sciendo

Economics and Business Review
https://www.ebr.edu.pl

Vol. 9 (2), 2023: 160–183
https://doi.org/10.18559/ebr.2023.2.738

How to fly to safety without overpaying for the ticket

Tomasz Kaczmarek[1]   Przemysław Grobelny[2]

Abstract

For most active investors treasury bonds (govs) provide diversification and thus reduce the risk of a portfolio. These features of govs become particularly desirable in times of elevated risk which materialize in the form of the flight-to-safety (FTS) phenomenon. The FTS for govs provides a shelter during market turbulence and is exceptionally beneficial for portfolio drawdown risk reduction. However, what if the unsatisfactory expected return from treasuries discourages higher bonds allocations? This research proposes a solution to this problem with Deep Target Volatility Equity-Bond Allocation (DTVEBA) that dynamically allocate portfolios between equity and treasuries. The strategy is driven by a state-of-the-art recurrent neural network (RNN) that predicts next-day market volatility. An analysis conducted over a twelve year out-of-sample period found that with DTVEBA an investor may reduce treasury allocation by two (three) times to get the same Sharpe (Calmar) ratio and overperforms the S&P500 index by 43% (115%).

Keywords

- asset allocation strategy
- target volatility
- flight-to-safety
- recurrent neural networks
- machine learning

JEL codes: C32, C45, C58, G11, G15

Article received 15 March 2023, accepted 16 June 2023.

This research was funded in part by National Science Centre, Poland, grant number 2021/41/N/HS4/02344. For the purpose of Open Access, the author has applied a CC-BY public copyright licence to any Author Accepted Manuscript (AAM) version arising from this submission.
The project was financed within the Regional Initiative for Excellence programme of the Minister of Science and Higher Education of Poland, years 2019-2022, grant no. 004/RID/2018/19, financing 3,000,000 PLN.

[1] Department of Investment and Financial Markets, Institute of Finance, Poznań University of Economics and Business, al. Niepodległości 10, 61-875 Poznań, Poland, corresponding author: tomasz.kaczmarek@ue.poznan.pl, https://orcid.org/0000-0003-1828-100X.

[2] Department of Investment and Financial Markets, Institute of Finance, Poznań University of Economics and Business, al. Niepodległości 10, 61-875 Poznań, Poland, przemyslaw.grobelny@ue.poznan.pl, https://orcid.org/0000-0003-1453-2844.

Suggested citation: Kaczmarek, T., & Grobelny, P. (2023). How to fly to safety without overpaying for the ticket. *Economics and Business Review*, *9*(2), 160–183. https://doi.org/10.18559/ebr.2023.2.738

Introduction

Flight-to-safety (FTS) is a financial market phenomenon occurring when investors reallocate portfolios from higher-risk investments (equity) to the safer alternatives such as high-grade government bonds (govs). FTS occurs during equity market turmoil and results in a negative temporal correlation between long-term bonds and equities (Baur & Lucey, 2009). This unique diversification benefit makes govs a desirable portfolio component. Unfortunately expected returns for govs are usually lower than for equities and sometimes even negative. The worst comes when interest rates are raised. Therefore, depending on investor preferences and market expectations investment in long-term bonds can be perceived either as an attractive asset class or as a cost of portfolio insurance against equity market turmoil.

This study aims to build a tactical allocation of equities and govs that simultaneously leverages long-term bonds' diversification benefits and reduces a bond investment costs. In other words it proposes how to reduce the average portfolio allocation in govs by shrinking their allocation on the stable market and increasing it during FTS periods. To achieve that the research concentrates on the relationship between FTS and market volatility. It is well documented that in times of elevated risk when investors fly to safety that the volatility of equity increases (Baele et al., 2020; Beber et al., 2009; Longstaff, 2004). Therefore accurate predictions for high (low) periods of market volatility can provide a signal for high (low) govs allocation.

The success of this strategy lies in the accurate prediction of market volatility. Predictability becomes especially important when markets are not stable which is typical for FTS episodes (Grabowski et al., 2023). Kaczmarek et al. (2022) predict market volatility with a multivariate recurrent neural network (RNN) and define market state with volatility predictions. They show that govs present safe-haven characteristics in periods of elevated market volatility. This research extends this direction and introduces Deep Target Volatility Equity-Bond Allocation (DTVEBA) that targets the desired level of equity volatility with RNN predictions and dynamically allocates portfolios between equity and bonds. With analysis conducted over a twenty year sample period it was found that RNN delivers sound predictions to reduce treasury allocation while maintaining its diversification benefits.

The contribution of this research is twofold. First, it extends studies about FTS events by demonstrating how to improve portfolio mean-variance relation by exposing it to FTS events when means for govs are low. Although the relationship between expected returns for equity and govs is conditional on multiple factors the benefits of having govs in the portfolio remain constant because of their positive impact on portfolio returns during FTS periods. This study utilizes deep recurrent neural networks to predict periods of high market volatility that may result in FTS events and allocate more heavily to govs only when it is necessary to provide shelter to the portfolio. Second, it contributes to portfolio risk management, specifically to volatility-targeting studies. It shows that RNN-based volatility forecasts create more efficient target volatility portfolios and effectively detect periods of high volatility when equity allocations should be low.

To this end the DTVEBA strategy is tested with the S&P500 index with twelve years of daily historical prices in the out-of-sample setting. Compared to fixed equity bond allocation an investor using the DTVEBA strategy may reduce treasury weight by two (three) times to get the same Sharpe (Calmar) ratio. Furthermore and depending on the desired equity volatility, the strategy overperforms the S&P500 index by 29%–68% or 106%–155% in terms of the Sharpe and Calmar ratios, respectively. Finally, the strategy delivers substantial benefits during a market stress period. During all the most significant drawdowns in the testing sample (2010–2021), the DTVEBA strategy overperforms the S&P500 index.

The article is organised as follows: Section 1 reviews the literature related to the role of volatility in portfolio risk management and demonstrates the theory that forms the basis for the DTVEBA strategy. Section 2 explains the model used to dynamically allocate portfolios between equity and govs based on predicted market volatility and demonstrates data used for the empirical study. In Section 3 the accuracy of volatility predictions and the performance of the DTVEBA strategy is discussed. The article concludes with recommendations for investors and portfolio managers, study limitations, and directions for further research.

1. Literature review

As evidenced by Fleming et al. (2001) volatility commands a pivotal position in portfolio risk management. Researchers provide empirical support for the superior performance of short-term volatility timing strategies over static portfolios with identical return and volatility expectations. Put differently their work attests to the substantial financial benefits that can be reaped from incorporating volatility forecasts into portfolio design. Moreira and Muir (2017) extend their findings by building volatility-managed portfolios that increase Sharpe ratios because of the unproportional relationship between changes in volatility and expected returns.

Market volatility forecasts significantly impact the performance of volatility-managed portfolios. Several widely accepted methodologies exist for nonlinear volatility modelling and prediction. ARCH models which revolve around daily returns are devised to gauge underlying market volatility (Bollerslev, 1986; Glosten et al., 1993). Another technique capitalizes on implied volatility which is contingent on option pricing. In this instance past returns or volatility are rendered unnecessary; nevertheless in selecting an appropriate option-pricing model becomes essential.

The third technique involves using measures that operate on intraday returns estimating unobserved integrated variance (Andersen & Bollerslev, 1998; Hansen & Lunde, 2005). Fleming et al. (2003) demonstrate the substantial economic value of applying the realized volatility in the context of investment decisions. Even though Realized Variance (RV) is a widely employed measure (Będowska-Sójka, 2018) its sensitivity to jumps in the variance is a recognized limitation. Consequently Barndorff-Nielsen and Shephard (2004) proposed a Realized Bipower Variation (RBV) alternative metric. Like its predecessor this measure relies on intraday returns but presents robustness against jumps. The standard approach in modelling any type of realized variance is to use ARFIMA models.

Each of these methods is not devoid of limitations because they demand varying assumptions concerning the distribution of the data at hand (Poon & Granger, 2005). Therefore, an alternative approach is to apply data mining or machine learning techniques for financial market modelling. Recurrent neural network (RNN) models, in particular, have been found to exhibit superior fitting capabilities on financial data series compared to parametric models. This is due to their inherent ability to decipher intricate data patterns without any pre-assumptions (Christensen et al., 2021; H. Y. Kim & Won, 2018; Y. Kim & Enke, 2018).

Kim and Enke (2016, 2018) construct a dynamic allocation strategy between equity and cash based on volatility forecasts. They compare different volatility prediction methods and demonstrate high economic gains in applying univariate recurrent neural networks to predict implied market volatility. However, Becker et al. (2015) visualize that multivariate models improve the volatility forecasting accuracy for portfolio allocation. Similarly in research on neural network-based modelling of variability H. Y. Kim and Won (2018) propose a hybrid long short-term memory (LSTM) multivariate model combining LSTM with various generalized autoregressive conditional heteroscedasticity (GARCH)-type models. Their solution outperforms all traditional volatility prediction techniques demonstrating recurrent neural networks‘ supremacy in predicting market volatility.

Applying market volatility forecasts to allocate a portfolio dynamically between equities and cash is an active field of research (Y. Kim & Enke, 2016, 2018; Perchet et al., 2016). This strategy's enhancement presents methods for reducing transaction costs through conditional allocation changes (Bongaerts et al., 2020; Zakamulin, 2019). Furthermore, Kaczmarek et al. (2022) use market volatility forecasts to determine the low and high market volatility states. They show that out of thirteen potential safe haven assets only govs demonstrate a negative correlation with equities in highly volatile markets and reveal safe haven properties. However, to the best of the authors' knowledge none of the related studies use volatility predictions to allocate portfolios between equity and govs. This study fills the gap by demonstrating the benefits of a dynamic allocation in both equity and govs conditioned on multivariate recurrent neural network predictions of stock market RBV.

The literature has intensely scrutinized the correlation between the returns on govs and equity. It results largely from the fact that these two assets are considered not only as complementary but also substitutes and both the level and dynamics of their return correlation are essential elements for asset allocation decisions (Boucher & Tokpavi, 2019). FTS periods are usually associated with substantial yet short-lived fluctuations in expected returns on equities and bonds. These changes are typically surrounded by active trading and/or risk transfer between different investors (Lehnert, 2022). As a result the demand for treasuries is growing in periods when investors, fearing the increasing volatility in the market, rebuild their portfolios towards less risky positions.

Due to the lower expected returns on govs rather than on equity investors pay the opportunity cost for implementing this strategy. In addition opposing demand and supply generated for stocks and bonds due to active FTS trading (followed by a move in the opposite direction) exacerbates the differences in expected returns that persist under low or standard volatility in the market. Thus the less dynamically investors react to market condition changes and the slower they restore the portfolio composition to the target volatility level, the greater the opportunity cost may be.

Empirical applications using data for U.S. govs and the S&P500 index show that when yields are low the strength of FTS from stocks to bonds weakens (Boucher & Tokpavi, 2019). Adrian et al. (2019) suggest that the effect of the FTS weakens also when market volatility is high. This may indicate that the implementation of risk-reducing strategies by investors is limited due to the opportunity cost and the proposed model based on reducing this factor by shortening the portfolio allocation period in safe-haven assets may contribute to the optimization of the FTS strategy.

2. Material and methods

The Deep Target Volatility Equity-Bond Allocation (DTVEBA) strategy is based on the target volatility (T.V.) framework supported by recurrent neural networks volatility predictions. The weight of equity (ω_e) in the target volatility approach is expressed by:

$$\omega_e = \min\left(\frac{\sigma_p}{\sigma_e}, 100\%\right) \tag{1}$$

where σ_p is target volatility and σ_e stands for standard deviation of the risky asset (Hocquard et al., 2013; Perchet et al., 2016). Equation (1) says that the weight of a risky asset is equal to the proportion between required (target) monthly volatility and the predicted volatility. The higher the forecasted volatility, the lower the equity weight. Typical T.V. strategy assumes that free cash ($100\% - \omega_e$) is allocated in a risk-free instrument. To build the DTVEBA strategy, this approach is modified, and equity (ω_e) is combined with long-term government bonds position ($100\% - \omega_e$). Long-term bonds are not risk-free assets but rather demonstrate flight-to-safety characteristics (Kaczmarek et al., 2022). In this way, DTVEBA is strongly (weakly) exposed to long-term bonds in periods of high (low) expected volatility.

The performance of T.V. strategies relies on volatility prediction accuracy. Kaczmarek et al. (2022) compare next-day bi-powered volatility prediction accuracy for five time-series methods by including three econometric models (ARFIMA, GARCH, GJR-GARCH) and two recurrent neural networks (univariate and multivariate versions of gated recurrent units (GRU). They demonstrate the supremacy of the multivariate GRU method that uses six explanatory variables: bi-powered realized volatility, S&P500 index, gold, crude, U.S. 3-year govs, and U.S. 3-year A.A. graded corporate bonds. This research proposes an extension to their approach that improves prediction accuracy by adding eight new explanatory variables to predict next-day bi-powered volatility.[3] First, it adds a volatility risk premium expressed with the relationship between the VIX index and realized volatility (Prokopczuk & Wese Simen, 2014). Second, the neural nets are trained with seven additional variates derived from econometric forecasting methods, namely: 1) adjusted conditional volatility from GARCH(1,1); 2) adjusted residuals from GARCH(1,1); 3) adjusted conditional volatility from EGARCH; 4) leverage effect EGARCH; 5) leverage EGARCH;

[3] Table 1 reports the DTVECA_10 strategy Sharpe ratio of 1.05 (2010–2021). Kaczmarek et al. (2022) reports Sharpe ratio for comparable strategy of 0.97 (Table 4, STRAT_CLEAN, 2010–2020).

6) adjusted conditional volatility from EWMA; and 7) adjusted residuals from EWMA (Hyup Roh, 2007; H. Y. Kim & Won, 2018).[4]

The simulation of the Deep Target Volatility Equity-Bond Allocation (DTVEBA) strategy is based on daily data from April 11, 2002, to December 31, 2021 which is split for a training period (up to 2009) and an out-of-sample testing period (2010-2021) with the yearly extending window. The data for realized bipower variation (RBV) are from the Oxford-Man Institute of Quantitative Finance Realized Library (Heber et al., 2009) where the sample starts from 2000. The data for explanatory variables are from Refinitiv Datastream.[5]

3. Results and discussion

The empirical experiment demonstrates how to dynamically allocate government bonds to portfolios based on predicted market volatility. The effectiveness of the proposed dynamic govs' allocation depends on the precision of the volatility forecasts. Therefore, firstly the quality of the applied forecasting model is demonstrated. Table 1 visualizes the results of Diebold-Mariano (D.M.) equal forecast accuracy tests for daily volatility of SPX500PI Index measured with realized bi-powered variation (BPVSPX500) and compares prediction accuracy for mean squared errors (MSE). Each row/column represents a different prediction method: 1) the base prediction method used to create DTVEBA strategy with all variates (ALL-VARIATES); 2) the same prediction model trained without hybrid, GARCH type input variates, namely COND_VOL_GARCH, RESID_GARCH, COND_VOL_EGARCH, LEV_EFFECT_EGARCH, LEV_EGARCH, COND_VOL_EWMA, RESID_EWMA (NO GARCH); 3) the same prediction model trained without VOL_RISK_PREM (NO RP); 4) the same prediction model trained without both hybrid, GARCH type input variates and

[4] Description of each variate and the Gated Recurrent Unit specification is available in appendix A and B, respectively.

[5] This study predicts the daily RBV of S&P500 Index. It follows variates selections for predicting realized daily market volatility after H. Y. Kim and Won (2018) who predict the next day volatility of KOSPI 200 stock index returns with 1) realized volatility of KOSPI 200 stock index returns, 2) KOSPI 200 INDEX log difference, 3) 3-year Korea Treasury Bond interest rate, 4) 3-year AA-grade corporate bond interest rate, 5) gold, 6) crude oil, 7) variates derived from GARCH and EGARCH models. The training period of this study is limited with data available for the variate with the shortest history. The Oxford-Man Realized Library (Heber et al., 2009) delivers data for realized volatility for major stock indices from 2000. Still the training period is shortened due to the limitation of data for the U.S. 3-year A.A. graded corporate bonds that were published from April 11, 2002. The inclusion of U.S. 3-year A.A. graded corporate bonds reduces the data sample only by two years and 3.5 months (or around 10% of the whole data sample) and has no significant impact on the study results.

VOL_RISK_PREM (NO GARCH & R.P.); and 5) naïve prediction based on the previous day observed realized bi-powered variation (NAÏVE).

Table 1. Diebold-Mariano (D.M.) equal forecast tests

	NO GARCH	NO RP	*NO GARCH & RP*	NAÏVE
ALL-VARIATES	2.4**	0.2	2.2**	3.3***
NO GARCH		–2.3**	0.7	2.5***
NO RP			2.2**	3.3***
NO GARCH & RP				2.3**

*** and ** denote a rejection of the null hypothesis at the 1% and 5% significance level—respectively.

Source: Own work.

The D.M. test show the significantly higher efficiency of neural network methods than NAÏVE. This is consistent with Hamid and Iqbal (2004) and Brooks (1998) who show significant economic benefits in using neural networks to forecast market volatility. Furthermore, when comparing the base hybrid prediction model combining GRU with GARCH (ALL-VARIATES) with the pure GRU model (NO-GARCH) the errors of the hybrid model are smaller than in the single models. It means that GRU can effectively learn temporal patterns of time-series data and the long-term phenomenon provided with input from GARCH and EGARCH models. The results from the S&P Index from the U.S. market supports the earlier finding of Hyup Roh (2007) and H. Y. Kim and Won (2018) who show a similar effect observed on the Korean market with the KOSPI index.

In contrast, the explanatory power of volatility risk premium turns out to be limited. Although Prokopczuk and Wese Simen (2014) find volatility risk premium as an essential determinant in predicting market volatility they concentrate on implied volatility. This research forecasts realized volatility because it demonstrates higher portfolio application usage (Fleming et al., 2003).[6] Thus the role of the risk premium in forecasting volatility is not constant and depends on the type of volatility being forecasted.

Figure 1 shows how closely the predicted daily volatility tracks the realized volatility. Gray shadows indicate days on which realized volatility is significantly higher than forecasts. Nevertheless volatility forecasts still closely follow realized volatility.

In DTVEBA the predicted volatility is used to dynamically allocate portfolio between equity and govs. The descriptive statistics for SPX500TR and USGOV10TR are demonstrated in Table A.2 in the appendix. The mean daily returns for SPX500TR are three times higher in the testing vs. training sample.

[6] Also results from Table 2 demonstrate that use of realized volatility instead of implied volatility increases performance of the strategy that allocates portfolio between equity and govs.

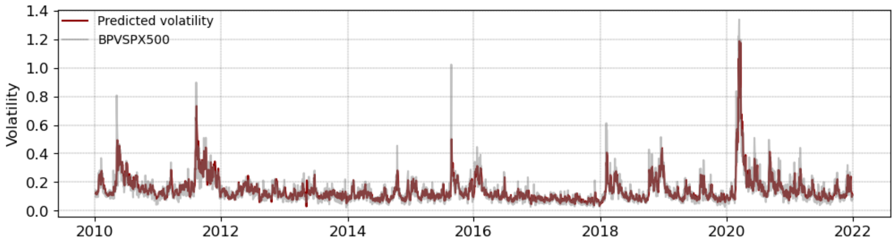

Figure 1. Predicted and realized daily volatility

Source: Own work.

In contrast the mean returns for USGOV10TR are constant in both periods. The standard deviation of both asset classes is higher in the training sample but the relationship between the volatility of equities and bonds is similar in the training and testing sample. These results demonstrate that the characteristics of equity and bonds are not the same in the training and testing period which may negatively impact the quality of volatility predictions. On the other hand the relationship between the volatility of equities and bonds stays persistent and the trade-off between the risk in equities and bonds is constant over time.

Figure 2 demonstrates compounded return of DTVEBA_10 strategy that targets 10% market volatility (top subfigure) and DTVEBA_10 strategy equity weight in time (bottom subfigure).[7] The compounded returns of DTVEBA_10 are compared with five benchmarks: 1) equity-only portfolio invested in

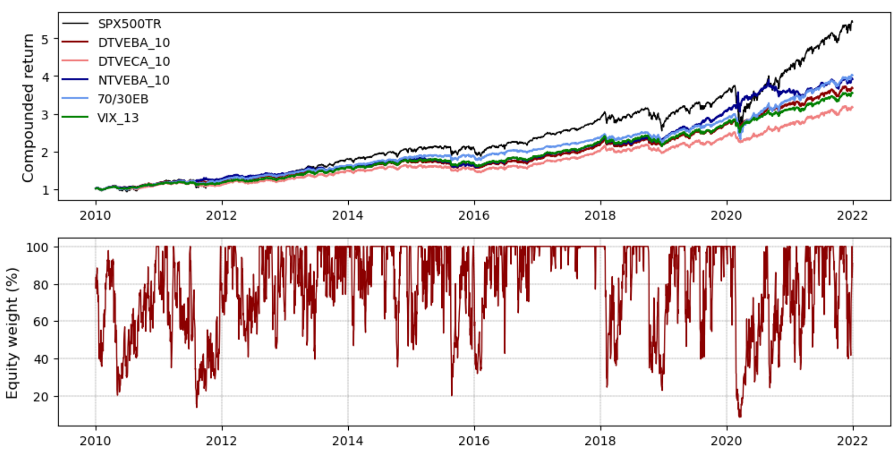

Figure 2. DTVEBA_10 strategy cumulative performance and equity allocation

Source: Own work.

[7] The 10% level is typical for other studies about the target volatility strategy, e.g., Y. Kim and Enke (2016, 2018), Kaczmarek et al. (2022).

S&P500 Total Return Index (SPX500TR); 2) target volatility strategy using cash instead of long-term bonds (DTVECA_10); 3) equity-bond target volatility strategy based on naïve prediction from previous day observed bi-powered variation (NTVEBA_10); 4) 70/30 equity-bond fixed allocation portfolio (70/30EB); and 5) equity-bond target volatility strategy based on VIX index that targets 13% implied volatility (VIX_13).

Figure 2 shows that the equity allocation in DTVEBA_10 changes dynamically and varies between 10-100%. Given that the sum of the weight of equity and govs in the strategy is 100% the bond allocation reaches values ranging from 0 to 90%. Furthermore, the line plots of compounded returns visualize that splitting the portfolio into stocks and bonds in DTVEBA_10 lowers the cumulative investment return relative to the S&P500TR Index. However, the volatility of DTVEBA_10 is lower which is easily seen during the COVID-19 sell-off.

Next, Table 2 provides detailed performance measures for all investment strategies. DTVEBA_10 outperforms all alternatives in terms of the Calmar ratio. The outperformance reaches 10.2% to 115.5% for DTVECA_10 and SPX500TR. The differences in the Sharpe ratio are less pronounced but the strategy still overperforms other alternatives.

Table 2. Equity/bonds mixed strategies

	SPX500TR	DTVEBA _10	DTVECA _10	NTVEB_10	70/30EB	VIX_13
Return	0.15	0.11	0.10	0.12	0.12	0.11
Std	0.17	0.09	0.09	0.10	0.11	0.09
Max drawdown	0.34	0.12	0.12	0.15	0.22	0.12
Max 1M loss	0.12	0.06	0.06	0.07	0.06	0.06
Sharpe	0.88	1.26*	1.06	1.21	1.06*	1.18*
Calmar	0.45	0.97	0.83	0.81	0.54	0.88
Av. equity share	1.00	0.78	0.78	0.79	0.70	0.78

The asterisk for Sharpe indicates that the difference with SPX500TR is statistically significant at the 5% level.

Source: Own work.

DTVEBA_10 and DTVECA_10 are two identical strategies, but use govs instead of cash allocation. The comparison of these investment alternatives demonstrates that the replacement of cash with long-term bonds is beneficial for overall target volatility strategy performance. Another two close cousins are DTVEBA_10 and NTVEBA_10. The evaluation of this pair which differs just in the volatility forecasting method demonstrates the benefits of RNN volatility predictions over the naïve approach.

70/30EB delivers a considerably lower Sharpe and Calmar ratio in relationship to DTVEBA_10. Although the fixed equity-bond strategy has on average eight p.p. higher average bond allocation than the target volatility alternative it still has a higher standard deviation and maximum drawdown. It means that DTVEBA_10 has lower risk and higher performance with higher (lower) equity (debt) allocation.

Finally, DTVEBA_10 is also compared with the target equity-bond allocation based on VIX. The average value of VIX is higher than the average realized volatility; therefore, this strategy needs to target 13% VIX to achieve the same average equity allocation as DTVEBA_10. In this setting, DTVEBA_10 overperforms VIX_13 by 8.5% (10.2%) in terms of the Sharpe (Calmar) ratio.

In detail: the low maximum drawdown of DTVEBA_10 shows that the strategy correctly detects market turmoil periods and benefits from the govs FTS phenomenon. These relationships are also confirmed when DTVEBA_10 to SPX500TR are compared, where the relative overperformance reaches 43% (115%) in terms of the Sharpe (Calmar) ratio.

The second test focuses on equity/debt allocation. It compares twelve DTVEBA strategies with target volatility ranging from 4% to 15% (TV4 to TV15). The average share of equity/debt is calculated for each of them and a peer fixed allocation strategy with the same proportions of equity and debt (34/66 to 92/8) is created. Table 3 presents the results.

Table 3. Target equity volatility with long term bonds vs. constant equity/long term bonds strategy

Panel A: Equity target volatility 4–9												
Strategy name	**TV4**	**34/66**	**TV5**	**43/57**	**TV6**	**51/49**	**TV7**	**60/40**	**TV8**	**67/33**	**TV9**	**73/27**
Return	0.07	0.08	0.08	0.09	0.09	0.10	0.09	0.11	0.10	0.12	0.11	0.12
Std	0.05	0.06	0.05	0.07	0.06	0.08	0.07	0.09	0.07	0.11	0.08	0.12
Max drawdown	0.07	0.09	0.07	0.11	0.08	0.15	0.09	0.18	0.11	0.21	0.11	0.24
Sharpe	1.43	1.44	1.48	1.36	1.46	1.25	1.40	1.16	1.33	1.09	1.29	1.04
Calmar	1.09	0.92	1.15	0.81	1.06	0.68	0.99	0.60	0.93	0.55	0.93	0.52
Equity share	0.34	0.34	0.43	0.43	0.51	0.51	0.60	0.60	0.67	0.67	0.73	0.73
Panel B: Equity target volatility 10–15												
Strategy name	**TV10**	**78/22**	**TV11**	**82/18**	**TV12**	**85/15**	**TV13**	**88/12**	**TV14**	**90/10**	**TV15**	**92/8**
Return	0.11	0.13	0.12	0.13	0.12	0.14	0.13	0.14	0.13	0.14	0.13	0.14
Std	0.09	0.13	0.10	0.14	0.10	0.14	0.11	0.15	0.11	0.15	0.12	0.16
Max drawdown	0.12	0.26	0.11	0.27	0.11	0.28	0.12	0.29	0.13	0.30	0.14	0.31
Sharpe	1.26	1.00	1.23	0.97	1.21	0.95	1.18	0.94	1.15	0.93	1.14	0.92
Calmar	0.97	0.50	1.08	0.49	1.12	0.48	1.06	0.47	1.00	0.47	0.95	0.46
Equity share	0.78	0.78	0.82	0.82	0.85	0.85	0.88	0.88	0.90	0.90	0.92	0.92

Source: Own work.

DTVEBA delivers higher Sharpe and Calmar ratios in all the cases analysed except for the most conservative of the examined portfolios. The largest over-performance is observed in pair TV12 and 82/18 strategy. The average Sharpe (Calmar) ratio for the DTVEBA strategies is 1.30 (1.03) and for fixed allocation strategies is 1.09 (0.58). With equal debt allocations the DTVEBA strate-

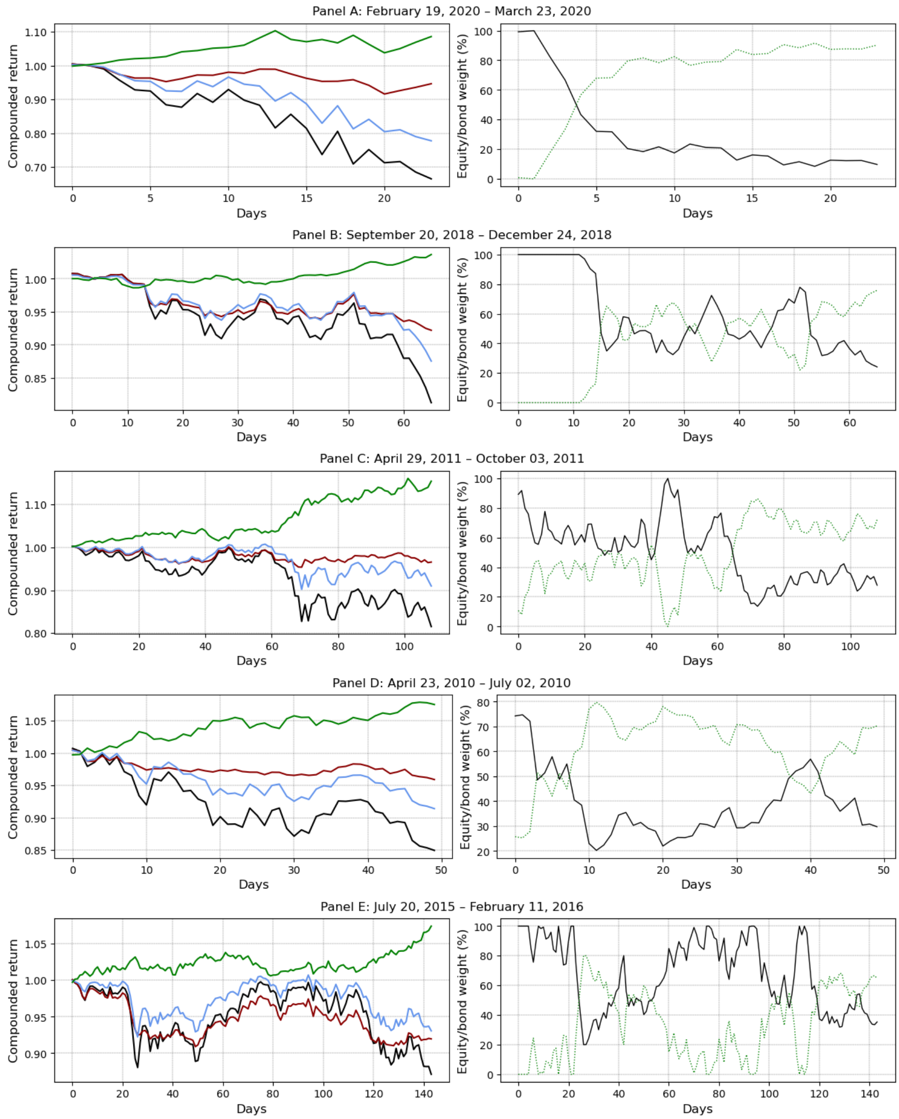

Figure 3. DTVEBA strategy results during the top five drawdowns of the S&P500 Total Return Index

Source: Own work.

gy overperforms the fixed allocation by 19.2% (Sharpe) and 77.5% (Calmar). From the other perspective if equal Sharpe and Calmar ratios with DTVEBA and fixed equity-debt allocation are sought significantly lower debt allocations are obtained. For example, 51/49EB and TV10 have the same Sharpe ratio but the debt allocation is more than twice lower. For the Calmar ratios all the DTVEBA strategies overperform 51/49EB. Based on all the Calmar observations it can can be concluded that obtaining the same Calmar ratios is possible when reducing the allocation in debt even more than three times.

Finally, DTVEBA's strategy performance is verified during the market stress conditions. Figure 3 presents the DTVEBA_10 strategy performance during the five most severe drawdowns in the testing sample (Table A.3 in the appendix shows the list of drawdowns). The figure is divided into five panels. Panels demonstrate different drawdown periods and consist of two subfigures. The left subfigure presents cumulative returns for the S&P500 Total Return Index (black line); fixed equity-bond 70/30EB allocation strategy (blue line); DTVEBA (red line); and USGOV10TR. The right subfigure visualizes the weight of equity (black line) and bonds (green dotted line) in the DTVEBA.

Figure 3 visualizes that in five of the most severe drawdown periods the strategy overperformed the S&P500 Index. In each case DTVEBA dynamically reduces equity and allocates most of the portfolio to govs that benefit from the FTS phenomenon.

Conclusions

This paper applies recurrent neural network predictions of realized market volatility to investment portfolio construction. It extends studies about target-volatility strategies by dynamic exposition to FTS events to enhance performance. Since equity markets perform better during low volatility high equity allocation in periods of low predicted volatility enhances portfolio performance. In contrast volatility increases during market turmoil; therefore high volatility forecasts reduce equity and increase govs' allocation thus providing a solid exposition for FTS events.

The conditional allocation to govs provides high flexibility to portfolio construction. With Deep Target Volatility Equity-Bond Allocation (DTVEBA) a constant exposition to govs is no longer obligatory to protect the portfolio against market turmoil by building exposition to govs and benefiting from FTS events. Instead the average govs allocation can be defined mainly through investor expectations about the relationship between expected returns from equity and govs. This flexibility is valuable during periods of low expected returns on government bonds when holding them in a portfolio may be solely related to affirm exposition to FTS events and to deliver lower expected returns. The effectiveness of this approach is verified in the fully out-of-sample approach. The DTVEBA strategy performance is measured by the Sharpe and Calmar ratios and compared to the S&P500 index as well as alternative portfolio diversification approaches.

To this end a considerably large-scale empirical analysis is conducted with twenty years of daily history spanning from 2002 to 2021. The strategy is tested with a 12-year window (2010–2021) that covers a few significant market selloff periods. The DTVEBA strategy targeting 10% volatility overperforms the S&P500 index by 43% and 115% in terms of the Sharpe and Calmar ratios and is exceptionally beneficial for portfolio drawdown risk reduction. The results show that a portfolio manager is better off by targeting volatility with the daily-adjusted decision than sticking to constant allocation.

The findings from this study are relevant to any investor that actively manages portfolio risk. It shows that market volatility estimations with a state-of-the-art recurrent neural networks model are essential to enhance portfolio performance and most importantly to manage its drawdown risk. Moreover, it demonstrates the innovatory approach to defining government bonds allocation where high govs allocation is required only before the FTS event. The results are also crucial for active investment advisors that with dynamic equity/govs allocation want to deliver an investment risk appropriate to investors' needs.

This study also has its limitations. First, the DTVEBA strategy is tested on a twelve year window—a maximum period where the input data is available and the model can be trained. Since deep neural networks master the analysis of complicated input data and discovery of sophisticated interactions the best performance is achieved with multivariate models consisting of many informative explanatory variables. Unfortunately, it is necessary to have data for all variables in order to start training the model so the variate with the shortest history defines the research sample. Furthermore, training the neural network requires many training and validation samples that cannot be used for out-of-sample testing. The research sample for this study starts in 2002, the training lasts eight years and the out-of-sample testing period starts in 2010. Second, the strategy is tested with the volatility of the S&P500 index. Although market volatility can be used as a good approximation for portfolio volatility the estimation of volatilities for each instrument in a portfolio may deliver more precise guidance concerning portfolio allocation. The last limitation of this study also represents an exciting direction for future research—namely research on controlling portfolio volatility.

Appendix A. Additional tables for the study

Appendix A consists of three tables that demonstrate the detailed description of variables, present descriptive statistics, and show the performance of DTVEBA_10 during the most severe five drawdowns for the S&P500 TR. It depicts a detailed description of the tables, the interpretation of which can be found in the main text.

Table A1 explains the variables used in the research. Panels organize the rows. Panel A presents the volatility measure for SPX500PI index. Panel B presents the variates used to predict volatility. Finally, panel C shows the total return indices used to calculate the target volatility strategies' performance. The table consists of four columns. The first column provides the name of each variable. The second column shows the symbol assigned to the variate. The third column describes how the variable is computed and measured. Finally, the fourth column indicates the source of the data used in this study for the given variable.

Next, Table A2 presents summary statistics for explanatory variables (Panel A) and strategy components in the training (Panel B) and testing period (Panel C). All variates are described in Table A1 The table columns present the mean (Mean); standard deviation (Std); skewness (Skew); excess kurtosis (Kurt); the number of daily observations, for which the given variable is available (Count); t-statistic of the Jarque-Bera test, for variable's normal distribution (Jarque-Bera), and the Augmented Dickey-Fuller test, for the presence of a unit root in a sample (ADF). *** and ** denote a rejection of the null hypothesis at the 1% and 5% significance level, respectively. The sample period runs from April 11, 2002, to December 31, 2021, and the testing period for the strategy starts in January 1, 2010.

Finally, Table A3 shows the most severe drawdowns in the testing sample. Results cover the period from January 1, 2010, to December 31, 2021. The table has six columns: 1) the first five worst drawdowns (Worst drawdown period); 2) the size of drawdown for S&P500 Total Return Index in % (SPX500TR Net drawdown); 3) the size of drawdown for Deep Target Volatility Equity-Bond Allocation (DTVEBA) equity target volatility strategy in % (target volatility = 10%) (DTVEBA_10 Net drawdown); 4) the start date of the drawdown (Peak date); 5) the end date of the drawdown (Valley date); and 6) the drawdown duration in days (Duration in days).

Table A.1. A detailed description of variables

Variate	Symbol	Description	Data source
Panel A: Volatility measures			
S&P 500 bi-powered variation	*BPVSPX500*	Bi-powered variation is calculated as $RBV_t(\Delta) = \frac{\pi}{2}\sum_{n=1}^{\frac{1}{\Delta}} \lvert r_{t,n}\rvert \lvert r_{t,n+1}\rvert$, where $r_{t,n}$ is n-th intraday return on day t and Δ denotes the frequency of intra-daily returns. We transform the bi-powered daily variation taken from the source into monthly standard deviation as $BPV = 22 \cdot \sqrt{RBV_t(\Delta)}$.	Oxford-Man*
Panel B: Explanatory variables			
S&P 500 COMPOSITE – PRICE INDEX	*SPX500PI*	One day log-return of (Standard and Poor's 500 Composite, price index), Datastream symbol: S&PCOMP	Datastream
Gold Bullion LBM $/t oz DELAY	*GOLDLOG*	One day log-return of (Gold Bullion London Bullion Market United States Dollar Per Metric Tonne Ounce Delay), Datastream symbol: GOLDBLN	Datastream
Crude Oil-WTI Spot Cushing U$/BBL	*CRUDELOG*	One day log-return of (Crude Oil-West Texas Intermediate Spot Cushing United States Dollar Per Barrel), Datastream symbol: CRUDOIL	Datastream
US GOVERNMENT BOND SERIES 3 YEAR – RED. YIELD	*USGOV3YI*	One day difference in yield of (United States Government Bond Series 3 Years), Datastream symbol: GBUS03Y	Datastream
RF US CORP BMK AA 3Y – RED. YIELD	*USCORPAA3YI*	One day difference in yield of (Refinitiv United States Corp Benchmark A.A. 3 Years), Datastream symbol: TRUCBYC	Datastream
Volatility risk premium (Prokopczuk & Wese Simen, 2014)	*VOL_RISK_PREM*	Quotient of CBOE SPX VOLATILITY VIX (NEW) – PRICE INDEX and *BPVSPX500*	Datastream
Adjusted conditional volatility from GARCH (1,1)	*COND_VOL_GARCH*	$\sigma_{t-1}^{2'} = \beta_1\sigma_{t-1}^2$ from the GARCH (1,1) model (Bollerslev, 1986) specification: $\sigma_t^2 = \alpha_0 + \alpha_1\varepsilon_{t-1}^2 + \beta_1\sigma_{t-1}^2$, where σ_t^2 represents square of volatility of returns and ε_t states for unpredictable error term. β_1 is estimated with the latest 500 observations in the training sample (Hyup Roh, 2007).	Own estimations
Adjusted residuals from GARCH(1,1)	*RESID_GARCH*	$\varepsilon_{t-1}^{2'} = \alpha_1\varepsilon_{t-1}^2$ from the GARCH (1,1) model as above. α_1 is estimated with the latest 500 observations in the training sample (Hyup Roh, 2007).	Own estimations

Adjusted conditional volatility from EGARCH	*COND_VOL_EGARCH*	$ln\sigma_{t-1}^{2'} = \beta ln\sigma_{t-1}^{2}$ from the EGARCH model (Nelson, 1991) specification: $ln\sigma_t^2 = \alpha + \beta ln\sigma_{t-1}^2 + \omega\left(\frac{\varepsilon_{t-1}}{\sigma_{t-1}}\right) + \gamma\left\vert\frac{\varepsilon_{t-1}}{\sigma_{t-1}} - \sqrt{\frac{2}{\pi}}\right\vert$ β is estimated with the latest 500 observations in the training sample (Hyup Roh, 2007)	Own estimations
Leverage effect EGARCH	*LEV_EFFECT_EGARCH*	$\gamma\left\vert\frac{\varepsilon_{t-1}}{\sigma_{t-1}} - \sqrt{\frac{2}{\pi}}\right\vert$ from the EGARCH model as above. γ is estimated with the latest 500 observations in the training sample (Hyup Roh, 2007).	Own estimations
Leverage EGARCH	*LEV_EGARCH*	$\omega\frac{\varepsilon_{t-1}}{\sigma_{t-1}}$ from the EGARCH model as above. ω is estimated with the latest 500 observations in the training sample (Hyup Roh, 2007).	Own estimations
Adjusted conditional volatility from EWMA	*COND_VOL_EWMA*	$\sigma_t^{2'} = \lambda\sigma_{t-1}^2$ from the EWMA model specification (RiskMetrics, by J.P. Morgan & Co.): $\sigma_t^2 = \lambda\sigma_{t-1}^2 + (1-\lambda)\varepsilon_{t-1}^2$ λ is estimated with the latest 500 observations in the training sample (Hyup Roh, 2007).	Own estimations
Adjusted residuals from EWMA	*RESID_EWMA*	$\varepsilon_{t-1}^{2'} = (1-\lambda)\varepsilon_{t-1}^2$ from the EWMA model as above. λ is estimated with the latest 500 observations in the training sample (Hyup Roh, 2007).	Own estimations
Panel C: Total return indices			
S&P 500 COMPOSITE - TOT RETURN IND	*SPX500TR*	One day total return of Standard and Poor's 500 Composite, Datastream symbol: S&PCOMP, total return index	Datastream
US BENCHMARK 10 YEAR DS GOVT. INDEX - TOT RETURN IND	*USGOV10TR*	One day total return of United States Benchmark 10 Year Datastream Government Index, Datastream symbol: BMUS10Y	Datastream

* Oxford-Man states for Oxford-Man Institute of Quantitative Finance, Realized Library (Heber et al., 2009).

Source: Own work.

Table A2. Descriptive statistics

	Mean	Std	Skew	Kurt	Count	Jarque-Bera	ADF
Panel A: Explanatory variables							
BPVSPX500	0.1614	0.1246	3.75	22.91	4951	119 618.2***	–7.5***
SPX500PI	0.0003	0.0123	–0.46	12.52	4951	32 431.9***	–17.0***
GOLDLOG	0.0004	0.0111	–0.49	5.33	4951	6 041.3***	–70.7***
CRUDELOG	0.0004	0.0269	0.52	18.76	4951	72 646.1***	–12.5***
USGOV3YI	–0.0003	0.0445	0.15	7.96	4951	13 053.3***	–23.7***
USCORPAA3YI	–0.0003	0.0327	0.10	424.72	4951	37 136 615.2***	–48.5***
VOL_RISK_PREM	1.4094	0.4782	1.47	4.75	4951	6 436.7***	–7.5***
COND_VOL_GARCH	1.2858	2.8780	8.65	105.71	4951	2 362 093.4***	–7.4***
RESID_GARCH	0.3654	1.3886	13.44	257.70	4951	13 821 016.6***	–7.7***
COND_VOL_EGARCH	–0.0699	0.9255	0.69	0.54	4951	454.8***	–6.2***
LEV_EFFECT_EGARCH	–0.1579	0.1192	–1.47	4.08	4951	5 199.4***	–16.2***
LEV_EGARCH	–0.0008	0.0062	–0.63	2.18	4951	1 302.8***	–31.3***
COND_VOL_EWMA	1.2766	2.9741	7.59	76.48	4951	1 251 508.9***	–7.3***
RESID_EWMA	0.2502	0.9511	13.44	257.67	4951	13 817 439.5***	–7.7***
Panel B: Strategy components (training period)							
USGOV10TR	0.0002	0.0052	0.11	3.45	1941	–	–
SPX500TR	0.0002	0.0142	0.11	9.33	1941	–	–
Panel C: Strategy components (testing period)							
USGOV10TR	0.0002	0.0043	–0.11	2.25	3010	–	–
SPX500TR	0.0006	0.0108	0.0006	0.0108	3010	–	–

Source: Own work.

Table A3. Target equity volatility strategy with long-term bonds during the most severe five drawdowns for the S&P500 TR

Worst drawdown period	SPX500TR Net drawdown	DTVEBA_10 Net drawdown	Peak date	Valley date	Duration in days
1	33.8	5.3	2020-02-19	2020-03-23	33
2	19.4	7.8	2018-09-20	2018-12-24	95
3	18.6	3.4	2011-04-29	2011-10-03	157
4	15.6	4.1	2010-04-23	2010-07-02	70
5	13.0	8.0	2015-07-20	2016-02-11	206

Source: Own work.

Appendix B. Volatility prediction model

B1. Development of RNN*s*

Recurrent Neural Network (RNN) is a class of neural network that is designed to perform tasks related to processing sequences. In contrast to a traditional feedforward network, a basic RNN has a distinguishing feature—backward connections. The most elementary form of an RNN with three nodes—input node x_t, output node y_t, and hidden node h_t—is depicted in Figure B1 (left). On the right side, a visualization of the unrolled network is demonstrated when the recurrent network is presented once per step. The recurrent neuron receives the input x_t and its output from the previous step y_{t-1} and takes two sets of weights at each time t: one for the input x_t and the other for the output from the previous time step, y_{t-1}. By defining the output at step t as y_t and the hidden state output as h_t, their relationship can be expressed mathematically (Géron, 2019):

$$h_{(t)} = \sigma\left(W_{hx}^T x_{(t)} + W_{hh}^T h_{(t-1)} + b_h\right) \tag{B1}$$

$$y_{(t)} = f_0\left(W_{yh}^T h_{(t)} + b_y\right) \tag{B2}$$

where two matrices of weights comprise the hidden state weights, W_{hx} and W_{hh}, along with the recurrent neurons; a matrix of weights called W_{yh}, encapsulates the output layer weights; bias vectors in the hidden layer, b_h; and the output layer, b_y; and lastly, an activation function denoted by σ. Ultimately, the function transfer from the hidden state to the output values is represented by f_0.

The basic configuration of a Recurrent Neural Network is limited in its ability to learn sequences of significant length. Consequently, it may be insufficient to train for complex tasks that involve long-term dependencies, as indicated by Bengio et al. (1994). To address this problem, multiple long-term memory cells have been developed. Gated Recurrent Units (GRU) were proposed by Cho et al. (2014) as a simplified version of LSTM that reduces computational costs without compromising performance, as confirmed by Greff et al. (2015) and Cong et al. (2020). The GRU architecture unites two state vectors into a singular vector $-h_{(t)}$. Additionally, it features a single control gate that manages both the input and forget gates. Unlike the LSTM, there is no output gate, and the state vector is produced entirely during each time step. As an alternative, a new gate controller is introduced $-r_{(t)}$. The formal computation details for GRU, which describe equations B3–B6, are as follows:

$$z_{(t)} = \sigma\left(W_{xz}^T x_{(t)} + W_{hz}^T h_{(t-1)} + b_z\right) \tag{B3}$$

$$r_{(t)} = \sigma\left(W_{xr}^T x_{(t)} + W_{hr}^T h_{(t-1)} + b_r\right) \tag{B4}$$

$$g_{(t)} = tanh\left(W_{xg}^T x_{(t)} + W_{hg}^T \left(r_{(t)} \otimes h_{(t-1)}\right) + b_g\right) \tag{B5}$$

$$h_{(t)} = z_{(t)} \otimes h_{(t-1)} + \left(1 - z_{(t)}\right) \otimes g_{(t)} \tag{B6}$$

B.2. Model training, validation, and testing

The neural network architecture in this study involves a stacked GRU composed of an initial layer for input, followed by two stacked hidden GRU layers, which subsequently feed their output to a final layer responsible for delivering a single volatility prediction value.

In order to achieve out-of-sample results, the data is split into two subsamples: training and testing. The former is employed to estimate values for the model's hyperparameters. As part of this process, the five-fold cross-validation technique is used. This technique splits the training dataset into five subsets and continuously searches for hyperparameters to minimize predictive mean squared error. The explored hyperparameter range includes: the number of neurons in the hidden layers; the batch size; the level of l2 regularisation on weights and bias; and a dropout ratio.

Machine learning algorithms are prone to overfitting. The neural network architecture and the training process involve four regularization techniques. Their goal is to reduce the probability of overfitting. The first of these approaches consists of the implementation of l2, also referred to as ridge regularization, on both the weights and bias. This widely-accepted and commonly utilized technique in the realm of machine learning helps to control overfitting by inflicting a penalty on the objective function (Gu et al., 2020). The second method employed is what is known as "dropout." This technique involves the random exclusion of certain neurons throughout the training phase (Cong et al., 2020). The parameters for both the l2 regularization and dropout methods are established via cross-validation. Thirdly, the "early-stopping" technique is implemented that stops training when the mean square error is no longer improving along with subsequent batches in the training process. Finally, as our fourth and final measure for regularization, an ensemble approach is implemented. With this technique predictions from separate training processes are averaged to get more reliable outcomes.

The training process of Recurrent Neural Networks is based on an effective adaptive method for stochastic gradient descent called "Adam". This algorithm, developed by Kingma and Ba (2015), evaluates the first and second moments of the gradients and generates adaptive learning rates for each parameter. All of the computations in the study are performed using the Python programming language, with the assistance of the Keras and TensorFlow libraries.

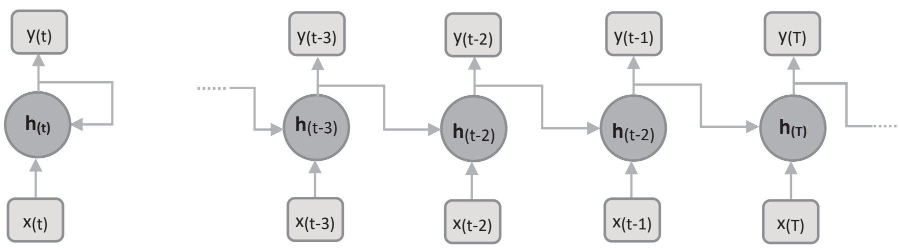

Figure B1. Main Recurrent Neural Networks structure (left side) and the unrolled representation (right side)

Source: Based on (Kaczmarek et al., 2022).

References

Adrian, T., Crump, R. K., & Vogt, E. (2019). Nonlinearity and flight-to-safety in the risk-return trade-off for stocks and bonds. *The Journal of Finance*, *74*(4), 1931–1973. https://doi.org/10.1111/jofi.12776

Andersen, T. G., & Bollerslev, T. (1998). Answering the skeptics: Yes, standard volatility models do provide accurate forecasts. *International Economic Review*, *39*(4), 885–905. https://doi.org/10.2307/2527343

Baele, L., Bekaert, G., Inghelbrecht, K., & Wei, M. (2020). Flights to safety. *The Review of Financial Studies*, *33*(2), 689–746. https://doi.org/10.1093/rfs/hhz055

Barndorff-Nielsen, O. E. (2004). Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics*, *2*(1), 1–37. https://doi.org/10.1093/jjfinec/nbh001

Baur, D. G., & Lucey, B. M. (2009). Flights and contagion—An empirical analysis of stock-bond correlations. *Journal of Financial Stability*, *5*(4), 339–352. https://doi.org/10.1016/j.jfs.2008.08.001

Beber, A., Brandt, M. W., & Kavajecz, K. A. (2009). Flight-to-quality or flight-to-liquidity? Evidence from the Euro-area bond market. *Review of Financial Studies*, *22*(3), 925–957. https://doi.org/10.1093/rfs/hhm088

Becker, R., Clements, A. E., Doolan, M. B., & Hurn, A. S. (2015). Selecting volatility forecasting models for portfolio allocation purposes. *International Journal of Forecasting*, *31*(3), 849–861. https://doi.org/10.1016/j.ijforecast.2013.11.007

Będowska-Sójka, B. (2018). Is intraday data useful for forecasting VaR? The evidence from EUR/PLN exchange rate. *Risk Management*, *20*(4), 326–346. https://doi.org/10.1057/s41283-018-0038-z

Bengio, Y., Simard, P., & Frasconi, P. (1994). Learning long-term dependencies with gradient descent is difficult. *IEEE Transactions on Neural Networks*, *5*(2), 157–166. https://doi.org/10.1109/72.279181

Bollerslev, T. (1986). Generalized autoregressive conditional heteroskedasticity. *Journal of Econometrics*, *31*(3), 307–327. https://doi.org/10.1016/0304-4076(86)90063-1

Bongaerts, D., Kang, X., & van Dijk, M. (2020). Conditional volatility targeting. *Financial Analysts Journal*, *76*(4), 54–71. https://doi.org/10.1080/0015198X.2020.1790853

Boucher, C., & Tokpavi, S. (2019). Stocks and bonds: Flight-to-safety for ever? *Journal of International Money and Finance*, *95*, 27–43. https://doi.org/10.1016/j.jimonfin.2019.03.002

Brooks, C. (1998). Predicting stock index volatility: Can market volume help? *Journal of Forecasting*, *17*(1), 59–80. https://doi.org/10.1002/(SICI)1099-131X(199801)17:1<59::AID-FOR676>3.0.CO;2-H

Cho, K., Van Merriënboer, B., Gulcehre, C., Bahdanau, D., Bougares, F., Schwenk, H., & Bengio, Y. (2014). Learning phrase representations using RNN encoder-decoder for statistical machine translation. *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, 1724–1734. Doha. Qatar. https://doi.org/10.3115/v1/d14-1179

Christensen, K., Siggaard, M., & Veliyev, B. (2021). *A machine learning approach to volatility forecasting*. Creates Research Papers. https://ideas.repec.org/p/aah/create/2021-03.html

Cong, L., Tang, K., Wang, J., & Zhang, Y. (2020, June 1). Deep sequence modeling: Development and applications in asset pricing. *The Journal of Financial Data Science*. https://doi.org/10.2139/ssrn.3646170

Fleming, J., Kirby, C., & Ostdiek, B. (2003). The economic value of volatility timing using "realized" volatility. *Journal of Financial Economics*, *67*(3), 473–509. https://doi.org/10.1016/S0304-405X(02)00259-3

Géron, A. (2019). *Hands-on machine learning with scikit-learn, keras, and tensorflow: Concepts, tools, and techniques to build intelligent systems*. O'Reilly Media, Incorporated. https://books.google.pl/books?id=OCS1twEACAAJ

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. https://doi.org/10.1111/j.1540-6261.1993.tb05128.x

Grabowski, W., Janus, J., & Stawasz-Grabowska, E. (2023). The influence of monetary policy response to the COVID-19 crisis on dynamic connectedness across financial markets in Central and Eastern Europe. *Entrepreneurial Business and Economics Review*, *11*(1), 7–28. https://doi.org/10.15678/EBER.2023.110101

Greff, K., Srivastava, R. K., Koutník, J., Steunebrink, B. R., & Schmidhuber, J. (2015). LSTM: A search space odyssey. *IEEE Transactions on Neural Networks and Learning Systems*, *28*(10), 2222–2232. https://doi.org/10.1109/TNNLS.2016.2582924

Gu, S., Kelly, B., & Xiu, D. (2020). Autoencoder asset pricing models. *Journal of Econometrics*, *222*(1), 429–450. https://doi.org/10.1016/j.jeconom.2020.07.009

Hamid, S. A., & Iqbal, Z. (2004). Using neural networks for forecasting volatility of S&P 500 Index futures prices. *Journal of Business Research*, *57*(10), 1116–1125. https://doi.org/10.1016/S0148-2963(03)00043-2

Hansen, P. R., & Lunde, A. (2005). A forecast comparison of volatility models: Does anything beat a GARCH(1,1)? *Journal of Applied Econometrics*, *20*(7), 873–889. https://doi.org/10.1002/jae.800

Heber, G., Lunde, A., Shephard, N., & Sheppard, K. K. (2009). *Oxford-Man Institute's realized library*. Oxford-Man Institute, University of Oxford, Library Version: 0.3.

Hocquard, A., Ng, S., & Papageorgiou, N. (2013). A constant-volatility framework for managing tail risk. *Journal of Portfolio Management*, *39*(2), 28–40. https://doi.org/10.3905/jpm.2013.39.2.028

Hyup Roh, T. (2007). Forecasting the volatility of stock price index. *Expert Systems with Applications*, *33*(4), 916–922. https://doi.org/10.1016/j.eswa.2006.08.001

Kaczmarek, T., Będowska-Sójka, B., Grobelny, P., & Perez, K. (2022). False safe haven assets: Evidence from the target volatility strategy based on recurrent neural network. *Research in International Business and Finance*, *60*, 101610. https://doi.org/10.1016/j.ribaf.2021.101610

Kim, H. Y., & Won, C. H. (2018). Forecasting the volatility of stock price index: A hybrid model integrating LSTM with multiple GARCH-type models. *Expert Systems with Applications*, *103*, 25–37. https://doi.org/10.1016/j.eswa.2018.03.002

Kim, Y., & Enke, D. (2016). Using neural networks to forecast volatility for an asset allocation strategy based on the target volatility. *Procedia Computer Science*, *95*, 281–286. https://doi.org/10.1016/J.PROCS.2016.09.335

Kim, Y., & Enke, D. (2018). A dynamic target volatility strategy for asset allocation using artificial neural networks. *The Engineering Economist*, *63*(4), 273–290. https://doi.org/10.1080/0013791X.2018.1461287

Kingma, D. P., & Ba, J. L. (2015, December 22). *Adam: A method for stochastic optimization*. 3rd International Conference on Learning Representations (ICLR). Conference Track Proceedings.

Lehnert, T. (2022). International review of financial analysis flight-to-safety and retail investor behavior. *International Review of Financial Analysis*, *81*, 102142. https://doi.org/10.1016/j.irfa.2022.102142

Longstaff, F. A. (2004). The flight-to-liquidity premium in U.S. treasury bond prices. *Journal of Business*, *77*(3), 511–526. https://doi.org/10.1086/386528

Moreira, A., & Muir, T. (2017). Volatility-managed portfolios. *The Journal of Finance*, *72*(4), 1611–1644. https://doi.org/10.1111/jofi.12513

Nelson, D. B. (1991). Conditional heteroskedasticity in asset returns: A new approach. *Econometrica*, *59*(2), 347–370. https://doi.org/10.2307/2938260

Perchet, R., de Carvalho, R. L., Heckel, T., & Moulin, P. (2016). Predicting the success of volatility targeting strategies: Application to equities and other asset classes. *The Journal of Alternative Investments*, *18*(3), 21–38. https://doi.org/10.3905/jai.2016.18.3.021

Poon, S. H., & Granger, C. (2005). Practical issues in forecasting volatility. *Financial Analysts Journal*, *61*(1), 45–56. https://doi.org/10.2469/faj.v61.n1.2683

Prokopczuk, M., & Wese Simen, C. (2014). The importance of the volatility risk premium for volatility forecasting. *Journal of Banking & Finance*, *40*, 303–320. https://doi.org/10.1016/j.jbankfin.2013.12.002

Zakamulin, V. (2019). Volatility weighting over time in the presence of transaction costs. *Journal of Wealth Management*, *21*(4), 33–45. https://doi.org/10.3905/jwm.2019.21.4.033