RiskOptimix

Mathematical background

Thomas van der Hulst
riskoptimix@gmail.com

August 2025

Contents

1 Introduction

This file contains the mathematical background of the RiskOptimix model. Specifically, this paper will discuss the use of different types of GARCH models in ensemble form, and two different multivariate GARCH models: Constant Conditional Correlation (CCC) and Dynamic Conditional Correlation (DCC).

All credits go to the curriculum of the course 'Financial Econometrics' at the Vrije Universiteit Amsterdam, which gave the foundation of my knowledge of financial econometrics and these types of models. Furthermore, it must be noted that certain sections of this paper are taken from my bachelor's thesis "*The effect of extreme volatility on the accuracy of GARCH models in VaR estimation*" which is available upon request. This thesis was based upon a publication by Angelidis et al. (2004). Lastly, it must be noted that this paper is the result of many evenings and weekends spent exploring GARCH models out of academic curiosity. While I've tried to be rigorous and follow the literature carefully, this remains an independent project that may contain errors or debatable methodological choices. Please read with appropriate skepticism, and feel free to reach out if you spot anything that seems off.

Throughout this paper, additional details and examples will be given using a basic portfolio given in Table 1, while it must be noted that this can be extended to a portfolio of any size and value. The data are obtained from Yahoo! Finance (2025).

Stock Ticker	**Portfolio Value ($)**
AAPL	5000.00
MSFT	3000.00

Table 1: This table represents a simple stock portfolio that will be used as an example throughout this paper. This portfolio contains $\$5000.00$ of Apple Inc. (AAPL) and $\$3000.00$ of Microsoft Corp. (MSFT) and $\$4000.00$.

2 Prices and returns

The movements of the stock prices of a particular stock are often modeled as a random walk process. Formally, the time series of prices $\{p_t\}_{t\in\mathbb{Z}}$ is given by $p_t = p_{t-1} + \varepsilon_t$, with ε_t following from a white noise sequence with $\mathbb{E}(\varepsilon_t|p_{t-1}) = 0$. Because of this random walk process, the price at time t would be the best forecast for the price at time $t+1$. Furthermore, since a random walk process implies that the time series is unit-root non-stationary, it follows that the time series itself is non-stationary. However, the variation in prices, commonly referred to as returns, is often considered stationary. Returns often exhibit volatility clustering, characterized by periods where large returns are followed by other large returns, and small returns are followed by small returns. This volatility clustering is useful for understanding the risk associated with a particular stock or portfolio. Predicting the volatility of future returns can be extremely useful for risk management and risk minimization.

To use the returns for volatility prediction, this paper uses the continuously compounded returns. This is denoted by $y_t = \ln(p_t/p_{t-1}) \times 100$, where p_t is the closing price of a stock at time t.

The prices and returns of the stocks in our portfolio over the period from January 1, 2010 to January 1, 2025 can be seen in Figures 1 and 2. These figures show interesting trends that motivate our chosen methodology. First, we see a rise in stock prices for both stocks over the period, with sharp drops around March 2020 and the beginning of 2022 due to the COVID-19 crisis. This drop goes in line with increased volatility in the stocks, which can be clearly observed in Figure 2. Here we see two important phenomena. The first is volatility clustering, where periods of low volatility tend to be followed by low volatility, and high volatility by high volatility. The second phenomenon is the leverage effect, where stock volatility tends to be higher following negative returns than positive returns, as observed by Black (1976). How we address these phenomena will be discussed in Section 3.

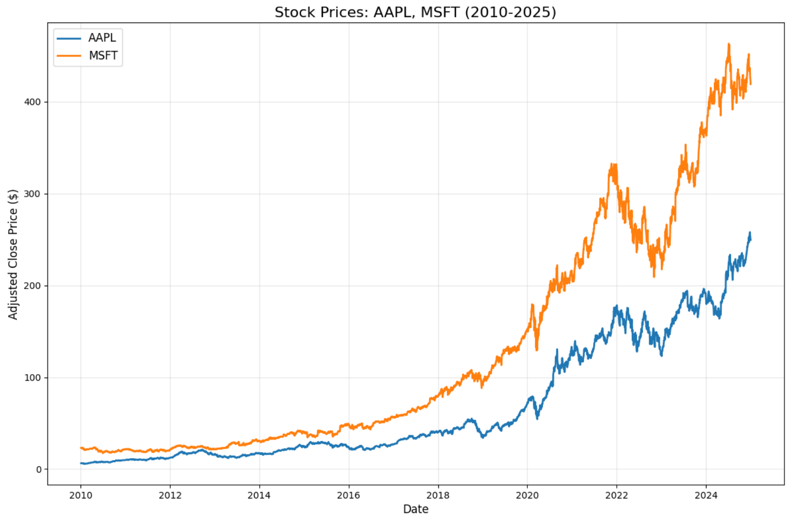

Figure 1: This figure contains the course of the stock prices of the stocks in the portfolio of Table 1 over the period from 2010-01-01 until 2025-01-01.

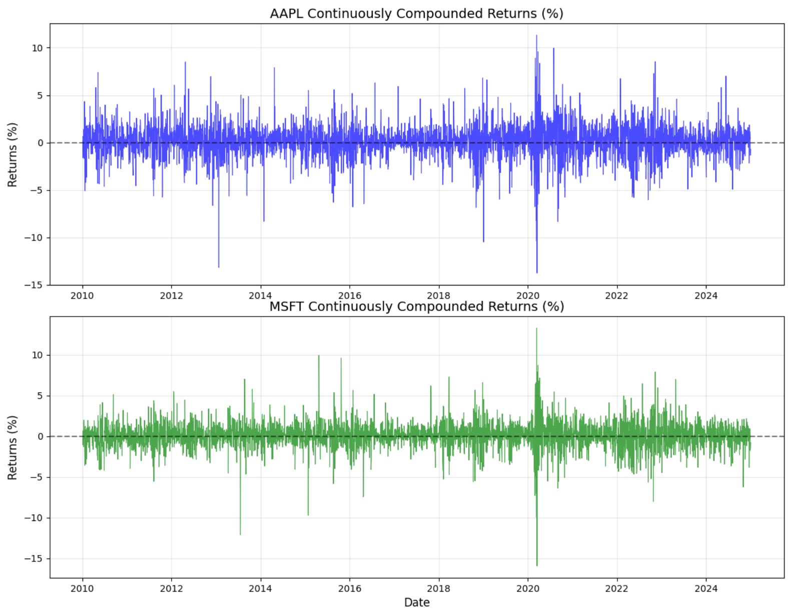

Figure 2: This figure contains the continuously compounded returns of the stocks in the portfolio of Table 1 over the period from 2010-01-01 until 2025-01-01.

3 GARCH models

3.1 ARCH models and extensions

As discussed in Section 2, one wishes to model the return process of the portfolio. For now, we focus on modelling the process of the returns of an individual stock. We model the returns process using an autoregressive structure as follows:

$$y_t = \mu + \sum_{i=1}^{k} \phi_i y_{t-i} + \varepsilon_t, \tag{1}$$

consisting of the predictable k-th order autoregressive (AR) process and the unpredictable component ε_t. In this AR process, μ represents the constant term that affects the unconditional mean $\mu/(1 - \sum_{i=1}^{k} \phi_i)$ under the assumption of a stationary process. The terms ϕ_i, $i = 1, \ldots k$ denote the coefficients of the k-order AR component of the model and capture the effect of past returns on the current return. The term $\varepsilon_t = z_t \sigma_t$ is seen as an unpredictable component and is modeled as an autoregressive conditional heteroscedastic (ARCH) process to capture volatility clustering and make predictions about the varying conditional variance (R. F. Engle, 1982). Here, z_t is assumed to be an independently and identically distributed random variable. Furthermore, z_t is assumed to have zero mean and unit variance $\{z_t\}_{t\in\mathbb{Z}} \sim \text{IID}(0, 1)$. The conditional variance of ε_t, denoted by σ_t^2, can be expressed in multiple forms, where R. F. Engle (1982) introduced an ARCH(q) model, given by

$$\sigma_t^2 = \omega + \sum_{i=1}^{q} \alpha_i \varepsilon_{t-i}^2,$$

with parameters $\omega > 0$ and $\alpha_1 \geq 0, \ldots, \alpha_q \geq 0$ determining the behavior of the conditional volatility. The restrictions set on these parameters ensure the positivity of σ_t^2. In this structural form, the volatility σ_t^2 depends on the lagged unpredictable component in (1).

A natural generalization of the ARCH process was proposed by Bollerslev (1986) to capture the effect of past conditional variances on the current conditional variances, denoted by the GARCH(p, q) model:

$$\sigma_t^2 = \omega + \sum_{i=1}^{q} \alpha_i \varepsilon_{t-i}^2 + \sum_{j=1}^{p} \beta_j \sigma_{t-j}^2. \tag{2}$$

Similar to the ARCH process, the restrictions $\omega > 0$, $\alpha_1 \geq 0, \ldots \alpha_q \geq 0$ and $\beta_0 \geq 0, \ldots \beta_p \geq 0$ are made to ensure the conditional variance is positive. Additionally to the ARCH model proposed by R. F. Engle (1982), the GARCH process also depends on the lagged volatility itself.

Furthermore, Bollerslev (1986) noted that for the GARCH(p, q) model, the problem of the ARCH model including a substantial number of lags is solved. This is because even the GARCH$(1, 1)$ process can be written by a ARCH(∞) process by unfolding (2):

$$\sigma_t^2 = \frac{\omega}{(1 - \beta_1)} + \alpha_1 \sum_{i=0}^{\infty} \beta_1^i \varepsilon_{t-1-i}^2.$$

The GARCH model is often used to explain the structure of the conditional variance. This is due to its ability to explain volatility clustering in the returns and its robustness against extreme results. At the same time, the GARCH model is seen as a symmetric model, as the conditional variance is impacted by the squared term ε_{t-i}^2. Therefore, positive and negative values have the same impact on σ_t^2. However, as was seen in Section 2, this relationship is often negative, where extreme volatility is more often the result of negative past returns than positive past returns. Because of this result, a model based on the GARCH(p,q) process may not correctly include this asymmetric property of the returns, and other models could help make better forecasts.

To account for the asymmetry in the time series, asymmetric ARCH models were introduced, where our models uses an exponential GARCH (EGARCH(p,q)) model and a threshold GARCH (TARCH(p,q)) model as extensions to the GARCH(p,q) model to capture the natural structure of the returns potentially better.

Along with other reasons, Nelson (1991) used the result of asymmetry of Black (1976) and proposed the EGARCH(p,q) model, which is given by

$$\ln(\sigma_t^2) = \omega + \sum_{i=1}^{q} \alpha_i \, g(z_{t-i}) + \sum_{j=1}^{p} \beta_j \ln(\sigma_{t-j}^2), \tag{3}$$

with $g(z_t) = \theta z_t + \gamma\left[|z_t| - \mathbb{E}|z_t|\right]$ (Nelson, 1991). Plugging in the function $g(z_t)$ in (3) and using the fact that $z_t = \frac{\varepsilon_t}{\sigma_t}$, the EGARCH$(p,q)$ model proposed by Nelson (1991) is rewritten to

$$\ln(\sigma_t^2) = \omega + \sum_{i=1}^{q} \left(\alpha_i \left(\left| \frac{\varepsilon_{t-i}}{\sigma_{t-i}} \right| - \mathbb{E}|z_{t-i}| \right) + \gamma_i \frac{\varepsilon_{t-i}}{\sigma_{t-i}} \right) + \sum_{j=1}^{p} \beta_j \ln(\sigma_{t-j}^2).$$

In this model, γ_i allows the conditional variance to respond differently to positive and negative returns, solving the GARCH model's problem of symmetry. Another problem that the model imposed by Nelson (1991) solves is the restrictions on the parameters in the GARCH model to have a positive conditional variance, where since the term $\ln(\sigma_t^2)$ is considered instead of σ_t^2, the conditional variance is ensured to be non-negative. The part of $\mathbb{E}|z_t|$ depends on the underlying chosen distribution of z_t and will be further discussed later.

Zakoian (1994) and Glosten et al. (1993) independently introduced the TARCH(p,q) model, given by

$$\sigma_t^2 = \omega + \sum_{i=1}^{q} \alpha_i \varepsilon_{t-i}^2 + \gamma \varepsilon_{t-1}^2 d_{t-1} + \sum_{j=1}^{p} \beta_j \sigma_{t-j}^2, \tag{4}$$

where $d_t = 1$ if $\varepsilon_t < 0$, and $d_t = 0$ otherwise. Similar but different to the EGARCH model, the TARCH model introduces the asymmetric effect by the additional structure $\gamma \varepsilon_{t-1}^2 d_{t-1}$, where the volatility is differently impacted if $d_t = 1$ or in other words if the returns are negative.

3.2 Distribution of the shocks

As was discussed in Subsection 3.1, the independent and identically distributed sequence of shocks z_t are assumed to follow some distribution. R. F. Engle (1982) assumed z_t follows a standard normal distribution and therefore has a density given by

$$f(z_t) = \frac{1}{\sqrt{2\pi}} \exp\left(-\frac{z_t^2}{2}\right).$$

Because returns could be possibly better explained by a distribution with fatter tails than a normal distribution, Bollerslev (1987) proposed that the sequence z_t follows a standardized t-distribution (Gosset, 1908). The degrees of freedom $\nu > 2$ here describe the thickness of the tails, where for increasing values of ν, the distribution tends to a normal distribution. This distribution has a density given by

$$f(z_t;\nu) = \frac{\Gamma((\nu+1)/2)}{\Gamma(\nu/2)\sqrt{\pi(\nu-2)}} \left(1 + \frac{z_t^2}{\nu-2}\right)^{-\frac{\nu+1}{2}},$$

where $\Gamma(\cdot)$ is the gamma function.

Nelson (1991) suggested the generalized error distribution (GED), which includes various distributions that feature heavier and thinner tails and was first proposed by Subbotin (1923). The tails of the GED are fatter than the tails of a normal distribution for $0 < \nu < 2$, where the GED obtains thinner tails for $\nu > 2$. The GED has a density given by

$$f(z_t; \nu) = \frac{\nu}{2^{1+\frac{1}{\nu}} \lambda \Gamma(\frac{1}{\nu})} \exp\left(-\frac{1}{2}\left|\frac{z_t}{\lambda}\right|\right),$$

where

$$\lambda = \left(\frac{\Gamma(1/\nu)}{2^{2/\nu}\Gamma(3/\nu)}\right)^{\frac{1}{2}}.$$

The model we use uses these three distributions discussed above to capture the structural effect of the returns.

The individual GARCH models are all estimated using the *arch* Python package (Sheppard, 2025). For those interested, the procedure of estimating these models using the method of Maximum Likelihood, including the derivations of the log-likelihood functions for each underlying distributional assumption, is discussed in my thesis "*The effect of extreme volatility on the accuracy of GARCH models in VaR estimation*" which is available upon request.

3.3 Ensemble GARCH approach for VaR estimation

As discussed in the previous sections, there are many options in choosing the appropriate GARCH model to model an individual stock. One must choose the lag order combination (p, q), the specific GARCH extension and the underlying assumed distribution of the shocks. The sections above discussed only a small subset of the possible choices, and it must be critically noted that many extensions and improvements have been made to dynamically model the stock returns.

Model specification and selection

Rather than selecting a single GARCH specification, we implement an ensemble approach that combines multiple GARCH models to leverage the strengths of different specifications. This is a commonly used approach in literature, where GARCH models are combined with each other or even other types of models such as machine learning models (see for example Aras (2021), Lahmiri and Boukadoum (2015) or Kakade et al. (2022)). Our ensemble considers 30 distinct model configurations, varying across multiple dimensions: (1) the GARCH model type (standard GARCH, EGARCH and TGARCH, discussed in Subsection 3.1), (2) the lag order combinations (p, q) including (1,1), (1,2), (2,1) and (2,2) for GARCH and (1,1), (1,2), (2,1) for EGARCH and TGARCH, and (3) the distributional assumption of the shocks z_t (normal, Student-t and GED, discussed in Subsection 3.2).

Rolling window VaR backtesting

To evaluate and weight these models, we use a rolling window backtesting procedure. Starting with an initial training window of at least 500 observations, we fit each GARCH specification to the training data and generate one-step-ahead Value-at-Risk (VaR) forecasts at the 5% confidence level. The window is then rolled forward by a fixed rebalancing frequency and the process repeats until the end of the sample period. This approach finds a balance between periodically refitting the GARCH models but keeping realistic trading conditions where a certain model is used for a longer period.

For each forecast, we calculate the VaR as:

$$VaR_t = (-\mu_t + z_\alpha \cdot \sigma_t),$$

where μ_t is the expected return, σ_t is the forecasted volatility from the GARCH model and z_α is the appropriate quantile of the assumed distribution.

Model evaluation criteria

Each model's performance is evaluated using multiple criteria:

1. **Christoffersen's Conditional Coverage Test**: This test, introduced by Christoffersen (1998), examines both the unconditional coverage (whether the observed exceedance rate matches the expected rate), a test for which was introduced by Kupiec (1995), and the independence of the exceedances. Models that fail these tests (p-values below 0.05) indicate poor VaR model specification.

2. **Asymmetric Quantile Loss Function**: We employ an asymmetric loss function that penalizes VaR exceedances more heavily than non-exceedances:

$$L_{\text{asym}}(r_t, \text{VaR}_t) = \begin{cases} \alpha \cdot (L_t - \text{VaR}_t) & \text{if } L_t > \text{VaR}_t \\ (1-\alpha) \cdot (\text{VaR}_t - L_t) & \text{otherwise,} \end{cases}$$

   where $L_t = -r_t$ is the realized loss and α is the confidence level (Koenker and Bassett Jr (1978)).

3. **Regulatory Loss Function**: We also implement a regulatory-style loss function that includes a magnitude penalty for severe exceedances, taking tail risk management into account. Here we take the asymmetric quantile loss function to obtain a "base loss" and amplify this loss by the magnitude penalty.

4. **Coverage Accuracy**: Lastly, we use the deviation between the empirical exceedance rate and the target confidence level to give a direct measure of the model calibration.

Ensemble weight determination

The ensemble weights are determined through a multi-criteria optimization approach. For the models passing the statistical tests (Christoffersen p-values > 0.05), we calculate weights based on four components:

- **Statistical validity** (40% weight): Based on the geometric mean of the unconditional and conditional coverage p-values.
- **Regulatory loss performance** (30% weight): Inverse of the average regulatory loss
- **Coverage accuracy** (20% weight): Inverse of absolute deviation from target exceedance rate
- **Asymmetric loss performance** (10% weight): Inverse of average asymmetric loss

The final weight for model i is then given by:

$$w_i = \frac{s_i}{\sum_{j=1}^{N} s_j},$$

where s_i is the weighted combination of the four component scores. It must be noted that these weights are hand-picked based on expected relevance. Improvements on these choices can therefore certainly be made.

If no models pass the statistical tests, we select the best five models based on a quality score using the same multi-criteria approach. This ensures the ensemble remains robust even when the market conditions cause poor individual performance.

Implementation and final forecasting

After determining the optimal weights through backtesting, we refit all selected models on the full dataset. The ensemble VaR forecast is then computed as the weighted average of the individual model forecasts.

4 Constant Conditional Correlation (CCC) model

The Constant Conditional Correlation (CCC) model, introduced by Bollerslev (1990), provides an effective solution to modeling multivariate volatility. The fundamental assumption of the CCC model is that while individual asset volatilities evolve dynamically over time, the correlation structure between assets remains constant. This simplification dramatically reduces the number of parameters to be estimated while still capturing important features of multivariate return dynamics.

For a portfolio of n assets, the CCC model decomposes the conditional covariance matrix $\mathbf{\Sigma}_t$ as:

$$\mathbf{\Sigma}_t = \mathbf{D}_t \mathbf{R} \mathbf{D}_t,$$

where $\mathbf{D}_t = \text{diag}(\sigma_{1,t}, \dots, \sigma_{n,t})$ is a diagonal matrix containing the conditional standard deviations of each asset, and $\mathbf{R}$ is the constant conditional correlation matrix. Each individual variance $\sigma^2_{i,t}$ follows a univariate GARCH process, as is shown in Section 3.

The correlation matrix $\mathbf{R}$ has the structure:

$$\mathbf{R} = \begin{bmatrix} 1 & \rho_{12} & \cdots & \rho_{1n} \\ \rho_{12} & 1 & \cdots & \rho_{2n} \\ \vdots & \vdots & \ddots & \vdots \\ \rho_{1n} & \rho_{2n} & \cdots & 1 \end{bmatrix}$$

where ρ_{ij} represents the constant conditional correlation between assets i and j. This ensures that $\text{Corr}(y_{it}, y_{jt}|\mathcal{Y}^{t-1}) = \rho_{ij}$ for all t, while the conditional covariance $\text{Cov}(y_{it}, y_{jt}|\mathcal{Y}^{t-1}) = \sigma_{ijt} = \sigma_{it}\sigma_{jt}\rho_{ij}$ varies through time via the dynamic volatilities.

4.1 Implementation using ensemble GARCH models

Our implementation uses the ensemble GARCH approach described in Subsection 3.3 to obtain robust estimates of the individual conditional volatilities. The CCC estimation procedure follows the following approach:

Stage 1: Univariate volatility modeling: For each asset i in the portfolio, we fit the ensemble of 30 GARCH specifications and determine the optimal weights through the rolling window VaR backtesting procedure (see Subsection 3.3 for details). The weighted ensemble provides:

- Conditional volatility series: $\hat{\sigma}_{it} = \sum_{m=1}^{M} w_{im} \hat{\sigma}_{it}^{(m)}$
- Standardized residuals: $\hat{z}_{it} = \varepsilon_{it}/\hat{\sigma}_{it}$,

where w_{im} is the weight for model m of asset i determined through the multi-criteria optimization.

Stage 2: Correlation matrix estimation: The constant correlation matrix $\mathbf{R}$ is estimated from the standardized residuals obtained from the ensemble models:

$$\hat{\rho}_{ij} = \frac{1}{T}\sum_{t=1}^{T} \hat{z}_{it}\hat{z}_{jt}.$$

Using this approach, the correlation estimation is based on properly filtered returns that account for the time-varying volatility dynamics captured by the ensemble models.

4.2 Forecasting and portfolio optimization

Given the fitted CCC model, the h-step ahead covariance matrix forecast is:

$$\hat{\mathbf{\Sigma}}_{t+h|t} = \hat{\mathbf{D}}_{t+h|t}\mathbf{R}\hat{\mathbf{D}}_{t+h|t},$$

where $\hat{\mathbf{D}}_{t+h|t}$ contains the h-step ahead volatility forecasts from each ensemble models. Here, the volatility forecasts for each individual asset is obtained with the ensemble approach by:

$$\hat{\sigma}_{i,t+h|t} = \sum_{m=1}^{M} w_{im}\hat{\sigma}^{(m)}_{i,t+h|t}.$$

This estimated CCC model is the foundation for the services that can be explored at RiskOptimix. An example of this is the 'optimal portfolio weights' service, where the forecasted covariance matrix serves as input to the minimum variance portfolio optimization:

$$\min_{\mathbf{w}} \mathbf{w}^T\hat{\mathbf{\Sigma}}_{t+h|t}\mathbf{w} \ \text{ subject to } \ \mathbf{w}^T\mathbf{1} = 1, \ w_i \geq 0.$$

4.3 Model advantages and limitations

The CCC model has significant computational advantages. It requires only $n(n-1)/2$ correlation parameters, next to the univariate GARCH parameters. However, the assumption of constant correlations can be restrictive in practice. Financial markets often exhibit time-varying correlations, particularly during crisis periods when correlations tend to increase (Wong, Vlaar, et al., 2003). We therefore chose to move from the CCC model to the Dynamic Conditional Correlation (DCC) model, which will be discussed in the next chapter.

4.4 Empirical results

To show the empirical application of the CCC model, we estimated simple GARCH$(1,1)$ models for all assets in our portfolio. We then calculated the conditional variance of each asset and the correlation between the assets. We used this to then compute the conditional covariance. All results can be seen in Figure 3. Here, the spike in conditional variance can be seen for both assets at the start of the COVID-19 crisis in 2020, and this logically translates to the conditional covariance. As expected, the correlation between the two assets here is constant and lies at 0.5016.

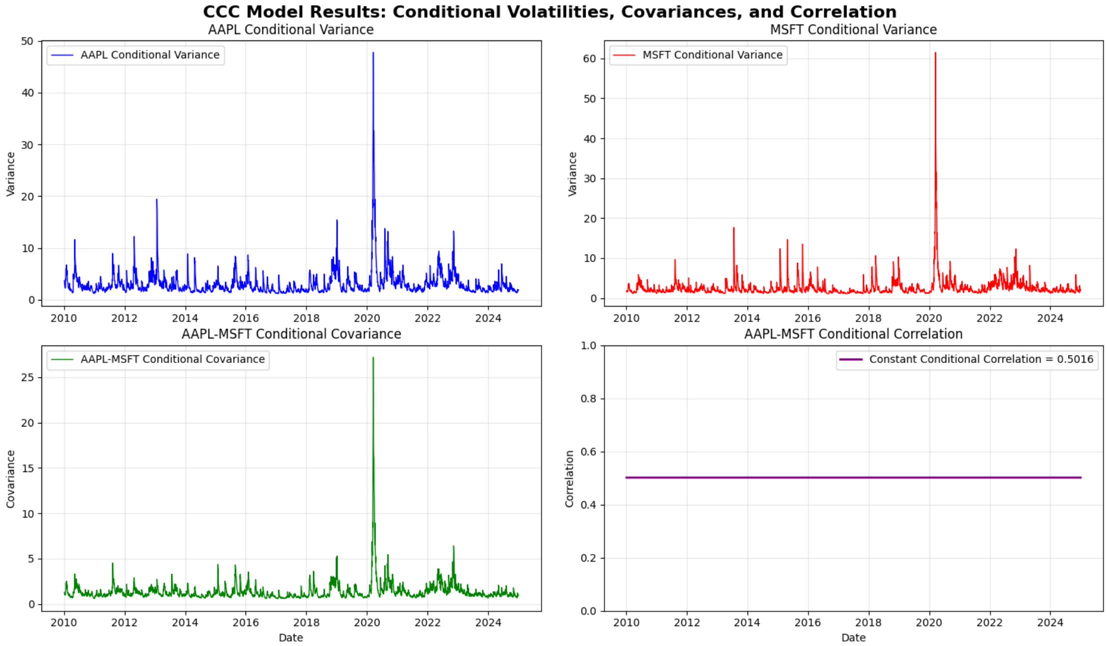

Figure 3: This figure contains the results of the estimated CCC model for the portfolio given in Table 1. Here the conditional variance of each asset is plotted, as well as the conditional covariance and the conditional correlation.

The "optimal" portfolio weights over time of our portfolio, based on the portfolio optimization discussed in Subsection 4.2 can be seen in Figure 4. While it can be seen that the optimal weights shift heavily over time, it can be seen that overall, the MSFT seems to be higher. This is also confirmed by calculating the average weights: AAPL: 0.4146 and MSFT: 0.5854.

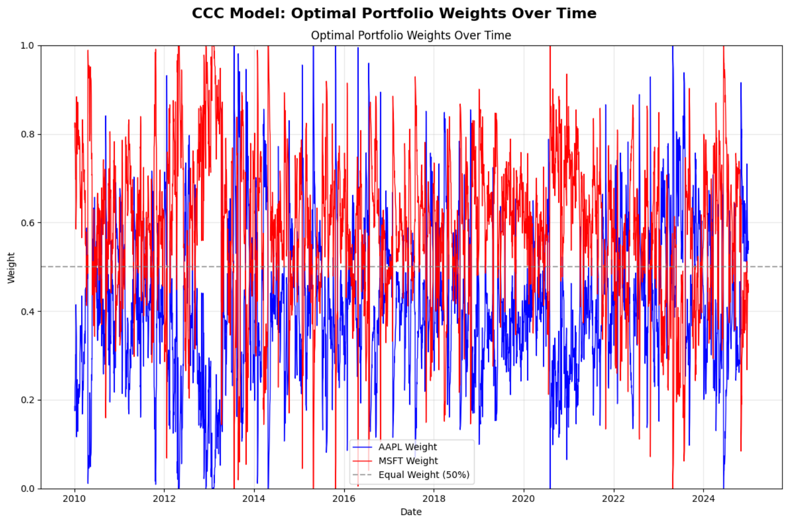

Figure 4: This figure contains the results of the portfolio optimization of the portfolio given in Table 1 discussed in Subsection 4.2, computed using the estimated CCC model.

5 Dynamic Conditional Correlation (DCC) Model

While the CCC model is computationally efficient, its assumption of constant correlations can be overly restrictive. The Dynamic Conditional Correlation (DCC) model, introduced by R. Engle (2002), addresses this limitation by allowing correlations to evolve dynamically while maintaining the two-stage estimation approach.

The DCC model maintains the same covariance decomposition as the CCC:

$$\mathbf{\Sigma}_t = \mathbf{D}_t \mathbf{R}_t \mathbf{D}_t,$$

where the crucial difference is that the correlation matrix $\mathbf{R}_t$ is now time-varying. The dynamics are introduced through an auxiliary process for the quasi-correlation matrix $\mathbf{Q}_t$:

$$\mathbf{Q}_t = (1 - \alpha - \beta)\bar{\mathbf{Q}} + \alpha \mathbf{z}_{t-1}\mathbf{z}_{t-1}^T + \beta \mathbf{Q}_{t-1},$$

where $\bar{\mathbf{Q}}$ is the unconditional covariance matrix of the standardized residuals $\mathbf{z}_t$ and α and β are scalar parameters for the dynamics. The correlation matrix is then obtained by rescaling:

$$\mathbf{R}_t = \text{diag}(\mathbf{Q}_t)^{-1/2}\mathbf{Q}_t \text{diag}(\mathbf{Q}_t)^{-1/2}.$$

The parameters must satisfy $\alpha \geq 0$, $\beta \geq 0$ and $\alpha + \beta < 1$ to ensure stationarity.

5.1 Estimation procedure

Our implementation follows the two-stage DCC estimation approach:

Stage 1: Univariate volatility estimation: This stage is identical to the CCC model. We use the ensemble GARCH approach to obtain robust estimates of individual volatilities and obtain the standardized residuals $\hat{z}_{it} = \varepsilon_{it}/\hat{\sigma}_{it}$ for each asset.

Stage 2: DCC parameter estimation: Given the standardized residuals from Stage 1, we estimate the DCC parameters (α, β) by maximizing the correlation component of the log-likelihood:

$$\mathcal{L}_c = -\frac{1}{2}\sum_{t=1}^{T}\left(\log|\mathbf{R}_t| + \mathbf{z}_t^T \mathbf{R}_t^{-1}\mathbf{z}_t - \mathbf{z}_t^T \mathbf{z}_t\right). \tag{5}$$

This component is obtained as follows. We start with the log-likelihood of the returns $\mathbf{r}_t$:

$$\mathcal{L} = -\frac{1}{2}\sum_{t=1}^{T}\left[n\log(2\pi) + \log|\mathbf{\Sigma}_t| + \mathbf{r}_t^T \mathbf{\Sigma}_t^{-1}\mathbf{r}_t\right].$$

Then noticing that $\mathbf{\Sigma}_t = \mathbf{D}_t \mathbf{R}_t \mathbf{D}_t$ this gives:

$$\begin{aligned}
\mathcal{L} &= -\frac{1}{2}\sum_{t=1}^{T}\left[n\log(2\pi) + 2\log|\mathbf{D}_t| + \log|\mathbf{R}_t| + \mathbf{r}_t^T \mathbf{D}_t^{-1}\mathbf{R}_t^{-1}\mathbf{D}_t^{-1}\mathbf{r}_t\right] \\
&= -\frac{1}{2}\sum_{t=1}^{T}\left[n\log(2\pi) + 2\log|\mathbf{D}_t| + \log|\mathbf{R}_t| + \mathbf{z}_t^T \mathbf{R}_t^{-1}\mathbf{z}_t\right].
\end{aligned}$$

This log-likelihood can be split into two parts: $\mathcal{L} = \mathcal{L}_v + \mathcal{L}_c$, where the part that depends only on $\mathbf{R}_t$, the correlation component, is then given by Equation (5).

The optimization is performed using Sequential Least Squares Programming (SLSQP) with constraints ensuring $\alpha, \beta > 0$ and $\alpha + \beta < 1$. The unconditional correlation matrix $\bar{\mathbf{Q}}$ is estimated as:

$$\bar{\mathbf{Q}} = \frac{1}{T} \sum_{t=1}^{T} \hat{\mathbf{z}}_t \hat{\mathbf{z}}_t^T.$$

5.2 Dynamic correlation generation and forecasting

Once the parameters are estimated, we generate the complete time series of correlation matrices by iterating through the DCC recursion. For forecasting, the h-step ahead correlation matrix is computed as:

$$\mathbf{Q}_{t+h|t} = (1 - \beta^h)\bar{\mathbf{Q}} + \beta^h \mathbf{Q}_t.$$

The forecasted covariance matrix combines the correlation forecast with volatility forecasts from the ensemble models:

$$\hat{\mathbf{\Sigma}}_{t+h|t} = \hat{\mathbf{D}}_{t+h|t} \hat{\mathbf{R}}_{t+h|t} \hat{\mathbf{D}}_{t+h|t}.$$

5.3 Empirical results

Similarly as for the CCC model in Subsection 4.4 we estimated GARCH(1,1) models for both assets and used this to estimate a DCC model. The results can be seen in Figure 5. The plots in the top row should be equal to the plots in Figure 3, but a difference should be observed in the plot of the conditional variance. That is, now not the constant conditional correlation is used, but instead the dynamic conditional correlation. The evolution of this dynamic correlation can be seen in the last plot, where it can, as expected, be seen that this evolves vividly over time. It is however interesting, but also logical, to see that the average of this estimated conditional correlation lies at almost the same value of 0.5015 of the constant correlation of the CCC model.

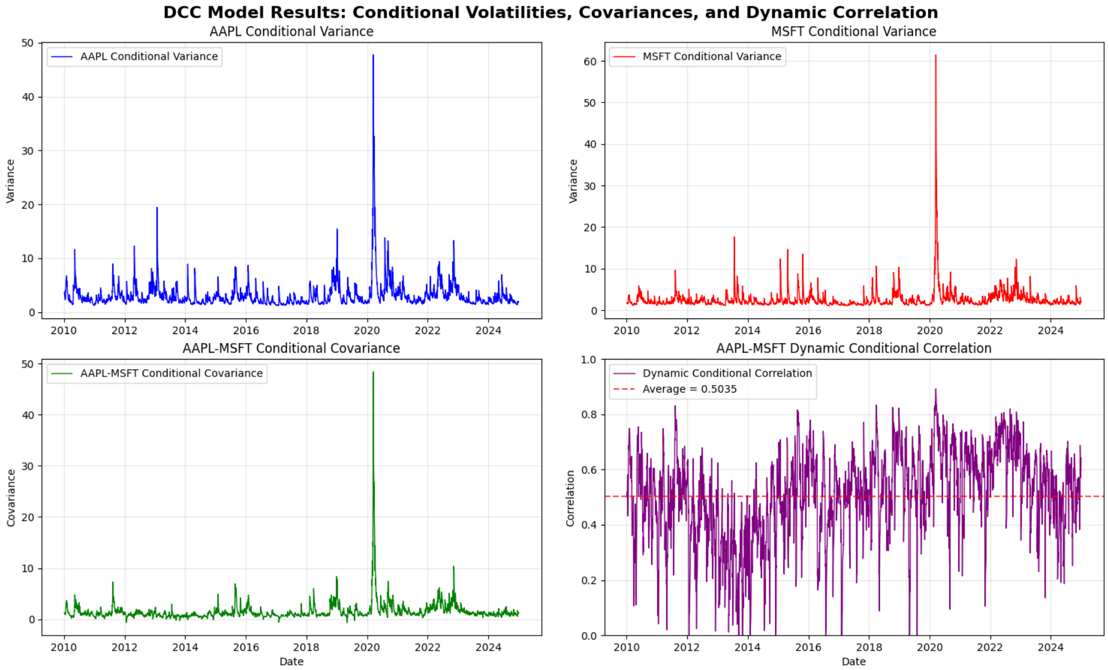

Figure 5: This figure contains the results of the estimated DCC model for the portfolio given in Table 1. Here the conditional variance of each asset is plotted, as well as the conditional covariance and the conditional correlation.

Calculating the optimal portfolio weights of our portfolio, this time using the DCC model, gives the results in Figure 6. These results look similar as the results in Figure 4 and the average weights are also almost equal to those obtained by the CCC model: AAPL: 0.4121 and MSFT: 0.5879. This shows that the models interestingly perform quite similarly over time. This is, of course, measured over a long period of time in close-up measures could reveal substantial results. Furthermore, the performance and the difference in performance of the models heavily depend on the specific setting of the experiment. It is however a nice way of showing the overall performance of the CCC and DCC model.

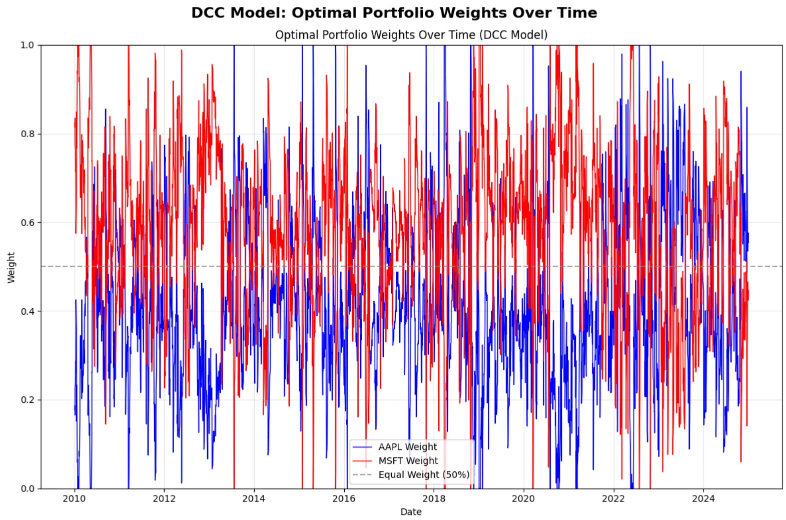

Figure 6: This figure contains the results of the portfolio optimization of the portfolio given in Table 1 discussed in Subsection 4.2, computed using the estimated DCC model.

6 Conclusion

This paper showed the foundation of the model behind RiskOptimix and its mathematical background. We hope this gave further insights into why this approach is used and why it can be useful.

The code that implements this dynamically for any given portfolio can be found on www.RiskOptimix.com/econometric-models.

For further questions, notes, or any insights, feel free to reach out to riskoptimix@gmail.com.

7 References

Angelidis, Timotheos, Alexandros Benos, and Stavros Degiannakis (2004). "The use of GARCH models in VaR estimation". In: *Statistical Methodology* 1.1-2, pp. 105–128. DOI: 10.1016/j.stamet.2004.08.004.

Aras, Serkan (2021). "Stacking hybrid GARCH models for forecasting Bitcoin volatility". In: *Expert Systems with Applications* 174, p. 114747.

Black, F. (1976). "Studies of Stock Market Volatility Changes". In: *Proceedings of the American Statistical Association, Business & Economic Statistics Section, 1976*. URL: https://cir.nii.ac.jp/crid/1570009749981528192.

Bollerslev, Tim (1986). "Generalized autoregressive conditional heteroskedasticity". In: *Journal of Econometrics* 31.3, pp. 307–327. ISSN: 0304-4076. DOI: https://doi.org/10.1016/0304-4076(86)90063-1. URL: https://www.sciencedirect.com/science/article/pii/0304407686900631.

— (1987). "A Conditionally Heteroskedastic Time Series Model for Speculative Prices and Rates of Return". In: *The Review of Economics and Statistics* 69.3, pp. 542–547. ISSN: 00346535, 15309142. URL: http://www.jstor.org/stable/1925546.

— (1990). "Modelling the coherence in short-run nominal exchange rates: a multivariate generalized ARCH model". In: *The review of economics and statistics*, pp. 498–505.

Christoffersen, Peter F. (1998). "Evaluating Interval Forecasts". In: *International Economic Review* 39.4, pp. 841–862. ISSN: 00206598, 14682354. URL: http://www.jstor.org/stable/2527341.

Engle, Robert (2002). "Dynamic conditional correlation: A simple class of multivariate generalized autoregressive conditional heteroskedasticity models". In: *Journal of business & economic statistics* 20.3, pp. 339–350.

Engle, Robert F. (1982). "Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation". In: *Econometrica* 50.4, pp. 987–1007. ISSN: 00129682, 14680262. URL: http://www.jstor.org/stable/1912773.

Glosten, Lawrence R, Ravi Jagannathan, and David E Runkle (1993). "On the relation between the expected value and the volatility of the nominal excess return on stocks". In: *The journal of finance* 48.5, pp. 1779–1801.

Gosset, WS (1908). "The probable error of a mean". In: *Biometrika* 6, pp. 1–25.

Kakade, Kshitij, Ishan Jain, and Aswini Kumar Mishra (2022). "Value-at-Risk forecasting: A hybrid ensemble learning GARCH-LSTM based approach". In: *Resources Policy* 78, p. 102903.

Koenker, Roger and Gilbert Bassett Jr (1978). "Regression quantiles". In: *Econometrica: journal of the Econometric Society*, pp. 33–50.

Kupiec (1995). *Techniques for verifying the accuracy of risk measurement models*. Vol. 95. 24. Division of Research and Statistics, Division of Monetary Affairs, Federal ...

Lahmiri, S and M Boukadoum (2015). "An ensemble system based on hybrid EGARCH-ANN with different distributional assumptions to predict S&P 500 intraday volatility". In: *Fluctuation and noise letters* 14.01, p. 1550001.

Nelson, Daniel B. (1991). "Conditional Heteroskedasticity in Asset Returns: A New Approach". In: *Econometrica* 59.2, pp. 347–370. ISSN: 00129682, 14680262. URL: http://www.jstor.org/stable/2938260.

Sheppard, Kevin (2025). *arch: ARCH models in Python*. Version 7.0.0. Python package for ARCH/GARCH and other volatility modeling. URL: https://github.com/bashtage/arch.

Subbotin, MT (1923). "On the law of frequency of errors". In: *Matematicheskii Sbornik* 31, pp. 296–301.

Wong, ASK, Peter JG Vlaar, et al. (2003). "Modelling time-varying correlations of financial markets". In: *De Nederlandsche Bank NV Research Memorandum* 739, p. 0319.

Yahoo! Finance (2025). *Yahoo! Finance - Stock Market Live, Quotes, Business & Finance News*. Accessed: August 26, 2025. URL: https://finance.yahoo.com.

Zakoian, Jean-Michel (1994). "Threshold heteroskedastic models". In: *Journal of Economic Dynamics and Control* 18.5, pp. 931–955. ISSN: 0165-1889. DOI: https://doi.org/10.1016/0165-1889(94)90039-6. URL: https://www.sciencedirect.com/science/article/pii/0165188994900396.