unconditional Probability (EUP), which is expressed as follows:

$$EUP = \mathbb{P}(excession) = \frac{\sum_{t=1}^{T} I_t}{T}$$

Where,

$$I_t = \begin{cases} 1 & r_{t+1} < -VaR_t(\alpha), \\ 0 & otherwise. \end{cases}$$

Then, we evaluate the responsiveness of our measure by introducing a conditional probability approach that considers the likelihood of a tail event occurring given that the VaR was exceeded on the previous date. This method is referred to as the Excession Conditional Probability (ECP).

$$ECP = \mathbb{P}(excession|excession\ at\ t-1) = \frac{\mathbb{P}(excession \cap excession\ at\ t-1)}{\mathbb{P}(excession\ at\ t-1)}$$

After the first indicator measuring the bias, Boudoukh et al. (1998) identify another negative attribute in the fact that tail events may gather. They therefore propose to appraise clustering with the Mean Absolute Percentage Error (MAPE). For their part, Hull & White (1998) state that the MAPE measures both bias and clustering, which is also true for our conditional probability measure. In a nutshell, the MAPE provides the average of the absolute value of the number of times the VaR is exceeded in a rolling window of 100 observations minus 1 or 5, respectively at a 1% or 5% level of confidence.

The findings of the benchmark analysis are presented in table 2. The results on the entire sample are displayed at the top of the table, followed by a breakdown of the results by asset class[5].

Insert Table 2

At a 5% confidence level, the composite VaR exhibits a lower level of bias (i.e., lower EUP) compared to all other methods. Moreover, the composite VaR demonstrates similar levels of ECP and MAPE to those of the BRW and HW VaRs, while outperforming the historical VaR in terms of these metrics. These findings are consistent on the entire sample and across all asset classes under consideration.

According to table 2, our analysis reveals that, at a 1% confidence level, the composite VaR exhibits comparable performance to the BRW and HW VaRs across all metrics and samples under considerations. Furthermore, all the aforementioned approaches demonstrate superior performance compared to the historical VaR. Nonetheless, figures 2 and 3 (cf. section

[5] More detailed tables are available upon request to the authors, which includes a comparison of results during crisis and non-crisis periods.

3.2) clearly depict that at a 1% confidence level, the composite VaR declines more rapidly than the historical VaR while also circumventing the undesirable peaks achieved by the BRW and HW methods. This reduced propensity of having high spikes is associated with the fact that both the ECP and MAPE metrics highlight the slight tendency of the composite VaR to cluster excessions. In summary, while the composite VaR's ECP and MAPE are higher than those of the BRW and HW VaRs, they remain lower than those of the historical VaR.

Our study indicates that the three weighted historical VaR metrics examined perform better than the commonly used historical VaR with respect to both bias and clustering. In addition to its efficiency in terms of speed, the composite VaR is effective in terms of bias, albeit demonstrating a slightly higher propensity to cluster excessions compared to the HW and BRW VaRs.

4) Conclusion

In this study, we propose a novel weighted historical VaR estimate that is simple, transparent, and computationally efficient. The method is based on combining the Normal and Laplace distributions. One of its core components offers insight into the nature of the tail risk and quantifies the likelihood of favoring either the Normal distribution or the Laplace distribution. To evaluate the effectiveness of our proposed VaR, we conduct an analysis based on common bias and clustering benchmarks. We compare it with the age-weighted and volatility-weighted historical methods proposed by Boudoukh et al. (1998) and Hull & White (1998) and the classical historical one.

Although our new approach offers rapidity, none of the evaluated methods can be definitely deemed as the most efficient VaR estimate. One notable advantage of our method is its low bias.

Acknowledgements

We are grateful to Thomas Charbonniaud, Zineb El Malki & Tanoh Franck Fabien Ziké for their preliminary tests, and to Camille Lardy for proof-reading this article.

References

Abad, P., & Benito, S. (2013). A detailed comparison of value at risk estimates. *Mathematics and Computers in Simulation*, *94*, 258-276.

Abad, P., Benito, S., & Lopez, C. (2014). A comprehensive review of Value at Risk methodologies. *The Spanish Review of Financial Economics*, *12*(1), 15-32.

Babat, O., Vera, J. C., & Zuluaga, L. F. (2018). Computing near-optimal Value-at-Risk portfolios using integer programming techniques. *European Journal of Operational Research*, *266*(1), 304-315.

Basel Committee on Banking Supervision. (2009). *Revisions to the Basel II market risk framework* (Bank for International Settlements).

Basel Committee on Banking Supervision. (2019). *Minimum capital requirements for market risk* (Bank for International Settlements).

Berkowitz, J., Christoffersen, P., & Pelletier, D. (2011). Evaluating Value-at-Risk Models with Desk-Level Data. *Management Science*, *57*(12), 2213-2227.

Bernard, C., Perchiazzo, A., & Vanduffel, S. (2022). Implied value-at-risk and model-free simulation. *Annals of Operations Research*.

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, *31*(3), 307-327.

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

Chan, S., Chu, J., Nadarajah, S., & Osterrieder, J. (2017). A Statistical Analysis of Cryptocurrencies. *Journal of Risk and Financial Management*, *10*(2), 1-23.

Chun, S. Y., Shapiro, A., & Uryasev, S. (2012). Conditional Value-at-Risk and Average Value-at-Risk : Estimation and Asymptotics. *Operations Research*, *60*(4), 739-756.

Engle, R. F. (1982). Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. *Econometrica*, *50*(4), 987-1007.

Engle, R. F., Lilien, D. M., & Robins, R. P. (1987). Estimating Time Varying Risk Premia in the Term Structure : The Arch-M Model. *Econometrica*, *55*(2), 391-407.

Engle, R. F., & Manganelli, S. (2004). CAViaR: Conditional Autoregressive Value at Risk by Regression Quantiles. *Journal of Business & Economic Statistics*, *22*(4), 367-381.

Engle, R. F., & Ng, V. K. (1993). Measuring and Testing the Impact of News on Volatility}. *The Journal of Finance*, *48*(5), 1749-1778.

Glosten, L. R., Jagannathan, R., & Runkle, D. E. (1993). On the Relation between the Expected Value and the Volatility of the Nominal Excess Return on Stocks. *The Journal of Finance*, *48*(5), 1779-1801.

Hoga, Y., & Demetrescu, M. (2022). Monitoring Value-at-Risk and Expected Shortfall Forecasts. *Management Science*, *In press*, 1-18.

Hrytsiuk, P., & Babych, T. (2020). The cryptocurrencies risk measure based on the Laplace distribution. *Proceedings of the Selected Papers of the Special Edition of International Conference on Monitoring, Modeling & Management of Emergent Economy (M3E2-MLPEED 2020), Odessa, Ukraine, CEUR-WS.org*, *2713*(25), 261-276.

Hull, J., & White, A. (1998). Incorporating volatility updating into the historical simulation method for VaR. *Journal of Risk*, *1*, 5-19.

Koundouri, P., Kourogenis, N., & Pittis, N. (2016). Statistical modeling of stock returns : Explanatory or descriptive ? A historical survey with some methodological reflections. *Journal of Economic Surveys*, *30*(1), 149-164.

Kupiec, P. H. (1995). Techniques for verifying the accuracy of risk measurement models. *The Journal of Derivatives*, *3*(2), 73-84.

Leippold, M., & Vasiljevic, N. (2019). Option-Implied Intrahorizon Value at Risk. *Management Science*, *66*(1), 397-414.

Leung, M., Li, Y., Pantelous, A. A., & Vigne, S. A. (2021). Bayesian Value-at-Risk backtesting : The case of annuity pricing. *European Journal of Operational Research*, *293*(2), 786-801.

Meng, X., & Taylor, J. W. (2020). Estimating value-at-risk and expected shortfall using the intraday low and range data. *European Journal of Operational Research*, *280*(1), 191-202.

Mercadier, M., & Strobel, F. (2021). A one-sided Vysochanskii-Petunin inequality with financial applications. *European Journal of Operational Research*, *295*(1), 374-377.

Nelson, D. B. . (1991). Conditional Heteroskedasticity in Asset Returns : A new approach. *Econometrica*, *59*(2), 347-370.

Punzo, A., & Bagnato, L. (2021). Modeling the cryptocurrency return distribution via Laplace scale mixtures. *Physica A: Statistical Mechanics and its Applications*, *563*(125354), 1-18.

Sentana, E. (1995). Quadratic ARCH Models. *The Review of Economic Studies*, *62*(4), 639-661.

Taylor, J. W. (2019). Forecasting Value at Risk and Expected Shortfall Using a Semiparametric Approach Based on the Asymmetric Laplace Distribution. *Journal of Business \& Economic Statistics*, *37*(1), 121-133.

Zhu, B., Ye, S., He, K., Chevallier, J., & Xie, R. (2019). Measuring the risk of European carbon market : An empirical mode decomposition-based value at risk approach. *Annals of Operations Research*, *281*, 373-395.

List of tables

	Equity	Currency	Commodity
No. of dates	6392	6392	6392
No. of assets	18	12	11
Mean	3.4E-04	7.9E-05	3.3E-04
Std. Dev.	0.014	0.007	0.020
Min	-0.158	-0.106	-0.434
Max	0.334	0.176	0.788

Notes. This table shows the average of the mean and standard deviation of each asset class daily returns and the overall minimum and maximum.

Table 1: Descriptive statistics summary

Confidence level		1%			5%		
Sample	VaR	EUP	ECP	MAPE	EUP	ECP	MAPE
All	Hist	0.0154	0.0666	1.3082	0.0553	0.1067	3.2187
	BRW	0.0133	0.0275	0.8516	0.0513	0.0745	1.8853
	HW	0.0101	0.0368	0.8388	0.0503	0.0827	2.2428
	Comp	0.0106	0.0504	0.954	0.0451	0.0796	2.3147
Equity	Hist	0.0157	0.0758	1.3676	0.0557	0.118	3.4153
	BRW	0.0135	0.0306	0.9071	0.0515	0.0836	2.0322
	HW	0.0107	0.0434	0.9051	0.0505	0.0951	2.4359
	Comp	0.0125	0.0635	1.065	0.0489	0.0934	2.4446
Currency	Hist	0.0152	0.059	1.2858	0.0547	0.0969	3.1429
	BRW	0.0127	0.0256	0.7802	0.0505	0.0624	1.7428
	HW	0.01	0.0229	0.7879	0.05	0.0638	2.0495
	Comp	0.0083	0.0325	0.8637	0.0416	0.0626	2.2237
Commodity	Hist	0.0152	0.0598	1.2354	0.0553	0.099	2.9796
	BRW	0.0135	0.0246	0.8386	0.0518	0.0726	1.8005
	HW	0.0093	0.0411	0.7859	0.0503	0.0832	2.1377
	Comp	0.0099	0.0485	0.8709	0.0429	0.0756	2.2014

Notes. This table shows the means of the benchmark values overall and for each asset class.

Table 2: Summary of the metrics at a 1% and 5% confidence levels

List of figures

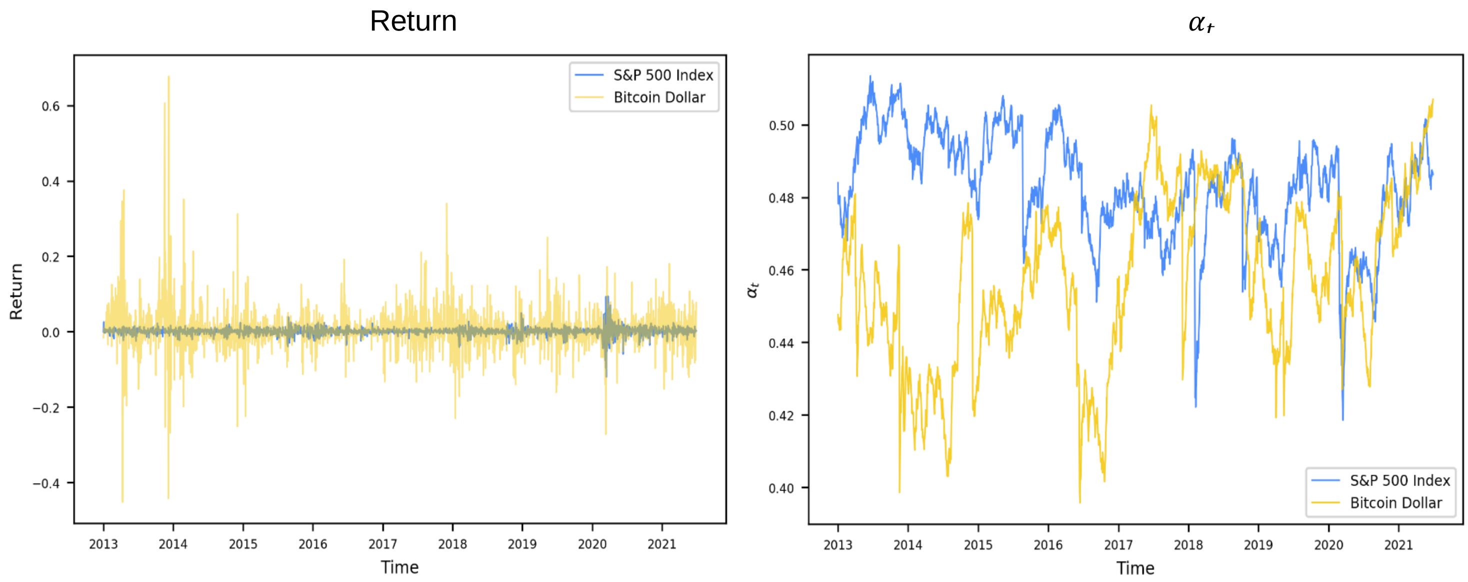

Notes. These charts display the daily returns (lhs) and α_t (rhs) of the S&P 500 Index (blue) and bitcoin dollar (yellow). An $\alpha_t < 0.5$ overweights the Laplace distribution.

Figure 1: Bitcoin/USD and S&P 500 Index

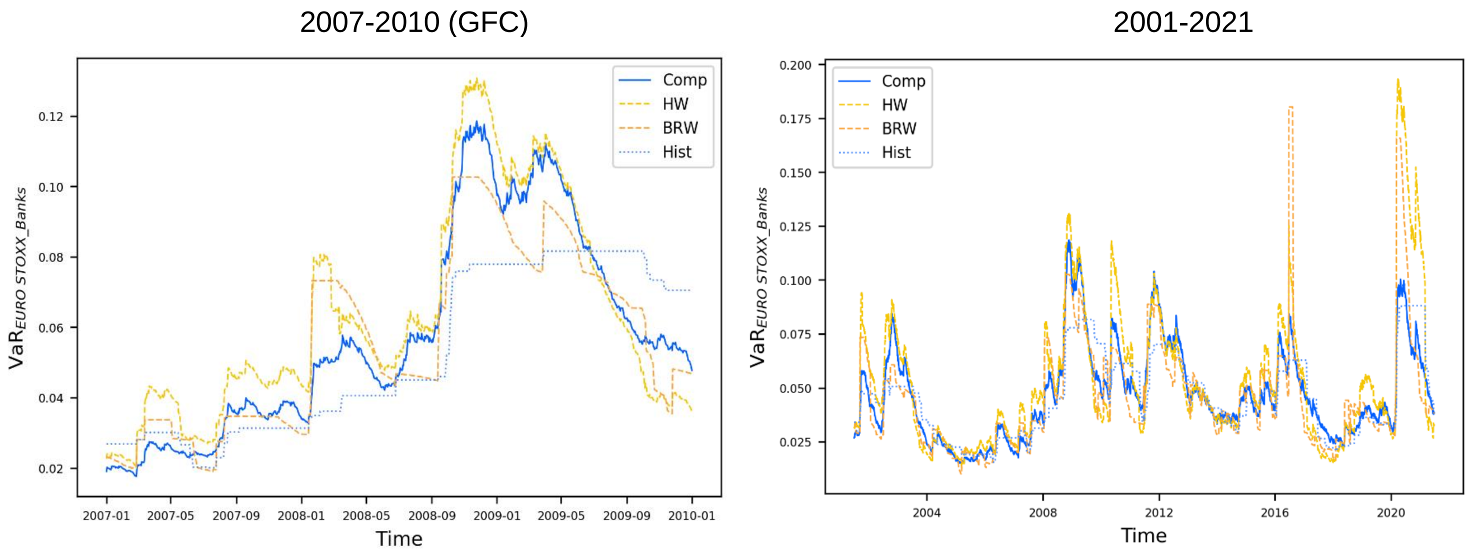

Notes. These graphs show the different daily VaR methods for the sector index: Euro Stoxx Bank during the GFC (lhs) and since mid-2001 (rhs) at a 1% confidence level. The composite method is more reactive than the historical VaR and avoids the undesirable peaks of the BRW VaR & HW VaR.

Figure 2: VaR Methods: Euro Stoxx Bank

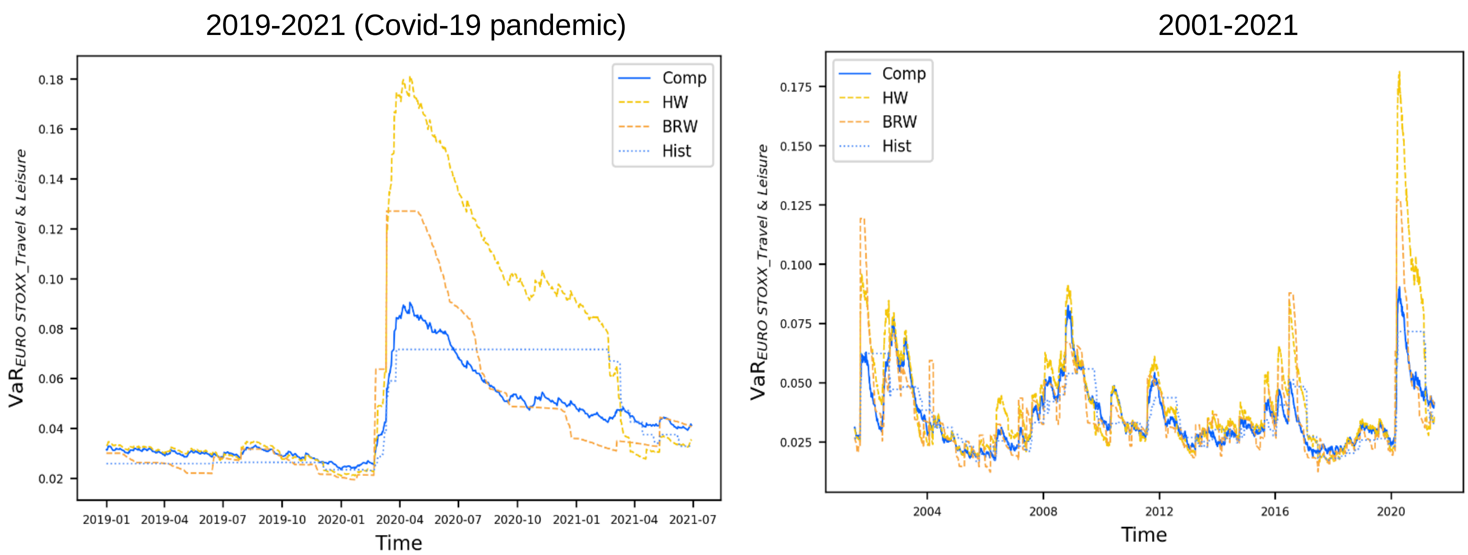

Notes. These graphs show the different daily VaR methods for the sector index: Euro Stoxx Travel & Leisure during the COVID-19 pandemic (lhs) and since mid-2001 (rhs) at a 1% confidence level. The composite method is more reactive than the historical VaR and avoids the undesirable peaks of the BRW VaR & HW VaR.

Figure 3: VaR Methods: Euro Stoxx Travel & Leisure