

Estimation and Backtesting of the Expected Shortfall and Value at Risk using Vine Copulas

An unconditional and conditional rolling window approach

Philipp Emanuel Maria Sommer

Thesis for the attainment of the academic degree

Master of Science (M.Sc.)

at the Department of Mathematics of the Technical University of Munich.

Examiner:

Prof. Claudia Czado Ph.D.

Supervisors:

Prof. Claudia Czado Ph.D., M.Sc. Karoline Bax

Submitted:

Munich, 13.04.2022

I hereby declare that this thesis is entirely the result of my own work except where otherwise indicated. I have only used the resources given in the list of references.

Munich, 13.04.2022

Philipp Emanuel Maria Sommer

Abstract

Accurate estimation of risk measures for financial portfolios is of great importance equally for financial institutions as well as regulators. Many existing methods lack the ability to adequately incorporate the high dimensional dependence structure of the financial portfolio.

In this research we capture the cross dependence of the assets using the flexible class of R-vine copulas and their trend and volatility univariately with ARMA-GARCH models. Given these two components we simulate portfolio level forecasts and follow a Monte Carlo approach to estimate various risk measures on the portfolio level. All of this is performed in a rolling window fashion. This approach extends the work of Maarouf (2021) as not only the Value at Risk (VaR) but also the expected shortfall (ES), which is the successor of the VaR as the primary market risk measure as of the Basel III accords, is considered. A detailed discussion of applicable backtesting strategies for the VaR and ES estimates is also provided.

Moreover, in this thesis we introduce a conditional estimation approach. We specify one or two additional market indices or other main market players that are also univariately modelled via ARMA-GARCH models and then incorporated in the D-vine copula which is used in the conditional setting for cross dependence modelling. Then we simulate the portfolio level forecasts conditionally on the market index or indices which leads to conditional portfolio level risk estimates. We then introduce a quantile based approach to observe the behavior of the risk measures given a certain state of the conditioning asset or assets. In particular, this can generate valuable insights in stress testing situations. An important part of the conditional approach is the conditional sampling from a D-vine copula. Thus, in order to facilitate this approach this thesis also introduces algorithms to sample from the rightmost or the two rightmost leaves of a D-vine copula.

The last part of the thesis covers multiple case studies on a Spanish stock portfolio where we successfully apply all the presented methods. Additionally the algorithms were optimized for computational efficiency. The complete code is publicly available and the developed R package *portvine* provides efficient implementations for all the risk measure estimation approaches proposed in this thesis.

Contents

Abstract	v
Introduction	1
I Methodological background	3
1 Notational conventions and general setup	5
2 Univariate financial time series	7
3 Risk measures	9
3.1 Definition and desired properties	9
3.2 Value at Risk (VaR)	9
3.3 Expected shortfall (ES)	10
4 Univariate financial time series models	13
4.1 Autoregressive models (AR)	13
4.2 Moving average models (MA)	14
4.3 Autoregressive moving average models (ARMA)	16
4.4 Generalized autoregressive conditional heteroscedasticity models (GARCH)	19
4.5 Assessment of model quality	20
5 Vine copulas	25
5.1 Bivariate dependence measures	25
5.2 Copula fundamentals	26
5.3 Bivariate copula families	28
5.4 Introduction to vine copulas	31
5.5 Regular (R) vine copulas	31
5.6 Drawable (D) vine copulas	33
5.7 Recursion for conditional distribution functions and estimation	33
6 Simulation from vine copulas	35
6.1 General simulation from multivariate distributions	35
6.2 Sampling from a D-vine copula	35
6.3 Sampling from a D-vine copula conditional on the rightmost leaf node	37
6.4 Sampling from a D-vine copula conditional on the two rightmost leaf nodes	45
7 Vine copula based risk measure estimation	47
7.1 Unconditional vine copula based risk measure estimation	47
7.2 Conditional vine copula based risk measure estimation using a single market index as conditioning variable	51
7.3 Conditional vine copula based risk measure estimation using two market indices as conditioning variables	55
8 Backtesting of the estimated risk measures	57
8.1 Traditional backtesting	57

8.2 Comparative backtesting	63
II Case studies	67
9 R package portvine	69
9.1 Introduction	69
9.2 Performance assessment	69
10 Unconditional Spanish stock portfolio analysis	77
10.1 The setup	77
10.2 Assessment of the marginal models	80
10.3 The fitted R-vines	85
10.4 Expected Shortfall estimators	89
10.5 Backtesting	93
11 Conditional Spanish stock portfolio analysis	103
11.1 One conditional asset	103
11.2 Two conditional assets	124
Conclusion	129
A Appendix	131
Bibliography	145

List of Figures

1.1	The densities of the standard normal ($\mathcal{N}(0, 1)$) and standard Student's t ($t_{\nu=2}(0, 1)$) distribution. $\nu > 0$ denotes the degree of freedom parameter for the Student's t distribution.	6
2.1	Illustration of a weakly stationary zero mean time series. The blue LOESS smoothing line helps to assess the constant mean assumption.	8
3.1	Illustration of the Value at Risk for the continuous random variable R_t^P using its distribution function $F_{R_t^P}$	10
4.1	Illustration of the sample paths and the corresponding empirical autocovariance function $\hat{\gamma}$ of three AR(1) models differing in their parameter ϕ	15
4.2	Illustration of the sample paths of three MA(1) models with differing parameters θ	16
4.3	Plots of the standardized residuals z_t over time for four different sampled residual series all with sample size 100.	21
4.4	Plots of the empirical ACF of the standardized residuals z_t for four different sampled residual series. The blue dashed lines indicate the 95% two sided confidence interval for the empirical ACF under the null hypothesis of white noise.	22
4.5	Normal Q-Q plots of the standardized residuals z_t for four different sampled residual series. The identity is given as blue line.	22
4.6	Plots of the p-values of Ljung-Box tests for different lags h based on the standardized residuals z_t for four different sampled residual series. The blue dashed line highlights the 0.05 confidence level.	23
5.1	Marginal normal contour plots of the presented bivariate copula families.	30
5.2	Exemplary R-vine tree sequence on 5 elements.	32
5.3	Exemplary D-vine tree sequence on 4 elements.	33
6.1	Illustration for Algorithm 2 in 4 dimensions.	38
6.2	Illustration of the tree sequence of a 4-dimensional D-vine with the market index I as the rightmost leaf node.	41
6.3	Show the build up of the matrices $\mathcal{V}, \mathcal{V}^2$ according to Algorithm 3 in a visual form for the 5 dimensional example.	44
7.1	Illustration of the one step ahead unconditional risk measure estimation approach using ARMA-GARCH marginal and R-vine dependence models.	48
7.2	Illustration of the parameters Γ and Ψ alongside the one step ahead unconditional risk measure estimation.	49
7.3	Illustration of the rolling parameters and the actual rolling of the respective windows of the rolling window approach.	49
7.4	Illustration of the one step ahead single conditional risk measure estimation approach using ARMA-GARCH marginal and D-vine dependence models.	52
8.1	Three zone approach for comparative backtesting with confidence level η as in Nolde and Ziegel (2017). The challenging model is assigned to one of the three regions <i>Passed</i> (green), <i>Failed</i> (red) or <i>Further investigation needed</i> (yellow) according to the test statistic given in Equation (8.24).	65

9.1	Illustration of the first two processing rounds of the 2/3 nested parallel strategy with regard to the vine windows. Here for illustrative purposes one assumes 4 vine windows with the size of 10 time units.	70
9.2	The influence of different simulation sample sizes and the respectively used parallel strategy on the runtime of the unconditional risk measure estimation implemented in the portvine package.	71
9.3	The influence of different numbers of variables and the respectively used parallel strategy on the runtime of the unconditional risk measure estimation implemented in the portvine package.	72
9.4	The influence of different sets of bivariate copula families and the respectively used parallel strategy on the runtime of the unconditional risk measure estimation implemented in the portvine package.	72
9.5	The influence of different simulation sample sizes and the respectively used parallel strategy on the runtime of the conditional risk measure estimation implemented in the portvine package. The observation for the strategy 10/25 and simulation sample size of 100000 corresponds to the estimation of only 1 conditional quantile as 2 conditional quantiles would come with the extensive memory requirements of >100Gb. From this fact we directly expect this observation to be roughly 33% faster.	73
9.6	The influence of the two simulation sample sizes 1000 and 10000 and the respectively used parallel strategy on the runtime of the conditional risk measure estimation implemented in the portvine package.	74
9.7	The influence of different numbers of variables and the respectively used parallel strategy on the runtime of the conditional risk measure estimation implemented in the portvine package.	75
9.8	The influence of different numbers of quantile levels α^I that are to be estimated and the respectively used parallel strategy on the runtime of the conditional risk measure estimation implemented in the portvine package.	75
9.9	The influence of different sets of bivariate copula families and the respectively used parallel strategy on the runtime of the conditional risk measure estimation implemented in the portvine package.	76
10.1	The daily log return series of the MSCI Spain with the time frames of interest highlighted. .	78
10.2	The daily log return series of the MSCI Spain and the individual assets.	79
10.3	The residual plots for the asset Iberdrola for the first marginal training window (01.01.2016 - 15.11.2018) and the first time frame (2016-2019).	81
10.4	The Ljung-Box heatmaps on the standardized residuals for the first time frame (2016-2019) and the first and last marginal training window (01.01.2016 - 15.11.2018 and 16.12.2016 - 31.10.2019).	82
10.5	The residual plots for the asset Iberdrola for the first marginal training window (02.04.2020 - 26.05.2021) and the second time frame (2020-2021).	83
10.6	The Ljung-Box heatmaps on the standardized residuals for the second time frame (2020-2021) and the first and last marginal training window (02.04.2020 - 26.05.2021 and 11.06.2020 - 04.08.2021).	84
10.7	The first trees of the R-vines corresponding to the first, sixth and twelfth vine window for the first time frame (2016-2019) in the unconditional setting.	86
10.8	The first trees of the R-vines corresponding to the first and fourth vine window for the second time frame (2020-2021) in the unconditional setting.	87
10.9	Comparison of the three ES estimators at confidence level $\alpha = 0.05$ over time for the first time frame (2016-2019).	89
10.10	Two views of the 3d plot relating the risk estimates of the three different estimators for the ES at confidence level $\alpha = 0.05$ for the first time frame (2016-2019).	90

10.11	Comparison of the three ES estimators at confidence level $\alpha = 0.05$ over time for the second time frame (2020-2021).	91
10.12	Two views of the 3d plot relating the risk estimates of the three different estimators for the ES at confidence level $\alpha = 0.05$ for the second time frame (2020-2021).	91
10.13	Visual comparison plots of the VaR and ES risk estimates at the two α levels 0.01 and 0.05 for the first and second time frame. The vine window parameters are $\Psi = 250$ and $\kappa = 25$ for the first time frame and $\Psi = 100$ and $\kappa = 25$ for the second time frame.	94
10.14	Exceedance plots of the VaR at α level 0.025 for the first and second time frame. The vine window parameters are $\Psi = 250$ and $\kappa = 25$ as well as $\Psi = 100$ and $\kappa = 25$ respectively.	95
10.15	Exemplary exceedance plot for the ES at α level 0.025 for the first time frame (2016-2019). The vine window parameters are $\Psi = 250$ and $\kappa = 25$	95
10.16	Comparison of the ES estimates at α level 0.05 based on the R-vine copula allowing for all parametric copula families and only Gaussian and Student's t ones for the first time frame (2016-2019).	100
10.17	Comparison of the ES estimates at α level 0.05 based on the R-vine copula allowing for all parametric copula families and only Gaussian and Student's t ones for the second time frame (2020-2021).	102
11.1	The residual plots for the conditional asset the SP500 index for the first marginal window and the first time frame (2016-2019).	104
11.2	The residual plots for the conditional asset the Eurostoxx 50 index for the first marginal window and the first time frame (2016-2019).	105
11.3	The first trees of the D-vines corresponding to the first and sixth/last vine window for the first time frame (2016-2019) in the single conditional setting. The conditional asset is the SP500	106
11.4	The first trees of the D-vines corresponding to the first and sixth/last vine window for the first time frame (2016-2019) in the single conditional setting. The conditional asset is the Eurostoxx 50	110
11.5	The first trees of the D-vine corresponding to the first and second/last vine windows for the second time frame (2020-2021) in the single conditional setting. The conditional asset is the SP500	112
11.6	The first trees of the D-vine corresponding to the first and second/last vine windows for the second time frame (2020-2021) in the single conditional setting. The conditional asset is the Eurostoxx 50	113
11.7	Behavior of the quantile based conditioning values on the log return scale of the conditioning asset SP500 for the first time frame (2016-2019).	114
11.8	Behavior of the quantile based conditioning values on the log return scale of the conditioning asset Eurostoxx 50 for the first time frame (2016-2019).	115
11.9	Behavior of the prior residual based conditioning values on the log return scale of the conditioning asset SP500 for the first time frame with a zoom at March 2019. The realized log returns of the conditional asset are drawn in grey and the conditioning values are drawn in violet.	116
11.10	Behavior of the realized residual based conditioning values on the log return scale of the conditioning asset SP500 for the first time frame with a zoom at March 2019. The realized log returns of the conditional asset are drawn in grey and the conditioning values are drawn in green.	116
11.11	Quantile based conditional ES estimates at $\alpha = 0.025$ with the conditioning asset SP500 for the first time frame (2016-2019).	117
11.12	Comparison of the quantile based conditional ES estimates at $\alpha = 0.025$ (blue) with the ones based on the realized residuals (green) for the first time frame (2016-2019). The conditioning asset is the SP500	118

11.13	Comparison of the quantile based conditional ES estimates at $\alpha = 0.025$ (blue) with the unconditional ES estimates (orange/gold) for the first time frame (2016-2019). The conditioning asset is the SP500 .	119
11.14	Quantile based conditional ES estimates at $\alpha = 0.025$ with the conditioning asset Eurostoxx 50 for the first time frame (2016-2019).	120
11.15	Comparison of the quantile based conditional ES estimates at $\alpha = 0.025$ (blue) with the ones based on the realized residuals (green) for the first time frame (2016-2019). The conditioning asset is the Eurostoxx 50 .	121
11.16	Comparison of the quantile based conditional ES estimates at $\alpha = 0.025$ (blue) with the unconditional ES estimates (orange/gold) for the first time frame (2016-2019). The conditioning asset is the Eurostoxx 50 .	122
11.17	Display of the conditional ES estimates at $\alpha = 0.025$ based on the prior residuals (violet) along side the realized portfolio log returns (grey) for the first time frame (2016-2019). The conditioning asset is the Eurostoxx 50 .	123
11.18	The first trees of the D-vine corresponding to the first and sixth/last vine window for the first time frame (2016-2019) in the double conditional setting. The conditional assets are the Eurostoxx 50 index and Iberdrola .	125
11.19	The first trees of the D-vine corresponding to the first and second/last vine window for the second time frame (2020-2021) in the double conditional setting. The conditional assets are the Eurostoxx 50 index and Iberdrola .	126
11.20	Comparisons of the quantile based conditional ES estimates at $\alpha = 0.025$ (blue) with the ones based on the realized residuals (green) for the first and second time frame. The conditioning assets are the Eurostoxx 50 index and Iberdrola .	127
A.1	The residual plots for the asset Banco Santander for the first marginal window and the first time frame (2016-2019).	131
A.2	The residual plots for the asset Inditex for the first marginal window and the first time frame (2016-2019).	131
A.3	The residual plots for the asset Cellnex Telecom for the first marginal window and the first time frame (2016-2019).	132
A.4	The residual plots for the asset Repsol Ypf for the first marginal window and the first time frame (2016-2019).	132
A.5	The residual plots for the asset Ferrovial for the first marginal window and the first time frame (2016-2019).	132
A.6	The residual plots for the asset Amadeus It Group for the first marginal window and the first time frame (2016-2019).	133
A.7	The residual plots for the asset Telefonica for the first marginal window and the first time frame (2016-2019).	133
A.8	The residual plots for the asset BBV Argentaria for the first marginal window and the first time frame (2016-2019).	133
A.9	The residual plots for the asset Banco Santander for the first marginal window and the second time frame (2020-2021).	134
A.10	The residual plots for the asset Inditex for the first marginal window and the second time frame (2020-2021).	134
A.11	The residual plots for the asset Cellnex Telecom for the first marginal window and the second time frame (2020-2021).	134
A.12	The residual plots for the asset Repsol Ypf for the first marginal window and the second time frame (2020-2021).	135
A.13	The residual plots for the asset Ferrovial for the first marginal window and the second time frame (2020-2021).	135
A.14	The residual plots for the asset Amadeus It Group for the first marginal window and the second time frame (2020-2021).	135

A.15	The residual plots for the asset Telefonica for the first marginal window and the second time frame (2020-2021).	136
A.16	The residual plots for the asset BBV Argentario for the first marginal window and the second time frame (2020-2021).	136
A.17	The residual plots for the conditional asset the SP500 index for the first marginal window and the second time frame (2020-2021).	136
A.18	The residual plots for the conditional asset the Eurostoxx 50 index for the first marginal window and the second time frame (2020-2021).	137
A.19	Behavior of the quantile based conditioning values of the conditioning asset SP500 for the second time frame (2020-2021).	140
A.20	Behavior of the quantile based conditioning values of the conditioning asset Eurostoxx 50 for the second time frame (2020-2021).	140
A.21	Quantile based conditional ES estimates at $\alpha = 0.025$ with the conditioning asset SP500 for the second time frame (2020-2021).	141
A.22	Comparison of the quantile based conditional ES estimates at $\alpha = 0.025$ (blue) with the ones based on the realized residuals (green) for the second time frame (2020-2021). The conditioning asset is the SP500.	141
A.23	Comparison of the quantile based conditional ES estimates at $\alpha = 0.025$ (blue) with the unconditional ES estimates (orange/gold) for the second time frame (2020-2021). The conditioning asset is the SP500.	142
A.24	Quantile based conditional ES estimates at $\alpha = 0.025$ with the conditioning asset Eurostoxx50 for the second time frame (2020-2021).	143
A.25	Comparison of the quantile based conditional ES estimates at $\alpha = 0.025$ (blue) with the ones based on the realized residuals (green) for the second time frame (2020-2021). The conditioning asset is the Eurostoxx 50.	143
A.26	Comparison of the quantile based conditional ES estimates at $\alpha = 0.025$ (blue) with the unconditional ES estimates (orange/gold) for the second time frame (2020-2021). The conditioning asset is the Eurostoxx 50.	144

List of Tables

5.1	A selection of bivariate archimedean single parametric copula families.	28
9.1	The influence of different simulation sample sizes and the respectively used parallel strategy on the runtime given in minutes of the unconditional risk measure estimation implemented in the portvine package. The number of variables is fixed to 10 and all parametric bivariate copula families are allowed for the fitting.	71
9.2	Overview of the runtime of some of the most interesting scenarios from the performance assessment of the unconditional risk measure estimation approach. In all displayed cases all parametric copula families were allowed, the observations per variable were fixed to $T = 1000$ and 5 marginal windows of size 50 as well as 10 vine windows of size 25 were used.	72
9.3	The influence of different simulation sample sizes and the respectively used parallel strategy on the runtime given in minutes of the conditional risk measure estimation implemented in the portvine package. The number of variables is fixed to 10 and all parametric bivariate copula families are allowed for the fitting. The entry for the strategy 10/25 and simulation sample size of 100000 corresponds to the estimation of only 1 conditional quantile as 2 conditional quantiles would come with the extensive memory requirements of >100Gb. From this fact we directly expect this entry to be roughly 33% smaller.	74
9.4	Overview of the runtime of some of the most interesting scenarios from the performance assessment of the conditional risk measure estimation approach. In all displayed cases all parametric copula families were allowed, the observations per variable were fixed to $T = 1000$ and 5 marginal windows of size 50 as well as 10 vine windows of size 25 were used. Besides the case in the fourth scenario, which would require memory of >100Gb for two quantile levels, always two quantile levels α^I were estimated.	75
10.1	The assets of the Spanish stock portfolio of interest with their respective market capitalization (EUR Billions) on the 29.10.2021.	77
10.2	The used copula families alongside their frequencies for the first, sixth and twelfth fitted R-vine for the first time frame (2016-2019) in the unconditional setting.	85
10.3	The used copula families alongside their frequencies for the first and fourth fitted R-vine for the second time frame (2020-2021) in the unconditional setting.	88
10.4	The results of the traditional backtests for the first time frame (2016-2019). In some cases (especially for small α levels) where no result was reported this was due to limitations of the ESR approach and its implementation. P-values smaller than 0.05 are marked bold	97
10.5	The results of the traditional backtests for the second time frame (2020-2021). In some cases (especially for small α levels) where no result was reported this was due to limitations of the ESR approach and its implementation, except for the one case where the exceedance residual test could not be performed as there were no exceedance residuals. P-values smaller than 0.05 are marked bold	98
10.6	The results i.e. $\Phi(\mathcal{T}_{comp})$ of the comparative backtests as of Definition 8.2.2 for the first time frame (2016-2019) for the ES at the two α levels 0.01 and 0.05 with respect to different rolling vine window parameters Ψ and κ . A value below 0.5 indicates a superior internal model but here in all cases no significant decision at confidence $\nu = 0.05$ can be supported.	99

10.7	The results i.e. $\Phi(\mathcal{T}_{comp})$ of the comparative backtests as of Definition 8.2.2 for the first time frame (2016-2019) for the ES at the two α levels 0.01 and 0.05 with respect to the copula families used in the R-vine. The parameters $\Psi = 250, \kappa = 50$ and all other parameters are fixed. A value below 0.5 indicates a superior internal model but here in all cases no significant decision at confidence $\nu = 0.05$ can be supported.	100
10.8	The results i.e. $\Phi(\mathcal{T}_{comp})$ of the comparative backtests as of Definition 8.2.2 for the second time frame (2020-2021) for the ES at the two α levels 0.01 and 0.05 with respect to different rolling vine window parameters Ψ and κ . A value below 0.5 indicates a superior internal model but here in all cases no significant decision at confidence $\nu = 0.05$ can be supported.	101
10.9	The results i.e. $\Phi(\mathcal{T}_{comp})$ of the comparative backtests as of Definition 8.2.2 for the second time frame (2020-2021) for the ES at the two α levels 0.01 and 0.05 with respect to the copula families used in the R-vine. The parameters $\Psi = 200, \kappa = 50$ and all other parameters are fixed. A value below 0.5 indicates a superior internal model but here in all cases no significant decision at confidence $\nu = 0.05$ can be supported.	102
11.1	The used copula families alongside their frequencies for the first and sixth fitted D-vine for the first time frame (2016-2019) in the single conditional setting. The conditional asset is the SP500	107
11.2	The used copula families alongside their empirical Kendall's tau associated with the conditional asset for the first and sixth fitted D-vine for the first time frame (2016-2019) in the single conditional setting. The conditional asset is the SP500	108
11.3	The used copula families alongside their frequencies for the first and sixth fitted D-vine for the first time frame (2016-2019) in the single conditional setting. The conditional asset is the Eurostoxx 50	109
11.4	The used copula families alongside their empirical Kendall's tau associated with the conditional asset for the first and sixth fitted D-vine for the first time frame (2016-2019) in the single conditional setting. The conditional asset is the Eurostoxx 50	111
A.1	The used copula families alongside their frequencies for the first and second fitted D-vine for the second time frame (2020-2021) in the single conditional setting. The conditional asset is the SP500.	138
A.2	The used copula families alongside their empirical Kendall's tau associated with the conditional asset for the first fitted D-vine for the second time frame (2020-2021) in the single conditional setting. The conditional asset is the SP500. ($j_1 = Iberdrola, j_2 = Ferrovial, j_3 = Inditex, j_4 = Repsol Ypf, j_5 = Amadeus It Group, j_6 = Banco Santander, j_7 = BBV Argentaria, j_8 = Telefonica, j_9 = Cellnex Telecom$)	138
A.3	The used copula families alongside their frequencies for the first and second fitted D-vine for the second time frame (2020-2021) in the single conditional setting. The conditional asset is the Eurostoxx 50.	139
A.4	The used copula families alongside their empirical Kendall's tau associated with the conditional asset for the first fitted D-vine for the second time frame (2020-2021) in the single conditional setting. The conditional asset is the Eurostoxx 50. ($j_1 = Amadeus It Group, j_2 = Inditex, j_3 = Ferrovial, j_4 = Banco Santander, j_5 = Repsol Ypf, j_6 = BBV Argentaria, j_7 = Telefonica, j_8 = Iberdrola, j_9 = Cellnex Telecom$)	139

Introduction

The effective assessment of risk is of great importance for financial institutions and thus also of critical concern for the corresponding regulators like the BaFin or the EBA. As currently evident accounting for market risk, that can manifest itself for example as high volatility clusters e.g. during the ongoing pandemic associated to the Covid-19 virus, is of utmost importance. Global guidelines for the assessment of risk in the financial context are constructed in the Basel accords. Their central tools in the assessment of risk are so called risk measures. To extend existing and propose new methods in order to effectively measure the risk of potentially large asset portfolios is the aim of this thesis. Moreover, we have developed and published a software package called *portvine* for the statistical programming language R to facilitate the usage of all proposed methods.

In this research having a fixed portfolio we model each of its assets given as a financial time series univariately in order to capture their trend and volatility with ARMA-GARCH models and then capture their cross dependence using the flexible class of R-vine copulas. Given these two components one can simulate portfolio level forecasts and follow a Monte Carlo approach to estimate various risk measures on the portfolio level. All of this is performed in a rolling window fashion. The central principle of this **unconditional risk measure estimation approach** presented in this thesis is illustrated in Figure A below. This

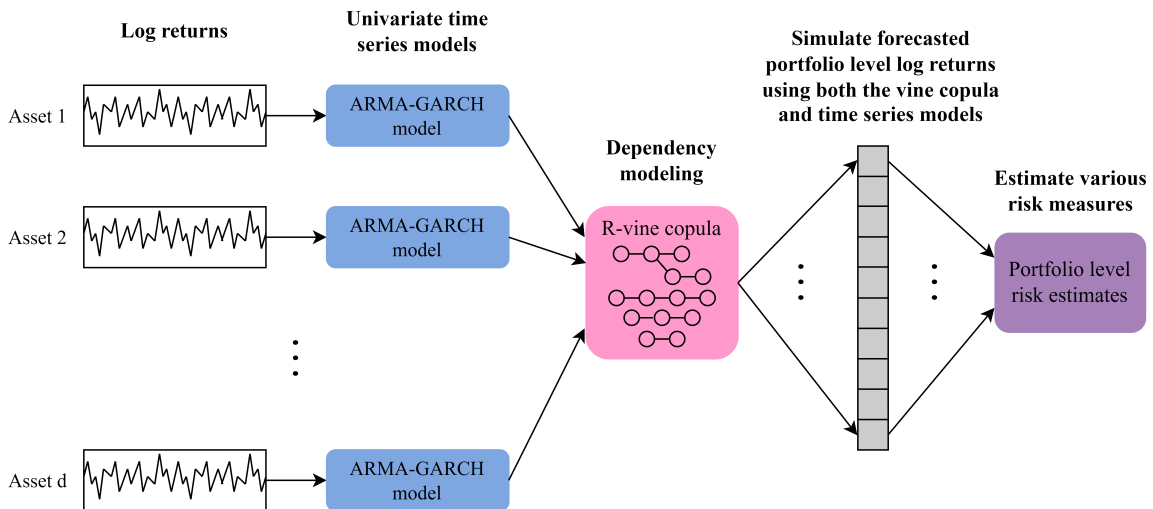


Figure A Illustration of the central principle of the unconditional risk estimation presented in this thesis.

approach extends the work of Maarouf (2021) as not only the Value at Risk (VaR) but also the expected shortfall as risk measure is considered. Additionally the algorithm was further optimized for computational efficiency. The expected shortfall as additional risk measure is especially important as it is the successor of the VaR as the primary market risk measure as of Basel Committee (2016, 2017) because it overcomes some major theoretical shortcomings of the VaR. Monstvilaitė (2016); Sharma and Sahni (2021) both propose the same unconditional risk measure estimation approach as Maarouf (2021) combining univariate ARMA-GARCH models with R-vine copula models as the flexible dependence model followed by a Monte Carlo approach to estimate the risk measures but also lack the coverage of the expected shortfall and do not provide software or optimized algorithms for the effective estimation. The same holds for Trucíos et al. (2019) who apply an approach based on Brechmann and Czardo (2013) only differing from the one of Maarouf (2021) in the usage of robust GARCH models instead of ARMA-GARCH models. They apply the approach to portfolios of cryptocurrencies and indeed cover the expected shortfall but they also neither

provide a detailed and optimized algorithm for the rolling window approach nor a publicly available implementation.

Moreover in this thesis we introduce a conditional estimation approach. One can specify one or two additional market indices or other main market players that are also univariately modelled via ARMA-GARCH models and then incorporated in the D-vine copula which is used in the conditional setting for cross dependence modelling. Then one simulates the portfolio level forecasts conditionally on the market index or indices which leads to conditional portfolio level risk estimates. Here as the conditioning values quantiles of the market index or indices are proposed which allows to observe the behavior of the risk measures conditioned on different states of the conditioning market index or indices. In particular this can generate valuable insights in stress testing situations. For one conditioning market index the central principle of this **conditional risk measure estimation approach** is illustrated in Figure B below. An important part of the conditional approach is the conditional sampling from a D-vine copula. Thus in order to facilitate this approach this thesis also introduces an algorithm to sample from the rightmost or the two rightmost leaves of a D-vine copula in Chapter 6.

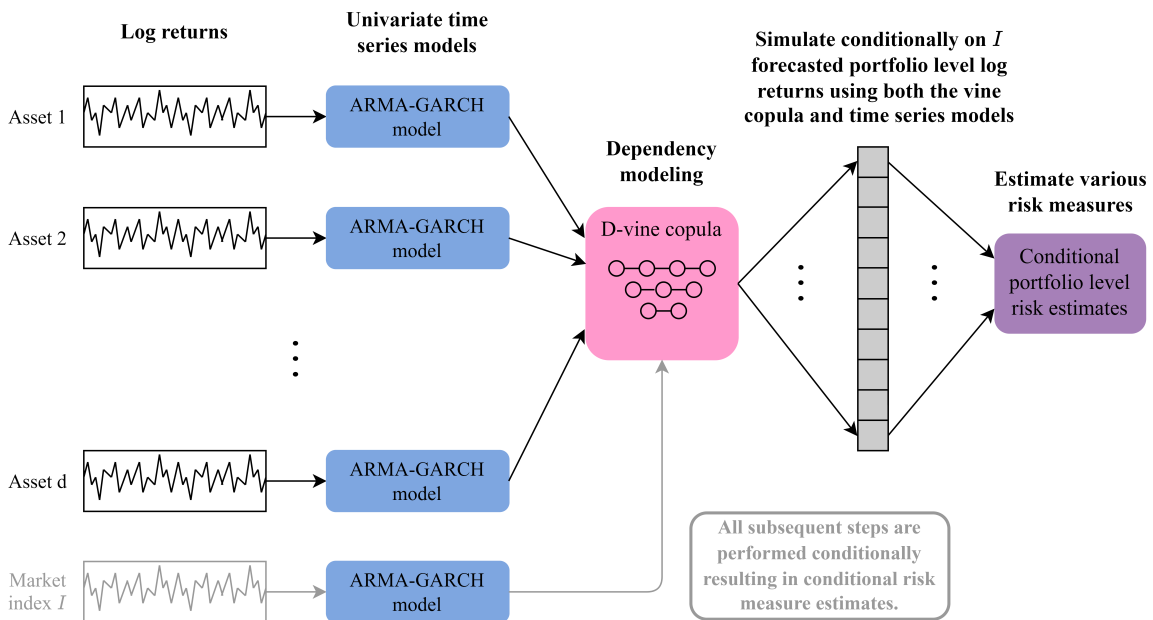


Figure B Illustration of the central principle of the conditional risk estimation presented in this thesis.

The thesis is split up into a methodological part and an applied part. The methodological part starts with Chapters 1 and 2 on the notational conventions and a primer on financial time series. After that in Chapter 3 the used risk measures are discussed in detail. Then Chapter 4 covers the needed theory on the univariate financial time series models. It is followed by Chapter 5 in which vine copula models are introduced and Chapter 6 covers the simulation from vine copulas where also the novel conditional simulation algorithms for D-vines can be found. Chapter 7 then formally presents the unconditional and conditional vine copula based risk measure estimation approaches and Chapter 8 wraps up the methodology part with a detailed walk through of applicable backtests for all considered risk measures. The applied part first presents the developed *portvine* R package in Chapter 9 which is publicly available and efficiently implements all the risk measure estimation approaches proposed in this thesis. This chapter also features performance measures of the implementations in the package. Finally Chapters 10 and 11 include case studies on a portfolio of Spanish stocks that show the practical application of all proposed methods.

Part I

Methodological background

1 Notational conventions and general setup

Regarding the mathematical notation this thesis will mostly adapt the one of Czado (2019). This includes uppercase random variables (RV) and lowercase notation for realizations of such i.e. $Y = y$. For absolutely continuous RVs the corresponding densities and distribution functions are existent and denoted by f and F respectively. The corresponding so called quantile function $Q = F^{-1}$ also exists. The same notation is used for the probability mass and distribution function of discrete RVs. In the following notation on distributions of random vectors, some essential univariate distributions and an important data transformation are shortly given following Czado (2019) and Spokoiny and Dickhaus (2015).

Definition 1.0.1 (Joint, marginal and conditional distributions). *Let $\mathbf{X} = (X_1, \dots, X_d)^T \in \mathbb{R}^d$ be a d dimensional absolutely continuous random vector. Then for the well defined joint, marginal and conditional densities and distribution functions the following notation is used:*

	density function	distribution function
Joint	$f(x_1, \dots, x_d)$	$F(x_1, \dots, x_d)$
Marginal	$f_j(x_j), j \in \{1, \dots, d\}$	$F_j(x_j), j \in \{1, \dots, d\}$
Conditional	$f_{j k}(x_j x_k), j \neq k$	$F_{j k}(x_j x_k), j \neq k$

Definition 1.0.2 (Uniform distribution). *An absolutely continuous RV $X \in \mathbb{R}$ is univariately uniformly distributed on the interval $[a, b]$ if it follows the density*

$$f(x) := \begin{cases} \frac{1}{b-a} & x \in [a, b] \\ 0 & \text{else} \end{cases} \quad (1.1)$$

and will be denoted as $X \sim \mathcal{U}(a, b)$.

Definition 1.0.3 (Normal distribution). *An absolutely continuous RV $X \in \mathbb{R}$ is univariately normally distributed with mean $\mu \in \mathbb{R}$ and variance $\sigma^2 > 0$ if it follows the density*

$$f(x; \mu, \sigma^2) := \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} \quad (1.2)$$

and will be denoted as $X \sim \mathcal{N}(\mu, \sigma^2)$. The widely used standard normal distribution is $\mathcal{N}(0, 1)$ and its density and distribution function are denoted as $\phi(\cdot)$ and $\Phi(\cdot)$ respectively.

Definition 1.0.4 (Student's t distribution). *An absolutely continuous RV $X \in \mathbb{R}$ is univariately Student's t distributed with mean $\mu \in \mathbb{R}$, variance $\sigma^2 > 0$ and degree of freedom parameter $\nu > 0$ if it follows the density*

$$f_\nu(x; \mu, \sigma^2) := \frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})\sqrt{\nu\pi\sigma^2}} \left\{ 1 + \left(\frac{x-\mu}{\sigma} \right)^2 \frac{1}{\nu} \right\}^{-\frac{\nu+1}{2}} \quad (1.3)$$

and will be denoted as $X \sim t_\nu(\mu, \sigma^2)$.

Notably the Student's t distribution has heavier tails than the normal distribution which is clearly visible in Figure 1.1. This behavior is the reason why this distribution is way more popular in the financial domain as empirically the data arising in finance often showed heavy tailed characteristics as outlined in Tsay (2010). A distribution that is especially important for upcoming hypothesis tests is the so called χ^2 -distribution.

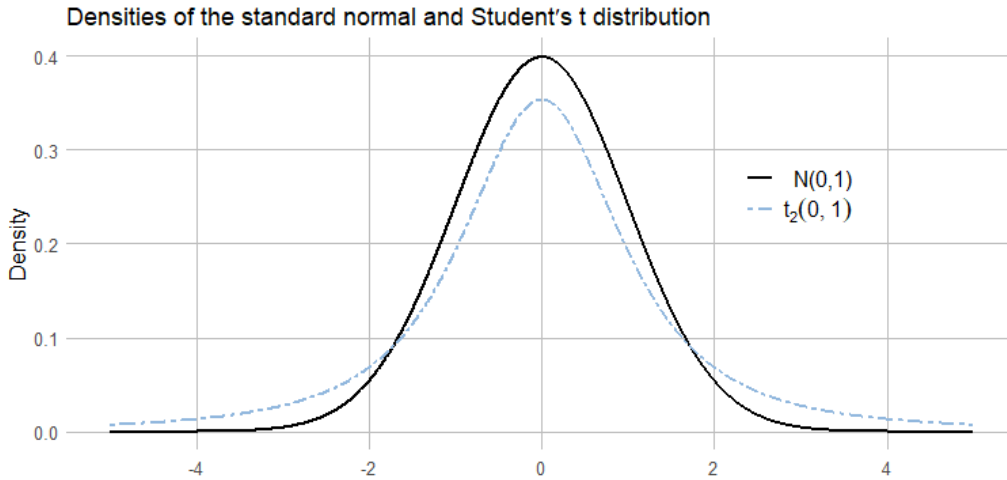


Figure 1.1 The densities of the standard normal ($\mathcal{N}(0, 1)$) and standard Student's t ($t_{\nu=2}(0, 1)$) distribution. $\nu > 0$ denotes the degree of freedom parameter for the Student's t distribution.

Definition 1.0.5 (χ^2 -distribution). For n i.i.d. samples of the standard normal distribution Y_i it holds that

$$\sum_{i=1}^n Y_i^2 \sim \chi^2(n) \quad (1.4)$$

where $\chi^2(n)$ denotes the χ^2 -distribution with n degrees of freedom.

Another important tool that will be of utmost importance in the main algorithms down the road is the so called probability integral transform that is defined next.

Definition 1.0.6 (Probability integral transform). Let X be a absolutely continuous RV with distribution function F . Then the transformation of realization x of X

$$u := F(x) \quad (1.5)$$

is the probability integral transform (**PIT**) at x .

It can be easily shown that the distribution of $U := F(X)$ is standard uniform i.e. $\mathcal{U}(0, 1)$ as

$$P(U \leq u) = P(F(x) \leq u) = P(X \leq F^{-1}(u)) = F(F^{-1}(u)) = u \quad \forall u \in [0, 1]$$

by the definition of the distribution function. The PIT will be used to transform the realizations that follow an in practice estimated distribution to the uniform scale. Of course for estimated or empirical distributions the resulting transformed values are only approximate uniform. This uniform scale will be called **copula-scale** from now on and be covered more thoroughly in the subsequent Sections. Notably also the reverse transformation from now on called the **inverse PIT** can be used to (approximately) transform realizations from the copula to the original scale. Hereby the possibly estimated quantile function $Q = F^{-1}$ is utilized.

2 Univariate financial time series

As evident from the introduction chapter at the core of all covered risk assessment approaches and algorithms in this thesis is a portfolio of interest. So it is reasonable to keep a consistent notation of its components. The portfolio of interest Ω consists of $d \in \mathbb{N}$ assets denoted by A_1, \dots, A_d . Additionally the weights w_1, \dots, w_d indicate the volume of each asset in the portfolio. For each of the portfolios assets its value or price $p_t^{A_j}$ over a discrete time frame of interest $t \in \{1, \dots, T\}$ is observed. In most financial studies instead of the prices one works with the log returns of the assets that are introduced next following Tsay (2010).

Definition 2.0.1 (Log returns). *For a fixed asset with prices p_t and p_{t-1} the log return or also called continuously compounded return at time t is defined as*

$$r_t = \ln \frac{p_t}{p_{t-1}} \quad (2.1)$$

This formula results from the infinite compounding frequency limit case of discrete compounding. For more details on discrete compounding or the incorporation of dividend payments consult Chapter 1 of Tsay (2010).

The statistical properties of the derived log returns are more tractable than the ones of simple returns $\left(\frac{p_t - p_{t-1}}{p_{t-1}} \right)$ and a nice additional property is that one can easily sum up log returns over a time interval to get valid multi-period log returns. But it has to be noted that the weighted sum of the log returns of a portfolio at time t is only an approximation of the full portfolio log return.

$$r_t^\Omega \approx \sum_{j=1}^d w_j * r_t^{A_j} \quad (2.2)$$

This widely used approximation will also be used in this thesis. So all in all the relevant portfolio information is given by the set

$$\Omega = \{w_j, r_t^{A_j} | t = 1, \dots, T; j = 1, \dots, d\}.$$

This set encompasses one univariate financial time series of log returns for each asset. Such a time series will be treated as a collection of RVs over time. An important prerequisite for many univariate time series models like the ones that will be introduced in Chapter 4 and thus now defined consulting Shumway and Stoffer (2017) is the stationarity of a time series .

Definition 2.0.2 (Stationary time series). *A time series is **strictly stationary** if the probabilistic behavior of an arbitrary collection of time points $\{X_{t_1}, \dots, X_{t_k}\}$ does not change for any possible time shift like $\{X_{h+t_1}, \dots, X_{h+t_k}\}$. So all joint distribution functions stay the same independent of arbitrary time shifts. A time series is **weakly stationary** if the process has finite variance, the mean value function $\mu_t = \mathbb{E}[X_t]$ is constant over time and the autocovariance function $\gamma(s, t) = Cov(X_s, X_t)$ depends only on the lagged time difference $|s - t|$. In particular for $s = t$ this means a constant variance i.e. $\gamma(s, s) = Var(X_s) = \sigma^2$.*

For the upcoming covered theory the notion of weak stationarity of the log return series is sufficient and without additional comments the term stationary will refer to weak stationarity. To check the stationarity of a given log return series one can use a visualization as in Figure 2.1. One should spot a constant variation around a constant value. Smoothing lines e.g. by loess smoothing can be a beneficial add on for the visualization. They help to assess the constant mean assumption and might uncover periodicity.

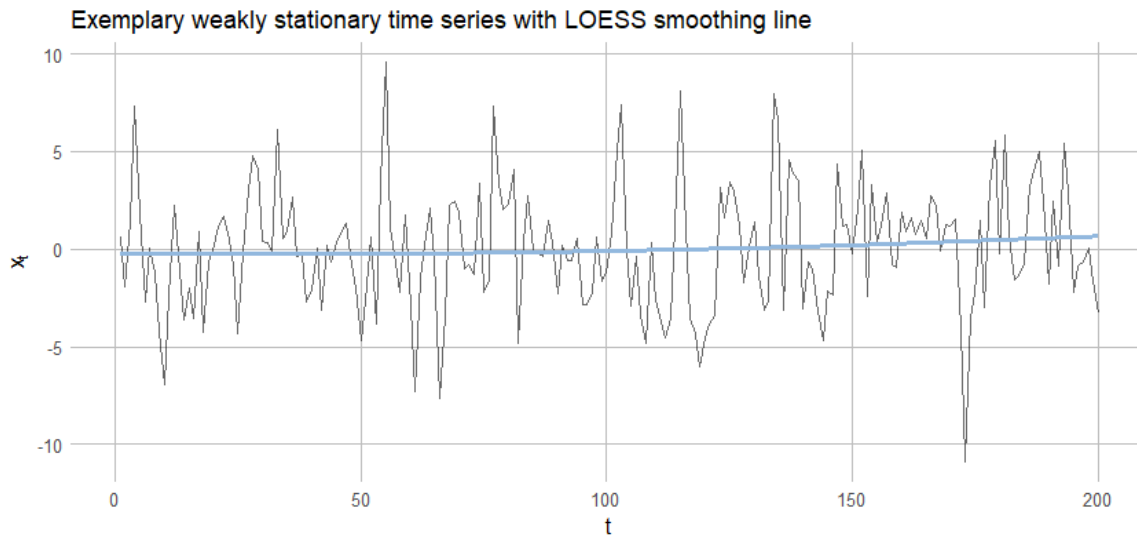


Figure 2.1 Illustration of a weakly stationary zero mean time series. The blue LOESS smoothing line helps to assess the constant mean assumption.

Of course one could also split the data into subsamples or bootstrap samples and then compare their empirical first two moments. For obviously non stationary time series one can have a look into detrending with regression methods or also differencing as outlined in Shumway and Stoffer (2017). When observing variance heterogeneity for example in the form of a cluster with considerably higher variance, which might be observable during a crash of the financial market, one will have to keep this in mind when modeling. The upcoming Section will indeed also present a class of models that can deal with these heterogeneities and thus time series that are not weakly stationary, but the mean function should still be constant over time. Before covering the univariate modeling of such time series of log returns in Chapter 4 it is reasonable with regard to the overall goal of risk assessment through backtesting to introduce the risk measures that will be used. Then it will also be even more obvious why the modeling of the log return series is needed. Thus the next section covers different measures of risk.

3 Risk measures

Before defining the actual risk measures and their estimators that will be used during this thesis a general definition of a risk measure will be given. This chapter follows Tasche (2002) and Pfaff (2016).

3.1 Definition and desired properties

Definition 3.1.1 (Risk measure). *Given a probability space with sample space \mathcal{O} , σ -Algebra \mathcal{F} and probability measure \mathcal{P} as well as a non-empty set of \mathcal{F} -measurable absolutely continuous RVs \mathcal{R} , then any mapping $\rho : \mathcal{R} \rightarrow \mathbb{R} \cup \infty$ is considered a valid risk measure.*

Here elements of \mathcal{R} might be seen as random variables at a certain time point for the value or log return of an asset. The mapping ρ then should return a reasonable real value indicating the risk of the asset at the specific point in time. Definition 3.1.1 will be met by every presented risk measure below. For quite some time now there is a strong case for using so called coherent risk measures which come with reasonable additional properties that a risk measure should satisfy and thus will be introduced next.

Definition 3.1.2 (Coherent risk measure). *A risk measure ρ as defined in Definition 3.1.1 is called coherent if it satisfies the following properties:*

1. *Monotonicity: $X, Y \in \mathcal{R}, X \leq Y \Rightarrow \rho(X) \geq \rho(Y)$.*
2. *Positive homogeneity: $X \in \mathcal{R}, h > 0, hX \in \mathcal{R} \Rightarrow \rho(hX) = h\rho(X)$.*
3. *Translation invariance: $X \in \mathcal{R}, a \in \mathbb{R}, a + X \in \mathcal{R} \Rightarrow \rho(a + X) = \rho(X) - a$*
4. *Subadditivity: $X, Y \in \mathcal{R}, X + Y \in \mathcal{R} \Rightarrow \rho(X + Y) \leq \rho(X) + \rho(Y)$*

While the consequences of the first three properties are quite obvious, the subadditivity has major implications on portfolio management that are not directly visible. Say one manages multiple portfolios and wants to estimate overall risk over all portfolios so the sum of all portfolios. Then when having risk estimates from a coherent risk measure for each portfolio one can get an upper bound on the overall risk by a simple summation. This cannot be achieved in general for risk measures lacking this fourth property and this could then for example lead to a discouragement regarding diversification. Having these general definitions covered one can move on and have a look at specific risk measures.

3.2 Value at Risk (VaR)

The most common risk measure in practice is still the Value at Risk. The measure was initially presented in the first version of the publication RiskMetrics by JP Morgan as pointed out by Pfaff (2016).

Definition 3.2.1 (Value at risk). *For a given asset or portfolio P , the continuous random variable R_t^P corresponding to the value of P at time t and with continuous distribution $F_{R_t^P}$ and quantile function $Q_{R_t^P}$ the value at risk (**VaR**) with confidence level $\alpha \in (0, 1)$ is defined as*

$$\mathbf{VaR}_\alpha^{P,t} := \sup\{r | F_{R_t^P}(r) \leq \alpha\} = Q_{R_t^P}(\alpha) \quad (3.1)$$

and can be interpreted as the value that R_t^P will only fall below with a probability of α . An illustration is given in Figure 3.1.

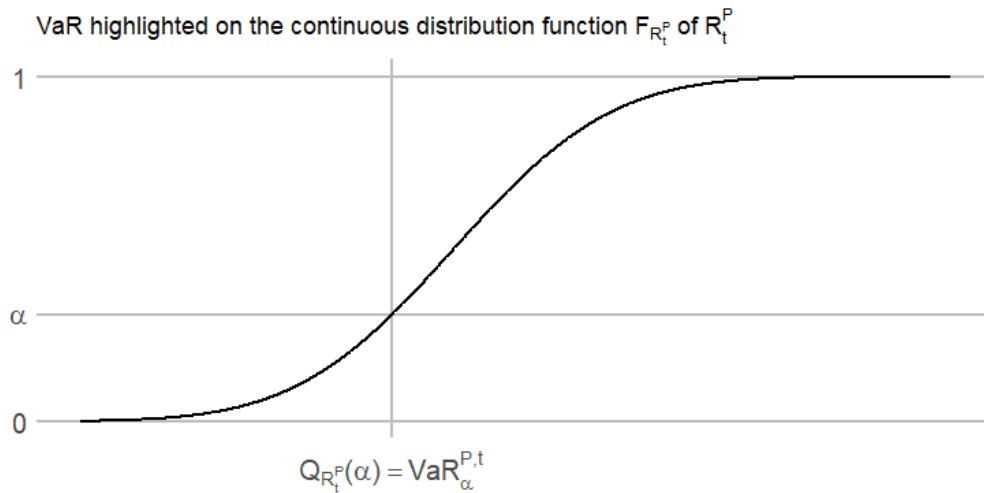


Figure 3.1 Illustration of the Value at Risk for the continuous random variable R_t^P using its distribution function $F_{R_t^P}$.

Here for the estimation of the VaR a Monte Carlo approach is utilized. The estimate for the $VaR_{\alpha}^{P,t}$ given samples $r_t^{P,s}$, for $s \in \{1, \dots, S\}$ from the random variable R_t^P is thus given straightforward by the empirical α level quantile of these samples as can be seen below in Equation (3.2).

$$\widehat{VaR}_{\alpha}^{P,t} = \widehat{VaR}(\{r_t^{P,s} | s \in \{1, \dots, S\}\}, \alpha) = \hat{Q}_{\{r_t^{P,s} | s \in \{1, \dots, S\}\}}(\alpha) \quad (3.2)$$

Here $\hat{Q}_{Set}(\alpha)$ denotes the standard empirical quantile function based on a *Set* of samples evaluated at the confidence level α .

Remark 3.2.1. *Shortcomings of the VaR*

- *The Value at Risk is indeed a valid risk measure as defined in Definition 3.1.1 but not a coherent one as it lacks the discussed subadditivity property as proven in Artzner et al. (1999) and Acerbi and Tasche (2002). This is a major weakness as this makes it hard to assess overall risk when managing multiple portfolios.*
- *The VaR is also law invariant which means that the distributions of the random variables of interest must not be identical in order to have the same VaR. This was in particular shown for heavy and light tailed distributions in Embrechts et al. (1997).*
- *The risk measure contains no information about the value of the asset or portfolio if its value falls below the VaR at the chosen confidence level. As for such a scenario the VaR is the upper bound for the portfolio value so for example corresponding to the best possible portfolio value in case a once in a hundred year crash happens. The more interesting question could be if such a crash happens what portfolio value may be expected. The answer to this question is actually estimated by the expected shortfall which will be discussed right below.*

3.3 Expected shortfall (ES)

Due to the mentioned shortcomings of the Value at Risk the banking regulators shifted their focus from this prominent risk measure with the third Basel Accords to the risk measure expected shortfall that will be introduced next Bayer and Dimitriadis (2020b).

Definition 3.3.1 (Expected shortfall). For a given asset or portfolio P , the continuous random variable R_t^P corresponding to the value of P at time t and with continuous distribution $F_{R_t^P}$ and quantile function $Q_{R_t^P}$ the **expected shortfall (ES)** with confidence level $\alpha \in (0, 1)$ is defined as

$$ES_\alpha^{P,t} = \mathbb{E}[R_t^P | R_t^P \leq \mathbf{VaR}_\alpha^{P,t}] = \frac{1}{\alpha} \int_0^\alpha \mathbf{VaR}_u^{P,t} du \quad (3.3)$$

and can be interpreted as the average value of R_t^P if the asset or portfolio drops in value below the $\mathbf{VaR}_\alpha^{P,t}$. The second characterization lends itself for the interpretation as the average value at risk for the lower confidence span of $(0, \alpha)$.

The two characterizations in Equation (3.3) can be connected as stated on page 5 of Bayer and Dimitriadis (2020b). One can do this by integration by substitution and by using the fact that the RV is continuous.

$$\begin{aligned} \mathbb{E}[R_t^P | R_t^P \leq \mathbf{VaR}_\alpha^{P,t}] &= \frac{1}{\alpha} \int_{-\infty}^{F_{R_t^P}^{-1}(\alpha)} r f_{R_t^P}(r) dr \\ &= \frac{1}{\alpha} \int_{-\infty}^{F_{R_t^P}^{-1}(\alpha)} F_{R_t^P}^{-1}(F_{R_t^P}(r)) F'_{R_t^P}(r) dr \\ &= \frac{1}{\alpha} \int_0^\alpha F_{R_t^P}^{-1}(r) dr \\ &= \frac{1}{\alpha} \int_0^\alpha \mathbf{VaR}_u^{P,t} du \end{aligned}$$

Contrary to the VaR the expected shortfall is indeed a coherent risk measure which was proven for example in Tasche (2002) and thus has the desired subadditivity property. So indeed the ES can mitigate the major shortcomings of the VaR but for example the law invariance weakness still applies to the ES. Moreover it should be noted that in the literature the expected shortfall is often called the **conditional Value at Risk (CVaR)** and the representation using the expected value for general also non absolutely continuous random variables is often called the **tail conditional expectation (TCE)** which is equivalent to the ES in the continuous case. The following decomposition shows that the ES is always smaller or equal to the VaR so actually more conservative as a risk measure. This comes from the non-negativity of the last expectation.

$$\begin{aligned} ES_\alpha^{P,t} &= \mathbb{E}[R_t^P | R_t^P \leq \mathbf{VaR}_\alpha^{P,t}] \\ &= \mathbb{E}[R_t^P + \mathbf{VaR}_\alpha^{P,t} - \mathbf{VaR}_\alpha^{P,t} | R_t^P \leq \mathbf{VaR}_\alpha^{P,t}] \\ &= \mathbb{E}[\mathbf{VaR}_\alpha^{P,t} | R_t^P \leq \mathbf{VaR}_\alpha^{P,t}] + \mathbb{E}[R_t^P - \mathbf{VaR}_\alpha^{P,t} | R_t^P \leq \mathbf{VaR}_\alpha^{P,t}] \\ &= \mathbf{VaR}_\alpha^{P,t} - \mathbb{E}[\mathbf{VaR}_\alpha^{P,t} - R_t^P | R_t^P \leq \mathbf{VaR}_\alpha^{P,t}] \end{aligned} \quad (3.4)$$

Again for the estimation of the expected shortfall a Monte Carlo approach is taken. The base setting is the same as for the VaR estimation, samples $r_t^{P,s}$, for $s \in \{1, \dots, S\}$ from the random variable R_t^P are given, but here the $ES_\alpha^{P,t}$ should be estimated. Other than for the VaR here three different estimators are presented below.

$$\begin{aligned}\widehat{ES}_{mean,\alpha}^{P,t} &= \widehat{ES}_{mean}(\{r_t^{P,s} | s \in \{1, \dots, S\}\}, \alpha) \\ &= \frac{1}{|\{r_t^{P,s} | s \in \{1, \dots, S\}, r_t^{P,s} \leq \mathbf{VaR}_\alpha^{P,t}\}|} \sum_{\{r_t^{P,s} | s \in \{1, \dots, S\}, r_t^{P,s} \leq \mathbf{VaR}_\alpha^{P,t}\}} r_t^{P,s}\end{aligned}\quad (3.5)$$

$$\begin{aligned}\widehat{ES}_{median,\alpha}^{P,t} &= \widehat{ES}_{median}(\{r_t^{P,s} | s \in \{1, \dots, S\}\}, \alpha) \\ &= \hat{Q}_{\{r_t^{P,s} | s \in \{1, \dots, S\}, r_t^{P,s} \leq \mathbf{VaR}_\alpha^{P,t}\}}(0.5)\end{aligned}\quad (3.6)$$

$$\begin{aligned}\widehat{ES}_{mc,n_mc,\alpha}^{P,t} &= \widehat{ES}_{mc}(\{r_t^{P,s} | s \in \{1, \dots, S\}\}, \alpha, n_mc) \\ &= \frac{1}{n_mc} \sum_{i=1}^{n_mc} \mathbf{VaR}_{u_i}^{P,t} \quad u_i \sim \mathcal{U}(0, \alpha) \text{ for } i \in \{1, \dots, n_mc\}\end{aligned}\quad (3.7)$$

The first two estimates are derived by just applying standard statistics for the expected value i.e. the empirical mean in (3.5) and median in (3.6) on the returns for which the condition that its value falls below the a priori computed VaR at the same confidence level α holds. Here the median based estimate might be more robust. These two estimating strategies are outlined in Pfaff (2016) and are the natural choice given the first representation of the ES in (3.3). The last estimate (3.7) is a Monte Carlo integration approach based on the second characterization of the ES in (3.3). Note that for this approach an additional parameter that specifies the number of Monte Carlo samples must be introduced. A comparison of their respective performances in practice is covered in the applied part of the thesis.

Having covered these two important risk measures as well as their estimators it is evident that for the risk measure estimation samples of the RVs of interest are needed. These RVs of interest are the future portfolio log returns. A crucial step in the process of computing these samples that will be covered in detail in Chapter 7 is the modelling and forecasting of each of the portfolios asset log return series. The models used for this will be called assetwise marginal time series models and be introduced in the next section.

4 Univariate financial time series models

The **assetwise marginal time series models** or short **marginal models** are called this way because for each assets log return series of the portfolio Ω an individual univariate time series model will be fitted in a rolling window fashion which will be described in detail in Chapter 7 and thus each asset can be viewed as a margin of the whole portfolio. So in this section this marginal modeling of one of the assets log return series $\{r_t^{A_j}, t = 1, \dots, T\}$ or in this section for simplicity $\{r_t, t = 1, \dots, T\}$ is discussed. Moreover starting this chapter the return series are assumed to be stationary according to Definition 2.0.2. The approach is the same for all log return series of the portfolio. The considered model class for the time series will be a so called ARMA-GARCH model. In this chapter first the components of the ARMA-GARCH models i.e. ARMA and GARCH models are introduced before stating the final model. Besides the model formulas also notes on parameter estimation, forecasting and useful properties are given. The chapter that follows mostly Shumway and Stoffer (2017) and Tsay (2010) concludes with a section on ways to assess model quality.

4.1 Autoregressive models (AR)

Autoregressive (AR) models of order p translate mathematically the intuitive idea that the current value r_t of the time series can be explained as a function of its last p values r_{t-1}, \dots, r_{t-p} . In the classic AR model this function is a linear one which results in the following formal definition.

Definition 4.1.1 (Autoregressive model). *An autoregressive model of order p for the time series $\{r_t, t = 1, \dots, T\}$ has the model formula*

$$r_t = \sum_{l=1}^p \phi_l r_{t-l} + w_t \quad (4.1)$$

with $w_t \stackrel{i.i.d.}{\sim} wn(0, \sigma^2)$ often called the **innovations** and wn called **white noise** being a zero mean, constant variance distribution. When the white noise distribution is a normal distribution one speaks of **Gaussian white noise**. The ϕ_l are the model parameters to be estimated. The model is abbreviated by **AR**(p).

From the definition it is evident that one models the time conditional mean of the time series while assuming constant variance over time. Having the coefficients ϕ_l one can define the so called AR polynomial

$$\phi(z) = 1 - \sum_{l=1}^p \phi_l z^l$$

with $z \in \mathbb{C}$ and then so called **causal** AR processes. A process is called causal if it does not depend on future observations like r_{t+1} at time t . An AR model is causal if and only if the roots of $\phi(z)$ lie outside the complex unit circle as outlined in Shumway and Stoffer (2017). In the following the considered AR models are assumed to be causal as future values of the time series are considered unknown.

For a better grasp of autoregressive models one has an exemplary look at the first order causal AR(1) model with $|\phi_1| = |\phi| < 1$ i.e. $r_t = \phi r_{t-1} + w_t$. The first two moments conditioned on the past observations are directly given by

$$\begin{aligned}\mathbb{E}[r_t | r_{t-1}, r_{t-2}, \dots] &= \phi r_{t-1} \\ \text{Var}(r_t | r_{t-1}, r_{t-2}, \dots) &= \text{Var}(w_t) = \sigma^2.\end{aligned}$$

Iterating backwards in time one can get the stationary linear process representation of the time series

$$\begin{aligned}r_t &= \phi r_{t-1} + w_t \\ &= \phi^2 r_{t-2} + \phi w_{t-1} + w_t \\ &\vdots \\ &= \sum_{l=0}^{\infty} \phi^l w_{t-l}.\end{aligned}\tag{4.2}$$

This is now useful for the calculation of the unconditional moments of the time series.

$$\begin{aligned}\mathbb{E}[r_t] &= \sum_{l=0}^{\infty} \phi^l \mathbb{E}[w_{t-l}] = 0 \\ \gamma(h) &= \text{Cov}(r_{t+h}, r_t) \\ &= \mathbb{E}\left[\left(\sum_{l=0}^{\infty} \phi^l w_{t+h-l}\right)\left(\sum_{k=0}^{\infty} \phi^k w_{t-k}\right)\right] \\ &= \sigma^2 \sum_{l=0}^{\infty} \phi^l \phi^{l+h} \\ &= \sigma^2 \frac{\phi^h}{1 - \phi^2}, \quad h \geq 0.\end{aligned}$$

Here the fact that $|\phi| < 1$ is used which leads to the convergence of the sum as a geometric series in the last line of the derivation of the autocovariance function. The zero mean assumption can be and is often relaxed by introducing a mean parameter ϕ_0 additive in the model formulation given in (4.1). The general behavior of AR models is not influenced by this and in the AR(1) model this leads to the mean $\phi_0/(1 - \phi_1)$. Also it is evident that a greater ϕ corresponds to a stronger autocovariance and variance structure following from the above given autocovariance function γ . This can be observed in Figure 4.1 that displays sample paths and their corresponding empirical autocovariance function $\hat{\gamma}$ of AR(1) models with $\phi \in \{0.1, 0.9, -0.9\}$. One can also observe that a negative ϕ leads contrary to the positive one not to a quite smooth sample path but to a very rough one. The alternating pattern in the empirical autocovariance functions is also no surprise when looking at the equation for the theoretical one above.

4.2 Moving average models (MA)

Other than the AR models **moving average (MA) models** of order q assume that the current value r_t can be explained via a function of the last q and the current noise terms w_t, \dots, w_{t-q} . Again in the classic MA models this function is assumed to be linear which results in the following formal definition.

Definition 4.2.1 (Moving average model). *A moving average model of order q for the time series $\{r_t, t = 1, \dots, T\}$ has the model formula*

$$r_t = \sum_{l=1}^q \theta_l w_{t-l} + w_t\tag{4.3}$$

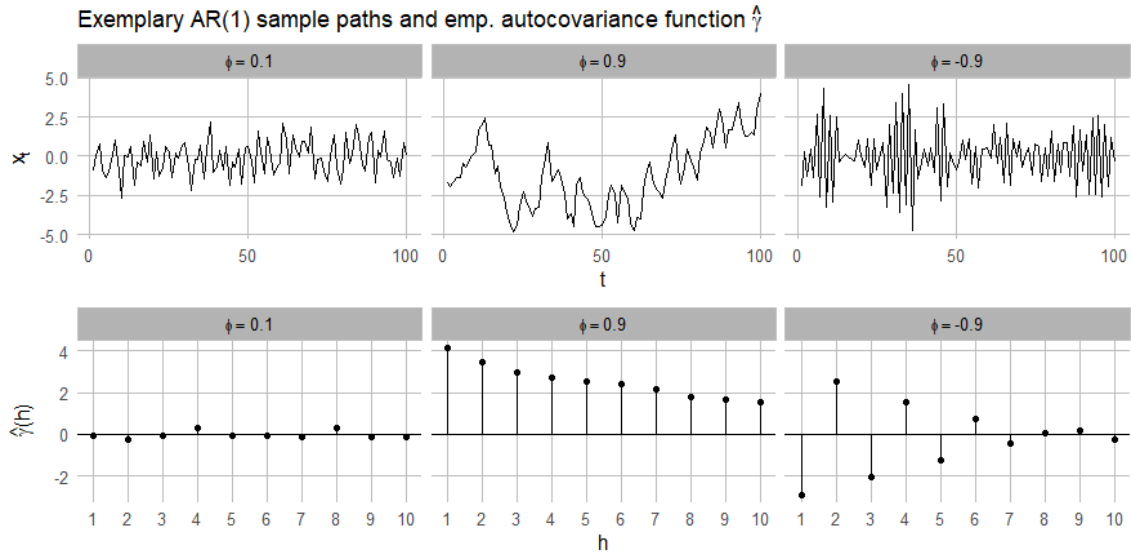


Figure 4.1 Illustration of the sample paths and the corresponding empirical autocovariance function $\hat{\gamma}$ of three AR(1) models differing in their parameter ϕ .

with $w_t \stackrel{i.i.d.}{\sim} wn(0, \sigma^2)$ again denoting a white noise distribution. The θ_l are the parameters. The model is abbreviated by $MA(q)$.

Again the conditional mean is modeled under the assumption of constant variance. Having the coefficients θ_l one can define the so called MA polynomial

$$\theta(z) = 1 + \sum_{l=1}^q \theta_l z^l$$

with $z \in \mathbb{C}$ and then the so called **invertible** MA processes. This is important as certain combinations of parameters and white noise specifications lead to the same and not distinguishable autocovariance structure. From these models always the model with an infinite AR representation is chosen. This model is then called the invertible process. A MA model is invertible if and only if the roots of $\theta(z)$ lie outside the complex unit circle as shown in Shumway and Stoffer (2017). In the following the considered MA models are assumed to be invertible.

Like for the AR models an exemplary look at the first order invertible MA(1) model with $|\theta_1| = |\theta| < 1$ i.e. $r_t = \theta w_{t-1} + w_t$ is taken. The first two unconditional moments conditioned on the observed past are trivially given from the model formulation. Also the unconditional moments are easily derivable.

$$\begin{aligned} \mathbb{E}[r_t] &= \theta \mathbb{E}[w_{t-1}] + \mathbb{E}[w_t] = 0 \\ \gamma(h) &= Cov(r_{t+h}, r_t) \\ &= \begin{cases} (1 + \theta^2)\sigma^2 & h = 0 \\ \theta\sigma^2 & h = 1 \\ 0 & h > 1 \end{cases} \end{aligned}$$

Notably the autocovariance function cuts off after the lag h corresponding to the order q . This will be useful and covered again when addressing model selection for ARMA models. So r_t in a MA model other than in an AR model for which r_t is correlated with all its predecessors is actually just correlated with its q predecessors. In the MA(1) case with its direct predecessor. Again the sample path of MA(1) models with positive θ is slightly smoother than their negative counterpart as depicted in Figure 4.2. Also the

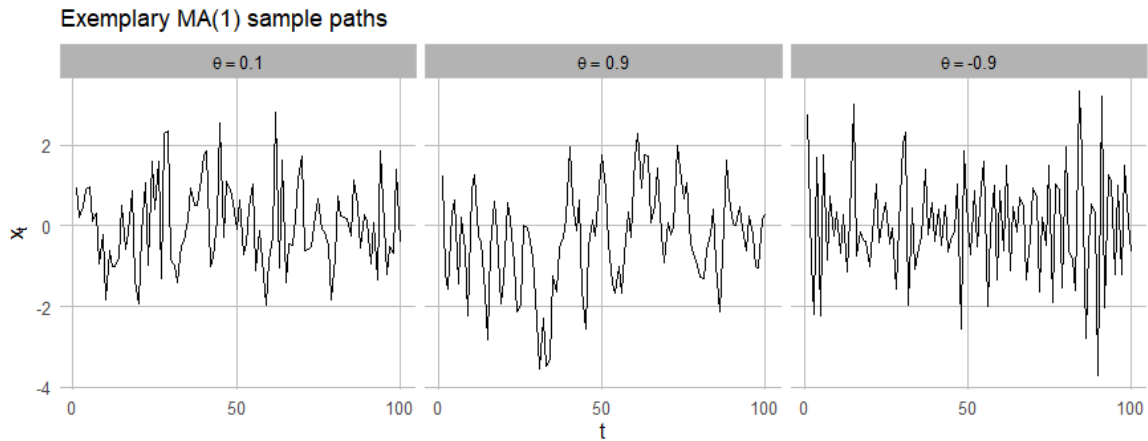


Figure 4.2 Illustration of the sample paths of three MA(1) models with differing parameters θ .

reduced variation for smaller θ can be spotted. As in the AR case one can and often does relax the zero mean assumption and introduces a mean parameter θ_0 additive in the model formulation given by Equation (4.3). This parameter is by linearity of the expectation directly the expected value of the process and its introduction does not influence the general properties of the MA model.

4.3 Autoregressive moving average models (ARMA)

Model definition

Having introduced AR and MA models separately one can combine the two models into one more flexible model the so called **autoregressive moving average models** short **ARMA model**.

Definition 4.3.1 (Autoregressive moving average models). *A autoregressive moving average model of AR order p and MA order q for the time series $\{r_t, z = 1, \dots, T\}$ has the model formula*

$$r_t = \phi_0 + \sum_{l=1}^p \phi_l r_{t-l} + w_t + \sum_{k=1}^q \theta_k w_{t-k} \quad (4.4)$$

with $w_t \stackrel{i.i.d.}{\sim} wn(0, \sigma^2)$ once again denoting a white noise distribution. The parameters are here θ_k, ϕ_l and note that here the mean parameter ϕ_0 is considered. The actual expected value of the time series is then $\phi_0 / (1 - \phi_1 - \dots - \phi_p)$. The model is abbreviated by **ARMA**(p, q).

In general especially for higher orders of p and q one should be aware of the problem of parameter redundancy. For the detection one should spot common roots in the respective AR and MA polynomials as parameter redundancy can result in wrong interpretations. An ARMA(p, q) model is considered causal if the respective AR(p) model is causal. The same applies for the invertability of the ARMA(p, q) model w.r.t. the invertability of the corresponding MA(q) model as stated in Shumway and Stoffer (2017).

AR and MA order selection

In order to fit such an ARMA model one has to determine the orders p and q . This because just fitting a large order model will lead to less efficient and less precise parameter estimates as outlined in Shumway and Stoffer (2017). Useful tools for choosing reasonable orders are the autocorrelation function (ACF) and the partial autocorrelation function (PACF). Of course in practice their efficient empirical estimators are used. The **autocorrelation function ACF** is basically a standardized version of the autocovariance function and measures the linear influence of r_s on r_t i.e. the serial correlation. The formula is given by

$$ACF(s, t) = \frac{\gamma(s, t)}{\sqrt{\gamma(s, s)\gamma(t, t)}} \quad (4.5)$$

and simplifies for stationary time series to

$$ACF(h) = \frac{\gamma(h)}{\gamma(0)} \quad (4.6)$$

So it holds that $ACF \in [-1, 1]$ with -1 and 1 indicating a perfect linear relationship as in the classical Pearson correlation and 0 indicating no linear dependence. Already indicated by the first order AR and MA models discussed the ACF of an AR model does tail off (mostly exponentially or in a sinusoidal fashion) and the ACF of an MA(q) model cuts off after $h = q$. Thus using ACF plots one can determine the MA order of an MA or ARMA model. For the choice of the AR order of a stationary process one uses the **partial autocorrelation function PACF**. For $h = 1$ one has $PACF(1) = ACF(1)$ and for subsequent $h > 1$ the partial autocorrelation is given by

$$PACF(h) = Cor(r_{t+h} - \hat{r}_{t+h}, r_t - \hat{r}_t) \quad (4.7)$$

were \hat{r}_t is linearly regressed on $r_{t+1}, \dots, r_{t+h-1}$ i.e. the fitted value from the regression

$$\hat{r}_t = \beta_1 r_{t+1} + \beta_2 r_{t+2} + \dots + \beta_{h-1} r_{t+h-1}$$

and were \hat{r}_{t+h} is linearly regressed on $r_{t+1}, \dots, r_{t+h-1}$ i.e. the fitted value from the regression

$$\hat{r}_{t+h} = \beta_1 r_{t+h-1} + \beta_2 r_{t+h-2} + \dots + \beta_{h-1} r_{t+1}.$$

So the PACF is basically the correlation of r_t and r_{t+h} with the linear dependence on the values between them removed on each of them.

One can follow Shumway et. al Shumway and Stoffer (2017) that show that in this case the PACF of an AR(p) model cuts off after $h = p$ and the MA model tails off. Thus this can be a valuable tool to determine the AR order p . Nevertheless for a general ARMA(p, q) model both ACF and PACF might tail off which can make the order selection difficult. A reasonable approach is then to start with low orders, assess and compare model quality with the tools that will be introduced at the end of this section and gradually increase the orders until the desired model quality is achieved or one detects the deficiency of the ARMA model to represent the data in a good enough way.

Estimation

For the estimation of AR(p), MA(q) and ARMA(p, q) models one assumes the order parameters p and q to be known. Although for AR models one has an optimal method of moments estimator through the so called Yule-Walker equations that are based on using empirical autocovariance functions of the observed process already for MA(1) models the method of moments estimation is inferior to the most prominent estimation method in the field which is **maximum likelihood estimation (MLE)**. Before starting the estimation one also has to specify a parametric distribution for the assumed white noise distribution. Common choices are the normal or Student's t distribution. As the white noise distribution is assumed to have a zero mean only the volatility parameter σ^2 has to be estimated alongside the θ and ϕ parameters. The estimation is then carried out by numerically e.g. by Newton-Raphson maximizing the overall likelihood of the ARMA(p, q) model

$$L(\phi_0, \phi_1, \dots, \phi_p, \theta_1, \dots, \theta_q, \sigma^2) = \prod_{t=1}^T f(r_t | r_{t-1}, \dots, r_1) \quad (4.8)$$

with respect to the unknown parameters and with f being the density of the specified white noise distribution. If the conditional distribution is Gaussian one has the mean r_t^{t-1} and variance P_t^{t-1} of the conditional distributions in Equation (4.8) where they represent the one step ahead forecast

$$r_t^{t-1} = \mathbb{E}[r_t | r_{t-1}, \dots, r_1]$$

and the prediction error variance of it

$$P_t^{t-1} = \mathbb{E}[(r_t - r_t^{t-1})^2]$$

respectively. They both are functions of the parameters to be estimated alone and can recursively be estimated with algorithms like the Innovations or Durbin-Levinson Algorithm as outlined in Shumway and Stoffer (2017) thus one can write down the overall likelihood as

$$L(\phi_0, \phi_1, \dots, \phi_p, \theta_1, \dots, \theta_q, \sigma^2) = (2\pi)^{-T/2} (P_1^0 P_2^1 \dots P_T^{T-1})^{-1/2} \exp\left(-\frac{1}{2} \sum_{t=1}^T \left[\frac{(r_t - r_t^{t-1})^2}{P_t^{t-1}}\right]\right).$$

The same approach is taken for the case where the conditional distributions in Equation (4.8) are from the Student's t distribution with degree of freedom ν . The overall likelihood is then given as

$$L(\phi_0, \phi_1, \dots, \phi_p, \theta_1, \dots, \theta_q, \sigma^2) = \left(\frac{\Gamma(\frac{\nu+1}{2})}{\Gamma(\frac{\nu}{2})\sqrt{\nu\pi}}\right)^T (P_1^0 P_2^1 \dots P_T^{T-1})^{-1/2} \prod_{t=1}^T \left\{1 + \left(\frac{r_t - r_t^{t-1}}{P_t^{t-1}}\right)^2 \frac{1}{\nu}\right\}^{-\frac{\nu+1}{2}}.$$

One speaks of **unconditional estimation** if one just maximizes without further restrictions the above likelihood with respect to all the unknown model parameters. The **conditional estimation** approach uses initial values obtained from the observed data to speed up the computations. Both the unconditional and conditional approach lead to optimal estimators for causal and invertible ARMA(p,q) models and as AR(p) and MA(q) models are just special ARMA models this applies also for their estimation. Moreover these estimation procedures lead to asymptotic results that allow the construction of confidence intervals and hypothesis test for the estimates. More details on the estimation of ARMA models as outlined above can be found in Shumway and Stoffer (2017).

Forecasting

Of course in this thesis one is very much interested in forecasting returns for a given fitted time series model. For the forecasting all model parameters are assumed to be estimated. The one step ahead forecast in the setting of causal and invertible ARMA(p,q) model is straightforward given by

$$\hat{r}_{T+1} = \hat{\phi}_0 + \sum_{l=1}^p \hat{\phi}_l r_{T+1-l} + \sum_{k=1}^q \hat{\theta}_k (r_{T+1-k} - \hat{r}_{T+1-k}) \quad (4.9)$$

Where the fitted residuals are used in the sum corresponding to the MA part. The one step ahead estimates \hat{r}_t are iteratively calculated via algorithms like the Innovations or the Durbin-Levinson algorithms as given in Shumway and Stoffer (2017). These algorithms assume as the starting value $\hat{r}_1 = \mathbb{E}[r_1] = \mu$ as a stationary time series with constant mean is assumed and additionally calculate the corresponding one step ahead estimated forecast errors. One can observe a gradual dampening of the influence of the AR coefficients and the MA coefficients disappearing according to their order q when looking at these iterative calculations. So the forecasts converge to the estimated mean of the series quite quickly at least if there is not a very strong autocovariance structure and the estimated variance grows towards the estimated variance of the whole process.

Until now all models covered had the assumption of constant variance in common. So the residuals are assumed to be homoscedastic. How to check this in practice will be covered in the last part of this section

that is concerned with quality assessment of time series models. In the practice of modelling log returns this assumption often does not hold and thus more flexible models that can account for heteroscedasticity are required. This leads to so called volatility models that will be introduced next.

4.4 Generalized autoregressive conditional heteroscedasticity models (GARCH)

Model definition

For volatility modeling using the **generalized autoregressive conditional heteroscedasticity (GARCH) model class** one splits the model formula into two components.

$$r_t = \mu_t + w_t \quad (4.10)$$

The so called **mean equation** μ_t is there to model the conditional mean as discussed for ARMA models. Actually often ARMA(p,q) models are used to parameterize the mean equation.

$$\mu_t = \phi_0 + \sum_{l=1}^p \phi_l r_{t-l} + \sum_{k=1}^q \theta_k w_{t-k} \quad (4.11)$$

with w_t being the innovation at time t . If the mean equation is parameterized this way one speaks of a so called **ARMA-GARCH** model. The parameterization of the innovations $w_t = r_t - \mu_t$ called the **volatility equation** is following a **GARCH(m, s)** model if it is given by

$$w_t = \sigma_t \epsilon_t, \quad \sigma_t^2 = \alpha_0 + \sum_{i=1}^m \alpha_i w_{t-i}^2 + \sum_{j=1}^s \beta_j \sigma_{t-j}^2 \quad (4.12)$$

So the innovations w_t are decomposed multiplicatively into a time dependent component σ_t and again a white noise component $\epsilon_t \stackrel{i.i.d.}{\sim} wn(0, 1)$. Additionally it must hold that $\alpha_i \geq 0$, $\beta_j \geq 0$ and $\sum_{i=1}^{max(m,s)} (\alpha_i + \beta_i) < 1$. The latter constraint implies a finite unconditional variance of the innovations w_t . For $s = 0$ the model reduces to an ARCH model. For details on this submodel class consult Chapter 3.4 of Tsay (2010).

Definition 4.4.1 (ARMA-GARCH model). *The model specified by combining Equations (4.10), (4.11) and (4.12) is called an ARMA(p,q)-GARCH(m,s) model. The orders p, q, m and s are fixed and the parameters ϕ_l , θ_k , α_i and β_j are to be estimated under discussed constraints.*

The parallels in the mean and volatility equations are quite evident so one can think of the GARCH model like an ARMA approach for the squared time dependent volatility component σ_t^2 of the innovations. Moreover one can conclude from the volatility equation that prior strong volatility gives rise to larger modelled volatility. This behavior is called **volatility clustering** and such periods of high volatility can be for example market crashes. Obviously this is a strength of the model. Additionally the model provides a functional insight describing the evolution of volatility. A weakness of the model however is that due to the square in the volatility equation it can not differentiate between positive and negative innovations. At this point it should also be noted that in most practical applications to financial time series data rather small order specifications often achieve the best results like the most common order specification ARMA(1,1)-GARCH(1,1).

Estimation

The estimation of ARMA-GARCH models is often performed using a **two pass estimation** method. First one focuses on the estimation of the parameters of the mean equation (4.11) and basically omits the volatility modeling. This results in the estimation of an ARMA model e.g. by MLE which was covered previously. After having fitted the mean equation one is left with a fitted residual series \hat{w}_t . The second

step is then to use the squared residual series \hat{w}_t^2 for the maximum likelihood estimation of the parameters of the volatility equation (4.12). As mentioned before the volatility equation is constructed like an ARMA model and thus the likelihood is given analogously to the ARMA case. This two pass method is of course just an approximation of the true parameters but empirical results especially on moderate or large sample sizes have shown very promising approximations. Alternatively one can also opt for a joint estimation of the parameters by considering the joint likelihood. This is computationally more demanding but the computational burden could be reduced by considering initial values from the two pass estimation method.

Forecasting

The forecasting of a fitted ARMA-GARCH model is quite similar to the one for ARMA models. The only difference is that one iteratively determines both the forecasts for the mean $\hat{\mu}_{T+1}$ and volatility $\hat{\sigma}_{T+1}^2$ like for ARMA models respectively as they both are time dependent. Exemplary for an ARMA(1,1)-GARCH(1,1) model the one step ahead volatility forecast is given by

$$\hat{\sigma}_{T+1}^2 = \hat{\alpha}_0 + \hat{\alpha}_1(r_T - \hat{r}_T)^2 + \hat{\beta}_1\hat{\sigma}_T^2$$

where like for the ARMA forecasting the estimates \hat{r}_t and $\hat{\sigma}_t^2$ are calculated iteratively with the initial values set to $\hat{r}_1 = \mu$ and $\hat{\sigma}_1^2 = 0$ as described in Tsay (2010) and Shumway and Stoffer (2017).

4.5 Assessment of model quality

The assessment of model quality for all the discussed time series models boils down to a thorough residual analysis. In order to work on a common scale one analyzes the **standardized residuals** which are calculated by

$$z_t = \frac{r_t - \hat{\mu}_t}{\hat{\sigma}_t} \quad (4.13)$$

where in the case of AR, MA and ARMA models the volatility estimate $\hat{\sigma}_t$ is of course not time dependent but constant. These standardized residuals $\{z_t | t = 1, \dots, T\}$ should for a well suited model as stated in the model definitions be i.i.d. samples of the specified white noise distribution. So the analysis of these standardized residuals w.r.t. no autocorrelation structure is the most important model quality assessment step for AR, MA and ARMA models. For ARMA-GARCH models this step is only the first of two as it only checks the adequacy of the mean equation. For the assessment of the volatility equation one has to additionally check for the nonexistence of autocorrelations in the squared standardized residuals $\{z_t^2 | t = 1, \dots, T\}$.

For the **residual analysis** one has multiple tools. The first and most straight forward one is to simply plot and visually inspect the residuals over time. This was done for four standardized residual series in Figure 4.3. For these examples one assumes to have fitted a hypothetical time series model with the standard Normal distribution as the specified target residual distribution. The exemplary standardized residual series all with total sample size 100 were simulated as follows.

1. I.i.d. normal: $z_t \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1)$.
2. I.i.d. Student's t: $z_t \stackrel{i.i.d.}{\sim} t_2(0, 1)$.
3. Correlated: z_t is a sample path from an AR(4) model with parameters $\phi_1 = 0.1, \phi_2 = \phi_3 = 0.3, \phi_4 = 0.2$ and zero mean.
4. Variance heterogeneity: $z_t \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 0.5)$ for $t \in \{1, \dots, 50\}$ and $z_t \stackrel{i.i.d.}{\sim} \mathcal{N}(0, 1 + (t - 51)/49)$ for $t \in \{51, \dots, 100\}$.




Figure 4.3 Plots of the standardized residuals z_t over time for four different sampled residual series all with sample size 100.

Thus for the assumed model only the first exemplary series is perfectly aligned with the model assumptions. In these basic exploratory plots already the autocorrelation of the correlated sample can be anticipated as well as the variance heterogeneity. Notably the i.i.d. normal and Student's t samples are not really distinguishable and both do not give rise to model quality concerns. The plots for the correlated sample suggest a poor modelling of the conditional mean which should be adjusted and the one showing severe variance heterogeneity suggests to either incorporate or to adjust the volatility equation.

Secondly as one does not want to spot any autocorrelation in the standardized residuals the visualization of the empirical $\widehat{ACF}(h)$ computed for the first lags h can be a valuable tool to check this assumption. Notably the empirical ACF based on T samples if the sample is white noise is under mild regularity assumptions and large T approximately normal distributed with zero mean and the standard deviation of \sqrt{T}^{-1} . This allows for the construction of 95% confidence bands around 0 that will be drawn in blue dashed lines in the following. These can assist in the process of detecting anomalous autocorrelation structure. Again the four exemplary samples are used for the empirical ACF plots in Figure 4.4. Here all plots besides the one of the correlated sample show no truly concerning autocovariance structure. For the correlated sample the findings are plain obvious and when observing such or slightly similar patterns one should reassess the conditional mean modeling.

Next one checks whether the distribution of the standardized residuals matches the specified one. This can be achieved by comparing the empirical sample quantiles with the theoretical ones. The approach and the corresponding plot is widely adapted and known as **Q-Q** plot. One plots the sample quantiles on the y axis and the corresponding theoretical ones on the x axis. So if the sample comes from the distribution one should observe a point cloud very close to the identity line. Notably this is only doable with a constant residual distribution. Thus whenever a volatility equation is included in the model this assessment method is not used. The Q-Q plots for the four exemplary residual series are given in Figure 4.5. For each plot the standard normal as the specified residual distribution is used for the theoretical quantiles. The slightly yet visible heavier tails for the Student's t distribution with degree of freedom $\nu = 2$ suggest the refitting with a heavier tailed residual distribution. In the variance heterogeneity case the heavy tails once again suggest a bad model fit.




Figure 4.4 Plots of the empirical ACF of the standardized residuals z_t for four different sampled residual series. The blue dashed lines indicate the 95% two sided confidence interval for the empirical ACF under the null hypothesis of white noise.

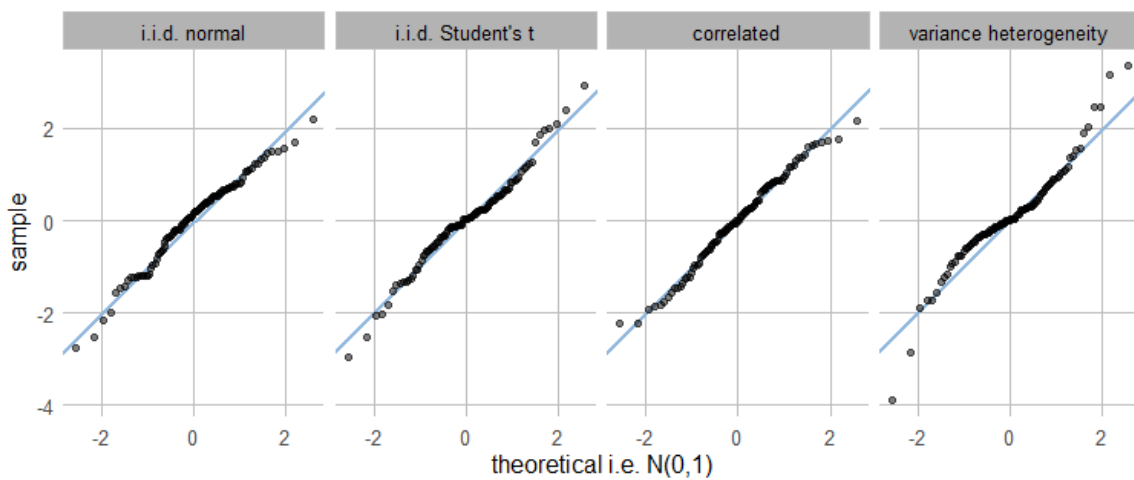


Figure 4.5 Normal Q-Q plots of the standardized residuals z_t for four different sampled residual series. The identity is given as blue line.

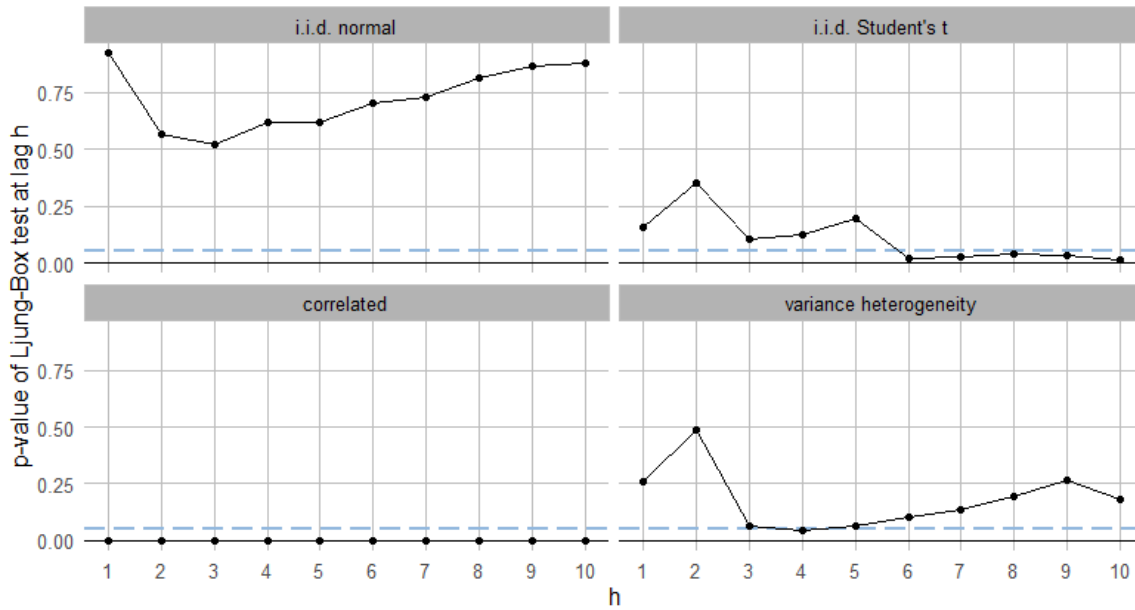


Figure 4.6 Plots of the p-values of Ljung-Box tests for different lags h based on the standardized residuals z_t for four different sampled residual series. The blue dashed line highlights the 0.05 confidence level.

In addition to plotting the empirical ACF one can consider a general hypothesis test for the existence of autocorrelation. The test that is called the **Ljung-Box** test is taking also the magnitude of the empirical ACF for multiple lags into consideration. For the residual series $\{z_t | t = 1, \dots, T\}$ the *Ljung-Box-Pierce Q-statistic* is given by

$$Q = T(T + 2) \sum_{h=1}^H \frac{\widehat{ACF}(h)^2}{T - h} \quad (4.14)$$

using the empirical ACF of the z_t . Under the null hypothesis of no autocorrelation up to lag H i.e. the $ACF(1) = \dots = ACF(H) = 0$ it holds that $Q \sim \chi_H^2$. Thus if Q would exceed the $(1 - \alpha)$ -quantile of the given χ^2 distribution one would reject the null hypothesis. In the use case of model quality assessment this will be a goodness of fit hypothesis test that indicates a good fit for high p-values i.e. one does not want to reject the null hypothesis. In practice often multiple values of H are tested. The resulting p-values can be displayed graphically like in Figure 4.6 for again the four exemplary residual series. On the left hand side one can spot once a picture perfect for the i.i.d. normal case and once the worst case scenario for the correlated sample. Both plots on the right side suggest not a perfect but sufficient behavior of the p-values. All in all one should use all the presented model assessment tools in order to perform a holistic residual analysis as every one of them has their strengths and weaknesses.

Until now only the model quality assessment given a single model was discussed. For the comparison of AR, MA, ARMA and also ARMA-GARCH models more or less the same principles apply. One should always opt for a parsimonious model w.r.t. the model parameters as already discussed above overfitting can result in inefficient estimates and wrong conclusions. In order to select both efficient which is mostly measured with the likelihood and sparse models the most common way to go is the usage of classical information criteria like the Akaike Information Criterion **AIC** of Akaike (1998) and the even more sparsity encouraging Bayesian Information Criterion **BIC** of Schwarz (1978).

Now given a univariate stationary log return series $\{r_t; t = 1, \dots, T\}$ one can construct powerful and most importantly well suited ARMA(p,q)-GARCH(m,s) models. After the successful assessment of model quality one can make one step or multi step forecasts that feature an estimate of the mean $\hat{\mu}$ and of the volatility $\hat{\sigma}^2$ for every forecasted time point. Moreover one has specified a known parametric distribution

for the residuals like the Student's t distribution. Using the PIT introduced in Definition 1.0.6 with this distribution one can transform the standardized residuals z of the model to the copula-scale. This copula data serves as the basis for the dependence modelling of multiple assets with vine copula models as described in Czado (2019). Moreover one can transform from the copula-scale to the log return scale by first applying the inverse PIT in order to transform to the standardized residual scale and secondly by adding the estimated mean $\hat{\mu}$ to the product of the estimated volatility $\hat{\sigma}^2$ and the standardized residual scale value that was calculated via the inverse PIT. These transformations to and from the copula scale using estimated and forecasted values based on a well suited ARMA-GARCH model will play a vital role in the risk estimation algorithms that are presented further down the road. Thus they will be discussed in more rigorous notion later. As already mentioned one does not restrict oneself to the univariate view of each of the portfolios assets but will apply multivariate dependence modeling with vine copulas to the whole portfolio. Thus the next chapter will start with a gentle introduction into the world of vine copulas before a closer look at their simulation is taken.

5 Vine copulas

This section will introduce the concept of **vine copulas** for high dimensional dependence modeling. For this we will closely follow Czado (2019).

5.1 Bivariate dependence measures

Although we are interested in high dimensional dependence modeling we first introduce **bivariate dependence measures**. These will then later be used in the modeling of higher dimensional dependencies. Arguably the most prominent bivariate dependence measure of all is the **Pearson product-moment correlation** more widely known as simply the Pearson correlation. It measures linear dependence between two random variables X_1 and X_2 and takes values in $[-1, 1]$.

Definition 5.1.1 (Pearson product-moment correlation). *For the random variables X_1 and X_2 with finite second moments one defines the **Pearson product-moment correlation coefficient** ρ as*

$$\rho := \rho(X_1, X_2) := Cor(X_1, X_2) = \frac{Cov(X_1, X_2)}{\sqrt{Var(X_1)}\sqrt{Var(X_2)}}. \quad (5.1)$$

It can be estimated using the empirical maximum likelihood estimators of the theoretical quantities.

Thus a value close to zero indicates that there is little evidence for a linear dependence, but even from a theoretical zero ρ no independence is implied. Values close the boundaries indicate a negative - close to -1 - and a positive - close to 1 - linear dependence respectively.

Remark 5.1.1. *Notably later one would like to access the bivariate linear dependence of two variables on the copula scale data i.e. the data u_{ij} lives in $[0, 1]$ with i denoting the observation number and j the variable number. As outlined in Czado (2019) for the estimation of the Pearson correlation of two variables U_1 and U_2 on the copula scale a prior transformation to the **normalized scale** through*

$$z_{ij} = \phi^{-1}(u_{ij}) \quad \forall i, j \quad (5.2)$$

is advisable.

But when building vine copulas later one is not only interested in the direct pairwise dependence. In particular one would also like to assess the dependence of variables with the effect of other variables already removed. A dependence measure that is constructed for this purpose is called the **partial correlation**.

Definition 5.1.2 (Partial correlation). *Given are the d random variables X_1, \dots, X_d that have zero mean and variances $\sigma_1^2, \dots, \sigma_d^2$ as well as the reduced index set $I_{-(i,j)}^d = \{1, \dots, d\} \setminus \{i, j\}$ with $i \neq j$. Then the **partial regression coefficients** $b_{i,j;I_{-(i,j)}^d}$ without loss of generality $i < j$ are defined as the quantities that minimize*

$$\mathbb{E}[(X_i - \sum_{j=2, j \neq i}^d a_{i,j;I_{-(i,j)}^d} X_j)^2] \quad (5.3)$$

*with respect to the coefficients $a_{i,j;I_{-(i,j)}^d}$. Then the **partial correlation** $\rho_{i,j;I_{-(i,j)}^d}$ is defined as*

$$\rho_{i,j;I_{-(i,j)}^d} = \text{sgn}(b_{i,j;I_{-(i,j)}^d}) \times \sqrt{b_{i,j;I_{-(i,j)}^d} \times b_{j,i;I_{-(i,j)}^d}}. \quad (5.4)$$

As shown by Yule and Kendall (1950) these can be calculated recursively using the following recursions.

Theorem 5.1.1 (Recursion for partial correlations). *The partial correlations as of Definition 5.1.2 satisfy the recursion*

$$\rho_{i,j;I_{-(i,j)}^d} = \frac{\rho_{i,j;I_{-(i,j)}^{d-1}} - \rho_{i,d;I_{-(i,j)}^{d-1}} \rho_{j,d;I_{-(i,j)}^{d-1}}}{\sqrt{1 - \rho_{i,d;I_{-(i,j)}^{d-1}}^2} \sqrt{1 - \rho_{j,d;I_{-(i,j)}^{d-1}}^2}}. \quad (5.5)$$

In practice however the most prominent estimation method as outlined in Kim (2015) is given by inverting the empirical covariance matrix of all variables to get the estimated pairwise partial correlations in the resulting matrix. The theory for this approach was given by Whittaker (1990). Notably for the case where the determinant of the covariance matrix is close to zero one uses the Moore-Penrose pseudo inverse for numerical stability.

5.2 Copula fundamentals

First up a short review of the most important data scales is given:

- **original scale:** The data is denoted with $x_{i,j}$ and lives on the whole real line.
- **copula scale:** The data is denoted with $u_{i,j}$ and lives on the interval $[0, 1]$.
- **normalized scale:** The data is denoted with $z_{i,j}$ and is computed via $\phi^{-1}(u_{i,j})$.

The transformation of i.i.d. data from the original to the copula scale is most often achieved with a so called **inference for margins (IFM)** approach where one estimates a univariate parametric marginal distribution function F and uses it to transform the data with

$$u_{i,j} = F(x_{i,j}). \quad (5.6)$$

To get back to the original scale one simply applies the inverse transformation. Besides the calculation of dependence measures like the partial correlation as given in Definition 5.1.2 the normalized scale is also often used for visualizations. This will be showcased in the next section.

Before introducing vine copulas one has to be familiar with the concept of a **copula** in general. A copula is basically a multivariate distribution function as given in Definition 1.0.1 but with the one constraint of all marginal distributions being the uniform distribution as given in Definition 1.0.2. This is formalized in the following definition.

Definition 5.2.1 (Copula). *A multivariate distribution function C as given in Definition 1.0.1 on the d -dimensional unit cube $[0, 1]^d$ is called a **copula** if and only if all marginal distributions are uniformly distributed. The corresponding copula density c is obtained by partial differentiation of the distribution function.*

$$c(u_1, \dots, u_d) = \frac{\partial^d}{\partial u_1 \dots \partial u_d} C(u_1, \dots, u_d) \quad \forall \mathbf{u} \in [0, 1]^d \quad (5.7)$$

A crucial result in the context of copulas and also integral to the construction of vine copulas is the **representation theorem of Sklar** which will now be given.

Theorem 5.2.1 (Sklar's theorem). *The joint distribution function F of a d -dimensional random vector \mathbf{X} can be expressed by the marginal distribution functions $F_i, i \in \{1, \dots, d\}$ and a d -dimensional copula C as*

$$F(x_1, \dots, x_d) = C(F_1(x_1), \dots, F_d(x_d)) \quad (5.8)$$

and in the same way one has the result for the corresponding densities.

$$f(x_1, \dots, x_d) = c(F_1(x_1), \dots, F_d(x_d)) f_1(x_1) \cdots f_d(x_d) \quad (5.9)$$

If the distribution F is absolutely continuous the copula C is unique. Moreover also the inverse statements hold i.e. the copula corresponding to a multivariate distribution can be written in terms of its marginal distributions as follows.

$$C(u_1, \dots, u_d) = F(F^{-1}(u_1), \dots, F^{-1}(u_d)) \quad (5.10)$$

Again the same can be stated for the respective densities.

$$c(u_1, \dots, u_d) = \frac{f(F^{-1}(u_1), \dots, F^{-1}(u_d))}{f_1(F_1^{-1}(u_1)) \cdots f_d(F_d^{-1}(u_d))} \quad (5.11)$$

The proof for this theorem can be found in Nelsen (2010). Sklar's theorem now allows for representations of conditional bivariate distribution in terms of their copula. As these will be important further later these representations are displayed in the following lemma.

Lemma 5.2.2 (Conditional distribution functions and densities of bivariate distributions in terms of their copula). *Using Theorem 5.2.1 one can rewrite the conditional bivariate quantities as*

$$F_{1|2}(x_1|x_2) = \frac{\partial}{\partial u_2} C_{12}(F_1(x_1), u_2) \Big|_{u_2=F_2(x_2)} \quad (5.12)$$

$$=: \frac{\partial}{\partial F_2(x_2)} C_{12}(F_1(x_1), F_2(x_2)) \quad (5.13)$$

$$f_{1|2}(x_1|x_2) = c_{12}(F_1(x_1), F_2(x_2)) f_1(x_1). \quad (5.14)$$

Proof. Using the definition of conditional densities and Theorem 5.2.1 one has

$$\begin{aligned} f_{1|2}(x_1|x_2) &= \frac{f_{12}(x_1, x_2)}{f_2(x_2)} \\ &= \frac{c_{12}(F_1(x_1), F_2(x_2)) f_1(x_1) f_2(x_2)}{f_2(x_2)} \\ &= c_{12}(F_1(x_1), F_2(x_2)) f_1(x_1) \end{aligned}$$

and using the identity

$$\begin{aligned} c_{12}(F_1(x_1), F_2(x_2)) f_1(x_1) &= \frac{\partial^2 C_{12}(u_1, u_2)}{\partial u_1 \partial u_2} \Big|_{u_1=F_1(x_1), u_2=F_2(x_2)} \frac{\partial u_1}{\partial x_1} \\ &= \frac{\partial}{\partial u_2} \left(\frac{\partial}{\partial x_1} C_{12}(F_1(x_1), u_2) \right) \Big|_{u_2=F_2(x_2)} \end{aligned}$$

one can derive that

$$\begin{aligned} F_{1|2}(x_1|x_2) &= \int_{-\infty}^{x_1} f_{1|2}(z_1|x_2) dz_1 \\ &= \int_{-\infty}^{x_1} \frac{\partial}{\partial u_2} \left(\frac{\partial}{\partial z_1} C_{12}(F_1(z_1), u_2) \right) \Big|_{u_2=F_2(x_2)} dz_1 \\ &= \frac{\partial}{\partial u_2} \left(\int_{-\infty}^{x_1} \frac{\partial}{\partial z_1} C_{12}(F_1(z_1), u_2) dz_1 \right) \Big|_{u_2=F_2(x_2)} \\ &= \frac{\partial}{\partial u_2} C_{12}(F_1(x_1), u_2) \Big|_{u_2=F_2(x_2)}. \end{aligned}$$

□

As for copulas one has univariate marginal distributions one can derive the associated conditional distribution function as

$$C_{1|2}(u_1|u_2) = \frac{\partial}{\partial u_2} C_{12}(u_1, u_2) \quad \forall u_1 \in [0, 1]. \quad (5.15)$$

Copula family	$C(u_1, u_2)$	parameter range	independence
Clayton	$(u_1^{-\delta} + u_2^{-\delta} - 1)^{-\frac{1}{\delta}}$	$0 < \delta < \infty$	$\delta \rightarrow 0$
Gumbel	$\exp(-[(-\ln u_1)^\delta + (-\ln u_2)^\delta]^{\frac{1}{\delta}})$	$\delta \geq 1$	$\delta = 1$
Frank	$-\frac{1}{\delta} \ln \left(\frac{1}{1-e^{-\delta}} ((1-e^{-\delta}) - (1-e^{-\delta u_1})(1-e^{-\delta u_2})) \right)$	$\delta \in \mathbb{R} \setminus \{0\}$	$\delta \rightarrow 0^+$
Joe	$1 - \left((1-u_1)^\delta + (1-u_2)^\delta - (1-u_1)^\delta (1-u_2)^\delta \right)^{\frac{1}{\delta}}$	$\delta \geq 1$	$\delta = 1$

Table 5.1 A selection of bivariate archimedean single parametric copula families.

These bivariate conditional distribution functions are of major importance in the vine copula theory and thus have a special name which is **h-functions**.

Definition 5.2.2 (h-functions of bivariate copulas). *For a bivariate copula C_{12} the corresponding **h-functions** are defined as*

$$h_{1|2}(u_1|u_2) = \frac{\partial}{\partial u_2} C_{12}(u_1, u_2) \quad (5.16)$$

$$h_{2|1}(u_2|u_1) = \frac{\partial}{\partial u_1} C_{12}(u_1, u_2) \quad (5.17)$$

for all $(u_1, u_2) \in [0, 1]^2$.

5.3 Bivariate copula families

Also of major importance for the vine copula theory are bivariate parametric copulas also called **bivariate copula families**. Without further ado the this section will introduce some of them. The simplest bivariate copula is without a doubt the **independence copula**. This copula actually has no parameters and is given by

$$C(u_1, u_2) = u_1 \times u_2$$

Copulas that are derived from elliptical distributions by applying the inverse statement in Theorem 5.2.1 represent the class of **elliptical copulas**. The two most important representatives of this class of copulas are the Gaussian and Student's t copula. The **bivariate Gaussian copula** is given as

$$C(u_1, u_2) = \Phi_2(\Phi^{-1}(u_1), \Phi^{-1}(u_2))$$

with Φ_2 being the two-dimensional standard normal distribution function. The **bivariate Student's t copula** is constructed via

$$C(u_1, u_2) = T_{\nu,2}(T_\nu^{-1}(u_1), T_\nu^{-1}(u_2))$$

with $T_{\nu,2}$ being the two-dimensional standard normal Student's t distribution function and T_ν the univariate standard Student's t distribution function both with degree of freedom ν . The parameters of the copulas directly arise from the location and scale parameters of the elliptical distributions. Another important class of copulas are the **Archimedean copulas** which are constructed via generator functions. Details on the construction can be found in Chapter 3 of Czado (2019). Prominent single parametric representatives of this class of copulas are displayed in Table 5.1.

Of course there are more parametric copula families like the **extreme-value copulas** and Chapter 3 of Czado (2019) gives a tour of all major parametric bivariate copula families. Additionally one can use rotations of these copulas like a 90 degree counterclockwise rotation by $C_{90}(u_1, u_2) = C(1 - u_1, 1 - u_2)$ in order to extend the possible dependence patterns that might be covered. Figure 5.1 shows **marginal normal contour plots** of all the bivariate copula families presented i.e. the marginal copula scale domains

are transformed by the inverse standard normal distribution function ϕ^{-1} for the sake of a better visualization in the tails. Each copula beside the independence copula is shown once with a strong and once with a rather weak dependence. These plots also show that by not restricting oneself to the class of elliptical copulas but also using Archimedean copulas one opens the door for much more flexible dependence patterns especially in the tails. This is also referred to as **tail dependence** whose modeling is considered crucial especially in financial risk modeling as stressed in Pfaff (2016). Now having an idea of bivariate parametric copula families one moves on towards rules to build a pair copula construction from data that uses exactly these bivariate building blocks as a basis.

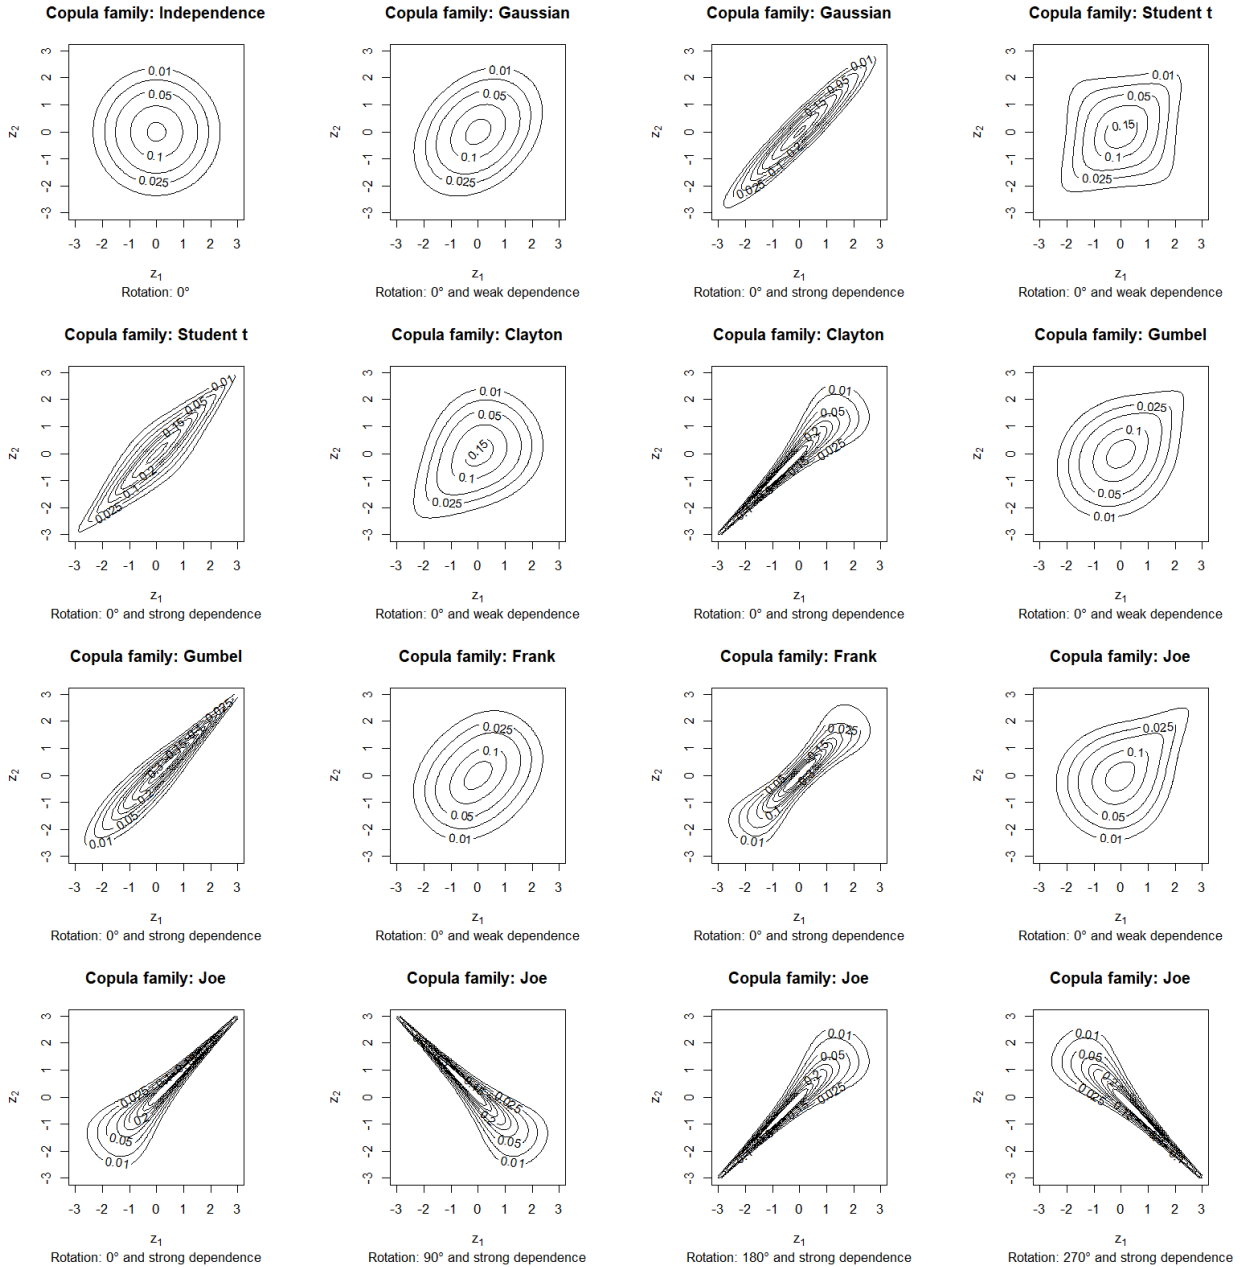


Figure 5.1 Marginal normal contour plots of the presented bivariate copula families.

5.4 Introduction to vine copulas

The fundamental idea in the theory of vine copulas is the one of **pair copula constructions** i.e. decomposing high dimensional distributions into bivariate building blocks by recursive conditioning. This idea can be well illustrated in the three-dimensional case where one uses Sklar's Theorem 5.2.1.

$$\begin{aligned}
 f(x_1, x_2, x_3) &= f_{3|12}(x_3|x_2, x_1) \times f_{2|1}(x_2|x_1) \times f_1(x_1) \\
 &= c_{13;2}(F_{1|2}(x_1|x_2), F_{3|2}(x_3|x_2); x_2) f_{3|2}(x_3|x_2) \times f_{2|1}(x_2|x_1) \times f_1(x_1) \\
 &= c_{13;2}(F_{1|2}(x_1|x_2), F_{3|2}(x_3|x_2); x_2) c_{23}(F_2(x_2), F_3(x_3)) f_3(x_3) \\
 &\quad \times c_{12}(F_1(x_1), F_2(x_2)) f_2(x_2) \times f_1(x_1)
 \end{aligned}$$

where $c_{13;2}$ denotes the associated copula density of the bivariate distribution of (X_1, X_3) given $X_2 = x_2$. This decomposition of an arbitrary three-dimensional density into only bivariate building blocks is however not unique as other - here in the three dimensional case exactly two more - conditioning orders could have been applied. One then speaks of a so called **pair copula decomposition** of some multivariate distribution if the copulas associated with the conditional distributions depend on a conditioning value like x_2 in $c_{13;2}$. If one ignores this dependence one is making the famous **simplifying assumption** i.e. mathematically speaking one assumes $\forall (u_i, u_j) \in [0, 1]^2$ and conditioning set D that $c_{ij|D}(u_i, u_j; X_d = x_d \forall d \in D) = c_{ij}(u_i, u_j)$. This simplifying assumption will be made for the rest of this thesis. For a three-dimensional copula such a pair copula construction simplifies even further as all marginal distributions are uniform.

$$c(u_1, u_2, u_3) = c_{13;2}(C_{1|2}(u_1|u_2), C_{3|2}(u_3|u_2)) \times c_{23}(u_2, u_3) \times c_{12}(u_1, u_2) \quad (5.18)$$

This formula could be used constructively if one uses the parametric bivariate copulas as this would allow to construct a potentially high dimensional - in this case three-dimensional - parametric density with these bivariate building blocks. The concept of regular vine copulas now formalizes this constructive procedure.

5.5 Regular (R) vine copulas

In order to achieve an algorithmic approach for the build up of pair copula constructions in d dimensions one utilizes tools from graph theory as introduced in Bedford and Cooke (2002). One thinks of the bivariate (conditional) copula building blocks as edges in a sequence of linked graphs. More specific these graphs are trees which means that all nodes are connected and there are no circles. What the for now mysterious nodes of these graphs are will become evident after the next definition. Figure 5.2 shows such a sequence of trees in 5 dimensions. Actually this sequence of trees in Figure 5.2 is a valid **regular vine (R-vine) tree sequence**. Such an R-vine tree sequence poses some important constraints on the tree sequence that allow to build valid pair copula constructions in high dimensions and is now defined formally.

Definition 5.5.1 (Regular vine tree sequence). *The set of trees $\mathcal{W} = (T_1, \dots, T_d)$ is called a **regular vine tree sequence** on d elements if*

1. $\forall j \in \{1, \dots, d\} : T_j$ is connected i.e. there exists a path from each node to each other node.
2. T_1 has the fixed node set $N_1 = \{1, \dots, d\}$.
3. Every node set N_j with $j \geq 2$ is exactly the edge set E_{j-1} .
4. For each $\{a, b\} \in E_j$ it has to hold that $|a \cap b| = 1$ for $j \geq 2$.

where property (4) is called the **proximity condition** that ensures that for each edge (a, b) in tree $T_j, j \geq 2$ between the nodes a and b , that are thus edges in the previous tree T_{j-1} , it holds that these edges a and b have a common node in T_{j-1} .

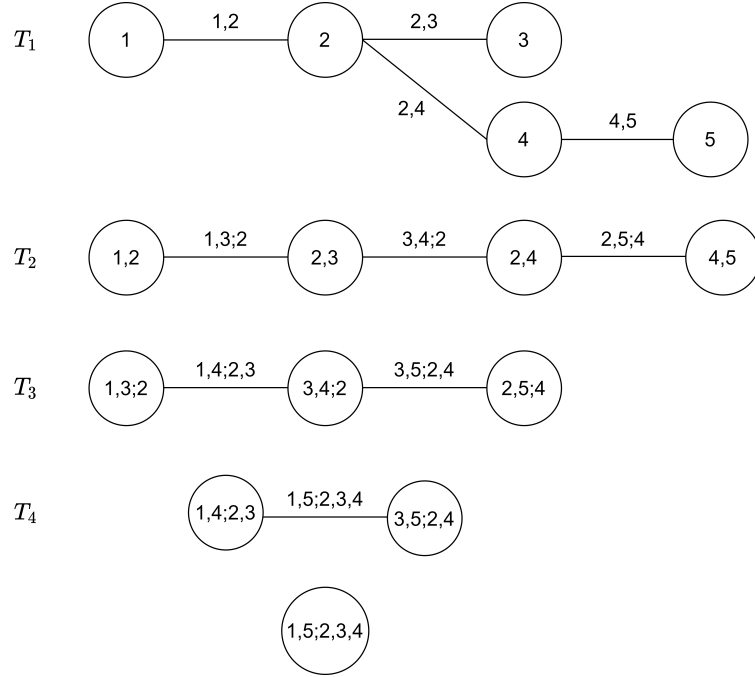


Figure 5.2 Exemplary R-vine tree sequence on 5 elements.

Before formally defining a regular vine copula some additional notation regarding such a tree sequence is needed.

Definition 5.5.2 (Complete union and conditioned sets). *Given a regular vine tree sequence \mathcal{W} as of Definition 5.5.1 and an edge $e = \{a, b\} \in E_j$ where the tree T_j is comprised of the edge set E_j and node set N_j one can define the following sets.*

- **complete union:** $A_e := \{i \in N_1 | \exists e_1 \in E_1, \dots, e_{j-1} \in E_{j-1} \text{ such that } i \in e_1 \in \dots \in e_{j-1} \in e\}$
- **conditioning set:** $D_e := A_a \cup A_b$
- **conditioned sets:** $C_{e,a} := A_a \setminus D_e, C_{e,b} := A_b \setminus D_e$ and $C_e := C_{e,a} \cap C_{e,b}$

A common abbreviation of the edge $e = (C_{e,a}, C_{e,b}; D_e)$ is $e = (e_a, e_b; D_e)$.

Now one can define a regular vine copula.

Definition 5.5.3 (Regular vine copula). *A d -dimensional copula for the random vector $\mathbf{U} = (U_1, \dots, U_d)$ is called a **regular vine (R-vine) copula** if one can determine \mathcal{W} and \mathcal{B} such that*

1. **R-vine tree sequence:** \mathcal{W} is a R-vine tree sequence on d elements.
2. **Bivariate copulas:** A set $\mathcal{B} = \{C_e | e \in E_j \text{ with } j = 1, \dots, d-1\}$ with C_e being symmetric bivariate copulas.
3. **Relationship between \mathcal{B} and \mathcal{W} :** Each edge $e = \{a, b\}$ in the trees of \mathcal{W} is associated with the corresponding copula C_e in \mathcal{B} such that it is the associated conditional distribution of $U_{C_{e,a}}$ and $U_{C_{e,b}}$ given $U_{D_e} = \mathbf{u}_{D_e}$. Moreover the simplifying assumption is made and thus $C_e(\cdot, \cdot)$ does not depend on the value of \mathbf{u}_{D_e} .

As one speaks of a copula of course all marginal distribution functions of the U_j are uniform distributions on $[0, 1]$. Additionally if the bivariate copulas contained in \mathcal{B} are parametric one collects the copula parameters of the respective pair copulas in \mathcal{B} in the set $\Theta(\mathcal{B})$.

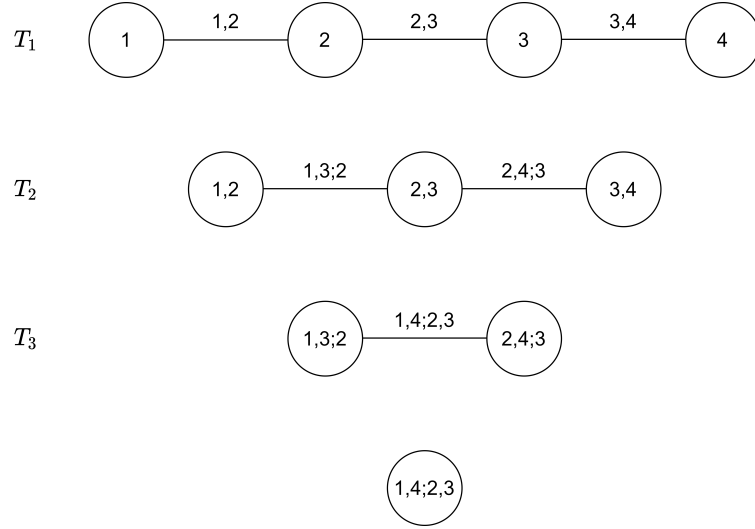


Figure 5.3 Exemplary D-vine tree sequence on 4 elements.

5.6 Drawable (D) vine copulas

Beside the general R-vine copulas there are also two prominent special cases namely the **drawable vines (D-vines)** and the **canonical vines (C-vines)**. In this thesis only the D-vines will be used beside the general R-vine. For details on C-vines a consultation of Chapter 4 of Czado (2019) is recommended.

Definition 5.6.1 (Drawable vine copula). *An R-vine copula is called a **drawable vine (D-vine)** if the following additional constraint holds for its tree sequence \mathcal{W} :*

$$\forall n \in N_j, j = 1, \dots, d-1 : |\{e \in E_j | n \in e\}| \geq 2. \quad (5.19)$$

This means graphically that each tree is just an ordered connected path with no loops.

Example 5.6.1. *An exemplary 4-dimensional D-vine tree sequence is displayed in Figure 5.3. The corresponding D-vine copula density is given by*

$$\begin{aligned} c_{1234}(u_1, u_2, u_3, u_4) &= c_{12}(u_1, u_2) \times c_{23}(u_2, u_3) \times c_{34}(u_3, u_4) \\ &\quad \times c_{13;2}(C_{1|2}(u_1|u_2), C_{3|2}(u_3|u_2)) \times c_{24;3}(C_{2|3}(u_2|u_3), C_{4|3}(u_4|u_3)) \\ &\quad \times c_{14;23}(C_{1|23}(u_1|u_2, u_3), C_{4|23}(u_4|u_2, u_3)). \end{aligned} \quad (5.20)$$

A notable fact is that due to the proximity condition if the ordering in the first tree is fixed all other trees are fixed too as outlined in Czado (2019). Moreover the nodes that are at the end of the paths are often called **leaf nodes**.

5.7 Recursion for conditional distribution functions and estimation

In order to decompose the conditional distributions appearing in Equation (5.20) for the estimation or simulation from such vine copulas the following theorem whose proof was first stated in Joe (1996) comes in very handy.

Theorem 5.7.1 (Recursion for conditional distribution functions). *The conditional distribution function $C_{j|D}$ with the conditioning set D can be decomposed with $i \in D$ and $D_{-i} = D \setminus \{i\}$ like*

$$C_{j|D}(u_j | \mathbf{u}_D) = h_{j|i;D_{-i}} \left(C_{j|D_{-i}}(u_j | u_{D_{-i}}) | C_{i|D_{-i}}(u_i | u_{D_{-i}}) \right). \quad (5.21)$$

So far the introduction to vine copulas was done in a theoretical manner. In practice one is obviously interested in estimating such vine copulas from copula scale data in order to model the dependence structure between random quantities of interest. The estimation procedure has to determine on the one side the vine tree structure \mathcal{W} and on the other side the set of bivariate copulas \mathcal{B} corresponding to the tree edges alongside its parameters $\Theta(\mathcal{B})$. The most prominent method in practice is currently the **sequential top-down estimation** as introduced in Dißmann et al. (2013). The estimation boils down to the following steps that are iterated over the trees in a top down fashion i.e. starting with T_1 .

1. Select the edges of the tree such that the sum of the edge weights, which are for example for the edge $\{a, b\}$ chosen to be an absolute empirical dependence measure of the data corresponding to a and b , is maximized. Note that the first tree is arbitrary but for the subsequent tree only edges that satisfy the proximity condition are allowed.
2. For each edge estimate the parameters of a specified set of bivariate copula families. This is mostly performed via a classical maximum likelihood estimation. Then choose the copula family with the corresponding estimated parameters for the edge that minimizes the AIC.
3. Generate pseudo observations from the fitted bivariate copulas to perform the above steps on this pseudo data representing the next tree level. For example for the edge $\{a, b; D\}$ in the next tree with conditioning set D one uses the pseudo observations generated from the parametric estimates of $C_{a|D}$ and $C_{b|D}$ for the estimation of the edge weight and the bivariate copula.

This sequentially estimates the tree structure \mathcal{W} and the pair copulas \mathcal{B} with parameters $\Theta(\mathcal{B})$. Although the algorithm might propagate misspecifications sequentially and makes only locally optimal choices in a greedy fashion it has been shown in simulation studies and in practice that it has a good performance. As the estimation procedure is not the primary focus of this thesis one refers to Chapter 7 of Czado (2019) as an excellent resource to get an overview of vine copula estimation techniques. The same goes for model selection. The most prominent tools in this setting are classical likelihood based information criteria like the **Akaike information criterion (AIC)** but now not on the pair copula but overall vine level.

6 Simulation from vine copulas

6.1 General simulation from multivariate distributions

For the risk measure estimation that will be proposed later a critical part will be to sample from a given vine copula. Thus now we assume to have selected and estimated the tree sequence \mathcal{W} as well as the bivariate copula families in the set \mathcal{B} and their respective parameters $\Theta(\mathcal{B})$. At the base of the vine copula simulation algorithms presented in this thesis is the so called **Rosenblatt transform** which was first introduced in Rosenblatt (1952) and allows to sample from a d -dimensional distribution using inverse conditional distributions.

Theorem 6.1.1 (Rosenblatt transform). *One obtains a multivariate random sample (x_1, \dots, x_d) from the d -dimensional distribution $F_{1,\dots,d}$ by iterating through the following steps.*

$$\begin{aligned} & \text{Sample } w_j \stackrel{i.i.d.}{\sim} \mathcal{U}(0, 1), j = 1, \dots, d \\ x_1 & := F_1^{-1}(w_1) \\ x_2 & := F_{2|1}^{-1}(w_2|x_1) \\ & \vdots \\ x_d & := F_{d|d-1,\dots,1}^{-1}(w_d|x_{d-1}, \dots, x_1) \end{aligned}$$

The proof of Theorem 6.1.1 can be performed via induction with the induction start being trivial as x_1 is obviously a sample from the marginal distribution of X_1 . Thus computing conditional distribution functions will be integral for sampling from them. For this endeavour the result in Theorem 5.7.1 comes in very handy.

6.2 Sampling from a D-vine copula

For a three dimensional R-vine copula which by the way has to be also a D-vine and has a fixed tree structure up to permutations of the node order we have the bivariate copulas C_{12} , C_{23} and $C_{13;2}$ with corresponding parameters θ_{12} , θ_{23} and $\theta_{13;2}$. In order to get the random sample (u_1, u_2, u_3) from this three-dimensional vine copula we then determine the needed conditional distributions according to the Rosenblatt transformation by applying the Theorem 5.7.1 and using the uniform samples w_1, w_2, w_3 .

$$\begin{aligned} u_1 & = w_1 \\ w_2 & = C_{2|1}(u_2|u_1) = h_{2|1}(u_2|u_1) \\ & \Rightarrow u_2 = h_{2|1}^{-1}(w_2|u_1) \\ w_3 & = C_{3|21}(u_3|u_2, u_1) \\ & = \frac{\partial}{\partial C_{1|2}(u_1|u_2)} C_{13;2}(C_{1|2}(u_1|u_2), C_{3|2}(u_3|u_2)) \\ & = h_{3|1;2}(h_{3|2}(u_3|u_2)|h_{1|2}(u_1|u_2)) \\ & \Rightarrow u_3 = h_{3|2}^{-1}(h_{3|1;2}^{-1}(w_3|h_{1|2}(u_1|u_2))|u_2) \end{aligned}$$

These tedious calculations of conditional copulas for the simulation from arbitrary D- or R-vines are performed using iterative algorithms. Firstly the algorithm presented in Czado (2019) for the sampling from D-vines will now be discussed. The algorithm that aims to calculate iteratively conditional distributions uses three $d \times d$ auxiliary matrices. The matrices V and V^2 contain the to be computed conditional distribution functions.

$$V := \begin{pmatrix} u_1 & u_2 & u_3 & u_4 & \dots \\ - & C_{2|1}(u_2|u_1) & C_{3|2}(u_3|u_2) & C_{4|3}(u_4|u_3) & \dots \\ - & - & C_{3|2,1}(u_3|u_2, u_1) & C_{4|3,2}(u_4|u_3, u_2) & \dots \\ - & - & - & C_{4|3:1}(u_4|u_{3:1}) & \dots \\ - & - & - & - & \ddots \end{pmatrix} \in \mathbb{R}^{d \times d} \quad (6.1)$$

$$V^2 := \begin{pmatrix} u_1 & u_2 & u_3 & u_4 & \dots \\ - & C_{1|2}(u_1|u_2) & C_{2|3}(u_2|u_3) & C_{3|4}(u_3|u_4) & \dots \\ - & - & C_{1|2,3}(u_1|u_2, u_3) & C_{2|3,4}(u_2|u_3, u_4) & \dots \\ - & - & - & C_{1|2:4}(u_1|u_{2:4}) & \dots \\ - & - & - & - & \ddots \end{pmatrix} \in \mathbb{R}^{d \times d} \quad (6.2)$$

The matrix Θ containing all the parameters of the bivariate copulas corresponding to the edges of the tree sequence.

$$\Theta := \begin{pmatrix} - & \theta_{1,2} & \theta_{2,3} & \theta_{3,4} & \dots \\ - & - & \theta_{3,1;2} & \theta_{4,2;3} & \dots \\ - & - & - & \theta_{4,1;3,2} & \dots \\ - & - & - & - & \ddots \end{pmatrix} \in \mathbb{R}^{d \times d} \quad (6.3)$$

The iterative algorithm that takes Θ as the input in order to calculate V and V^2 is given in Algorithm 1.

Algorithm 1: Sampling from D-vine copulas (Taken from Czado (2019))

Input : Parameter matrix Θ for the d -dimensional D-vine $C_{1,\dots,d}(\cdot)$

Output: A d -dimensional sample u_1, \dots, u_d from $C_{1,\dots,d}(\cdot)$

Sample $w_j \stackrel{i.i.d.}{\sim} \mathcal{U}(0, 1)$, $j \in \{1, \dots, d\}$

$v_{1,1} = w_1$, $v_{1,1}^2 = w_1$

for $j \leftarrow 2$ **to** d **do**

$v_{j,j} = w_j$

for $k \leftarrow j - 1$ **to** 1 **do**

$v_{k,j} = h_{j|j-k;(j-k+1):(j-1)}^{-1}(v_{k+1,j}|v_{k,j-1}^2, \theta_{j|j-k;(j-k+1):(j-1)})$

if $j < d$ **then**

$v_{k+1,j}^2 = h_{j-k|j;(j-k+1):(j-1)}(v_{k,j-1}^2|v_{k,j}, \theta_{j-k|j;(j-k+1):(j-1)})$

end

$v_{1,j}^2 = v_{1,j}$

end

$u_i = v_{1,i}$ **for** $j \in \{1, \dots, d\}$

For regular vines the approach is quite similar and is facilitated by the introduction of one further auxiliary matrix. This algorithm will however not be discussed in detail here as the focus of the remainder of this section will stay with D-vine copulas. The general algorithm for unconditional R-vine sampling as for D-vines in Algorithm 1 can be found in Czado (2019) and a formal proof of it is given in Chapter 5 of Dißmann (2010).

6.3 Sampling from a D-vine copula conditional on the rightmost leaf node

Here we want to first construct and then sample conditionally from a specific D-vine copula. We assume now to have the d assets A_1, \dots, A_d and a market index I with corresponding i.i.d. copula data for the fitting. So we want to capture their dependence using a $d + 1$ dimensional D-vine copula.

Determine the D-vine ordering

As for the fitting of a D-vine copula as described in the previous subsections an ordering is needed. This ordering directly specifies the vine tree structure \mathcal{W} . Here an ordering of the following scheme is proposed, using a permutation j_1, \dots, j_d of the indices $1, \dots, d$.

$$A_{j_d} - A_{j_{d-1}} - \dots - A_{j_1} - I \quad (6.4)$$

This ordering allows for the following practical interpretation. Asset A_{j_1} is most dependent on the market index I , the asset A_{j_2} is most dependent on the asset A_{j_1} after the effect of the market index I is accounted for, A_{j_3} is most dependent on A_{j_2} after accounting for A_{j_1} and I and so on. Here an approach using the Pearson and partial correlations as given in Definition 5.1.2 is proposed in Algorithm 2 in order to determine a sensible ordering for the fitting process.

Example 6.3.1. *Algorithm 2 can be well explained using the 4 dimensional copula $C_{j_3, j_2, j_1, I}$. We can decompose its density into the bivariate densities $c_{j_1, I}$, c_{j_2, j_1} , c_{j_3, j_2} , $c_{j_2, I; j_1}$, $c_{j_3, j_1; j_2}$ and $c_{j_3, I; j_2, j_1}$ just like in Example 5.6.1. Figure 6.1 illustrates this D-vine where each edge corresponds to one of these bivariate densities. Now we would like to iteratively determine the permutation j_1 up to j_d in such a way that the overall likelihood which can be factorized with the mentioned bivariate densities is as high as possible. The idea is to estimate weights of the edges corresponding to the bivariate copula densities. These weights are in the case of the first tree levels the absolute Pearson correlation coefficients estimated on the normalized scale data and for the other tree level that have conditioning sets we use the partial correlation to account for the conditioning set i.e. the weight for edge $j_3, j_1; j_2$ is $|\rho_{j_3, j_1; j_2}|$ estimated again on the normalized scale data. The transformation of the copula data to the normalized scale is facilitated through Equation (5.2). Then we determine the next j_k such that the sum of all the edge weights corresponding to the additional edges that were not yet specified but contain only j_k, j_{k-1}, \dots, I is maximized. The considered edges in each step are highlighted in different colors in Figure 6.1.*

This greedy process does of course not guarantee a global optimum as it makes local choices. Especially for high dimensions however this approach can be simplified by introducing the cutoff depth c_{depth} . This implies that only all Pearson or partial correlation coefficients corresponding to the edges up to the specified depth will be used for the calculation of the ordering. This cutoff value is indicated in Figure 6.1 by a dotted line so all edges below the line are not considered. Example 6.3.2 showcases Algorithm 2 in 5 dimensions.

Example 6.3.2. *This example showcases how Algorithm 2 determines the desired permutation j_1, \dots, j_d in 5 dimensions i.e. $d = 4$. Here we will not use the cutoff depth so it defaults to d . Thus we would like to determine j_1 up to j_4 as I is fixed as the rightmost leaf node. Given the normalized data z_1, \dots, z_4, z^I , j_1 is straight forwardly given as the variable index that maximizes the absolute Pearson correlation coefficient between the market index I and itself estimated on the normalized scale i.e. $|\rho_{j_1, I}|$ estimated using z_{j_1} and*

Algorithm 2: Determine D-vine ordering including a market index as the rightmost leaf node

Input : A_1, \dots, A_d and a market index I with corresponding i.i.d. copula data u_1, \dots, u_d, u^I as well as a cutoff depth $c_{\text{depth}} \in \{1, \dots, d\}$ defaulting to d

Output: j_1, \dots, j_d permutation for the indices $1, \dots, d$ specifying an ordering as in (6.4)

Transform the copula data to the normalized scale using Equation (5.2). The data on the normalized scale is denoted by z_1, \dots, z_d, z^I with the convention $z_{j_0} = z^I$ for simplicity in the algorithm below.

$j_1 \leftarrow \delta \in \{1, \dots, d\}$ s.t. $|\rho_{\delta, I}| \geq |\rho_{\delta', I}|, \forall \delta' \in \{1, \dots, d\}$

for $l \leftarrow 2$ **to** d **do**

| $j_l = \operatorname{argmax}_{\delta \in \{1, \dots, d\} \setminus \{j_1, \dots, j_{l-1}\}} \sum_{k=\max(l-c_{\text{depth}}, 0)}^{l-1} |\rho_{\delta, j_k; j_{k+1}:j_{l-1}}|$

end

Here $\rho_{a,b;D}$ denotes the Pearson (if D is the empty set) or partial correlation coefficient between z_a and z_b with conditioning set D and it is estimated always using the data on the normalized scale.

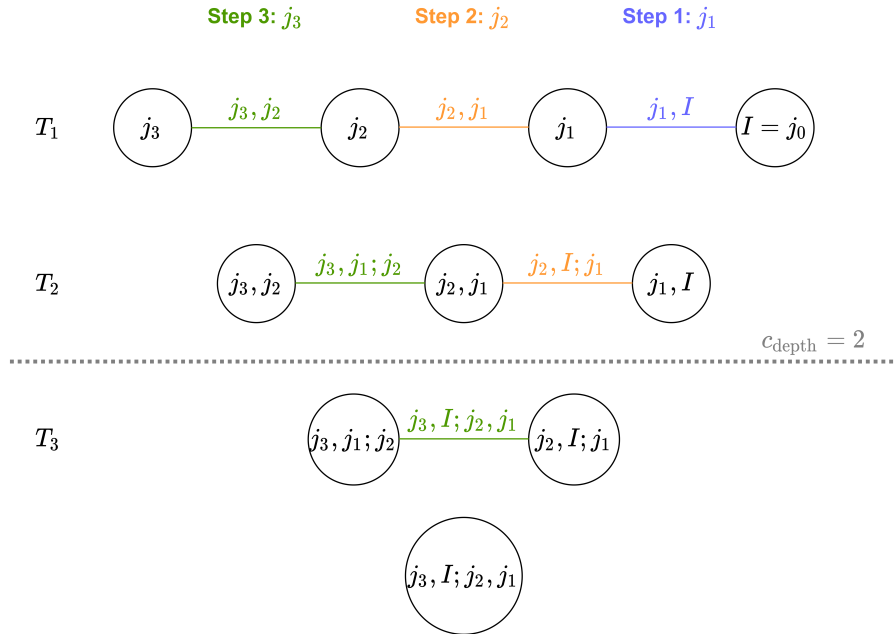


Figure 6.1 Illustration for Algorithm 2 in 4 dimensions.

z^I is maximized with respect to j_1 . We then enter the for loop in the algorithm to determine the remaining j_2, j_3, j_4 .

$$\begin{aligned}
l = 2 : \quad j_2 &= \operatorname{argmax}_{\delta \in \{1,2,3,4\} \setminus \{j_1\}} |\rho_{\delta, j_0=I; j_1: j_1}| + |\rho_{\delta, j_1; j_2: j_1}| \\
&= \operatorname{argmax}_{\delta \in \{1,2,3,4\} \setminus \{j_1\}} |\rho_{\delta, I; j_1}| + |\rho_{\delta, j_1; \emptyset}| \\
&= \operatorname{argmax}_{\delta \in \{1,2,3,4\} \setminus \{j_1\}} |\rho_{\delta, I; j_1}| + |\rho_{\delta, j_1}| \\
l = 3 : \quad j_3 &= \operatorname{argmax}_{\delta \in \{1,2,3,4\} \setminus \{j_1, j_2\}} |\rho_{\delta, j_0=I; j_1: j_2}| + |\rho_{\delta, j_1; j_2: j_2}| + |\rho_{\delta, j_2; j_3: j_2}| \\
&= \operatorname{argmax}_{\delta \in \{1,2,3,4\} \setminus \{j_1, j_2\}} |\rho_{\delta, I; j_1, j_2}| + |\rho_{\delta, j_1; j_2}| + |\rho_{\delta, j_2}| \\
l = 4 : \quad j_4 &= \{1, 2, 3, 4\} \setminus \{j_1, j_2, j_3\}
\end{aligned}$$

Sampling algorithm conditioned on the rightmost leaf node

Now assume the D-vine copula $C_{j_1, \dots, j_d|I}$ is fitted. The next step is to have an algorithm at hand to sample from $C_{j_1, \dots, j_d|I}(\cdot|u^I)$ for u^I being a fixed realisation of the market index I on the copula scale. Before writing down a generic algorithm we start with an illustrative example.

Example 6.3.3. Assume the ordering $A_1 - A_2 - A_3 - I$ with fitted D-vine copula $C_{1,2,3,I}$. So the following pair copulas were fitted: $C_{1,2}, C_{2,3}, C_{3,I}, C_{1,3;2}, C_{2,I;3}, C_{1,I;2,3}$. The exemplary 4-dimensional D-vine is displayed in Figure 6.2. As in the regular vine case the simulation approach will be based on a Rosenblatt transform as given in Theorem 6.1.1. The goal is to conditionally sample conditioned on the index I thus we start the Rosenblatt transform from the back and sample our way to the left i.e. first sample u_3 from the inverse of the conditional copula $C_{3|I}$, then sample u_2 from the inverse of $C_{2|3,I}$ and so on.

The detailed computations are elaborated now down below for this example.

$$\begin{aligned}
 & u^I \text{ fixed} \\
 & w_1, w_2, w_3 \stackrel{i.i.d.}{\sim} \mathcal{U}(0, 1) \\
 w_3 &= C_{3|I}(u_3|u^I) \\
 &= h_{3|I}(u_3|u^I) \\
 &\Rightarrow u_3 = h_{3|I}^{-1}(w_3|u^I), \text{ so } u_3 \text{ is a sample from } C_{3|I=u^I}. \\
 w_2 &= C_{2|3,I}(u_2|u_3, u^I) \\
 &= \frac{\partial}{\partial C_{I|3}(u^I|u_3)} C_{2,I;3}(C_{2|3}(u_2|u_3), C_{I|3}(u^I|u_3)) \\
 &= \frac{\partial}{\partial v_2} C_{2,I;3}(h_{2|3}(u_2|u_3), v_2) |_{v_2=h_{I|3}(u^I|u_3)} \\
 &= h_{2|I;3}(h_{2|3}(u_2|u_3), h_{I|3}(u^I|u_3)) \\
 &\Rightarrow u_2 = h_{2|3}^{-1}(h_{2|I;3}^{-1}(w_2|h_{I|3}(u^I|u_3))|u_3), \text{ so } u_2 \text{ is a sample from } C_{2|3=u_3, I=u^I}. \\
 w_1 &= C_{1|2,3,I}(u_1|u_2, u_3, u^I) \\
 &= \frac{\partial}{\partial C_{I|2,3}(u^I|u_2, u_3)} C_{1,I;2,3}(C_{1|2,3}(u_1|u_2, u_3), C_{I|2,3}(u^I|u_2, u_3)) \\
 &= h_{1|I;2,3}(C_{1|2,3}(u_1|u_2, u_3)|C_{I|2,3}(u^I|u_2, u_3)) \\
 &= h_{1|I;2,3}(h_{1|3;2}(C_{1|2}(u_1|u_2)|C_{3|2}(u_3, u_2))|C_{I|2,3}(u^I|u_2, u_3)) \\
 &= h_{1|I;2,3}(h_{1|3;2}(h_{1|2}(u_1|u_2)|h_{3|2}(u_3, u_2))|C_{I|2,3}(u^I|u_2, u_3)) \\
 &\Rightarrow u_1 = h_{1|2}^{-1}(h_{1|3;2}^{-1}(h_{1|I;2,3}^{-1}(w_1|C_{I|2,3}(u^I|u_2, u_3))|h_{3|2}(u_3, u_2))|u_2), \\
 &\text{so with } C_{I|2,3}(u^I|u_2, u_3) = h_{I|2,3}(h_{I|3}(u^I|u_3)|h_{2|3}(u_2, u_3)) \\
 &u_1 \text{ is a sample from } C_{1|2=u_2, 3=u_3, I=u^I}.
 \end{aligned}$$

Finally (u_1, u_2, u_3) is a realization of the conditional distribution of (U_1, U_2, U_3) given $I = u^I$ and it can be easily checked that all the needed bivariate copulas are estimated by the assumed D-vine. So for a 4 dimensional D-Vine copula one can sample conditional on the last instance backwards with the formulas given above. It is also evident that this can be extended to arbitrary $d + 1$ dimensions. But it is also quite obvious in the example that the formulas blow up quite quickly and thus a generic sampling algorithm like the one already established for the unconditional sampling from D-vines has to be stated. This is the next step.

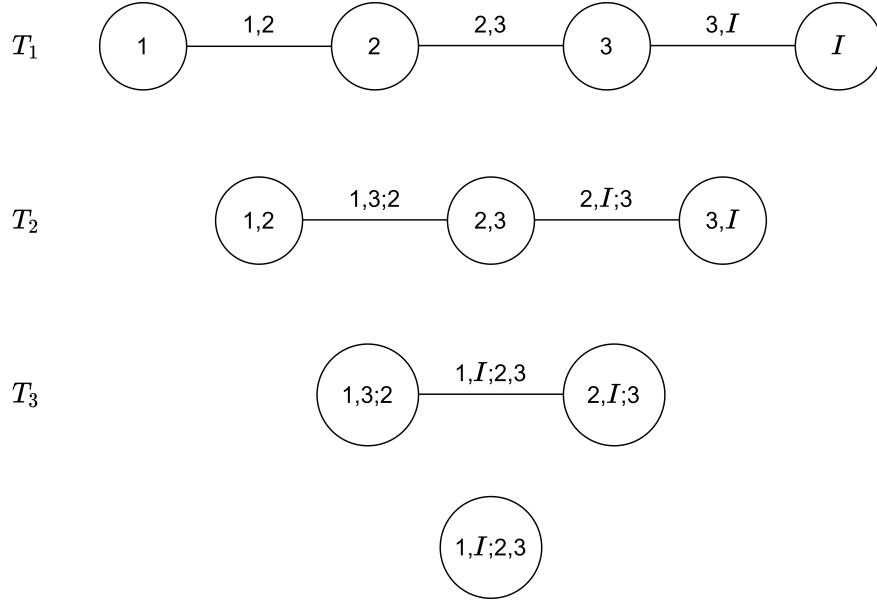


Figure 6.2 Illustration of the tree sequence of a 4-dimensional D-vine with the market index I as the rightmost leaf node.

Now for simplicity of notation in the algorithm the market index I will be denoted by the $d+1$ component e.g. $u^I = u_{d+1}$. The fitted parameter matrix Θ , the notation is adapted from Algorithm 1 and the auxiliary matrices \mathcal{V} and \mathcal{V}^2 containing the to be computed conditional distribution functions are given below.

$$\Theta := \begin{pmatrix} - & \theta_{1,2} & \theta_{2,3} & \theta_{3,4} & \dots \\ - & - & \theta_{3,1;2} & \theta_{4,2;3} & \dots \\ - & - & - & \theta_{4,1;3,2} & \dots \\ - & - & - & - & \ddots \end{pmatrix} \in \mathbb{R}^{d+1 \times d+1} \quad (6.5)$$

$$\mathcal{V} := \begin{pmatrix} u_{d+1} & u_d & u_{d-1} & u_{d-2} & \dots \\ - & C_{d-1|d}(u_{d-1}|u_d) & C_{d-2|d-1}(u_{d-2}|u_{d-1}) & C_{d-3|d-2}(u_{d-3}|u_{d-2}) & \dots \\ - & - & C_{d-2|d-1,d}(u_{d-2}|u_{d-1}, u_d) & C_{d-3|d-2,d-1}(u_{d-3}|u_{d-2}, u_{d-1}) & \dots \\ - & - & - & C_{d-3|(d-2):d}(u_{d-3}|u_{(d-2):d}) & \dots \\ - & - & - & - & \ddots \end{pmatrix} \in \mathbb{R}^{d+1 \times d+1} \quad (6.6)$$

$$\mathcal{V}^2 := \begin{pmatrix} u_{d+1} & u_d & u_{d-1} & u_{d-2} & \dots \\ - & C_{d|d-1}(u_d|u_{d-1}) & C_{d-1|d-2}(u_{d-1}|u_{d-2}) & C_{d-2|d-3}(u_{d-2}|u_{d-3}) & \dots \\ - & - & C_{d|d-1,d-2}(u_d|u_{d-1}, u_{d-2}) & C_{d-1|d-2,d-3}(u_{d-1}|u_{d-2}, u_{d-3}) & \dots \\ - & - & - & C_{d|(d-1):(d-3)}(u_d|u_{(d-1):(d-3)}) & \dots \\ - & - & - & - & \ddots \end{pmatrix} \in \mathbb{R}^{d+1 \times d+1} \quad (6.7)$$

Note that the above defined matrices $\mathcal{V}, \mathcal{V}^2$ are not the same as the auxiliary matrices V, V^2 of Algorithm 1. Now the formal algorithm to iteratively derive those matrices and thus also the conditional sample in the first row of the final \mathcal{V} matrix is proposed.

Algorithm 3: Sampling from D-vine copulas conditional on the last component

Input : Parameter matrix Θ for the $d+1$ dimensional D-vine and the fixed conditioning value u_{d+1} for the last component

Output: A d dimensional sample u_1, \dots, u_d from $C_{1, \dots, d|d+1}(\cdot|u_{d+1})$

Sample $w_j \stackrel{i.i.d.}{\sim} \mathcal{U}(0, 1)$, $j \in \{1, \dots, d\}$

$v_{1,1} = u_{d+1}$, $v_{1,1}^2 = u_{d+1}$

for $j \leftarrow 2$ **to** $d + 1$ **do**

$a = d + 2 - j$

$v_{j,j} = w_a$

for $k \leftarrow j - 1$ **to** 1 **do**

$v_{k,j} = h_{a|a+k;(a+1):(a+k-1)}^{-1}(v_{k+1,j}|v_{k,j-1}^2, \theta_{a|a+k;(a+1):(a+k-1)})$

if $j < d + 1$ **then**

$v_{k+1,j}^2 = h_{a+k|a;(a+1):(a+k-1)}(v_{k,j-1}^2|v_{k,j}, \theta_{a+k|a;(a+1):(a+k-1)})$

end

$v_{1,j}^2 = v_{1,j}$

end

$u_i = v_{1,d+2-j}$ for $j \in \{1, \dots, d\}$

Example 6.3.4. In order to grasp Algorithm 3 better it will be applied to a 5 dimensional example. So $d = 4$ and the conditioning component is the fifth. The parameter specifications contained in Θ will be omitted for this display to not crowd the computations too much.

$$u_5 = u^I \text{ given}$$

$$w_1, w_2, w_3, w_4 \stackrel{i.i.d.}{\sim} \mathcal{U}(0, 1)$$

$$v_{1,1} = u_5, v_{1,1}^2 = u_5$$

$$\text{Get } u_4 : v_{2,2} = w_4$$

$$v_{1,2} = h_{4|5}^{-1}(v_{2,2}|v_{1,1}^2) = u_4$$

$$\text{correct as } w_4 = C_{4|5}(u_4|u_5)$$

$$\text{Get } u_3 : v_{2,2}^2 = h_{5|4}(v_{1,1}^2|v_{1,2})$$

$$v_{1,2}^2 = v_{1,2}$$

$$v_{3,3} = w_3$$

$$v_{2,3} = h_{3|5;4}^{-1}(v_{3,3}|v_{2,2}^2)$$

$$v_{3,3}^2 = h_{5|3;4}(v_{2,2}^2|v_{2,3})$$

$$v_{1,3} = h_{3|4}^{-1}(v_{2,3}|v_{1,2}^2) = u_3$$

$$\text{correct as } w_3 = C_{3|4,5}(u_3|u_4, u_5) = h_{3|5;4}(C_{3|4}(u_3|u_4)|C_{5|4}(u_5, u_4))$$

$$\text{Get } u_2 : v_{2,3}^2 = h_{4|3}(v_{1,2}^2|v_{1,3})$$

$$v_{1,3}^2 = v_{1,3}$$

$$v_{4,4} = w_2$$

$$v_{3,4} = h_{2|5;3,4}^{-1}(v_{4,4}|v_{3,3}^2)$$

$$v_{4,4}^2 = h_{5|2;3,4}(v_{3,3}^2|v_{3,4})$$

$$v_{2,4} = h_{2|4;3}^{-1}(v_{3,4}|v_{2,3}^2)$$

$$v_{3,4}^2 = h_{4|2;3}(v_{2,3}^2|v_{2,4})$$

$$v_{1,4} = h_{2|3}^{-1}(v_{2,4}|v_{1,3}^2) = u_2$$

$$\text{correct as } w_2 = C_{2|3,5}(u_2|u_{3:5}) = h_{2|5;3,4}(h_{2|4;3}(h_{2|3}(u_2|u_3)|C_{4|3}(u_4|u_3))|C_{5|3,4}(u_5|u_3, u_4))$$

$$\text{Get } u_1 : v_{2,4}^2 = h_{3|2}(v_{1,3}^2|v_{1,4})$$

$$v_{1,4}^2 = v_{1,4}$$

$$v_{5,5} = w_1$$

$$v_{4,5} = h_{1|5;2,4}^{-1}(v_{5,5}|v_{4,4}^2)$$

$$v_{3,5} = h_{1|4;2,3}^{-1}(v_{4,5}|v_{3,4}^2)$$

$$v_{2,5} = h_{1|3;2}^{-1}(v_{3,5}|v_{2,4}^2)$$

$$v_{1,5} = h_{1|2}^{-1}(v_{2,5}|v_{1,4}^2) = u_1$$

$$\text{correct as } w_1 = C_{1|2,5}(u_1|u_{2:5})$$

$$= h_{1|5;2,4}(h_{1|4;2,3}(h_{1|3;2}(h_{1|2}(u_1|u_2)|C_{3|2}(u_3|u_2))|C_{4|3,2}(u_4|u_3, u_2))|C_{5|2,4}(u_5|u_{2:4}))$$

$$\text{for completeness: } v_{1,5}^2 = v_{1,5}$$

In order to understand the algorithm even better Figure 6.3 might help. It shows visually how the algorithm builds up the matrices. First the diagonal of \mathcal{V} is initialized with the conditional value and the uniform samples. Then one follows the arrows corresponding to calls of (inverse) h-functions of both matrices, always alternating between the two, from left to right in order to calculate all the needed quantities to finally end

up with the desired sample in the first rows of both matrices.

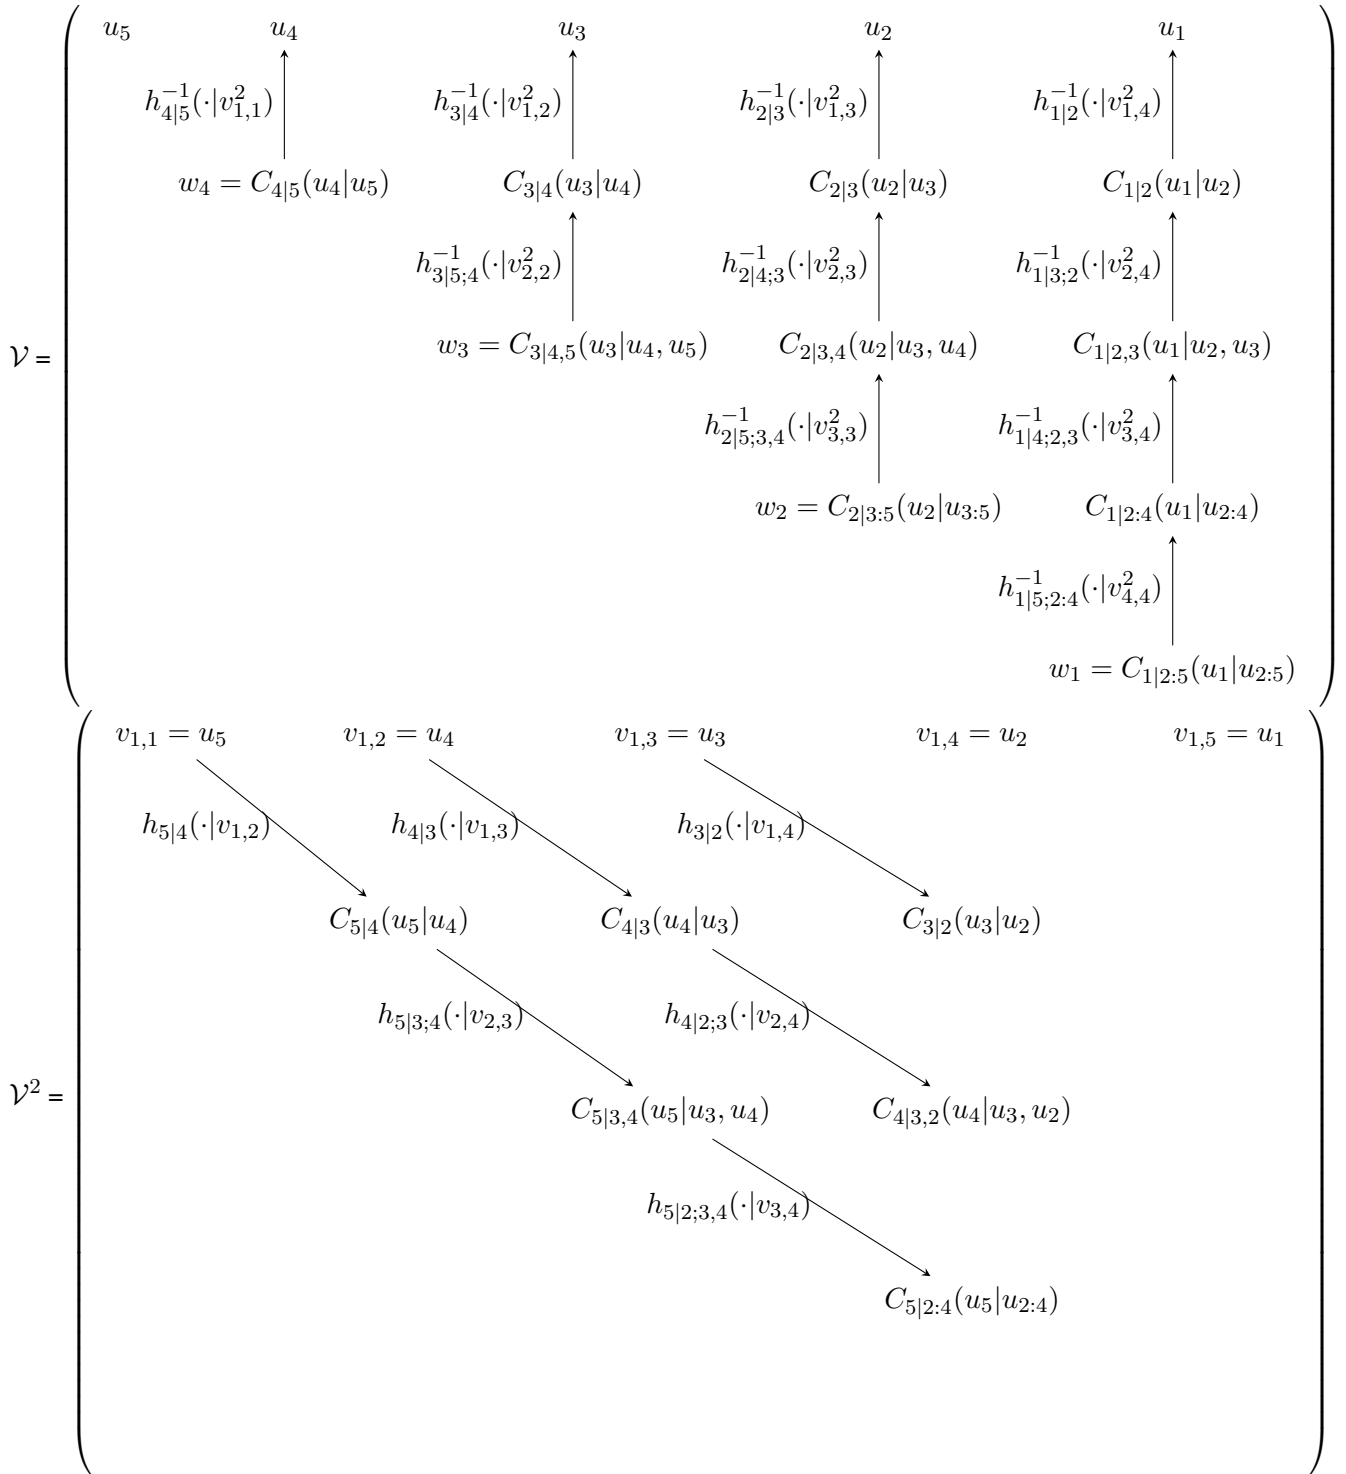


Figure 6.3 Show the build up of the matrices $\mathcal{V}, \mathcal{V}^2$ according to Algorithm 3 in a visual form for the 5 dimensional example.

So these detailed walk-throughs of Algorithm 3 display how the recursions build up the correct conditional distribution functions. This algorithm can be naturally extended to sample from the last n components, but it has to be noted that due to the curse of dimensionality the conditioning structure gets harder to grasp and to interpret.

6.4 Sampling from a D-vine copula conditional on the two rightmost leaf nodes

Here besides the case with one conditioning market index also the case with 2 market indices I_1, I_2 will be covered. This results in slightly modified algorithms compared to the ones given in Algorithms 2 and 3. The modified ordering scheme that has to be determined in order to fit the now $d + 2$ dimensional D-vine copula now not only has to select a permutation on the asset indices but also a permutation \tilde{j}_1, \tilde{j}_2 on the market index's indices 1, 2.

$$A_{j_d} - A_{j_{d-1}} - \dots - A_{j_1} - I_{\tilde{j}_1} - I_{\tilde{j}_2} \quad (6.8)$$

The interpretation of the ordering stays more or less the same to the single market index case. Again an approach based on the Pearson and partial correlation coefficients like presented in Algorithm 2 is proposed in Algorithm 4 in order to select a sensible ordering for the fitting of the D-vine copula.

Algorithm 4: Determine D-vine ordering including 2 market indices

Input : A_1, \dots, A_d and 2 market indices I_1, I_2 with corresponding i.i.d. copula data

$u_1, \dots, u_d, u^{I_1}, u^{I_2}$ as well as a cutoff depth $c_{\text{depth}} \in \{1, \dots, d + 1\}$ defaulting to $d + 1$

Output: j_1, \dots, j_d permutation for the indices $1, \dots, d$ and \tilde{j}_1, \tilde{j}_2 permutation for the market indices specifying an ordering as in (6.8)

Transform the copula data to the normalized scale using Equation (5.2). The data on the normalized scale is denoted by $z_1, \dots, z_d, z^{I_1}, z^{I_2}$ with the conventions $z_{j_0} = z^{I_1}, z_{j_{-1}} = z^{I_2}$ to simplify the algorithm below with respect to the notation.

$$\tilde{j}_2 = 3 - \tilde{j}_1$$

$$j_1, \tilde{j}_1 \leftarrow \delta \in \{1, \dots, d\}, k \in \{1, 2\} \text{ s.t. } |\rho_{\delta, I_k}| \geq |\rho_{\delta', I_{k'}}|, \forall \delta' \in \{1, \dots, d\}, k' \in \{1, 2\}$$

for $l \leftarrow 2$ **to** d **do**

$$j_l = \operatorname{argmax}_{\delta \in \{1, \dots, d\} \setminus \{j_1, \dots, j_{l-1}\}} \sum_{k=\max(l-c_{\text{depth}}, -1)}^{l-1} |\rho_{\delta, j_k; j_{k+1}:j_{l-1}}|$$

end

Here $\rho_{a,b;D}$ denotes the Pearson (if D is the empty set) or partial correlation coefficient between z_a and z_b and it is estimated always using the data on the normalized scale.

The next step is again to have an algorithm at hand to sample from now $C_{j_1, \dots, j_d | I_1, I_2}(\cdot | u^{I_1}, u^{I_2})$ for u^{I_1}, u^{I_2} being fixed realisations of the market indices I_1, I_2 on the copula-scale. For simplicity the market indices $I_{\tilde{j}_1}, I_{\tilde{j}_2}$ will be denoted by the $d + 1, d + 2$ component e.g. $u^{I_{\tilde{j}_1}} = u_{d+1}$. The to be constructed auxiliary matrices $\mathcal{V}, \mathcal{V}^2$ of course are just shifted so one has to add 1 to all the indices and extend them by one dimension, the formal algorithm is given as Algorithm 5.

Algorithm 5: Sampling from D-vine copulas conditional on the last 2 components

Input : Parameter matrix Θ for the $d+2$ dimensional D-vine and the fixed conditioning values

u_{d+2}, u_{d+1} for the last 2 components

Output: A d dimensional sample u_1, \dots, u_d from $C_{j_1, \dots, j_d | I_1, I_2}(\cdot | u^{I_1}, u^{I_2})$

Sample $w_j \stackrel{i.i.d.}{\sim} \mathcal{U}(0, 1), j \in \{1, \dots, d\}$

$v_{1,1} = u_{d+2}, v_{1,1}^2 = u_{d+2}$

$v_{1,2} = u_{d+1}, v_{1,2}^2 = u_{d+1}$

$v_{2,2}^2 = h_{d+2|d+1}(v_{1,1}^2 | v_{1,2})$

for $j \leftarrow 3$ **to** $d + 2$ **do**

$a = d + 3 - j$

$v_{j,j} = w_a$

for $k \leftarrow j - 1$ **to** 1 **do**

$v_{k,j} = h_{a|a+k;(a+1):(a+k-1)}^{-1}(v_{k+1,j} | v_{k,j-1}^2, \theta_{a|a+k;(a+1):(a+k-1)})$

if $j < d + 1$ **then**

$v_{k+1,j}^2 = h_{a+k|a;(a+1):(a+k-1)}(v_{k,j-1}^2 | v_{k,j}, \theta_{a+k|a;(a+1):(a+k-1)})$

end

$v_{1,j}^2 = v_{1,j}$

end

$u_i = v_{1,d+3-j}$ for $j \in \{1, \dots, d\}$

7 Vine copula based risk measure estimation

Now the overall conditional and unconditional estimation approaches for risk measures are presented whose integral parts were discussed in the previous sections. Before stating the algorithms in the final rolling window fashion first the one step ahead risk measure estimation procedures that are at the very basis of the rolling window approach are covered respectively.

7.1 Unconditional vine copula based risk measure estimation

First some additional notation is needed. Given a portfolio $\Omega = \{w_j, r_t^{A_j} | t = 1, \dots, T; j = 1, \dots, d\}$.

- $\Gamma < T \in \mathbb{N}$: **Length of the fitting window for the marginal models.** Here one assumes that this is also the time index of the most recent asset return observation when starting the rolling window approach. So for the first marginal time series model one uses all the available historic data. In Figure 7.2 the associated window to Γ is highlighted in blue.
- $\Psi \leq \Gamma \in \mathbb{N}$: **Length of the fitting window for the vine copula model** in order to capture dependencies.
- $S \in \mathbb{N}$: **Number of simulated log returns** for the risk measure estimation.

Figure 7.1 illustrates the flow of the data along the applied methods and transformations for the one step ahead unconditional risk measure estimation approach.

- Starting with the log returns of the d assets $r_t^{A_j}$ one first fits an adequate ARMA-GARCH model for each of the assets log return series. An ARMA(1,1)-GARCH(1,1) is probably often a good starting point as also mentioned in Tsay (2010) and with the quality assessment tools covered in Section 4.5 one can adjust the models to be adequate.
- The next step is to work with the fitted standardized residuals of the univariate models for the time frame of the vine training window. In order to model the dependence between assets using the now assumed i.i.d. standardized residuals one transforms them to the copula scale using the PIT with the fitted marginal residual distributions.
- The obtained copula scale data is then used to estimate a R-vine copula. Afterwards one can simulate S samples from this vine copula of course on the copula scale.
- In order to retransform the copula scale residuals to forecasted log returns of the assets one uses the inverse PIT and the forecasted mean and volatility estimates of the one step ahead time unit.
- Using the weights one can then aggregate the individual log return samples to portfolio level log return samples.
- From the resulting vector of size S it is then straightforward to derive risk measure estimates for the one step ahead time unit.

For this approach to be efficient not only a good modeling of the margins is needed but one also has the implicit assumption that the dependence structure assumed via the vine copula does not change for the upcoming one step ahead time unit. This one step approach can then be applied repeatedly to estimate

One step ahead unconditional risk measure estimation

from time unit Γ to $\Gamma + 1$




Figure 7.1 Illustration of the one step ahead unconditional risk measure estimation approach using ARMA-GARCH marginal and R-vine dependence models.

risk measures for the whole interval spanned by $\Gamma + 1$ and T , which is mostly required in backtesting




Figure 7.2 Illustration of the parameters Γ and Ψ alongside the one step ahead unconditional risk measure estimation.

applications. This repeated one step approach for the interval and the parameters Γ and Ψ is illustrated in Figure 7.2. Already for moderately large intervals of interest, like 100, a huge number of models has to be fitted which is a computational burden. Thus a rolling window approach is proposed in order to refit the models only after a certain time interval of usage. This of course might come to the price of potentially larger model misfits e.g. the dependence structure is assumed to stay constant over the whole interval. To extend the approach to a rolling window two additional parameters have to be specified.

- $\gamma \leq (T - \Gamma) \in \mathbb{N}$: **Length of the forecasting window of the marginal models.**
- $\kappa \leq \gamma \in \mathbb{N}$: **Length of the usage window of the vine copula model.**

It might be reasonable to pick them such that $(T - \Gamma) \bmod \gamma \equiv 0$ and $\gamma \bmod \kappa \equiv 0$ will have to hold. An illustration of the parameters and the rolling of the respective windows of the rolling window approach is given in Figure 7.3. Now the general approach for the unconditional risk measure estimation in a rolling




Figure 7.3 Illustration of the rolling parameters and the actual rolling of the respective windows of the rolling window approach.

window fashion is stated in Algorithm 6.

Algorithm 6: The rolling window unconditional risk measure estimation approach

Input : Portfolio Ω , parameters $\Gamma, \gamma, \Psi, \kappa, S$ and risk measure estimator $RM(\cdot|\alpha)$ at level α

Output: Risk measure estimates \hat{m}_t for $t \in \{\Gamma + 1, \dots, T\}$

First estimate the marginal models in a rolling window fashion.

for $marginal_window \leftarrow 1$ **to** $\lceil (T - \Gamma)/\gamma \rceil$ **do**

for $j \leftarrow 1$ **to** d **do**

 Fit an appropriate marginal model i.e. here an ARMA-GARCH model on

$$r_{1+\gamma*(marginal_window-1)}^{A_j}, \dots, r_{\Gamma+\gamma*(marginal_window-1)}^{A_j}$$

 Estimate from the model the conditional means and standard deviations given below.

$$\hat{\mu}_{1+\Gamma-\Psi+\gamma*(marginal_window-1)}^{A_j}, \dots, \hat{\mu}_{\min[T, \Gamma+\gamma*(marginal_window)]}^{A_j}$$

$$\hat{\sigma}_{1+\Gamma-\Psi+\gamma*(marginal_window-1)}^{A_j}, \dots, \hat{\sigma}_{\min[T, \Gamma+\gamma*(marginal_window)]}^{A_j}$$

 Calculate the standardized residuals using the formula $z_t^{A_j} = \frac{r_t^{A_j} - \hat{\mu}_t^{A_j}}{\hat{\sigma}_t^{A_j}}$ for the window in order

 to get $z_{1+\Gamma-\Psi+\gamma*(marginal_window-1)}^{A_j}, \dots, z_{\min[T, \Gamma+\gamma*(marginal_window)]}^{A_j}$

 They should, under the assumption of a correct marginal model, be approximate i.i.d. with marginal distribution function F_j . Convert those values on the residual scale to the copula scale by applying F_j to each of the values resulting in

$$u_{1+\Gamma-\Psi+\gamma*(marginal_window-1)}^{A_j}, \dots, u_{\min[T, \Gamma+\gamma*(marginal_window)]}^{A_j}$$

end

end

Estimate the dependence structure and the risk measures by simulation in a rolling window fashion.

for $vine_window \leftarrow 1$ **to** $\lceil (T - \Gamma)/\kappa \rceil$ **do**

 Use only $u_t^{A_j}, \hat{\mu}_t^{A_j}, \hat{\sigma}_t^{A_j}$ from the marginal window $\lceil \kappa * vine_window/\gamma \rceil$.

 Fit a regular vine copula $C_{1, \dots, d}$ using the copula data

$$u_{1+\Gamma-\Psi+\kappa*(vine_window-1)}^{A_j}, \dots, u_{\Gamma+\kappa*(vine_window-1)}^{A_j}, \text{ for } j \in \{1, \dots, d\}.$$

for $k \leftarrow 1$ **to** κ **if** $vine_window < \lceil (T - \Gamma)/\kappa \rceil \vee (T - \Gamma) \bmod \kappa \equiv 0$ **else to** $(T - \Gamma) \bmod \kappa$ **do**

 Simulate from $C_{1, \dots, d}$ multivariate samples $\hat{u}_s = (\hat{u}_s^{A_1}, \dots, \hat{u}_s^{A_d})$, for $s \in 1, \dots, S$

 Then transform the samples from copula \rightarrow residual \rightarrow return scale.

$$\hat{r}_{\Gamma+\kappa*(vine_window-1)+k}^{A_j, s} = \hat{\mu}_{\Gamma+\kappa*(vine_window-1)+k}^{A_j} + \hat{\sigma}_{\Gamma+\kappa*(vine_window-1)+k}^{A_j} * F_j^{-1}(\hat{u}_s^{A_j}), \text{ for } j \in \{1, \dots, d\} \text{ and } s \in \{1, \dots, S\}$$

 Get the weighted full portfolio return by a simple weighted sum.

$$\hat{r}_{\Gamma+\kappa*(vine_window-1)+k}^{\Omega, s} = \sum_{j=1}^d w_j * \hat{r}_{\Gamma+\kappa*(vine_window-1)+k}^{A_j, s}, \text{ for } s \in \{1, \dots, S\}$$

 Having all these samples one can estimate various risk measures.

$$\hat{m}_{\Gamma+\kappa*(vine_window-1)+k} = RM(\{\hat{r}_{\Gamma+\kappa*(vine_window-1)+k}^{\Omega, s} | s \in \{1, \dots, S\}\}|\alpha)$$

end

end

7.2 Conditional vine copula based risk measure estimation using a single market index as conditioning variable

For the single conditional variable risk measure estimation approach some further notation is required.

- **Market index** I that is observed over the same time scale as the portfolio Ω , the log returns are again denoted by r_t^I .
- Let j_1, \dots, j_d be a to be determined **permutation** of the indices of the assets $1, \dots, d$.
- $\alpha^I \in (0, 1)$: **Confidence level of the estimated quantile from the marginal market index distribution**. This confidence level is the conditioning value for the final risk measure estimate on the copula scale. How this fits into the estimation process can be found in detail below.

Actually the whole inference for margins approach that includes fitting the individual marginal models and transforming their standardized residuals to the copula-scale is almost identical in the conditional case with the only slight difference that the same procedures are applied to the market index I too. But then instead of a regular vine copula model a D-vine copula model with ordering $A_{j_d} - A_{j_{d-1}} - \dots - A_{j_1} - I$ is fitted. The ordering is determined as outlined in Algorithm 2. After that the conditional character of the approach starts to shine through as now one estimates a prespecified α^I level quantile of the marginal market index distribution on the copula scale. As on the copula scale the marginals are all uniform the quantile is actually given by the confidence level α^I itself. Then conditionally on that quantile one again simulates multivariate samples from $C_{j_1, \dots, j_d|I}(\cdot|\alpha^I)$. The retransformation to the return scale is then performed also on the conditioning value α^I as well as for all assets. As the sample is now conditional also the estimated risk measure is conditional on the market index evaluated at the respective estimated quantile on the return scale. This approach is illustrated for the one step ahead case in Figure 7.4 and formally written down in a rolling window fashion in Algorithm 7. Thus the resulting risk measure allows to analyze its robustness with respect to states of the overall or some specific market represented by I . These states are controlled via the α^I level. This might be applied to **stress testing** situations were one for example would like to assess the effect of a really bad performing market index I or asset on the risk estimates. This would then correspond in this setting to estimating conditional risk measures with a small α^I level. It might also be a more conservative risk measure for small α^I levels in the case of a generally positive dependence of the portfolio with the market index I . We will call this conditional strategy the **quantile based** one. Additionally to this approach based on a quantile level one can also use the copula scale residual of the index I from the prior time unit i.e. to estimate the risk measure at time unit t use u_{t-1}^I as the conditioning value at the copula scale. The only change to the quantile based approach is that one simulates conditionally from

$$C_{j_1, \dots, j_d|I}(\cdot|u_{t-1}^I) \quad (7.1)$$

and the forecasted conditioning value of I on the log return scale can be determined with

$$\hat{\mu}_t^I + \hat{\sigma}_t^I * F_I^{-1}(u_{t-1}^I). \quad (7.2)$$

The resulting estimated risk measure will emulate the behavior of conditioning on the predicted market index log return series and might be a valuable comparison to the estimated quantile based conditional risk measures. Opposed to the **quantile based conditional risk measures** we will call these **prior residual conditional risk measures**. It should be noted that the conditional series based on the fitted residuals of the time unit before will most likely exaggerate sudden high volatility situations. Imagine the univariate marginal time series model can not anticipate a sudden price drop of the conditioning asset at time $t - 1$. This leads to a very small copula scale residual u_{t-1}^I . By then conditioning at time t on this very harsh drop at the prior time unit one will probably in many situations exaggerate the drop of the estimated risk measures assuming without loss of generality a generally positive dependence of the portfolio with the market index I . In this case of a sudden drop this might even be beneficial by resulting in more conservative estimates at time t but the same applies to large copula scale residuals at time $t - 1$ in

One step ahead conditional risk measure estimation

from time unit Γ to $\Gamma + 1$




Figure 7.4 Illustration of the one step ahead single conditional risk measure estimation approach using ARMA-GARCH marginal and D-vine dependence models.

Algorithm 7: The rolling window conditional risk measure estimation approach using D-vines

Input : Portfolio Ω , market index I , parameters $\Gamma, \gamma, \Psi, \kappa, S, \alpha^I$ and risk measure estimator $RM(\cdot|\alpha)$ at level α

Output: Conditional risk measure estimates $\hat{m}_{t|I}$ for $t \in \{\Gamma + 1, \dots, T\}$

First estimate the marginal models in a rolling window fashion.

for *marginal_window* $\leftarrow 1$ **to** $\lceil (T - \Gamma)/\gamma \rceil$ **do**

for $j \leftarrow 1$ **to** d **do**

 Fit an appropriate marginal model i.e. here an ARMA-GARCH model on

$$r_{1+\gamma*(\text{marginal_window}-1)}^{A_j}, \dots, r_{\Gamma+\gamma*(\text{marginal_window}-1)}^{A_j}$$

 Estimate from the model the conditional means and standard deviations given below.

$$\hat{\mu}_{1+\Gamma-\Psi+\gamma*(\text{marginal_window}-1)}^{A_j}, \dots, \hat{\mu}_{\min[T, \Gamma+\gamma*(\text{marginal_window})]}^{A_j}$$

$$\hat{\sigma}_{1+\Gamma-\Psi+\gamma*(\text{marginal_window}-1)}^{A_j}, \dots, \hat{\sigma}_{\min[T, \Gamma+\gamma*(\text{marginal_window})]}^{A_j}$$

 Calculate the standardized residuals using the formula $z_t^{A_j} = \frac{r_t^{A_j} - \hat{\mu}_t^{A_j}}{\hat{\sigma}_t^{A_j}}$ for the window in order

 to get $z_{1+\Gamma-\Psi+\gamma*(\text{marginal_window}-1)}^{A_j}, \dots, z_{\min[T, \Gamma+\gamma*(\text{marginal_window})]}^{A_j}$

 As outlined in Algorithm 6 transform from residual to the copula scale by applying F_j to each of the values resulting in $u_{1+\Gamma-\Psi+\gamma*(\text{marginal_window}-1)}^{A_j}, \dots, u_{\min[T, \Gamma+\gamma*(\text{marginal_window})]}^{A_j}$

end

 Apply the same modeling procedure to the index returns r_t^I in order to estimate the $u_t^I, \hat{\mu}_t^I$ and $\hat{\sigma}_t^I$

end

Estimate the dependence structure and the risk measures by simulation in a rolling window fashion.

for *vine_window* $\leftarrow 1$ **to** $\lceil (T - \Gamma)/\kappa \rceil$ **do**

 Use only $u_t, \hat{\mu}_t, \hat{\sigma}_t$ from the marginal window $\lceil \kappa * \text{vine_window} / \gamma \rceil$.

 The copula data for the vine window is given by

$$u_{1+\Gamma-\Psi+\kappa*(\text{vine_window}-1)}^A, \dots, u_{\Gamma+\kappa*(\text{vine_window}-1)}^A, \text{ for } A \in \{A_1, \dots, A_d, I\}.$$

 Determine the ordering $A_{j_d} - A_{j_{d-1}} - \dots - A_{j_1} - I$ i.e. the permutation of indices j_1, \dots, j_d as outlined in Algorithm 2 in order to fit the D-vine.

 Fit the corresponding D-vine copula $C_{j_1, \dots, j_d, I}$ using the copula data and the ordering.

for $k \leftarrow 1$ **to** κ **if** *vine_window* $< \lceil (T - \Gamma)/\kappa \rceil \vee (T - \Gamma) \bmod \kappa \equiv 0$ **else to** $(T - \Gamma) \bmod \kappa$ **do**

 Simulate from $C_{j_1, \dots, j_d, I}(\cdot|\alpha^I)$ multivariate samples $\hat{u}_s = (\hat{u}_s^{A_1}, \dots, \hat{u}_s^{A_d})$, for $s \in \{1, \dots, S\}$

 Then transform the samples and the conditioning quantile α^I from copula \rightarrow residual \rightarrow return scale.

$$\hat{r}_{\Gamma+\kappa*(\text{vine_window}-1)+k}^{A_j, s} = \hat{\mu}_{\Gamma+\kappa*(\text{vine_window}-1)+k}^{A_j} + \hat{\sigma}_{\Gamma+\kappa*(\text{vine_window}-1)+k}^{A_j} * F_j^{-1}(\hat{u}_s^{A_j}), \text{ for } j \in \{1, \dots, d\} \text{ and } s \in \{1, \dots, S\}$$

$$\hat{r}_{\Gamma+\kappa*(\text{vine_window}-1)+k}^{I, \alpha^I} = \hat{\mu}_{\Gamma+\kappa*(\text{vine_window}-1)+k}^I + \hat{\sigma}_{\Gamma+\kappa*(\text{vine_window}-1)+k}^I * F_I^{-1}(\alpha^I)$$

 Get the weighted full portfolio return by a simple weighted sum.

$$\hat{r}_{\Gamma+\kappa*(\text{vine_window}-1)+k}^{\Omega, s} = \sum_{j=1}^d w_j * \hat{r}_{\Gamma+\kappa*(\text{vine_window}-1)+k}^{A_j, s}, \text{ for } s \in \{1, \dots, S\}$$

 Having all these samples one can estimate various conditional risk measures.

$$\hat{m}_{\Gamma+\kappa*(\text{vine_window}-1)+k|I=\hat{r}_{\Gamma+\kappa*(\text{vine_window}-1)+k}^{I, \alpha^I}} = RM(\{\hat{r}_{\Gamma+\kappa*(\text{vine_window}-1)+k}^{\Omega, s} | s \in \{1, \dots, S\}\}|\alpha)$$

end

end

the opposite way. Thus it has to be examined in practice how strong this volatility exaggeration might play out. Another non quantile based conditional risk measure estimation approach is the one calculating the **realized residual conditional risk measures** which are based on the actual observed residuals which means one conditions on the actually realized value of the conditional asset. The only change with respect to the quantile based approach is here that one simulates conditionally from

$$C_{j_1, \dots, j_d | I}(\cdot | F_I(\frac{r_t^I - \hat{\mu}_t^I}{\hat{\sigma}_t^I})) \quad (7.3)$$

in order to estimate the risk at time t , notably by using r_t^I which is a look in the future and not a forecast like $\hat{\mu}_t^I$ and $\hat{\sigma}_t^I$. The conditioning value on the log return scale is then exactly the realized log return of the conditional asset. This is an oracle estimator as it suffers from information leakage but might be a valuable comparison to the quantile based conditional risk measures.

7.3 Conditional vine copula based risk measure estimation using two market indices as conditioning variables

Extending Algorithm 7 to the case of two conditioning market indices I_1, I_2 is actually quite straight forward. Perform the inference for margins/ marginal models block in the algorithm equivalently for all assets and market indices. Then in each vine window determine the permutation indices of the assets j_1, \dots, j_d and the market indices \tilde{j}_1, \tilde{j}_2 that specify the ordering

$$A_{j_d} - \dots - A_{j_1} - I_{\tilde{j}_1} - I_{\tilde{j}_2}$$

for the $d + 2$ dimensional D-vine copula using Algorithm 4 and fit the D-vine. The only crucial change left is to get both conditioning values for the conditional sampling as outlined in Algorithm 5. This is done by setting the conditioning copula scale value of $I_{\tilde{j}_1}$ to α^I as in the single conditional approach and then to set the conditioning copula scale value of $I_{\tilde{j}_2}$ to $h_{d+2|d+1}^{-1}(\alpha^I | \alpha^I)$. This mimics the α^I level quantile of the bivariate copula between the market indices on the copula scale adjusting for their dependence. Then sample the asset returns conditioning on these sampled market indices according to Algorithm 5 and proceed as in the single conditional case. The prior residual and realized residual based conditional risk measures can be calculated in exactly the same manner as in the single conditional setting. This now allows for the stress testing of the risk measures with respect to two market indices.

8 Backtesting of the estimated risk measures

Now having estimated the risk measures in a rolling window fashion we want to assess the quality of the estimation process by comparing the estimated risk estimates for the portfolios log returns with the realized log returns. This assessment of the risk measurement procedure is not only crucial for regulators but also for internal evaluation of the methods used in the risk assessment and is mostly framed as statistical hypothesis test. This chapter will first cover and discuss the applicability to the presented estimation procedure of various traditional backtests that are framed like classical goodness of fit tests and then comparative backtests that provide tools for the comparison of different risk estimation procedures.

8.1 Traditional backtesting

The term traditional backtest here is adapted from Nolde and Ziegel (2017) and refers to a goodness of fit type hypothesis test with the null hypothesis:

$$H_0 : \text{"The risk measure estimation procedure is correct."}$$

So one considers the procedure to be adequate if one does not reject the null hypothesis at a certain confidence level. Like most goodness of fit tests such tests do not control for the power i.e. when not rejecting the null hypothesis one can not be sure that the null hypothesis is correct. In the following traditional backtests for the two discussed risk measures VaR and ES are covered.

Value at Risk backtesting

The de facto standard backtests for the Value at Risk (VaR) were introduced in Christoffersen (1998) and Christoffersen et al. (2001) and build upon the idea of so called **violation processes**. Given the VaR estimates $\{\hat{m}_{VaR\alpha}^t; t = \Gamma + 1, \dots, T\}$ at level α and the corresponding realized portfolio log returns $\{r_t; t = \Gamma + 1, \dots, T\}$ the violation process is defined as

$$I_t^\alpha = \mathbb{1}_{\{r_t < \hat{m}_{VaR\alpha}^t\}} \quad (8.1)$$

The first test is an unconditional likelihood ratio test in the sense of a traditional backtest also called **unconditional coverage test**. The unconditional here means that one assumes the violation process to have independent observations with respect to when they appear. This independence assumption will later be dropped when introducing a conditional approach. Nevertheless there is a place for unconditional coverage tests as pointed out in Acerbi and Székely (2014). They stress that the often preferred practice in the industry is a visual check for independence as well as violation clusters and then applying an unconditional coverage test. The hypothesis for a traditional backtest translates in this unconditional setting to the following Definition.

Definition 8.1.1 (LR test of unconditional coverage). *In order to test a violation process $\{I_t^\alpha; t = \Gamma + 1, \dots, T\}$ of VaR estimates calculated with confidence level α , that is assumed to arise independently, w.r.t. **unconditional coverage** the following hypotheses should be tested against each other.*

$$\begin{aligned} H_0 : \quad \mathbb{E}[I_t^\alpha] &= \alpha \\ H_1 : \quad \mathbb{E}[I_t^\alpha] &\neq \alpha \end{aligned}$$

The unconditional expected values used in these hypothesis explain the name of the test. The respective likelihoods then of course are derived from the binomial distribution and the binomial coefficient can be omitted due to the construction of a likelihood ratio test.

$$L_0(\alpha; I_{\Gamma+1}^\alpha, \dots, I_T^\alpha) = (1 - \alpha)^{\sum_{t=\Gamma+1}^T (1 - I_t^\alpha)} (\alpha)^{\sum_{t=\Gamma+1}^T I_t^\alpha}$$

$$L_1(\pi; I_{\Gamma+1}^\alpha, \dots, I_T^\alpha) = (1 - \pi)^{\sum_{t=\Gamma+1}^T (1 - I_t^\alpha)} (\pi)^{\sum_{t=\Gamma+1}^T I_t^\alpha}$$

The final test statistic along its asymptotic distribution is then given as

$$LR_{uc}^\alpha = -2 \log \left(\frac{L_0(\alpha; I_{\Gamma+1}^\alpha, \dots, I_T^\alpha)}{L_1(\hat{\pi}; I_{\Gamma+1}^\alpha, \dots, I_T^\alpha)} \right) \sim \chi^2(1) \quad (8.2)$$

with the maximum likelihood estimate $\hat{\pi} = \frac{1}{|\{I_t^\alpha; t=\Gamma+1, \dots, T\}|} \sum_{t=\Gamma+1}^T I_t^\alpha$.

Notably an equivalent test was stated by Kupiec as outlined in Roccioletti (2016). Obviously following from the independence assumption this test does not control for clustered violations with respect to the time they appear. In order to fix this potential issue one first states a likelihood ratio test of independence that will then be combined with the above test of unconditional coverage to get a conditional coverage test.

Definition 8.1.2 (LR test of independence). *Here the independence of a violation process $\{I_t^\alpha; t = \Gamma + 1, \dots, T\}$ resembles the null hypothesis which is tested against a first-order Markov alternative. Mathematically this translates to hypotheses in terms of transition probabilities*

$$H_0 : P(I_t^\alpha = 0 | I_{t-1}^\alpha = 1) = P(I_t^\alpha = 1 | I_{t-1}^\alpha = 1) \forall t$$

$$H_1 : P(I_t^\alpha = 0 | I_{t-1}^\alpha = 1) \neq P(I_t^\alpha = 1 | I_{t-1}^\alpha = 1) \forall t.$$

As a violation process has only the two states $\{0, 1\}$ the transition probability matrix from time $t - 1$ to t of the corresponding binary first-order Markov chain aligned with the alternative hypothesis is given as

$$\Pi_1 = \begin{pmatrix} 1 - \pi_{01} & \pi_{01} \\ 1 - \pi_{11} & \pi_{11} \end{pmatrix}$$

with $\pi_{ij} = P(I_t^\alpha = j | I_{t-1}^\alpha = i)$. This leads to the approximate likelihood of the alternative hypothesis

$$L_1(\Pi_1; I_{\Gamma+1}^\alpha, \dots, I_T^\alpha) = (1 - \pi_{01})^{n_{00}} \pi_{01}^{n_{01}} (1 - \pi_{11})^{n_{10}} \pi_{11}^{n_{11}}$$

with n_{ij} being the number of observed transitions from state i at time $t - 1$ to state j at time t . The corresponding maximum likelihood estimate of Π_1 is straightforwardly given as

$$\hat{\Pi}_1 = \begin{pmatrix} \frac{n_{00}}{n_{00} + n_{01}} & \frac{n_{01}}{n_{00} + n_{01}} \\ \frac{n_{10}}{n_{10} + n_{11}} & \frac{n_{11}}{n_{10} + n_{11}} \end{pmatrix}$$

Under the null hypothesis i.e. independence the probability transition matrix is parameterized by

$$\Pi_0 = \begin{pmatrix} 1 - \pi_0 & \pi_0 \\ 1 - \pi_0 & \pi_0 \end{pmatrix}$$

where $\pi_0 = P(I_t^\alpha = 0 | I_{t-1}^\alpha = 1) = P(I_t^\alpha = 1 | I_{t-1}^\alpha = 1)$ with corresponding likelihood

$$L_0(\Pi_0; I_{\Gamma+1}^\alpha, \dots, I_T^\alpha) = (1 - \pi_0)^{n_{00} + n_{10}} \pi_0^{n_{01} + n_{11}}$$

and maximum likelihood estimate $\hat{\pi}_0 = \frac{n_{01}+n_{11}}{n_{00}+n_{10}+n_{01}+n_{11}}$. The final test statistic along its asymptotic distribution is then given as

$$LR_{ind} = -2\log\left(\frac{L_0(\hat{\Pi}_0; I_{\Gamma+1}^\alpha, \dots, I_T^\alpha)}{L_1(\hat{\Pi}_1; I_{\Gamma+1}^\alpha, \dots, I_T^\alpha)}\right) \sim \chi^2(1) \quad (8.3)$$

As a remark one notes that this independence test does not depend on the confidence level α of the risk measure estimation. So it truly just tests for independence. The two LR type tests can now be framed as a single test for conditional coverage i.e. a test that both tests the independence of the VaR estimates and the raw number of violations.

Definition 8.1.3 (LR test for conditional coverage). *The previous two LR tests for unconditional coverage and independence can be used jointly to construct a likelihood ratio test for **conditional coverage** in the spirit of a traditional backtest. Again a violation process $\{I_t^\alpha; t = \Gamma+1, \dots, T\}$ of VaR estimates calculated with confidence level α is needed. The conditional coverage test statistic along its asymptotic distribution is then given as*

$$LR_{cc}^\alpha = LR_{uc}^\alpha + LR_{ind} \sim \chi^2(2) \quad (8.4)$$

Both the presented unconditional and conditional coverage tests are widely adapted in practice as they do just require the VaR estimates and the respective realizations and no additional information like the estimation method. This is especially important for regulators that get reported just these risk estimates. Furthermore they are quite easy to carry out.

Expected Shortfall backtesting

Of course by construction violation processes as the basis of backtests make no sense with the expected shortfall as the ES is not quantile based. Thus the backtesting of the ES needs a different approach and other than in the VaR setting not a predominant method has evolved. Therefore in the following different approaches are introduced and evaluated regarding their applicability to the estimation process outlined in this thesis.

Exceedance residual tests

The first backtest for the expected shortfall proposed in McNeil and Frey (2000) is focused on how the ES estimates compare with the realized log returns in the case of a violation with respect to the VaR. This way of thinking about the problem of backtesting is very natural when looking at the Definition of the ES in Equation (3.3). This idea is formalized by introducing the so called **exceedance residuals** given the ES estimates $\{\hat{m}_{ES_\alpha}^t; t = \Gamma+1, \dots, T\}$, VaR estimates $\{\hat{m}_{VaR_\alpha}^t; t = \Gamma+1, \dots, T\}$ at level α and the corresponding realized portfolio log returns $\{r_t; t = \Gamma+1, \dots, T\}$:

$$\{r_t - \hat{m}_{ES_\alpha}^t; t = \Gamma+1, \dots, T \text{ and } r_t < \hat{m}_{VaR_\alpha}^t\} \quad (8.5)$$

These exceedance residuals should under correct estimation of both risk measures have a zero mean. This leads to the following test definition. Here one interprets that $\hat{m}_{VaR_\alpha}^t$ and $\hat{m}_{ES_\alpha}^t$ are random quantities.

Definition 8.1.4 (Exceedance residual backtest). *Again in the spirit of a traditional backtest the **two sided exceedance residual test** tests the hypotheses*

$$\begin{aligned} H_0 &: \text{The exceedance residuals as in Equation (8.5) have zero mean.} \\ H_1 &: \text{Not } H_0 \end{aligned}$$

*In order to test this null hypothesis H_0 against the alternative H_1 one utilizes the distribution independent bootstrap test proposed on page 224 of Efron and Tibshirani (1994). From the regulators view however a two sided test is not necessary as a too conservative estimate and thus stronger capital requirements should not be punished. Therefore also a **one sided exceedance residual test** with the null hypothesis*

corresponding to the exceedance residuals having a mean greater or equal to zero can alternatively or additionally be performed. These tests can also be performed on the standardized exceedance residuals and McNeil and Frey report similar results for both approaches. Additionally the here discussed tests on the unstandardized exceedance residuals are also called **simple exceedance residual tests**.

As in the presented risk measure estimation algorithms of this thesis no portfolio level volatility estimates are computed the choice is easy as one can only use the exceedance residuals as defined in (8.5) and not the standardized ones. Moreover it has to be mentioned that this backtest uses additionally the VaR estimates. This is for the presented estimation process no problem but in a regulatory framework where the financial institution might only report the ES estimates the test is not applicable for the regulator. Moreover a bad estimation of the VaR could potentially harm the backtest of perfectly estimated ES estimates as the backtest is basically a joint one. This weakness of additional estimates is shared across many proposed ES backtesting procedures in the literature with some even calling for specific fixed estimation models. The next proposed backtest will also have this weakness.

Conditional calibration tests

For the discussion of conditional calibration tests following Nolde and Ziegel (2017) some additional notation is needed. A vector of risk measures ρ_k as defined in Definition 3.1.1 is denoted by

$$\boldsymbol{\rho} = (\rho_1, \dots, \rho_k) \quad (8.6)$$

Having $\boldsymbol{\rho}$ one can now define properties of such a vector of risk measures.

Definition 8.1.5 (Identifiable risk measures). *A vector of risk measures $\boldsymbol{\rho}$ is **identifiable** if and only if there exists a function $V : \mathbb{R}^k \times \mathbb{R} \rightarrow \mathbb{R}^k$ such that it holds for all $X \in \mathcal{R}$ (see Definition 3.1.1) that*

$$\mathbb{E}[V(m, X)] = 0 \Leftrightarrow m = \boldsymbol{\rho}(X) \quad (8.7)$$

V is then called the **identification function**.

One can think of a identification function for a certain risk measure as one that returns negative values for violations of the measure and positive ones if no violation occurs in a weighted fashion such that for the theoretical risk measure an expected value of 0 is achieved. For the VaR at confidence level α one has identifiability for the identification function

$$V(m, x) = \alpha - \mathbb{1}_{\{x < m\}} \quad (8.8)$$

and the reader may spot the conceptual similarity with the null hypothesis of the unconditional coverage test in Definition 8.1.1. For the vector (VaR, ES) both at confidence level α one has identifiability for the identification function

$$V((m_{VaR}, m_{ES}), x) = \left(m_{ES} - m_{VaR} - \frac{1}{\alpha} \mathbb{1}_{\{x < m_{VaR}\}} (m_{VaR} - x) \right) \quad (8.9)$$

Note that again like in the exceedance residual tests the ES is not treated on its own but jointly with the VaR this comes from the fact that on its own the ES is not identifiable as proven in Fissler and Ziegel (2016). These identification functions then allow the formal definition of calibration.

Definition 8.1.6 (Calibrated risk measures). *A sequence of estimated risk measures $\{\hat{m}_t; t \in \mathbb{N}\}$ for X_t is called **calibrated for $\boldsymbol{\rho}$ on average** if $\forall t \in \mathbb{N}$ it holds almost surely that*

$$\mathbb{E}[V(\hat{m}_t, X_t)] = 0 \quad (8.10)$$

and **conditionally calibrated for $\boldsymbol{\rho}$** if it holds almost surely $\forall t \in \mathbb{N}$ that

$$\mathbb{E}[V(\hat{m}_t, X_t) | \mathcal{F}_{t-1}] = 0 \quad (8.11)$$

where \mathcal{F}_{t-1} is the σ -algebra that encompasses all the information at hand at time $t - 1$ and ϱ is identifiable with V .

The notion of conditional calibration is a stronger one as due to the tower property of conditional expectations the conditional calibration implies calibration on average while the converse is generally not true. Nolde and Ziegel (2017) thus argue that in a dynamic risk management context the construction of backtests upon the stronger notion of conditional calibration is reasonable. This leads to the following backtest for conditional calibration that again is in the spirit of a traditional backtest.

Definition 8.1.7 (Conditional calibration backtest). *For a k dimensional vector of risk measures ϱ and a fixed estimation procedure one can formulate the test hypotheses for **conditional calibration** as*

$$\begin{aligned} H_0 : & \text{ The sequence of risk measure estimates is conditionally calibrated for } \varrho \\ H_1 : & \text{ Not } H_0 \end{aligned}$$

From the fact that

$$\mathbb{E}[V(\hat{m}_t, X_t) | \mathcal{F}_{t-1}] = 0 \Leftrightarrow \mathbb{E}[H_t^\top V(\hat{m}_t, X_t)] = 0$$

holds almost surely for all \mathcal{F}_{t-1} -measurable functions H_t that are $\mathbb{R}^{k \times q}$ valued one can construct a test statistic in the spirit of a classical Wald test. The \mathcal{F}_{t-1} -measurable function H_t is also called the test function. For estimated risk measures \hat{m}_t in the interval $t \in \{1, \dots, n\}$, corresponding realizations x_t and identification function V the general test statistic is of the form

$$\mathcal{T}_{cc} = n \left(\frac{1}{n} \sum_{t=1}^n H_t^\top V(\hat{m}_t, x_t) \right)^\top \hat{\Sigma}_n \left(\frac{1}{n} \sum_{t=1}^n H_t^\top V(\hat{m}_t, x_t) \right) \quad (8.12)$$

with

$$\hat{\Sigma}_n = \frac{1}{n} \sum_{t=1}^n [H_t^\top V(\hat{m}_t, x_t)][H_t^\top V(\hat{m}_t, x_t)]^\top \quad (8.13)$$

Under the null hypothesis H_0 one has under mild regularity conditions that

$$\mathcal{T}_{cc} \xrightarrow{\mathcal{D}} \chi^2(q) \text{ as } n \rightarrow \infty$$

which allows to construct an asymptotic test.

For the choice $H_t = 1$ one calls the resulting tests **simple conditional calibration tests**. Other choices of test functions will not be considered here as the choice is yet somewhat subjective. Nolde and Ziegel even stress that a misspecification of H_t could lead to a slow convergence to the asymptotic test distribution. Moreover they often require volatility estimates that as discussed before are not available at a portfolio level for the here presented estimation method.

Notably these conditional calibration tests are closely related to the backtests that were covered so far. The simple conditional calibration test of the VaR translates also to a test of the number of exceedances like the unconditional coverage test by Christoffersen. Likewise for a certain test function H_t the conditional calibration test is almost identical to the exceedance residual test proposed by McNeil and Frey. As the proposed conditional calibration tests are two sided tests again from a regulators perspective a one sided test that does not punish arbitrary conservative estimates is desirable. In this case one aims to test for the so called **conditional super-calibration** i.e. $\mathbb{E}[V(\hat{m}_t, X_t) | \mathcal{F}_{t-1}] \geq 0$. The construction of these one sided tests is in line with the two sided tests above. The only difference is that one tests each dimension of the risk measure vector ϱ individually which leads to the problem of multiple testing. To deal with this either the classical Bonferroni or the Hommel correction is applied as proposed in Nolde and Ziegel (2017). As already mentioned the weakness of joint testing for the ES still holds here but the next class of ES backtests that will be covered actually does not suffer from this weakness.

ES regression tests

Quite recently Bayer and Dimitriadis (2020b) have introduced the first ES backtests that solely rely on estimates for the ES and no auxiliary information like VaR estimates that are not mandatorily reported to the regulators. This overcomes the major weakness of the backtests for the ES presented so far. They use so called **regression based testing**. In particular they regress the actual realizations X_t on the expected shortfall estimates \hat{m}_{ES}^t at confidence level α

$$X_t = \gamma_1 + \gamma_2 \hat{m}_{ES}^t + \epsilon_t \quad (8.14)$$

with $ES_\alpha^t(\epsilon_t | \mathcal{F}_{t-1}) = 0$ almost surely. \mathcal{F}_{t-1} again denotes the σ -algebra containing the history up to $t - 1$. For correctly specified risk estimates the intercept parameter γ_1 should be 0 and the slope parameter γ_2 should be 1 almost surely. The naive way of proceeding would be to fit this simple linear regression model and to test these parameters via a Wald-test. But here again it comes into play that there is no identification function as covered in Definition 8.1.5 for the ES alone. This makes maximum likelihood as well as generalized methods of moments estimation infeasible in a stand alone manner. But based on the joint identification functions that were introduced in Fissler and Ziegel (2016) for the pair (VaR, ES) Bayer and Dimitriadis (2020b) outline a joint regression technique that extends Equation (8.14) with the second regression

$$X_t = \beta_1 + \beta_2 \hat{m}_{VaR}^t + \epsilon_t^* \quad (8.15)$$

with $VaR_\alpha^t(\epsilon_t^* | \mathcal{F}_{t-1}) = 0$ almost surely. This allows the joint estimation of all the parameters $\gamma_1, \gamma_2, \beta_1, \beta_2$. Note that here still VaR estimates are needed which is generally undesirable. Thus the first proposed backtest is called the **auxiliary expected shortfall regression (ESR) backtest** as it still requires auxiliary information.

Definition 8.1.8 (Auxiliary ESR backtest). *The setup is the regression system specified by Equations (8.14) and (8.15) and all their respective variables. Due to the joint identifiability of the VaR and ES joint parameter estimation is feasible and one tests for the correct specification of the ES Equation (8.14) via the hypotheses*

$$\begin{aligned} H_0 : (\gamma_1, \gamma_2) &= (0, 1) \text{ and arbitrary } \beta_1, \beta_2 \\ H_1 : (\gamma_1, \gamma_2) &\neq (0, 1) \text{ and arbitrary } \beta_1, \beta_2 \end{aligned}$$

The test is then a classical Wald-test with statistic

$$\mathcal{T}_{AESR} = n((\hat{\gamma}_1, \hat{\gamma}_2) - (0, 1)) \hat{\Sigma}^{-1} ((\hat{\gamma}_1, \hat{\gamma}_2) - (0, 1))' \quad (8.16)$$

and asymptotic covariance estimator $\hat{\Sigma}$ for the reduced parameter vector (γ_1, γ_2) . $\hat{\Sigma}$ is estimated via the joint regression framework for VaR and ES by Dimitriadis and Bayer (2019) and its involved structure as well as its asymptotic theory is described in detail in Section 1.4 of Bayer and Dimitriadis (2020b). Here n denotes the sample size.

Now in order to get rid of the dependence on the VaR estimates they are just substituted by the respective ES estimates which leads to the so called **strict ESR backtest**.

Definition 8.1.9 (Strict ESR backtest). *The setup is the regression system specified by the Equation (8.14) and*

$$X_t = \beta_1 + \beta_2 \hat{m}_{ES}^t + \epsilon_t^* \quad (8.17)$$

and all their respective variables. Again estimation is performed like in the auxiliary ESR backtest and one tests for the correct specification of the ES Equation (8.14) via the same hypotheses

$$H_0 : (\gamma_1, \gamma_2) = (0, 1) \text{ and arbitrary } \beta_1, \beta_2$$

$$H_1 : (\gamma_1, \gamma_2) \neq (0, 1) \text{ and arbitrary } \beta_1, \beta_2$$

The test is then once again a classical Wald-test with statistic

$$\mathcal{T}_{SESR} = n((\hat{\gamma}_1, \hat{\gamma}_2) - (0, 1))\hat{\Sigma}^{-1}((\hat{\gamma}_1, \hat{\gamma}_2) - (0, 1))' \quad (8.18)$$

and asymptotic covariance estimator $\hat{\Sigma}$ for the reduced parameter vector (γ_1, γ_2) . $\hat{\Sigma}$ is estimated via the joint regression framework for VaR and ES by Dimitriadis and Bayer (2019) and its involved structure as well as its asymptotic theory is described in detail in Section 1.4 of Bayer and Dimitriadis (2020b). Here n denotes the sample size.

Note that through the usage of ES estimates in Equation (8.17) potential misspecifications if the VaR and ES are not colinear can arise. However Bayer and Dimitriadis (2020b) provide asymptotic theory for this misspecification case that assures consistency of the parameter estimates under mild regularity conditions. This is achieved by posing constraints on the covariance structure during estimation. The strict ESR backtest still has one shortcoming which is that it is a two sided test and as worked out for the presented ES backtests a one sided test is very interesting for the regulator that does not want to punish conservative risk estimation. Therefore the third and final ESR test enters the stage.

Definition 8.1.10 (One sided intercept ESR backtest). *As the name **one sided intercept ESR backtest** suggests one regresses the risk estimation errors on just an intercept parameter in the primary ES regression. This leads to the following regression setup*

$$X_t - \hat{m}_{ES}^t = \beta_1 + \beta_2 \hat{m}_{ES}^t + \epsilon_t^* \quad (8.19)$$

$$X_t - \hat{m}_{ES}^t = \gamma_1 + \epsilon_t \quad (8.20)$$

with $ES_\alpha^t(\epsilon_t | \mathcal{F}_{t-1}) = 0$ and $VaR_\alpha^t(\epsilon_t^* | \mathcal{F}_{t-1}) = 0$ almost surely. Having this setup one can define two sided as well as one sided hypotheses

$$H_0^{\text{two sided}} : \gamma_1 = 0 \text{ and arbitrary } \beta_1, \beta_2$$

$$H_0^{\text{one sided}} : \gamma_1 \geq 0 \text{ and arbitrary } \beta_1, \beta_2$$

which are tested via t-tests based on the estimated covariance against their alternatives.

Note that this approach is equivalent to just fixing the slope parameter of the strict ESR backtest to one. Bayer and Dimitriadis (2020b) verify the performance of these ESR backtests in an extensive simulation study. All three of these ESR backtests are applicable to the estimation approach presented in this thesis as no additional information, like volatility estimates, except for the risk estimates is needed for the backtests. This makes these backtests very appealing for portfolio level risk estimates as well as for the regulator. This concludes the tour of some state of the art ES backtests in the spirit of traditional backtests and the discussion on whether they are applicable in this particular setting.

8.2 Comparative backtesting

In order to compare the performance of certain estimation procedures traditional backtests framed like goodness of fit tests do not suffice. Thus comparative backtests following Nolde and Ziegel (2017) will be introduced now. At the core of such tests are the so called **scoring functions**.

Definition 8.2.1 (Consistent scoring functions). *A mapping $S : \mathbb{R}^q \times \mathbb{R} \rightarrow \mathbb{R}$ is called scoring function. If for the scoring function and the k dimensional vector of risk measures ϱ it holds for all $X \in \mathcal{R}$ (see Definition 3.1.1) and for all $\mathbf{m} = (m_1, \dots, m_k) \neq \varrho(X)$ that*

$$\mathbb{E}[S(\varrho(X), X)] < \mathbb{E}[S(\mathbf{m}, X)] \quad (8.21)$$

*then S is **strictly consistent**. If such a strictly consistent scoring function exists for a vector of risk measures ϱ one calls ϱ **elicitable**.*

While the Value at Risk is itself elicitable the expected shortfall is not. But the combination of (VaR, ES) is indeed elicitable as first shown by Fissler and Ziegel (2016). Thus again in order to perform comparative backtesting for the ES in the sense of Nolde and Ziegel (2017) one has to use the joint workaround. One can think of a strictly consistent scoring function as one that is minimized for an optimal estimation procedure. Thus having such a scoring function one can directly compare different risk estimation approaches in absolute terms. Thereby the procedure with the smallest mean scoring value is considered the best estimation procedure. The choice of the scoring function is however not unique. In the following the strictly consistent suggestions of Nolde and Ziegel (2017) will be presented. For the VaR at confidence level α the scoring function

$$S(m, x) = (\mathbb{1}_{\{x < m\}} - \alpha)m - \mathbb{1}_{\{x < m\}}x \quad (8.22)$$

was proposed. This piecewise linear scoring function is well known under the name **pinball loss** and is heavily used in quantile regression settings. As the VaR is basically a quantile forecast this fits nicely also in this context. It basically returns the weighted distance to the estimated quantile where violations of the estimate are weighted stronger for small confidence levels α . For the pair (VaR, ES) the following scoring function is suggested in Nolde and Ziegel (2017):

$$S((m_{VaR}, m_{ES}), x) = \mathbb{1}_{\{x < m_{VaR}\}} \frac{x - m_{VaR}}{m_{ES}} + \alpha \left(\frac{m_{VaR}}{m_{ES}} - 1 + \log(-m_{ES}) \right) \quad (8.23)$$

This of course always assumes negative risk estimates. For a comparative test one needs two estimation procedures here called the **standard and internal model**. The aim is to challenge the standard model with the internal model. Their risk measure estimates are denoted as $\hat{\eta}^*$ and $\hat{\eta}$ respectively. The next step is then to state the comparative backtest according to Nolde and Ziegel (2017).

Definition 8.2.2 (Comparative backtest). *Having the internal model challenging the standard model for a risk measure estimation task of an elicitable vector of risk measures ϱ one can state the two null hypotheses for **comparative backtesting**.*

H_0^- : *The internal model estimates the risk measures at least as well as the standard one.*

H_0^+ : *The internal model estimates the risk measures at most as well as the standard one.*

In this comparative backtesting setting H_0^- is again a goodness of fit test that one does not want to reject. In order to control the type I error i.e. preferring an inferior internal model over the standard model, H_0^+ is introduced which one then aims to reject. Nolde and Ziegel (2017) worked out that under mild regularity assumptions both these hypotheses can be asymptotically tested with the following test statistic for n observed risk estimates for each model.

$$\mathcal{T}_{comp} = \frac{\Delta_n \bar{S}}{\hat{\sigma}_n / \sqrt{n}} \xrightarrow{\mathcal{D}} \mathcal{N}(0, 1) \quad (8.24)$$

$$\Delta_n \bar{S} := \frac{1}{n} \sum_{t=1}^n [S(\hat{\eta}_t, x_t) - S(\hat{\eta}_t^*, x_t)] \quad (8.25)$$




Figure 8.1 Three zone approach for comparative backtesting with confidence level η as in Nolde and Ziegel (2017). The challenging model is assigned to one of the three regions *Passed* (green), *Failed* (red) or *Further investigation needed* (yellow) according to the test statistic given in Equation (8.24).

where $\hat{\sigma}_n$ is the square root of the empirical variance estimate of the vector of score function differences that is summed up in Equation (8.25). H_0^- is rejected at confidence level η when $1 - \Phi(\mathcal{T}_{comp}) \leq \eta$ and H_0^+ when $\Phi(\mathcal{T}_{comp}) \leq \eta$.

Testing these hypotheses simultaneously leads to the **three zone approach** displayed in Figure 8.1. So while the *Passed* (H_0^+ is rejected) and *Failed* (H_0^- is rejected) regions suggest that there is a superior or inferior estimation approach at a certain level of confidence the *Further investigation needed* region (no hypothesis can be rejected) allows the interpretation as the situation that lacks evidence for a definitive decision at the respective level of confidence. As outlined above the choice of the scoring function that is used at the core of the comparative backtest is not unique. In the practical part of this thesis the ones for the VaR and the pair (VaR, ES) given in Equations (8.22) and (8.23) are used. Again the joint treatment of VaR and ES is a weakness as a perfectly good ES estimation procedure could be not identified because of a bad estimation of the VaR. This comparative backtesting strategy is not intended to be used instead of traditional backtests that are especially interesting from a regulators standpoint but should be a tool for the assessment of new estimation techniques in comparison with existing approaches. Thus when backtesting a new model one should first have a look into traditional backtesting and when passing these tests one can move on and challenge existing or alternative models. This closes the theoretical part of the thesis and now the discussed theory as well as the proposed methods are put to the test in practice in thorough case studies.

Part II

Case studies

9 R package portvine

9.1 Introduction

In order to perform the estimation techniques that were proposed in this thesis we have developed the *portvine* package for the statistical computing language R, R Core Team (2021). The thoroughly tested and documented package was build on the shoulders of giants. Most importantly one has to name the packages *rvinecopulib* of Nagler and Vatter (2021) and *rugarch* of Ghalanos (2022). These facilitate the effective estimation of vine copula and ARMA-GARCH models respectively. The package is publicly available on GitHub and will be submitted to CRAN. It furthermore features extensive documentation with the main API documentation at <https://emanuelsommer.github.io/portvine/> and all package related references are given there. For the further analysis of the risk measure estimates derived from the *portvine* package mainly packages from the *tidyverse* from Wickham et al. (2019) were used for data wrangling and visualizations. For the traditional backtesting of the ES we will use the *esback* R package of Bayer and Dimitriadis (2020a). Notably all the source code for the upcoming analyses and the theoretical visualizations presented up to now are also publicly available on GitHub at <https://github.com/EmanuelSommer/PortvineThesis>. Before starting with the case studies we will have a look at the performance of the *portvine* package.

9.2 Performance assessment

As evident in the Algorithms 6 and 7 the unconditional as well as the conditional estimation approach is highly parallelizable. This should definitely be exploited on modern tech stacks that often allow the use of many cores at once. In particular the *portvine* package even allows for **nested parallel processing** i.e. two levels of parallel processes. These levels are the following.

1. The first level allows to fit the univariate time series models for each asset in parallel and to fit each vine window in parallel.
2. The second level allows if specified that within each vine window the time units, here in the case studies days, are processed using parallel processes. This is illustrated in Figure 9.1.

From here on we will use the notation p_1/p_2 for the possible parallel strategies where p_1 specifies the number of parallel processes/cores on the first level and p_2 the number of parallel processes/cores every first level process can in turn call. This means 1/1 is the default sequential case, 5/1 does only use the first level parallelization and 8/2 would require 16 cores as every first level parallel process can use 2 parallel processes while iterating over the time units/days. An illustration of the nested parallelization of the vine windows can be found in Figure 9.1.

Moreover the conditional sampling algorithms specified in Algorithms 3 and 5 were implemented in the low level programming language C++ which reduced the runtime by a factor of roughly 30 and memory consumption by a factor of roughly 200. This is important as especially for large sampling sizes (the parameter S) the runtime is mainly driven by the sampling.

In the following some measurements that drastically show the significance of using the power of parallelization are shown. The measurements were performed on a Linux Cluster of the Leibniz Supercomputing Centre. The provided computational resources are gratefully acknowledged. The used R version is 3.6.0. For details on the extensive possibilities to set up the parallel strategy on your tech stack have a look at the package documentation. For all following performance measures we have the fixed and realistic parameters.

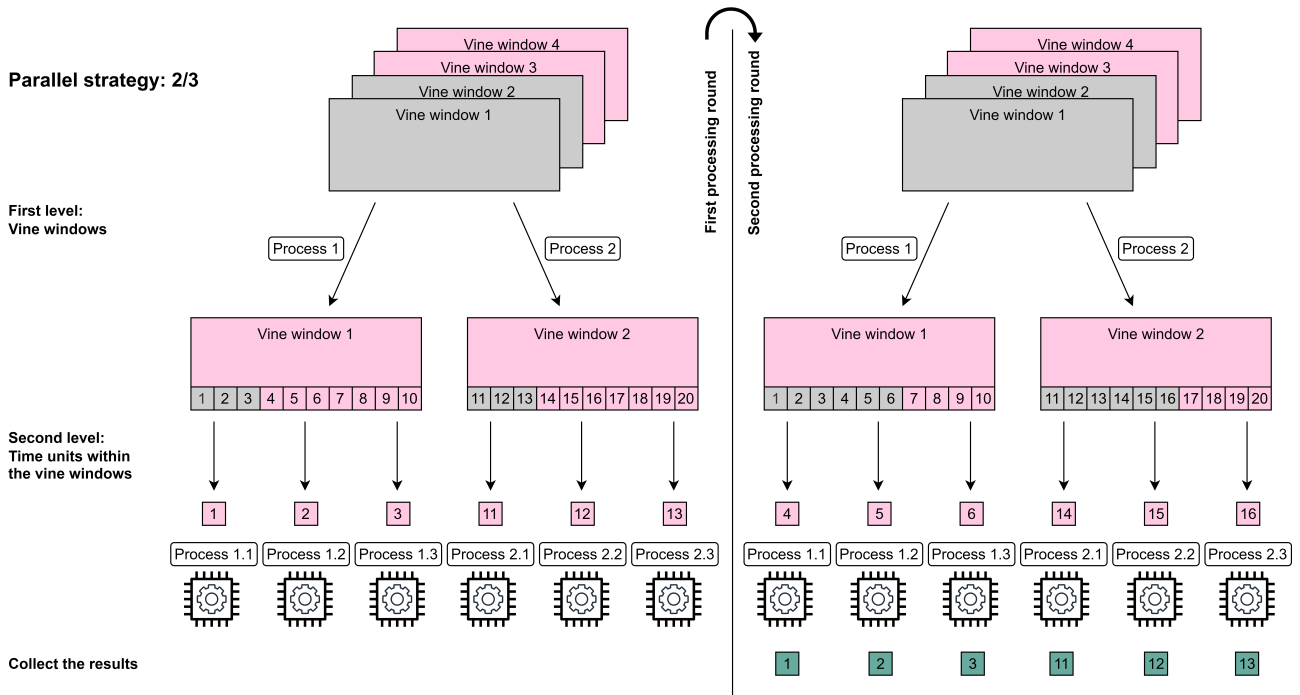


Figure 9.1 Illustration of the first two processing rounds of the 2/3 nested parallel strategy with regard to the vine windows. Here for illustrative purposes one assumes 4 vine windows with the size of 10 time units.

- The data set contains 1000 observations so roughly 4 financial years i.e. $T = 1000$.
- The marginal time series training windows size Γ is set to 750.
- The forecasting window of the marginal time series has size $\gamma = 50$, which leads to 5 marginal windows.
- The vine copula training windows size Ψ is set to 200.
- The vine copula usage window has size $\kappa = 25$, which leads to 10 vine windows.
- We estimate the VaR and the expected shortfall with mean estimation as specified in Equation (3.5).
- We estimate these risk measures at the common confidence levels α 1% and 2.5%.

Thus in this setting the highest reasonable parallel strategy is 10/25 as then one would parallelize each vine window and each time unit/day within these windows. This of course requires quite a strong setup that can handle 250 parallel processes and the extensive memory requirements that come with this level of parallelization.

Performance of the unconditional approach

Firstly in order to explore the influence of different simulation sample sizes S we further fix the number of variables in the data set to 10 and use all available parametric bivariate copula families for the fitting process. This influence on the runtime is important as in practice we would like to have really large number of samples of the portfolios log returns because the risk measure is then estimated upon this sample. Figure 9.2 and the corresponding Table 9.1 display the runtimes of different simulation sample sizes and the respectively used parallel strategy. Especially for 100000 samples the use of 40 cores with 10/4 instead of 20 cores with 10/2 pays off almost linearly. The maximum strategy of 10/25 achieves a great runtime of under 7 minutes. For the reasonable simulation sample size of 10000 we see that with already 20 cores we can achieve a quite quick runtime of roughly 5 minutes. Also it seems reasonable to assign

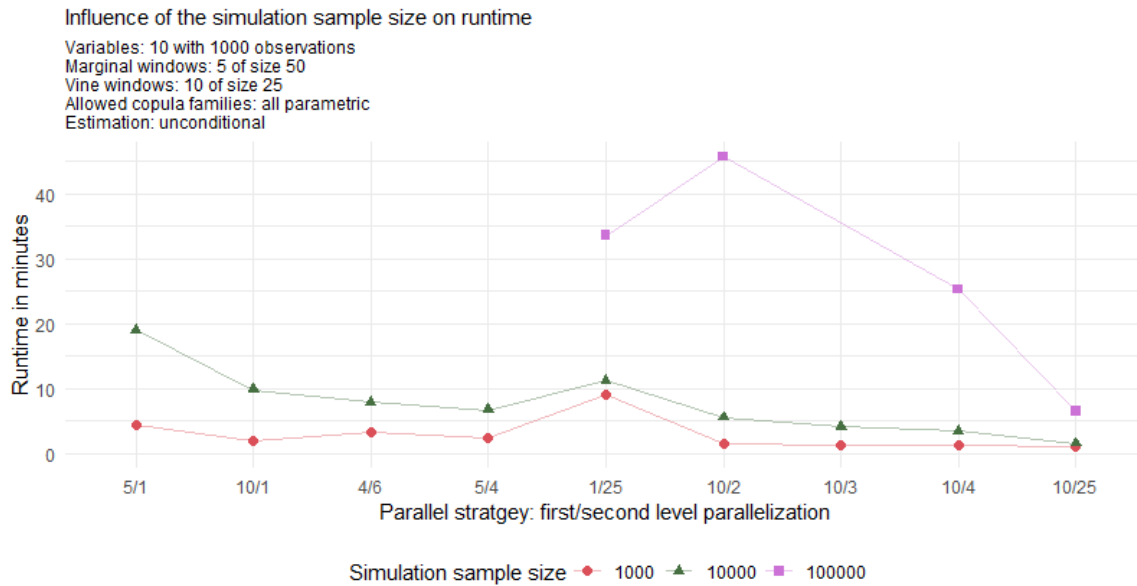


Figure 9.2 The influence of different simulation sample sizes and the respectively used parallel strategy on the runtime of the unconditional risk measure estimation implemented in the portvine package.

Simulation samples S	Parallel strategy									
	5/1	10/1	4/6	5/4	1/25	10/2	10/3	10/4	10/25	
1000	4.33	2.07	3.38	2.47	9.12	1.56	1.42	1.35	1.11	
10000	18.98	9.82	7.94	6.72	11.31	5.60	4.20	3.54	1.57	
100000	-	-	-	-	33.57	45.66	-	25.26	6.58	

Table 9.1 The influence of different simulation sample sizes and the respectively used parallel strategy on the **runtime given in minutes** of the unconditional risk measure estimation implemented in the portvine package. The number of variables is fixed to 10 and all parametric bivariate copula families are allowed for the fitting.

as many parallel processes as possible up to the number of vine windows to the first level of parallelization and then to assign as many additional processes to the second level of parallelization. This however changes for really large simulation sample sizes where it can be seen that assigning lots of processes to the second level might pay off. This is evident when looking at the strategy 1/25.

The next comparison which now explores the influence of including more variables in the data set fixes the simulation sample size $S = 10000$. The results indicate an almost linear behavior between the number of variables and the runtimes are given in Figure 9.3. The results are not surprising as the computational effort is obviously larger for more variables in the data set.

The last comparison in the unconditional case will be the one of different sets of bivariate copulas that are allowed. The simulation sample size S is again fixed to 10000 and the number of variables to 10. The results are depicted in Figure 9.4. We can definitely detect a reduced runtime by using only single parametric copula families, which are here given by the Gaussian, Frank, Joe, Clayton and Gumbel bivariate copulas. By using only the Gaussian or Student's t bivariate copulas it is indicated that we can reduce the runtime even further.

Table 9.2 gives a quick overview of some of the most interesting scenarios. We see that this unconditional approach can be easily scaled up to large portfolios using a modern set up.

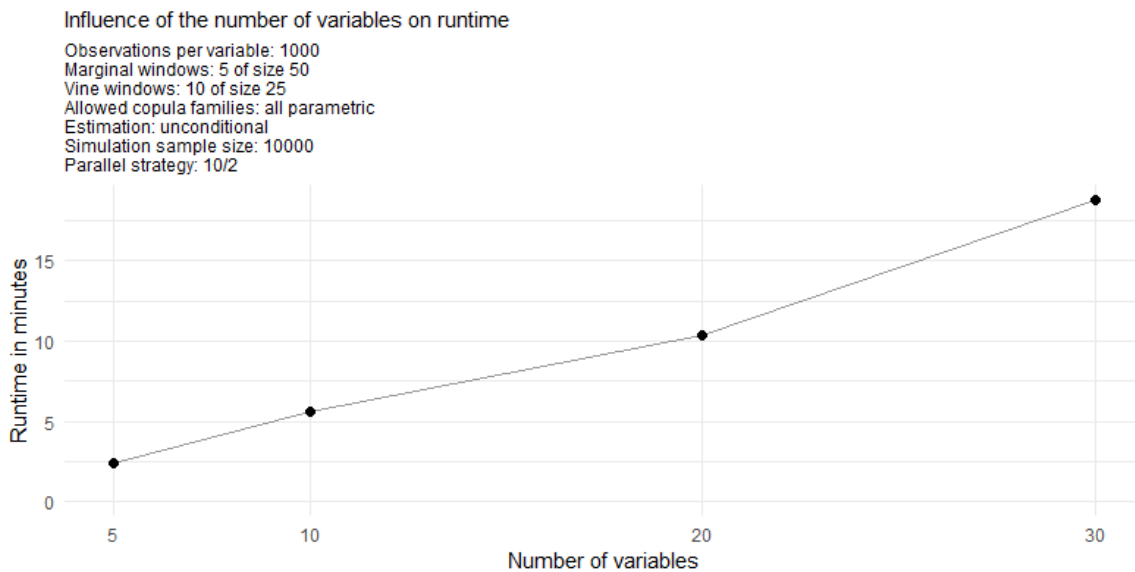


Figure 9.3 The influence of different numbers of variables and the respectively used parallel strategy on the runtime of the unconditional risk measure estimation implemented in the portvine package.

Number of variables d	Simulation sample size S	Cores used	Runtime (minutes)
10	10000	20	5.60
10	100000	20	45.66
10	10000	250	1.57
10	100000	250	6.58
30	10000	20	18.80

Table 9.2 Overview of the runtime of some of the most interesting scenarios from the performance assessment of the **unconditional** risk measure estimation approach. In all displayed cases all parametric copula families were allowed, the observations per variable were fixed to $T = 1000$ and 5 marginal windows of size 50 as well as 10 vine windows of size 25 were used.

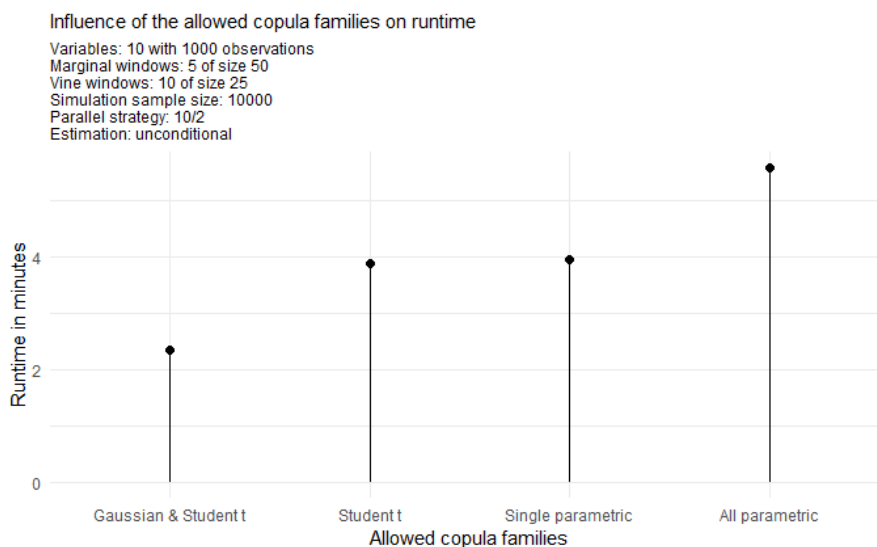


Figure 9.4 The influence of different sets of bivariate copula families and the respectively used parallel strategy on the runtime of the unconditional risk measure estimation implemented in the portvine package.

Performance of the conditional approach

As the conditional approach performs almost identical for a single or two conditional variables we display just the performance measures of the single conditional case. Furthermore we will estimate at the two quantile levels α^I 5% and 50%. The conditional risk measure estimation approach is expected to have longer runtimes mostly due to the fact that for each specified quantile level α^I and the additional prior residual risk estimates one has to sample S conditional samples. For the here specified two quantile levels this means that one samples 3 times as many conditional samples as unconditional samples in the unconditional risk estimation approach. Moreover the overhead of the conditional estimation is slightly higher. These expectations are indeed met in the experiments. Again we start with the exploration of the influence that is induced by different simulation sample sizes S on the runtime. For this the number of variables is fixed to 10 and all parametric bivariate copula families are allowed for the D-vine fitting. Figure 9.5 and the corresponding Table 9.3 display the runtimes of different simulation sample sizes and the respectively used parallel strategy. As annotated on the plot the observation for the maximal strategy

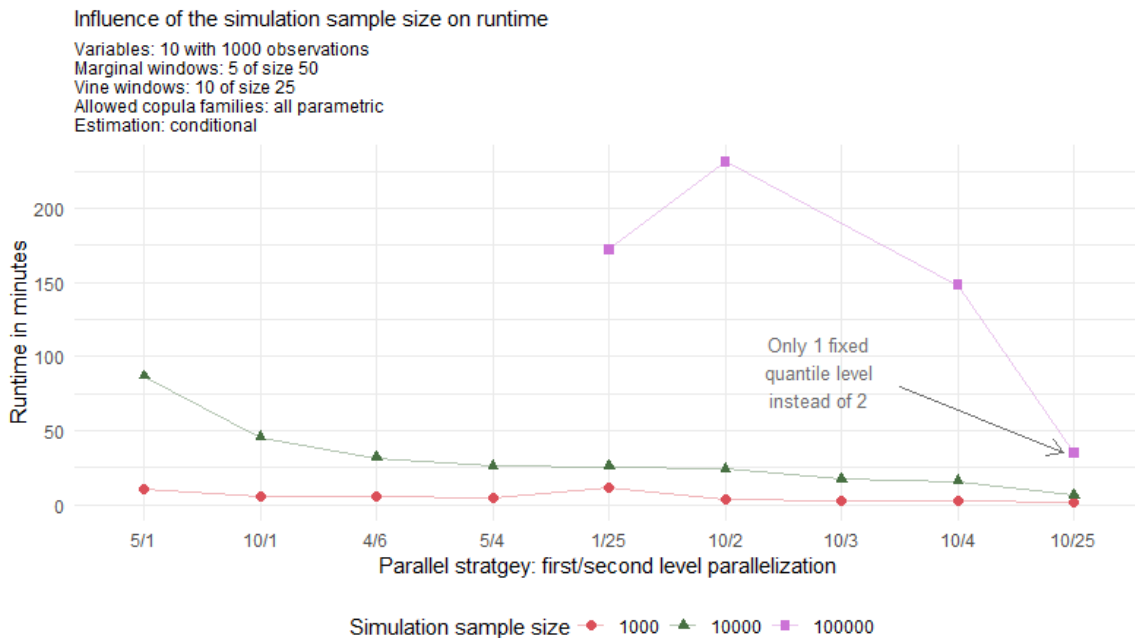


Figure 9.5 The influence of different simulation sample sizes and the respectively used parallel strategy on the runtime of the conditional risk measure estimation implemented in the portvine package. The observation for the strategy 10/25 and simulation sample size of 100000 corresponds to the estimation of only 1 conditional quantile as 2 conditional quantiles would come with the extensive memory requirements of >100Gb. From this fact we directly expect this observation to be roughly 33% faster.

10/25 and simulation sample size of 100000 corresponds to the estimation of only 1 conditional quantile as 2 conditional quantiles would come with the extensive memory requirements of >100Gb. From this fact we directly expect this estimation to be roughly 33% faster. But even after adjusting for this expected 33% faster result we observe a drop of another 44% in runtime for this particular parallel strategy. To get a better sense for the smaller simulation sample sizes 10000 and 1000 the filtered visualization is given in Figure 9.6. Once again the time saving from the slowest runtime which is roughly 1.5 hours to the fastest runtime which is less than 7 minutes is clearly visible.

The next comparison which now again explores the influence of including more variables in the data set fixes the simulation sample size $S = 10000$. The results of this comparison are given in Figure 9.7. The results are again not surprising as the computational effort is obviously larger for more variables in the data set.

Simulation samples S	Parallel strategy								
	5/1	10/1	4/6	5/4	1/25	10/2	10/3	10/4	10/25
1000	10.86	5.65	5.85	4.36	10.88	3.48	2.67	2.61	1.57
10000	86.77	45.37	31.51	25.71	25.55	24.11	16.97	15.80	6.22
100000	-	-	-	-	171.93	231.20	-	148.02	34.67

Table 9.3 The influence of different simulation sample sizes and the respectively used parallel strategy on the **runtime given in minutes** of the conditional risk measure estimation implemented in the portvine package. The number of variables is fixed to 10 and all parametric bivariate copula families are allowed for the fitting. The entry for the strategy 10/25 and simulation sample size of 100000 corresponds to the estimation of only 1 conditional quantile as 2 conditional quantiles would come with the extensive memory requirements of >100Gb. From this fact we directly expect this entry to be roughly 33% smaller.

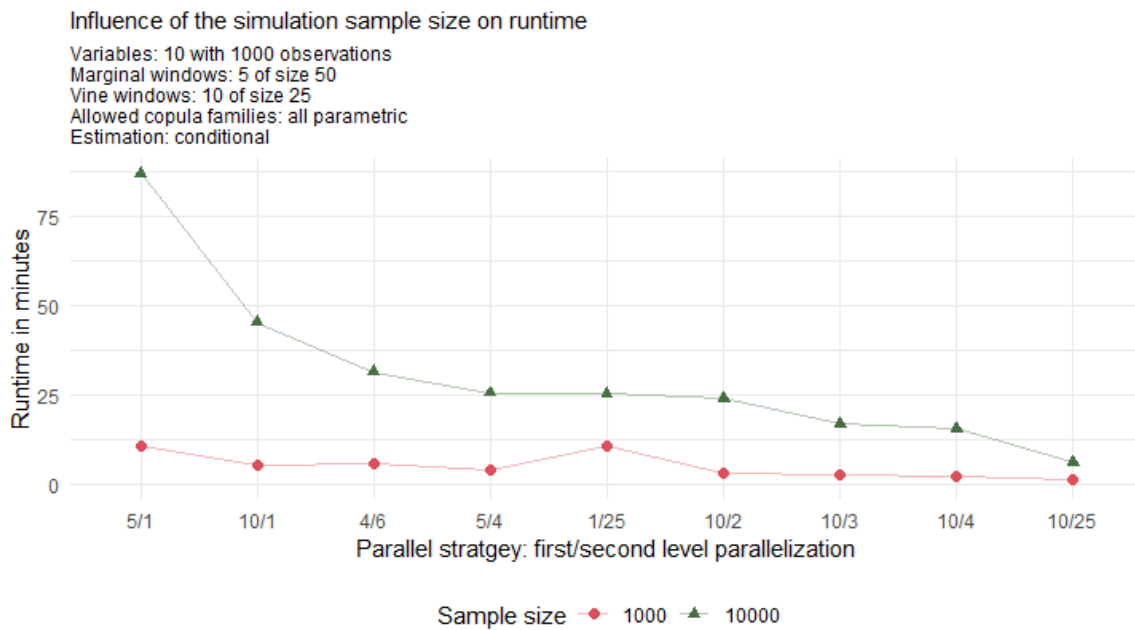


Figure 9.6 The influence of the two simulation sample sizes 1000 and 10000 and the respectively used parallel strategy on the runtime of the conditional risk measure estimation implemented in the portvine package.

In the conditional case it makes sense to also examine the influence of more quantile levels α^I that are to be estimated on the runtime. From the framing of the algorithm it should have a linear influence. This expectation is supported by Figure 9.8.

The last comparison will be again the one of different sets of bivariate copulas that are allowed. The simulation sample size S is again fixed to 10000 and the number of variables to 10. The results are depicted in Figure 9.9. We can clearly see an almost identical pattern as in the unconditional case.

Table 9.4 again gives a quick overview of some of the most interesting scenarios. We see due to the fact that the conditional approach requires more samples one needs a stronger hardware setup in order to apply the approach with high simulation sample sizes of >10000 on large portfolios.

All in all we can conclude that by enabling really flexible parallelization the *portvine* package is able to save a lot of runtime even on reasonably sized tech stacks with for example 20 parallel processes but can also be used on high performance computing infrastructure. This enables to base the risk measure estimation of large portfolios on large sample sizes which is desirable for the practical application.

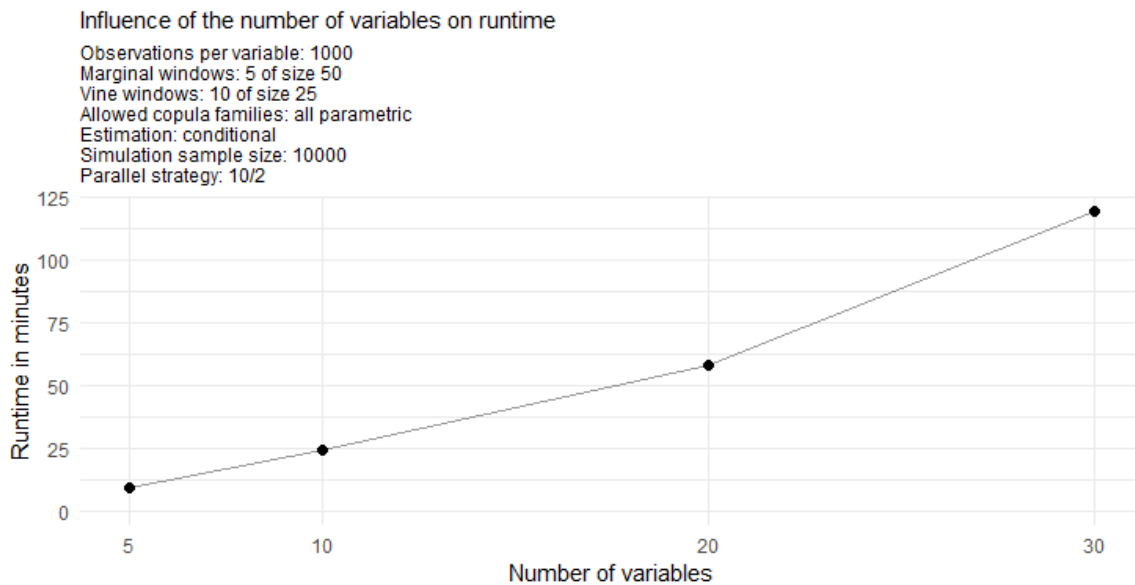


Figure 9.7 The influence of different numbers of variables and the respectively used parallel strategy on the runtime of the conditional risk measure estimation implemented in the portvine package.

Number of variables d	Simulation sample size S	Cores used	Runtime (minutes)
10	10000	20	24.11
10	100000	20	231.20
10	10000	250	6.22
10	100000	250	34.67 (only 1 quantile level)
30	10000	20	119.25

Table 9.4 Overview of the runtime of some of the most interesting scenarios from the performance assessment of the **conditional** risk measure estimation approach. In all displayed cases all parametric copula families were allowed, the observations per variable were fixed to $T = 1000$ and 5 marginal windows of size 50 as well as 10 vine windows of size 25 were used. Besides the case in the fourth scenario, which would require memory of $>100\text{Gb}$ for two quantile levels, always two quantile levels α^I were estimated.

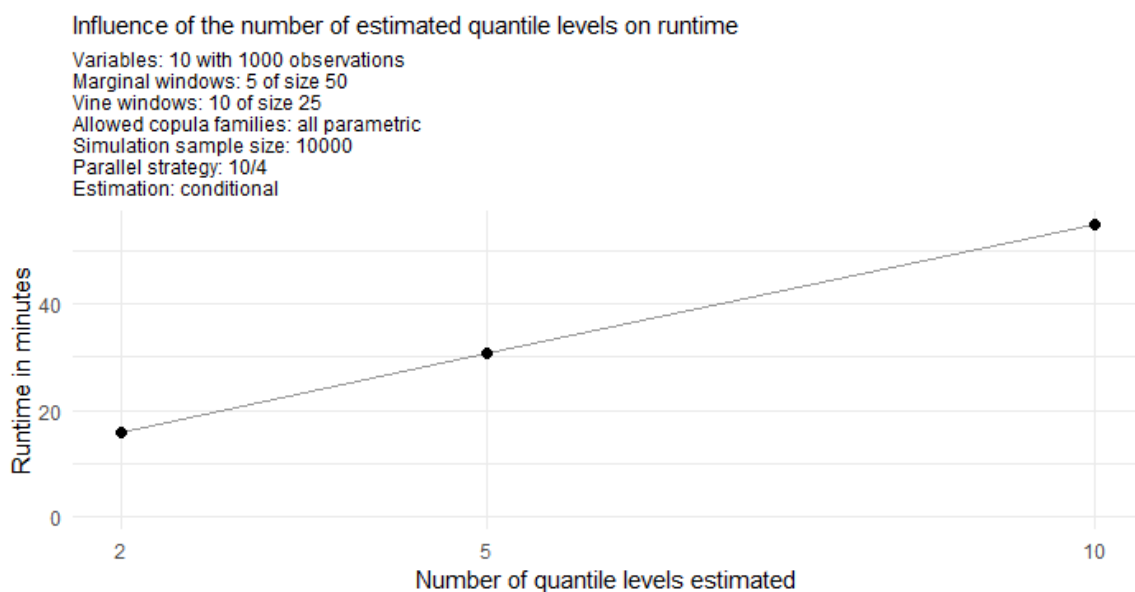


Figure 9.8 The influence of different numbers of quantile levels α^I that are to be estimated and the respectively used parallel strategy on the runtime of the conditional risk measure estimation implemented in the portvine package.




Figure 9.9 The influence of different sets of bivariate copula families and the respectively used parallel strategy on the runtime of the conditional risk measure estimation implemented in the portvine package.

10 Unconditional Spanish stock portfolio analysis

In this chapter we will first present the portfolio of Spanish stocks that will be used for the case studies as well as the two time frames of interest in Section 10.1. After that we will assess the quality of the fitted marginal models in Section 10.2. The next step is then to analyze the fitted R-vines in Section 10.3. Section 10.4 provides a comparison of the three different estimators for the ES that were presented in Chapter 3. Thereafter we visualize and backtest the estimated risk measures in Section 10.5. We first perform the traditional backtests on all models considering various specifications right after the initial visualizations of the risk measures in Section 10.5 and then compare the different models with respect to different vine training parameters and the set of used bivariate copula families for the D-vine. This is done towards the end of Section 10.5 using comparative backtesting with a focus on the ES.

10.1 The setup

At the core of all following analysis will be a portfolio of nine Spanish stocks i.e. $d = 9$. These are the nine top constituents of the MSCI Spain stock market index. The nine assets along their corresponding weights w_j are given below in Table 10.1. The weight corresponds to the respective market capitalization (EUR billions) of the asset on the 29th October 2021.

Asset	Iberdrola	Banco Santander	BBV Argentaria	Inditex	Cellnex Telecom
Weight	60.48	56.82	40.42	34.08	27.09
Asset	Amadeus It Group	Telefonica	Repsol Ypf	Ferrovial	
Weight	26.06	19.37	16.04	12.99	

Table 10.1 The assets of the Spanish stock portfolio of interest with their respective market capitalization (EUR Billions) on the 29.10.2021.

We will have a look at the following two time frames.

- **Time frame 1:** 01.01.2016 - 31.12.2019, $T = 1042$ daily log returns and a training window of $\Gamma = 750$.
- **Time frame 2:** 02.04.2020 - 13.10.2021, $T = 400$ daily log returns and a training window of $\Gamma = 300$.

They are both highlighted on the time series plot of the daily log returns of the MSCI Spain in Figure 10.1. We can observe greater volatility during the stock market selloffs 2015-2016 maybe due to Chinese stock market turbulence, the EU debt crisis and the Brexit votum as well as for the first 'Covid-19 year' 2020. The second time frame with the pandemic in place allows backtest our estimation approach in a higher volatility setting.

The time series of the individual assets of the portfolio along side the one of the MSCI Spain are given in Figure 10.2 and indeed they look stationary and volatility clusters can be observed. Thus an approach using only ARMA models would probably not perform well. The forecasting window size for the marginal time series models is fixed to $\gamma = 50$ which corresponds to roughly 2 financial months. This results in six rolling marginal windows i.e. refits in the first time frame and in two rolling marginal windows in the second time frame. The default marginal model is an ARMA(1,1)-GARCH(1,1) model with the residual white noise distribution being a skewed Student's t distribution. The assessment of the quality of marginal model fit follows right after this section. For the parameters specifying the repeated R-vine fitting and refitting four different specifications were examined. These are all possible combinations of the two parameters

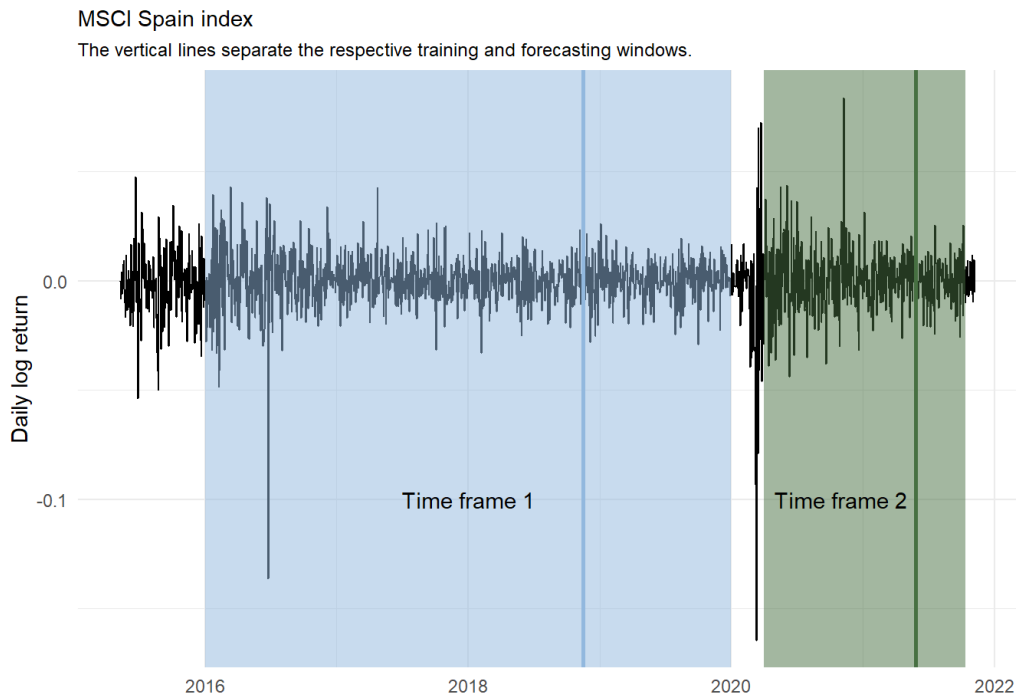


Figure 10.1 The daily log return series of the MSCI Spain with the time frames of interest highlighted.

$\Psi \in \{250, 500\}$ and $\kappa \in \{25, 50\}$ for the first time frame and $\Psi \in \{100, 200\}$ and $\kappa \in \{25, 50\}$ for the second time frame. We will try to determine the best specification later via comparative backtesting. For the R-vine fitting all parametric bivariate copula families implemented in the *rvinecopulib* package were allowed. A comparison with models restricted to only Gaussian or Student's t components will be performed later also via comparative backtesting. For all models the VaR and the ES were estimated for the α levels 0.01, 0.025, 0.05. For the ES all presented estimators in Chapter 3 were used and a comparison of the estimators will also be given later. Moreover all estimated risk measures were based on a realistic simulation sample size of $S = 100000$ samples.





Figure 10.2 The daily log return series of the MSCI Spain and the individual assets.


10.2 Assessment of the marginal models

We will use the tools presented in Section 4.5 in order to assess the quality of the marginal models. The first time frame with 6 marginal models corresponding to the marginal windows for each asset will be analyzed first. It is actually quite realistic to only have a look at the residual analysis of the first marginal model as at the beginning of the rolling window approach only this analysis would be possible. Here only the residual plots corresponding to the asset Iberdrola will be displayed in Figure 10.4. All other residual plots for the first time frame and first marginal window can be found in the Appendix A. Overall all residual plots were not indicating any model misfit. To have a better overview over all assets at once for a certain marginal window one can have a look at the heatmaps displayed in Figure 10.4 that show the results of the Ljung-Box tests for different lags and all considered assets at once. Once again no concerning misfit can be detected for any of the six marginal windows. Thus we will stick to the default model for all marginal models for the first time frame.

The same analysis is then also performed on the second time frame. The exemplary residual plots for the asset Iberdrola are displayed in Figure 10.5 and the Ljung-Box heatmaps for both marginal windows in Figure 10.6. Again the other residual plots are collected in the Appendix A. All in all, for the second time frame no severe model misfit was detected with the fit of the series corresponding to the Banco Santander asset, whose residual plots can be found in Figure A.9, being the most questionable. As it is a close call we will also for the second time frame stick to the default ARMA(1,1)-GARCH(1,1) model in all cases. Thus, the next step is to have a look at the fitted R-vine copulas.




(a) Residual plots for the mean equation.




(b) Residual plots for the volatility equation.

Figure 10.3 The residual plots for the asset Iberdrola for the first marginal training window (01.01.2016 - 15.11.2018) and the first time frame (2016-2019).




(a) The first rolling window.




(b) The last rolling window.

Figure 10.4 The Ljung-Box heatmaps on the standardized residuals for the first time frame (2016-2019) and the first and last marginal training window (01.01.2016 - 15.11.2018 and 16.12.2016 - 31.10.2019).




(a) Residual plots for the mean equation.




(b) Residual plots for the volatility equation.

Figure 10.5 The residual plots for the asset Iberdrola for the first marginal training window (02.04.2020 - 26.05.2021) and the second time frame (2020-2021).



(a) The first rolling window.



(b) The last rolling window.

Figure 10.6 The Ljung-Box heatmaps on the standardized residuals for the second time frame (2020-2021) and the first and last marginal training window (02.04.2020 - 26.05.2021 and 11.06.2020 - 04.08.2021).

10.3 The fitted R-vines

Now we analyze the fitted R-vines. The most interesting questions that arise in this context are which bivariate building blocks were fitted in which structure and how strong their dependence is. Also it might be interesting to detect changing dependence patterns over time. This can be facilitated by visualizing the first tree level. As a large number of R-vines was fitted for all models in this section we will focus only on one model per time frame.

First time frame (2016-2019) Again one starts with the first time frame and as the model we choose the one with $\Psi = 250$ and $\kappa = 25$. Thus 12 R-vine copula models were fitted and Figure 10.7 displays the first tree of the first, sixth and twelfth R-vine. The fitted bivariate copula family and the empirical Kendall's tau are displayed on each edge.

The strongest bivariate dependence is the one between the two banks Banco Santander and the BBV Argentaria. These two assets and the Telefonica stock seem to have a central role for the given portfolio. Moreover, Table 10.2 shows that quite a significant proportion of components were non Gaussian.

Copula family	Number of usage	Copula family	Number of usage
Gaussian	3	Independence	7
Student's t	2	Clayton	5
Frank	12	Gumbel	5
Joe	1	BB7	1

(a) The first vine window.


Copula family	Number of usage	Copula family	Number of usage
Gaussian	3	Independence	7
Student's t	5	Clayton	5
Frank	8	Gumbel	6
Joe	1	BB8	1

(b) The sixth vine window.

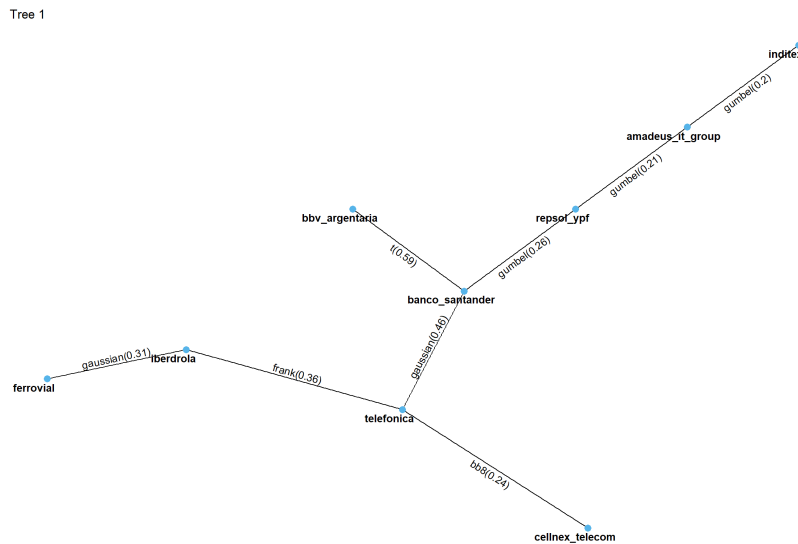
Copula family	Number of usage	Copula family	Number of usage
Gaussian	6	Independence	9
Student's t	5	Clayton	1
Frank	4	Gumbel	5
Joe	5	BB8	1

(c) The twelfth vine window.


Table 10.2 The used copula families alongside their frequencies for the first, sixth and twelfth fitted R-vine for the first time frame (2016-2019) in the unconditional setting.



(a) The first vine window.




(b) The sixth vine window .




(c) The twelfth vine window.

Figure 10.7 The first trees of the R-vines corresponding to the first, sixth and twelfth vine window for the first time frame (2016-2019) in the unconditional setting.

Second time frame (2020-2021) For the second time frame, we choose the model with $\Psi = 200$ and $\kappa = 25$. This leads to 4 fitted R-vine copulas. The first trees of the first and fourth R-vine copula are displayed in Figure 10.8. The central assets are the same as in the first time frame.



(a) The first vine window.



(b) The sixth vine window .

Figure 10.8 The first trees of the R-vines corresponding to the first and fourth vine window for the second time frame (2020-2021) in the unconditional setting.

Table 10.3 shows that for the second time frame a lot more copulas with a focus on modelling tail dependence e.g. the BB copulas were used. This might be connected to the generally higher volatility for this time frame.

Copula family	Number of usage	Copula family	Number of usage
Gaussian	1	Independence	8
Student's t	3	Clayton	4
Frank	4	Gumbel	7
Joe	3	BB1	3
BB7	2	BB8	1

(a) The first vine window.

Copula family	Number of usage	Copula family	Number of usage
Gaussian	5	Independence	9
Student's t	5	BB1	2
Frank	5	Gumbel	5
Joe	3	BB8	2

(b) The fourth vine window.

Table 10.3 The used copula families alongside their frequencies for the first and fourth fitted R-vine for the second time frame (2020-2021) in the unconditional setting.

All in all, one can detect for both time frames that the refitting of the dependence model is indeed necessary as the structure and composition of the fitted R-vines changes over time. An adequate frequency of the refitting can be determined using comparative backtesting. Also we detected a lot of non Gaussian or Student's t components especially for the second time frame so the comparison of this more flexible class of R-vine allowing for all parametric copula families with the restricted classes only allowing Gaussian and/or Student's t components is of interest. This comparison will also be performed using the comparative backtests.

10.4 Expected Shortfall estimators

As we present three different estimators for the ES namely the mean, median and Monte Carlo integration based estimators in Chapter 3 a comparison of the resulting estimates seems reasonable.

First time frame (2016-2019) Again we start with the first time frame and have only a look at one of the four fitted models. Here again the one with the parameters $\Psi = 250$ and $\kappa = 25$ is used. Figure 10.9 shows how the three estimators for the ES at confidence level $\alpha = 0.05$ behave over time. The Monte Carlo integration based estimator is based on 10000 Monte Carlo samples and the resulting ES estimates are almost identical to the ones corresponding to the mean estimator. Only the median based estimator differs from the other two as it is always less conservative than the two. As the VaR exceedances are apparently left skewed the median is more robust resulting in less conservative estimates. This left skew of the VaR exceedances is most likely the usual case in practice. Figure 10.10 shows that overall the three






Figure 10.9 Comparison of the three ES estimators at confidence level $\alpha = 0.05$ over time for the first time frame (2016-2019).

estimators are almost colinear as we can roughly observe an identity line.



(a)



(b)

Figure 10.10 Two views of the 3d plot relating the risk estimates of the three different estimators for the ES at confidence level $\alpha = 0.05$ for the first time frame (2016-2019).

Second time frame (2020-2021) Figure 10.11 displays how the three estimators for the ES at confidence level $\alpha = 0.05$ behave over time for the second time frame. Here again the model with parameters $\Psi = 200$ and $\kappa = 25$ is chosen. We observe the exact same pattern as for the first time frame. This also holds for




Figure 10.11 Comparison of the three ES estimators at confidence level $\alpha = 0.05$ over time for the second time frame (2020-2021).

the 3d plot displayed in Figure 10.12 again simply relating the estimated risk measures for the time frame without the time.

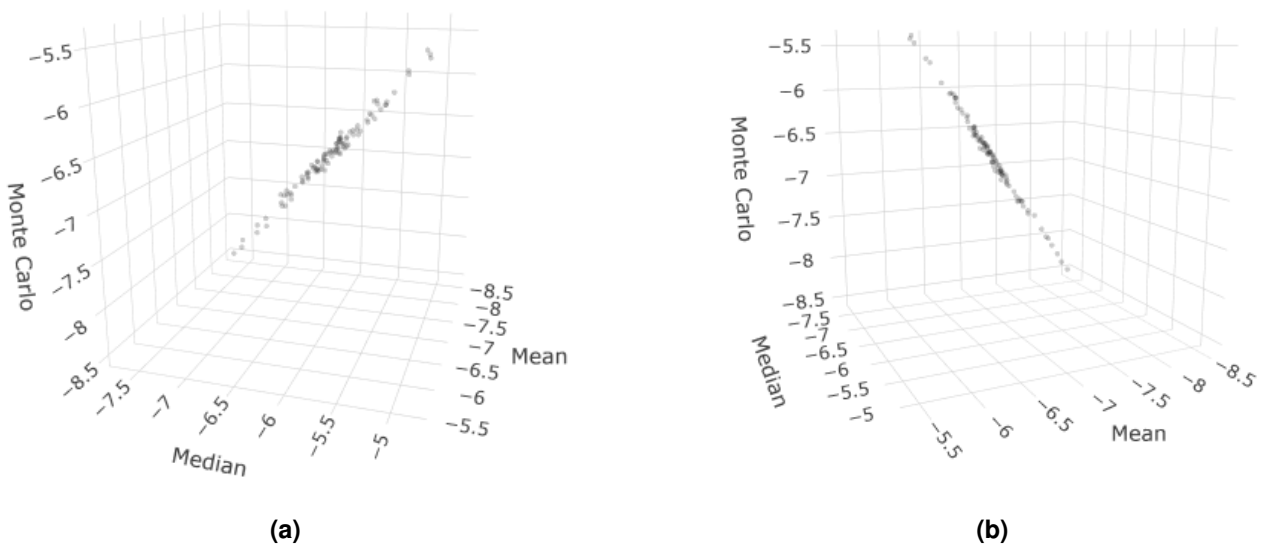


Figure 10.12 Two views of the 3d plot relating the risk estimates of the three different estimators for the ES at confidence level $\alpha = 0.05$ for the second time frame (2020-2021).

We will from now on only use the ES estimator based on the mean as defined in Equation (3.5). This more conservative choice would probably also be reasonable for a regulator. Having performed this com-

parison we can move on to the highly interesting section covering the backtesting of the estimated risk measures.


10.5 Backtesting

Before one backtests the estimated risk measures, one can use visualizations to get a rough sense for the behavior of the estimated risk measures. Figure 10.13 provides a visual comparison of the VaR and ES estimates for the two time frames. Obviously the ES estimates are always smaller than the VaR which is clear by definition and estimation. Moreover, these plots show that the implicit assumption of colinearity of the VaR and ES that Bayer and Dimitriadis (2020b) use for the strict ESR backtests seems to be very reasonable. The sudden drop in the risk measures in the beginning of May 2021 in Figure 10.13b is due to the fact that there the marginal window ends and the marginal models as well as the vine copula were adjusted.


Quite prominent visualizations of the VaR are exceedance plots that highlight the exceedances. Figure 10.14 shows an exemplary exceedance plot for each time frame.

For a similar visualization of the ES it is reasonable to not mark the exceedances of the ES but again the ones of the corresponding VaR which is more natural given that the ES is defined as the expected value given the VaR is exceeded. Such an ES exceedance plot is exemplarily given in Figure 10.15 for the first time frame (2016-2019).

These descriptive plots may directly show flaws in the estimation as well as clustered failing of the approach. This can then be analyzed further in the business setting. Here no odd behavior is visible and we can proceed with the traditional backtests.




(a) The first time frame (2016-2019).




(b) The second time frame (2020-2021).

Figure 10.13 Visual comparison plots of the VaR and ES risk estimates at the two α levels 0.01 and 0.05 for the first and second time frame. The vine window parameters are $\Psi = 250$ and $\kappa = 25$ for the first time frame and $\Psi = 100$ and $\kappa = 25$ for the second time frame.



(a) The first time frame (2016-2019).



(b) The second time frame (2020-2021).

Figure 10.14 Exceedance plots of the VaR at α level 0.025 for the first and second time frame. The vine window parameters are $\Psi = 250$ and $\kappa = 25$ as well as $\Psi = 100$ and $\kappa = 25$ respectively.




Figure 10.15 Exemplary exceedance plot for the ES at α level 0.025 for the first time frame (2016-2019). The vine window parameters are $\Psi = 250$ and $\kappa = 25$.

Traditional backtesting

As outlined in Chapter 8 we want to test all our estimated risk measures with traditional backtests. These are goodness of fit type tests where we would not like to reject the null hypothesis which states that the risk estimation procedure is correct. In Table 10.4 we collect all the p-values from all the presented backtests in Chapter 8 for both the VaR and ES at the confidence levels $\alpha \in \{0.01, 0.025, 0.05\}$ and for all four models of the first time frame. In some cases (especially for small α levels) where no result was reported this is due to limitations of the ESR approach and its implementation. More specifically the smallest simulation Bayer and Dimitriadis (2020b) have tested is the one with $\alpha = 0.025$ and a forecasting window of size 250. Thus whenever the ESR approach could not be performed the results are omitted. We will from now on consider the **common significance level of 5%** to evaluate the resulting p-values of the traditional backtests. For the **first time frame** the results are really pleasing as at every α level and for every successfully performed backtest one does not reject the null hypothesis of having a correct risk estimation procedure at the significance level of 5%.

The respective results of the models for the **second time frame** are collected in Table 10.5. Here again some ESR backtests could again not be performed as the forecasting window of size 100 in the second time frame is smaller than the smallest test simulation of the authors of the ESR approach. Thus, for the second time frame the ESR approach reaches its limits even quicker. Moreover, in two cases the exceedance residual backtests could not be performed due to the fact that there were no exceedance residuals. This suggests a conservative enough estimation. For the second time frame and the α level of 0.01 actually the risk estimation was in some cases too conservative this is also indicated by the four two sided backtests that failed i.e. one had to reject the null hypothesis. The p-values of the corresponding one sided backtests were then consequently as the estimation was too conservative very large. All other backtests could mostly quite clearly not be rejected and thus the results are also very pleasing for the second time frame.

All in all every considered model passed almost all the traditional backtests for both VaR and ES on all considered confidence levels α with only four exceptions where the estimation was too conservative resulting in the rejection of the two sided test. This suggests that the proposed unconditional risk measure estimation approach for both considered risk measures performs really good in practice. Yet it is quite obvious that from these results no favourable model can be chosen. For this endeavour we will have a look at comparative backtesting next.

Vine training length Ψ : 250, Vine window size κ : 25				
Risk measure	Backtest	α : 0.01	α : 0.025	α : 0.05
VaR	unconditional coverage	0.27	0.11	0.71
VaR	conditional coverage	0.49	0.16	0.37
ES	exceedance residuals (two sided)	0.99	0.86	0.18
ES	exceedance residuals (one sided)	0.70	0.49	0.08
ES	simple conditional calibration (two sided)	0.64	0.39	0.49
ES	simple conditional super-calibration (one sided)	0.69	0.45	0.49
ES	strict ESR test (two sided)	-	0.39	0.22
ES	one sided intercept ESR test	-	0.29	0.15
Vine training length Ψ : 250, Vine window size κ : 50				
Risk measure	Backtest	α : 0.01	α : 0.025	α : 0.05
VaR	unconditional coverage	0.27	0.054	0.71
VaR	conditional coverage	0.49	0.08	0.37
ES	exceedance residuals (two sided)	0.90	0.94	0.16
ES	exceedance residuals (one sided)	0.65	0.54	0.08
ES	simple conditional calibration (two sided)	0.64	0.27	0.42
ES	simple conditional super-calibration (one sided)	0.59	0.38	0.43
ES	strict ESR test (two sided)	-	0.49	0.35
ES	one sided intercept ESR test	-	0.24	0.14
Vine training length Ψ : 500, Vine window size κ : 25				
Risk measure	Backtest	α : 0.01	α : 0.025	α : 0.05
VaR	unconditional coverage	0.27	0.11	0.87
VaR	conditional coverage	0.49	0.16	0.49
ES	exceedance residuals (two sided)	1.00	0.66	0.16
ES	exceedance residuals (one sided)	0.68	0.72	0.08
ES	simple conditional calibration (two sided)	0.56	0.35	0.32
ES	simple conditional super-calibration (one sided)	1.00	0.74	0.65
ES	strict ESR test (two sided)	-	0.48	0.52
ES	one sided intercept ESR test	-	0.32	0.26
Vine training length Ψ : 500, Vine window size κ : 50				
Risk measure	Backtest	α : 0.01	α : 0.025	α : 0.05
VaR	unconditional coverage	0.27	0.11	0.87
VaR	conditional coverage	0.49	0.16	0.49
ES	exceedance residuals (two sided)	1.00	0.69	0.12
ES	exceedance residuals (one sided)	0.68	0.71	0.06
ES	simple conditional calibration (two sided)	0.57	0.36	0.30
ES	simple conditional super-calibration (one sided)	0.95	0.73	0.65
ES	strict ESR test (two sided)	-	0.73	0.56
ES	one sided intercept ESR test	-	0.39	0.27

Table 10.4 The results of the traditional backtests for the first time frame (2016-2019). In some cases (especially for small α levels) where no result was reported this was due to limitations of the ESR approach and its implementation. P-values smaller than 0.05 are marked **bold**.

Vine training length Ψ : 100, Vine window size κ : 25				
Risk measure	Backtest	α : 0.01	α : 0.025	α : 0.05
VaR	unconditional coverage	0.38	0.16	0.20
VaR	conditional coverage	0.65	0.28	0.22
ES	exceedance residuals (two sided)	0.00	0.32	1.00
ES	exceedance residuals (one sided)	1.00	0.87	0.54
ES	simple conditional calibration (two sided)	0.06	0.32	0.54
ES	simple conditional super-calibration (one sided)	1.00	0.99	0.60
ES	strict ESR test (two sided)	-	-	0.69
ES	one sided intercept ESR test	-	-	0.37
Vine training length Ψ : 100, Vine window size κ : 50				
Risk measure	Backtest	α : 0.01	α : 0.025	α : 0.05
VaR	unconditional coverage	1.00	0.16	0.20
VaR	conditional coverage	0.99	0.28	0.22
ES	exceedance residuals (two sided)	-	0.19	0.49
ES	exceedance residuals (one sided)	-	0.88	0.77
ES	simple conditional calibration (two sided)	0.00	0.11	0.46
ES	simple conditional super-calibration (one sided)	1.00	1.00	0.78
ES	strict ESR test (two sided)	-	-	0.93
ES	one sided intercept ESR test	-	-	0.41
Vine training length Ψ : 200, Vine window size κ : 25				
Risk measure	Backtest	α : 0.01	α : 0.025	α : 0.05
VaR	unconditional coverage	0.38	0.16	0.20
VaR	conditional coverage	0.65	0.28	0.22
ES	exceedance residuals (two sided)	0.00	0.35	0.98
ES	exceedance residuals (one sided)	1.00	0.81	0.54
ES	simple conditional calibration (two sided)	0.07	0.33	0.54
ES	simple conditional super-calibration (one sided)	1.00	1.00	0.57
ES	strict ESR test (two sided)	-	-	0.69
ES	one sided intercept ESR test	-	-	0.38
Vine training length Ψ : 200, Vine window size κ : 50				
Risk measure	Backtest	α : 0.01	α : 0.025	α : 0.05
VaR	unconditional coverage	1.00	0.38	0.20
VaR	conditional coverage	0.99	0.57	0.22
ES	exceedance residuals (two sided)	-	1.00	0.39
ES	exceedance residuals (one sided)	-	0.67	0.81
ES	simple conditional calibration (two sided)	0.00	0.07	0.44
ES	simple conditional super-calibration (one sided)	1.00	1.00	0.83
ES	strict ESR test (two sided)	-	-	0.98
ES	one sided intercept ESR test	-	-	0.43

Table 10.5 The results of the traditional backtests for the second time frame (2020-2021). In some cases (especially for small α levels) where no result was reported this was due to limitations of the ESR approach and its implementation, except for the one case where the exceedance residual test could not be performed as there were no exceedance residuals. P-values smaller than 0.05 are marked **bold**.

Comparative backtesting

First time frame (2016-2019)

Vine window/training related parameters First up we would like to determine which specification of the two vine window/training parameters $\Psi \in \{250, 500\}$ and $\kappa \in \{25, 50\}$ is the preferred one for the first time frame. We will base our decision on comparative backtests for the ES. The results of the comparative backtests as of Definition 8.2.2 for the ES at the two α levels 0.01 and 0.05 are displayed in Table 10.6. To evaluate the given transformed test statistic $\Phi(\mathcal{T}_{comp})$ one should remember the three zone approach illustrated in Figure 8.1. Thus a value below 0.5 indicates a superior internal model. A significant decision at the common significance level $\nu = 0.05$ can only be supported if $\Phi(\mathcal{T}_{comp}) < \nu$ or $\Phi(\mathcal{T}_{comp}) > 1 - \nu$. It follows that here no definitive decision can be supported in all comparisons. However we can have a look at indications. We for example observe that the bigger vine training length of $\Psi = 500$ seems to be superior to the smaller one of $\Psi = 250$. Moreover the smaller window size of the vines $\kappa = 25$ seems not to be necessary as the comparative backtests even suggest that the parameter κ set to 50 performs better. So from the indications one would probably use the model with vine training size $\Psi = 500$ and window size $\kappa = 50$.

$\alpha = 0.01$ internal model	standard model			
	$\Psi = 250, \kappa = 25$	$\Psi = 500, \kappa = 25$	$\Psi = 250, \kappa = 50$	$\Psi = 500, \kappa = 50$
$\Psi = 250, \kappa = 25$		0.8127	0.7170	0.8456
$\Psi = 500, \kappa = 25$	0.1873		0.2757	0.5677
$\Psi = 250, \kappa = 50$	0.2830	0.7243		0.7648
$\Psi = 500, \kappa = 50$	0.1544	0.4323	0.2352	

(a) ES estimated at confidence level $\alpha = 0.01$.

$\alpha = 0.05$ internal model	standard model			
	$\Psi = 250, \kappa = 25$	$\Psi = 500, \kappa = 25$	$\Psi = 250, \kappa = 50$	$\Psi = 500, \kappa = 50$
$\Psi = 250, \kappa = 25$		0.6957	0.6546	0.6657
$\Psi = 500, \kappa = 25$	0.3043		0.3421	0.3670
$\Psi = 250, \kappa = 50$	0.3454	0.6579		0.6227
$\Psi = 500, \kappa = 50$	0.3343	0.6330	0.3773	

(b) ES estimated at confidence level $\alpha = 0.05$.

Table 10.6 The results i.e. $\Phi(\mathcal{T}_{comp})$ of the comparative backtests as of Definition 8.2.2 for the first time frame (2016-2019) for the ES at the two α levels 0.01 and 0.05 with respect to different rolling vine window parameters Ψ and κ . A value below 0.5 indicates a superior internal model but here in all cases no significant decision at confidence $\nu = 0.05$ can be supported.

Allowed copula families As already mentioned before we have also fixed the vine parameters Ψ and κ to 250 and 50 respectively to compare how the variation in the set of allowed bivariate copula families influences the risk estimation quality. In particular we examine the **three cases** where **only Gaussian** components, **Student's t and Gaussian** components and **all parametric copula families** are allowed. The results of the comparative backtest on the ES again at the two α levels 0.01 and 0.05 are displayed in Table 10.7. Again for the significance level $\nu = 0.05$ no clear decision is supported however here for the first time frame the indications suggest that the model with Gaussian and Student's t components performs the best. Figure 10.16 shows both the ES estimates at α level 0.05 based on the R-vine copula allowing for all parametric copula families and only Gaussian and Student's t ones. There we can observe that both models estimate the risk for the first time frame quite similar. The only obvious difference is the third exceedance for the general R-vine. It will be interesting whether a more flexible model allowing for all parametric copula families and accounting better for tail dependence is superior to the Gaussian and

$\alpha = 0.01$ internal model	standard model		
	All parametric copula families	Gaussian & Student's t	Gaussian
All parametric copula families		0.7002	0.5801
Gaussian & Student's t	0.2998		0.1880
Gaussian	0.4199	0.8120	

(a) ES estimated at confidence level $\alpha = 0.01$.

$\alpha = 0.05$ internal model	standard model		
	All parametric copula families	Gaussian & Student's t	Gaussian
All parametric copula families		0.8160	0.6848
Gaussian & Student's t	0.1840		0.1154
Gaussian	0.3152	0.8846	

(b) ES estimated at confidence level $\alpha = 0.05$.

Table 10.7 The results i.e. $\Phi(\mathcal{T}_{comp})$ of the comparative backtests as of Definition 8.2.2 for the first time frame (2016-2019) for the ES at the two α levels 0.01 and 0.05 with respect to the copula families used in the R-vine. The parameters $\Psi = 250, \kappa = 50$ and all other parameters are fixed. A value below 0.5 indicates a superior internal model but here in all cases no significant decision at confidence $\nu = 0.05$ can be supported.




Figure 10.16 Comparison of the ES estimates at α level 0.05 based on the R-vine copula allowing for all parametric copula families and only Gaussian and Student's t ones for the first time frame (2016-2019).

Student's t vine in the second time frame that is set in an overall more volatile market situation.

Second time frame (2020-2021)

Vine window/training related parameters Now we would like to determine which specification of the two vine window/training parameters $\Psi \in \{100, 200\}$ and $\kappa \in \{25, 50\}$ is the suggested one for the second time frame. We will again base our decision on comparative backtests for the ES. The results of the comparative backtests for the ES at the two α levels 0.01 and 0.05 are displayed in Table 10.8 in the same manner as above for the first time frame. As for the first time frame at the significance level $\nu = 0.05$

$\alpha = 0.01$ internal model	standard model			
	$\Psi = 100, \kappa = 25$	$\Psi = 200, \kappa = 25$	$\Psi = 100, \kappa = 50$	$\Psi = 200, \kappa = 50$
$\Psi = 100, \kappa = 25$		0.6760	0.0686	0.7091
$\Psi = 200, \kappa = 25$	0.3240		0.1999	0.7862
$\Psi = 100, \kappa = 50$	0.9315	0.8001		0.7983
$\Psi = 200, \kappa = 50$	0.2909	0.2138	0.2017	

(a) ES estimated at confidence level $\alpha = 0.01$.

$\alpha = 0.05$ internal model	standard model			
	$\Psi = 100, \kappa = 25$	$\Psi = 200, \kappa = 25$	$\Psi = 100, \kappa = 50$	$\Psi = 200, \kappa = 50$
$\Psi = 100, \kappa = 25$		0.8260	0.0981	0.8454
$\Psi = 200, \kappa = 25$	0.1740		0.1153	0.9069
$\Psi = 100, \kappa = 50$	0.9019	0.8847		0.8898
$\Psi = 200, \kappa = 50$	0.1546	0.0931	0.1102	

(b) ES estimated at confidence level $\alpha = 0.05$.

Table 10.8 The results i.e. $\Phi(\mathcal{T}_{comp})$ of the comparative backtests as of Definition 8.2.2 for the second time frame (2020-2021) for the ES at the two α levels 0.01 and 0.05 with respect to different rolling vine window parameters Ψ and κ . A value below 0.5 indicates a superior internal model but here in all cases no significant decision at confidence $\nu = 0.05$ can be supported.

we can not support a definitive decision. The indications however are the same as in the first time frame. The longer vine training window of $\Psi = 200$ and the longer vine usage of $\kappa = 50$ is suggested quite clearly by the results of the comparative backtests on the ES.

Allowed copula families With exactly these suggested parameters fixed we again examine the three cases where only Gaussian components, Student's t and Gaussian components and all parametric copula families are allowed now for the second time frame. The results of the comparative backtest on the ES again at the two α levels 0.01 and 0.05 are displayed in Table 10.9. Here we again can not make a definitive decision at the $\nu = 0.05$ significance level but especially for $\alpha = 0.05$ there is a clear indication that the model allowing for all parametric copula families results in a superior risk measure estimation than the ones restricted to Gaussian and/or Student's t components. The higher usage of copulas accounting for tail dependence seems to pay off in the higher volatility setting of the second time frame. Visually however the difference between the models is difficult to detect as visible in Figure 10.17 that shows the ES estimates at $\alpha = 0.05$ of the model allowing for all parametric copula families and the one restricted to Gaussian and Student's t components side by side.

This concludes the portrayal of the unconditional approach using the portfolio of Spanish stocks and we move on to the conditional approach.

$\alpha = 0.01$ internal model	standard model		
	All parametric copula families	Gaussian & Student's t	Gaussian
All parametric copula families		0.3903	0.4444
Gaussian & Student's t	0.6097		0.6114
Gaussian	0.5556	0.3886	

(a) ES estimated at confidence level $\alpha = 0.01$.


$\alpha = 0.05$ internal model	standard model		
	All parametric copula families	Gaussian & Student's t	Gaussian
All parametric copula families		0.1005	0.1000
Gaussian & Student's t	0.8995		0.7078
Gaussian	0.9000	0.2922	

(b) ES estimated at confidence level $\alpha = 0.05$.

Table 10.9 The results i.e. $\Phi(\mathcal{T}_{comp})$ of the comparative backtests as of Definition 8.2.2 for the second time frame (2020-2021) for the ES at the two α levels 0.01 and 0.05 with respect to the copula families used in the R-vine. The parameters $\Psi = 200, \kappa = 50$ and all other parameters are fixed. A value below 0.5 indicates a superior internal model but here in all cases no significant decision at confidence $\nu = 0.05$ can be supported.

Risk measure: ES with alpha level 0.05, all parametric copula families

Exceedances of the VaR are highlighted in red.



Risk measure: ES with alpha level 0.05, only Gaussian and Student's t copulas

Exceedances of the VaR are highlighted in red.




Figure 10.17 Comparison of the ES estimates at α level 0.05 based on the R-vine copula allowing for all parametric copula families and only Gaussian and Student's t ones for the second time frame (2020-2021).

11 Conditional Spanish stock portfolio analysis

This chapter is split up into the first more detailed Section 11.1 covering the single conditional risk measure estimation approach and the shorter Section 11.2 shortly displaying the conditional approach using two conditional assets. Section 11.1 thereby is organized as follows. First up the extended setup with the two considered conditional assets the SP500 and Eurostoxx 50 is introduced. Then the additional marginal models are checked for an appropriate fit in the next subsection of Section 11.1. This is followed by a detailed look at the fitted D-vine copulas for both time frames and conditional assets. The next part of Section 11.1 covers the behavior of the conditioning values. It is important to have a good grasp how the conditioning values behave for the quantile based, prior residual and realized residual strategy in order to interpret the resulting conditional risk measures correctly. Finally the last part of Section 11.1 displays the conditional risk measure estimates. In particular one compares the quantile based conditional risk measure estimates with the ones based on the prior residuals and realized residuals as well as with the unconditional risk measure estimates in order to understand which influence and how big the influence of the conditioning asset is on the portfolio level risk measure estimates.

11.1 One conditional asset


The setup

To showcase the conditional risk measure estimation based on a single conditional asset we will use the same base portfolio of Spanish stocks as given in Table 10.1 with the corresponding weights. The two time frames of interest as they were highlighted in Figure 10.1 will also be the same. The additional two conditional assets that will be considered are the **SP500** and the **Eurostoxx 50 index**. While the former is an index indicating the state of the American economy the latter does the same for the European economy. As our portfolio consists of Spanish stocks we expect the influence of the Eurostoxx 50 index to be greater than the one of the SP500. Some assets are also part of the Eurostoxx index, namely the Banco Santander, the Amadeus It Group, Iberdrola and Inditex. However they cumulatively only account for roughly 5% of the index. The parameters concerning the marginal time series models will be the same as for the unconditional approach and the parameters concerning the here D-vine copula models will be fixed in the conditional case to the values that were suggested by the comparative backtesting in the unconditional setting. These were $\Psi = 500$ and $\kappa = 50$ for the first time frame and $\Psi = 200$ and $\kappa = 50$ for the second time frame. Also the simulation sample size is again set to $S = 100000$. Moreover all parametric copula families implemented in the *rvinecopulib* were allowed in the D-vine fitting. The default marginal model is also for the conditional assets an ARMA(1,1)-GARCH(1,1) model with the residual white noise distribution being a skewed Student's t distribution. The quality assessment of the additional two time series models for each time frame will follow right below. For all models the VaR and the ES were estimated for the α levels 0.01, 0.025, 0.05. The focus during this conditional analysis will however lie on the ES estimated at the level $\alpha = 0.025$.

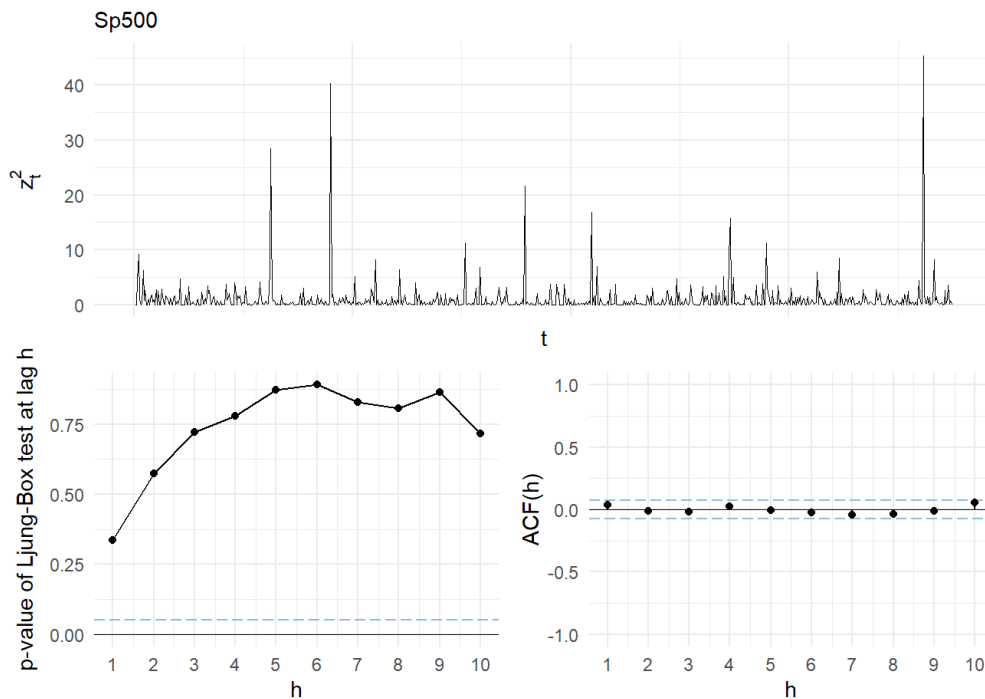
Quality assessment of the additional marginal models

Again the tools of Section 4.5 are used to assess the quality of the four additional marginal time series models. We will perform the residual analysis for the model corresponding to the first marginal model. The residual plots for the first time frame are given in Figure 11.1 for the SP500 and in Figure 11.2 for the Eurostoxx 50 index. All plots for both the mean and volatility equation do not look suspicious and thus the

fit of the ARMA(1,1)-GARCH(1,1) models is assumed to be good. The residual plots for the second time frame are given in the Appendix as Figures A.17 and A.18. The mean equation of the SP500 asset shows a slightly questionable fit but as no severe misfit is visible and all other residual plots indicate good fits we will also stick to the default model for both marginal models of the second time frame.




(a) Residual plots for the mean equation.




(b) Residual plots for the volatility equation.

Figure 11.1 The residual plots for the conditional asset the **SP500** index for the first marginal window and the first time frame (2016-2019).



(a) Residual plots for the mean equation.



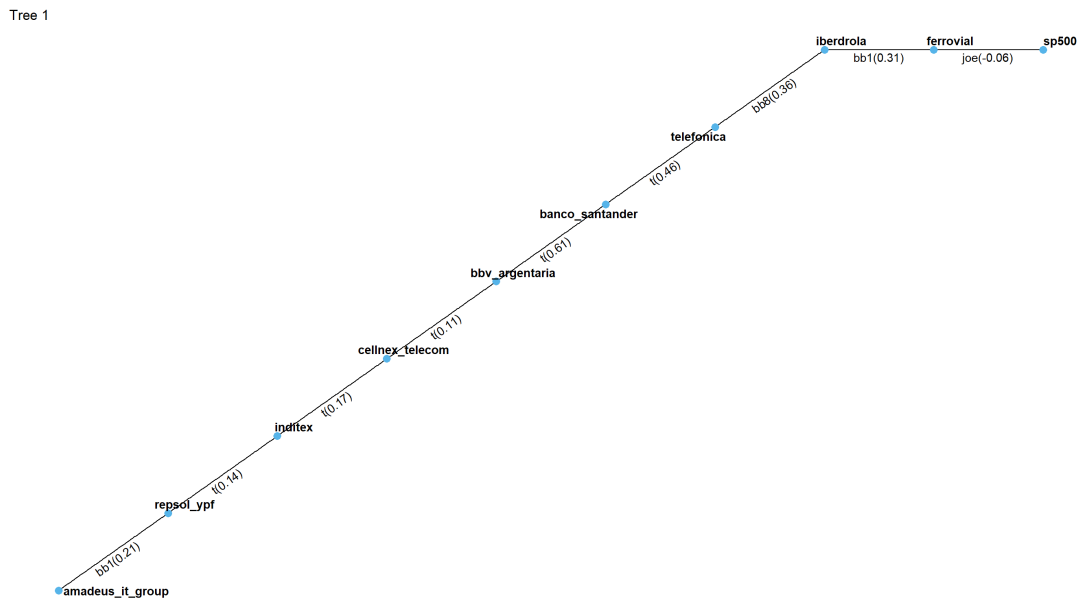
(b) Residual plots for the volatility equation.

Figure 11.2 The residual plots for the conditional asset the **Eurostoxx 50 index** for the first marginal window and the first time frame (2016-2019).


The fitted D-vines

Now we analyze the fitted D-vines. Here we would like to see how Algorithm 2 has determined the ordering and how strong the dependencies are. Also it is interesting to detect changing patterns over time. This is again facilitated by visualizing the first tree level. Furthermore also the used bivariate building blocks are of interest and especially the ones associated with the conditioning market index. These are the bivariate copulas associated with the rightmost edge in each tree/path of the D-vine.

The first time frame (2016-2019) We start with the first time frame and the SP500 as the conditional asset. The first trees of the D-vines corresponding to the first and last vine window alongside the edge copula families and the associated empirical Kendall's tau are displayed in Figure 11.3. The really weak



(a) The first vine window.



(b) The sixth/last vine window.

Figure 11.3 The first trees of the D-vines corresponding to the first and sixth/last vine window for the first time frame (2016-2019) in the single conditional setting. The conditional asset is the **SP500**.

dependence between the conditional SP500 asset and the rightmost portfolio asset in both cases is actually surprising. Although the American market is one of the most important export markets outside of Europe for Spain, the index seems to have very little influence on our portfolio. Consequently also the ordering of the other stocks changes very strongly over time. Table 11.1 displays the used bivariate copula families with their frequency of usage. Especially in the sixth vine window we can detect a quite large number of used independence copulas. This might have to do with the weak influence of the SP500. In order to even better understand how the conditional asset influences our D-vine we look at the used copula families alongside their empirical Kendall's tau directly associated with the conditional asset i.e. the ones associated with the rightmost edge in each tree level. They are given in Table 11.2 and indeed for each tree level the dependence is very weak and quite some independence copulas arise directly in the context of the conditional asset for the sixth vine window.

Copula family	Number of usage	Copula family	Number of usage
Gaussian	2	Independence	6
Student's t	8	Clayton	4
Frank	11	Gumbel	2
Joe	6	BB1	2
BB8	4		

(a) The first vine window.

Copula family	Number of usage	Copula family	Number of usage
Gaussian	3	Independence	11
Student's t	7	Clayton	5
Frank	9	Gumbel	4
BB1	3	BB8	3

(b) The sixth vine window.

Table 11.1 The used copula families alongside their frequencies for the first and sixth fitted D-vine for the first time frame (2016-2019) in the single conditional setting. The conditional asset is the **SP500**.

Tree level	Copula	Copula family	Empirical Kendall's tau
1	$C_{j_1, I}$	Joe	-0.0593
2	$C_{j_2, I; j_1}$	Student's t	-0.0061
3	$C_{j_3, I; j_1, j_2}$	Joe	0.0455
4	$C_{j_4, I; j_1, j_3}$	Joe	0.0308
5	$C_{j_5, I; j_1, j_4}$	Joe	0.0268
6	$C_{j_6, I; j_1, j_5}$	Independence	0
7	$C_{j_7, I; j_1, j_6}$	Independence	0
8	$C_{j_8, I; j_1, j_7}$	Joe	0.0703
9	$C_{j_9, I; j_1, j_8}$	Frank	0.0479

(a) The first vine window. (j_1 = Ferrovial, j_2 = Iberdrola, j_3 = Telefonica, j_4 = Banco Santander, j_5 = BBV Argentaria, j_6 = Cellnex Telecom, j_7 = Inditex, j_8 = Repsol Ypf, j_9 = Amadeus It Group)

Tree level	Copula	Copula family	Empirical Kendall's tau
1	$C_{j_1, I}$	Student's t	0.0425
2	$C_{j_2, I; j_1}$	Independence	0
3	$C_{j_3, I; j_1, j_2}$	Independence	0
4	$C_{j_4, I; j_1, j_3}$	Independence	0
5	$C_{j_5, I; j_1, j_4}$	Clayton	-0.0389
6	$C_{j_6, I; j_1, j_5}$	Clayton	-0.0344
7	$C_{j_7, I; j_1, j_6}$	Gaussian	0.0422
8	$C_{j_8, I; j_1, j_7}$	Clayton	-0.0472
9	$C_{j_9, I; j_1, j_8}$	Independence	0

(b) The sixth vine window. (j_1 = Repsol Ypf, j_2 = Banco Santander, j_3 = BBV Argentaria, j_4 = Telefonica, j_5 = Amadeus It Group, j_6 = Ferrovial, j_7 = Inditex, j_8 = Iberdrola, j_9 = Cellnex Telecom)

Table 11.2 The used copula families alongside their empirical Kendall's tau associated with the conditional asset for the first and sixth fitted D-vine for the first time frame (2016-2019) in the single conditional setting. The conditional asset is the **SP500**.

Now we have a look of the case with the Eurostoxx 50 index as the conditional asset for the first time frame. Here the first trees of the D-vines corresponding to the first and last vine window alongside the edge copula families and the associated empirical Kendall's tau are displayed in Figure 11.4. The story here is quite different from the one before. We can observe a quite strong positive bivariate dependence between the conditional asset the Eurostoxx 50 index and the rightmost portfolio asset which is over time stable the bank Banco Santander. Also the rest of the ordering especially on the right side is quite stable. This indicates that the conditional risk measures based on this dependence structure will be quite strongly influenced by the conditional asset. Table 11.3 shows again the used bivariate components alongside their usage frequency. Visibly in the last used D-vine a lot less Gaussian and Student's t components were used. The flexibility to adapt the dependence model facilitated by the rolling window approach is here displayed once again. Table 11.4 shows that not only in the first tree level the dependence on the conditional asset was quite strong, but for example for the first vine window even in the last tree level the fitted Frank copula had a corresponding empirical Kendall's tau of 0.2335.


Copula family	Number of usage	Copula family	Number of usage
Gaussian	4	Independence	5
Student's t	10	Clayton	3
Frank	9	Gumbel	5
Joe	2	BB1	1
BB8	6		

(a) The first vine window.


Copula family	Number of usage	Copula family	Number of usage
Gaussian	3	Independence	7
Student's t	5	Clayton	1
Frank	11	Gumbel	8
BB1	5	BB8	3
Joe	2		

(b) The sixth vine window.

Table 11.3 The used copula families alongside their frequencies for the first and sixth fitted D-vine for the first time frame (2016-2019) in the single conditional setting. The conditional asset is the **Eurostoxx 50**.



(a) The first vine window.



(b) The sixth/last vine window.

Figure 11.4 The first trees of the D-vines corresponding to the first and sixth/last vine window for the first time frame (2016-2019) in the single conditional setting. The conditional asset is the **Eurostoxx 50**.

Tree level	Copula	Copula family	Empirical Kendall's tau
1	$C_{j_1, I}$	Student's t	0.5236
2	$C_{j_2, I; j_1}$	Gaussian	0.1225
3	$C_{j_3, I; j_1, j_2}$	Student's t	0.1966
4	$C_{j_4, I; j_1: j_3}$	Gumbel	0.1200
5	$C_{j_5, I; j_1: j_4}$	Frank	0.2296
6	$C_{j_6, I; j_1: j_5}$	Gumbel	0.1104
7	$C_{j_7, I; j_1: j_6}$	Clayton	0.0869
8	$C_{j_8, I; j_1: j_7}$	Gaussian	0.0925
9	$C_{j_9, I; j_1: j_8}$	Frank	0.2335

(a) The first vine window. (j_1 = Banco Santander, j_2 = BBV Argentaria, j_3 = Telefonica, j_4 = Repsol Ypf, j_5 = Inditex, j_6 = Iberdrola, j_7 = Cellnex Telecom, j_8 = Ferrovial, j_9 = Amadeus It Group)


Tree level	Copula	Copula family	Empirical Kendall's tau
1	$C_{j_1, I}$	BB1	0.5015
2	$C_{j_2, I; j_1}$	Gumbel	0.0948
3	$C_{j_3, I; j_1, j_2}$	BB1	0.1765
4	$C_{j_4, I; j_1: j_3}$	Frank	0.2133
5	$C_{j_5, I; j_1: j_4}$	BB1	0.2921
6	$C_{j_6, I; j_1: j_5}$	Frank	0.2063
7	$C_{j_7, I; j_1: j_6}$	Gaussian	0.1875
8	$C_{j_8, I; j_1: j_7}$	Gumbel	0.0898
9	$C_{j_9, I; j_1: j_8}$	Independence	0

(b) The sixth vine window. (j_1 = Banco Santander, j_2 = BBV Argentaria, j_3 = Telefonica, j_4 = Repsol Ypf, j_5 = Amadeus It Group, j_6 = Inditex, j_7 = Ferrovial, j_8 = Iberdrola, j_9 = Cellnex Telecom)

Table 11.4 The used copula families alongside their empirical Kendall's tau associated with the conditional asset for the first and sixth fitted D-vine for the first time frame (2016-2019) in the single conditional setting. The conditional asset is the **Eurostoxx 50**.


the Amadeus It Group, Inditex and Ferrovial then followed by the Banco Santander asset. So it seems like during the Covid-19 pandemic the dependence between the general European market was stronger with software, clothing and construction companies than with the banks that had the strongest dependence right before the pandemic. This could be due to the fact that the direct consequences of the pandemic or economic downtrends show themselves quicker in industries like the clothing industry as people might be discouraged to spend their money on leisure items. The Tables A.3 and A.4, which display again the used bivariate copula families with the respective frequencies and the copula families corresponding to the edges directly associated with the conditional asset, can be found in the appendix. They again show the higher usage of copulas stressing the modelling of tail dependence and strong dependencies associated with the conditional asset throughout the tree levels.

Tree 1



(a) The first vine window.

Tree 1



(b) The second/last vine window.

Figure 11.6 The first trees of the D-vine corresponding to the first and second/last vine windows for the second time frame (2020-2021) in the single conditional setting. The conditional asset is the **Eurostoxx 50**.

Behavior of the conditioning values

Having a good grasp of the fitted dependence structure, it is reasonable to know how the actual conditioning values on which the conditional risk measure estimates are based upon behave. In particular the three major strategies will be compared. The quantile based, the prior residual and the realized residual approach. As the strategies are conceptually the same for both time frames here only the first time frame is considered. The plots showing the important quantile based conditioning values for both conditioning assets on the second time frame however can be found in the appendix in Figures A.19 and A.20.

The first conditioning values of interest are the **quantile based** ones. Basically through the specification of the quantile level α^I we would like to mitigate a bad performance of the conditional asset with values close to 0 and a good performance over the time frame of interest with values closer to 1. We can then retransform this conditional quantile level by the inverse PIT as discussed in the methodology part to obtain the conditional values on the original log return scale. We visualize these conditioning values on the log return scale alongside the realized time series of the conditional asset. This allows us to see whether the quantile based conditioning values actually do what we expect. Figures 11.7 and 11.8 show the quantile based conditioning values for both considered conditional assets. In both cases we can observe that the quantile based approach produces forecasted time series that are mitigating various states of the conditioning asset. For example the conditioning values based on the $\alpha^I = 0.05$ provide somewhat of a worst case scenario at least for the SP500 asset.




Figure 11.7 Behavior of the **quantile based conditioning values** on the log return scale of the conditioning asset **SP500** for the first time frame (2016-2019).

The **prior residual** and **realized residual** strategies are displayed only exemplary for the SP500 as the conditioning asset because the behavior is the same for the Eurostoxx 50 index. The conditioning values corresponding to the prior residual strategy as of Equations (7.1) and (7.2) are given in Figure 11.9. We can definitely see a close mitigation of the realized time series. There are however slight delays and some exaggerations visible. This already in the theoretical part anticipated behavior is especially problematic

Quantile based conditioning values of the Eurostoxx 50

The realized log returns of the Eurostoxx 50 are given in grey.




Figure 11.8 Behavior of the **quantile based conditioning values** on the log return scale of the conditioning asset **Eurostoxx 50** for the first time frame (2016-2019).

with respect to exaggerations towards higher log returns. But a look at the corresponding conditional risk measure estimates will later show whether this strategy can be a valid one especially for comparisons.

Speaking of comparisons, the realized residual based conditioning values as of Equation (7.3) should exactly mitigate the realized time series of the conditioning asset in order to compare the resulting oracle risk measure estimates with the ones resulting from the quantile based approach. This is supported by Figure 11.10 which shows the realized residual based conditioning values on the log return scale of the conditioning asset SP500 for the first time frame.




Figure 11.9 Behavior of the **prior residual based conditioning values** on the log return scale of the conditioning asset **SP500** for the first time frame with a zoom at March 2019. The realized log returns of the conditional asset are drawn in grey and the conditioning values are drawn in violet.




Figure 11.10 Behavior of the **realized residual based conditioning values** on the log return scale of the conditioning asset **SP500** for the first time frame with a zoom at March 2019. The realized log returns of the conditional asset are drawn in grey and the conditioning values are drawn in green.

The conditional risk measure estimates

Now finally we want to have a look at the actual conditional risk measure estimates. As mentioned at the beginning of this section we will have a look at the ES estimated at the level $\alpha = 0.025$. We will first focus on the quantile based conditional risk measures and will compare these with the realized residual based risk measures and the unconditional risk measures. Therefore we consider unconditional risk measures on the same portfolio and with all parameters held equal to the conditional estimation once allowing for the general unconditional estimation based on a R-vine copula and once allowing only for a D-vine in the unconditional approach. This leads to 3 comparisons of the conditional quantile based ES estimates for each conditional asset and time frame.

Again starting with the first time frame we look at the ES estimates conditioned on the **SP500**. Figure 11.11 shows the raw quantile based ES estimates. Quite evidently as also anticipated by the weak de-




Figure 11.11 Quantile based conditional ES estimates at $\alpha = 0.025$ with the conditioning asset **SP500** for the first time frame (2016-2019).

pendence in the corresponding D-vines the resulting conditional ES series lie quite close to each other although the biggest difference in the quantile level is quite large with 0.45. So as the portfolio level ES seems to be quite robust against strong market downtrends of the American market one could argue that the portfolio might be a good hedge against the risk of a downtrend in the American market. Also the comparison with the conditional ES estimates based on the realized residuals displayed in Figure 11.12 paints the same picture. Although much more volatile the realized residual based ES estimates are very close to the quantile based ones.

Moreover the comparisons with the unconditional ES estimates given in Figure 11.13 show that all quantile based conditional ES estimates are quite close to the unconditional estimates. This not surprising behavior shows how the conditional risk measure estimation approach not only provides the final conditional risk measure estimates but with the possibility to interpret the dependence model we can provide insights why the conditional risk measures behave like they do. The results for the second time frame considering the SP500 as the conditional asset are almost identical to the ones presented here for the first time frame. Thus the visualizations concerning the second time frame can be found in the Appendix A.

Quantile based conditional Expected Shortfall

Alpha level: 2.5%

Conditional variable: SP500

The realized portfolio log returns are given in grey.

The ES estimates conditioned on the **realized residuals** are drawn in **green**.





Figure 11.12 Comparison of the **quantile based conditional ES estimates** at $\alpha = 0.025$ (blue) with the ones based on the **realized residuals** (green) for the first time frame (2016-2019). The conditioning asset is the **SP500**.

All in all the portfolio manager might use these results to argue that the portfolio at hand is well protected against market risks based on fluctuations of the American market.

Quantile based conditional Expected Shortfall


Alpha level: 2.5%
Conditional variable: SP500
The realized portfolio log returns are given in grey.
The unconditional D-vine based ES estimates are drawn in orange.



(a) Unconditional approach using a D-vine.

Quantile based conditional Expected Shortfall

Alpha level: 2.5%
Conditional variable: SP500
The realized portfolio log returns are given in grey.
The unconditional R-vine based ES estimates are drawn in gold.



(b) Unconditional approach using a R-vine.

Figure 11.13 Comparison of the **quantile based conditional ES estimates** at $\alpha = 0.025$ (blue) with the **unconditional ES estimates** (orange/gold) for the first time frame (2016-2019). The conditioning asset is the **SP500**.

Next up we look at the ES estimates conditioned on quantiles resembling states of the European market i.e. the **Eurostoxx 50 index** for the first time frame. Figure 11.14 shows the raw quantile based ES estimates. Opposed to the ones shown before with the SP500 as the conditioning asset here we can see




Figure 11.14 Quantile based conditional ES estimates at $\alpha = 0.025$ with the conditioning asset **Eurostoxx 50** for the first time frame (2016-2019).

that indeed the different series corresponding to the quantile levels are clearly distinguishable and thus states of the European economy have a serious effect on the estimated ES. This is also evident when looking at the comparison with the ES estimates based on the realized residuals given in Figure 11.15. The visualization emphasises that the portfolio is indeed vulnerable to market risk based on fluctuations of the European market.

Valuable insights are derived from Figure 11.16 comparing the conditional ES estimates at the different quantile levels with the unconditional ES estimates at the same α level. Notably it does not seem to have a huge influence whether one uses a D-vine or a more flexible R-vine in the unconditional case. The most interesting observation we can make here is that the unconditional ES series is very close to the one corresponding to the quantile level $\alpha^I = 0.1$. Thereby the D-vine based unconditional ES series is especially at the end of the series a bit closer to the quantile based conditional ES series than the R-vine based one. Thus we can argue that by using only the unconditional approach to estimate the ES at the confidence level $\alpha = 0.025$ we have already accounted for a quite bad general European market. However if we would like to account for a really bad European market for example at the quantile level 0.05 we would have to further adjust our risk estimates to be more conservative as indicated by the series corresponding to the quantile level $\alpha^I = 0.05$. These conditional risk measure estimates could then also be used to further analyze the effects on capital requirements or other key metrics. The results for the second time frame are again very similar to the ones displayed here for the first time frame. The corresponding visualizations can be found in the Appendix A.

Quantile based conditional Expected Shortfall

Alpha level: 2.5%

Conditional variable: Eurostoxx 50

The realized portfolio log returns are given in grey.

The ES estimates conditioned on the **realized residuals** are drawn in **green**.





Figure 11.15 Comparison of the **quantile based conditional ES estimates** at $\alpha = 0.025$ (blue) with the ones based on the **realized residuals** (green) for the first time frame (2016-2019). The conditioning asset is the **Eurostoxx 50**.

Quantile based conditional Expected Shortfall


Alpha level: 2.5%
Conditional variable: Eurostoxx 50
The realized portfolio log returns are given in grey.
The unconditional D-vine based ES estimates are drawn in orange.



(a) Unconditional approach using a D-vine.

Quantile based conditional Expected Shortfall

Alpha level: 2.5%
Conditional variable: Eurostoxx 50
The realized portfolio log returns are given in grey.
The unconditional R-vine based ES estimates are drawn in gold.



(b) Unconditional approach using a R-vine.

Figure 11.16 Comparison of the **quantile based conditional ES estimates** at $\alpha = 0.025$ (blue) with the **unconditional ES estimates** (orange/gold) for the first time frame (2016-2019). The conditioning asset is the **Eurostoxx 50**.

Yet we have not had a look at the conditional risk measure estimates based on the prior residual strategy. As an example Figure 11.17 displays the conditional ES estimates with the conditional asset being the Eurostoxx 50 for the first time frame. At the first glance the conditional risk measure estimates seem to be




Figure 11.17 Display of the conditional ES estimates at $\alpha = 0.025$ based on the **prior residuals** (violet) along side the realized portfolio log returns (grey) for the first time frame (2016-2019). The conditioning asset is the **Eurostoxx 50**.

very close to the actual realized portfolio returns which might be highly desirable. A closer look however uncovers that often the risk measures are not at all conservative enough. Moreover the comparison with the quantile based conditional risk measure estimates is not fair as they use forecasted quantities for a whole marginal window while the prior residual strategy has access to the realized data of the day before. We also performed traditional backtest on the risk measures based on the prior residuals and due to the already visible and theoretically anticipated exaggerations the resulting risk measures failed almost all traditional backtests. Just for comparison the risk measures based on the realized residuals performed very well on the traditional backtests, which is however as they are oracle estimators not too surprising. As further research one could try to truncate the prior day residuals for example with the ReLU function $f(z) = \max(0, z)$, which is prominently known from deep learning. But this could of course lead to too conservative risk measure estimates.

All in all the strategy based on the quantiles clearly seems to be the most promising one. This comes mainly from its good interpretability and robustness. Especially the comparison with the unconditionally estimated risk measures can lead to a good understanding how well the market risk based on fluctuations of the conditional asset is already accounted for in the unconditional approach.


11.2 Two conditional assets

After having covered the case with one conditional asset in much detail we would like to shortly display the very similar approach that involves two conditional assets. At first we considered the Spanish stock portfolio with the indices SP500 and Eurostoxx 50 as the two conditional assets. But as we have seen in the single conditional case the SP500 had almost no influence which led to the fact that the double conditional case basically ended up to be almost identical to the single conditional case with just the Eurostoxx 50 index considered. Thus we take the **Eurostoxx 50** index as the one conditional asset and the asset with the highest market capitalization **Iberdrola** as the second one. The remaining stocks of our Spanish portfolio will represent the base portfolio of interest. Besides that all parameters and time frames will be held exactly like in the single conditional case.

As the portfolio involves only stocks we have already covered we do not need to assess the quality of fit of the ARMA(1,1)-GARCH(1,1) models again. This means we can right away have a look at the fitted D-vines. The D-vine ordering is now based on Algorithm 4 and thus it will be also interesting which ordering the two conditional assets will have. Figure 11.18 shows the first trees of the D-vines corresponding to the first and last vine window of the first time frame. We can observe that over time the ordering was very much stable. Taking the second rightmost position the Eurostoxx 50 index has the strongest bivariate dependency with again the Banco Santander asset. Also the rest of the ordering reminds of the single conditional case.


Figure 11.19 displays the first trees of the D-vines corresponding to the first and last vine window of the second time frame. Also for the second time frame the Eurostoxx 50 index takes the second rightmost position but this time with the strongest bivariate dependence being with the Amadeus It Group which has a more important role during the pandemic as we have uncovered also in previous analyses.

Tree 1



(a) The first vine window.


Tree 1



(b) The sixth/last vine window.


Figure 11.18 The first trees of the D-vine corresponding to the first and sixth/last vine window for the first time frame (2016-2019) in the double conditional setting. The conditional assets are the **Eurostoxx 50 index** and **Iberdrola**.

Tree 1



(a) The first vine window.


Tree 1




(b) The sixth/last vine window.

Figure 11.19 The first trees of the D-vine corresponding to the first and second/last vine window for the second time frame (2020-2021) in the double conditional setting. The conditional assets are the **Eurostoxx 50 index and Iberdrola**.

Subsequently we have a short look at the estimated conditional risk estimates and as the procedure is also the same as in the single conditional case we will just present the comparison of the quantile based conditional ES estimates at level $\alpha = 0.025$ with the ones obtained by the realized residual strategy in Figure 11.20 for both time frames.



(a) The first time frame (2016-2019).



(b) The second time frame (2020-2021).

Figure 11.20 Comparisons of the **quantile based conditional ES estimates** at $\alpha = 0.025$ (blue) with the ones based on the **realized residuals** (green) for the first and second time frame. The conditioning assets are the **Eurostoxx 50 index and Iberdrola**.

As already evident from the fitted D-vines of course the two conditional assets have a major influence on the ES estimates. The influence is also stronger than for example the influence of just the Eurostoxx 50 index as the conditioning asset as both conditional assets have a strong dependence with the base portfolio. All additional analyses and comparisons performed for the single conditional case can be performed as well for the double conditional approach and thus it provides an extended toolset to uncover and measure the influence of important market players or indices simultaneously on the risk measures of interest.

Conclusion

In this master thesis, we have extended existing and proposed new methods to effectively measure the market risk of potentially large financial portfolios. Our proposed unconditional risk measure estimation approach is based on univariately modelling each asset given as a financial time series in order to capture its trend and volatility with ARMA-GARCH models, then using the flexible class of R-vine copulas to model the cross dependence within the portfolio and to finally estimate various risk measures on the portfolio level following a Monte Carlo approach. The whole estimation is performed in a rolling window fashion. This approach extends existing approaches like the one of Maarouf (2021) by incorporating the estimation of the expected shortfall, which as of the Basel III accords is the primary risk measure for market risk.

Furthermore we introduced a novel conditional estimation approach. Having one or two additional market indices or other main market players we modelled them also univariately and incorporated them in our dependence model, which in the conditional case is a D-vine copula. Then we simulated conditionally on the index or indices portfolio level forecasts based on which we estimated conditional risk measures. We proposed as the conditioning values quantiles of the conditioning asset or assets, which allows to observe the behavior of the risk measures conditioned on different states of the conditioning market index or indices.

To facilitate the conditional approach we have introduced novel algorithms to not only build up a D-vine with a fixed position of the conditioning asset or assets but most importantly algorithms to sample conditionally on the rightmost or two rightmost leafs from a D-vine copula. In addition a detailed discussion of applicable backtesting strategies for both the VaR but most notably the ES was provided.

All of these approaches were put to the test in multiple case studies using the developed R package *portvine*, which efficiently implements all proposed risk measure estimation approaches and is publicly available. The implemented risk measure estimation algorithms were also optimized for computational efficiency and for example allow as a consequence to be processed highly parallelized. We used a portfolio based on the top constituents of the MSCI Spain index for the case studies and the main findings were the following.


- In almost all examined cases the unconditional risk measure estimation procedure proposed in this thesis passed all considered traditional backtests for both risk measures of interest. Only in very few cases the estimation was slightly too conservative which is generally the desirable way to fail a backtest.
- For the considered time frame 2, that is set during the Covid-19 pandemic (02.04.2020-13.10.2021) and exhibits an increased market volatility, using comparative backtests we found clear indications that using the flexible class of R-vine copulas provides superior risk measure estimation than using vine copulas that only consider Gaussian or Student's t components for our proposed unconditional estimation approach.
- Concerning the conditional risk measure estimation approach, we have shown that not only the resulting conditional risk measure estimates but also the analysis of the dependence model i.e. the D-vines can create valuable insights about the influence of the conditional asset(s) on the risk measures of interest.
- The comparison of the quantile based conditional risk measure estimates with the unconditionally estimated risk measures can lead to a good understanding how well the market risk based on fluctuations of the conditional asset(s) is already accounted for in the unconditional approach. In this manner we for example found strong indications that the considered portfolio of Spanish stocks might be a good hedge against fluctuations in the American market.

All in all, both the unconditional and conditional estimation approaches proposed in this theses performed really well in the presented case studies. Having published the R package *portvine* it is now easy to try out and incorporate these methods for risk measure estimation yourself.


Further research is possible in multiple directions. For instance, one could explore different conditional strategies like simulations with fixed conditional series, although fixed conditioning simulations would probably result in an explosion of runtime. Also the extension of the conditional approach to other vine classes as the dependence model would be possible if one can provide a generic conditional sampling algorithm for the vine class. Finally, only a hand full of use cases of the proposed conditional risk measure estimates were presented in this thesis and one could introduce further comparisons or metrics involving those, that help portfolio managers making better decisions concerning the assessment of market risk.

A Appendix

Plots for the assessment of quality for the marginal models




(a) Residual plots for the mean equation.




(b) Residual plots for the volatility equation.

Figure A.1 The residual plots for the asset Banco Santander for the first marginal window and the first time frame (2016-2019).

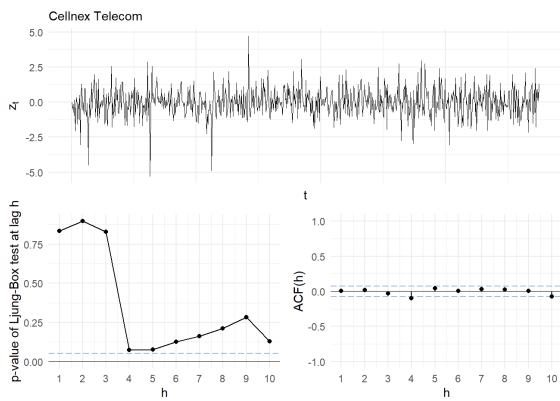


(a) Residual plots for the mean equation.

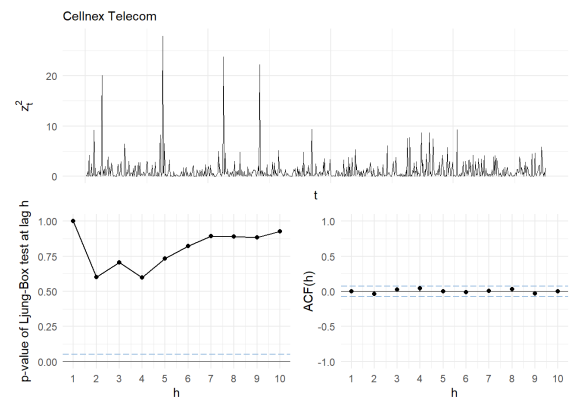


(b) Residual plots for the volatility equation.

Figure A.2 The residual plots for the asset Inditex for the first marginal window and the first time frame (2016-2019).




(a) Residual plots for the mean equation.




(b) Residual plots for the volatility equation.

Figure A.3 The residual plots for the asset Cellnex Telecom for the first marginal window and the first time frame (2016-2019).




(a) Residual plots for the mean equation.

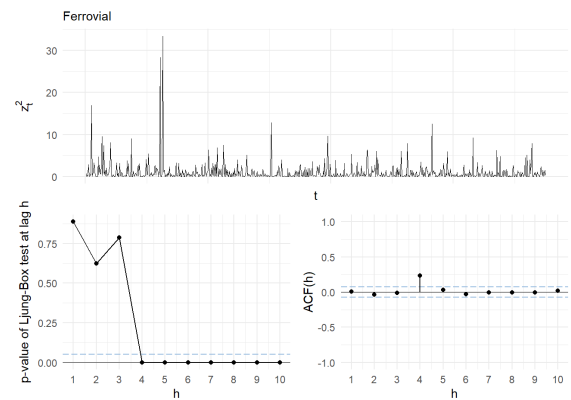


(b) Residual plots for the volatility equation.

Figure A.4 The residual plots for the asset Repsol Ypf for the first marginal window and the first time frame (2016-2019).




(a) Residual plots for the mean equation.

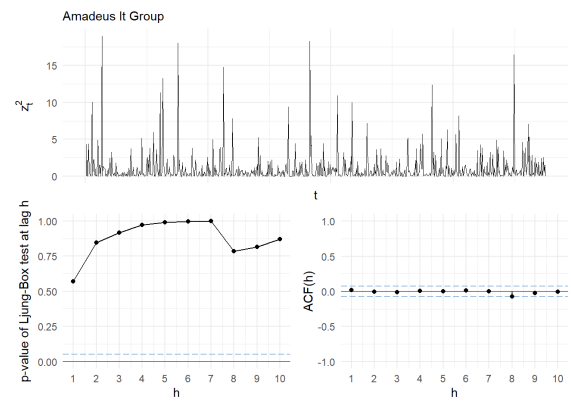


(b) Residual plots for the volatility equation.

Figure A.5 The residual plots for the asset Ferrovial for the first marginal window and the first time frame (2016-2019).




(a) Residual plots for the mean equation.




(b) Residual plots for the volatility equation.

Figure A.6 The residual plots for the asset Amadeus It Group for the first marginal window and the first time frame (2016-2019).




(a) Residual plots for the mean equation.




(b) Residual plots for the volatility equation.

Figure A.7 The residual plots for the asset Telefonica for the first marginal window and the first time frame (2016-2019).




(a) Residual plots for the mean equation.




(b) Residual plots for the volatility equation.

Figure A.8 The residual plots for the asset BBV Argentina for the first marginal window and the first time frame (2016-2019).




(a) Residual plots for the mean equation.




(b) Residual plots for the volatility equation.

Figure A.9 The residual plots for the asset Banco Santander for the first marginal window and the second time frame (2020-2021).




(a) Residual plots for the mean equation.




(b) Residual plots for the volatility equation.

Figure A.10 The residual plots for the asset Inditex for the first marginal window and the second time frame (2020-2021).

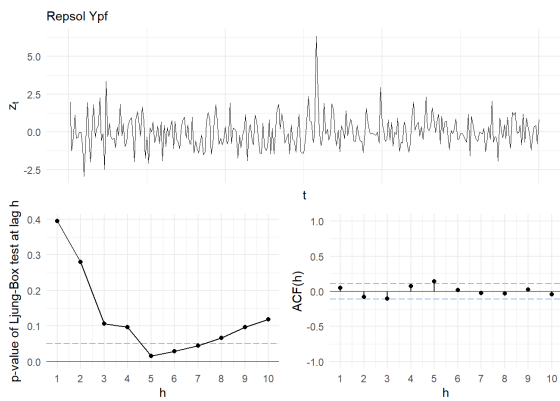


(a) Residual plots for the mean equation.

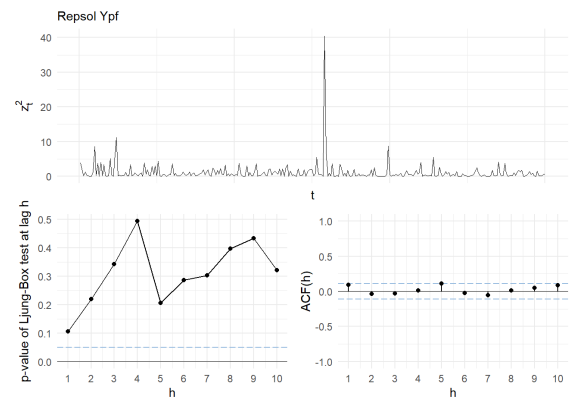


(b) Residual plots for the volatility equation.

Figure A.11 The residual plots for the asset Cellnex Telecom for the first marginal window and the second time frame (2020-2021).




(a) Residual plots for the mean equation.




(b) Residual plots for the volatility equation.

Figure A.12 The residual plots for the asset Repsol Ypf for the first marginal window and the second time frame (2020-2021).




(a) Residual plots for the mean equation.




(b) Residual plots for the volatility equation.

Figure A.13 The residual plots for the asset Ferrovial for the first marginal window and the second time frame (2020-2021).




(a) Residual plots for the mean equation.




(b) Residual plots for the volatility equation.

Figure A.14 The residual plots for the asset Amadeus It Group for the first marginal window and the second time frame (2020-2021).




(a) Residual plots for the mean equation.




(b) Residual plots for the volatility equation.

Figure A.15 The residual plots for the asset Telefonica for the first marginal window and the second time frame (2020-2021).




(a) Residual plots for the mean equation.




(b) Residual plots for the volatility equation.

Figure A.16 The residual plots for the asset BBV Argentaria for the first marginal window and the second time frame (2020-2021).




(a) Residual plots for the mean equation.




(b) Residual plots for the volatility equation.

Figure A.17 The residual plots for the conditional asset the SP500 index for the first marginal window and the second time frame (2020-2021).



(a) Residual plots for the mean equation.



(b) Residual plots for the volatility equation.

Figure A.18 The residual plots for the conditional asset the Eurostoxx 50 index for the first marginal window and the second time frame (2020-2021).

Additional material corresponding to the fitted D-vines in the single conditional case

Copula family	Number of usage	Copula family	Number of usage
Gaussian	2	Independence	12
Student's t	4	Clayton	5
Frank	6	Gumbel	2
Joe	7	BB1	3
BB7	3	BB8	1

(a) The first vine window.

Copula family	Number of usage	Copula family	Number of usage
Gaussian	5	Independence	10
Student's t	5	Clayton	2
Frank	7	Gumbel	4
BB1	2	BB7	3
BB8	1	Joe	6

(b) The second vine window.

Table A.1 The used copula families alongside their frequencies for the first and second fitted D-vine for the second time frame (2020-2021) in the single conditional setting. The conditional asset is the SP500.

Tree level	Copula	Copula family	Empirical Kendall's tau
1	$C_{j_1, I}$	Joe	-0.0566
2	$C_{j_2, I; j_1}$	Independence	0
3	$C_{j_3, I; j_1, j_2}$	Joe	-0.0675
4	$C_{j_4, I; j_1: j_3}$	Joe	0.0972
5	$C_{j_5, I; j_1: j_4}$	Independence	0
6	$C_{j_6, I; j_1: j_5}$	Independence	0
7	$C_{j_7, I; j_1: j_6}$	Independence	0
8	$C_{j_8, I; j_1: j_7}$	Independence	0
9	$C_{j_9, I; j_1: j_8}$	Independence	0

Table A.2 The used copula families alongside their empirical Kendall's tau associated with the conditional asset for the first fitted D-vine for the second time frame (2020-2021) in the single conditional setting. The conditional asset is the SP500. (j_1 = Iberdrola, j_2 = Ferrovial, j_3 = Inditex, j_4 = Repsol Ypf, j_5 = Amadeus It Group, j_6 = Banco Santander, j_7 = BBV Argentaria, j_8 = Telefonica, j_9 = Cellnex Telecom)

Copula family	Number of usage	Copula family	Number of usage
Gaussian	5	Independence	6
Student's t	4	Clayton	2
Frank	11	Gumbel	6
Joe	5	BB1	1
BB7	5		

(a) The first vine window.

Copula family	Number of usage	Copula family	Number of usage
Gaussian	5	Independence	7
Student's t	5	Clayton	5
Frank	6	Gumbel	6
BB1	2	BB7	4
BB8	1	Joe	4

(b) The second vine window.

Table A.3 The used copula families alongside their frequencies for the first and second fitted D-vine for the second time frame (2020-2021) in the single conditional setting. The conditional asset is the Eurostoxx 50.

Tree level	Copula	Copula family	Empirical Kendall's tau
1	$C_{j_1, I}$	Gaussian	0.4461
2	$C_{j_2, I; j_1}$	Frank	0.3078
3	$C_{j_3, I; j_1, j_2}$	Frank	0.1747
4	$C_{j_4, I; j_1, j_3}$	Gumbel	0.1528
5	$C_{j_5, I; j_1, j_4}$	Independence	0
6	$C_{j_6, I; j_1, j_5}$	Gaussian	-0.0902
7	$C_{j_7, I; j_1, j_6}$	Independence	0
8	$C_{j_8, I; j_1, j_7}$	Frank	0.3420
9	$C_{j_9, I; j_1, j_8}$	Joe	0.0404

Table A.4 The used copula families alongside their empirical Kendall's tau associated with the conditional asset for the first fitted D-vine for the second time frame (2020-2021) in the single conditional setting. The conditional asset is the Eurostoxx 50. (j_1 = Amadeus It Group, j_2 = Inditex, j_3 = Ferrovial, j_4 = Banco Santander, j_5 = Repsol Ypf, j_6 = BBV Argentaria, j_7 = Telefonica, j_8 = Iberdrola, j_9 = Cellnex Telecom)

Behavior of the quantile based conditioning values for the second time frame (2020-2021)




Figure A.19 Behavior of the **quantile based conditioning values** of the conditioning asset SP500 for the second time frame (2020-2021).




Figure A.20 Behavior of the **quantile based conditioning values** of the conditioning asset Eurostoxx 50 for the second time frame (2020-2021).

Visualizations of the single conditional risk measure estimates for the second time frame (2020-2021)




Figure A.21 Quantile based conditional ES estimates at $\alpha = 0.025$ with the conditioning asset SP500 for the second time frame (2020-2021).






Figure A.22 Comparison of the quantile based conditional ES estimates at $\alpha = 0.025$ (blue) with the ones based on the **realized residuals** (green) for the second time frame (2020-2021). The conditioning asset is the SP500.



(a) Unconditional approach using a D-vine.



(b) Unconditional approach using a R-vine.

Figure A.23 Comparison of the **quantile based conditional ES estimates** at $\alpha = 0.025$ (blue) with the **unconditional ES estimates** (orange/gold) for the second time frame (2020-2021). The conditioning asset is the SP500.




Figure A.24 Quantile based conditional ES estimates at $\alpha = 0.025$ with the conditioning asset Eurostoxx50 for the second time frame (2020-2021).






Figure A.25 Comparison of the quantile based conditional ES estimates at $\alpha = 0.025$ (blue) with the ones based on the realized residuals (green) for the second time frame (2020-2021). The conditioning asset is the Eurostoxx 50.



(a) Unconditional approach using a D-vine.



(b) Unconditional approach using a R-vine.

Figure A.26 Comparison of the **quantile based conditional ES estimates** at $\alpha = 0.025$ (blue) with the **unconditional ES estimates** (orange/gold) for the second time frame (2020-2021). The conditioning asset is the Eurostoxx 50.

Bibliography

- Acerbi, C. and Székely, B. (2014). Backtesting expected shortfall introducing three model-independent, non-parametric back-test methodologies for expected shortfall.
- Acerbi, C. and Tasche, D. (2002). Expected shortfall: A natural coherent alternative to value at risk. *Economic Notes*, 31:379–388.
- Akaike, H. (1998). Information Theory and an Extension of the Maximum Likelihood Principle. In Parzen, E., Tanabe, K., and Kitagawa, G., editors, *Selected Papers of Hirotugu Akaike*, Springer Series in Statistics, pages 199–213. Springer New York. DOI: 10.1007/978-1-4612-1694-0_15.
- Artzner, P., Delbaen, F., Eber, J.-M., and Heath, D. (1999). Coherent measures of risk. *Mathematical Finance*, 9:203–228.
- Basel Committee (2016). Minimum capital requirements for market risk. Available at <https://www.bis.org/publ/bcbs265.pdf>. Basel Committee on Banking Supervision Consultative Document.
- Basel Committee (2017). Pillar 3 disclosure requirements – consolidated and enhanced framework. Available at <https://www.bis.org/publ/bcbs265.pdf>. Basel Committee on Banking Supervision Consultative Document.
- Bayer, S. and Dimitriadis, T. (2020a). *esback: Expected Shortfall Backtesting*. R package version 0.3.0.
- Bayer, S. and Dimitriadis, T. (2020b). Regression-Based Expected Shortfall Backtesting. *Journal of Financial Econometrics*.
- Bedford, T. and Cooke, R. M. (2002). Vines—a new graphical model for dependent random variables. *The Annals of Statistics*, 30(4):1031 – 1068.
- Brechmann, E. C. and Czado, C. (2013). Risk management with high-dimensional vine copulas: An analysis of the Euro Stoxx 50. *Statistics & Risk Modeling*, 30(4):307–342.
- Christoffersen, P., Hahn, J., and Inoue, A. (2001). Testing and comparing value-at-risk measures. *Journal of Empirical Finance*, 8.
- Christoffersen, P. F. (1998). Evaluating interval forecasts. *International Economic Review*, 39(4).
- Czado, C. (2019). *Analyzing Dependent Data with Vine Copulas - A Practical Guide With R*. Springer, Berlin, Heidelberg.
- Dimitriadis, T. and Bayer, S. (2019). A joint quantile and expected shortfall regression framework. *Electronic Journal of Statistics*, 13(1):1823 – 1871.
- Dißmann, J., Brechmann, E. C., Czado, C., and Kurowicka, D. (2013). Selecting and estimating regular vine copulae and application to financial returns. *Computational Statistics and Data Analysis*, 59:52–69.
- Dißmann, J. F. (2010). Statistical inference for regular vines and application. Diploma thesis, Technical University Munich.
- Efron, B. and Tibshirani, R. (1994). *An Introduction to the Bootstrap*. Chapman & Hall/CRC Monographs on Statistics & Applied Probability. Taylor & Francis.

- Embrechts, P., Klüppelberg, C., and Mikosch, T. (1997). *Modelling Extremal Events for Insurance and Finance*. Springer-Verlag.
- Fissler, T. and Ziegel, J. F. (2016). Higher order elicibility and osband's principle. *The Annals of Statistics*, 44(4).
- Ghalanos, A. (2022). *rugarch: Univariate GARCH models*. R package version 1.4-6.
- Joe, H. (1996). Families of m -variate distributions with given margins and $m(m-1)/2$ bivariate dependence parameters. *Lecture Notes-Monograph Series*, 28:120–141.
- Kim, S. (2015). Ppcor: An r package for a fast calculation to semi-partial correlation coefficients. *Communications for Statistical Applications and Methods*, 22:665–674.
- Maarouf, M. (2021). Backtesting value-at-risk of financial data using vine copulas. Master's thesis, Technical University Munich.
- McNeil, A. J. and Frey, R. (2000). Estimation of tail-related risk measures for heteroscedastic financial time series: an extreme value approach. *Journal of Empirical Finance*, 7(3).
- Monstvilaite, M. (2016). Portfolio value-at-risk using regular vine copulas. Master's thesis, Stockholms universitet.
- Nagler, T. and Vatter, T. (2021). *rvinecopulib: High Performance Algorithms for Vine Copula Modeling*. R package version 0.6.1.1.1.
- Nelsen, R. B. (2010). *An Introduction to Copulas*. Springer Publishing Company, Incorporated.
- Nolde, N. and Ziegel, J. F. (2017). Elicibility and backtesting: Perspectives for banking regulation. *The Annals of Applied Statistics*, 11(4).
- Pfaff, B. (2016). *Financial Risk Modelling and Portfolio Optimization with R*. Wiley.
- R Core Team (2021). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria.
- Roccioletti, S. (2016). *Backtesting Value at Risk and Expected Shortfall*. BestMasters. Springer Fachmedien Wiesbaden.
- Rosenblatt, M. (1952). Remarks on a Multivariate Transformation. *The Annals of Mathematical Statistics*, 23(3):470 – 472.
- Schwarz, G. (1978). Estimating the Dimension of a Model. *The Annals of Statistics*, 6(2):461 – 464.
- Sharma, C. and Sahni, N. (2021). A mutual information based r-vine copula strategy to estimate var in high frequency stock market data. *PLOS ONE*, 16(6):1–16.
- Shumway, R. and Stoffer, D. (2017). *Time Series Analysis and Its Applications: With R Examples*. Springer Texts in Statistics. Springer International Publishing, 4 edition.
- Spokoiny, V. and Dickhaus, T. (2015). *Basics of Modern Mathematical Statistics*. Springer Texts in Statistics. Springer Berlin Heidelberg.
- Tasche, D. (2002). Expected shortfall and beyond. *Journal of Banking and Finance*, 26:1519–1533.
- Trucíos, C., Tiwari, A., and Alqahtani, F. (2019). Value-at-risk and expected shortfall in cryptocurrencies' portfolio: A vine copula-based approach. *SSRN Electronic Journal*.
- Tsay, R. (2010). *Analysis of Financial Time Series*. CourseSmart. Wiley.

Whittaker, J. (1990). *Graphical models in applied multivariate statistics*. Wiley Publishing.

Wickham, H., Averick, M., Bryan, J., Chang, W., McGowan, L. D., François, R., Grolemund, G., Hayes, A., Henry, L., Hester, J., Kuhn, M., Pedersen, T. L., Miller, E., Bache, S. M., Müller, K., Ooms, J., Robinson, D., Seidel, D. P., Spinu, V., Takahashi, K., Vaughan, D., Wilke, C., Woo, K., and Yutani, H. (2019). Welcome to the tidyverse. *Journal of Open Source Software*, 4(43):1686.

Yule, G. U. and Kendall, M. G. (1950). *An introduction to the theory of statistics*. Charles Griffin London, 14th edition.