PEARL

04 University of Plymouth Research Theses

01 Research Theses Main Collection

2022

Portfolio Optimisation Using Genetic Algorithms and Copula Financial Models

Alotaibi, Tahani

http://hdl.handle.net/10026.1/19632

http://dx.doi.org/10.24382/631 University of Plymouth

All content in PEARL is protected by copyright law. Author manuscripts are made available in accordance with publisher policies. Please cite only the published version using the details provided on the item record or document. In the absence of an open licence (e.g. Creative Commons), permissions for further reuse of content should be sought from the publisher or author.

Copyright Statement

This copy of the thesis has been supplied on condition that anyone who consults it is understood to recognise that its copyright rests with its author and that no quotation from the thesis and no information derived from it may be published without the author's prior consent.



Portfolio Optimisation Using Genetic Algorithms and Copula Financial Models

by

Tahani Sultan Alotaibi

A thesis submitted to the University of Plymouth in partial fulfilment for the degree of

DOCTOR OF PHILOSOPHY

School of Engineering, Computing and Mathematics

September 2022

Acknowledgements

F^{IRST} and foremost, I would like to express my immense gratitude and appreciation to my first supervisor and director of this study Dr Matthew Craven, whose support throughout the process has been immeasurable, and without whom the completion of this study would not have been possible. His generous dedication of his valuable time to support this process, his valuable and timely feedback, his knowledgeable comments and suggestions to improve the quality of this thesis, are all greatly appreciated. It has been an absolute honour to be supervised by Dr Matthew Craven.

I also wish to gratefully thank my second and third supervisors, Dr Luciana Dalla Valle and Dr Panagiotis Tziogkidis (University of Macedonia), for helping me along this path and pointing me towards useful guides and invaluable advice. Thank you also to my fourth supervisor Dr Colin Christopher, for his support.

I would also like to express my heartfelt gratitude to my ex-supervisor Professor Stephen Huggett for his encouragement and his open door. From the moment I met him, he has motivated, encouraged, guided and supported me to pursue my goals.

My special thanks also go to my ex-supervisors Dr David Graham and Dr Ana Paula Palacios, for their guidance and assistance. I also would like to thank Dr Malgorzata Wojtys for her assistance and useful suggestions.

I am very grateful to the members of my thesis committee, external examiner Professor Konstantinos Syriopoulos and internal examiner Professor David McMullan, for providing me with valuable comments on this thesis. I also would like to thank Prof Alison Raby as the Chair of the committee.

My thanks also go to all my friends and colleagues in the School of Computing, Electronics and Mathematics, and particularly to Zainab Bu Sinnah who was always there when I needed assistance. Also, a huge thank you to Shaqra University and the Saudi Cultural Bureau for their generosity in financially supporting this PhD journey. My deepest gratitude goes to my fantastic parents, who have been a constant source of support in every aspect of my life, including my studies. Words cannot express my debt to my parents, who have not only given me life but have supported and encouraged every single thing I have done.

Extra special thanks go to my supportive and amazing husband, who deserves my sincere gratitude for his unconditional love, understanding and forbearance. Unique and infinite love also to our two beautiful children Fahad and Lna, who I hope will be compensated for some of what this PhD journey has taken from them by its potential success.

My loving thanks to all my brothers and sisters for their continued support, unconditional love and encouragement. I could not have done this without your love, prayers and support. I would like to dedicate this thesis to my supportive family.

Tahani

Author's Declaration

A^T no time during the registration for the degree of Doctor of Philosophy has the author been registered for any other University award without prior agreement of the Doctoral College Quality Sub-Committee. Work submitted for this research degree at University of Plymouth has not formed part of any other degree either at University of Plymouth or at another establishment.

Presentations:

- Large Scale Portfolio Optimisation Using Heuristic Algorithms Built Around Copula Financial Models, poster presented at Postgraduate Society Research Showcase, 4th December 2017, University of Plymouth.
- Efficient Frontiers in Portfolio Optimisation with Minimum Proportion Constraints, poster presented at the Genetic and Evolutionary Computation Conference held in Prague (Czech Republic), 13th–17th July 2019.
- Constructing Class-Based Portfolios on Gulf Markets with Metaheuristics, presented at the OR'61 conference, University of Kent, Canterbury (UK), 3rd–5th September 2019.

Workshops and Conferences Attended:

- "Evostar 2018" conference, Parma, Italy, 4th-6th April 2018.
- "Heuristics and Approximation Algorithms" NATCOR PhD workshop at the University of Nottingham, UK, 9th-13th April 2018.

Publications:

Alotaibi, T. S. & Craven, M. J., Efficient Frontiers in Portfolio Optimisation with Minimum Proportion Constraints, in "Proceedings of the Companion Publication of the 2019 Annual Conference on Genetic and Evolutionary Computation", ACM (2019), 358-359.

Alotaibi, T. S., Tziogkidis, P. & Craven, M. J., *Constructing Class-Based Portfolios on Gulf Markets with Metaheuristics,* extended abstract presented at OR'61 (2019).

Three further publications are forthcoming, based on the three main chapters in this thesis. These will be submitted to international journals shortly.

Word count for the main body of this thesis: 42042

Signed:

Date: 20/04/2022

Portfolio Optimisation Using Genetic Algorithms and Copula Financial Models

Tahani Sultan Alotaibi

Abstract

P^{ORTFOLIO} optimisation aims to efficiently find optimal proportions of portfolio assets, given certain constraints, and has been well-studied (Jin et al. 2016). While portfolio optimisation ascertains asset combinations most suited to investor requirements, numerous real-world problems impact its simplicity, e.g., investor preferences. Trading restrictions are also commonly faced, and must be met. However, in adding constraints to Markowitz's basic mean-variance model, problem complexity increases, causing difficulties for exact optimisation approaches to find large problem solutions inside reasonable timeframes. The concern of this study concern is portfolio optimisation under three real-world constraints, using Markowitz's mean-variance portfolio optimisation problem to obtain optimal risk-return trade-off (the efficient frontier).

A genetic algorithm (GA) is formulated, efficiently solving cardinality-constrained portfolio optimisation problems and generating the same quality of results as the literature. It is then applied to minimum proportion cardinality-constrained problems, and experimental results are presented based on historical daily financial markets of five benchmark datasets from five recognised indices of the OR-library. Results show that increasing cardinality (K) from 2 to 10 gives higher solution quality while computational time increases significantly with K. Additionally, applying a high minimum proportion leads to shape differences in efficient frontier curves. Also, solutions tend to be restricted to smaller portions of efficient frontiers, producing lower convergence and spread.

The high-dimensional problem of constructing optimal risky portfolios concerns investors looking for maximum reward to variability ratios. Thus, we consider a GA operating on portfolio optimisation problems subject to class constraints, rather than cardinality and minimum proportion constraints. GA usage in constructing optimal portfolios in Gulf Cooperation Council (GCC) stock markets, incorporating precious metals and oil, is investigated. Results show that combining stock assets, oil and precious metals reduces risk, creating further investment and diversification opportunities. Shifting yearly frontiers arising from the GCC dataset are also investigated, showing that inward/upward shifts denote higher expected returns for the same risk levels, though economic events may affect frontiers.

Finally, Worst-Case Value-at-Risk is considered as an alternative risk measure using copulas on the GCC dataset, allowing a more accurate investment risk assessment to be applied, compared with traditional methods such as Value-at-Risk. To ameliorate the large computation times required to find solutions to these problems under the required quantities of assets, the researcher utilised the University of Plymouth High Performance Computer cluster.

Contents

Ac	cknov	vledgements	iii
Αι	uthor	's Declaration	v
Ał	ostra	ct	vii
Li	st of	Figures	xiii
Li	st of	Tables	xvii
Li	st of	Acronyms	xix
1	Intro	oduction	1
	1.1	Methods for Computing Solutions to Hard Problems	3
	1.2	The Motivation, Contributions and Aims of The Thesis	10
	1.3	Objectives of the Thesis	13
	1.4	Overview of Chapters	13
2	Opt	imisation of Portfolio Risk	15
	2.1	Introduction	15
	2.2	Markowitz Mean-Variance Problem	15
	2.3	Quadratic Programming	17
	2.4	Application of QP to the Markowitz Mean-Variance Problem	18
	2.5	Benchmark Datasets of the Portfolio Optimisation Problem	19
	2.6	Practical Constraints	20

		2.6.1	Cardinality Constraint	20
		2.6.2	Class Constraint	21
		2.6.3	Rounded or Minimal Lots	21
		2.6.4	Bounds Upon Proportions	22
		2.6.5	Pre-assignment Constraint	22
		2.6.6	Literature on the Effects of Imposing Constraints	22
	2.7	Measu	ures of Risk	23
		2.7.1	Pre-Markowitz Risk Measures	24
		2.7.2	Modern Portfolio Theory (MPT)	24
		2.7.3	Value at Risk	24
		2.7.4	Coherent Risk Measures	27
	2.8	Summ	nary of the Chapter	29
3	Port	tfolio O	potimisation with Cardinality and Minimum Proportion Constrain	ts 30
	3.1	Introd	uction	30
	3.1 3.2	Introd	uction	30 30
	3.1 3.2	Introde Algorit	uction	30 30 31
	3.1 3.2	Introdu Algorit 3.2.1	uction	30 30 31 31
	3.13.23.3	Introdu Algorit 3.2.1 3.2.2 GA or	uction	30 30 31 31 38
	3.13.23.3	Introdu Algorit 3.2.1 3.2.2 GA or 3.3.1	uction	30 30 31 31 38 39
	3.13.23.3	Introdu Algorit 3.2.1 3.2.2 GA or 3.3.1 3.3.2	uction	 30 30 31 31 38 39 39 39
	 3.1 3.2 3.3 3.4 	Introdu Algorit 3.2.1 3.2.2 GA or 3.3.1 3.3.2 GA or	uction	 30 30 31 31 38 39 39 44
	3.13.23.33.4	Introdu Algorit 3.2.1 3.2.2 GA or 3.3.1 3.3.2 GA or 3.4.1	uction	 30 30 31 31 38 39 39 44 45
	3.13.23.33.4	Introdu Algorit 3.2.1 3.2.2 GA or 3.3.1 3.3.2 GA or 3.4.1	uction	 30 30 31 31 38 39 39 44 45 49

	3.5	GA on the Cardinality Constrained Problem with Minimum Asset Propor-
		tions
		3.5.1 Experimental Results 57
		3.5.2 Further Runs with Non-Zero Minimum Proportions 62
		3.5.3 Further Experiments of Increasing x_{\min}
	3.6	Summary of the Chapter
4	Con	structing Class-Based Portfolios in Gulf Markets with a Genetic Algo-
	rithr	n 83
	4.1	Introduction
	4.2	Contextual Background
	4.3	Literature Review
	4.4	Data Description
	4.5	Approach and Algorithm Setup 100
		4.5.1 Formulation and Constraint Handling
		4.5.2 Determining the Step Penalty
		4.5.3 Initial Configurations
	4.6	Results
		4.6.1 Random Initial Configurations
		4.6.2 Robustness Exercise
		4.6.3 Comparing the Current Approach with the Solely Cardinality-Constrained
		Approach
	4.7	An Investigation of Shifting Yearly Frontiers
	4.8	Summary of the Chapter
5	Port	folio Optimisation Using the Copula-Based WCVaR Approach 152

5.1	Introdu		152
5.2	Copula	as	155
	5.2.1	The Fréchet-Hoeffding bounds	157
5.3	Deper	dence Measures	157
	5.3.1	Pearson's Correlation Coefficient ρ	157
	5.3.2	Kendall's τ Coefficient	158
	5.3.3	Spearman's Rank Correlation Coefficient ρ_s	158
5.4	Popula	ar Copula Families	158
5.5	Tail De	ependence	165
5.6	Copula	a Inference	165
	5.6.1	Maximum Likelihood Method	165
	5.6.2	Inference Functions for Margins (IFM)	166
	5.6.3	Canonical Maximum Likelihood Method (CML)	166
	5.6.4	Test of Multivariate Independence	166
	5.6.5	Goodness-of-fit Test	167
5.7	The G	ARCH (p,q) Model	167
	5.7.1	Basic Properties	168
5.8	Copula	a-GARCH Model	169
5.9	Worst	Case Conditional Value at Risk	170
	5.9.1	Worst Case GARCH-Copula CVaR Portfolio Optimisation	171
	5.9.2	An Example of Worst Case GARCH-Copula CVaR	171
	5.9.3	Computing the WCVaR and VaR for a Large Dataset	177
5.10	Summ	ary of the Chapter	178

Re	References			
	6.2	Further Work	184	
	6.1	Main Findings and Conclusions	180	

List of Figures

2.1	Computed EF and the EF of Beasley to dataset (D5) of Beasley (2000).	19
2.2	Sample R code to compute the VaR of ^GSPC during the period 01/01/200 13/04/2017.	0- 26
2.3	Distribution of returns for ^GSPC ticker with VaR marked with a dotted line.	26
3.1	Common representations used by GAs.	34
3.2	GA Output for dataset (D1), 50 points.	41
3.3	Heatmap for dataset (D1), 50 points.	41
3.4	GA Output for dataset (D5), 100 points.	43
3.5	Heatmap for dataset (D5), 100 points. Axis labels have been removed for clarity due to the large numbers of points and assets.	43
3.6	GA output (left) and solution-generation plot (right) for dataset (D2) with	
	<i>K</i> = 5	53
3.7	Heatmap for dataset (D2) with $K = 5$	54
3.8	GA Output for dataset (D5) with $K = 2, 5, 10. \dots$	55
3.9	Plots for dataset (D4) with $K = 3$ and minimum proportion $x_{\min} = 0.3$.	59
3.10) Plots for dataset (D4) with $K = 3$ and minimum proportion $x_{\min} = 0.01$.	60
3.11	Plots for dataset (D4) with $K = 3$ and minimum proportion $x_{\min} = 0.3$ (Alotaibi & Craven 2019).	61
3.12	Plots for dataset (D3) with different values of <i>K</i> and minimum proportion $x_{\min} = 0.01.$	64

3.13	Plots for dataset (D5) with different values of <i>K</i> and minimum proportion $x_{\min} = 0.01.$	67
3.14	Plots for dataset (D3) with different values of <i>K</i> and minimum proportion $x_{\min} = 0.1.$	71
3.15	Plots for dataset (D3) with $K = 2$ and minimum proportions $x_{min} = 0.01$, 0.1, 0.2, 0.25 and 0.3.	73
3.16	Plots for dataset (D3) with $K = 3$ and respective minimum proportions $x_{\min} = 0.01, 0.1, 0.2, 0.25, and 0.3.$	76
3.17	Plots for dataset (D5) with $K = 5$ and minimum proportions $x_{\min} = 0.01$, 0.05, 0.1, 0.15, 0.18 and 0.2.	79
4.1	Time series plots of prices for oil, gold, silver and copper commodities (black lines) and the respective returns series (blue lines).	97
4.2	Time series plots of prices for the GCC stock indices (black lines) and the respective returns series (blue lines).	98
4.3	Output with a step penalty of 500 for GCC dataset, oil composition with 5 classes and $K = 20$ using one run at random.	106
4.4	Output with a step penalty of 10000 for GCC dataset, oil composition with 5 classes and $K = 20$ using one run at random.	106
4.5	Output with a step penalty of 2000 for GCC dataset, oil composition with 5 classes and $K = 20$ using one run at random.	107
4.6	GA output for the GCC dataset, oil and precious metals with 5 classes and $K = 20$ using random ICs	112
4.7	GA output for the GCC dataset, oil and precious metals with 5 classes and $K = 50$ using random ICs	113
4.8	Plots for comparison between values of <i>K</i> using five classes and random ICs.	115

4.9 GA output for GCC dataset , oil and precious metals w K = 20 using the same IC.	ith 5 classes and	20
4.10 GA output for GCC dataset , oil and precious metals w K = 50 using the same IC.	ith 5 classes and 12	22
4.11 Plots for comparison between values of K using five classical structures of the second structure of the second structure	asses and fixed IC. 12	24
4.12 Plots using three classes with $K = 20.$		28
4.13 Plots using seven classes with $K = 20. \dots \dots \dots$	13	0
4.14 Plots using ten classes with $K = 20$.	13	31
4.15 Plots using all thirteen classes with $K = 20$	13	3
4.16 Outputs for the GCC dataset, 50 points.	13	6
4.17 Plot of unconstrained efficient frontiers (UEFs) of the 0	GCC dataset, per	
year	13	17
4.18 GA output for the GCC dataset, with 5 classes and $K =$	= 20 14	.3
4.19 Oil and precious metals composition with 5 classes an	$d K = 20. \dots 14$	5
4.20 Portfolios with number of classes used superimpose	d in the relevant	
position with 5 classes and $K = 20. \dots \dots \dots$	14	7
4.21 Mean Annualised Sharpe Ratios	14	9
5.1 Normal copula.		60
5.2 Student's t-copula.		60
5.3 Clayton copula.		;2
5.4 Frank copula.		;2
5.5 Gumbel copula		63
5.6 A plot of simulated time series data from GARCH(1,1).		;9
5.7 Performance of price movements of each index over th	e historical record. 17	'2

5.8	Scatterplots between each pair of index returns in the lower-triangular	
	panels, histograms of each marginal in the diagonal and Pearson's cor-	
	relation coefficients in the upper-triangular panels.	173
5.9	The Student's t-quantile functions of the residuals of GARCH(1,1) for	
	each marginal.	173
5.10	Plots of autocorrelation functions (ACF)	174
5.11	Scatterplots between pseudo-observations of each pair of indexes in the	
	lower-triangular panels, histograms of each marginal in the diagonal and	
	Pearson's correlation coefficients in the upper-triangular panels.	175
5.12	? Trajectory of WCVaR and minimum-variance portfolio values.	176

List of Tables

3.1	The effect of increasing the number of generations and the 'run' parameter.	48
3.2	The effect of varying the number of points. The statistics for the collec-	
	tion of runs on 50 points are reproduced from those of Table 3.5	49
3.3	Changing the value of the cardinality constraint, <i>K</i> , from 2 up to 10 on	
	dataset (D1)	50
3.4	Changing the value of <i>K</i> from 2 to 10 on dataset (D2).	50
3.5	Changing the value of <i>K</i> from 2 to 10 on dataset (D3).	51
3.6	Changing the value of <i>K</i> from 2 to 10 on dataset (D4).	52
3.7	Changing the value of <i>K</i> from 2 to 10 on dataset (D5)	52
3.8	Running dataset (D3) with $x_{\min} = 0.01$ and the values of K shown	63
3.9	Dataset (D5) with $x_{\min} = 0.01$ and the values of K shown	66
3.10	CD and RR measurements for datasets (D3) [Table 3.8] and (D5) [Table	
	3.9] with $x_{\min} = 0.01$ and the values of <i>K</i> shown	68
3.11	Running dataset (D3) with $x_{\min} = 0.1$ and the values of <i>K</i> shown.	69
3.12	Values of CD and RR for dataset (D3) with $x_{\min} = 0.1$ and varying K	69
3.13	Running dataset (D3) with $K = 2$ and the values of x_{\min} shown.	72
3.14	Dataset (D3) with $K = 3$ and the values of x_{\min} shown.	74
3.15	CD and RR measurements for (D3) and $K = 2$, $K = 3$ with x_{\min} ranging	
	from 0.01 to 0.3.	75
3.16	Running dataset (D5) with $K = 5$ and the values of x_{\min} shown.	77
3.17	CD and RR measurements for (D5), $K = 5$, with minimum proportions	
	ranging from 0.01 to 0.3.	78

4.1	Stock market characteristics and macroeconomic data for GCC in 2020 (with the exception of oil revenue, the latest data of which is from 2019). Data from various sources, including World Bank and The Global Econ-	
	omy. Structure of the table replicated from (Maghyereh et al. 2017).	87
4.2	Number of assets in each class.	95
4.3	Descriptive statistics, from 1st January 2008 to 31st January 2018 for the returns of the GCC stock market indices.	99
4.4	Descriptive statistics, from 1st January 2008 to 31st January 2018 for the returns of oil, gold, silver and copper.	100
4.5	Statistics produced by running with the values of <i>K</i> shown.	109
4.6	CD measurement for GCC dataset and number of points that include oil and metals with the values of <i>K</i> shown.	110
4.7	Statistics produced by running with the values of <i>K</i> shown. The ICs were produced at random, with 0_n denoting a sequence of <i>n</i> zeros	117
4.8	CD measurement for GCC dataset and number of points that include oil and metals with the values of <i>K</i> shown.	118
4.9	Running with the number of classes shown.	126
4.10	CD measurement for GCC dataset and number of points that include oil and metals.	127
4.11	Statistics produced by running algorithm for each year	141
4.12	CD measurement for GCC dataset and number of points that include oil and metals for each year.	141
5.1	Characteristics of alternative copula structures	164
5.2	Comparison between traditional multivariate normal method and WC-VaR technique.	176
5.3	Generating VaR and WCVaR using WCVaR technique for the GCC datase	t.178

List of Acronyms

- M-V Mean-Variance
- LAM Limited Asset Markowitz
- EA Evolutionary Algorithm
- GA Genetic Algorithm
- SA Simulated Annealing
- TS Tabu Search
- ACO Ant Colony Optimisation
- **QP** Quadratic Programming
- **EF** Efficient Frontier
- **CCEF** Cardinality Constrained Efficient Frontier
- **UEF** Unconstrained Efficient Frontier
- MOEA Multi-Objective Evolutionary Algorithm
- GCC Gulf Cooperation Council
- HPC High Performance Computing
- VaR Value at risk
- MAD Mean Absolute Deviation
- ES Expected shortfall
- IGD Inverted Generational Distance
- **CD** Crowding Distance

- RR Range of Return
- **IC** Initial Configuration
- CVaR Conditional Value at Risk
- WCVaR Worst Case Conditional Value at Risk
- WCCVaR Worst Case CVaR
- cdf Cumulative Distribution Function
- pdf Probability Density Function
- **IFM** Inference Functions for Margins
- CML Canonical Maximum Likelihood Method
- GARCH Generalized AutoRegressive Conditional Heteroscedasticity
- **ACF** Autocorrelation Functions

Chapter 1

Introduction

A Portfolio is a collection of assets - stocks, commodities, bonds, and so on. Portfolio optimisation aims to efficiently choose optimal proportions of these portfolio assets, subject to given constraints. This is a well-studied area of finance (Jin et al. 2016). The first modern approach, known as Modern Portfolio Theory (MPT), solved modern "portfolio problems" and was introduced by Markowitz (Markowitz 1952). This mean-variance (M-V) model of a financial portfolio supposes that the historical market of the assets influences their future market (Lwin et al. 2014).

Portfolio management, in conceptual terms, identifies which combination of financial assets best fits an investor's requirements. But this simple concept is affected by a wide range of real-world problems, not least the needs and preferences of an investor. Most often, an investor prefers safety over risk where their income streams are concerned. However, if returns are expected to be high enough, a riskier option may be preferred. Investors have varying levels of risk aversion, which affect the level of risk they are willing to accept. However, what can actually be achieved is dependent on the market, and the range of assets available to them. Therefore, portfolio management is also considered to be the identification of combinations which an investor can adopt in order to find the ideal return-risk trade-off (Maringer 2008). In spite of their different attitudes and requirements, investors invariably want to maximise their returns at the same time as minimising their risk (Banihashemi & Navidi 2017).

When Markowitz published his theory of portfolio selection in 1952, an answer was finally provided to the question of which choices an investor should make regarding the allocation of funds. By first applying expected return statistical measures and standard deviation, Markowitz showed the quantification of a security's risk and return. He then

advised that return and risk be considered together, as a return-risk compromise or trade-off, thereby offering a different approach to fund allocation (Kolm et al. 2014). The Markowitz model, or (M-V) model, was proposed in response to his belief that investors always want maximised returns with minimised risk. The outcome of his model is the concept of the efficient frontier (Banihashemi & Navidi 2017). The standard Markowitz problem (Section 2.2) is a quadratic optimisation problem, which can naturally be solved by quadratic programming to find the efficient frontier. Then, risk is minimised on the basis of an expected level of return, or return is maximised on the basis of an accepted level of risk.

In the real world, however, various constraints (Section 2.6) are commonly used in order to make the problem reflect reality more closely (in that asset management policies or regulations may be satisfied (Skolpadungket et al. 2007)). Such constraints or restrictions are, for example, bounds upon the proportions of the assets or the number of assets used. Since Markowitz's (M-V) model requires an efficient frontier to be traced – an illustration of the return-risk trade-off in the form of a continuous curve – and since quadratic programming finds this frontier easily, his approach is widely recognised and commonly applied (Chang et al. 2000). Indeed, the significance of Markowitz's (M-V) model is without question. Despite this, there have been a large number of disputes to some of its basic assumptions (Jin et al. 2016).

For example, the model disregards various restrictions found in the real world, such as the cost of transactions and tax, or investment decisions based on strategy or personal circumstances. It makes the assumption that any fraction may be used in the trade of an asset. It suggests investors choose the maximum possible amount of assets, in order to use diversification to minimise risk. But an investor often deals with various restrictions, such as limit on capital to be assigned to a particular industry or asset, and in reality, investors tend to prefer a controlled number of assets (Jin et al. 2016).

Rapid development has meant that commercials solvers (CPLEX, for example) can now handle certain constraints for various optimisation problems without prohibitive computational cost. How complex a portfolio optimisation problem is often rests on the constraints (Maringer 2008). Standard quadratic programming is used in the basic (M-V) model, and many tools have been used to find optimal solutions to the problem. But financial constraints in the real world cause the complexity level to rise significantly. Cardinality constraint, for example, renders the problem non-convex due to the limitations on asset numbers. Thus it is not always appropriate to find optimal solutions via exact methods, and as a result many researchers have concentrated on heuristics for constrained portfolio optimisation problems (Jin et al. 2016).

The problem's complexity increases upon addition of constraints to the basic (M-V) model, which causes difficulty for exact optimisation approaches to deliver solutions to large problems in a practical amount of time. Therefore, meta-heuristics and hybrid meta-heuristics have been investigated. While these prove efficient in finding optimal solutions within practical amounts of time, they are not able to guarantee solution optimality (Lwin et al. 2014).

1.1 Methods for Computing Solutions to Hard Problems

Introducing constraints often renders the resulting portfolio problems NP-complete or NP-hard (Mansini & Speranza 1999). Essentially, this means the problems cannot be solved in polynomial time (Garey & Johnson 1979). Thus, computing solutions to such problems often leads researchers to turn to approximate algorithms such as meta-heuristics (or hybrids of such). Meta-heuristics are useful in terms of achieving optimal or near-optimal solutions in a reasonable amount of time (Gulpinar et al. 2010) (although the precise meaning of "reasonable" is ill-defined, being dependent upon problem, industry, purpose of solver, and so on; commonly it is used to mean polynomial time).

Cesarone et al. (2013) introduced a state of the art method to solve a Limited Asset Markowitz (LAM) model. Their notion of a LAM model is equivalent to our formulation of Equations 2.1–2.4 in Section 2.2 if we include a cardinality constraint and bounds upon proportions. They investigated a method of obtaining optimal solutions of what were, at that time in 2009, the OR-Library (Beasley 2000) benchmark problems. They also ran their approach on five other self-created datasets (EuroStoxx50 - 47 assets, FTSE 100 - 76 assets rather than the 81 of OR-Library, MIBTEL - 221 assets, S&P500 - 476 assets and NASDAQ - 2191 assets). The authors compared their results to the

work of Di Gaspero et al. (2007) and others.

Some researchers (for example, the works of Brito & Vicente (2013) and Fieldsend et al. (2004)) attempt to solve constrained problems by incorporating constraints into the problem definition or into the objective itself, the latter leading to creation of more objectives as a trade-off with the number of constraints. In this process, different characterisations of hard problems may be considered as rewritings of the problems with differing numbers of objectives. These objectives often conflict. Equally, sometimes, the number of objectives may be reduced via rewriting. The latter may be achieved through scalarisation (for example, with the λ formulation of the Markowitz problem). Below we review works which fall into both categories of approach.

Some examples of work which treat the problem as a single-objective problem (also known as the weighted problem from the weight λ) are as follows. The work of Chang et al. (2000) considered the problem with the cardinality and minimum proportion constraints imposed, and constructed three heuristic algorithms (genetic algorithm (GA), simulated annealing (SA), and tabu search (TS)) to attempt to find optimal solutions. First, the authors ran all three algorithms on the unconstrained portfolio problem and found that the GA exhibited the best performance in terms of mean percentage error (compared with the solution produced by quadratic programming - QP), followed by SA and, finally, TS the worst performance. The mean time taken for all three algorithms rose sharply with the number of assets in each dataset (D1)–(D5) (these datasets are defined in Section 2.5). Second, the same algorithms were then ran on the cardinality constrained portfolio problem with minimum proportions, revealing the same ordering of algorithm quality in terms of the average percentage error (1.49% for the GA, 1.85% for SA and 2.05% for TS). The authors also provided results for a pooling of the solutions obtained by all three methods. This work is examined further in Chapter 3. It may be worthwhile to mention that this work was the first work using the benchmark datasets (D1)–(D5) introduced by the third author, Beasley, in Beasley (2000).

There are also many works which treat the problem with two objectives. The work of Mansini & Speranza (1999) illustrated that the problem with minimum lot constraint is NP-complete, and introduced mixed integer linear model heuristic algorithms to find a

solution close to the optimal one. They applied three different such heuristics on real data from the Milan Stock Exchange during the two periods 1989-1991 and 1992-1994. Their experiments claimed that these heuristics achieve very good solutions in a reasonable amount of time, claiming 0.12% average error.

Lin & Liu (2008) proposed three possible models for portfolio selection problems with minimum lot constraint and introduced a GA to solve the problem. On the other hand, Di Gaspero et al. (2011) suggested a hybrid technique to incorporate a local search metaheuristic with a quadratic programming procedure. They attempted to solve the problem with cardinality, quantity and pre-assignment constraints. A metaheuristic local search method searches sections of the search space and identifies cases. These cases are then solved with a QP algorithm. This approach purportedly compares well with other solvers (e.g., CPLEX) in terms of accuracy compared with the unconstrained efficient frontier, outperforming them for larger datasets.

However Armañanzas & Lozano (2005) discussed the three well-known optimisation techniques of greedy search, simulated annealing (SA), and ant colony optimisation (ACO) for solving the portfolio problem with cardinality and min/max proportion constraints explicitly as a multi-objective problem. They found that the cardinality constrained efficient frontiers (CCEFs) produced by the SA and ACO approaches are better than those of greedy search. Intuitively, this is unsurprising due to the construction of all three methods. More interestingly, perhaps, is that some points on the CCEF were only reachable by SA and some by greedy search in the particular variants used by the authors. ACO seemed to capture the points at the top right of the CCEF effectively, whereas SA seems to be more effective at capturing points at the bottom left of the CCEF.

Ruiz-Torrubiano & Suárez (2010) offered a variety of methods to solve the problem of portfolio selection with cardinality constraints, including SA and various estimation of distribution algorithm (EDA) family members. They compared these methods with a random assortment recombination genetic algorithm (RAR-GA) using set representation and bespoke crossover operators and mutation designed to maintain the solution candidates' cardinality. They also introduced pre-processing techniques which give

considerable improvements in the proposed algorithms' efficiency and quality, and remove products with low probability of inclusion in the final solution. Using SA for the solution of the combinatorial portion of optimisation problems results in the identification of portfolios of equal quality as those found using RAR-GA, with computational costs similar or lower. However, EDA algorithms used alone give poor results, and do not handle high-dimensional problems effectively. When probability distributions evolve from an increased number of products in the optimisation universe, estimation and sampling becomes harder. It may be possible to identify and remove from the universe any assets which are unlikely to feature in the optimal portfolio, using the solution to a relaxed optimisation problem. Metaheuristic methods improve search efficiency considerably, particularly in the case of EDA (above all population-based incremental learning (PBIL)) which become similarly viable to RAR-GA and SA. Indeed, this pruning can result in exact solutions being found with less computational cost than the other methods (EDA, RAR-GA and SA), for problems such as Nikkei and Hang Seng.

Anagnostopoulos & Mamanis (2011*a*) compared five well-known advanced multi-objective evolutionary algorithms (MOEAs) on the mean variance model with cardinality constraint. The algorithms compared were the Niched Pareto Genetic Algorithm 2 (NPGA2), Non-dominated Sorting Genetic Algorithm II (NSGA-II), Pareto Envelope-based Selection Algorithm (PESA), Strength Pareto Evolutionary Algorithm 2 (SPEA2), and e-Multiobjective Evolutionary Algorithm (e-MOEA). All MOEAs were tested on the now familiar datasets (D1)–(D5) for the values K = 5 and K = 10 with lower and upper proportion bounds 0.01 and 1 respectively. Their results showed that SPEA2 was the strongest algorithm on this problem, and all MOEAs outperformed an single-objective algorithm similar to that of Chang et al. (2000). Skolpadungket et al. (2007) applied various multi objective genetic algorithms for finding a solution of portfolio optimisation with cardinality constraints, floor and round lot constraints.

The work of Lwin et al. (2014) proposed an MOEA (MODEwAwL) to solve extended Markowitz's mean-variance portfolio optimisation model with cardinality, quantity, preassignment and round lot constraints. The MOEA proposed is influenced by a learningguided heuristic. Their algorithm was shown to compare well, in terms of solution quality and time, with the algorithms NSGA-II and SPEA2 above as well as two more algorithms (PEAS-II and PAES). They also achieved an efficient solution of up to 1318 assets (dataset (D7)).

Jin et al. (2016) employed exact solver to establish optimal solutions or lower bound, in constrained portfolio optimisation models. Their aim was to establish these solutions or bounds for extended constraint OR-library benchmark datasets, taken from five recognised indices. They began with an investigation of the consequences of round-lot, pre-assignment and class constraints using the Markowitz model for cardinality and quantity constraint, in order to examine increased problem difficulty. They then analysed several constraint settings. Their aim was to increase awareness of M-V models' current state using CPLEX and subsets containing practical constraints. They outline the fundamental M-V model, along with the extended constraints, then review the models identified in the literature by initially discussing the consequences of round-lot, pre-assignment and class constraints, using the M-V model for cardinality and quantity constraint with regards to computational cost and solution quality. They then demonstrate an experimental CPLEX 12.6 framework in order to determine solutions for the most frequently used constrained M-V models identified in the literature. Their findings demonstrate that round-lot, class constraints and pre-assignment do not impact significantly on the cost of solution for cardinality and quantity constrained problems. They confirm that cardinality constraint contributes to the difficulty of a problem, and that imposing settings as specific as possible on them makes it easier to obtain optimal solutions, so that a model which is completely constrained counterbalances the impact of cardinality constraint.

Oh et al. (2006) proposed a portfolio algorithm that constructs portfolios using the important portfolio beta β_p (a measure of systemic srisk) in order to assess its strengths. Selecting stocks using their market capitalisation, and optimising them with regards to β_p standard deviation, their β -G portfolio algorithm (a new portfolio algorithm) is claimed to produce the following findings: firstly, that it performs particularly well in stable markets, regardless of bullish or bearish conditions; and secondly, that its applications to the short-term in particular are often outstanding, especially with strategies of small numbers of stocks making up the portfolio and, particularly, on short training periods, which further relates the results to market stability.

Li et al. (2006) proposed an optimal lot solution for formulating discrete-feature constrained M-V portfolio selections with concave transaction costs, by means of an exact algorithm. Their method employs a contour-domain and convergent Lagrangian approach to solve the aforementioned kinds of problems using exploitation of some of the M-V formulation's featuresalong with the considered portfolio model. Their process of cutting the contour-domain reduced the Lagrangian dual search's duality gap, and assured convergence. Using Hong Kong stock market data, the proposed algorithm was shown to be efficient.

Shaw et al. (2008) use portfolio security limitation to examine the portfolio selection problem, which they formulate with cardinality-constrained quadratic programming and their own Lagrangian relaxation approach. Unlike many Lagrangian approaches which take advantage of the constraints' special structure, theirs takes its advantage from the objective function's structure. Tracking major markets such as the S&P 100, S&P 500, FTSE 250, and FTSE 100 with their algorithm, they compute promising results which suggest that optimal solutions can be produced for 250 asset problems and below. For problems between 250 and 500 assets, improved heuristic solutions can usually be produced.

Di Tollo & Roli (2008) present a summary of literature regarding the application of metaheuristics for portfolio selection problems. They identified several key guidelines for designing portfolio selection problems metaheuristic solvers. Firstly, that there is no such thing as a best method. Secondly, that developing an effective metaheuristic solver depends on skilfully combining various factors - such as which algorithm is chosen as one of the model's functions, as well as a phase of development during which there is systematic evaluation of the different design choices. They note that very often a sole heuristic is insufficient to meet practical requirements, and that therefore hybrids are frequently required. Their discussion of the consulted studies show that while metaheuristic and hybrid approaches have potential to solve the portfolio selection problems, there are algorithm design and modelling concerns which require attention if effective tools are to be implemented.

Guijarro (2018) suggested a new measure to identify the cardinality constrained frontier

in mean-variance portfolio optimization problems, by introducing a perceptual methodology that focuses on the optimal frontier rather than the optimal portfolio. By perceptually focusing on analysis of the cardinality constrained frontier, they searched for the most similarity to the mean-variance unconstrained frontier, producing an objective measurement which translates the perceptual concept and employs a genetic algorithm to deal with the combinatorial complexity that cardinality constraints create. Obtaining results from five real-world stock exchange indices, Guijarro (2018) showed that the similarity ratio successfully diversifies low cardinality values investments, and significantly reduces the number of securities to be considered, rendering included assets invariant to specific return levels.

Liagkouras (2019) introduced an algorithm intended to address the problem of limitations in the capacity of evolutionary techniques caused by search space. He included a custom coding structure which maintains invariant processing time however large the dataset to be tested, dramatically reducing the time required for these large-scale combinatorials to be considered. This new method performs better than either of the 2 MOEAs it was tested against, for the solution of portfolio optimisation problems with cardinality constraints and various measures of risk. Furthermore, processing time is reduced significantly.

Lim et al. (2020) attempted the design of an optimal portfolio, by building on existing products using a search which identifies stocks that have further growth potential. They conducted an analysis of previous risk-adjusted returns, using as their momentum the inertia in these, and scored each stock's relative values to inform their estimations. In addition, they employed a genetic algorithm to optimise portfolios. Testing their GA model against S&P500 and KOSPI200, their asset valuation and momentum strategy proved capable of searching out a high-return, low-risk, optimal portfolio which does not require expert analysis or domain knowledge in order to grow.

Akbay et al. (2020) considered the problem of cardinality constrained optimisation using a combination of parallel variable search algorithm (which determines the assets to be held) and quadratic programming (which calculates each asset's proportion). The variable neighbourhood search algorithm's preliminary solution and search pool was determined using λ -based risk coverage rate and the quadratic programming then determines the capital ratios of the assets selected. This method was tested against 5 well-known sets of data, and the authors found that the solution compares very well with other approaches, and is particularly efficient in the case of low-risk portfolios.

Kalayci et al. (2020) acknowledged that the efficiency of approaches using exact solution and cardinality constraints to optimise portfolios is limited. They therefore focused on metaheuristics and other approximation-based approaches, although optimality is not certain with these methods, suggesting that solutions nearing optimal may be reached far more quickly. Their algorithm is a metaheuristic hybrid, consisting of borrowed elements (elitism mechanisms, modification rates and gaussian formulation) from genetic algorithms and the optimisation of artificial bee colonies and continuous ant colonies, respectively. It proves to be as effective as the advanced algorithms they review, when tested against 7 benchmark problems.

Recently, Kizys et al. (2022) proposed a simheuristic algorithm, combining Monte Carlo simulation techniques with the variable neighborhood search metaheuristic, for solving the stochastic portfolio optimisation problem. According to the authors, the methodology has the ability to solve real-sized stochastic instances in reasonably short amounts of time. The method is also shown to yield significantly improved results even in conditions of low stochasticity, as compared with classical metaheuristic approaches which generate deterministic portfolio optimisation problem solutions.

1.2 The Motivation, Contributions and Aims of The Thesis

This section discusses the motivation and contributions of this thesis.

Constraints are a major factor in portfolio optimisation problem complexity (Maringer 2008). An investment manager (or investors) may be constrained by number of assets, and minimum/maximum capital to be assigned to any one industry or asset.

However, as far as we are aware, most work in the literature considers the problem of finding the efficient frontier portfolio optimisation problem with including cardinality constraints at only values 5 and 10, lower proportion bound at 0.01 and upper proportion bound 1. This motivates our research to more challenging portfolio optimisation prob-

lems to impose new different values of cardinality constraints like 2,3,4 and minimum proportion constraints such as 0.05, 0.1, 0.15, 0.18, 0.2, 0.25, 0.3.

On the other hand, solvers have been used for varying types of portfolio optimisation problems of assets with varying constraints (e.g., Lwin et al. (2014)). However, these approaches assume the underlying joint distribution of assets is normal and also ignore the fact that investors may care about characteristics of the asset distribution other than mean and standard deviation. In addition, asset distributions of real-world financial data are rarely normal and exhibit so-called "heavy tails" behaviour (e.g., financial crashes).

Copulas model the marginal of each asset separately (which may be any distribution) and also the interdependencies between assets (see Patton (2004), Kakouris & Rustem (2014), Krzemienowski & Szymczyk (2016) for an interesting approach using Monte-Carlo methods, and Concepcion Ausin & Lopes (2010) for an approach using GARCH models). Many datasets, especially high-frequency financial data, are large and require significant processing capabilities and algorithms. We wish to work on this rapidly-evolving area of research, bringing in applications in business such as portfolio optimisation in Gulf Cooperation Council (GCC) countries.

In spite of oil's importance in the GCC markets, only limited comparative portfolio optimisation problem studies were found in the literature. As well as over the last 20 years, the availability of commodity markets to investors has increased, offering new diversification opportunities for investment portfolios. These opportunities require sufficient understanding of the correlation between financial and commodity markets, if they are to prove lucrative (Mirović et al. 2017). As commodities correlate very little or even negatively with equities, the employment of commodities can be an appropriate way to construct portfolios (Gorton & Rouwenhorst 2006). This motivates the use of approximate algorithms such as meta-heuristics that may provide a sufficiently good solution in order to construct portfolios in GCC markets.

The main contribution of this thesis is to address and solve the above issues. This was addressed in the following way.

- A genetic algorithm was formulated in order to test the portfolio optimisation problem with cardinality constraints and ensure that the same magnitude of results is produced as in the literature. This was then used on the minimum proportion cardinality constraint problem;
- A parallel MPI code was created to address the large number of generations required to find solutions to these types of problems, the large quantity of assets needed, and the considerable computation time taken. This code was implemented on the High Performance Computing (HPC) cluster at the University of Plymouth;
- 3. This thesis makes substantial contributions to using algorithms with combinations of class, cardinality and minimum proportion constraints in the context of portfolio optimisation and, also, in the integration analysis of Gulf country stock markets, oil and precious metals. An extensive review of the literature has confirmed that no study has previously investigated optimal portfolio construction in all GCC markets, with the inclusion of oil and precious metals into these portfolios. A greater sample size of GCC stock markets was used in this study;
- Copula-based Worst Case Conditional Value at Risk (WCVaR) was applied to the GCC dataset, so as to make an accurate assessment of the risk to investment ratio.

In this thesis the main constraints discussed are the cardinality constraint limiting the number of assets, minimum proportions and class constraints. To begin with, a parallel genetic algorithm in R to replicate existing results on unconstrained benchmark datasets is developed and an investigation is made of it using the HPC cluster. Copulas and packages that enable computation with novel forms of risk measure (standard portfolio problems assume a normal asset return distribution, which does not conform to the real world, leaving a gap in risk realisation) are then investigated. The study also looks at business applications, potentially providing real-world impact. This is done by constructing optimal GCC portfolios that include oil, at least one precious metal (gold, silver or copper) and stocks of companies in the Gulf market. This PhD project makes use of interdisciplinary topics from financial mathematics, metaheuristics, high performance computing, probabilistic, business and economic applications. Given the interdisciplinary nature of this study, the literature review provided consults the most relevant literature rather than providing a comprehensive review of each field.

In the next section, the objectives of the thesis are given.

1.3 Objectives of the Thesis

The main objective of this thesis is to provide an extension to the basic portfolio optimisation model. This is done by taking new cardinality constraint values and new values of minimum proportion into consideration. Optimal portfolios on the GCC dataset - a real world dataset of large problem size - are created using class and other relevant constraints and the results are analysed. The WCVaR risk measure, an industry alternative measure, is investigated and shown to accurately assess the risk to investors.

1.4 Overview of Chapters

This thesis comprises six chapters, including this introduction. Following this, the structure of the thesis is as follows.

In Chapter 2, a mathematical definition of the basic Markowitz Mean-Variance problem and the application of the quadratic programming method in solving this problem are introduced. All necessary tools are presented in the following sections pointing out a number of practical constraints commonly faced by investors and datasets utilised for computational analysis in this thesis. At the end, some measures of risks are reviewed.

Chapter 3 discusses the techniques and methodology which were employed to achieve our research objectives. This chapter provides a framework that may be used to solve portfolio optimisation problems, on five different datasets, employing the GA with new values of minimum proportion constraints and various new cardinality values. Empirical results and analysis are both provided.

In Chapter 4, the extended mean-variance model with class constraint is used to construct optimal portfolios on the GCC dataset (produced via daily returns data from

1.4. OVERVIEW OF CHAPTERS

January 2008 to January 2019 from the GCC countries), and to investigate the usefulness of including oil and precious metals in hedging equity portfolios. Following this, a discussion is given on fluctuations between year-to-year frontiers.

In Chapter 5, copulas based on WCVaR are applied to the GCC dataset in order to assess the risk to investment accurately.

In Chapter 6, the main conclusions and major findings of the study are summarised and suggestions for further research can be found at the end of this chapter.
Chapter 2

Optimisation of Portfolio Risk

2.1 Introduction

Portfolio optimisation and the diversification of portfolios have long been influential for financial market understanding and development, as well as for making financial decisions (Kolm et al. 2014). Despite different investor attitudes, the primary concerns are generally maximising return and minimising risk simultaneously (Banihashemi & Navidi 2017). This is often achieved by recognising the financial asset combination which best fits the needs of the investor (Maringer 2008). It was Markowitz (1952) who first gave the portfolio theory as a strategy for investments, taking return and risk into consideration. His model – called the Markowitz model, or mean-variance (M-V) model – attempts to reduce the risk (variance) for a given level of return, resulting in a frontier area known as efficient frontier (Banihashemi & Navidi 2017). Portfolio management is also understood as the identification of combinations which an investor can use to make an ideal trade-off between projected profit and its related risks.

This chapter presents some basic concepts that will be used in later chapters. Markowitz mean-variance problem, quadratic programming, benchmark datasets of the portfolio optimisation problem, application of QP to the Markowitz mean-variance problem, a number of practical constraints commonly faced in the real world, and finally common measures of risk.

2.2 Markowitz Mean-Variance Problem

Markowitz proposed a portfolio (i.e., a vector $\mathbf{x} = (x_1, ..., x_n)$) with *n* assets, where x_i is the weight of asset *i* and μ_i is the expected return of asset number *i*. The expected

return of the portfolio may be written as the scalar product

$$\overline{R} = (\mathbf{x}, \mu) = \sum_{i=1}^{n} \mu_i x_i,$$
(2.1)

with its variance (commonly referred to as 'risk') as

$$r = \sigma^2 = Var[\overline{R}] = \sum_{i=1}^n \sum_{j=1}^n \sigma_{ij} x_i x_j.$$
(2.2)

Note that σ_{ij} is the covariance between the returns for assets *i* and *j*. It is expected that the weights are normalised, giving that

$$\sum_{i=1}^{n} x_{i} = 1$$

 $0 \le x_{i} \le 1, \qquad i = 1, 2, ..., n.$ (2.3)

The Markowitz M-V problem is then to solve the following optimisation problem:

Find $\min r$ subject to the constraints

$$\overline{R} \ge R_0$$

$$\sum_{i=1}^n x_i = 1$$

$$x_i \ge 0 \text{ for } i = 1, \dots, n$$
(2.4)

That is, we must minimise risk subject to a certain desired level of return, R_0 , being attained. Equivalently, we may seek to maximise expected return subject to a given level of risk (Lwin et al. 2014). The M-V model of Markowitz focuses on obtaining the efficient frontier (EF), which is the set of optimal portfolios that respect the above objectives. This is also known as the unconstrained portfolio optimisation problem (Hochreiter 2010), and the standard method of solving this problem is by quadratic programming. The constraint that $x_i \ge 0$ for i = 1, ..., n prohibits short-selling (i.e., asset *i* would be short sold if $x_i < 0$). As above, this is a standard constraint in this type of

problem.

2.3 Quadratic Programming

Quadratic programming (QP) is the process of solving an optimisation problem, where the objective is to maximise or minimise a quadratic objective function subject to linear inequality and/or equality constraints. The definitions and theorems in this subsection are from the work (Werner & Sotskov 2006).

Definition 1. If $A = (a_{ij})$ is a matrix of order $n \times n$ and $X^T = (x_1, x_2, ..., x_n)$, then the term

$$Q(x) = X^T A X \tag{2.5}$$

is called a quadratic form. We have

$$Q(X) = Q(x_1, x_2, \dots, x_n) = (x_1, x_2, \dots, x_n) \begin{pmatrix} a_{11} & a_{12} & \cdots & a_{1n} \\ a_{21} & a_{22} & \cdots & a_{2n} \\ a_{n1} & a_{n2} & \cdots & a_{nn} \end{pmatrix} \begin{pmatrix} x_1 \\ x_2 \\ \vdots \\ x_n \end{pmatrix}$$
$$= \sum_{i=1}^n \sum_{j=1}^n a_{ij} x_i x_j.$$

Lemma 1. Let *A* be a matrix of order $n \times n$. Then the quadratic form X^TAX can be written as a quadratic form X^TA^*X , for a symmetric matrix A^* where

$$A^* = \frac{1}{2} \left(A + A^T \right).$$
 (2.6)

Definition 2. A square matrix *A* of order $n \times n$ and its associated quadratic form Q(x) are said to be

- 1. positive definite if $Q(x) = X^T A X > 0$ for all $X^T = (x_1, x_2, ..., x_n) \neq (0, 0, ..., 0)$;
- *2.* positive semi-definite if $Q(x) = X^T A X \ge 0$ for all $x \in \mathbb{R}^n$;
- 3. negative definite if $Q(x) = X^T A X < 0$ for all $X^T = (x_1, x_2, ..., x_n) \neq (0, 0, ..., 0)$;
- 4. negative semi-definite if $Q(x) = X^T A X \leq 0$ for all $x \in \mathbb{R}^n$;

5. indefinite if they are neither positive semi-definite nor negative semi-definite.

A simple case of quadratic programming is when the matrix *Q* is positive definite.

2.4 Application of QP to the Markowitz Mean-Variance Problem

In this section, an experimental investigation was conducted to generate the unconstrained efficient frontier (UEF). The sample code uses the solve.QP R command to solve a quadratic programming problem for each successive point on the EF. The standard representation of an optimisation in solve.QP is to minimise $\frac{1}{2}D^TX - d^TX$ subject to the constraint $AX \ge b$ where *D* is a matrix and *d* is a real-valued, *n*-dimensional vector and *X* is the vector of optimisation variables x_1, \ldots, x_n (also known as the weights). Therefore some conversion of our problem is necessary to put it into the above format. To convert the Markowitz problem into one suitable for the QP solver in R, it has to be converted to a matrix problem. The constraints are given below.

$$\begin{pmatrix} 1 & 1 & \cdots & 1 \\ \mu_{1} & \mu_{2} & \cdots & \mu_{n} \\ 1 & 0 & \cdots & 0 \\ 0 & 1 & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ 0 & 0 & \cdots & 1 \end{pmatrix} \begin{pmatrix} x_{1} \\ x_{2} \\ x_{3} \\ x_{4} \\ \vdots \\ x_{n} \end{pmatrix}^{"} \geq " \begin{pmatrix} 1 \\ R_{0} \\ 0 \\ 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix} \Longrightarrow \begin{array}{l} \sum_{i=1}^{n} x_{i} \mu_{i} \geq R_{0} \\ \Longrightarrow & x_{1} \geq 0 \\ 0 \\ \vdots \\ 0 \end{pmatrix} \Longrightarrow \begin{array}{l} x_{2} \geq R_{0} \\ x_{1} \geq R_{0} \\ x_{2} \geq R_{0} \\ x_{2} \geq R_{0} \\ \vdots \\ 0 \\ \vdots \\ 0 \end{pmatrix} \Longrightarrow \begin{array}{l} x_{2} \geq R_{0} \\ x_{2} \geq R_{0} \\ x_{3} \\ x_{4} \\ \vdots \\ 0 \\ \end{array}$$

The problem then becomes: find $\min \mathbf{x}^T \sigma \mathbf{x}$ (variance of expected returns) such that $A\mathbf{x}^{"} \geq \mathbf{b}$ (constraints). The addition of quotes to the ' \geq ' symbol indicates that the first constraint is an equality and the other constraints are inequalities.

Figure 2.1 shows the output of the code using 2000 EF points. The precision of computation in the code is set to that of the results of Beasley (2000); i.e., ten digit precision. The figure shows an efficient frontier computed for the dataset (D5) of OR-Library (see the next subsection) and compared with the solution provided by the author. The curve gives the best possible trade-off between return and risk, and represents the set of optimal portfolios. It can be seen that the curve of the computed solution matches that





Figure 2.1: Computed EF and the EF of Beasley to dataset (D5) of Beasley (2000).

2.5 Benchmark Datasets of the Portfolio Optimisation Problem

Availability of benchmark datasets is important to evaluate computational results for a class of problems. Based on these results, researchers can compare the efficiency of their algorithms and solutions quality (Cesarone et al. 2013). The most common datasets based on real-life data for portfolio optimisation problems were first introduced by Chang et al. (2000). The authors examined five test datasets based on five different capital market indices. They considered the Hang Seng (Hong Kong), DAX 100 (Germany), FTSE 100 (UK), S&P 100 (USA) and Nikkei 225 (Japan). They obtained 291 values for each stock to compute weekly returns and covariances from March 1992 to September 1997. The arithmetic mean of returns is used, which is identical to what is used in the standard Markowitz model. It can be seen that the size of the above datasets (D1) to (D5) are 31, 85, 89, 98 and 225 assets respectively. Collectively, these datasets are known as OR-Library (Beasley (2000)) and form standard benchmark datasets - these are widely used by authors to test portfolio optimisers (e.g., Lwin

et al. (2014)). The covariance matrices and return vectors of these datasets can be found in files 'port1', 'port2', ..., 'port5' available from the above reference. It should also be said that all datasets (D1–D5) have been used solely for the cardinality constrained portfolio optimisation problems considered in Chapter 3. In Chapters 4 and 5 other datasets are used.

2.6 Practical Constraints

Adding practical constraints to the M-V problem often makes the problem much more difficult (Lwin et al. 2014). In these situations, algorithms which utilise stochastic means are commonly preferred (e.g., genetic algorithms, tabu search or simulated annealing) (Chang et al. 2000). In the following subsections, common constraints which have been used in practical applications will be discussed.

2.6.1 Cardinality Constraint

In practice, the cardinality constraint limits the number of assets, *K*, from the total number of assets (*n*) that form a portfolio. The objective of this constraint is for investors to limit their choices of assets via this restriction. Let z_i be a binary variable associated to whether asset *i* is held in the portfolio. For example, if we have n = 7 assets in total and we restrict to K = 3 then a suitable *z*-vector might be $\mathbf{z} = (0,0,1,0,0,1,1)$, meaning that only assets 3, 6 and 7 are held. Thus the constraint is commonly, in the works of Chang et al. (2000) and Jin et al. (2016) for example, written as:

$$\sum_{i=1}^{n} z_i = K,$$
$$z_i \in \{0,1\} \text{ for } i = 1, \dots, n.$$

In some work (e.g., Maringer & Kellerer (2003)), the cardinality constraint is considered to be

$$\sum_{i=1}^n z_i \le K$$

However, some authors (e.g., (Chang et al. 2000, p. 1276)) believe that dealing with the cardinality constraint as an inequality makes computation in an algorithm which purports to solve such a problem more practicable.

2.6.2 Class Constraint

Investors may separate assets into disjoint sets, called classes, C_1, \ldots, C_M , where $\bigcup_{i=1}^M = \{1, \ldots, N\}$ and $C_i \cap C_j = \emptyset$ for all $i \neq j$. These classes share common characteristics (for instance, oil stocks, utility stocks, and so on). These asset classes are often chosen because they are negatively correlated (Jin et al. 2016). The goal here is to make the portfolio more diversified and/or "safe". Let $L_m > 0$ be the lower weight limit and $U_m \leq 1$ the upper weight limit of class *m* for $m = 1, \ldots, M$ and x_i the weight of asset *i*. Thus the class constraints are commonly written as (e.g., in Jin et al. (2016))

$$L_m \leq \sum_{i \in C_m} x_i \leq U_m \text{ for } m = 1, \dots, M.$$

According to Jin et al. (2016), this constraint is linear and thus does not render the Markowitz problem NP-hard by itself. Jin et al. (2016) used different settings for the lower bound and partitioned the assets of OR-Library dataset (D4) into distinct classes in order to test the class constraint. Since class constraint is linear, the expectation was that it ought not to compound problem difficulty. It was also expected that difficulty would be reduced by the amount of classes, since the initial problem was divided into subproblems with reduced dimension. The authors reported that using fewer classes resulted in an increased amount of computational effort being required. Also, as predicted, not a lot of difficulty with regards to computational cost is generated by imposing a class constraint on a cardinality constrained model, although the number of EF points correctly approximated may be affected. The authors also reported that class constraints do not result in significant differences in the cost of solving this modified cardinality constrained problem.

2.6.3 Rounded or Minimal Lots

The rounded lot constraint determines the number of any asset in the portfolio to be an exact multiple of minimum lots, where a minimum lot is the minimum quantity of an asset that may be traded (Skolpadungket et al. 2007). Let y_i be a positive integer and l_i the minimum lot for each asset. Then the round lot constraint is formulated as in Jin et al. (2016):

$$x_i = y_i \cdot l_i$$
 for $i = 1, ..., n$.

2.6.4 Bounds Upon Proportions

The bounds upon proportions constraint defines the minimum proportion ε and maximum proportion δ for each asset in the portfolio. It can be stated as follows (Chang et al. 2000):

$$\varepsilon z_i \leq x_i \leq \delta z_i$$
 for $i = 1, \ldots, n$,

Such constraints are also known as quantity constraints (Jin et al. 2016). Recall the definition of z_i in Section 2.6.1.

2.6.5 Pre-assignment Constraint

This constraint is used to represent subjective preferences of an investor. Suppose that an investor intends to determine the proportion of specific assets in their portfolio and that the set of all such desired assets is denoted by (the pre-assignment set) *P*. Let a binary variable s_i denote whether asset *i* is included in *P* (i.e., $s_i = 1$ for all $i \in P$ and 0 otherwise). Clearly, we also require that (Jin et al. 2016):

$$s_i \leq z_i$$
 for $i = 1, \ldots, n$.

2.6.6 Literature on the Effects of Imposing Constraints

When imposing any of these constraints, it is important to ask what happens when each of them is applied to the standard Markowitz problem. It is almost certain that the problem will become more difficult when such constraints are applied. As mentioned in Section 2.6.1, some authors have imposed the cardinality constraint as an inequality (i.e., the number of assets chosen is $\leq K$) as in practicality this makes the problem more manageable. For example, (Chang et al. 2000) and (Lwin et al. 2014) use the equality formulation, whereas (Anagnostopoulos & Mamanis 2011*a*) and (Maringer & Kellerer 2003) use the formulation of an inequality.

There is some work on the comparative effects of applying single constraints from the

above list to portfolio problems (Jin et al. 2016). In their work, the authors took the Markowitz problem with cardinality constraint K = 10 and quantity constraint $\varepsilon = 0.01$, and applied the pre-assignment, rounded lot and class constraints in sequence. They found that the choice of assets that were pre-assigned directly affected the solution quality (which the authors measured in terms of average percentage error with the unconstrained EF). However, on the other hand, the authors state that the time taken to compute solutions with large numbers of pre-assigned assets is discernibly lower when compared to smaller numbers of pre-assigned assets. This seems intuitively clear since there are fewer assets left for any solver to find. In addition, pre-assigning assets which would be in the unconstrained efficient frontier leads to better solution quality. In terms of the round lot constraint, the authors found the effect of applying this constraint (through testing with three distinct sizes of lot) was small in terms of the "optimal rate". This measure is simply the percentage of points on the constrained efficient frontier that appeared in the correct positions from the computed solution. Finally, in terms of the class constraint, increasing the lower weight limit L_m for a particular class C_m results in a decrease in optimal rate.

2.7 Measures of Risk

Risk measurement is important to measure and predict the risk of financial institutions. The recent global financial crisis has indicated the importance of risk measures as tools which can avoid or limit risk at an early stage. Risk measures may be classified into four phases (Mitra & Ji 2010):

- 1. Pre-Markowitz risk measures.
- 2. Modern Portfolio Theory.
- 3. Value at Risk and related risk measures.
- 4. Risk Measures based on coherent risk.

2.7.1 Pre-Markowitz Risk Measures

A risk measure ρ is a mapping from \mathcal{G} , a linear space of measurable functions, to the real line

$$\rho:\mathscr{G}\to\mathbb{R}.$$

For the definition of a measurable function and a measurable space, the reader may consult (Halmos 2013) for details.

In the pre-Markowitz era, risk measurement and diversification were considered. Back in the time of Daniel Bernoulli in 1738, it was argued that decisions about risk may be evaluated by the expected return (Mitra & Ji 2010). Some economists had used variance to measure risk (Fisher 1906, Tobin 1958). The work of Graham (2003) suggested minimisation of risk strategies such as taking into account the margin of safety and also portfolio diversification (Mitra & Ji 2010).

2.7.2 Modern Portfolio Theory (MPT)

Despite risk measures being known prior to Markowitz's risk measure, he was the first to formalise asset risk, return and diversification in a mathematical framework. MPT was introduced in Markowitz (1952), for which he was later won a Nobel prize in Economics. For more details, refer to Section 2.2.

2.7.3 Value at Risk

Value at risk (VaR) measures level of market risk in firms or an investment portfolio over a given period of time (Daníelsson 2011). Furthermore, the Basel Committee on Banking Supervision acknowledge VaR as the basis of measuring risk because it is one of the vital risk measurements in the banking sector (Bernoulli 1954). It is used by investors to determine the potential losses in portfolios (Mitra & Ji 2010). VaR is defined as follows:

Definition 3 (Value-at-Risk). (Daníelsson 2011, p.76) The loss on a trading portfolio with a given probability p of losses such that p equals or exceeds VaR in a given trading period and a (1-p) probability of losses that are lower than the VaR.

According to Mitra & Ji (2010) and Terpezan Tabără (n. d.), there are three primary

methods used to estimate VaR. One is by using the familiar variance-covariance matrix of the risk factors (assets) (Mitra & Ji 2010). The other two methods follow (Harper 2017, Farid 2010):

• Variance-Covariance: It is assumed that daily returns adhere to a normal distribution, though the presence of "fat tails" mean that this simplification does not always match reality. Variances and correlations are calculated for all risk factors, giving this method the well-known variance-covariance matrix. An estimate for VaR is then produced, using the covariance matrix and each asset's importance.

• Historical Simulation: This is one of the more instructive ways to calculate p% VaR. It assumes that the returns distribution in the future will be identical to that of past data; in other words, it assumes that the returns distribution does not change with time. For instance, if we have past data then we first calculate the continuous returns over all days. Second, order them by magnitude from the smallest to the largest; and third, find the appropriate position that has p% of the distribution of returns on the left and (100 - p)% on the right.

We shall calculate the 1% VaR of S&P500 (ticker name ^GSPC) from 01/01/2000 until 13/04/2017. The risk measured here is from the amount of money which could be lost by an investor. Initially, every day we calculate the difference in the price with the day before, and so the daily returns. It is then instructive to plot a histogram of these daily returns; in the R language this may be done by using the command hist.

VaR is the value that has 1% of the distribution on the left and this number means that we are 99% confident that losses in the next day will not exceed a given number. The R code in Figure 2.2 gives the calculation of this number, which is approximately 3.5%. Thus there is a probability of 1% that losses in the next day will exceed 3.5%. Figure 2.3 gives the plot produced by this R code.

• Monte-Carlo Simulation: This type of simulation typically relies upon many random experiments according to a given model. In the case of finance, a model

```
library(tseries)
library(zoo)
price = get.hist.quote(instrument = "^gspc",
start = "2000-01-01", quote="AdjClose")
#time series starts 03/01/2000
y=diff(log(price))
y=coredata(y)
sy=sort(y)
n=length(sy)
p=0.01
VaR=sy[n*p]
VaR
```

[1] -0.03501744

Figure 2.2: Sample R code to compute the VaR of ^GSPC during the period 01/01/2000-13/04/2017.



Figure 2.3: Distribution of returns for ^GSPC ticker with VaR marked with a dotted line.

of future asset returns is assumed and this is simulated through running randomised trials or experiments. The method aims to arrive at a distribution of losses, through which the VaR may be extracted for various probabilities (as with historical simulation). The main difference between this approach and historical simulation is that the hypothetical trials are generated by a statistical distribution (the distribution according to the application) rather than historical data. The typical approach of a simulation is as follows:

- Generate a set of simulated prices at random.
- Compute daily return series.
- Proceed as with a historical simulation as above.

It should be mentioned, however, that differences in opinion as to what a suitable risk metric is, given certain circumstances or assumptions, still exist. Below, some alternative measures of risk are reviewed.

2.7.4 Coherent Risk Measures

The works of Artzner et al. (1999) and Mitra & Ji (2010) classified a coherent risk measure as a risk measure $\rho : \mathscr{G} \to \mathbb{R}$ which must satisfy the following four properties:

- Risk is monotonic: if X ≥ Y then, almost surely, ρ(X) ≤ ρ(Y). It means if the value of portfolio X is at least the value of portfolio Y, then almost surely the risk of X is at most the risk of Y.
- Risk is homogeneous: for all X ∈ 𝒢, λ ∈ ℝ₊ such that λX ∈ 𝒢, we have ρ(λX) = λρ(X). That is, the risk of an enlarged portfolio depends directly on its enlargement factor λ.
- 3. Riskless translation invariance: for all $X \in \mathscr{G}$ and $x \in \mathbb{R}$ such that $X + x \in \mathscr{G}$. We have $\rho(X + x) = \rho(X) x$. Broadly, it ensures that the addition of a given amount of cash to the portfolio *X* decreases the risk by an identical amount.
- Risk is sub-additive: ρ(X + Y) ≤ ρ(X) + ρ(Y). This means that the risk of two combined portfolios *X* and *Y* is at most the risks of each separate portfolio. In other words, diversification in a portfolio helps decrease its risk.

A coherent risk measure also has the obvious property that $\rho(0) = 0$. It is accepted (Mitra & Ji 2010) that risk measures obeying the above four properties are desirable for reasons of good behaviour of investments (properties 1, 2 and 3) and portfolio

diversification (property 4). It is well known that value at risk (Section 2.7.3) is not sub-additive, and so cannot be a coherent risk measure. Neither is variance coherent. Some examples of coherent risk measures are as follows (Ahmed et al. 2007).

Semi-Absolute Deviation

The semi-absolute deviation is equal to one half the absolute deviation. It decreases the number of constraints with respect to the absolute deviation. This risk measure has been applied to portfolio problems subject to the minimum lot constraint and transaction costs which are concave (Mansini & Speranza 1999). The definition of the measure is as follows:

Definition 4 (Liu & Qin (2012)). Let ξ be an uncertain variable with finite expected value *e*. Let $(\xi - e)^- = \min(\xi - e, 0)$ and $(\xi - e)^+ = \max(\xi - e, 0)$. Then the semi-absolute deviation of ξ is defined as

$$Sa[\xi] = E[|(\xi - e)^{-}|].$$
(2.7)

According to Chiodi et al. (2003), this may be shown to be equivalent to the following model of risk.

Mean Absolute Deviation

The mean absolute deviation (MAD) is posited as an alternative to the mean-variance model. The reason why this approach had appeared is that the computational difficulty that was faced when an attempt is made to solve a large-scale quadratic programming problem (Konno & Yamazaki 1991). The work of Konno & Yamazaki (1991) defines the mean absolute deviation risk to eliminate difficulties associated with the mean variance model. This is stated as

$$E\left(\left|\sum_{j=1}^n R_j x_j - E\left(\sum_{j=1}^n R_j x_j\right)\right|\right),$$

where R_j is the rate of return and x_j is the weight for asset *j*. The authors claim that this risk measure allows the portfolio problem to be solved by a linear program instead of by quadratic programming and it purportedly has positive features compared to the standard Markowitz M-V model (i.e., no covariance matrix is required and the computation time is more practicable). The problem they tackle is otherwise unconstrained. On the other hand, Simaan (1997) argues that when covariance matrix is omitted this will actually cause greater estimation risk, and that this risk is greater than the benefits thus obtained.

Konno et al. (1993) considered the possible asymmetry of return and expanded the MAD approach to consist of skewness in the objective function. Subsequently, Konno & Suzuki (1995) extended a mean-variance objective function to include skewness.

Expected Shortfall

Expected shortfall (ES) is a coherent alternative to the non-coherent VaR. Expected shortfall is also known as conditional value at risk (CVaR). This is defined as follows.

Definition 5. (Daníelsson 2011, p.85) Expected loss conditional on VaR being violated (i.e., expected profit or loss, *Q*, when it is lower than negative VaR):

$$ES = -E[Q \mid Q \le -VaR(p)] \tag{2.8}$$

In Chapter 5, recent risk measure Worst Case Conditional Value at Risk (WCVaR) is also considered as an alternative and is described in detail.

2.8 Summary of the Chapter

In this chapter, Markowitz mean-variance problem employed in this work has been introduced. Additionally, The basic concepts have been reviewed, including a taxonomy of constraints used in portfolio construction (for instance, cardinality constraint, minimum proportion constraints and class constraint which will be applied in this thesis). In addition, a detailed descriptions of the benchmark datasets of the portfolio optimisation problem used in this study for computational analysis are provided. Finally, various risks measures are also presented.

In the next chapter, optimal solutions of portfolio problems considering Markowitz's Mean-Variance portfolio optimisation model with various cardinality values and new values of minimum proportion constraint are investigated.

Chapter 3

Portfolio Optimisation with Cardinality and Minimum Proportion Constraints

3.1 Introduction

This chapter investigates the optimal solution of portfolio problems with Evolutionary Algorithms (EA), which mimic natural evolution. The most important feature of an EA is finding an optimal solution which can be used to attain the Efficient Frontier (EF). The most famous forms of EA are Genetic Algorithms (GA) and evolution strategies. In this chapter, algorithms to solve hard portfolio problems such as hillclimbing algorithms and genetic algorithms, and how they work, are discussed.

This chapter is organised as follows: Firstly, unconstrained portfolio problems are discussed, beginning with some experiments with a GA on the unconstrained portfolio problem. Then, experiments with a GA on the standard and new values of cardinality constraint are performed and analysed. Finally, a presentation and test is then given of the results accomplished by applying a GA to solve the cardinality constrained portfolio optimisation problem with minimum asset proportions. Throughout, the meaning and practical applications of the results are detailed. In the following section, we give an overview of algorithms, genetic or otherwise, which purport to solve hard problems (and, in particular, portfolio problems).

3.2 Algorithms to Solve Hard Portfolio Problems

Evolutionary Algorithms (EAs) have become popular tools since the mid-1960s. They mimic the ideas of natural evolution. However, the most important advantage of EAs is that a population of candidate solutions can be evolved to attain a good approxima-

tion of the Efficient Frontier (Neumann & Witt 2013). EAs use simulated evolution to find solution of complicated problems. EAs have a variety of types, such as genetic algorithms and evolution strategies which are the most basic forms of EAs. In addition, they are most commonly used for optimisation (Whitley 2001). In this chapter a brief overview is presented of hill climbing and genetic algorithm approaches proposed in the literature, how they work, and their basic concepts.

3.2.1 Hillclimbing Algorithms

Hillclimbing is an iterative algorithm, which starts with candidate solution to a problem. After that, the algorithm attempts to obtain a better solution by changing the solution gradually. Then, the process repeats until no better candidate solutions can be found. According to Luke (2013) there are four phases to reach an optimal solution:

- 1. Initialisation procedure: provide one or more initial candidate solutions.
- 2. Assessment procedure: assess the quality of a candidate solution.
- 3. Make a copy of a candidate solution.
- Tweak a candidate solution, which randomly generates a slightly different candidate.

If the quality of the modified solution is better than that of the copied solution then it becomes the solution to be copied by step 3.

3.2.2 Genetic Algorithms

Nowadays, genetic algorithms find the efficient frontier more efficiently, leading them to be widely applied in industry and science (Samanipour & Jelovica 2020). A Genetic Algorithm (GA) is an EA that takes a population of strings (in the most general presentation, a string of symbols from a given problem representation) that encode possible solutions and performs a sequence of copying strings, crossing over strings with others, and mutating strings to create a child population of individuals that are more fit. This child population is then ranked in best-fitness-first order and then replaces the parent population. This process continues, gradually increasing the fitness until some

termination condition is invoked. Such termination conditions, among others, could be a certain number of populations being created, a given time for algorithmic execution being reached or a target fitness being reached. How the fitness is measured (by a fitness function applied to each individual in a given population), the operations performed in the sequence and the representation of possible solutions typically depend upon the problem itself, the application, convenience and speed. The work of Goldberg (1989) gives further details for the interested reader.

As an application, the work of Arriaga & Valenzuela-Rendón (2012) claimed that Steepest Ascent Hill Climbing gives effective solutions for the portfolio selection problem with cardinality constraint and rounded lots (see Section 2.6.3). Being a hillclimber, this means that the algorithm of the authors is simple and uses few system resources. They show that the performance is similar to a GA (of which details are not provided). However, in many problems this is usually not true as hillclimbers tend to suffer from local minima to a greater extent than GAs.

There are examples of researchers using multi-objective EAs (such as NSGA-II, SPEA2, PESA2), but it is always when two (or more) objectives are employed to treat the problem; however, in our work we will not treat the problem with two objectives. The NSGA-II (Non-dominated Sorting Genetic Algorithm), suggested by Deb et al. (2002), bases its binary tournament selection on crowding distance. Mutation and crossover are performed using polynomial mutation operators and simulated binary crossover. Suggested by Zitzler et al. (2001), the SPEA2 (Strength Pareto Evolutionary Algorithm) uses density estimation, fine-grained fitness assignment and archive truncation, along with polynomial mutation operators and simulated binary crossover. The PESA2 (Pareto envelope-baSed Evolutionary Algorithm), suggested by Corne et al. (2001), also uses polynomial mutation operators and simulated binary crossover, combined with fitness-assigning hyper-boxes. The above multi-objective EAs are generally among the most widely used EAs on a variety of problem types.

The remaining content in this section gives an overview of possible configurations of GA operators and parameters. Justification of which operators and parameters were used is given in Section 3.3.1.

Representations

The most important decision in order to implement a genetic algorithm is how to represent the solution. Solutions could be represented in different ways (see Figure 3.1). In this section, some of the most commonly representations will be presented as follows. Representations are matched to the problem being solved, e.g., binary representation and knapsack problem (Puchinger et al. 2010).

- Binary Representation: This representation is that of a bit ('binary digit') string; that is, entries of the string are the Boolean decision variables 0 or 1. Decision variables represent, for example, 'yes' or 'no'. One popular use is in the knapsack problem (Puchinger et al. 2010), for example, where a '1' indicates that an item has been chosen for inclusion in the knapsack and a '0' indicates no inclusion. In Figure 3.1 (a), for instance, the individual has the bit 0 at positions 1, 2 and 5, and the bit 1 at positions 3 and 4.
- Real-Valued Representation: A real-valued representation is used for problems where real-valued strings are required. For example, portfolio optimisation is one such problem. Each decision variable is continuous, rather than discrete (as above). We note that, in fact "real-valued" exists as an approximation because a machine can only hold a real value up to a given number of digits. For the sake of expediency, we will overlook this point.
- Permutation Representation: For some problems, the solution can be presented by an order of elements. In this case a permutation representation is used. An example of this type is the traveling salesman problem which consists of a salesman and a set of cities. The salesman has to travel to each city, starting from a certain one and coming back to the same city exactly once. The total length of the journey has to be minimised. In this case, a permutation corresponds to a route around the cities.

As above, a real-valued representation, with each value representing an asset proportion, was employed on the portfolio problem.



Figure 3.1: Common representations used by GAs.

Genetic Operators

Genetic operators are used in GAs to form new offspring from the existing population, and take the forms of selection, mutation and crossover. It is acknowledged that there may be other variants to those covered below. We, in addition, prescribe that (unless otherwise stated) the use of the word 'random' may be taken to mean 'uniform at random'. Some common forms of genetic operators are as follows.

Selection

In this process, parents are selected to create offspring for the next generation. There are many ways to select the best parents (for example, roulette wheel selection, tour-nament selection, and so on).

- Roulette Wheel Selection: This is used in GAs to select parents according to their fitness. Therefore, the better individuals have a higher chance of mating and propagating their features to the next generation. Assuming the fitnesses of each population member have been evaluated and the members are sorted in decending order according to fitness, this process is implemented by the following steps:
 - 1. Calculate the sum, $S = \sum f(\mathbf{x}_i)$, of the fitnesses, f, of all individuals (vectors) \mathbf{x}_i in the population.
 - 2. Generate a random number $a \in (0, S)$.
 - 3. Choose the individual \mathbf{x}_m from the population with *m* the minimal number such that $\sum_{i=1}^{m} f(\mathbf{x}_i) > a$.

The idea behind this procedure is similar to that of a roulette wheel spin. A roulette wheel is typically divided into parts, each of which we will use to represent an individual from the population. The area of each part is proportional to the

fitness of the individual it represents. When the wheel is spun, the probability of it landing upon a given individual (or area) is thus also proportional to the fitnesses of the individuals in the population. The random number *a* represents the spin (or, indeed, several such spins). The individual chosen is then given by the area landed upon by the wheel. This process is repeated to select as many individuals as needed.

- Tournament Selection: Choose a given number of individuals at random from the population; the individual with the best fitness from those chosen is then chosen to become a parent for crossover. Repeat this process as needed.
- Selection by Linear Ranking: Selection comes after ranking. Assuming the ranking of best individual at rank N and the worst at rank 1, the selection probability associated with each individual at position *i* is (by Shukla et al. (2015))

$$p_i = \frac{1}{N} \left(n^- + (n^+ - n^-) \cdot \frac{i - 1}{N - 1} \right), \tag{3.1}$$

where n^-/N and n^+/N are the probabilities of selecting the worst and best individuals respectively.

It is worth mentioning that many selection operators contain a note of elitism. Elitism is where the individual which is ranked at the top of the population is always selected at least once. This guarantees that the best fitness values over subsequent populations are increasing.

Mutation

Mutation is the genetic operator which takes a single individual in the population, and applies a small change. It is used to preserve genetic diversity, as generally small changes only affect a small number of positions in the individual. There are various mutation types. For example:

1. Bit flip: This mutation method flips one or more bits at random. It is used with binary representations as the bit '0' may be flipped to a '1' and vice-versa.

- Swap Mutation: In this method, two position in an individual are picked at random. The values in their positions are then exchanged.
- 3. Scramble Mutation: A subset of positions for a given individual are chosen. The values at those positions are randomly rearranged. A particular type of scrambling is that of inversion, where a contiguous subset of positions is chosen after which their ordering is reversed.
- 4. Mutation at a Random Position: Given a real-valued representation and individual vector **v** of length *n*, choose a random position in the vector. Then assign a value from the uniform distribution $U(u_{\min}, u_{\max})$ to that position in **v**.

Crossover

Crossover is a genetic operator that selects one or more parents to produce one or more offspring. In this section some of the most popularly used crossover operators are presented.

- Single-point crossover: A random crossover point is selected. The first part of the first parent (to the left of the crossover point) is swapped with the second part of the second parent (to the right of the crossover point) to get two new offspring.
- Multi-point crossover: This is a generalization of the one-point crossover in which alternating portions (slices) are exchanged to make new offspring.
- Uniform crossover: In this operator individuals are not divided into slices but each position in the individual is treated separately. In fact in this method, a coin is flipped for each individual to be selected or not in the offspring.
- Laplace crossover: Deep & Thakur (2007) proposed a crossover operator called Laplace Crossover (LC) which uses the Laplace distribution. Beginning with two parent solutions $x^{(1)} = (x_1^{(1)}, \ldots, x_n^{(1)})$ and $x^{(2)} = (x_1^{(2)}, \ldots, x_n^{(2)})$ as real-valued equal length vectors, i.e., $x_i^{(k)} \in \mathbb{R}$ for $k \in \{1, 2\}, i \in \{1, \ldots, n\}$, the operator produces two

offspring
$$y^{(1)} = (y_1^{(1)}, \dots, y_n^{(1)})$$
 and $y^{(2)} = (y_1^{(2)}, \dots, y_n^{(2)})$ as

$$\begin{aligned} y_i^{(1)} &= x_i^{(1)} + \beta |x_i^{(1)} - x_i^{(2)}| \\ y_i^{(2)} &= x_i^{(2)} + \beta |x_i^{(1)} - x_i^{(2)}| \end{aligned}$$
(3.2)

where β is a random number given by an inverted Laplace distribution function

$$\beta = \begin{cases} a - b \ln(u), u \le \frac{1}{2} \\ a + b \ln(u), u > \frac{1}{2} \end{cases},$$
(3.3)

for $a \in \mathbb{R}$, b > 0 and $u \in [0, 1]$ a uniformly distributed random number. The operator u is generated initially and then fixed during the crossover operation.

Blend crossover: Eshelman & Schaffer (1993) proposed a blend crossover (*BLX* – α) operator for a fixed α. Let x⁽¹⁾, x⁽²⁾ be two parent solutions. The operator randomly selects an individual in the range [x⁽¹⁾ – α(x⁽²⁾ – x⁽¹⁾), x⁽²⁾ + α(x⁽²⁾ – x⁽¹⁾)]. Then, a single offspring is given as

$$y = (1 - \gamma)x^{(1)} + \gamma x^{(2)},$$
 (3.4)

where $\gamma = (1 + 2\alpha)u - \alpha$ is a uniformly distributed factor for a fixed value of α , and $u \in [0,1]$ is a random number. According to the authors, $\alpha = 0.5$ is (in some sense) the "best" value of α with respect to algorithm performance.

Termination Criteria and Convergence

For genetic algorithms, there is a need to determine when to stop the search. Termination criteria govern the end of an evolutionary search process. Four such criteria are as follows, the first three from Bozorg-Haddad et al. (2017) and the final criterion from Jain et al. (2001).

1. Maximal runtime: The maximum time criterion terminates the algorithm after a given processing time is consumed. Then, the best solution achieved within that time is expressed.

- 2. Maximal number of generations: In this criterion, the algorithm stops when a predefined number of generations (seen as iterations of the algorithm) have been processed.
- The improvement rate of the best solution: This criterion stops the algorithm when there is no improvement in objective value (the objective being something which must be minimised or maximised).
- A given objective value attained: This criterion is met when the best objective value f^t_{*} at generation t reaches or exceeds a given limit f.

Note that there are other termination criteria, which are omitted since they are out of the scope of this discussion. Other EA parameters include population size, mutation and crossover rates, and so on.

The local optimum, in the case of optimisation problems, is the optimal maximal or minimal solution among a nearby collection of possible solutions. Such optima, of which there may be many, are referred to as being less optimal than the global optimum. They are often more practical to reach than the more desirable global optimum, and typically require crossover-type operations to escape. Local optimality has been used by (Paquete et al. 2007) to form the design and analytical basis of their algorithms, which were iterative. Both of these two improved algorithms define their search neighbourhood using feasible solution sets, and use outperformance relations as criteria for acceptance. The concept of the local optimum is central to stochastic local search algorithms' development, and many metaheuristic methods aim to conquer local optimality problems (Paquete et al. 2007). Next, we specialise our discussion to a GA coded in R, giving experimental results.

3.3 GA on the Unconstrained Problem

The unconstrained portfolio optimisation problem is given by the following formulation: Minimise

$$\sum_{i=1}^{n} \sum_{j=1}^{n} \sigma_{ij} x_i x_j.$$
(3.5)

Subject to

$$\overline{R} = (\mathbf{x}, \mu) = \sum_{i=1}^{n} \mu_i x_i,$$
(3.6)

$$\overline{R} \ge R_0 \tag{3.7}$$

$$\sum_{i=1}^{n} x_i = 1$$

 $0 \le x_i \le 1, \qquad i = 1, 2, \dots, n.$ (3.8)

where x_i is the weight of asset *i*, μ_i is the expected return of asset number *i*, \overline{R} is the expected return of the portfolio, R_0 is the desired return and *r* is the variance (risk). Note that σ_{ij} is the covariance between the returns for assets *i* and *j*.

3.3.1 Implementation

As a preliminary experiment, a parallelised genetic algorithm was coded in R to solve the unconstrained portfolio optimisation problem (specifically, the instances given by the constituent data sets of Beasley (2000)). The operators used were linear ranking selection, random position mutation and Laplace crossover, with the GA parameters of crossover probability 0.8, mutation probability 0.1 and maximum number of iterations 5000. Portfolio vectors are normalised. The R package 'GA' was used to provide the GA functionality.

From initial experiments, the above GA parameters seemed to result in more accurate solutions. The initial population used were random vectors with entries in the range 0 to 1 and the population size was 100. A 'waiting period' (i.e., the maximum number of generations the GA runs for whilst no improvement in cost is observed) of 500 was also used. The GA solves in parallel a portfolio optimisation problem for each point on the EF, generating an efficient frontier. The method is parallel and runs on multiple cores (8 cores for the unconstrained experiments).

3.3.2 EF Construction

In this subsection some results obtained using GA to construct EF are shown in Figures 3.2–3.5. Figure 3.2 shows output on dataset (D1), with 31 assets. An EF consisting of

fifty points was used. The left hand sub-figure gives a comparison between the computed EF from the GA (green dots) and Beasley's benchmark solution (in red); observe that the computed solution closely matches the prescribed solution. This verifies the correctness of the GA in this case. The EF comprises 50 points which appear in an ordered manner. Point 1 corresponds to the bottom left of the EF and point 50 on the top right of the EF.

The right hand sub-figure provides an indication of how difficult it is for the GA to solve the problem for some particular points on the EF. The red line gives the mean number of generations proceeded across all EF points. A low number of generations taken by the GA on the points at the bottom and the top of the EF indicates solving the problem is easy for the GA. The points in the middle of the EF (roughly points 15 to 31, counting from the bottom left of the EF upwards) seems to occupy the GA for a large number of generations, reaching over 4000 generations for two points. This may indicate that these points are more difficult for the GA to approximate.

From a finance perspective, the points corresponding to a lower number of generations also correspond to portfolios with fewer assets. Intuitively, this may be because there may not be a large choice of assets which offer a very low level of risk (for a low return) and very few choices that offer higher risk (at a higher return). Also, according to intuition, the middle of the EF may be more popular for investors. This section of the EF can be said to offer a fair return for a reasonable level of risk. This is where there are a larger number of combinations of assets from the complete universe of all assets (those in dataset (D1) in this case), suggesting a larger number of generations for the GA to reach a final decision. The number of combinations are given by appropriate binomial coefficients.

Moving on to Figure 3.3, this diagram shows a heatmap with rows corresponding to asset number and columns referring to the point number on the EF (with point 1 correspond to the bottom left of the EF). The rainbow colours on the heatmap represent the weights of each asset at that position in its weight vector, with the colour key representing a weight of zero by dark red and a weight of 1 by purple. Reading the heatmap from the top to the bottom, we see a clear change in distribution of assets as the point



Figure 3.2: GA Output for dataset (D1), 50 points.



Figure 3.3: Heatmap for dataset (D1), 50 points.

positions on the EF change. For instance, 'result.1' shows assets x16 and x26, and assets in the range x27-x30 have higher weights. These form the investment which generates the return and risk profile for point 1 on the EF. The highest point at the top

3.3. GA ON THE UNCONSTRAINED PROBLEM

right of the EF is given by 'result.50' on the heatmap, this showing complete investment in asset x5. In other words, the investor would put all the capital into the highest returning asset to obtain the highest return on their portfolio. Correspondingly, x5 is the highest returning asset from this dataset. It should be noted that the heatmap gives a general picture of the portfolio makeup for each position on the EF. On the heatmap are large areas of dark red; these correspond to very small (and possibly anomalous) weights and do not seem to be efficiently distinguished by the heatmap. Further insight into the data generated would be required for an efficient distinction.

Figures 3.4 and 3.5 give analogous output for dataset (D5) on 100 points. The main differences in this run is that EF points for which the GA exhibits increased runtime to find an acceptable solution (and thus, notionally, increased difficulty) is from approximately the points numbered 30 to 60 (shown on the right of Figure 3.4).

Some initial tests concerning the dataset (D5) on the scalability of the algorithm with respect to the number of cores were conducted. One hundred points were used, and the algorithm was run on an Intel i7 laptop running at 2.2GHz with 8GB of memory for these experiments. Times taken ranged from 30029.89 seconds (approximately 8 hours and 21 minutes) for two cores, through to 11105.43 seconds for four cores, and finally 5948.14 seconds (approximately 1 hour and 39 minutes) for eight cores. This shows a substantial speed improvement as the number of cores is increased, suggesting the algorithm is scalable.



Figure 3.4: GA Output for dataset (D5), 100 points.



Figure 3.5: Heatmap for dataset (D5), 100 points. Axis labels have been removed for clarity due to the large numbers of points and assets.

In the next subsection, the Markowitz M-V model is extended with a real-world constraint (cardinality constraint) and this new portfolio optimisation problem is solved using the GA.

3.4 GA on the Cardinality Constrained Problem

In order to extend the unconstrained problem to the cardinality constrained problem let: z_i be a binary variable associated to whether asset *i* is held in the portfolio. The cardinality constrained portfolio optimisation problem is (Chang et al. 2000):

Minimise

$$\sum_{i=1}^{n} \sum_{j=1}^{n} \sigma_{ij} x_i x_j.$$
(3.9)

Subject to

$$\overline{R} = (\mathbf{x}, \mu) = \sum_{i=1}^{n} \mu_i x_i,$$
(3.10)

$$\overline{R} \ge R_0 \tag{3.11}$$

$$\sum_{i=1}^{n} x_i = 1$$

 $0 \le x_i \le 1, \qquad i = 1, 2, ..., n.$ (3.12)

$$\sum_{i=1}^{n} z_i \leq K$$

$$z_i \in [0,1], \quad i = 1, 2, ..., n. \quad (3.13)$$

Inequality (3.13) limits the number of assets to a maximum of *K*. The optimal solution to the above continuous problem is known as the Cardinality-Constrained Efficient Frontier (CCEF). It should be mentioned that we solve the CCEF for $\sum_{i} z_i \leq K$ rather than $\sum_i z_i = K$ - this relaxed constraint is generally used. Some authors use strictly *K* assets (ie., $\sum_{i=1}^{n} z_i = K$) (Chang et al. (2000), Jin et al. (2014), Woodside-Oriakhi et al. (2011), Xu et al. (2010)) while some others consider portfolio with the relaxed constraint above ($\sum_{i=1}^{n} z_i \leq K$) (Schaerf (2002), Ruiz-Torrubiano & Suárez (2010)) (see Subsection 2.6.1).

3.4.1 Computational Experiments

Performance Measure

In order to measure GA performance throughout this thesis, the term IGD is used to refer to Inverted Generational Distance (Chen et al. 2012). It shows diversity and convergence of solutions to the UEF. The smaller value of IGD indicates the better set of solutions. This measure gives the average minimum distance between the computed solution from the GA CCEF from each point in the (unconstrained) UEF. It can be obtained from the formula

$$IGD = \frac{1}{p} \sum_{s^* \in S_{UEF}} \min\{d_{s^*s} | s \in S_{CEF}\}$$
(3.14)

where *p* refers to the number of points in the reference efficient frontier and S_{CEF} denotes the constrained efficient frontier, S_{UEF} refers to the unconstrained efficient frontier, s^* and *s* represent the solutions in the CCEF and UEF respectively and d_{s^*s} is the Euclidean distance between solutions s^* and *s*.

The measure (also known as an indicator) CC_{IGD} refers to the IGD between the computed GA solution and the CCEF of Cesarone et al. (2013) for the given dataset and value of *K*. Similarly, U_{IGD} refers to the IGD between the computed GA solution and the unconstrained solution, with ND_{IGD} has come to be used to refer to the sifted (nondominated) versions of those comparators used for the computation of CC_{IGD}. The work of Graham & Craven (2021), defines the *sifted EF* (dominance) to be the set of all non-dominated points. For instance, if we have two points x_1 , x_2 and x_2 is located on the right of x_1 then x_1 dominates x_2 . When this is found to be true for each pair of points, it is considered to be dominated. Thus, the sifted EF is produced by considering all non-dominated points. Note there are various other measures of GA solution quality (e.g., generational distance, hypervolume, spread (Chen et al. (2012), Deb et al. (2002))) - these are not used in this thesis.

Constraint Handling

There are four main classifications of methods for constraint handling. The first method separates constraints from objectives, which adds complexity and computational cost. The second and third methods, which involve genetic operators that produce feasible solutions, or repair methods, reduce search space size and can improve algorithm performance (Samanipour & Jelovica 2020). However, repair methods integrated with operators such as in (Mitchell et al. 2003), (Samanipour & Jelovica 2020) and (Salcedo-Sanz 2009) could be implemented. A fourth method in metaheuristics (and, in particular, GAs) is through the imposition of penalties. A penalty function is the most popular method in GAs to handle constraints, which penalises candidate solutions which violate constraints. The penalty function is used to, effectively, transform the constrained optimisation problem into the unconstrained one. This can be achieved by incorporating a penalty function into the objective function, that consists of a measure of violation of the constraints. The measure of violation is non-zero when the constraints are violated and is zero in the region where constraints are not violated. In this section, a 'penalty function' term is added to the objective function which is minimising risk, the unconstrained problems are configured. The penalty function is put upon those weights x_i that violate that constraint. Let $n_{>0}(\mathbf{x})$ be the number of non-zero weights x_i in the weight vector \mathbf{x} (in practicality, zero is implemented in the algorithm as a small tolerance value). Then the penalty is

$$10^2 (n_{>0}(\mathbf{x}) - \mathbf{K})^2.$$
 (3.15)

Thus the penalty is proportional to the number of weight violations. The choice of the factor 10^2 was arrived at through experimentation. For example, if we have cardinality constraint K = 5 and eight assets which have non-zero weights, then the total penalty would be 900.

Parameter Settings

We used a parallel GA to find the cardinality constrained efficient frontiers (CCEFs) of the datasets in Beasley (2000). The GA is coded in R, using the 'GA', 'foreach' and 'Rmpi' packages, and attempts to optimise each CCEF point (i.e., (risk, return) pair) separately. The advantage of this approach is that it is easily parallelisable and may run on as many cores as desired. All results presented in this work were run on 64 cores of an HPC cluster at the home institution. Note the CCEFs of Cesarone et al. (2013), provided for an illustrative comparison, have a 1% minimum proportion. To recap, this means that, if an asset is selected, then it must have weight of at least 0.01 (see Section 2.6.4 for a full definition).

The GA operators used were linear ranking selection, random position mutation and Laplace crossover, with crossover probability 0.5, mutation probability 0.4 and maximum number of iterations 2000*n* (*n* being the number of assets). These algorithm settings were chosen experimentally as they seem to provide the most appropriate performance. The population size was 50, and a given number of points were optimised to approximate the true CCEF that gives the best possible trade off between risk and return. In all cases in this chapter, thirty runs on each specified dataset were performed on the HPC cluster.

In the next subsection, some tests will be performed to illustrate the finding of effective GA parameters.

Preliminary Tests

First, a test to ascertain reasonable values of two of the main parameters (the maximum number of GA iterations (generations) and the GA 'run' parameter) is performed. The latter parameter is simply the number of iterations for which the GA will run if no improvement in objective function value is observed.

As stated previously, the CCEF of Cesarone et al. (2013) was created with a 1% minimum proportion taken into account (something not implemented in this subsection). Hence, strictly speaking, the values taken by the measures CC_{IGD} (the IGD between the computed GA solution and the CCEF of Cesarone et al. (2013)) and ND_{IGD} (where the sifted (non-dominated) versions of those comparators used for the computation of CC_{IGD}) are given for reference only. The results of this test are given below. Table 3.1 shows the effect of increasing the number of generations and the 'run' parameter. In this table, dataset (D2) and a value of K = 2 have been used. The values of the measures are means/SDs, respectively, from thirty runs on fifty points. The number of GA generations is denoted g. It is clear that increasing both the values of g and 'run' generally cause the IGD measurements to decrease, in terms of mean and standard deviation. In addition, it is also clear that the amount of time taken increases markedly.

The transition of the value of 'run' from 1000 to 100*n* causes the CC_{IGD} to decrease to around a quarter of that of the former value. Changing 'run' to 200*n* does not seem to substantially affect the approximation errors but the time taken approximately doubles. Thus it was decided to proceed with g = 1000n (which is the level set in Lwin et al. (2014) and Chang et al. (2000)) and run = 100*n* for all future tests. Note that the CC_{IGD} measurement tends to be smaller than its unconstrained counterpart. The non-dominated IGD is often greater than the other two measurements, due to the average distance between non-dominated points (i.e., the leftmost points in the compared solutions) being greater than that of the solution with all points (dominated or not) included.

g	Run	U _{IGD}	CC _{IGD}	ND _{IGD}	Time
1000	1000	1.480e-3/6.284e-4	1.375e-3/6.439e-4	1.503e-3/7.415e-4	30.10/2.17
500n	1000	1.042e-3/5.665e-4	9.118e-4/5.828e-4	1.252e-3/7.753e-4	82.83/18.20
1000n	1000	8.862e-4/2.582e-4	7.529e-4/2.680e-4	1.086e-4/4.842e-4	94.06/25.17
1000n	100n	3.805e-4/4.706e-5	2.067e-4/5.077e-5	3.246e-4/8.283e-5	665.23/90.74
1000n	200n	3.158e-4/1.972e-5	1.303e-4/2.360e-5	2.377e-4/7.269e-5	1251.12/150.38

Table 3.1: The effect of increasing the number of generations and the 'run' parameter.

Now, we will examine the GA runs on dataset (D3) with cardinality constraint K = 5. These are depicted in Table 3.2.

Points	U _{IGD}	CC _{IGD}	ND _{IGD}	Time
10				480.06/120.02
10	2.9176-4/9.3666-5	2.8156-4/9.4556-5	2.8436-4/6.6436-3	460.06/139.02
20	1.954e-4/4.487e-5	1.820e-4/4.592e-5	2.081e-4/3.919e-5	538.62/143.87
50	1.220e-4/2.098e-5	1.059e-4/2.155e-5	1.600e-4/3.960e-5	656.31/111.82
100	8.441e-5/1.371e-5	6.627e-5/1.426e-5	1.439e-4/4.763e-5	845.03/106.14
200	6.340e-5/6.386e-6	4.380e-5/6.778e-6	1.192e-4/3.934e-5	1374.39/127.09
500	5.090e-5/2.724e-6	3.031e-5/2.662e-6	1.014e-4/2.822e-5	2871.30/172.94

Table 3.2: The effect of varying the number of points. The statistics for the collection of runs on 50 points are reproduced from those of Table 3.5.

The above table illustrates that increasing the number of points from ten up to five hundred affects a monotonic decrease in all IGD measurements. The measurement ND_{IGD} still decreases monotonically but less rapidly than the other two IGD measurements, while the time taken, as expected, increases. It was decided to fix the number of points at fifty for remainder of this chapter (unless otherwise stated).

3.4.2 Experimental Results

In this section, the computational results of testing the algorithm on the datasets (D1)-(D5) with cardinalities K = 2,3,4 as well as the "standard" values of K = 5,10 (Jin et al. 2016) are presented. The results are shown on Tables 3.3–3.7. It is worth reminding the reader that the minimum proportion constraint is not used at present and so a comparison with the work detailed in Jin et al. (2016) is not meaningful at this stage. Thus only an illustrative comparison is given.

Table 3.3 shows that when K = 2, the least computational time, the mean results have high U_{IGD}, CC_{IGD} and ND_{IGD} errors. However, as *K* increases, the number of assets available increases and so the EF converges to the UEF. For example, when K = 10 the mean is 4.966e-5 and the standard deviation is 1.139e-8 while the time consumption is three times that when K = 2. The table also indicates that standard deviation alters greatly at different *K* values. in the case of K = 2, the standard deviation of the U_{IGD} is 2.325e-5 and of the CC_{IGD} is 2.684e-5 and of the ND_{IGD} is 2.901e-5, whereas those values for K = 10 are 1.139e-8 and 1.109e-8, respectively. Note that the value of ND_{IGD} for K = 10 was not available due to an issue with sifting the points of the solution and the CCEF but it is not expected to be commensurate with other readings judging by other tables.

K	U_{IGD}	CC _{IGD}	ND_{IGD}	Time
2	2.510e-4/2.325e-5	1.333e-4/2.684e-5	1.549e-4/2.901e-5	75.18/14.99
3	1.402e-4/6.814e-6	9.804e-5/6.929e-6	1.570e-4/2.939e-5	82.12/18.33
4	9.853e-5/6.156e-6	8.552e-5/5.600e-6	1.452e-4/5.587e-5	98.80/20.68
5	7.045e-5/4.834e-6	6.701e-5/4.453e-6	9.802e-5/1.859e-5	119.44/25.48
10	4.966e-5/1.139e-8	4.952e-5/1.109e-8	-	216.14/13.59

Table 3.3: Changing the value of the cardinality constraint, *K*, from 2 up to 10 on dataset (D1).

Table 3.4 below again demonstrates that the lowest computation time, when K = 2, gives mean results with high U_{IGD}, CC_{IGD} and ND_{IGD} errors. An increase in the *K* value again sees the number of assets available increase. For instance, when K = 10 the mean is 6.312e-5 and the standard deviation is 1.992e-6. The time consumed is double that when K = 2. The table also shows that standard deviation again alters significantly at different *K* values; for example, when K = 2, the standard deviation of the U_{IGD} is 4.706e-5 and of the CC_{IGD} is 5.077e-5 and of the ND_{IGD} is 8.283e-5, whereas the corresponding values for K = 10 are 1.992e-6, 1.990e-6 and 2.944e-5, respectively.

K	U _{IGD}	CC _{IGD}	ND_{IGD}	Time
2	3.805e-4/4.706e-5	2.067e-4/5.077e-5	3.246e-4/8.283e-5	665.23/90.74
3	2.239e-4/2.287e-5	1.405e-4/2.311e-5	2.781e-4/7.803e-5	647.11/112.25
4	1.505e-4/1.538e-5	1.055e-4/1.577e-5	2.279e-4/7.880e-5	705.27/129.94
5	1.194e-4/7.082e-6	9.304e-5/7.443e-6	2.319e-4/6.269e-5	862.57/153.76
10	6.312e-5/1.992e-6	6.025e-5/1.990e-6	1.273e-4/2.944e-5	1360.09/129.46

Table 3.4: Changing the value of K from 2 to 10 on dataset (D2).
Table 3.5 below demonstrates that increasing the value of *K* results in longer mean computation time, with K = 2 resulting in mean computation time of 649.88, compared with 1147.61 for K = 10. But this is now in contrast to the standard deviation computation time, which is 103.87 for K = 2, 127.67 for K = 4 and 89.88 for K = 10. Once more, the errors in U_{IGD}, CC_{IGD} and ND_{IGD} are higher when the *K* value is lower. At K = 2, the mean U_{IGD} is 4.136e-4, compared with the standard deviation of 1.176e-4. Likewise, the CC_{IGD} for K = 2 is 3.158e-4, whereas the standard deviation is 1.251e-4. The mean ND_{IGD} for K = 2 is 2.376e-4, compared with standard deviation of 5.912e-5. However, when K = 10, the mean U_{IGD} value is 4.578e-5, compared with standard deviation of 2.032e-6, CC_{IGD} is 4.257e-5 compared with 2.016e-6 and ND_{IGD} is 1.057e-4 compared with 3.508e-5.

K	U _{IGD}	CC _{IGD}	ND _{IGD}	Time
2	4.136e-4/1.176e-4	3.158e-4/1.251e-4	2.376e-4/5.912e-5	649.88/103.87
3	2.308e-4/5.666e-5	1.880e-4/6.091e-5	2.125e-4/3.907e-5	620.53/106.83
4	1.547e-4/2.647e-5	1.295e-4/2.778e-5	1.459e-4/3.665e-5	659.01/127.67
5	1.220e-4/2.098e-5	1.059e-4/2.155e-5	1.600e-4/3.960e-5	656.31/111.82
10	4.578e-5/2.032e-6	4.257e-5/2.016e-6	1.057e-4/3.508e-5	1147.61/89.88

Table 3.5: Changing the value of K from 2 to 10 on dataset (D3).

Table 3.6 shows a partial return to the trend of an increase in the *K* value leading to an increase in computation time. The lowest mean time is now at K = 4 (842.59), with very similar results for K = 3 (861.08), then K = 2 and K = 5 both around 900, and K = 10 by far the largest value at 1507.39. This is reflected in the standard deviation results, where K = 10 is the largest (152.61) and K = 2 is the smallest (125.14). Errors in U_{IGD}, CC_{IGD} and ND_{IGD} are: K = 2 the mean U_{IGD} is 5.356e-4, and the standard deviation is 1.117e-4. The mean CC_{IGD} for K = 2 is 3.466e-4, and the standard deviation is 1.156e-4, and for the mean ND_{IGD} K = 2 is 4.031e-4, and standard deviation is 5.138e-6, and mean CC_{IGD} is 6.438e-5 and standard deviation is 5.098e-6. Similarly to Table 3.3, the value of ND_{IGD} for K = 10 was not available due to an issue with sifting the points of the solution and the CCEF. It is thought, however, that a representative value may be

K	U _{IGD}	CC _{IGD}	ND _{IGD}	Time
2	5.356e-4/1.117e-4	3.466e-4/1.156e-4	4.031e-4/1.186e-4	898.56/129.83
3	3.563e-4/8.174e-5	2.660e-4/8.471e-5	3.353e-4/1.166e-4	861.08/140.17
4	2.550e-4/6.759e-5	2.044e-4/6.982e-5	2.390e-4/7.466e-5	842.59/125.14
5	1.825e-4/5.190e-5	1.494e-4/5.310e-5	1.602e-4/3.990e-5	900.30/134.89
10	7.031e-5/5.138e-6	6.438e-5/5.098e-6	-	1507.39/152.61

interpolated via observation of the surrounding values in this table.

Table 3.6: Changing the value of K from 2 to 10 on dataset (D4).

Table 3.7 demonstrates a similar relationship between *K* value and computation time as Table 3.5 , with the mean time being highest at K = 10 (17693.02) and lowest at K = 2 (6828.14). The standard deviation in computation time for K = 5 is the largest. However, the standard deviation in time for K = 2 is also the shortest compared to the other values of *K*. The highest U_{IGD}, CC_{IGD} and ND_{IGD} values are again seen at K = 2, with mean values of 2.006e-4, 9.438e-5 and 1.679e-4 with standard deviations 1.212e-5, 9.505e-6 and 4.197e-5, respectively.

K	U_{IGD}	CC _{IGD}	ND _{IGD}	Time
2	2.006e-4/1.212e-5	9.438e-5/9.505e-6	1.679e-4/4.197e-5	6828.14/839.50
3	1.165e-4/5.320e-6	7.034e-5/5.242e-6	1.207e-4/2.758e-5	7287.20/1198.81
4	8.189e-5/5.408e-6	5.932e-5/5.178e-6	1.194e-4/3.749e-5	8914.39/2213.69
5	6.098e-5/3.138e-6	4.910e-5/3.166e-6	1.046e-4/3.095e-5	11094.33/2477.08
10	2.594e-5/7.735e-7	2.573e-5/7.609e-7	3.429e-5/3.325e-6	17693.02/1529.95

Table 3.7: Changing the value of K from 2 to 10 on dataset (D5).

Overall, the results given in Tables 3.3-3.7 indicate that increasing the value of the cardinality constraint *K* causes a decrease in the IGD measurements. The computational time increases considerably with *K*.

Figure 3.6 (a) gives a comparison between the computed EF from the GA (green dots), Beasley's solution (in red) and the Cesarone CCEF (in blue) for the dataset (D2) with K = 5. It is worth mentioning that 50 GA points are compared with the closest 50 points out of 2000 points. The green dots agree closely with Beasley's solution at the top of the curve. Towards the bottom of the curve, the agreement is less close. For example, there is an area of non-alignment around the return values of 0.006 and 0.005, with generally fewer examples occurring between 0.002 and 0.004, and a point of considerable non-alignment at approximately 0.0035. Note that there may be portfolios (green dots) that lie outside the bounds of the figure due the presence of poor quality solutions in the EA output. The companion figure (Figure 3.6(b)) provides an indication of how difficult it is for the GA to solve the problem for some particular points on the EF. The red line gives the mean number of generations taken across all points, and the number of iterations fluctuates significantly between approximately 10000 to over 45000.



Figure 3.6: GA output (left) and solution-generation plot (right) for dataset (D2) with K = 5.

Figure 3.7 gives the heatmap for the same dataset and *K* value ((D2), K = 5). A clear change in distribution of assets can be seen: for instance, result 48 (i.e., point number 48) shows that assets x13 and x38 have high weights. The highest point at the the top right of the EF is given by result 50 on the heatmap. This shows complete investment in asset 38. This shows that the GA chooses fewer than K = 5 assets for some portfolios (which is in keeping with the definition of the problem, where $\sum_{i=1}^{n} z_i \leq K$).



Figure 3.7: Heatmap for dataset (D2) with K = 5.

Figure 3.8 shows a comparison between three different values of *K* (2, 5 and 10) in dataset (D5). It is clear that the value of *K* affects point position when K = 2 (Figure 3.8 (a)), due to the distances of the GA solution (green dots) from the Cesarone CCEF (blue line). Agreement with Beasley's solution (red line) is demonstrated at returns of 0.004, but as the returns decrease the agreement generally becomes less close. There is a further area of agreement between the GA solution and the CCEF at returns of just over 0.0015, and again at just above zero. The risk increases dramatically at return values of 0.004, with the risk level at 0.0016, compared with the lowest risk level at returns of around 0.0015. As the value of *K* increases, the GA solution draws

closer to Beasley's solution, as seen in the closer approximation of the GA solution and the Cesarone CCEF in Figure 3.8 (b), where the CCEF only moves away from that of Beasley and Cesarone significantly at returns of 0.0025, 0.001 and zero, and in Figure 3.8 (c), where the agreement is almost complete at all points, except for a very slight non-alignment at returns of 0.001 and 0.0016.



Figure 3.8: GA Output for dataset (D5) with K = 2, 5, 10.

In the next subsection, the inclusion of a high minimum proportion constraint into cardinality constrained portfolio optimisation problem is examined. Results of the GA with this new constraint are analysed. As far as we know, there is a paucity of research in the problem with high minimum proportions.

3.5 GA on the Cardinality Constrained Problem with Minimum Asset Proportions

The cardinality constrained portfolio optimisation problem with minimum asset proportions is as follows (Chang et al. 2000):

Compute min
$$r = \sigma^2 = \sum_{i=1}^n \sum_{j=1}^n \sigma_{ij} x_i x_j$$

subject to

$$\overline{R} = (\mathbf{x}, \mu) = \sum_{i=1}^{n} \mu_i x_i$$

$$\overline{R} \ge R_0$$

$$\sum_{i=1}^{n} x_i = 1$$
(3.16)

$$\sum_{i=1}^{n} z_i \le K \tag{3.17}$$

$$x_{\min} z_i \le x_i \le x_{\max} z_i \text{ for } i = 1, 2, \dots, n$$
 (3.18)

$$z_i \in \{0, 1\}$$
 for $i = 1, \dots, n$,

where x_i is the weight of asset *i*, μ_i is the expected return of asset number *i*, \overline{R} is the expected return of the portfolio and *r* is the variance (risk). Note that σ_{ij} is the covariance between the returns for assets *i* and *j*. Inequality (3.17) limits the number of assets to a maximum of *K* and (3.18) imposes the lower x_{\min} and upper x_{\max} bounds on assets. The lower bound thus gives the minimum proportion constraint (in practice, $x_{\max} = 1$) and is a common practical restriction.

At this point, the measures CC_{IGD} and ND_{IGD} begin to take effect, and the comparison with Cesarone et al. (2013) is no longer an illustrative one.

Here, penalty function is the method most commonly used in GAs to handle constraints. Using a penalty method is a way of replacing constrained optimisation problems via a sequence of unconstrained problems. Ideally, the solution of these unconstrained problems converge to the same solution as the constrained problem. By adding a

'penalty function' term to the objective function (which is minimising risk), the unconstrained problems are configured. Where constraints are violated the violation measure is nonzero, while it is zero for areas of non-violation. To include the minimum proportion constraint, a penalty function is put upon those weights x_i that violate that constraint:

$$10^{5} \sum_{x_{i} < x_{\min}} |x_{i} - x_{\min}|$$
(3.19)

such that x_i are in the $\leq K$ weights chosen. Thus the penalty is proportional to the sum of the violations of each weight. The choice of the factor 10^5 was arrived at through experimentation. In practice, this penalty is relatively small for small *n* (e.g., (D1)).

Under this constraint, the cardinality K and the minimum proportion x_{\min} are subject to an implied practical constraint

$$x_{\min}\sum z_i \le 1 \tag{3.20}$$

(where $\sum z_i \leq K$). That is, as *K* increases, x_{\min} must decrease in turn. For example, a situation where K = 10 and $x_{\min} = 0.2$ would be impossible, as a weights vector would then be $[0.2, 0.2, \dots, 0.2]$ and $\sum (w_i z_i) = 2 > 1$, thus violating the usual portfolio optimisation constraint. For example, if K = 4 and $x_{\min} = 0.24$ then a valid weights vector would be [0.25, 0.26, 0.24, 0.25].

3.5.1 Experimental Results

In this subsection, the computational results of testing the algorithm on the dataset (D4) with cardinalities K = 3 as well as minimum proportions $x_{\min} = 0.01$ and 0.3 are presented. Figure 3.9 is an example output of dataset (D4), with cardinality K = 3 and minimum proportion $x_{\min} = 0.3$. One hundred EF points were computed. The blue curve is the Cesarone et al. (2013) curve with minimum proportion $x_{\min} = 0.01$ for comparison. The operators used were linear ranking selection, random position mutation and Laplace crossover, with crossover probability 0.5, mutation probability 0.4 and maximum number of iterations 2000n. The run took around five times longer than normal, and it has been found running the GA on each point for 2000n generations (rather than 1000n generations) gives better results. It can be seen from Figure 3.9 (a) that there is an area in the middle of the EF where the computed solution from the GA

(green dots) seems to break away from the Cesarone et al. (2013) solution (the figure is scaled only upon the GA solutions found). When the return value is between 0.004 and 0.009 the GA solution (green dots) tends to drift further from the UEF and Cesarone CCEF, and the risk level increases. For example, at returns of around 0.008 the GA solution is furthest away from Beasley, and the risk has increased from its lowest value (less than 0.0005 at the bottom of the curve) to over 0.0015.

In Figure 3.9 (b), it is evident that the points from approximately point 38 of the EF upward seems to occupy the GA for a large number of generations, reaching over 190000 for several points. This may illustrate that these points are more difficult for the GA to approximate. An EF consists of 100 points and a termination criterion of 2000*n* generations were used. A population size of 50 was also specified in order to reduce the computation time.

Following on, Figure 3.9 (c) implies that the cardinality constraint is met less often than when running the GA without the minimum proportion constraint. This conclusion conforms to our prediction that the imposition of the additional minimum proportion constraint increases the difficulty of the problem and therefore the optimisation.

Finally, Figure 3.9 (d) shows a typical cost-generation curve containing a number of local cost function minima (e.g., from around 60,000-85,000 generations). The curve begins levelling out at around 50,000 generations, the cost value decreasing by around 200 over the next 150,000 generations. This may be an indication that the GA is nearing a "good" solution (with the optimal solution for $x_{min} = 0.3$ being unknown).



Figure 3.9: Plots for dataset (D4) with K = 3 and minimum proportion $x_{\min} = 0.3$.

Figure 3.10 provides plots of the GA with minimum proportion $x_{min} = 0.01$ (as in Cesarone et al. (2013)) for dataset (D4) with K = 3. Five hundred points were used. The left subfigure shows all the points, and the right the sifted (non-dominated) points. It can be seen that the points at the top of each curve, above return values of 0.009, are fewer and align more closely with Cesarone et al. (2013) curve. On the other hand, below return values of 0.009 there are significantly more points and they align much less frequently with the blue curve, as well as demonstrating higher risk. For example, the furthest points from the Cesarone curve are at returns of around 0.007, where the risk is approximately 0.0017.

Figure 3.10 (b) shows that there is a good overall agreement along the Cesarone curve, except for a couple of points at return values of 0.004 and 0.007 (though these are still very close).



Figure 3.10: Plots for dataset (D4) with K = 3 and minimum proportion $x_{\min} = 0.01$.

Figure 3.11 (left subfigure) gives the GA solution for dataset (D4) with K = 3 and minimum proportion $x_{\min} = 0.3$ plotted against the Cesarone solution of minimum proportion 0.01. The experiment took approximately 15 hours to run. The right subfigure gives an example of the sifted GA output (originally 500 points before sifting). This work has been published at the GECCO 2019 Conference (Alotaibi & Craven 2019). It can be seen that there is definite clustering behaviour from the sifted GA-computed solution (green dots) at three locations in the CCEF of Cesarone et al. (2013), suggesting that the spread of green dots along the CCEF is somehow restricted by virtue of the relatively high minimum proportion constraint. Unsurprisingly, the green dots are further away from the CCEF as given risk/return values are hard to achieve using a minimum proportion of 0.3. Comparing this result with Figure 3.10 the GA solution seems to less closely approximate the Cesarone solution because of the increase in minimum proportion from 0.01 to 0.3. The number of points is lower and they are more widely spread in the middle of the curve between the return values of 0.003 and 0.007. in addition, it is evident that the number of points at the top right is relatively low while on the left at the bottom of the curve there is a large amount of points. This can be interpreted as meaning that there is a low number of portfolios with relatively high return and risk, while there are far more portfolios with relatively low return and risk.



Figure 3.11: Plots for dataset (D4) with K = 3 and minimum proportion $x_{\min} = 0.3$ (Alotaibi & Craven 2019).

3.5.2 Further Runs with Non-Zero Minimum Proportions

In the following results, fifty points are used for each run, with thirty runs used for each set of experiments. Results are grouped by minimum proportion (x_{\min}) . Mean crowding distance (mean of a vector of distances) and difference in returns are often used as measures of spread in GAs. The crowding distance value of a solution measures the density of solutions surrounding a particular solution. It is computed by sorting the set of solutions in ascending order according to objective function values. The crowding distance is the average distance of two neighbours of each point along each of the objectives (Raquel & Naval Jr 2005). Also, the range of return is the difference between the minimum and maximum return values from the sifted solution produced by the GA (this is, in some way, a measure of spread of the solution).

GA Solutions with $x_{\min} = 0.01$

There have been several studies in the literature reporting the effect of applying cardinality and minimum proportions constraints in terms of solution quality (IGD) and computational cost. Previous studies have concentrated on $x_{\min} = 0.01$. (Jin et al. 2016, Table 4) reported that the mean IGD value for (D3) with $K \le 10$ and $x_{\min} = 0.01$ is 3.24e - 5 (recall, that is 3.24×10^{-5}) which is similar to what we found in Table 3.8, where the IGD is 5.078e - 5 under the same constraints and for the same number of points. Also, for (D5) with $K \le 10$ and $x_{\min} = 0.01$ the IGD value on Table 4 of the above work is 2.35e - 5. This is almost as the same as what we found in Table 3.9 (which is 2.735e - 5). The work of Chang et al. (2000) stated IGDs for (D3) and (D5) of 4.45e - 5and 2.48e - 5 respectively, whereas Lwin et al. (2014) reported the IGDs are 5.93e - 5and 3.86e - 5 for (D3) and (D5) with K = 10 and $x_{\min} = 0.01$.

Table 3.8 gives the results of the experiment implemented on dataset (D3) with different values of *K*. The lower and upper bounds of the weights are set as $x_{\min} = 0.01$, $x_{\max} = 1$. The results indicate that the approximation error of U_{IGD}, CC_{IGD} and ND_{IGD} decreases when the value of *K* increases. The time taken also increases when the value of *K* increases; for example, the algorithm spent much more computational time when K = 10, at approximately double that when K = 2. There are several studies in the literature which report on the case of (D3) with K = 10 and $x_{\min} = 0.01$. As stated above

Jin et al. (2016) found that IGD = 3.24e-5. Jin et al. (2016) exhibited a numerical study of different models in literature. Ruiz-Torrubiano & Suárez (2010) found that IGD = 3.24e-5. Their finding is consistent with those of Maringer & Kellerer (2003), who also reported IGD = 3.24e-5. Similarly, we found that IGD = 5.078e-5, which is in good agreement with the results of these past studies.

Comparing the above results with the case of only cardinality constraint on (D3), Table 3.5, it can be seen that when the minimum proportion $x_{\min} = 0.01$ is applied, the mean of all IGD measurements is very slightly higher than those from the cardinality constraint case. However, there is a larger difference in standard deviation on all IGD measurements. In terms of computational time, it is approximately three times more than those from the cardinality constraint case when imposing minimum proportion for all values of *K*.

K	U _{IGD}	CC _{IGD}	ND _{IGD}	Time
2	4.936e-4/1.386e-4	3.955e-4/1.443e-4	3.563e-4/1.637e-4	1479.39/213.74
3	2.816e-4/9.237e-5	2.402e-4/9.565e-5	2.760e-4/1.133e-4	1507.59/174.32
4	2.009e-4/5.789e-5	1.772e-4/5.665e-5	2.181e-4/8.974e-5	1577.49/173.60
5	1.442e-4/4.281e-5	1.283e-4/4.391e-5	1.911e-4/5.697e-5	1660.56/207.07
10	5.387e-5/3.456e-6	5.078e-5/3.551e-6	1.019e-4/2.769e-5	2182.39/577.45

Table 3.8: Running dataset (D3) with $x_{\min} = 0.01$ and the values of K shown.

Figure 3.12 provides the GA solution for dataset (D3) with K = 2,3,4,5 and 10 and minimum proportion $x_{\min} = 0.01$ plotted against the Cesarone et al. (2013) solution of minimum proportion 0.01. The subfigure (a) gives an example of the sifted GA output. It is noticeable that it has an unusual shape, whereas subfigure (e) shows the GA solution is much closer to that of Cesarone (and also the UEF, since the two solutions almost coalesce in this case), and therefore a far smoother curve. Figure 3.12 (b) shows that when K = 3, the GA solution aligns a little better with the UEF, though it still only agrees at the top of the curve, between return values of 0.007 and 0.008, with risk levels of 0.0006 and upwards. Even here, there are a few exceptions, such as at returns of about 0.0075, where non-alignment is seen. Towards the bottom of the curve, at return values of below 0.007, the non-alignment is significant, and points are



Figure 3.12: Plots for dataset (D3) with different values of *K* and minimum proportion $x_{\min} = 0.01$.

clustered together around return values of approximately 0.006, whereas the points are more spread out towards the lower return values. In Figure 3.12 (c), when K = 4, alignment with the Cesarone CCEF is seen throughout the return values of 0.006 up to 0.008. Although the risk for returns of between 0.005 and 0.006 shows some higher value instances than in Figure 3.12 (b), there is also a greater number of points with lower risk levels, and a larger concentration of points towards the bottom of the curve. The agreement in Figure 3.12 (d) (K = 5) is quite similar to 3.12 (c) (K = 4), with good agreement again at return values above 0.0065. However, the agreement below 0.006 now shows significantly better overall agreement, with risk levels decreasing to just 0.0002-0.0003 at return values of 0.002-0.004.

When K = 10, in Figure 3.12 (e), the closest agreement is seen between the computed CCEF, Beasley's UEF and Cesarone's CCEF, with excellent alignment throughout return values of 0.0055 and 0.008, and only slight non-alignment below returns of 0.0055. This risk now aligns with that of the UEF in roughly above a half of the curve, meaning that returns of 0.008 have attached risk of 0.0014. Returns of 0.002-0.004 have much lower risk at around 0.0002-0.0003, with the GA solution exhibiting a (comparatively small) difference with the lower part of the UEF.

Table 3.9 shows the results from running dataset (D5) with cardinalities from K = 2 to K = 10 and $x_{\min} = 0.01$. Comparing with Table 3.7, the mean of U_{IGD} and CC_{IGD} measurements are slightly increased in every case for this value of x_{\min} . Thus, it is clear that the minimum proportion constraint (as it is small) does not make a big difference to the solution quality IGD measurements, though it can affect the computational time a great deal. Compared with Table 3.7, the computation time roughly doubles in each case when the minimum proportion constraint is applied; this may be because the GA spends the extra time attempting to satisfy this constraint. From this table, a mean IGD of 2.735e-5 was found for K = 10. This is consistent with the studies of Maringer & Kellerer (2003) and Ruiz-Torrubiano & Suárez (2010), which showed a mean of IGD = 2.35e-5 for (D5).

K	U _{IGD}	CC _{IGD}	ND _{IGD}	Time
2 3 4 5	2.015e-4/1.257e-5 1.243e-4/8.955e-6 8.580e-5/4.975e-6 6.381e-5/3.786e-6 2.766e-5/1.217e-6	9.588e-5/1.067e-5 7.772e-5/8.039e-6 6.323e-5/4.795e-6 5.719e-5/3.542e-6 2.735e-5/1.171e-6	1.780e-4/4.792e-5 1.366e-4/4.270e-5 1.198e-4/2.858e-5 9.455e-5/2.238e-5 4.070e-5/5.164e-6	15774.54/1643.76 16020.63/2079.98 16764.75/2670.54 19501.44/2333.95 33219 89/4271 14

3.5. GA ON THE CARDINALITY CONSTRAINED PROBLEM WITH MINIMUM ASSET PROPORTIONS

Table 3.9: Dataset (D5) with $x_{\min} = 0.01$ and the values of K shown.

Figure 3.13 provides the GA solution for dataset (D5) with K = 2, 3, 4, 5 and 10 and minimum proportion $x_{\min} = 0.01$ plotted against the Cesarone et al. (2013) solution of minimum proportion 0.01. What is striking in this figure is that although the number of assets is large (i.e., 225), the GA solution is getting closer to UEF as *K* increases. In fact, the GA CCEF and UEF are very close when K = 10 (subfigure (e)). Comparing this with Figure 3.12, it is clear that better agreement is seen when K = 10, resulting in almost total alignment, except for a couple of small areas between return values of 0.000 and 0.002. At lower *K* values the results are similar to those in Figure 3.12, with the green dots moving further away from the curve as *K* decreases.

Based on the above results both cases where $x_{\min} = 0.01$ are of the same order of accuracy of Tables 4 and 5 of Jin et al. (2016). (The times exhibited on Table 4 of Jin et al. (2016) cannot be compared with the work exhibited here since the computational hardware is different, insofar as the CPU used and also the number of cores used. The aforementioned work also gives the time per point, which may be scaled, but the non-comparability is still unaffected).



Figure 3.13: Plots for dataset (D5) with different values of *K* and minimum proportion $x_{\min} = 0.01$.

Table 3.10 below provides a comparison between crowding distance (CD) and range of return (RR) measurements for datasets (D3) and (D5) with $x_{min} = 0.01$. Respectively, these correspond to the data in Tables 3.8 and 3.9. There are similar levels of decrease in the mean of CD measurement and increase in the mean of RR measurement for both datasets. An interpretation may be that increasing the number of assets, *n*, does not affect the spread of points across the curve covered by the GA solution a great deal.

	D3		D5	
K	CD	RR	CD	RR
2 3 4 5 10	8.165e-4/1.769e-4 6.553e-4/8.191e-5 5.821e-4/5.111e-5 5.349e-4/4.594e-5 4.310e-4/2.617e-5	3.578e-3/9.028e-4 4.114e-3/7.383e-4 4.349e-3/7.229e-4 4.655e-3/5.198e-4 5.287e-3/4.135e-4	7.499e-4/1.400e-4 5.892e-4/9.084e-5 4.727e-4/5.448e-5 4.006e-4/3.978e-5 2.703e-4/1.645e-5	3.319e-3/4.934e-4 3.407e-3/3.952e-4 3.454e-3/3.886e-4 3.590e-3/2.546e-4 3.839e-3/8.010e-5

Table 3.10: CD and RR measurements for datasets (D3) [Table 3.8] and (D5) [Table 3.9] with $x_{\min} = 0.01$ and the values of *K* shown.

It is worth noting that, for small values of K, the proportions used are likely to be above 1%. This will make the difference at the bottom left of the CCEF only since a larger number of assets are used to access the small risk areas of the CCEF.

In the next part, the minimum proportion is increased from $x_{\min} = 0.01$ to $x_{\min} = 0.1$.

GA Solutions with $x_{min} = 0.1$

Here, to recap, the runs treated in this section are indicative only as (Cesarone et al. (2013)) uses $x_{\min} = 0.01$. For a value of K = 10, the minimum proportion $x_{\min} = 0.1$ implies that the EF will consist of only a single solution with very tight constraints. However, we do not get a perfect solution (as we are using an approximative method, an EA) and so some weights may not strictly obey the constraint $w_i \ge x_{\min}$.

Table 3.11 gives the results after applying the minimum proportion $x_{min} = 0.1$ to dataset (D3). Comparing with the results from Table 3.8 it could be argued that the hardness of problem has increased. That is, comparing with results from Table 3.8, it can be

seen that the values of the IGD measurements are commensurate with those where $x_{\min} = 0.01$ but the time taken has more than doubled.

K	U _{IGD}	CC _{IGD}	ND _{IGD}	Time
2	3.787e-4/6.273e-5	2.758e-4/6.773e-5	2.541e-4/5.871e-5	3659.31/341.53
3	2.256e-4/5.509e-5	1.794e-4/5.840e-5	2.358e-4/6.027e-5	3603.05/434.33
4	1.650e-4/2.917e-5	1.383e-4/3.047e-5	1.824e-4/4.400e-5	3628.71/285.05
5	1.332e-4/2.527e-5	1.158e-4/2.615e-5	1.916e-4/4.469e-5	3897.14/458.43
10	6.837e-5/4.902e-6	6.513e-5/4.956e-6	1.249e-4/3.079e-5	5848.68/427.40

Table 3.11: Running dataset (D3) with $x_{\min} = 0.1$ and the values of K shown.

Moving on to Table 3.12, where the CD and RR measurements are presented, we can see that when the value of K increases the standard deviation of CD and RR measurements generally decrease and also the mean of CD decreases while the mean of RR increases.

K	CD	RR
_		
2	8.606e-4/1.568e-4	4.146e-3/6.590e-4
3	6.772e-4/8.832e-5	4.373e-3/6.799e-4
4	6.432e-4/8.753e-5	4.492e-3/5.953e-4
5	5.803e-4/6.304e-5	4.540e-3/3.936e-4
10	5.502e-4/5.494e-5	5.212e-3/4.442e-4

Table 3.12: Values of CD and RR for dataset (D3) with $x_{\min} = 0.1$ and varying K.

Figure 3.14 provides the GA solution for dataset (D3) with K = 2,3,4,5 and 10 and minimum proportion $x_{min} = 0.1$ plotted against the Cesarone et al. (2013) solution of minimum proportion 0.01. The subfigure (a) gives an example of the sifted GA solution. It is noticeable that it has an unusual shape (expected, since K = 2), whereas the other subfigures have smoother curves.

Figure 3.14 (b) shows that when K = 3, the GA solution aligns a little better with the UEF and the Cesarone CCEF, though it still only agrees at the top of the curve, above return values of 0.007, with risk levels of 0.0008 and upwards. Towards the bottom of the curve, at return values of below 0.007, the non-alignment is significant, and points are more spread out towards the lower return values.

In Figure 3.14 (c), when K = 4, alignment with Cesarone CCEF is seen throughout the return values of just above 0.006. The agreement in Figure 3.14 (d) (K = 5) is quite similar to 3.14 (c) (K = 4), with good agreement at return value above 0.006. However, the agreement below 0.006 now shows significantly better overall agreement, with risk levels decreasing to just 0.0002-0.0004 at return values of 0.002-0.005. The greatest area where the GA solution differs from that of Cesarone CCEF is between return values of 0.004 and 0.0055, where risk is between 0.0004 and 0.0006. When K = 10, in Figure 3.14 (e), the closest agreement is seen between the GA solution, Beasley's UEF and Cesarone's CCEF, with alignment throughout return values of 0.0055 and 0.008, and the only non-alignment occurring below returns of 0.0055. The risk now aligns with that of the UEF, meaning that returns of 0.008 have attached risk of 0.0016, whereas returns of 0.002-0.004 have much lower risk at around 0.0002-0.0003. To conclude, with reference to both Figures 3.12 and 3.14, it is evident that they share the same trends. However, the GA solution is slightly further from UEF when the minimum proportion increases to 0.1.



Figure 3.14: Plots for dataset (D3) with different values of *K* and minimum proportion $x_{\min} = 0.1$.

3.5.3 Further Experiments of Increasing *x*_{min}

In this section we fix a dataset and value of *K*, and increase x_{\min} . Note that the cluster had a time limit of three days on any computational work. Thus, where applicable (e.g., $x_{\min} > 0.01$ and small *K*), experiments were run in batches and the statistics collected.

Here, Table 3.13 illustrates the performance of running the algorithm on (D3) with K = 2 and different values of minimum proportion. Each line of this table encapsulates fifty points for each of thirty runs. The data corresponding to $x_{\min} = 0.01$ and $x_{\max} = 0.1$ is taken from Tables 3.8 and 3.11 respectively. It is apparent that when $x_{\min} = 0.01$, the mean and standard deviation of all IGD measurements are higher than those when $x_{\min} = 0.1, 0.2, 0.25$ (except the mean value of ND_{IGD} when $x_{\min} = 0.25$). At $x_{\min} = 0.3$ the mean and standard deviation of all IGD measurements increase significantly compared with other values of x_{\min} , for example the mean ND_{IGD} of 8.072e-4 while the standard deviation slightly increases compared with $x_{\min} = 0.1, 0.2, 0.25$. With respect to runtime, it is evident that when x_{\min} is increased, the time taken increases rapidly.

<i>x</i> _{min}	U _{IGD}	CC _{IGD}	ND _{IGD}	Time
0.01	4.936e-4/1.386e-4	3.955e-4/1.443e-4	3.563e-4/1.637e-4	1479.39/213.74
0.1	3.787e-4/6.273e-5	2.758e-4/6.773e-5	2.541e-4/5.871e-5	3659.31/341.53
0.2	3.509e-4/7.439e-5	2.426e-4/7.987e-5	2.772e-4/8.399e-5	6925.63/133.35
0.25	4.239e-4/6.384e-5	3.199e-4/6.771e-5	4.138e-4/1.494e-4	7275.06/87.57
0.3	8.088e-4/3.905e-4	7.233e-4/4.037e-4	8.072e-4/3.173e-4	7353.84/98.13

Table 3.13: Running dataset (D3) with K = 2 and the values of x_{\min} shown.

Figure 3.15 shows the efficient frontiers of testing the problem with different values of minimum proportion (0.01, 0.1, 0.2, 0.25 and 0.3) on (D3) with K = 2. Note that the computational results comparing our solution with Cesarone et al. (2013) seem to spread away from the Cesarone CCEF when x_{min} increases. This may be explained by the fact that imposing these minimum proportions makes the optimal points difficult to approximate. A better agreement of GA solution can be seen when $x_{min} = 0.01$ (subfigure (a)) compared to the other subfigures. At higher x_{min} values, the green dots move further away from the Cesarone CCEF as the minimum proportion x_{min} increases.



Figure 3.15: Plots for dataset (D3) with K = 2 and minimum proportions $x_{\min} = 0.01$, 0.1, 0.2, 0.25 and 0.3.

Table 3.14 shows the results of running the algorithm on (D3) with K = 3 and various minimum proportions. The results of the IGD measurements are similar to those reported in Table 3.13. The IGD measurements (when K = 3) are slightly smaller than those when K = 2 (Table 3.13), which is to be expected since the problem is known to be harder the smaller the value of K. Also, it can be seen that time taken is similar in both tables.

<i>x</i> _{min}	U_{IGD}	CC _{IGD}	ND _{IGD}	Time
0.01	3.087e-4/1.151e-4	2.689e-4/1.187e-4	2.750e-4/1.029e-4	1471.48/144.17
0.1	2.256e-4/5.509e-5	1.794e-4/5.840e-5	2.358e-4/6.027e-5	3603.05/434.33
0.2	2.253e-4/2.267e-5	1.766e-4/2.436e-5	2.674e-4/7.108e-5	6843.36/130.98
0.25	3.221e-4/6.173e-5	2.794e-4/6.579e-5	3.794e-4/8.998e-5	7277.73/119.54
0.3	5.99e-4/2.743e-4	5.643e-4/2.781e-4	6.643e-4/3.167e-4	7377.33/102.26

Table 3.14: Dataset (D3) with K = 3 and the values of x_{\min} shown.

Table 3.15 presents CD and RR measurements for (D3) and K = 2, K = 3 with x_{min} ranging from 0.01 to 0.3. In the cases of $x_{min} = 0.01$, 0.1 and 0.2 the mean CD measurements display very little difference (the order of magnitude in each case is identical). However, at $x_{min} = 0.25$ the mean CD measurement is around 120 times what it was in the row above. When we progress to $x_{min} = 0.3$, the size of the mean CD measurement increases by approximately 39 times (for K = 2) and four times (for K = 3). This may occur as the problem is effectively overconstrained: when the minimum proportion is increased that leads to restrict the solution to increasingly small portions of the curve, resulting in less convergence and lower spread.

	<i>K</i> =	= 2	<i>K</i> =	= 3
x _{min}	CD	RR	CD	RR
0.01 0.1 0.2 0.25 0.3	8.165e-4/1.769e-4 8.606e-4/1.568e-4 9.584e-4/1.704e-4 0.117/0.416 4.558/3.913	3.578e-3/9.028e-4 4.146e-3/6.590e-4 4.341e-3/7.395e-4 4.894e-3/7.335e-4 4.874e-3/1.090e-3	6.792e-4/1.135e-4 6.772e-4/8.832e-5 8.855e-4/2.013e-4 0.652/1.773 2.394/2.969	4.337e-3/9.299e-4 4.373e-3/6.799e-4 4.687e-3/6.146e-4 4.834e-3/6.612e-4 5.247e-3/6.968e-4

3.5. GA ON THE CARDINALITY CONSTRAINED PROBLEM WITH MINIMUM ASSET PROPORTIONS

Table 3.15: CD and RR measurements for (D3) and K = 2, K = 3 with x_{min} ranging from 0.01 to 0.3.

Figure 3.16 shows a comparison between 5 different values of x_{\min} (0.01, 0.1, 0.2, 0.25 and 0.3) in dataset (D3) with K = 3 plotted against the Cesarone et al. (2013) solution of minimum proportion 0.01. It is clear that the value of x_{\min} affects point position when $x_{\min} = 0.01$ (Figure 3.16 (a)), and it can be seen also that all points at the top of the curve are very close to the UEF, while in the middle under the value of return of 0.006 they move further from the Cesarone CCEF. In subfigures (b) and (c), as x_{\min} are increased to 0.1 and 0.2 respectively, above the value of return of 0.007 the points align closely with the UEF. Then, the risk increases dramatically under return values of 0.007 compared to the previous subfigure. Subfigure (d) shows the GA solution with $x_{\min} = 0.25$. Here, a big gap can be seen in the central portion of the plot, between return values of return 0.004 and 0.006. Finally, subfigure (e) illustrates that increasing the minimum proportion to 0.3 leads to a restriction of the GA solution, and scaling issues resulting in many points not being included in the relevant plot. Comparing this with Figure 3.15, it is clear that the GA solution is closer to the UEF when K = 3.



Figure 3.16: Plots for dataset (D3) with K = 3 and respective minimum proportions $x_{\min} = 0.01, 0.1, 0.2, 0.25, and 0.3.$

Table 3.16 shows the results obtained on dataset (D5) with K = 5 and different values of x_{min} . It is evident that increasing the minimum proportions leads to a slight increase in IGD measurements. However, an increase in computational time also results when larger minimum proportions are used, due to the increased number of computations. For example, when 0.01 minimum proportion is used, a relatively low mean computation time of 19501.44 is observed, whereas when using 0.2 minimum proportion, this time increases to 82200.30.

Looking at the same examples, it is clear that the distance between the GA solution and the UEF is smaller when the minimum proportion is smaller (U_{IGD} of 6.381e-5 for $x_{min} = 0.01$, compared with 2.723e-4 for $x_{min} = 0.2$). The standard deviation for U_{IGD} is higher for $x_{min} = 0.2$ than for $x_{min} = 0.01$, indicating greater variation in the values when using higher minimum proportion. The results are similar for the CC_{IGD} and ND_{IGD}, with $x_{min} = 0.01$ producing a mean distance of 5.719e-5 and 9.455e-5, compared with 2.617e-4 and 4.629e-4 for $x_{min} = 0.2$, therefore indicating better agreement with the UEF at $x_{min} = 0.01$ than at $x_{min} = 0.2$. Again, the standard deviation is higher for $x_{min}=0.2$ than for $x_{min} = 0.01$, indicating more variation with higher value of minimum proportion. In conclusion, it is observed that as the minimum proportion used decreases, the GA solution moves closer to the UEF. There is also a clear trend in time, which grows as the minimum proportion is increased. As already stated, the above evidence points to satisfaction of the belief that imposition of a high minimum proportion constraint results in increased problem difficulty (Lwin et al. 2014).

<i>x</i> _{min}	U _{IGD}	CC _{IGD}	ND _{IGD}	Time
0.01	6.381e-5/3.786e-6	5.719e-5/3.542e-6	9.455e-5/2.238e-5	19501.44/2333.95
0.05	6.832e-5/4.373e-6	5.630e-5/4.332e-6	9.997e-5/1.714e-5	36427.79/3570.22
0.1	8.089e-5/5.922e-6	6.822e-5/5.644e-6	1.213e-4/2.541e-5	58766.00/5058.05
0.15	9.595e-5/7.665e-6	8.364e-5/7.637e-6	1.623e-4/5.676e-5	78896.26/3655.38
0.18	1.452e-4/1.979e-5	1.331e-4/1.994e-5	2.759e-4/1.156e-4	82053.57/3260.68
0.2	2.723e-4/1.255e-4	2.617e-4/1.263e-4	4.629e-4/2.524e-4	82200.30/3211.28

Table 3.16: Running dataset (D5) with K = 5 and the values of x_{\min} shown.

Table 3.17 shows the CD and RR measurements for the (D5) dataset at different x_{min} values. The trend is that a higher x_{min} value results in a higher mean CD. The standard deviation in CD becomes higher when x_{min} increases; for example, with $x_{min} = 0.01$ it is 3.978e-5 and with $x_{min} = 0.2$ it is 2.456e-2. This is clearer in Figure 3.17 where, as x_{min} increases, the distance between (GA CCEF) points generally increases, causing these points to spread further from the Cesarone et al. (2013) curve (more details in the next paragraph). Moving on to the range of return, with $x_{min} = 0.2$ the mean value is lower whereas when $x_{min} = 0.05$ this RR value is the largest. For the standard deviation of RR, the trend is that a higher x_{min} value results in a higher standard deviation.

<i>x</i> _{min}	CD	RR
0.01	4.006e-4/3.978e-5	3.590e-3/2.546e-4
0.05	4.301e-4/3.845e-5	3.622e-3/2.229e-4
0.1	4.936e-4/6.126e-5	3.563e-3/2.666e-4
0.15	5.543e-4/6.908e-5	3.502e-3/3.739e-4
0.18	6.401e-4/9.703e-5	3.160e-3/6.409e-4
0.2	6.819e-3/2.456e-2	2.690e-3/8.914e-4

Table 3.17: CD and RR measurements for (D5), K = 5, with minimum proportions ranging from 0.01 to 0.3.

Figure 3.17 provides a comparison between 6 different values of x_{\min} (0.01, 0.05, 0.1, 0.15, 0.18 and 0.2.) in the dataset of (D5) with K = 5 plotted against the Cesarone et al. (2013) solution of minimum proportion 0.01. It can be seen that from subfigure 3.17 (a) onwards, the agreement with the Cesarone CCEF becomes generally less close, which confirms that increasing the value of the minimum proportion leads to higher risk. For example, comparing subfigures (c), when $x_{\min} = 0.1$, and (e), when $x_{\min} = 0.18$, it can be seen with values of return between 0.000 and 0.0025 that the risk is between 0.0003-0.0006 and 0.0003-0.0008 respectively.



Figure 3.17: Plots for dataset (D5) with K = 5 and minimum proportions $x_{\min} = 0.01$, 0.05, 0.1, 0.15, 0.18 and 0.2.

3.6 Summary of the Chapter

This chapter has discussed the detailed GA for solving mean-variance cardinality constrained portfolio optimisation problems, firstly, in the case with no minimum proportion constraints and then, secondly, in the case with high minimum proportion constraints. Results were given for five benchmark datasets. As detailed at the beginning of this chapter, there is a literature gap in the case of high minimum proportions and also a gap in the use of metaheuristics approximating solutions for small *K*. Addressing these gaps is the main contribution of this chapter.

At the beginning of the chapter, the GA was applied to the unconstrained portfolio optimisation problem to find the EF. Then, in the second part of this work, the cardinality constraint that limits the maximum number of assets in a portfolio was considered. The obtained efficient frontier is a combination of portfolios with varying numbers of assets.

The GA was then implemented to solve the portfolio optimisation problem with new varying values of *K*. To compare with the computed EF from the GA, Beasley's solution and the Cesarone CCEF were given for the datasets (D1)-(D5) with cardinalities K = 2,3,4 as well as the "standard" values of K = 5,10. It was found that increasing the value of the cardinality constraint *K* causes a decrease in the IGD measurements and the computational time increases considerably with *K*. Empirical results demonstrate that comparatively small cardinalities (e.g., K = 10) can achieve performance which, in practicality, does not vary markedly from an unconstrained frontier (whereas for K = 2,3 or 4, for example, this conclusion does not hold). As noted in Section 3.4, when the basic model is extended with the cardinality constraint, the problem becomes more difficult.

Finally, the GA was used to solve the portfolio optimisation problem for different cardinalities and high minimum proportion constraints. It was found that increasing the minimum proportions leads to an increase in IGD measurements and that led to restricting the solution to increasingly small portions of the GA CCEF, resulting in less convergence and lower spread. Further research indicates that increasing the minimum proportion constraint causes a considerable increase in computational time, thus increasing problem difficulty as is shown in Section 3.5. Also, there are differences that occur in the shape of efficient frontier curves when applying such constraints.

It is acknowledged that there are other objective functions that may be used in the GA and problem formulation than the standard risk and return. For example, there is the Sharpe ratio, i.e., the ratio of return and the square root of the risk (see Section 4.7 for further details). This should be maximised; however, the point is made here that the Sharpe ratio is a measure of portfolio performance rather than a classical objective function which replaces either risk or return (although both are clearly still calculated). Including this as an objective would typically result in a radically different efficient frontier than is commonly observable from that of the standard M-V formulation. Otherwise, it is possible to include the Sharpe ratio as a third objective and to perform multi-objective optimisation using NSGA-II, for example. This is an extension on the type of optimiser used in the current work.

There are some limitations in using the GA algorithm, such as local optima that arise during the evolution process. Also, genetic algorithms do have a major limitation, which is the coding of the fitness function to achieve higher fitness and the production of better solutions. Choosing the wrong fitness function can have significant effects, including rendering the solution to a problem unobtainable, or worse. Finally, although various GA operators were tried, and the final combination of operators were arrived at through copious experimentation, there may be more specific (possibly unknown) problem-dependent operators which may affect the algorithm.

It has been demonstrated that the GA code is adaptable and produces results comparable to the literature (Chang et al. 2000, Jin et al. 2016, Maringer & Kellerer 2003, Ruiz-Torrubiano & Suárez 2010, Cesarone et al. 2013). This naturally leads to two sets of questions which provide a departure to Chapters 4 and 5 respectively:

- In Chapter 4 we extend the mean variance portfolio optimisation problem with class constraint and test the operation using difficult real datasets with problems of larger sizes. The following research questions are addressed subject to the above objective.
 - 1. What is the impact of using other types of constraints, such as class con-

straint, on the real-world application of Gulf stock markets?

- 2. Which types of assets (stocks, oil or precious metals) significantly affect the optimal portfolio construction?
- 3. How do other commodities such as oil and precious metals influence the optimal portfolios in GCC countries?
- 4. What is the effect of including multiple classes in portfolio optimisation? What are the benefits and risks of allocating investors' funds across different classes?
- 5. How is it important for investors to diversify their portfolio when allocating their funds across different classes?
- Variance is used as the risk measure in portfolio optimisation models, but there are numerous advantages both theoretically and practically to the use of other risk measures and so it is somewhat surprising that these other measures are not more commonly used. Therefore, an alternative approach is taken into consideration, in which portfolio holdings resulting from risk measures other than variance are compared. The Worst Case Conditional Value at Risk (WCVaR) based on copulas accurately assesses risk to investment. A copula is a tool that allows the separation of distributions of individual assets from their dependence structure and is also used to model nonlinear dependence. This is beneficial because the marginal distributions of assets can be assumed to follow any distribution and copulas can be used to derive the joint multivariate distribution, accounting for the marginals dependence structure. Therefore, using a copula in financial data more accurately represents the joint distribution of losses and returns.

The WCVaR as a recent risk measure is examined in Chapter 5 to address the following research questions.

- 1. What is the advantage of using WCVaR based on copula as a risk measure?
- 2. How accurate can WCVaR be?
- 3. To what extent is WCVaR better than traditional methods?

Chapter 4

Constructing Class-Based Portfolios in Gulf Markets with a Genetic Algorithm

4.1 Introduction

Optimal portfolio construction is a challenging topic that concerns both academics and practitioners who want to maintain profitable wealth. Many models have been developed in the past in order to identify the optimal portfolio, yet there is still no agreement about the best portfolio construction methods. The optimal portfolio therefore still remains one of the most pertinent topics in finance. However, there is an increasing attention on including specific constraints that reflect the actions that traders commonly perform on assets in the market; for example, the class constraint (Jin et al. 2016). It is generally agreed that imposing a class constraint allows for a safer and more diversified portfolio, by limiting the proportion of investments made in assets with shared characteristics.

This chapter examines the application of a Genetic Algorithm (GA) to solve the class constrained portfolio optimisation problem on a large number of assets, comparing the results with those obtained by the cardinality constrained approach of Chapter 3.

GAs are metaheuristic optimisers, where the minimum or maximum of a function is not able to be established analytically. A population of potential solutions is evolved using an iterative approach inspired by natural selection, and the applications for this heuristic are widespread, including finance (Lin et al. 2004). These GAs are simple yet powerful, and have shown potential to improve or even replace other methods. Where many techniques for solving computational problems require searches of numerous conceivable solutions, genetic algorithms can solve hard problems more efficiently and more reliably (Chang et al. 2000), due to their application of parallelism, searching a small yet meaningful portion of the possible search space. This means they can be connected to and interfaced with existing models and simulations (for further information on problems upon which GAs have been employed, see Section 1.1), and hybridised easily so that problem-specific data can be input. This makes GAs very efficient as optimisation algorithms, and they are particularly useful in finding approximate solutions to optimisation problems (Drake & Marks 2002). In this chapter, we examine portfolio optimisation problems over the assets in Gulf Cooperation Council countries.

The Gulf Cooperation Council (GCC) countries have oil based economies, resulting in great dependence on oil exports (Maghyereh et al. 2017). The share prices of the GCC's oil based companies are influenced by the rise and fall of oil prices, in that when oil prices increase, this positively influences the share prices of the oil companies in the GCC countries, which in turn improves the performance of banks and other industries. Oil, like precious metals, is traded in US\$, even in the emerging Gulf markets. Yet, despite the importance of oil in these markets, a search of the relevant literature yielded limited comparative studies on portfolio optimisation problems in the GCC markets.

In this work we focus on constructing optimal portfolios that include oil, at least one precious metal (gold, silver or copper) and stocks of firms in the Gulf market to minimise risk based on a given level of return. We use an extended mean-variance model with class constraint and method developed earlier in Chapter 3. We examine the return and risk, and investigate the usefulness of including oil and precious metals in hedging equity portfolios, using daily data from January 2008 to January 2019 for the GCC countries. Also, we examine time shifting frontiers on GCC markets.

We find that lower risk can be achieved with the inclusion of oil and metals, given their potential as hedging and diversifying instruments. Thus, diversification opportunities are also expanded, and we can yield even better returns when other commodities are potentially added. Since the GCC countries are oil producers it could be useful to Gulf investors to invest in oil. In addition, it is noticed that an inward/upward shift of a portfolio's frontier denotes that higher returns can be expected for the same risk levels. We also find that economic events affect frontiers.

This chapter is structured as follows:

First of all, Section 4.2 provides some background on the GCC markets. After that, a summary of the relevant literature is reviewed in Section 4.3, after which in Section 4.4 detailed descriptions of the data used in the computational analysis of this study are presented, together with an outline of GCC stock markets, oil and precious metals. This data is described and some preliminary analysis is conducted. Then, methods used in this chapter are discussed in Section 4.5. Finally, the empirical results are presented and analysed in Sections 4.6–4.7, followed by a summary and conclusions in Section 4.8.

4.2 Contextual Background

The GCC is a regional intergovernmental political and economic union consisting of six countries: Bahrain, Kuwait, Oman, Qatar, Saudi Arabia, and the UAE. It was established in 1981, and is characterised by its countries predominantly being net oil exporters and therefore vulnerable to oil shocks, as well as being located in an area of frequent political turmoil which affects its equity markets (Siriopoulos et al. 2021).

The stock markets of GCC countries are small and relatively new, and are different from emerging and developed markets, due to being segmented both from each other and globally (Siriopoulos et al. 2021), accounting for only about 0.8% of global capitalisation (Maghyereh et al. 2017). Many transmission channels exist between the GCC countries and the US, including significant trade flows. For example, the US consumes 25 per cent of the oil produced in the GCC, and the GCC governments invest huge amounts in US companies and assets, particularly weaponry, logistics and military training. These ties mean that the majority of GCC block peg their currencies to the US dollar (Awartani et al. 2013). In addition, GCC countries secure against fluctuations in the dollar by holding gold reserves, of which they have about 1.3% of global reserves (Maghyereh et al. 2017). As in Looney (2020), GCC countries held 30.56% of the world's verified oil reserves and produced 22.76% of its crude oil, with the oil production and reserves in Saudi Arabia being the largest market in the GCC and the second biggest in the world, accounting for about 12.5% of the world's crude oil. The other gulf countries produce roughly 4.1% (UAE), 3% (Kuwait), 2% (Qatar), 1.1%

(Oman) and 0.06% (Bahrain).

The limitations of the GCC stock markets result from structural and regulatory disadvantages which include comparatively small numbers of listed companies, limited diversification in the sector, significant institutional holdings and constant inadequacies in financial and banking systems. In more recent years, a variety of changes have been made to supervisory, regulatory and legal systems in order to improve market transparency. Crucially, GCC markets have begun to open their operations to investors internationally, as well as to improve their liquidity (Awartani et al. 2013). However, their segmentation and lack of integration with global markets does provide an excellent opportunity for portfolio diversification and hedging, both regionally and globally, made even more attractive to investors by the region's high growth levels immediately before the financial crisis (Awartani et al. 2013).

Corporations which have a direct relation to oil also have a large proportion of the total capitalisation of the stock indices in GCC markets, with the exception of the Kuwait stock exchange (Maghyereh et al. 2017). Table 4.1 shows stock market indicators from GCC countries, along with selected macroeconomic data. Bahrain's market is the least liquid, and also the smallest. In contrast, the Saudi market, which is the oldest, is clearly the leader in the block. It is the most liquid, and is ranked as one of the emerging global stock market's top ten, as well as having the highest turnover ratio. Saudi Arabia is also the largest capital market, responsible for 80% of the block's total traded capital. The UAE stock exchange and the Qatar financial market, the next two largest, constitute 15.1%. Kuwait's stock exchange makes up the next 3.49%, and Oman and Bahrain contribute less than 2% of the total. It is evident that the oil has a significant influence on the economic activity and therefore the revenues of GCC countries, ranging from just over 2% of GDP in Bahrain to 42% in Kuwait. It is also clear that trade balance is where a large part of this contribution comes from, with fuel exports providing over 67% of total merchandise value in Saudi Arabia and over 92% in Kuwait. Furthermore, fluctuations in oil price numerically affect production costs and then in turn impact companies' cash flows and profits, which is then reflected in stock prices. This leads to uncertainty in all economic sectors which has a negative effect on companies' investments (Maghyereh et al. 2017). It can therefore be expected that
oil prices have a significant effect on GCC stock market movements (Maghyereh & Al-Kandari 2007).

Indicators	Bahrain	Kuwait	Oman	Qatar	S. Arabia	UAE
Stock Market Indicators						
Start of trading	1987	1952	1989	1997	1935	1988
Start of electronic trading	1989	1995	1989	2002	1988	2000
Number of listed companies	42	163	111	47	207	130
Market capitalisation of listed						
domestic companies	04.0	100	10 5		0400.4	004.0
(CUIT. US\$ DD)	24.6	106	16.5	165.4	2429.1	294.8
domestic companies						
(% of GDP)	70.9	100	22.3	114.5	347	82.2
Stocks traded, total					• • •	•
value (% of GDP)	1.7	12.4	1.4	16	69.3	10.5
Stocks traded, turnover ratio						
of domestic shares (%)	2.4	12.4	6.3	14	20	12.7
Foreign investment through						
mutual funds	Y	Y	Y	Y	Y	Y
Macroeconomic Indicators						
GDP (curr. US \$ bn)	34.73	105.96	73.97	144.41	700.12	358.87
GDP p. capita (curr. US \$)	20,410	24,812	14,485	50,124	20,110	36,285
Oil revenue (% of GDP)	2.15	42.14	24.88	16.91	24.24	16.20
Fuel exports (% of						
merchandise exports)	NA	92.88	NA	81.81	67.62	71.45
Foreign direct investment,	0.00	0.50	0.07	1 00	0 77	F F 4
net inflows (% of GDP)	2.90	-0.59	3.87	-1.69	0.77	5.54
roreign arect investment,	-0 50	2.08	0 72	1 80	0 70	5 28
net outliows (/o of GDT)	-0.55	2.50	0.72	1.09	0.70	5.20

Table 4.1: Stock market characteristics and macroeconomic data for GCC in 2020 (with the exception of oil revenue, the latest data of which is from 2019). Data from various sources, including World Bank and The Global Economy. Structure of the table replicated from (Maghyereh et al. 2017).

Indeed, conducting research in GCC countries is interesting, according to (Arouri et al. 2011). The first reason is that, as major oil exporters, their stock markets are vulnerable to fluctuations in oil prices and also to reactions to these fluctuations. The second is that these countries provide a favourable area for diversification of portfolios. Therefore, it is essential that empirical studies are focussed on GCC countries so that investors

may make thoughtful investment decisions. Equally, these studies will be important for policy makers who seek to regulate the stock market more successfully.

In Section 4.3 a summary of the relevant literature is provided.

4.3 Literature Review

Constrained portfolio optimisation is part of portfolio theory that is applied to assess optimal investor capital proportions for each asset, to result in maximal return for minimal risk. The application consists of various constraints, for example cardinality, bounds upon the proportions and asset class constraints (Chang et al. (2000), Fernández & Gómez (2007)).

The original model due to Markowitz suffers from a major drawback as it does not allow for real-world considerations, such as transaction costs, cardinality constraint, management costs and short sales. In addition, it ignores the preferences of investors, does not account for tax, and assumes perfect markets (Merton (1980), DeMiguel et al. (2009)).

Many methods have been used to understand portfolio theory. One of the tools is the minimum-variance portfolio, which uses as its input parameter the variance-covariance matrix, as mentioned in (Merton (1980), Nelson (1992), DeMiguel et al. (2009), Behr et al. (2013)). From the financial point of view, Unlike expected returns, the variance-covariance estimates are stable in time, meaning they can be forecast reliably. Behr et al. (2013) use weighted constraints to improve the variance-covariance optimisation method. As well as achieving a lower turnover and more stable Sharpe ratio, this method has a substantial impact on reducing risk.

Statman & Clark (2013) identified two gaps between investors' preferred portfolios and optimised M-V portfolios: inaccurate estimates and investor preferences outside of high expected returns and low risk. They proposed a mean-variance "efficient range" - an alternative to the efficient frontier - defined as portfolio locations which accept both the inaccurate estimates and investor preferences. They call for investor judgement in deciding sensible ranges and boundaries, claiming that locating the exact M-V efficient frontier is a "pretension".

Portfolio optimisation usually seeks the basic requirement of prediction of portfolio growth, for investors in financial market. The receding horizon control, or model predictive control, uses handle hard constraints to optimise portfolios, by applying these on the inputs. Advance knowledge of statistical return distribution is not required for this method. In financial applications, the major advantage of this method is that observing actual market behaviour informs each decision stage, allowing fresh estimates to be drawn (Merton (1980), DeMiguel et al. (2009)).

Non-single time horizons can also be applied to the construction of portfolios, using probability or non-uniform distributions over different horizons to produce a horizonwise estimation of risk. This allows for the fact that when transaction costs are not zero, agents do not trade continuously. Instead, they are offered a trading horizon when purchasing shares, which remains unknown to them and to econometricians. This assumption makes sure that investors concerns are focused on wealth distribution at various future dates when liquidation and realising of returns can occur (Ben Mabrouk 2020).

In order to understand commodity usage in portfolios, it is necessary to identify their specific attributes as compared with other concepts. The first distinction is that commodities are physical goods, such as precious metals, oil and so on. Commodities may be invested in via mutual funds, with futures, or directly. They also have growth potential, and are used to hedge against inflation (Zaremba (2015), Fabozzi et al. (2008)). Using commodities in portfolio diversification is not new, although commodities can be seen in a negative light. Despite the fact that commodities are unable to yield market beta that increases with time, they do have useful characteristics: for example, inhomogeneity. Intra-correlation between commodities can be low, for example in the case of crude oil which is the major component of our present study portfolio. Difficulty is found in buying, selling and storing commodities for investment, due to their relation to real-world conditions such as supply and demand. However, commodities represent one of the simplest methods of portfolio diversification (Jensen et al. (2000), Satyanarayan & Varangis (1996)).

Rubbaniy et al. (2021) analyse the fear index for COVID-19 using wavelet coherence,

4.3. LITERATURE REVIEW

along with an investigation of soft commodity safe-haven properties (Jan 2020-Apr 2021) using futures prices and soft commodities' spot. They find that safe-haven behaviour is evident in each of the soft commodities sampled, and also that each commodity's nature affects its safe-haven properties. Using M-V portfolio analysis, they show that portfolios including commodity futures are subject to lower risk than stockonly portfolios, concluding that during the COVID-19 period certain soft commodities could be suitable as safe-haven features. They observe that the nature of these soft commodities makes each suitable for different lengths of investment, so that some may be used as safe-havens for the short-term, but others may only be effective over the long-term. Furthermore, they observe that this behaviour is also affected by whether the investment is in futures or spot markets.

Wang et al. (2021) recently showed the effect of energy commodity futures on portfolio diversification strategy, concluding that their method produces a return increase simultaneously to a portfolio volatility decrease. They were also able to show that their diversification resulted in higher returns for portfolios, compared with traditional ones, without the need for increased risk. The authors also indicated that the role of energy futures in portfolio optimisation deserves further study. This motivates our study here, which proposes oil as an important diversification tool in portfolio optimisation, alongside other resources such as precious metals.

It is known that oil, and petroleum in general, are significant parts of financial transactions in GCC countries. But these commodities have shown high price volatility and uncertain behaviours in recent decades, due to various social, political and pandemic situations (Ben Mabrouk (2020), Su et al. (2021), Umar et al. (2021)).

Gorton & Rouwenhorst (2006) found that a better efficient frontier can be produced by the inclusion of commodities, due to the fact that commodity futures and stocks tend to be negatively correlated. Therefore, if more commodities are included, a better risk-return trade-off can be achieved.

Commodities can be considered portfolio hedges because of daily market chaos such as investing, buying and selling. Inflation can be seen as a major disadvantage, which requires hedging. Since commodity prices tend to rise with inflation, they can offer protection from its effects (Abanomey & Mathur (1999), Fabozzi et al. (2008)).

One of the main concepts in portfolio optimisation theory as well as practice is cardinality. In simple terms, it describes the selection of assets in limited numbers, so that existing asset value maximises the target function, while remaining below the cardinal constraint (Chan & Kroese 2010). Researchers have proposed the addition of upper and lower limits to the variables of the Markowitz model, in the form of the cardinality constrained mean-variance model (Fernández & Gómez 2007). Chang et al. (2000) determined the efficient frontier in the presence of the cardinal constraint by applying genetic, tabu search, and simulated annealing algorithms. Methods like these have demonstrated efficient frontier displacement. The major cause of the displacement was shown to be the input parameters added into the model along with retail algorithms as encoded constraints.

Several studies, such as Fortenbery & Hauser (1990), Satyanarayan & Varangis (1996) and Abanomey & Mathur (1999) have concluded that an upward shift of the efficient frontier may occur when commodity futures are included as part of the investment universe. Jensen et al. (2000) have examined the performance of commodity futures to assess their general attractiveness as a stand-alone investment, and as a portfolio component. Futures were shown to be a relatively poor stand-alone investment, since they have both lower returns and higher standard deviation than stocks. In a portfolio context, however, return/risk optimisation (over a range of risk levels) gave substantial weight to commodity futures over the full sample. Such allocation significantly enhanced the portfolios' returns. Fabozzi et al. (2008) is also a good reference for readers interested in commodities application and effect on the efficient frontier shifting.

Wang et al. (2011) relaxed the limitations of various previous studies - the martingale approach, an assumed complete market and particular asset return probability distribution - to develop a calculation process to derive the dynamic portfolio frontier and its resultant dynamic asset allocation. Drawing the dynamic and static portfolio frontiers on the same graph, they were able to show that the former is more efficient.

Scherer & He (2008) have shown that commodity inclusion in portfolios can induce an efficient frontier shift and improve the portfolio. They demonstrate that investors are

better able to build optimal portfolios using the tangency portfolio and risk-free assets. The next section provides the sources of data and the types of data under which portfolio optimisation will be conducted. Moreover, information about oil, precious metal prices and GCC stock markets is introduced.

4.4 Data Description

The historical data of daily stock closing prices was obtained from the Thomson Datastream service via the Plymouth Business School and the logarithmic return was used to calculate the rate of return on assets. The arithmetic mean of returns is used, which is identical to what is used in the standard Markowitz model. It is calculated by totaling the return values in a series of numbers and then dividing by the number of terms in the series. It should be mentioned that stocks with missing values at the beginning of the period of the study were dropped from the dataset. The dataset consists of daily return observations of stock prices of 496 Gulf companies from the GCC economies that are Saudi Arabia, Qatar, the United Arab Emirates, Oman, Bahrain and Kuwait. In addition, oil prices and the prices of three precious metals (gold, silver and copper) are included. Applying an Augmented Dickey Fuller (ADF) test to the returns of each asset individually for the total number of observations in the dataset resulted in uniform stationarity of returns (with and without trends considered). Overall statistics fell below those for individual years - this difference in size of test statistic is due to the sample size (Cheung & Lai 1995) and to the characteristics of the dataset. While the data is suitable for the application of the non-parametric methods of this chapter, the reader is referred to Section 5.9.3 for some issues that may affect the implementation of the methods therein.

The data covers the period from January 2008 to January 2019 (2894 observations), which is an eventful period: beginning in the year of the global financial crisis, and covering the oil crisis when prices saw a significant drop. However, the price of gold was rising at that time and this provided a balance to the negative impact of the crisis, particularly in GCC countries where economic growth was actually significant (Maghyereh et al. 2017). The study period continues through gold's price peak out 2010 and the fall that followed, and onwards through 2014, the beginning of the oil crisis. These two

periods have impacted GCC oil exporters negatively, along with central bank reserves, adding to the already turbulent political climate of the Middle East and creating added insecurity for investment (Maghyereh et al. 2017). As a result, the subject at hand is an interesting case to investigate.

We use this data because the GCC markets are developing and have a significant growth economy this may help investors to locate profitable investment opportunities. In addition, GCC countries are oil producers which affects on their economies. Moreover, there is a lack of literature in constructing optimal portfolios in GCC markets.

Moreover, it can be said that nominal returns are those which give the nominal return value with no adjustment for inflation. In comparison, real returns are those which are adjusted for inflation. These give a comparative value with past return values (Martin 2021). Nominal values or return values (those adjusted for inflation) can each be used to measure returns, though there are different arguments for each. The main case for the use of real return values (inflation-adjusted) is that they provide a real-world perspective, which has the benefit of allowing returns to be analysed and contrasted with past performance (Jones & Wilson 2006). It is argued that this makes it simpler to comprehend if stocks' and commodities' real return values have grown over a given period, since some apparent growth is merely attributable to inflation. When returns are inflation-adjusted, this can be identified and an asset's actual performance can be assessed.

Nominal values, on the other hand, reflect the daily prices and values we encounter (Hooks 1989). People making purchases have to deal with what the cost is today, and are looking to the future prices of the asset (Jones & Wilson 2006). Those in favour of nominal value of returns therefore argue that the they give the clearest picture, since investors are not necessarily influenced by an asset's value in the past. Which method is most effective can be concluded on the basis of changes in autocorrelation (the relationship between an asset's past and present value) between each of the two values. Strong autocorrelation is indicated by substantial differences between past and present value, whereas insignificant differences would indicate weak autocorrelation (Martin 2021).

Research suggests that the use of nominal values tend to result in strong autocorrelation between stocks in the long term, indicated by high price fluctuation as a result of which autocorrelation rises (Fama & French 1988). If not adjusted for inflation, returns are likely to rise in autocorrelation since the inflation rate affects the difference in values (Fama & French 1988). It can be concluded therefore that autocorrelation cannot be reduced by real returns, but is rather increased. However, inflation-adjusted nominal returns will likely produce weak autocorrelation, which causes correlation to rise across categories of assets since each is divided by the same rate of inflation (Schmidt & Curry 2022).

In conclusion, it seems that practical reasons make real (inflation-adjusted) returns the better choice, since GCC countries use multiple currencies and it would be impractical to adjust every one of these for inflation. So nominal returns are the better choice.

In this study we divided the sample into 13 classes: 11 classes comprising stocks of firms, a class with solely oil, and a final class with 3 precious metals. The number and type of assets in each class is given in Table 4.2. The grouping was determined by the Global Industry Classification Standard (GICS) which is a standardised classification system for equities that used by investors and established indices (or investment community). The GICS was developed by Standard & Poor's and Morgan Stanley Capital International in 1999 to provide a global standard for classifying firms into sectors and industries. The GICS structure contains 11 sectors, 24 industry groups, 68 industries and 157 sub-industries. All Companies are categorised regarding their principal business activity (similar operating characteristics). The sectors are: Energy, Materials, Industrials, Consumer Discretionary, Consumer Staples, Health Care, Financials, Information Technology, Communication Services, Utilities and Real Estate.

We use the Global Industry Classification Standard because some sectors are procyclical (e.g., the financial sector) while others are counter-cyclical or do not depend on the business cycle (e.g., pharmaceuticals) (Basu 1996). Pro-cyclical stocks refer to a positive correlation between any stock's price and the overall state of the economy. In other words, any economic quantity that tends to move in the same direction as the economy, increasing in expansion and declining in a recession. In contrast, any stock's

4.4. DATA DESCRIPTION

Class	Type of Asset	Number of Assets
1	Energy	20

Materials

Industrials

Consumer Discretionary Services

Consumer Staples

Healthcare

Financial

Information Technology

Communication Services

Utilities

Real Estate

Oil

Precious Metals

price are negatively correlated with the overall state of the economy are said to be counter-cyclical stocks.

Table 4.2: Number of assets in each class.

Figure 4.1 shows the time series plots for the commodities of oil, gold, silver and copper and the respective returns series, over the period 2008-19. The primary axes of the plot for each commodity correspond to prices and returns. Oil prices show a dramatic upward trend throughout 2008, reaching a peak towards the end of the year before dropping sharply in 2009 as the effects of the financial crisis take hold (Bjørnland 2009). The returns show high volatility during that period, and also later in 2015-16. Gold and silver prices increase slightly at the beginning of the timeframe, reaching initial peaks in 2011, followed by periods of significant fluctuation before levelling out at a much lower price by 2013-14. Like oil, these commodities also demonstrate high volatility in 2009. Copper shows a similar trend to oil, with high prices pre-2009, then a dramatic drop, but then a relatively quick recovery leading to a new peak in 2011. Prices then decrease gradually until the end of the study period.

Figure 4.2 shows the stock indices and respective returns of six GCC countries over the same period. In all the plots we observe volatility in returns during the first year of the study. Immediately post-2009 a tranquil period is observed, broken by the oil price decline in 2015 which leads to another period of high volatility across the GCC stock exchanges, coinciding with oil returns , demonstrating strong asymmetry that showed across GCC exchanges, albeit in varying degrees. Gold price momentum is preserved during the financial crisis, however, in agreement with the theory of gold's safe-haven potential in periods of turmoil.





(a) Crude Oil





(c) Silver



Figure 4.1: Time series plots of prices for oil, gold, silver and copper commodities (black lines) and the respective returns series (blue lines).



Figure 4.2: Time series plots of prices for the GCC stock indices (black lines) and the respective returns series (blue lines).

The descriptive statistics for the returns of GCC stock market indexes and for the returns of oil, gold, silver and copper from 1st January 2008 to 31st January 2019 are presented in Tables 4.3 and 4.4, respectively. Each table reports the mean return, standard deviation, skewness and kurtosis, along with minimum and maximum returns. The various stock market indexes (Table 4.3) show that UAE, Saudi Arabia and Qatar have the highest standard deviation, at 1.516%, 1.327% and 1.235%, respectively. The highest mean returns across the period are also seen in Qatar (0.033%), UAE (0.014%) and Saudi Arabia (0.001%), though the maximum returns demonstrate a different trend, with UAE far exceeding the other 2 markets at 48.912% compared with just 9.513% for Saudi Arabia and 8.696% for Qatar. However, the minimum returns in these three markets follow the same trend, being more than three times higher for UAE than for Saudi Arabia and Qatar. Across the GCC markets, UAE is the only example of positive skewness, and its kurtosis levels are more than 20 times that of the other markets. Moreover, a Jarque-Bera test was conducted for checking normality. According to (Table 4.3), it is notable that there is no normality in return data in GCC countries where the p-value is 0.00001 < 0.05 which means rejecting the null hypothesis. Generally, the values of the Jarque-Bera test are very high, which means that the distributions are very far from the normal distribution for all GCC countries.

S	Saudi Arabia	Bahrain	Kuwait	Oman	Qatar	UAE
Mean	0.001	-0.022	-0.020	-0.022	0.033	0.014
Maximum	9.513	2.792	6.517	8.370	8.696	48.911
Minimum	-9.812	-4.800	-6.718	-8.331	-8.807	-30.574
Std. dev	1.326	0.529	0.996	0.987	1.234	1.516
Skewness	-0.517	-0.801	-0.478	-0.816	-0.438	8.689
Kurtosis	12.083	7.111	8.346	19.799	11.103	432.408
Jarque-Bera test	17678	6387	8482	47439	14911	22512560

Table 4.3: Descriptive statistics, from 1st January 2008 to 31st January 2018 for the returns of the GCC stock market indices.

Across the assets (Table 4.4), the standard deviations show that oil is riskier than any of the other assets (2.416%), though its minimum return is very similar to that of silver (-12.038% and -12.187%, respectively). Oil's maximum return, however, is

considerably higher than that of the other assets (17.837%, compared with 14.642% for silver, 12.441% for copper and 7.106% for gold). The skewness for oil and copper across the period is positive, while gold and silver show negative skewness, which suggests that gold and silver are more likely to experience losses rather than gains. With regards to the kurtosis, all the returns show kurtosis values of above three, which would be normal distribution, indicating extreme movements. According to the Jarque-Bera test, it is notable that there is no normality in return data for oil, gold, silver and cooper for the p-value 0.00001 < 0.05 which means rejecting the null hypothesis "there is normality".

Oil	Gold	Silver	Copper
0.009	0.0224	0.0228	0.011
17.837	7.106	14.642	12.441
-12.038	-9.663	-12.187	-9.806
2.416	1.144	1.986	1.693
0.342	-0.293	-0.107	0.070
5.314	6.140	4.864	4.307
3449	4573	2848	2231
	Oil 0.009 17.837 -12.038 2.416 0.342 5.314 3449	OilGold0.0090.022417.8377.106-12.038-9.6632.4161.1440.342-0.2935.3146.14034494573	OilGoldSilver0.0090.02240.022817.8377.10614.642-12.038-9.663-12.1872.4161.1441.9860.342-0.293-0.1075.3146.1404.864344945732848

Table 4.4: Descriptive statistics, from 1st January 2008 to 31st January 2018 for the returns of oil, gold, silver and copper.

4.5 Approach and Algorithm Setup

This chapter considers the basic mean variance portfolio optimisation problem with class constraints for diversifying optimal portfolios, which were constructed to contain Gulf market stocks, oil and at least one precious metal. The use of class constraints was first established by Chang et al. (2000) and is used to constrain the total percentage invested in assets with shared features, which results in a safer, more diversified portfolio. To construct the optimal class-based portfolios, an approach that has been discussed earlier in Chapter 3 has been adopted. This research examines the GA on the extended model of mean variance by focusing on the class constraint to identify efficient frontier which offers the best trade-off risk against the return. As referred to previously, a genetic algorithm is an optimisation method which mimics natural genetics. The algorithm relies on the three genetic operators of crossover, mutation and

selection of the fittest (see Section 3.2.2 for details). Also, as specified previously, the algorithm is employed because it provides high-quality solutions (although possibly sub-optimal) in a reasonable amount of time.

We generate unconstrained portfolios (the total set of which is named the unconstrained efficient frontier) and constrained portfolios with class constraints in some of the combined portfolios, in particular precious metals and oil. We compare the unconstrained and class constrained efficient frontiers. The resulting optimal portfolios, being designed for GCC investors, are partitioned by class constraints.

Class constraint is referred to in the work of Ruiz-Torrubiano & Suárez (2010), and applied in three studies by Anagnostopoulos & Mamanis (2010), Anagnostopoulos & Mamanis (2011*a*) and Anagnostopoulos & Mamanis (2011*b*) and also applied in studies by Chang et al. (2000) and Jin et al. (2016). It can also be formed by dividing assets into subsets with shared characteristics, then performing optimisation on each class's best representative (Pai & Michel 2009).

Jin et al. (2016) tested class constraints with the quantity and cardinality constrained Markowitz problem. The authors state that, because the class constraint is linear, then this does not increase the difficulty of the unconstrained quadratic problem. It was also expected that difficulty would be reduced by the amount of classes, since the initial problem was divided into subproblems with reduced dimension. As expected, fewer classes result in the situation needing far more computational effort. Conversely, more classes reduce the computational effort but also have the side effect of limiting the search. This trade-off may be subject to the decision maker's preferences.

Also as predicted, not a lot of difficulty with regards to computational cost is generated by class constraint on a cardinality constrained model, although optimality may be affected. Jin et al. (2016) report that class constraints do not result in significant differences in the cost of solving the quantity and cardinality constrained problem. However, as mentioned previously, this chapter examines the application of a GA on the problem for the practical interest of solutions realisable via a GA. The adopted approach for this research is similar to that detailed in Section 3.4.1, except that the algorithm has been adjusted in order to include class constraints rather than cardinality constraints or a minimum proportion constraint.

With respect to the experiments, all comparisons are with the UEF (i.e., the QPgenerated solution), which was computed on the GA of Section 3.3 using 2000 points and settings of 20000 generations (with a 'runs' parameter of 5000). The GA population parameters used are identical to those of Section 3.4.1. That is, the GA operators used were linear ranking selection, random position mutation and Laplace crossover, with crossover probability 0.5, and mutation probability 0.4. The population size was 50.

In order to measure GA performance, IGD (Inverted Generational Distance (Chen et al. 2012) is used. It measures the average minimum distance from the computed GA solution to each point in the unconstrained efficient frontier (UEF). The notation U_{IGD} refers to the IGD between the computed GA solution and the unconstrained solution, with ND_{IGD} used to refer to the sifted (non-dominated) versions of those comparators used for the computation of GA solution. The term ND_{IGD} defines the *sifted EF* (dominance) to be the set of all non-dominated points. For more details, the reader refers to Section 3.4.1. On each experiment in this chapter, sixty-four cores of the HPC cluster at the University of Plymouth were used.

It is worth mentioning that in this chapter, the cardinality constraints are not applied except for in Section 4.6.3. Thus the CCEF is unknown and we only know the UEF, so can only measure the proximity of the solution found from that. Thus any IGD measurements are only given with respect to the UEF.

The problem formulation, constraints and the constraint-handling procedure is given in the next subsection, with initial GA settings subsequently.

4.5.1 Formulation and Constraint Handling

In this section, we first present the unconstrained portfolio optimisation problem as follows:

Minimise

$$\sum_{i=1}^{n} \sum_{j=1}^{n} \sigma_{ij} x_i x_j.$$
 (4.1)

Subject to

$$\overline{R} = (\mathbf{x}, \mu) = \sum_{i=1}^{n} \mu_i x_i,$$
(4.2)

$$\overline{R} \ge R_0 \tag{4.3}$$

$$\sum_{i=1}^{n} x_{i} = 1$$

$$0 \le x_{i} \le 1, \qquad i = 1, 2, ..., n,$$
(4.4)

where x_i is the weight of asset *i*, μ_i is the expected return of asset number *i*, \overline{R} is the expected return of the portfolio and *r* is the variance (risk). Note that σ_{ij} is the covariance between the returns for assets *i* and *j*. Finding the set of all optimal solutions of the above problem (4.1)–(4.4) traces out the efficient frontier.

Second, we present class constraint as:

Let C_m , m = 1, ..., M, be M sets of assets that are considered to be mutually exclusive, i.e., $C_i \cap C_j = \emptyset$ for all $i \neq j$. Class constraints limit the proportion of the portfolio that can be invested in assets in each class. Let L_m be the lower weight limit and U_m be the upper weight limit for class m – thus the class constraints are written as in Chang et al. (2000) and Jin et al. (2016):

$$L_m \leq \sum_{i \in C_m} w_i \leq U_m, \ m = 1, \dots, M.$$

It is noticeable that we choose $L_m = 0$ and $U_m = 1$ for all *m*, since, as above, minimum proportions are not used in this chapter.

It should also be re-emphasised that the algorithm is metaheuristic (see Section 3.2.2) and so is naturally an approximation algorithm. Exact results, for example class 2 having at most five assets selected, are *aims* that the algorithm attempts to satisfy but are not guaranteed to be present in the final result. It should also be mentioned that the covariance matrix is a square matrix of dimension 500 and thus the evaluation of the risk measure (which requires matrix multiplication) tends to be computationally costly.

As we saw in Section 3.5, a penalty method is the most popular method in GAs to handle constraints. Penalties are used to replace constrained optimisation problems via a sequence of unconstrained problems. We use the penalty method since it proves to be computationally efficient. Ideally, the solution of these unconstrained problems converge to the same solution as the constrained problem. By adding a 'penalty function' term to the objective function, the constrained optimisation problem is transformed into an unconstrained one. Where constraints are violated the violation measure is nonzero, while it is zero for areas of non-violation. See, for example, this constrained optimisation problem:

Find min
$$f(x)$$
 (4.5)

such that
$$g_i(x) \le 0, i \in I = \{1, 2, \dots, n\},$$
 (4.6)

where $g_i(x)$ is any constraint. In practice, a penalty function is included in order to combine the constraints and the objective function into a form from which the unconstrained problems can be practically solved. In this case, in order to encourage the algorithm to use oil as much as possible in portfolios because GCC countries are oil-based markets and it is assumed that investors should invest heavily in oil then they should hedge their portfolio, a step penalty was applied:

$$F(x,\varepsilon) = f(x) + \varepsilon \sum_{i \in I} g_i(x), \qquad (4.7)$$

where ε is a penalty parameter. Various values of the above penalty are tested in Section 4.6.2.

The fitness function used is

$$f(\mathbf{x}) = -(\mathbf{x}^T \Sigma \mathbf{x} + M_1 (\mathbf{x} \cdot \boldsymbol{\mu})^2), \qquad (4.8)$$

where the value of M_1 is 100000 in the current chapter. The penalty applied was

$$f_p = \frac{\sum_{i=1}^{M} M_2(\#\{w_i \in C_i : w_i > \text{tol}\} - K_i)^2}{n \cdot n_{\text{class}}},$$
(4.9)

where *M* is the number of classes (13 - see Table 4.2), K_i is the number of assets desired in class C_i (so $\sum_i K_i = K$), tol is the tolerance, 10^{-9} , by which a weight is counted (due to arithmetic precision), n_{class} is the number of classes aimed for and M_2 is a constant (tuned to 100 by copious experimentation). If the summand for class i = 12 is zero (i.e., no oil selected), then a step penalty of 2000 is added to the sum f_p . The value of 2000 for the "lack of oil" step penalty was reached experimentally as follows.

4.5.2 Determining the Step Penalty

In order to decide the appropriate value of step penalty here are some tests that have been executed. We examine the GCC dataset with different values of step penalty, illustrating that oil composition for example, step penalty of 500, 2000 and 10000. The settings were: K = 20 and 5 classes, using one run at random. The experiments proceeded with number of generations 100n, run = 30n and the number of efficient frontier points as 50. Figures 4.3 and 4.4 show the effect of the step penalty of 500 and 10000 respectively. It has been found that the number of points (each point represents a feasible portfolio) which include oil is 10 out 50 points for a step penalty of 500, and 13 out of 50 points for a step penalty of 10000. However, the highest number of portfolios that include oil was found to be 16 for a step penalty of 2000 (Figure 4.5). It is noted that the number of oil tends to be smaller when we proceed experiments with step penalties more or less than step penalty of 2000. Also, as we increase the penalty past that point, the computed EF tends to move away from the UEF. Therefore, it was decided to execute with step penalty of 2000 for all future tests in this chapter.



Figure 4.3: Output with a step penalty of 500 for GCC dataset, oil composition with 5 classes and K = 20 using one run at random.



Figure 4.4: Output with a step penalty of 10000 for GCC dataset, oil composition with 5 classes and K = 20 using one run at random.

The values of the class constraints are set according to the procedure below.

4.5.3 Initial Configurations

Here, an initial configuration (IC), $\mathbf{c} = (c_1, c_2, \dots, c_{13})$, is an ordered vector consisting of thirteen class limits (due to the set up of the data into M = 13 classes) of which the sum is equal to the cardinality constraint: $\sum_{i=1}^{13} c_i = K$. Each entry, c_i , of \mathbf{c} gives the subcardinality constraint on class *i* in effect throughout the run of the GA. For example, the



Figure 4.5: Output with a step penalty of 2000 for GCC dataset, oil composition with 5 classes and K = 20 using one run at random.

initial configuration

$$\mathbf{c} = (0, 0, 2, 2, 0, 3, 0, 2, 4, 3, 0, 3, 1) \tag{4.10}$$

means that the GA should include at most one asset from class 13 (oil), at most two assets from classes 3, 4 and 8, at most three assets from class 6, 10 and 12, and finally a maximum of four assets from class 9 in the portfolio. The GA includes a facility to generate random initial configurations which sum to a given cardinality *K*, such that $1 \le c_{12} \le 3$ (i.e., precious metals) and $c_{13} \le 1$ (i.e., oil). This ensures that at most three precious metals (class 12) and oil (class 13) are chosen. There is control over how many classes are populated in the initial configuration. The example initial configuration (4.10) was created using this method. Note, as previously stated, that each component of the initial configuration may be considered as an 'aim' and so there is a non-zero probability that no oil or precious metals will be in the portfolio.

From observations of algorithm runs, it is noted that the numbers of assets assigned to each class by the GA are not typically equal to c_i (which is the number of assets desired in each class according to the initial configuration), and thus do not typically sum to *K*. This is not unusual when using an approximation method such as a GA (Chang et al. 2000). There is also the question of counting the number of assets assigned in each class, with an asset included when its proportion reaches some threshold.

This is important with the current data representation - continuous data - since the question of what exactly equals zero is much more involved than it would be with an integer representation, for example. In this study, the threshold of 10^{-9} was chosen. In practice, there seemed to be a very small difference in numbers of assets assigned to given classes when this threshold was changed.

The next section presents our experimental results.

4.6 Results

This section discusses the empirical results of our research. First of all, Section 4.6.1 presents varied initial configurations after each run. Then, an indication arising from fixed initial configuration is discussed in Section 4.6.2 to check robustness of our approach. Next, portfolio compositions when the number of classes in the initial configuration varies is investigated in Section 4.6.2. Finally, the current approach with the solely cardinality-constrained approach is compared in Section 4.6.3.

4.6.1 Random Initial Configurations

In this section, we propose a simulation which uses the case of random initial configuration. Experiments for the number of GA generations to run were conducted, giving a limit of 100n (i.e., 50,000 generations). The 'runs' parameter was set to 30n (i.e., 15000 generations). Due to the size of the matrices involved (dimension n = 500) and implementation of the fitness function (which has to compute a large number of operations), the above parameters were chosen to offer a good compromise between accuracy and time.

For each value of *K*, thirty experiments were run (with each using a distinct random IC). The results of implementing the GA in the GCC dataset, with oil and precious metals, are presented in Table 4.5. The number of classes used in the IC is also set at five out of the possible thirteen (including classes 12 - oil - and 13 - precious metals), with the 3 stock classes selected at random each time the algorithm was run. The values of the measures are written in the form of mean followed by standard deviation (e.g., 3.508e-5/1.111e-5 refers to mean 3.508e-5 and standard deviation 1.111e-5). The average computational time of all the *K* values are relatively consistent, ranging from 80365.91

to 84426.52 seconds, and K = 50 is the fastest. The lowest mean errors in U_{IGD} and ND_{IGD}, are found in K = 20, which also has the least standard deviation. K = 50 has the largest standard deviation, followed by K = 30 and K = 40, a pattern which generally also applies to the mean U_{IGD} and ND_{IGD} errors for each of these K values. Note that, for this chapter the ND_{IGD} measurement refers to the IGD between the sifted GA solution and UEF co-ords (the latter of which are all non-dominated, by definition, and so do not need to be sifted). The plotted results of the various K values in the table can be seen in Figure 4.8.

Classes = 5, thirty runs

K	U _{IGD}	ND _{IGD}	Time
20	3.508e-5/1.111e-5	6.037e-5/3.392e-5	84426.52/7678.74
30	3.630e-5/1.450e-5	9.060e-5/6.447e-5	83114.68/7906.15
40	3.783e-5/1.326e-5	9.545e-5/6.026e-5	82413.36/5936.60
50	3.598e-5/1.501e-5	1.195e-4/6.568e-5	80365.91/7401.84

Table 4.5: Statistics produced by running with the values of K shown.

Table 4.6, however, shows the crowding distance measurement (CD) and number of points that include oil and/or metals for the GCC dataset at different *K* values. To recap, the crowding distance is the average distance of two neighbours of each point along each of the objectives (Raquel & Naval Jr 2005). The mean crowding distance measures the density of solutions surrounding a particular solution. The trend is that a higher *K* value results in a lower mean CD. The standard deviation in CD is very different, however, with K = 20 giving the lowest standard deviation CD, followed by K = 50, then K = 30 and K = 40. This time, the greatest number of points that include oil and metals is seen when K = 20, and this number decreases as the *K* value increases.

K	CD	Oil/Metals
20	1.129e-4/1.620e-5	36.13/2.57
30	9.974e-5/2.067e-5	34.30/3.58
40	9.624e-5/2.236e-5	33.73/3.05
50	8.696e-5/1.751e-5	33.67/2.93

Table 4.6: CD measurement for GCC dataset and number of points that include oil and metals with the values of *K* shown.

Now, in order to produce efficient frontiers using random initial configurations we investigate the return-risk relationship when the number of assets is 20 and 50 respectively, which is graphically depicted in Figures 4.6 and 4.7.

First, Figure 4.6, produced by running the code 30 times using 30 random selections of IC and plotting typical results for K = 20, demonstrates the effect that the IC has on the results. Figure 4.6 (a) shows the relationship between the GA solution (green dots) and the quadratic UEF (red line) at K = 20. While the green dots align with the UEF at the top of the curve, there is an area of considerable non-alignment between return values of 3e-04 and 6e-04, as well as less agreement at return values of under 3e-04. Figure 4.6 (b) shows that the highest number of assets used is about 41.

Looking at Figure 4.6 (c), it is clear that a lower number of classes, for example 3-6 per point, yield far larger returns, in exchange for increased risk. In addition, points with over 9 classes never reach above a 7e-04 return, and points between return values of 3e-04 and 6e-04 are generally furthest from the UEF. This behaviour is also seen

in Figures 4.6 (d), (e) and (f), where various combinations of oil and metal composition (with stock assets) were used to compare risk return ratios. Stock assets with oil included (Figure 4.6 (e)) are seen to achieve returns of 0 to 1e-03, along with good agreement with the UEF above the return of 8e-04. With random IC, we have portfolios with oil and metal (but mainly oil) accessing the top-right part of the frontier.



Figure 4.6: GA output for the GCC dataset, oil and precious metals with 5 classes and K = 20 using random ICs.

Figure 4.7 shows the effect of using 30 random ICs at K = 50. It is clear that, at lower return values, the results are in fairly close agreement. The results quite closely match the UEF as return values increase, and tend to agree less when return values are low.



Figure 4.7: GA output for the GCC dataset, oil and precious metals with 5 classes and K = 50 using random ICs.

Overall, it is evident that the top right of the EF curve includes oil and stocks assets only. The addition of metals mitigates the risk, and opens new investment opportunities that helps investors to diversify their portfolios.

Comparison Between Results Using Different Values of *K*

A comparison between results using different values of *K* when random ICs can be seen in Figure 4.8. It is evident that when K = 20 and K = 40, the change in IC has little effect on the general shape of the curve. For random ICs all four plots demonstrate a relatively consistent curve (Figure 4.8). Looking at the number of classes, the points with the highest class value (12) in Figure 4.8 (a), when using the random IC and K = 20, are grouped between return values of just under 2e-04 and 4e-04, so that with random ICs there are potential returns of 4e-04 for portfolios including more classes.

Overall, it can be seen that a lower number of classes tends to produce both increased return and increased risk. However, the IC affects these relations.



So far this section has focused on a varied initial configuration. The following section will discuss other approaches.

4.6.2 Robustness Exercise

In this section we perform a robustness exercise to check for the robustness in our approach. It is shown that the results do not change considerably when investigating fixed IC. However, with random IC we have more mixed portfolios on the top-right corner of the curve. Also, we examine portfolio compositions when the number of classes in the initial configuration is varied. Thus we will decide to proceed with the random IC approach and 5 classes for all remaining experiments in this chapter.

Potential Effects Arising from Fixed Initial Configuration

The objective of this section is to determine whether fixing initial configuration can affect the results and then present comparison between the findings obtained and results of Section 4.6.1.

For each experiment in this section, an initial configuration was produced by the algorithm at random, and thirty experiments were run with this configuration. This is because results of EA runs by themselves provide limited information.

Below are results for K = 20, 30, 40 and 50 with, for each fixed value of K, the same IC used for every run. It should be noted that K has to be at least equal to the number of classes. The results of running the detailed GA on the GCC dataset, again with oil and precious metals, are presented in Table 4.7. The number of classes used in the IC was set at five out of the possible thirteen (again including classes 12 - oil - and 13 - precious metals), with the number of assets from each stock class and the precious metals class being selected randomly each time. The results indicate that, for these experiments, a longer computational time does not necessarily lead to lower approximation error: for example, in the case of K = 20, which has the longest computational time, the mean results have relatively high U_{IGD} and ND_{IGD} errors, meaning that this K value yields a higher error than values with shorter computational time, such as K = 40, which yields a far lower error (comparative U_{IGD} errors are 3.214e-5 for K = 20 and 3.033e-5 for K = 40, and ND_{IGD} errors are 4.807e-5 and 4.692e-5, respectively). It can be said that

is because, as *K* increases, the number of assets available increases and so the EF converges to the UEF. We get different EFs when random ICs are used. However, in the case of K = 30, the shortest mean computational time (76683.82) does yield the results which are furthest from the UEF, with a mean U_{IGD} error of 4.872e-5 and ND_{IGD} error of 1.280e-4, compared with the closest results to the UEF which are that of K = 40 (U_{IGD} of 3.033e-5 and ND_{IGD} of 4.692e-5).

The table also demonstrates that standard deviation differs greatly at different *K* values. For example, in the case of K = 30, the standard deviation of the U_{IGD} is 2.489e-6 and of the ND_{IGD} is 5.904e-5, whereas those values for K = 40 are 1.475e-6 and 9.267e-6, respectively. The plotted results for various *K* values can be seen in Figure 4.11.

This table for fixed initial configuration (Table 4.7) shows quite different statistics from those of the previous table (Table 4.5). It seems that the standard deviations of U_{IGD} and ND_{IGD} are of an order of magnitude higher than those of Table 4.7. This may be expected as the initial configuration is different for each run in Table 4.5.

# Classes = 5, thirty runs					
K	IC	U _{IGD}	ND _{IGD}	Time	
20 30 40 50	$\begin{array}{c}(0_5,6,6,0_3,5,1,2)\\(11,0_4,9,0_4,8,1,1)\\(13,13,0_4,12,0_4,1,1)\\(0_2,16,0,16,0_5,16,1,1)\end{array}$	3.214e-5/1.545e-6 4.872e-5/2.489e-6 3.033e-5/1.475e-6 4.078e-5/2.963e-6	4.807e-5/7.653e-6 1.280e-4/5.904e-5 4.692e-5/9.267e-6 1.560e-4/1.472e-5	91987.14/8323.92 76683.82/5594.05 88878.68/7813.28 85557.58/8371.11	

Table 4.7: Statistics produced by running with the values of *K* shown. The ICs were produced at random, with 0_n denoting a sequence of *n* zeros.

Furthermore, this section aims to experimentally investigate the crowding distance measurements (CD) and number of points that include oil and metals for the GCC dataset at different *K* values as it is shown in Table 4.8. It is clear that the lowest mean and standard deviation crowding distances are seen in the case of K = 50 (7.212e-5 / 4.287e-6), with the mean oil and metal inclusion of 31.53. With K = 40 the mean CD is largest; however, standard deviation CD is largest when K = 30. The number of points including oil and metals is highest when K = 40 and lowest when K = 50.

K	IC	CD	Oil/Metals
20	$(0_5, 6, 6, 0_3, 5, 1, 2) (11, 0_4, 9, 0_4, 8, 1, 1) (13, 13, 0_4, 12, 0_4, 1, 1)$	1.208e-4/1.292e-5	33.50/1.943
30		1.009e-4/2.283e-5	34.27/1.76
40		1.211e-4/1.529e-5	38.10/2.76

Table 4.8: CD measurement for GCC dataset and number of points that include oil and metals with the values of *K* shown.

In order to produce efficient frontiers using a fixed initial configuration, we investigate the return-risk relationship when the number of assets is 20 and 50 respectively, which is graphically depicted in Figures 4.9 and 4.10 below.

In Figure 4.9, it is shown that the relationship between the GA solution (green dots) and the quadratic UEF (red line), at K = 20. While the green dots align with the UEF at the top of the curve, there is an area of considerable non-alignment between return values of 3e-04 and 6e-04, as well as less agreement at return values of under 3e-04. Figure 4.9 (b) shows that the largest number of assets used is about 45, although a considerable portion of the points are using around 20 assets or less.

Looking at Figure 4.9 (c), it is clear that a lower number of classes, for example 3-6 per point, yield far larger returns, in exchange for increased risk. In addition, points with over 9 classes never reach above a 6e-04 return, and again, points between return values of 3e-04 and 6e-04 are generally furthest from the UEF. This behaviour is also seen in Figures 4.9 (d), (e) and (f), where various combinations of oil and metal composition (with stock assets) were used to compare risk return ratios. It is evident that when oil and metal are combined with the stock assets (Figure 4.9 (d), the risk and return both decrease in comparison to using just the stock assets, except for at returns of between 5e-04 and 6e-04, where the risk is relatively lower than at, for example, returns of 4e-04. When stock assets are combined with oil only (Figure 4.9 (e)) the results are similar, generally showing that oil returns of 3e-04 to 5e-04 are subject to higher risk that returns of under 3e-04. In this case, there is no oil return above 5e-04 at all, and the higher return points consist solely of stock assets.

In Figure 4.9 (f) the metal only composition appears very similar to the oil/metal combination, offering again a lower risk at 5e-04 to 6e-04 than at slightly lower returns. Overall, then, it is clear that the top of the EF curve is always occupied by stock assets only, and adding oil and metals to these points creates new investment opportunities for portfolios diversification.

However, stock assets with oil only (Figure 4.6 (e)) were seen to achieve far higher returns than these with a fixed IC (Figure 4.9 (e)) though the agreement with the UEF is the same. However, the opposite is true for stock assets with oil and metal, where the EF agreement with the UEF does remain the same, but the maximum returns for oil/metal are decreased. The results more closely match the UEF as return values increase, and tend to agree less when return values are low. Note that this is in contrast to the effect found at K = 20, where the mean of 30 random ICs (Figure 4.6) caused the results to agree more closely with the UEF than this fixed IC does (Figure 4.9).



Figure 4.9: GA output for GCC dataset , oil and precious metals with 5 classes and K = 20 using the same IC.

Turning now to Figure 4.10, with a change to K = 50 (compared with K = 20 in Figure 4.9), it is clear that the points are far less in line with the UEF curve under returns of around 3e-04, meaning that even very small returns are subject to greatly increased risk. This is consistent across all three combinations of stock assets with oil, with metal and with both.

Figure 4.10 (b) illustrates that the greatest number of assets used is about 90. The value of *K* clearly helps to pull the number of assets down. To compare, in the case of Figure 4.9 (b) the value of *K* is lower, giving a better restriction and resulting in a lower number of assets being used, whereas in Figure 4.10 (b) the general number of assets is higher in line with the higher value of *K*. Also, it is clear that, at lower return values, the results with the random IC (Figure 4.7) were in slightly closer agreement than these with the fixed IC (Figure 4.10).

Over all, we can see that near the very top we do not have any oil, purely because investors can get better returns by just buying the top 1- or 2- best performing assets. However, lower down we tend to get more oil as when we go down the EF more classes of assets are used (11 or 12 classes).



Figure 4.10: GA output for GCC dataset , oil and precious metals with 5 classes and K = 50 using the same IC.
Comparison Between Results Using Different Values of *K*

Moving onto Figure 4.11, it is clear that the value of *K* affects point position when considering stock assets combined with oil and metals. While close agreement with the UEF curve is demonstrated at returns of above 8e-04 in all cases, with *K* values of 30 and 50 the risk increases at return values of 4e-04 or less (and dramatically so when return is lower than 1e-04), so that the risk at 1e-04 is approximately 0.00027 and 0.00019 at K = 30 and K = 50, respectively. Interestingly, although the general trend at K = 40 is not dissimilar to K = 20 at return values above 1e-04, the risk at the return value of 0e+00 (i.e., zero return) is very similar to that at K = 50, so that with K = 20 the value of risk is 0.00017, and with K = 50 the value of risk is 0.00019.

These plots show that the shape of the curve differs for *K* values with a fixed IC (Figure 4.11), unlike Figure 4.7 which showed that for K = 20 and K = 40, the change in IC had little effect on the general shape. However, now at a fixed IC the shapes of the curve when K = 30 and K = 50 are distinct from those when K = 20 and K = 40 (Figure 4.11). However, for random ICs shown previously, all four plots demonstrated a relatively consistent curve (Figure 4.8), which is similar to what is now seen when K = 20 and K = 40 at a fixed IC (Figures 4.11 (a), (b), (e) and (f)).

Looking at the number of classes, though, some further differences are evident. For example, the points with the highest class value (12) that we saw in Figure 4.8 (a), when using the random IC and K = 20, were grouped between return values of just under 2e-04 and 4e-04, whereas now, when using a fixed IC at the same *K* value (Figure 4.11 (a)) they are spread over a wider return value range of just over 0e+00 and 4e-04, indicating that with random ICs the potential return for portfolios including more classes is greater than with a fixed IC. The same is seen in a comparison of Figures 4.8 (c) and 4.11 (c), where K = 30. The comparison for K = 50 does not follow this trend, however, as can be seen in a comparison of Figures (4.8 (g)) and (4.11 (g)) where we previously saw that using random ICs decreased the potential return for points with the largest amount of classes. Overall, from the comparison of Figures 4.8 and 4.11, it can be seen that a lower number of classes tends to produce both increased return and increased risk. However, the IC affects these relations.



Figure 4.11: Plots for comparison between values of K using five classes and fixed IC.

Therefore, graphical results obtained from fixed ICs are similar to those obtained from random ICs when using this approach (but not identical). The main difference, for example, is that the top right corner of the curve contains more mixed portfolios when using random ICs. The next subsection shows the consequences of changing the number of classes used.

Portfolio Compositions when the Number of Classes in the Initial Configuration is Varied

This section presents results for experiments conducted on the effects of changing the number of classes used in the IC. We allow for the initial configuration to include assets from various number of classes. Note that in the case of 13 classes, the solution always include assets from all classes. Practically, this may help investors to increase their profits or decrease the risk when they include more classes in their portfolios which do not mean an increase in the number of assets in those portfolios. We fix a value of K = 20 as an example for this section.

An objective of the algorithm encouraged by the penalty function is to to have oil and at least one precious metal (i.e., classes 12 and 13 are used). Statistics for the mean and standard deviation distances between the UEF and the GA-produced EF for each number of classes are given (U_{IGD}), along with those for the sifted solution (recall, a sifted solution is a solution in which dominated points are removed) in comparison with the UEF and the computation time (Table 4.9). The IGD values for five classes have been taken from Table 4.5.

As Table 4.9 shows, increasing the number of classes leads to a decrease in U_{IGD} , and therefore a result that more closely agrees with the UEF. However, an increase in computational time also results when larger numbers of classes are used, due to the increased number of computations. For example, when 3 classes are used, a relatively low mean computation time of 75956.52 is observed, whereas when using 13 classes, this time increases to 87747.74. Looking at the same examples, it is clear that the distance between the GA-produced EF and the UEF is larger when the number of classes is smaller (U_{IGD} of 5.901e-5 for 3 classes, compared with 1.439e-5 for 13). The standard deviation for U_{IGD} is also higher for 3 classes than for 13, indicating

greater variation in the values when using fewer classes. The results are similar for the ND_{IGD} , with 3 classes producing a mean distance of 1.416e-4, compared with 2.718e-5 for 13 classes, therefore indicating better agreement with the UEF at 13 classes than at 3. Again, the standard deviation is higher for 3 classes than for 13, indicating more variation with fewer classes.

In conclusion, it is observed that as the number of classes used increases, the EF moves closer to the UEF. There is also a clear trend in time, which grows as the number of classes is increased, and is likely due to the larger numbers of combinations that are used.

		- j ,	
# Classes	U _{IGD}	ND _{IGD}	Time
3	5 901e-5/1 677e-5	1 416e-4 /7 961e-5	75956 52/8446 56
4	3.695e-5/1.191e-5	6.179e-5/3.651e-5	78315.74/8554.07
5	3.508e-5/1.111e-5	6.037e-5/3.392e-5	84426.52/7678.74
6	2.807e-5/7.490e-6	4.096e-5/1.277e-5	85063.73/5688.57
7	2.329e-5/6.976e-6	3.895e-5/1.244e-5	85482.84/7680.98
8	1.996e-5/4.678e-6	3.274e-5/1.060e-5	88419.56/6272.53
9	1.661e-5/3.110e-6	3.004e-5/7.957e-6	84045.52/5731.47
10	1.739e-5/3.181e-6	2.804e-5/6.607e-6	85932.95/7034.52
11	1.481e-5/2.304e-6	2.554e-5/6.209e-6	85650.53/8703.93
12	1.438e-5/1.647e-6	2.495e-5/5.127e-6	85269.59/5614.04
13	1.439e-5/1.684e-6	2.718e-5/5.781e-6	87747.74/8057.60

#	K	= 20,	thirty	runs,	random	IC
---	---	-------	--------	-------	--------	----

Table 4.9: Running with the number of classes shown.

On Table 4.10, it can be seen that the crowding distances when different numbers of classes are used, and presented the inclusion of oil and metals in the classes. The values for five classes have been taken from Table 4.5. As expected, using 13 classes produces the best result, with a crowding distance of 8.961e-5, compared with 1.097e-4 for 3 classes. Also, it is clear that as the number of classes used decreases, the number of points include oil and metals increases. The mean oil/metals inclusion at 12 and 13 classes is 31.067 and 31.267 out of 50 points, which are the lowest values in the table, whereas, with 3 classes, oil and metal are included in a mean of 37.733 out of 50 points. This confirms the idea that oil is a risky asset. The risk of oil is also

# Classes	CD	Oil/Metals
	1 007 1/0 0 10 F	07 700/0 010
3	1.09/e-4/2.948e-5	37.733/2.612
4	1.211e-4/2.296e-5	35.233/2.979
5	1.129e-4/1.620e-5	36.130/2.570
6	1.084e-4/1.364e-5	34.724/2.776
7	1.042e-4/1.267e-5	34.767/2.763
8	9.635e-5/1.079e-5	33.500/1.943
9	8.966e-5/7.407e-6	32.400/3.092
10	9.507e-5/1.055e-5	32.233/2.208
11	8.914e-5/6.908e-6	31.933/1.946
12	8.954e-5/6.566e-6	31.067/2.333
13	8.961e-5/7.010e-6	31.267/2.164

demonstrated by Table 4.9, in which we see that the higher numbers of classes (which contain fewer oil and metal points) result in closer agreement with the UEF.

Table 4.10: CD measurement for GCC dataset and number of points that include oil and metals.

Having discussed the effects of changing the number of classes used in the IC, the final part of this section shows plots of increasing numbers of classes (3, 7, 10 and 13, respectively). Figures 4.12- 4.15 show the relationship between GA solution and the (quadratic) UEF, the significant assets used in each case, and the heatmaps for assets. Plots show the use of 3, 7, 10 and 13 classes respectively with K = 20, with the left plot depicting the EFs; middle plot the significant assets, and the right plot the heatmap. Note that, for readability, the heatmap does not show the 500 asset numbers or the point numbers. However, the numbers given in the text are the result of calculation.

Portfolio Compositions Using 3 Classes

Figure 4.12(a) gives a comparison between the computed GA solution (green dots) and the quadratic UEF (in red), with 3 classes and an approximation of K = 20. While the GA solution aligns with the UEF at the top of the curve, returns from 8e-04 begin to show less agreement. The trend continues through an area of greatest risk (in the "bulge") at around 4e-04, right down to 2e-04 after which the GA solution begins to align with the UEF once more. Figure 4.12(b) shows that the largest number of assets used is about 47. The red line gives the mean number of significant assets across all

points. Figure 4.12(c) gives the heatmap for 500 assets across 50 points when the number of classes is 3, showing a clear change in distribution of assets. For instance, result 43 (i.e., point number 43) shows assets x150, x210 and x360 have high weights (although, due to the number of assets, this is difficult to see from the heatmap). The highest point at the top right of the EF is given by result 50 on the heatmap. It shows complete investment in asset 310. In other words, the investor would put all the capital into the highest returning asset to obtain the highest return on their portfolio.



(c) Heatmap for 500 assets, 50 points

Figure 4.12: Plots using three classes with K = 20.

Portfolio Compositions Using 7 Classes

Figure 4.13 (a) gives a second comparison between the computed GA solution (green dots) and the quadratic UEF (in red), this time with 7 classes, and again K = 20. Now, the GA solution matches much more closely with the UEF than it did with 3 classes (Figure 4.12 (a)), though a similar overall trend is observed, with agreement being closest at the top of the curve, and the greatest risk (in the "bulge") showing at a return of around 4e-04. With 7 classes, however, we also notice that a second area of relatively high risk is seen at zero return. Figure 4.13 (b) shows that with a larger number of classes the largest number of assets used is slightly lower, at about 41. Figure 4.13 (c) gives another heatmap for 500 assets and 50 points, this time when there are 7 classes, and again a change in distribution of assets is seen. For example, result 49 shows assets x151 and x315 have high weights. The highest point at the top right of the EF is given by result 50 on the heatmap, which shows complete investment in asset 360.

Portfolio Compositions Using 10 Classes

Figure 4.14 (a) gives another comparison between the computed GA solution (green dots) and the quadratic UEF (in red), with 10 classes this time, and again K = 20. While a similar trend of greatest risk at around 4e-04 is observed, as in the cases of 3 classes (Figure 4.12) and 7 classes (Figure 4.13), the overall agreement between the GA solution and the UEF is far closer, indicating an overall lowering of risk down to a maximum of under 0.0001. The plot of the significant assets (Figure 4.14(b)) now demonstrates a largest number of assets of about 48, making it more similar to Figure 4.12 than 4.13. Figure 4.14 (c) again gives the heatmap for 500 assets and 50 points, now using 10 classes, again showing a change in distribution of assets. Result 2, for instance, shows assets x176, x178 and x415 have high weights. The highest point at the top right of the EF is provided by result 50 on the heatmap, showing complete investment in asset 315.



(c) Heatmap for 500 assets, 50 points

Figure 4.13: Plots using seven classes with K = 20.



(c) Heatmap for 500 assets, 50 points

Figure 4.14: Plots using ten classes with K = 20.

Portfolio Compositions Using 13 Classes

Figure 4.15 (a) once more gives the comparison between the computed GA solution (green dots) and the quadratic UEF (in red), using all 13 classes. Again, K = 20 is used. This time, we see that agreement between the GA and UEF are very close, with an exception only in the area around 4e-04. Figure 4.15 (b) again demonstrates a largest number of assets of about 48, and the red line again gives the mean number of significant assets across all points. Figure 4.15 (c) gives the heatmap for 500 assets and 50 points when the number of classes is 13, showing that a clear change in distribution of assets for instance result 30 shows assets x33, x193 and x315 have high weights. The highest point at the top right of the EF is given by result 50 on the heatmap. It shows complete investment in asset 315.



(c) Heatmap for 500 assets, 50 points

Figure 4.15: Plots using all thirteen classes with K = 20.

The following points have emerged from the present investigation:

- It is clear that the best agreement between the computed GA solution and the quadratic UEF is seen when using all 13 classes, and that as the number of classes increases, the agreement between the computed GA and the quadratic UEF improves.
- 2. The number of significant assets does not change greatly with the variation in

number of classes, though there is a slight dip when using 7 classes.

- 3. A comparison of the four heatmaps shows that between points 1-16 there is a varying distribution, between 17-25 the distribution of assets is almost zero and from point 26 onwards the distribution is very similar.
- 4. As the number of classes decreases, the deviation away from the curve in terms of risk is very large (Figure 4.12 (a)).
- 5. It is worth noting, though, that requiring all classes to be part of the portfolio seems counter-intuitive. The reason is that such information is captured in the stock index of each country, although oil and precious metals are included there. According to (Mensi et al. 2015), investors tend to prefer diversification and the number of industries they will choose thus depends on the diversification opportunities they have in their list of choices. This acts as a balance to point 1 above, showing that there is a trade-off between these two concerns.

In the next subsection, solutions obtained by the current and the approach from Chapter 3 which uses solely the cardinality-constrained approach are compared.

4.6.3 Comparing the Current Approach with the Solely Cardinality-Constrained Approach

In this section, the approach is compared with that used in Chapter 3 to solve the cardinality-constrained problem. This time, we assume that it is being optimised with solely the value of *K* of concern. This is now applied to the data in this current chapter. Here, we present the results of a parallel GA to find the cardinality constrained efficient frontiers (CCEFs) of the GCC dataset. The num_iterations parameter was set to 2000n (as in Chapter 3) to ensure consistency with the cardinality constrained method for this small range of experiments. However, based upon preliminary runs, it was identified that the 'run' parameter needed to be reduced from 100n to 50n due to the time taken. This time issue is due to the size of the dataset; running the cardinality constrained approach on that size of dataset results in an average computation time of over three days per experiment. For example, it can be seen in Table 3.7 that the time taken with D5 and K = 5 is (11094.33) while the time taken in the GCC dataset with the same

value of K is 138934.21. As the maximum time allowed on the cluster is three days, the settings were reduced. This rendered accuracy less important, but still sufficient for an illustrative comparison to be performed. All other parameter settings are the same as those in Section 3.4.1.

We will now move on to consider Figure 4.16 as an example output. Fifty points were computed, running K = 20 and including portfolios with no classes (since the class constraint is now absent and has been replaced with the cardinality constraint). It was found that K = 20 assets enabled a close approximation of the UEF (Figure 4.16 (c)), and there is a clear difference between this plot and Figure 4.15 (a) where there is generally a very close agreement between our GA solution and that of the UEF, but with an exception in the area between 2e-04 and 7e-04, due to the presence of different constraints. However, upon running with K = 10 and K = 5, Figures 4.16 (b) and (a) were produced. These figures show, again, a very close agreement between our GA solution goes further in the area of the familiar 'bulge' around 2e-04 and 8e-04 (subfigure (b)), 3e-04 and 5e-04 (subfigure (a)), indicating higher risk compared to Figure 4.16 (c). It should be noted that very small values of K (e.g., K = 5) would be unlikely to be useful to a GCC investor who wishes to diversify their portfolio, since the presence of oil and a precious metal would account for two of the assets from an already small number.



Figure 4.16: Outputs for the GCC dataset, 50 points.

In the next section, the detailed GCC data is partitioned by year and the algorithm is tested on each of the years in order to investigate shifting frontiers.

4.7 An Investigation of Shifting Yearly Frontiers

This section studies the time shifting frontiers by using GCC countries as a sample from the period 1st Jan 2008 to 31st Jan 2019.

Given an EF for a particular year, the EF for the next year (or, indeed, for any year) is likely to be distinct. The term *shifting frontier* refers to the different positions an EF may take, whether it be 2D movement or tilting/stretching. An outward (right) shift of the frontier, for example, would mean that risk has increased. The concept is used to

illustrate changes in the fundamental risk-return relationships from year to year. Also, when investors change behaviour and preferences the frontier may shift to reflect those changes.

An experimental investigation was conducted to show unconstrained efficient frontiers (UEFs) calculated per year of the data. In Figure 4.17, each UEF corresponds to one year of the period 2008-2018, for the GCC dataset. By the figure, it is clear that the end positions of each respective UEF do indeed change. Also, it is evident that the observations every year are different. It can be seen that in the first part of the timeframe, until 2010, there had been some improvements (higher returns for the same level of risk - i.e., the successive UEFs generally shift upwards). There are enhanced levels of risk in both 2010 and 2016.



Figure 4.17: Plot of unconstrained efficient frontiers (UEFs) of the GCC dataset, per year.

The Gulf stock markets had sharp fluctuations during 2008 following the global financial

crisis, as the indices of these markets fell and trading volumes decreased. Therefore, the investment opportunities were quite constrained, which pushed the frontier down, meaning that for additional risk the return was relatively low. For example, it is noticed that, looking horizontally, risk increased from 0.001 to 0.011 for the same level of return at 0.006 in the year 2008, or, looking vertically, that return increased overall from 0.006 to 0.007 for a risk level of 0.005 from 2008 to 2009.

In 2009, the plot shows that assets of a higher amount of return were possible compared with 2008. For example at value of risk 0.008, it can be obtained higher return at 0.008 in 2009 compared to 0.006 in 2008, and there are other instances in the timeframe (2010 and 2016) when even higher returns were available, though this was only possible through purchasing assets of higher volatility (risk), as demonstrated in the example of 2010, where return values increased from 0.008 to 0.010 in return for a risk increase of about 0.005 to over 0.030.

In 2010, the levels of returns attainable by the UEF exceeded those attainable by the 2008 UEF by some 0.002-0.004. Also, it can be seen that low risk opportunities were available. Risk-return relationships along both UEFs differ - with low risk, however, the relationship is significantly smaller and the two UEFs (and, indeed, all UEFs) seem to coincide. At a risk of 0.015 onward, a higher level of return can be achieved but with more risk. In 2010 a new window of opportunity is opened with investments made achieving higher return with higher risk. This means the option of investment may be better in 2010. In 2010, gold prices started increasing and one of the most prominent factors that contributed to these higher prices was that investors may have began resorting to gold as a haven and a low-risk investment preferred to currencies, stocks and oil. Moreover, gold is relatively safe in times of economic turmoil, and when there is a state of uncertainty about the global economy.

In addition to this 2010 activity, the Middle East region also witnessed political and security tensions, which may have caused the demand for gold to rise (Ottaway & Hamzawy 2011). It seems that various factors affected the price of gold at this time, namely: that a slowdown was emerging in the economic data of a number of countries, including emerging markets and especially those of China; and that the Ameri-

can economy was experiencing a remarkable slowdown, out of which the weakened dollar, which is used in assessing gold may help increase the buying of gold, making it an attractive investment because it is priced in dollars (Bedoui et al. 2018, Maghyereh et al. 2017). For example, for a risk value of 0.007 there are different levels of return (the maximum possible of which was a value of 0.008 in 2010).

On the other hand, to achieve a return of 0.006 there were different levels of risk available across the UEFs. For example, in 2010 this return level would require relatively low risk, while in 2008, to achieve this level of return the risk would be higher. Obviously, this means there would no good quality portfolios because of the financial crisis. If we compare 2008 and 2010 with the same risk we had a different level of return, meaning that we had better investment options to achieve a high level of performance. This plot shows that a desired level of return, dependent upon the year, may be achieved with a very minimal risk or a very high one. That is, the plot shows how risk differs over time assuming a given level of returns.

However, a decline followed this period until 2014, though afterwards improvements began to be seen again and have continued year on year. It is remarkable how much investment opportunities change over time. It also shows that the degree of risk aversion of investors may also be changing and this obviously is affected by the business cycle of financial crisis and economic growth, i.e., some investors change their preferences from year to year, and there are also investors who are risk averse. These considerations all add to the fact that what is considered an optimal portfolio in 2010 is quite different to that in 2014, and we now have more opportunities for the same risk.

In 2018, a decline was seen in the Gulf stock markets due to several key reasons, foremost among which is the concern among investors as a result of political and economic crises in the gulf countries (Charfeddine & Al Refai 2019). Added to this was uncertainty about the global economy, which prompted the global stock markets to experience large losses, which in turn affected the Gulf markets, along with the continued fluctuation in oil prices which led to further variation in the Gulf markets. Oil prices continued to put pressure on the Gulf, as well as international markets, and it is well-known that changes in oil prices play a pivotal role in influencing financial markets, in

particular the stock markets (Chen & Xu 2019). Theoretically, in general, it can be said that in the event of low oil prices, the cost of production will decrease, which reflects positively on consumers through being able to obtain goods at favourable prices and thus increase demand and overall spending. Also, producing companies are able to reap more profits and greater distribution, which reflects positively on the share prices of these companies, leading to higher market indicators and improved performance. Furthermore, the improvement and increase in corporate cash flow leads to an increase in investment, which means an increase in production and employment, and consequently an improvement in stock prices and an increase in economic growth as a whole. In the case of high oil prices, however, the increase impacts negatively on the performance of companies due to the higher costs of commodity production and higher prices.

Overall, our findings confirm that an inward/upward shift of the frontier in a portfolio would indicate that the same returns can be achieved with less risk, or, that with the same level of risk, expected returns are higher. Frontiers can tilt up/down while keeping the lower left level similar if not the same, in other words they can pivot. In addition, it is apparent that frontiers may be affected by economic events. Note that the shifting frontier highlights that it may not be advisable to use too long a period in a portfolio selection, since the risk preferences of investors can change significantly over the period and therefore, the portfolio weights also change significantly (e.g. the ones allocated to safer investments). The period of data indicates that the efficient frontier (and further, an efficient portfolio on the efficient frontier) depends upon which subperiod is chosen. For example, an investor who wishes to invest over a 10-year period may take a sequence of 1-year efficient portfolios, changing composition after each year, or a 10-year efficient portfolio. This depends upon the appetite for risk and level of uncertainty - different portfolios make sense over different periods.

Table 4.11 gives the mean computed U_{IGD} and ND_{IGD} statistics, along with the standard deviations, when the algorithm is run for each of the ten years of this study (2008-2018). It also gives the computation time for each of those years. Table 4.12 gives the mean crowding distance for each year, along with its standard deviation, and also the mean oil and metal inclusion in the points.

Year	U _{IGD}	ND_{IGD}	Time
1	1.763e-04/1.231e-05	2.812e-04/8.994e-05	74348.90/5651.79
2	1.180e-04/1.458e-05	1.714e-04/3.921e-05	71252.60/3821.75
3	3.443e-04/2.180e-05	4.397e-04/3.431e-05	71447.54/3209.00
4	7.132e-05/1.643e-05	1.938e-04/8.696e-05	69231.36/6209.14
5	2.424e-04/2.985e-05	6.671e-04/2.792e-04	71525.82/6796.32
6	2.726e-04/3.993e-05	4.811e-04/7.561e-05	66711.29/3645.29
7	2.281e-04/2.108e-05	4.747e-04/1.433e-04	70207.58/3580.21
8	1.118e-04/1.290e-05	3.920e-04/1.140e-04	74499.48/4078.42
9	2.190e-04/1.109e-05	2.899e-04/3.970e-05	74526.99/4768.78
10	1.192e-04/1.377e-05	2.335e-04/7.416e-05	72572.38/4629.89
11	1.696e-04/1.889e-05	3.552e-04/1.075e-04	71337.13/4128.17

Table 4.11: Statistics produced by running algorithm for each year.

Year	CD	Oil/Metals
1	1.122e-03/1.393e-04	26.700/3.075
2	9.400e-04/5.873e-05	46.533/1.737
3	2.814e-03/2.084e-04	40.133/2.460
4	5.125e-04/8.567e-05	35.633/2.671
5	5.039e-04/1.104e-04	42.133/2.360
6	8.250e-04/9.703e-05	44.367/2.428
7	1.310e-03/7.418e-05	30.800/3.367
8	3.867e-04/3.995e-05	30.767/5.029
9	1.925e-03/1.327e-04	40.433/3.411
10	6.576e-04/6.749e-05	40.833/3.291
11	6.925e-04/5.212e-05	33.833/3.724

Table 4.12: CD measurement for GCC dataset and number of points that include oil and metals for each year.

Figures 4.18, 4.19 and 4.20 below show the efficient frontiers of testing the problem on the GCC dataset, with 5 classes and K = 20, followed by the corresponding oil and precious metals compositions, and, finally, the portfolios with number of classes actually used superimposed.

Figure 4.18 (a) gives the comparison between the computed GA solution (green dots) and the quadratic UEF (in red), with 5 classes and K = 20 in 2008. It is clear that

agreement between the GA and UEF is very close, with the exception only of the areas between returns of 0.004 and 0.006, where agreement is still quite close, and at the bottom of the curve at returns of 0 and 0.001. Figure 4.18 (b) gives the comparison once more between the computed GA solution (green dots) and the quadratic UEF (in red), again with 5 classes and K = 20, this time in 2010. Again, agreement between the GA and UEF is very close, but this time the exceptions are in the area between returns of 0.005 and 0.007, where agreement is still relatively close, and the area at the bottom of the curve at returns of 0.001 and 0.002, where agreement is furthest, though at returns of 0 the agreement is now very close.

Figure 4.18 (c) gives the same comparison, now for 2014, between the computed GA solution (green dots) and the quadratic UEF (in red), again with 5 classes and K = 20 in 2014. It is clear here that there is agreement at the top of the curve and at around returns of between 0.002 and 0.004, as well as at returns of 0. The area of greatest risk (in the "bulge") is at around 0.001, after which the GA solution aligns quite closely with the UEF once more.

Figure 4.18 (d) gives the same comparison between the computed GA solution (green dots) and the quadratic UEF (in red), again with 5 classes and K = 20 in 2018. While the GA solution aligns with the UEF at the top of the curve as in the other plots, returns from 0.003 downwards show less agreement. The trend continues through an area of greatest risk (in the "bulge") at around 0.001, but even greater than it was in 2014, after which the GA solution begins to align with the UEF again. It is noticed that in 2018 much lower risk seen at 0.004 and return reaches 0.008 while in 2010 the return is almost the same with much higher risk.



Figure 4.18: GA output for the GCC dataset, with 5 classes and K = 20.

Figure 4.19 (a) gives the comparison between the computed GA solution when using oil and metals in the classes (green dots) and when not using oil or metals (red dots) and the quadratic UEF (in red), again with 5 classes and K = 20, in 2008. It is clear that when using oil and metals, such as is indicated by the green points near the bottom, lower risk can be achieved. The reds near the top indicate a tendency for higher return but not when using oil or metals.

Figure 4.19 (b) gives the comparison again between the computed GA solution when using oil and metals in the classes (green dots) and when not using oil or metals (red dots) and the quadratic UEF (in red), again with 5 classes and K = 20, in 2010. Now quite a different result can be seen, whereby agreement at the top of the curve appears to exist whether using oil and metals or not, indicating that high returns can be obtained with oil and metals as well as without. Also, an increased investing into oil and metals rather than excluding them can be seen in 2010. It may investors interested in investing

in oil and metals after financial crisis. The frontier has shifted and in the same time it can be seen there is a clear change in preference of investors because portfolios that include oil and metals seem to perform much better compared to ones that do not include oil and metals. There is a change in investors' preferences to move from no oil and metals to inclusion of oil and metals. Again, though, the tendency for the inclusion of oil and metals to reduce risk is most clearly seen at the bottom of the curve.

Figure 4.19 (c) once again gives the comparison again between the computed GA solution when using oil and metals in the classes (green dots) and when not using oil or metals (red dots) and the quadratic UEF (in red), again with 5 classes and K = 20, in 2014. Similarly to 2010, agreement at the top of the curve appears to exist both when using oil and metals and when not, though this time the incidences of oil and metal inclusion creating higher returns are lower than in 2010. Again, the tendency for the inclusion of oil and metals to reduce risk is most clearly seen at the bottom of the curve, but this time there is a little more risk around returns of 0.001 than was seen in 2008 and 2010. The higher concentration of red points near the top again indicate a tendency for higher return and higher risk when not using oil or metals.

Figure 4.19 (d) gives the final comparison between the computed GA solution when using oil and metals in the classes (green dots) and when not using oil or metals (red dots) and the quadratic UEF (in red), again with 5 classes and K = 20, in 2018. This time the top of the curve shows far more green dots, indicating that higher returns can be obtained when using oil and metals in the classes, though the agreement is not as close as in plots a, b and c. In this 2018 case, the tendency for the inclusion of oil and metals to reduce risk is most clearly seen at returns of between 0.003 and 0.002, and the risk around returns of 0.001 is by far the highest of the 4 plots.

Overall, greens dots near the bottom of the EF show that by using oil and metals we can generally achieve lower risk (2018 being the exception - where there is a "bulge" - which is relatively high as a proportion of the highest risk, around 0.004, possible). The red dots near the top of the EF tend to indicate higher return and higher risk but not by using oil or metals. Furthermore, we can approximate the top of the EF very well with few assets, whereas the bottom of the EF requires more assets to be included.

Moreover, the frontier has shifted over these plots and at the same time it can be seen there is a clear change in risk. For example, for the same level of return at 0.006, the risk has increased from 0.002 in 2008 to 0.0018 in 2018. This may result in investor preference moving from no oil and metals to the inclusion of oil and metals.



Figure 4.19: Oil and precious metals composition with 5 classes and K = 20.

Figure 4.20 (a) gives the portfolios with number of classes used superimposed in the relevant position with 5 classes and K = 20 in 2008. It is clear that a lower number of classes - given by the superimposed numbers - for example, 3-6 actual classes used per point, yield far larger returns, in exchange for increased risk. This matches experience that fewer assets are required to approximate this top part of the EF. In addition, points with over 10 classes never reach above a 0.006 return, and points between return values of 0.004 and 0.006 and under 0.0019 are the furthest from the UEF, though still in quite close agreement.

Figure 4.20 (b) gives the portfolios with number of classes used superimposed in the relevant position, with EA settings being K = 20 across five classes, in 2010. It is again evident that a lower number of classes, particularly just 3 per point, yields far larger returns, again in exchange for increased risk. This time, points with over 10 classes do not yield returns over 0.004, and the points furthest from the curve are those with return values of 0.005 to 0.008 and around 0.001. At a risk of 0.015 onward, a higher level of return can be achieved but with more risk, with using oil and metals introducing comparatively more points at the top right corner of the EF (compared to the situation in 2008). In 2010 a new window of opportunity is opened with investments made using few classes (three) achieving higher return with higher risk. This means the option of investment may be better in 2010 as the number of classes is fewer. This is evidences by points at the top right of the EF being harder to come by in later years (subfigures (c) and (d)).

Figure 4.20 (c) superimposes in the relevant position the number of classes for each portfolio, with 5 classes and K = 20 in 2014. This time a larger number of classes yields higher returns than in 2008 and 2010, although in general the trend is still for larger numbers of classes to yield lower returns, as can be seen by the concentration of points with 10-12 classes at the bottom of the curve. Note, though, the two points containing 9 classes also found at the bottom of the curve, which differentiate this figure from (a), (b) and (d).

Figure 4.20 (d) again superimposes in the relevant position the number of classes for each portfolio, with 5 classes and K = 20, this time in 2018. In a similar way to 2014, a larger number of classes yields high returns than in 2008 and 2010, although this time the agreement with the UEF is not as close around returns of 0.004-0.007. Again, the general trend is for larger numbers of classes to yield lower returns, and now the points with 11-13 classes are both at the bottom of the curve. In 2014 the rate of return above 0.005 achieved is almost the same as in 2018, with three times higher risk than in 2018. Thus, the opportunity of investment in 2018 may be much better than in 2014 because the same level of return can be achieved at much lower risk.



Figure 4.20: Portfolios with number of classes used superimposed in the relevant position with 5 classes and K = 20.

Sharpe Ratios

The Sharpe ratio will be used to evaluate portfolio performance, since this is the simplest method with which to do so for risk-adjusted portfolios. The method was developed by Sharpe in 1966 and used to evaluate mutual fund performance. To calculate it, a portfolio's or an asset's excess return is determined, for a set period. This is then divided by a measure of its volatility - the standard deviation of the portfolio.

A simplified version of the Sharpe Ratio, neglecting the risk free rate since it is not applicable, is calculated as

Sharpe Ratio =
$$\frac{R_p - r_f}{r_p}$$
, (4.11)

where R_p is the portfolio return, r_f the risk free rate (which, as above, we take to be

zero) and r_p the portfolio risk (volatility, standard deviation of return).

This ratio provides clarity on the relationship between return and risk, demonstrating the amount of excess return that can be obtained in exchange for the extra risk. According to (Baldridge & Curry 2022), a Sharpe ratio of 1-2 is generally considered "good", 2-3 is "very good", and 3+ being "excellent": in other words, the higher the ratio, the higher the return in relation to an asset's or a portfolio's risk. This makes higher Sharpe Ratio investments generally more desirable. Figure 4.21 illustrates mean annualised Sharpe ratios over 50 points for the years 2008, 2010, 2014 and 2018.

Looking first at Figure 4.21(a) (i.e., that for 2008) it is clear that the highest Sharpe ratio is found with portfolio 24, at a ratio of over 4. This is therefore the superior portfolio regarding risk-adjusted returns. It contains only stock assets (classes 1-2, 4-7, 9 and 11), and does not contain any precious metals or oil. Looking now to 2010 (Figure 4.21(b)), it is clear that the highest Sharpe ratio of the mean annualised Ratios of 50+ points belongs to portfolio 18, at over 4 once more. This portfolio again contains stock assets (this time from classes 1-7 and 11), along with one precious metal asset. Once more, no oil is included.

For 2014 (Figure 4.21(c)), 3.5 is the highest mean annualised Sharpe Ratio and is found in portfolio 20. Therefore, portfolio 20 is the superior portfolio regarding risk-adjusted returns. It consists of stock assets (classes 1-7, 9 and 11) along with oil, but no precious metals. The graph for 2018, Figure 4.21(d), demonstrates that portfolio 17 has the highest ratio, at more than 4. It is therefore the superior portfolio with regards to risk-adjusted returns. It is constructed only from stock assets (classes 1-7, 10 and 11), with no precious metals or oil included.

Generally, it is evident that the highest performing portfolios in the four time periods considered (2008, 2010, 2014 & 2018) contained assets from classes 1 and 2, 4-7 and 11. Class 8 was not included in any of them. This suggests that those classes may offer a far better investment than others. It is worth noting that oil was only included in the better portfolio of 2014, and one precious metal in the superior portfolio of 2010. It is striking that, for the years consulted, the highest portfolio performance was seen in portfolio numbers between 17 and 24. This means that it is the middle of the efficient



frontiers where the highest performing portfolios are located.

Figure 4.21: Mean Annualised Sharpe Ratios.

4.8 Summary of the Chapter

This chapter considered portfolio optimisation problem with class constraints in the construction of diversified optimal portfolios. These were comprised of oil, one or more precious metals, and Gulf market stocks, over the period 2008 to 2018. A GA was then used on this data to provide approximations to the UEF in each case. Empirical evidence demonstrates that diversifying portfolios with these commodities reduces the risk of the portfolios, applying Markowitz's mean-variance framework.

In this chapter, all experiments used an initial configuration (IC), or ordered vector, of 5 classes out of the total possible 13. Each of these 13 classes corresponds to a particular sector, so for our purposes the IC included class 12 (oil) and 13 (precious metals) each time. Experiments were executed for several different numbers of assets, with K = 20, 30, 40 and 50. For each K value, a random IC was applied. As a result, more mixed portfolios appear in the top-right corner of the EF curve, and additionally, computation time averages were more consistent than when IC is fixed. The results indicate that the number of available assets increases with the value of K, which improves convergence between EF and UEF considerably. In addition, one of the most notable findings from the experiments is that risk is reduced when precious metals and oil combine with stock assets, compared with using stock assets on their own and this creates new investment opportunities for portfolios diversification.

It was also found that a decrease in U_{IGD} occurs when the number of classes increases, which moves the EF nearer the UEF. This larger number of combinations is likely responsible for the greater computational time, though. In contrast, a decrease in the number of classes leads to a rise in the number of points include oil and metals, which is useful information for portfolio diversification and therefore for investors.

We also examined the results for the cardinality constrained efficient frontiers (CCEFs), without using the class constraint, on the GCC dataset. In this case, good agreement with UEF was found when K = 20, but less so when K = 10. Taking a cardinality of K = 5 caused the curve to move away from the UEF, echoing the similar conclusion from Chapter 3. In simple terms, better returns and reduced risk were observed when portfolios contain 20 assets, in comparison to when they contain 5.

4.8. SUMMARY OF THE CHAPTER

This chapter concluded by investigating shifting yearly frontiers, unconstrained efficient frontiers (UEFs) calculated and investigated for each year of the dataset for GCC countries. The computed GA solution was then compared with the quadratic UEF with 5 classes, which we presented using K = 20 for the years 2008, 2010, 2014 and 2018. It is noticed that an inward/upward shift of a portfolio's frontier denotes that less risk is necessary for the same level of return, or in other words, that higher returns can be expected for the same risk levels. Frontiers can tilt up/down while keeping the lower left level similar if not the same, in other words they can pivot around the bottom left extremity. It is also evident that economic events may affect frontiers.

It is worth noting that the shifting frontier suggests that it may be advisable to limit the length of period used in portfolio selection, as portfolio weights can alter significantly over a given period, in line with investor risk preferences. Lower risk was achieved by including oil and metals, as seen in the bottom of the curve. The top of the EF curve was approximated very well with a smaller number of assets, while more assets are required at the bottom. We also found that the smaller numbers of classes, for example 3-6 classes (see Figure 4.20), were able to yield much greater returns albeit with a risk increase.

The findings of our analysis allows us to draw the following conclusions.

- The inclusion of oil and metals allows lower risk to be achieved, meaning that these can be applied as diversifying and hedging instruments. With the opportunities for expanded diversification, other commodities may be potentially added to obtain even better returns.
- It is expected that investment in oil would be attempted by investors in the Gulf, thereby justifying its step constraints, though other step penalties are available if oil is not chosen.
- Portfolio managers should try to diversify portfolios using oil and precious metals, while looking at individual classes of assets which meet their portfolio needs.

In the next chapter, the portfolio optimisation using a copula-based WCVaR approach is investigated.

Chapter 5

Portfolio Optimisation Using the Copula-Based WCVaR Approach

5.1 Introduction

A major component in portfolio optimisation problems is risk, and the most common risk measure used by practitioners and researchers is variance (Markowitz 1952). A model known as Mean-Variance (M-V) depends on a strict assumption that assets returns are multivariate normally distributed. However, this assumption is not supported in practice, and various studies have therefore made proposals for alternative risk measures that seek to move beyond the M-V model's limitations (Hoe et al. (2010), Konno & Yamazaki (1991), Young (1998)).

Calculating risk can be done in various other ways, such as Value at Risk (VaR) and Conditional Value at Risk (CVaR), as well as Worst Case Conditional Value at Risk (WCVaR). The latter considers asset dependence structures using copulas. VaR is a standard recommended by the Basel committee, but has been recently criticised (Kakouris & Rustem 2014). The first reason is VaR's lack of satisfactory sub-additivity, making it an incoherent risk measure. It can also have multiple local minimums, and is not convex (Salahi et al. 2013). The second is that the percentile of distribution tail loss is inadequate, according to Kakouris & Rustem (2014). This criticism is also used by Szegö (2005) to argue that "VaR does not measure risk", and he advises alternatives such as CVaR, since this is the distribution expectation above VaR and is impacted by distribution tail fatness, giving a better picture of distribution tail loss.

Marimoutou et al. (2009) used conditional and unconditional Extreme Value Theory models to estimate VaR in both short and long oil market trading positions. They com-

5.1. INTRODUCTION

pared their model with Historical Simulation and Filtered Historical Simulation methods and GARCH, demonstrating that significant improvements are achieved with conditional Extreme Value Theory and Filtered Historical Simulation in contrast to conventional approaches. They also note that the GARCH (1,1)-t model has the potential to yield similarly improved results, particularly for left tail. Their results corroborate that a filtering process is an important factor in success when using standard methods such as Historical Simulation and Extreme Value Theory.

To achieve a linear or convex problem for risk management and portfolio optimisation, Rockafellar et al. (2000) & Rockafellar & Uryasev (2002) suggest a minimisation formulation. Some assumptions are required for the CVaR calculation in their formula, concerning the assets' underlying distribution. Zhu & Fukushima (2009) describe this as a hypercube, ellipsoidal set or other uncertainty domain, or a multivariate distribution.

The most common multivariate distribution is the Gaussian, since calibration is easy and efficient simulation algorithms are available. Elliptical distributions also have these advantages to a point, and Credit Risk often uses student-t distribution (Chan & Kroese 2010). But the Gaussian distribution has disadvantages, such as its symmetry, indicating that losses and gains have equal probability. Further, stronger comovements of assets are seen in crisis than in prosperity, as regards financial markets (Hu 2006). Another drawback is that dependence measurements use linear correlation, which may not adequately measure dependence in non-linear asymmetric comovements (Kakouris & Rustem 2014). Mixture distributions can mitigate elliptical distributions' underlying symmetry and their limitations (Kakouris & Rustem 2014), and have been used by Hu (2006) and Smillie (2008) in the bivariate case. Copulas are defined as multivariate distribution functions with one-dimensional margins distributed equally on the [0,1] closed interval (Cherubini et al. (2004) and Nelsen (2007)). Univariate cumulative distributions of random variables can replace the uniform margins (Cherubini et al. (2004) and Nelsen (2007)). Thus, consideration is not on dependency between random variables, but on their marginal distributions, making them more flexible in comparison to standard distributions, due to the separation of the multivariate dependency and univariate distribution selection (Kakouris & Rustem 2014).

Kakouris & Rustem (2014) are motivated by Hu (2006) and Zhu & Fukushima (2009) to bring in mixture copulas to extend CVaR to WCVaR in a worst case scenario. Their Archimedean copulas each characterise a different kind of dependency. Mixing them allows Kakouris & Rustem (2014) a greater spectrum of dependencies, from which they can derive CVaR and WCVaR for copulas. Mixing copulas can result in a solution to convex optimisation problems in the WCVaR framework. Copulas also give distribution selection more flexibility, in the CVaR case.

For robust results in portfolio optimisation and risk measurement, dependency should be considered (Kakouris & Rustem 2014). The authors propose copulas and mixture copulas as one way to achieve this, and compare them against Gaussian CVaR and Worst Case Markowitz. Empirical analysis reveals that WCVaR outperforms them in crisis periods.

Robust portfolio optimisation has also been combined with copula models by (Sabino da Silva & Ziegelman 2017), in a Worst Case CVaR framework (WCCVaR). They made comparisons with a Gaussian Copula-CVaR portfolio (GCCVaR), alongside the S&P 500 index and an equally weighted portfolio (1/N). Better hedges were found with the copula basis than the 1/N portfolio, against losses. Better downside risk statistics were found in portfolios with the WCCVaR approach than the GCCVaR in rebalancing periods, as well as more profitability when daily or weekly rebalancing.

Chukwudum (2018) examined the shaping of the extremal dependence structure with a combination of copulas and Generalised Pareto distribution, demonstrating the effect that tail dependence has on risk measures, reinsurance net premium and risk allocation. Applied to the insurance sector in Nigeria, it was shown that correlating losses increases risk measure value, as well as impacting risk allocation. It was also shown that the use of the Clayton copula generally lowered reinsurance premiums, whereas with the Gumbel copula these were generally increased, compared with premiums hypothesised independently.

The objective of this chapter is to apply the WCVaR based on copulas on the GCC dataset which was introduced in the previous chapter, in order to accurately assess risk to investment.

This chapter is organised as follows: first, the definition of copula and popular copula families are introduced. Then, dependence measures (Pearson's correlation coefficient, Kendall's τ coefficient and Spearman's rank correlation coefficient) are discussed. Different methods of copula inference are introduced. WCVaR using the copula-GARCH model is explained in detail. Finally, the WCVaR and VaR for the large GCC dataset are computed and then the results obtained are compared.

5.2 Copulas

Copulas are tools which permit distributions of individual assets to be separated from their dependence structure. They are also used to model nonlinear dependence. The benefits of this are that any distribution can be assumed when considering marginal distributions of assets, with copulas being used to obtain the joint multivariate distribution, taking marginals dependence structure into account. Moreover, in general, if we consider the loss distribution of financial data, the left and the right tail represent profit and loss respectively and they are often not the same (i.e., one tail may be heavier than the other). Therefore, a symmetric distribution, like the normal distribution may, not be appropriate. Using a copula in this situation is a more flexible choice because, depending on the type of copula and on the parameter values, we can have different shapes of distributions. Thus a more accurate representation of joint distribution of returns and losses. If we consider several assets, a multivariate distribution describes simultaneously what happens with all assets and in a copula setting the marginal of each asset may be different (Joe 1997, Nelsen 2007).

In the definitions contained in this section, we use \sim to mean "distributed as" (for example, $U_k \sim U[0,1]$ means that U_k conforms to the uniform distribution over [0,1]).

Definition 6 (Copula). (Mai & Scherer 2017) A function $C : [0,1]^d \to [0,1]$ is called a (*d*-dimensional) copula if there is a probability space $(\Omega, \mathscr{F}, \mathbb{P})$ supporting a random vector (U_1, \ldots, U_d) such that $U_k \sim U[0,1]$ for all $k = 1, \ldots, d$ and

$$C(u_1,\ldots,u_d)=\mathbb{P}(U_1\leq u_1,\ldots,U_d\leq u_d),$$

where $u_1, ..., u_d \in [0, 1]$.

Theorem 2 (Sklar's Theorem). (Daníelsson 2011, p.27) Let F be the distribution of X, G the distribution of Y and H the joint distribution of (X,Y). Assume that F and G are continuous. Then there exists a unique copula C such that

$$H(X,Y) = C(F(X), G(Y)).$$
 (5.1)

Sklar's Theorem describes how to use copulas to derive the joint multivariate distribution and embedding the dependence structure between the marginals.

Definition 7. (Blitzstein & Hwang 2014) A bivariate function $f : \mathbb{R} \times \mathbb{R} \to \mathbb{R}$ is a joint (bivariate) probability density function if and only if

• $f(x,y) \ge 0$ for all $x \in \mathbb{R}$ and $y \in \mathbb{R}$,

•
$$\int_{-\infty}^{\infty} \int_{-\infty}^{\infty} f(x,y) dx dy = 1.$$

Definition 8. (Kenney & Keeping 1951, p. 92) The bivariate normal distribution is the statistical distribution with probability density function

$$f(x,y) = \frac{1}{2\pi\sigma_1\sigma_2\sqrt{1-\rho^2}} \exp\left[\frac{-z}{2(1-\rho^2)}\right]$$
(5.2)

where

$$z = \frac{(x - \mu_1)^2}{\sigma_1^2} - \frac{2\rho(x - \mu_1)(y - \mu_2)}{\sigma_1\sigma_2} + \frac{(y - \mu_2)}{\sigma_2^2}$$

and

$$\rho = cor(x, y) = \frac{V_{12}}{\sigma_1 \sigma_2}$$

is the correlation of x and y, V_{12} is the covariance, μ_1 and μ_2 are the means of x and y respectively and σ_1 and σ_2 are the standard deviations of x and y respectively.

5.2.1 The Fréchet-Hoeffding bounds

Any multivariate copula $C_{\theta}(u_1, \dots, u_d)$, where θ is copula dependence parameter vector, should satisfy the Fréchet-Hoeffding bounds such that, for every $(u_1, \dots, u_d) \in [0, 1]^d$:

$$W \leq C_{\theta}(u_1, \cdots, u_d) \leq M.$$

W is called the lower Fréchet-Hoeffding bound, that represents perfect negative dependence between variates and is defined as:

$$W = \max\left(1 - n + \sum_{i=1}^{n} u_i, 0\right).$$

M is called the upper Fréchet-Hoeffding bound, that represents perfect positive dependence between variates and is defined as:

$$M=\min(u_1,\cdots,u_n).$$

5.3 Dependence Measures

The main dependence measures are:

- Pearson's Correlation Coefficient ρ
- Kendall's τ Coefficient
- Spearman's Rank Correlation Coefficient ρ_s

5.3.1 Pearson's Correlation Coefficient ρ

Pearson's correlation coefficient is the most popular indicator to examine dependence between two random variables *X* and *Y*, defined as:

$$\rho_{XY} = \frac{Cov(X,Y)}{\sqrt{Var(X)}\sqrt{Var(Y)}}.$$
(5.3)

5.3.2 Kendall's τ Coefficient

Kendall's τ is a measure of rank correlation. Assume that (X_1, Y_1) and (X_2, Y_2) are two independent pairs of random variables drawn from the random vector (X, Y). Kendall's τ is defined by:

$$\tau(X,Y) = \Pr[(X_1 - X_2)((Y_1 - Y_2) > 0] - \Pr[(X_1 - X_2)((Y_1 - Y_2) < 0].$$
(5.4)

The term $Pr[(X_1 - X_2)(Y_1 - Y_2) > 0]$ refers to Pr[concordance] while the term $Pr[(X_1 - X_2)(Y_1 - Y_2) < 0]$ refers to Pr[discordance], and thus $\tau(X, Y)$ is a measure of the relative difference between the probability of concordance and the probability of discordance.

5.3.3 Spearman's Rank Correlation Coefficient ρ_s

Spearman's ρ_s is a statistical measure of the strength of a monotonic relationship between ranked variables. If *X* and *Y* are two random variables with continuous distribution functions F_1 and F_2 , respectively and *F* is the joint distribution function, then Spearman's ρ_s is:

$$\rho_s(X,Y) = \rho(F_1(X), F_2(Y)).$$
(5.5)

5.4 Popular Copula Families

This Section gives basic definitions of the most commonly applied copula families. These comprise elliptical copulas and Archimedean copulas.

Elliptical Copula

Elliptical copulas are the copulas of elliptical distributions. Elliptical distribution families are broadly used in statistics and economics, principally in finance (Frahm et al. 2003). They are used because they share many of the tractable properties of the multivariate normal distribution (Cheng et al. 2007). Additionally, they are used commonly to model dependencies among a great number of marginals. Frahm et al. (2003) studied the dependence structure generated by elliptical distributions. Owen & Rabinovitch (1983) discuss the relation between elliptical distributions and portfolio theory and its empirical applications. de Melo Mendes & de Souza (2004) used a new approach based on
copulas to calculate financial risk. In this Section, we restrict the discussion to the most commonly applied elliptical copulas in the literature: Gaussian copula and Student's tcopula.

Normal Copula

Definition 9. (Daníelsson 2011, p. 25), (Trivedi et al. 2007) Let Φ be the cumulative distribution function of a univariate normal distribution, Φ^{-1} be its inverse function and θ be dependence parameter. Let $U, V \sim U(0, 1)$ and $\Phi_G(u, v)$ be the bivariate normal distribution function with correlation *G*. The normal Gaussian copula is defined as:

$$C(U,V;\boldsymbol{\theta}) = \Phi_G(\Phi^{-1}(u), \Phi^{-1}(v);\boldsymbol{\theta})$$

$$= \int_{-\infty}^{\Phi^{-1}(u)} \int_{-\infty}^{\Phi^{-1}(v)} \frac{1}{2\pi(1-\theta^2)^{\frac{1}{2}}} \left(-\frac{(s^2-2\theta st+t^2)}{2(1-\theta^2)} \right) ds dt.$$
(5.6)

The normal copula is radially symmetric and has dependence level equal in the upper and lower tail. The normal copula can represent both negative and positive dependence for equal degrees. The whole Fréchet lower and upper bounds are reached.

Figure 5.1 shows a contour plot of normal copula. The contour plot is a way to display a three-dimensional surface in two-dimensions. It represents lines that join points of equal probability density. With respect to dependence between two marginals, it is shown that high values of one marginal are associated with high values of the other marginal; that indicates a positive dependence between two marginals. In addition, the copula parameter $\theta = 0.7$ corresponds to Kendall's $\tau = 0.5$. From Table 5.1 it is clear how to obtain Kendall's τ from the copula parameter.

Student's t-Copula

Definition 10. (Demarta & McNeil 2005) The Student's t-copula is given by

$$C_{\nu,P}^{t}(U,V) = \int_{-\infty}^{t_{\nu}^{-1}(u)} \int_{-\infty}^{t_{\nu}^{-1}(v)} \frac{\Gamma(\frac{\nu+2}{2})}{\Gamma(\frac{\nu}{2})\sqrt{(\pi\nu)^{2}|P|}} \left(1 + \frac{X'P^{-1}X}{\nu}\right)^{-\frac{\nu+2}{2}} dX;$$
(5.7)

where $X = (X_1, X_2)'$ is the two-dimensional random vector that has a (non-singular) multivariate *t* distribution with degrees of freedom *v*, *P* is the correlation matrix, t_v is the



Figure 5.1: Normal copula.

density function and t_v^{-1} is the quantile function of a standard univariate t_v distribution.

The Student's t-copula includes the whole Fréchet lower and upper bounds. The Student's t-copula is flexible in that it can describe both negative and positive dependence and is radially symmetric, has dependence level equal in the upper and lower tails. Figure 5.2 shows the contour plot of the Student's t-copula with $\theta = 0.7$ and degrees of freedom v = 3.



Figure 5.2: Student's t-copula.

Archimedean Copulas

Definition 11. (Ruppert 2011) A d-dimensional copula is called Archimedean if

$$C(u_1, ..., u_d) = \phi^{-1} \{ \phi(u_1) + ... + \phi(u_d) \}$$
(5.8)

where the function ϕ is the generator of the copula and satisfies

- φ is a continuous, strictly decreasing, and convex function mapping [0,1] onto [0,∞];
- $\phi(0) = \infty;$
- $\phi(1) = 0.$

There are many families of Archimedean copulas, but we will focus on three commonly used (e.g., Kakouris & Rustem (2014) provides a copula formulation for VaR and ES): the Clayton, Frank and Gumbel copula. All definitions below are from Nelsen (2007) and Ruppert (2011).

Definition 12 (Clayton Copula). The Clayton copula has generator $\phi_{\theta}(t) = \frac{1}{\theta}(t^{-\theta} - 1)$, $\theta > 0$ and is given by

$$C_{\theta}(u,v) = \left[\max \left(u^{-\theta} + v^{-\theta} - 1, 0 \right) \right]^{-\frac{1}{\theta}}.$$
(5.9)

The Clayton copula is asymmetric and it expresses strong left tail dependence and weak right tail dependence. The Clayton copula does not allow negative dependence. Figure 5.3 shows the contour plot of Clayton copula with $\theta = 2$.

Definition 13 (Frank Copula). The Frank copula with generator $\phi_{\theta}(t) = -\ln\left(\frac{e^{-\theta t}-1}{e^{-\theta}-1}\right)$, $-\infty < \theta < \infty$ has the form

$$C_{\theta}(u,v) = -\frac{1}{\theta} \ln\left(1 + \frac{(e^{-\theta u} - 1)(e^{-\theta v} - 1)}{e^{-\theta} - 1}\right).$$
 (5.10)

The Frank copula is radially symmetric and is comprehensive, since it reaches the lower and upper Fréchet boundaries. It is able to model positive and negative depen-



Figure 5.3: Clayton copula.

dence. However, it is more useful for data in which tail dependence is relatively weak. Figure 5.4 shows the contour plot of Frank copula with $\theta = 5.74$.



Figure 5.4: Frank copula.

Definition 14 (Gumbel Copula). *The Gumbel copula having generator* $\phi_{\theta}(t) = (-\ln t)^{\theta}$, $\theta \ge 1$, *is written as*

$$C_{\theta}(u,v) = \exp\left(-\left[\left(-\ln u\right)^{\theta} + \left(-\ln v\right)^{\theta}\right]^{\frac{1}{\theta}}\right).$$
(5.11)

The Gumbel copula is not comprehensive and is radially asymmetric. It is the opposite to the Clayton, since it is well suited to represent strong right tail dependence and weak left tail dependence. The Gumbel copula is similar to the Clayton, since it cannot account for negative dependence. Figure 5.5 shows the contour plot of the Gumbel copula with $\theta = 2$.



Figure 5.5: Gumbel copula

					I
Kendall's $ au$ and range	$rac{2}{\pi} \operatorname{arcsin}(heta)$ -1 $\leq au \leq 1$	$rac{2}{\pi} rcsin(heta)$ $-1 \leq au \leq 1$	0 < au < 1	$1-rac{1}{ heta}$ $0\leq au < 1$	$-1 \leq au \leq 1$
	$1 \leq \theta \leq 1$ $\theta = 0$ is independence	$\[\] 1 \leq \theta \leq 1 \] \[\] \theta = 0 \]$ independence	$0 < \theta < \infty$ $\theta \rightarrow 0$ is independence	$0 \le \theta < \infty$ $\theta = 0$ is independence	$-\infty < \theta < \infty$ $\theta \rightarrow 0$ is independence
Archimedean generation function ψ	N/A	N/A	$\phi(t)=rac{1}{ heta}(t^{- heta}-1)$	$\phi(t) = (-ln)^{ heta}$	$\phi(t) = -ln \left[\frac{e^{-\theta} - 1}{e^{-\theta} - 1} \right]$
Dependence structure characteristics	Radially symmetric, weak tail dependencies, left and right tail dependencies go to zero at extremes	Radially symmetric, weak tail dependencies, left and right tail dependencies go to zero at extremes	Radially asymmetric, strong left tail dependence and weak right tail dependence, right tail dependence goes to zero at right extreme	Radially asymmetric, weak left tail dependence, strong right tail dependence, left tail dependence goes to zero at left extreme	Radially symmetric, very weak tail dependencies left and right tail dependencies go to zero at extremes
Copula	Gaussian	Student's t	Clayton	Gumbel	Frank

5.5 Tail Dependence

The tail dependence is used for describing the association between the upper or lower tail of two or more random variables (AghaKouchak et al. 2013). The tail dependence can be defined as follows.

Let u, v be random variables, λ_U and λ_L are upper tail dependence and lower tail dependence, respectively and they are defined by Trivedi et al. (2007):

$$\lambda_L = \lim_{\nu \to 0^+} \frac{C(\nu, \nu)}{\nu} \tag{5.12}$$

$$\lambda_U = \lim_{\nu \to 1^-} \frac{S(\nu, \nu)}{1 - \nu}$$
(5.13)

where S(v,v) = Pr[U > v, V > v] is the joint survival function with $U = F_1^{-1}(x)$, $V = F_2^{-1}(Y)$. F_1^{-1} and F_2^{-1} denote the generalized inverse distribution functions of *X* and *Y* respectively (Trivedi et al. 2007).

5.6 Copula Inference

5.6.1 Maximum Likelihood Method

Suppose $\{(X_{i1}, \dots, X_{ip})^T, i = 1, \dots, n\}$ are *n* independent realisations from a multivariate distribution. Let F_i be the cumulative distribution functions cdfs and f_i be the probability density functions pdfs of the *p* margins with $j = 1, \dots, p$ and let *c* be a copula density. Let $\theta = (\beta^T, \alpha^T)$ be the parameter vector to be estimated, β be the vector of marginal parameters and α be the vector of copula parameters. The loglikelihood function can be defined as follows:

$$l(\theta) = \sum_{i=1}^{n} \log c\{F_1(X_{i1}; \beta), \cdots, F_p(X_{ip}; \beta); \alpha\} + \sum_{i=1}^{n} \sum_{j=1}^{p} \log\{f_i(X_{ij}; \beta)\}$$
(5.14)

The ML estimator of θ is

$$\hat{\theta}_{ML} = argmax_{\theta \in \Delta} l(\theta), \tag{5.15}$$

where Δ is the parameter space (Yan et al. 2007).

5.6.2 Inference Functions for Margins (IFM)

The number of parameters rises when the dimension p gets large. Thus, the optimisation problem becomes more difficult. Joe & Xu (1996) presented the inference functions for margins (IFM) method. This is a two-step method which estimates the marginal parameter β as follows:

$$\hat{\beta}_{IFM} = argmax_{\beta} \sum_{i=1}^{n} \sum_{j=1}^{p} \log f_i(X_{ij}; \beta)$$
(5.16)

After that the association parameters α is estimated as:

$$\hat{\alpha}_{IFM} = \operatorname{argmax}_{\alpha} \sum_{i=1}^{n} \log c \left((F_1(X_{i1}; \hat{\beta}_{IFM}), \cdots, F_p(X_{ip}; \hat{\beta}_{IFM}); \alpha) \right)$$
(5.17)

As a result there is small number of parameters for each maximisation task. This leads to diminished computational difficulty (Yan et al. 2007).

5.6.3 Canonical Maximum Likelihood Method (CML)

The canonical maximum likelihood method is used to estimate the dependence parameter α without specifying the marginal distributions. The canonical maximum likelihood method converts the observations (X_{i1}, \dots, X_{ip}) into pseudo-observations with uniform marginals (U_{i1}, \dots, U_{ip}). The CML method estimates α as:

$$\hat{\alpha}_{CML} = \arg\max_{\alpha} \sum_{i=1}^{n} \log c\left(U_{i1}, \cdots, U_{ip}; \alpha\right)$$
(5.18)

In order to construct a copula model from data, test of multivariate independence and goodness-of-fit tests can be used (Yan et al. 2007).

5.6.4 Test of Multivariate Independence

A test of multivariate independence is testing independence among the components of *X* (Kojadinovic et al. 2010):

$$H_0: C = \Pi$$
 against $H_1: C \neq \Pi$,

where $\Pi(\mathbf{u}) = \Pi_{i=1}^{d} u_i$, $\mathbf{u} \in [0,1]^{d}$, is the independence or product copula. Let $X = (X_1, \dots, X_d)$ be a random vector, F_1, \dots, F_d the continuous marginal cumulative distribution functions (cdfs) and the cumulative distribution functions (cdf) *H* of *X* as follow:

$$H(x) = C(F_1(x_1), \cdots, F_d(x_d)), \quad X \in \mathbb{R}^d$$
 (5.19)

The test of multivariate independence is based on the statistic:

$$I_n = \int_{[0,1]^d} n \left[C_n(\mathbf{u}) - \prod_{k=1}^d u_i \right]^2 d\mathbf{u},$$
 (5.20)

where **u** = (u_1, \dots, u_d) , with $u_i \in [0, 1], i = 1, \dots, d$.

5.6.5 Goodness-of-fit Test

The goodness-of-fit tests has the following hypotheses:

$$H_0: C \in \mathbf{C}$$
 vs $H_1: C \notin \mathbf{C}$

for a parametric copula family $\mathbf{C} = \{C_{\theta}\}$, where θ is the copula dependence parameter. It is used to fit an appropriate parametric copula family to data. This test depends on the statistic

$$S_n = \sum_{i=1}^n \left[C_n(\mathbf{U}_i) - C_{\theta_n}(\mathbf{U}_i) \right]^2,$$
 (5.21)

where C_n is the empirical copula and C_{θ_n} is an estimator of *C* under the hypothesis that $H_0: C \in \{C_\theta\}$ holds and $\mathbf{U}_i = \frac{R_i}{n+1}, i \in 1, ..., n$ are pseudo-observations from *C* (Kojadi-novic et al. 2010).

5.7 The GARCH (p,q) Model

The GARCH model is a time series model that is probably the most widely used in finance. It is used to evaluate the volatility of returns in financial markets. It is known that stock returns exhibit heavy-tailed probability distributions. This is might be because the conditional variance is variable. Using GARCH model can capture fat tails and volatility clustering within financial time series. Bollerslev (1986) developed the Generalized AutoRegressive Conditional Heteroscedasticity (GARCH) model as a generalization of the ARCH model proposed by Engle (1982).

A GARCH(p,q) model takes the following form:

$$\varepsilon_t = W_t \sqrt{h_t}$$

$$h_t = \alpha_0 + \sum_{i=1}^p \alpha_i \varepsilon_{t-i}^2 + \sum_{j=1}^p \beta_j h_{t-i},$$
(5.22)

where ε_t is return at time *t*, W_t is residual, $h_t = VaR[\varepsilon_t | \mathscr{F}_{t-1}]$ where \mathscr{F}_{t-1} is the information up to time t - 1; $\alpha_0 > 0$, α_i , β_j are model parameters.

The GARCH (1,1) Case

The GARCH (1,1) model is the most popular model used in empirical finance. It can be written as:

$$\varepsilon_t = W_t \sqrt{h_t}$$

$$h_t = \alpha_0 + \alpha_1 \varepsilon_{t-1}^2 + \beta_1 h_{t-1}.$$
(5.23)

Figure 5.6 shows simulated time series data ε_t , $t = 1, \dots, n$, from a GARCH(1,1):

5.7.1 Basic Properties

 Uniqueness and stationarity. As stated in Fryzlewicz (2007) the system of equations 5.22 has a unique and stationary solution if

$$\sum_{i=1}^{p} \alpha_i + \sum_{j=1}^{p} \beta_j < 1.$$
 (5.24)

• Mean zero. This is the most important property of GARCH models and it states that the mean of ε_t is zero:

 $E[\mathbf{\varepsilon}_t]=0.$

Lack of serial correlation. For h > 0, ε_t is not correlated to ε_{t+h}:
 E[ε_t, ε_{t+h}] = 0.



Figure 5.6: A plot of simulated time series data from GARCH(1,1).

· Unconditional variance. It is given as follows:

$$E[\varepsilon_t^2] = \frac{\alpha_0}{1 - \sum_{i=1}^R \alpha_i + \beta_i},$$
(5.25)

where R = max(p,q).

• Heavy tails of ε_t . It can be shown that ε_t does not have all finite moments. In particular, $E(h_t^q)$ is infinite with q > 0. Therefore, ε_t has heavy tails.

5.8 Copula-GARCH Model

The copula-GARCH model is used to measure financial risks like portfolio risk and volatility over time. This model aims to describe asymmetric dependence and complex non-linear relations between assets. Copulas have been used to estimate risk measures in financial markets such as (VaR) which is the risk measure extensively used by financial analysts (Messaoud & Aloui 2015). There are several studies that have used copula for determining VaR. For example, Huang et al. (2009) used the copula-GARCH model to estimate VaR of portfolios. The authors found that the copula-GARCH model is flexible in separating marginal distributions from the dependence structure which is more effective in describing high volatility. This is contrary to traditional methods that computed VaR using the variance-covariance and the Monte Carlo approaches. Em-

brechts et al. (2005) introduced a methodology that is based on the theory of copulas to calculate VaR in the worst case scenarios. Cherubini & Luciano (2001) used copula functions to estimate VaR and allocate capital. They applied their approach to the data of five years of daily returns on two stock market indices: FTSE100 and S&P100. They found that copulas separate the marginal distributions from the dependence between returns results in joint probabilities of extreme losses. Jondeau & Rockinger (2006) proposed a new measure of conditional dependence using copula functions with GARCH process. Fortin & Kuzmics (2002) used convex linear combinations of copula with GARCH-model to estimate VaR on a set of European stock indices: FSTE and DAX. Messaoud & Aloui (2015) use the GJR-GARCH based on a student t-copula to estimate the VaR and CVaR of a portfolio and describe asymmetric dependence. The authors consider daily returns of market indices from four countries: Egypt, Malaysia, South Africa, and Turkey.

5.9 Worst Case Conditional Value at Risk

There are many measures to calculate risk. A common used risk measure is VaR, but it has been criticised because it is not a coherent measure of risk, as well as being a percentile of loss distribution which does not describe possible losses on the tails of distribution. That leads Szegö (2005) to propose the CVaR as an alternative measure. CVaR is a coherent measure and it gives some better properties of loss on the tail of distribution than VaR (Quaranta & Zaffaroni (2008), Bertsimas et al. (2004), Rockafellar & Uryasev (2002)). Kakouris & Rustem (2014) presented the copula formulation of the CVaR of a portfolio. The authors extended the CVaR to WCVaR using mixture copulas. In this section VaR, CVaR and WCVaR will be discussed as follows:

Definition 15. Let $\mathbf{w} \in \mathbb{W} \subset \mathbb{R}^m$ be a decision vector, $\mathbf{u} \in \mathbb{I}^n$ a random vector, $\tilde{g}(\mathbf{w}, \mathbf{u})$ the cost function and $\mathbf{F}(\mathbf{x}) = (F_1(x_1), ..., F_n(x_n))^T$ a set of marginal distributions where $\mathbf{u} = \mathbf{F}(\mathbf{x})$. Also, assume that \mathbf{u} follows a continuous distribution with copula function $C(\cdot)$. Then for a confidence level β , *VaR* is defined as

$$VaR_{\beta}(w) = min\{\alpha \in \mathbb{R} : C(\boldsymbol{u} | \tilde{g}(\boldsymbol{w}, \boldsymbol{u}) \le \alpha) \ge \beta\}.$$
(5.26)

Definition 16. Given w, u, F(x) and $\tilde{g}(w, u)$ as in definition 15, $CVaR_{\beta}$ for a confidence

level β is defined as

$$CVaR_{\beta}(\mathbf{w}) = \frac{1}{1-\beta} \int_{\tilde{g}(\mathbf{w},\mathbf{u}) \ge VaR_{\beta}(\mathbf{w})} \tilde{g}(\mathbf{w},\mathbf{u})c(\mathbf{u})d\mathbf{u}.$$
(5.27)

Definition 17. The WCVaR for fixed $\mathbf{w} \in \mathbb{W}$ and confidence level β , with respect to C is defined as

$$WCVaR_{\beta}(\mathbf{W}) = sup_{c(\cdot)\in C}CVaR_{\beta}(\mathbf{W}).$$
(5.28)

5.9.1 Worst Case GARCH-Copula CVaR Portfolio Optimisation

In order to apply the Worst Case Copula-CVaR Portfolio Optimisation Sabino da Silva & Ziegelman (2017) we need firstly to model the marginals with the GARCH model assuming skew t-distributed innovations. Then, we extract the residuals and transform them into pseudo-observations. Next, we estimate a t-copula to model the dependencies between marginals. Then, we generate a high number of scenarios (e.g simulations) using the dependence structure that we obtained from estimating the copula. After that, we transform these scenarios into quantiles of the t-distribution. Then, we derive the simulated daily asset log-returns using the standard deviations which have been estimated by the GARCH model. Finally, to optimise the portfolio weights we apply the WCVaR technique to find the optimal portfolio.

5.9.2 An Example of Worst Case GARCH-Copula CVaR

The following example demonstrates the computation of VaR and CVaR (Expected Shortfall) for a portfolio using multivariate copula simulation with GARCH marginals. We downloaded data from Yahoo Finance, and imported the closing prices of six market indexes that we modelled with GARCH processes. The indexes are SPY: Large Cap US (S&P 500), EEM: Emerging Markets Equity, TLT: 20+ Year Treas Bond (iShares Barclays), COY: US High-Yield Bond, GSP: Commodities broad (iPath S&P GSCI Total Return Index) and RWR: Real estate (REIT Index). We use daily data over the period from 01-10-2009 to 24-06-2013. We then applied the methodology that we described above in Section 5.9.1 on these data.

Figure 5.7 shows the relative price movements of each index. The initial level of each index has been normalised to unity to facilitate the comparison of relative performance over the historical record.



Figure 5.7: Performance of price movements of each index over the historical record.

We characterise individually the distribution of returns of each index. Since copulas allow us to capture dependences between the marginals, we investigated the existence of correlations between the returns using scatterplots and correlation coefficients.

Figure 5.8 illustrates scatterplots between each pair of index returns in the lowertriangular panels, histograms of each marginal in the diagonal and Pearson's correlation coefficients in the upper-triangular panels. For instance it is evident that there is a positive strong correlation between SPY and EEM as well as between SPY and RWR and between EEM and RWR, whereas we notice a negative correlation between SPY and TLT. However, we can see that there is a weak correlation for example between COY and GSP.

Then, we specify and estimate GARCH(1,1) models for each marginal with t-distributed innovation processes and we extracted the standardised residuals. Figure 5.9 shows the Student's t-quantile functions of the residuals, allowing us to check if the assumptions of the GARCH models are met.



Figure 5.8: Scatterplots between each pair of index returns in the lower-triangular panels, histograms of each marginal in the diagonal and Pearson's correlation coefficients in the upper-triangular panels.



Figure 5.9: The Student's t-quantile functions of the residuals of GARCH(1,1) for each marginal.

On Figure 5.10, the plots demonstrate the autocorrelation functions (ACF) of the squared, mean-adjusted values. The ACF calculates a single time series' similarity against a de-

layed version (with a delayed copy) of itself (Welsh 1999). It is clear from the plots that there is serial correlation in the squared values thus the GARCH model is appropriate.



Figure 5.10: Plots of autocorrelation functions (ACF).

Figure 5.11 was produced after transforming the standardised residuals into pseudoobservations, which are constrained in order to belong to the set [0,1]. Figure 5.11 illustrates scatterplots between pseudo-observations of each pair of indexes in the lower-triangular panels, histograms of each marginal in the diagonal and Pearson's correlation coefficients in the upper-triangular panels. Comparing Figure 5.11 with Figure 5.8 it can be seen that there is a similar behaviour in terms of dependence between marginals.

The dependence structure between marginals is estimated using a t-copula. Then, the dependence determined by the estimated t-copula was used for generating N = 10000 random variates for the pseudo-uniformly distributed variables. After that, we use these quantiles in conjunction with standard deviations to calculate the N = 10000 portfolio return scenarios.

Next, we transformed the uniform variates to daily centered returns via the inverse cdf of each index using Student's t distributions for each marginals. The multivariate simulations from the copula model can be used to compute the Value-at-Risk (VaR) and



Figure 5.11: Scatterplots between pseudo-observations of each pair of indexes in the lower-triangular panels, histograms of each marginal in the diagonal and Pearson's correlation coefficients in the upper-triangular panels.

Expected Shortfall (CVaR) of a sample portfolio. We could also find optimal portfolio weights that give us a minimum risk for a certain level of return. We can do this using the WCVaR technique, which allows us to find optimal portfolio weights that give us a minimum risk for a certain level of return. Using the WCVaR method we obtained the following weights: 0.37916193 for SPY, 0 for EEM, 0.52727024 for TLT, 0.03467091 for COY, 0.05889691 for GSP and 0 for RWR. We find that VaR and CVaR are 1.340576 and 1.675549 respectively. We also calculated VaR and CVaR for a multivariate normal as it is a traditional method and we obtained optimal weights of 1.087546 and 1.267103 respectively.

In order to compare the WCVaR technique with the multivariate normal, we also calculated VaR and CVaR adopting randomly chosen weights (0.1 for SPY, 0.2 for EEM, 0.3 for TLT, 0.2 for COY, 0.1 for GSP and 0.1 for RWR). VaR is 1.641517 and CVaR is 2.122687. Finally, we compared the results obtained above with data generated from a multivariate normal distribution, getting a VaR and CVaR of 1.619638 and 1.836164 respectively.

Table 5.2 shows the comparison between the traditional multivariate normal method

and the WCVaR technique. These results reveal that the traditional multivariate normal method gives lower VaR and CVaR and underestimates risks. While the WCVaR method takes into proper account the strong dependence in the tails using a Student's t-copula, the traditional multivariate normal fails to do that, yielding lower risk measures.

	With (Optimal Weights	With Random Weights		
	WCVaR	Multivariate Normal	WCVaR	Multivariate Normal	
VaR CVaR	1.340576 1.675549	1.087546 1.267103	1.641517 2.122687	1.619638 1.836164	

Table 5.2: Comparison between traditional multivariate normal method and WCVaR technique.

Figure 5.12 shows the comparison between the WCVaR and minimum-variance portfolio values. It is clear that the WCVaR and the minimum-variance are very similar in the period 01-10-2009 to 01-09-2011, while the minimum variance measure exceeds the WCVaR from 02-09-2011 to 01-10-2012. From 02-10-2012 to 01-06-2013, there is a sharp increase in the value of WCVaR, which outperforms the minimum-variance.



Figure 5.12: Trajectory of WCVaR and minimum-variance portfolio values.

5.9.3 Computing the WCVaR and VaR for a Large Dataset

In this section we apply the WCVaR method and VaR to the GCC dataset which includes stock markets from Saudi Arabia, Qatar, the United Arab Emirates, Oman, Bahrain and Kuwait - more details are shown in Chapter 4. The daily returns of the GCC stock index are considered and taken from the Thomson Datastream service via the Plymouth Business School. The sample period of the index data runs from 01/01/2008 to 31/01/2019 and comprises a total of 2894 observations. Since the dataset used in this section is very large, sixty-four cores of the HPC cluster at the University of Plymouth were used to process the experiment. Table 5.3 shows the results of the VaR and WCVaR measures for the GCC dataset. Since the dataset is very large (496 assets), it was partitioned by class constraints and the methodology was applied to each class individually; the number of assets assigned to each class is shown in Table 5.3. We focused on the following eight classes: Energy (20 assets), Materials (68 assets), Industrials (52 assets), Healthcare (9 assets), Information Technology (5 assets), Communication Services (13 assets), Utilities (7 assets), and Real Estate (47 assets).

In Table 5.3, classes 4, 5, 7, 12 and 13 are not included. In particular, classes 4, 5 and 7 are affected by multicollinearity, which arises when two or more variables are linearly related and there is a lack of orthogonality between them (for more information the reader is referred to Alin (2010)). In classes 4, 5 and 7 of this dataset we noticed that the values of some assets were identical. In addition, class 4 has several missing data, which prevents the application of the methodology. Moreover, class 7 has over 200 assets, making the application of the methodology computationally unfeasible. Also, classes 12 and 13 consist of only 3 and 1 assets respectively, which are datasets that are too small for the application of the methodology. The WCVaR and VaR are calculated at 10% confidence level. It can be seen that the values of WCVaR are higher than those of VaR for all classes.

When we compare between various classes, we can see that Communication Services has the highest WCVaR which means it is the most risky class, while Real Estate has the lowest WCVaR, meaning that it can be interpreted as less risky than the other classes. Also, it is noticed that classes 2, 3 and 11 which have the largest number of

assets have the smallest values of WCVaR. Therefore, the value of WCVaR decreases when the number of assets increases.

To conclude, applying WCVaR on a large dataset is very computationally demanding and requires a HPC. However, the WCVaR provides us with a more accurate estimation method than the traditional VaR to measure risks.

Class	Type of Asset	Number of assets	VaR	WCVaR
1	Energy	20	0.4749	0.6532
2	Materials	68	0.1702	0.2218
3	Industrials	52	0.3766	0.5712
6	Healthcare	9	1.1005	1.4873
8	Information Technology	5	1.0342	1.8063
9	Communication Services	13	1.3899	2.2619
10	Utilities	7	1.6091	2.1049
11	Real Estate	47	0.0128	0.0173

Table 5.3: Generating VaR and WCVaR using WCVaR technique for the GCC dataset.

5.10 Summary of the Chapter

In this chapter an advanced risk measure (WCVaR) was used to calculate the risk of portfolio in an accurate way. The WCVaR technique involves two main components, namely: copulas which allow us to calculate the dependence structure of assets, and time series analysis like GARCH models to analyse the marginals. The approach was applied firstly on a small dataset, including the prices of 6 assets collected over the period between 2009 and 2013, then on the large GCC dataset, including the prices of 500 assets collected over the period between 2008 and 2019. Then, we evaluated the performance of the WCVaR and compared it against the VaR measure. The results showed that the WCVaR outperformed VaR and measured risk more accurately. Moreover, it was found that including more assets in the portfolio leads to a decrease in the WCVaR value. On the other hand, the use of the WCVaR approach is in general computationally demanding, making it challenging to work with large datasets and requiring powerful computers like HPCs. The main advantage of the WCVaR method is that it measures risk more accurately compared to other traditional methods. Hence, this ap-

proach can be very beneficial to decision-makers, investors, financial institutions and banks, in their decision-making regarding minimising the overall risk of their portfolio investments.

In the next chapter, the content of the thesis is summarised and a list of potential further work is detailed.

Chapter 6

Conclusions and Future Work

This thesis is concerned with the portfolio optimisation problem subject to three different real-world constraints. We consider Markowitz's mean-variance portfolio optimisation model to find the efficient frontier which is the best trade-off between risk and return. We extend the standard mean-variance model to include cardinality constraints, minimum proportions and class constraints. We spot the differences that occur in the shape of efficient frontier curves when applying such constraints. The genetic algorithm is tested against five OR-library benchmark datasets, taken from five recognised worldwide indices.

This thesis can be outlined in three parts. The first part shows the results of applying a genetic algorithm to solve the problem of portfolio selection with different values of cardinality constraints and minimum proportions. In the second part, the results of using a genetic algorithm to solve the portfolio optimisation problem with class constraint and construct optimal portfolios on a Gulf Cooperation Council (GCC) dataset are investigated. In the third part of the thesis, we examine the WCVaR based on copula as an alternative risk measure. All applications of the developed GA were successfully performed on the high performance cluster at the University of Plymouth.

We provide the main results and future work on portfolio optimisation problems as follows.

6.1 Main Findings and Conclusions

In Chapter 3, new values of cardinality K (2, 3, 4), different to those mostly studied in the literature (5 and 10) for portfolio optimisation, were tested and analysed. The genetic algorithm was used to identify the optimal solutions for the benchmark datasets in the OR-library. The main reason of using a genetic algorithm was that it is easily parallelisable and may run on as many cores as desired. A comparison between the computed EF from the GA, Beasley's solution and the Cesarone CCEF (Cardinality Constrained Efficient Frontier) were given. Experimental results showed that increasing the value of K from 2 to 10 on datasets (D1)-(D5) from Beasley (2000) causes a decrease in the IGD measurements (i.e., the solution quality is higher) while the computational time increases significantly with K.

In the second part of the chapter, GA results assuming several distinct values of nonzero minimum proportions ($x_{min} = 0.05, 0.1, 0.15, 0.18, 0.2, 0.25, and 0.3$) were examined on the datasets (D1)-(D5). To the best of the author's knowledge, there is a literature gap in the case of high minimum proportions and also a gap in the use of metaheuristics approximating solutions for small *K*. The results showed that there was a large effect of increasing the minimum proportions on both IGD measurements and consumed time. It was found that increasing the minimum proportions leads to an increase in IGD measurements. A considerable increase in computational time also results when larger minimum proportions were used.

A comparison between the CCEFs produced with different values of x_{\min} and varying values of *K* plotted against the Cesarone et al. (2013) solution of minimum proportion 0.01 in different datasets was provided. It was found that at higher x_{\min} values, the GA solution moves further away from the Cesarone curve and the solution tends to be restricted to smaller portions of the CCEF, resulting in lower convergence and lower spread. The agreement with the CCEF of Cesarone becomes generally less close, which confirms that increasing the value of the minimum proportion may lead to increasing problem difficulty. Also, higher minimum proportions tend to lead to a restriction of the GA solution, and scaling issues, resulting in many points not being included in the relevant plot.

In the case of K = 10 and $x_{min} = 0.01$ (i.e., those used by many authors), the result of IGD measure was compared against work conducted by Jin et al. (2016), Maringer & Kellerer (2003), Ruiz-Torrubiano & Suárez (2010). It was shown that there was a close correspondence with the experimental results obtained by the other approaches.

In Chapter 4, the portfolio optimisation problem with class constraints was considered

6.1. MAIN FINDINGS AND CONCLUSIONS

for diversifying optimal portfolios, which were constructed to contain Gulf market stocks, oil and at least one precious metal, for the period between 2008 and 2018. A GA was then used on this data to provide approximations to the UEF (unconstrained efficient frontier) in each case. Empirical evidence demonstrates that diversifying portfolios with these commodities reduces the risk of the portfolios, applying Markowitz's meanvariance framework. The experiments in this chapter all used an ordered vector, or initial configuration (IC), comprising 5 classes from the total 13 available. Each time, classes 12 and 13 (oil and precious metals) were included. Using different values of K (20, 30, 40 and 50), experiments were conducted with random IC applied. The results show that a larger number of mixed portfolios appear at the top-right of the EF curve, and that random ICs produce more consistent computation time averages for all the K values, ranging from 80365.91 to 84426.52 seconds, compared with a fixed IC. It is also evident that an increase in the K value corresponds to an increase in the number of available assets, so that the approximation between the UEF by the GA-produced EF is considerably improved. More importantly, risk is shown to be reduced when stock assets are combined with oil and precious metals, creating further opportunities for investment and diversification of portfolios. Increasing the number of classes leads to a decrease in U_{IGD} (that is, there is closer agreement between the EF and UEF). However, an increase in computation time also occurs, which may be due to the higher number of combinations. Decreasing the number of classes, on the other hand, results in increased amounts of oil and metals, which will be of interest to investors.

Results were also examined for CCEFs on the GCC dataset, when not using class constraint. Good agreement with UEF was noted when K = 20, but less so when K = 10. The CCEF curve moved away from the UEF further at K = 5, so agreeing with the conclusion seen in Chapter 3. In simple terms, better returns and reduced risk were observed when portfolios contain 20 assets, in comparison to when they contain 5.

We ended the chapter with an investigation of shifting annual frontiers for each year of the GCC dataset. We compared the (quadratic) UEF with 5 classes with our computed GA solution, and presented the results for the years 2008, 2010, 2014 and 2018. In each case we used K = 20. It was evident that an upward or inward shift in a frontier

6.1. MAIN FINDINGS AND CONCLUSIONS

indicates that higher returns can be achieved while keeping risk levels the same, or that the same level of return may be achieved with less risk. It is possible to maintain the lower left level (i.e., risk is minimised but return is also minimised) while the frontiers tilt upwards or downwards. However, economic events can impact frontiers. The shifting frontier may be a reason for investors to constrain the time period for their portfolio selection, since time is a factor in portfolio weights and investors' risk preferences may also be affected. Including oil and metals in portfolios allowed a lower level of risk to be achieved, as demonstrated at the bottom of the curve. More assets were required to approximate the bottom of the UEF curve, whereas this was achieved with fewer assets towards the top. It was also found that using fewer classes, for example 3-6 classes per point, resulted in a risk increase, but far greater returns.

Including oil and metals in portfolios results in lower levels of risk, which makes their inclusion suitable for use as hedging and diversifying instruments. With the opportunities for expanded diversification, other commodities may be potentially added to obtain even better returns. Our recommendation is that portfolio managers diversify with oil and precious metals, taking into consideration the classes of assets which best fit their portfolio requirements.

In Chapter 5, we used an advanced risk measure, WCVaR (Worse Case Value-at-Risk) to calculate the risk of portfolio in an accurate way. The WCVaR technique involves two main components, namely: copulas which allow us to calculate the dependence structure of assets, and time series analysis, such as that with GARCH models, to analyse the marginals. Firstly, we applied this technique on a small dataset, including the prices of 6 assets collected over the period between 2009 and 2013, and then on the large GCC dataset, including the prices of 500 assets collected over the period between 2008 and 2019. Then, we evaluated the performance of the WCVaR and compared it against the VaR measure. The WCVaR and VaR for the GCC dataset were calculated at the 10% confidence level, the methodology applied separately to each class of the dataset. It was found that the values of WCVaR were higher than those of VaR for all classes. Thus, the WCVaR outperformed VaR and measured risk more accurately. Also, a larger numbers of assets yielded smaller values of WCVaR. Therefore, the value of WCVaR decreased when the number of assets increased.

6.2. FURTHER WORK

Using WCVaR on a big dataset was very computationally demanding and required time on the HPC cluster. However, the WCVaR provides a more accurate risk estimation than the traditional VaR.

6.2 Further Work

It is planned to publish the main works from Chapters 3, 4 and 5 in suitable high impact factor journals. Such journals include Journal of the Operational Research Society, European Journal of Operational Research, Applied Mathematics and Computation, and Energy Economics. Based on this thesis and its investigations, the research will be extended to include the following:

- Chapter 3 uses a genetic algorithm, but there are other evolutionary algorithms which may be investigated (e.g., particle swarm or ant colony optimisers). Also, local optima that arise during the evolution process could be further analysed and techniques be developed to assist a GA with overcoming them. In addition, visualisation of GA results and processes (e.g., via an extended multi-objective approach or a single-objective version in the style of the work of Walter et al. (2020)) may be a useful pursuit. Multi-objective approaches (e.g., the NSGA-II algorithm of Deb et al. (2002)) with suitably-extended portfolio optimisation problems are also another applicable area.
- For Chapter 4, there are several avenues of further work that may be pursued. First, there are other ways of dealing with classes besides the way they were dealt with in the chapter. For example, individual classes could be dealt with by themselves, that is, decompose the problem into smaller sub-problems. For example, if no assets are desired from classes 3, 5, 8 and 9 then the relevant sub-problems may be disregarded. Time-wise, this may be faster than the approach detailed in the chapter (since the matrices involved in the cost/fitness function are smaller), but the complexity of the algorithm and the combinations of sub-problems which contribute to the problem may be higher. Alternatively, a *surrogate* for a subproblem may be used also: if it is assumed that the correlations between assets in certain classes can be approximated by that of, say, 15-20 stocks, then this

could be factored into the calculation to reduce complexity. A pre-assignment constraint may also be possible, enabling oil and precious metals to be assigned to the portfolios no matter what the results.

Experimenting with the algorithms developed to allow for short-selling (investing without owning an asset by borrowing it) can also be considered. This requires traders to abide by certain budget constraint rules. Here it is assumed that there is an initial endowment of wealth and that we can short-sell (Bongini et al. 2002) up to a given proportion of remaining wealth, determined by how the marginal account is set. Also, it would be interesting to create efficient frontier for 5-year period to compare the dynamics of the efficient frontier for 1-year, 5-year, and 10-year.

Finally, the work of Graham & Craven (2021) could be used to shed light on the results of the GA and, in particular, the source of the apparent "bulge" of the GA-produced efficient frontier away from the quadratic shape of the unconstrained efficient frontier.

In Chapter 5 we considered only the t-copula, and it would be interesting to extend this work using other copula families or a mixture of copulas (Hu 2006, Kakouris & Rustem 2014). A mixture of copulas is combinations of a set of different types of copulas. They are a flexible method in the case the data has a lot of asymmetries and they consider various patterns of dependence in order to determine an appropriate measure of dependence to achieve robust results (Kakouris & Rustem 2014). Then, it would be interesting to compare the results of this method to that which were presented in this chapter. These new methods can also be applied to other datasets, like cryptocurrency datasets (Bitcoin, for example) or data in other fields (for example, environmental risk data such as floods or natural hazards). Additional suggestions include using different types of models for the marginals such as Generalized Additive Models (GAM) instead of GARCH models (Hastie & Tibshirani 1990, Coussement et al. 2010).

List of references

- Abanomey, W. S. & Mathur, I. (1999), 'The hedging benefits of commodity futures in international portfolio diversification', *The Journal of Alternative Investments* 2(3), 51–62.
- AghaKouchak, A., Sellars, S. & Sorooshian, S. (2013), 'Methods of tail dependence estimation', *Extremes in a Changing Climate* pp. 163–179.
- Ahmed, S., Çakmak, U. & Shapiro, A. (2007), 'Coherent risk measures in inventory problems', *European Journal of Operational Research* 182(1), 226–238.
- Akbay, M. A., Kalayci, C. B. & Polat, O. (2020), 'A parallel variable neighborhood search algorithm with quadratic programming for cardinality constrained portfolio optimization', *Knowledge-Based Systems* p. 105944.
- Alin, A. (2010), 'Multicollinearity', Wiley Interdisciplinary Reviews: Computational Statistics 2(3), 370–374.
- Alotaibi, T. S. & Craven, M. J. (2019), Efficient frontiers in portfolio optimisation with minimum proportion constraints, *in* 'Proceedings of the Genetic and Evolutionary Computation Conference', ACM, pp. 358–359.
- Anagnostopoulos, K. P. & Mamanis, G. (2010), 'A portfolio optimization model with three objectives and discrete variables', *Computers & Operations Research* 37(7), 1285–1297.
- Anagnostopoulos, K. P. & Mamanis, G. (2011*a*), 'The mean-variance cardinality constrained portfolio optimization problem: An experimental evaluation of five multiobjective evolutionary algorithms', *Expert Systems with Applications* **38**(11), 14208– 14217.

Anagnostopoulos, K. P. & Mamanis, G. (2011b), 'Multiobjective evolutionary algorithms

for complex portfolio optimization problems', *Computational Management Science* **8**(3), 259–279.

- Armañanzas, R. & Lozano, J. A. (2005), A multiobjective approach to the portfolio optimization problem, *in* 'Evolutionary Computation, 2005. The 2005 IEEE Congress on', Vol. 2, IEEE, pp. 1388–1395.
- Arouri, M. E. H., Bellalah, M. & Nguyen, D. K. (2011), 'Further evidence on the responses of stock prices in GCC countries to oil price shocks', *International Journal* of Business 16(1), 89.
- Arriaga, J. & Valenzuela-Rendón, M. (2012), Steepest ascent hill climbing for portfolio selection, *in* 'European Conference on the Applications of Evolutionary Computation', Springer, pp. 145–154.
- Artzner, P., Delbaen, F., Eber, J.-M. & Heath, D. (1999), 'Coherent measures of risk', *Mathematical finance* **9**(3), 203–228.
- Awartani, B., Maghyereh, A. I. & Al Shiab, M. (2013), 'Directional spillovers from the US and the Saudi market to equities in the gulf cooperation council countries', *Journal of International Financial Markets, Institutions and Money* **27**, 224–242.
- Baldridge, R. & Curry, B. (2022), 'Understanding the sharpe ratio', *available at https:* //www.forbes.com/advisor/investing/sharpe-ratio/. Date accessed: 02/08/2022.
- Banihashemi, S. & Navidi, S. (2017), 'Portfolio performance evaluation in mean-CVaR framework: A comparison with non-parametric methods value at risk in mean-var analysis', *Operations Research Perspectives* **4**, 21–28.
- Basu, S. (1996), 'Procyclical productivity: Increasing returns or cyclical utilization?', The Quarterly Journal of Economics 111(3), 719–751.
- Beasley, J. E. (2000), 'OR-Library', *available at http://people.brunel.ac.uk/* ~mastjjb/jeb/orlib/portinfo.html.

- Bedoui, R., Braeik, S., Goutte, S. & Guesmi, K. (2018), 'On the study of conditional dependence structure between oil, gold and USD exchange rates', *International Review of Financial Analysis* **59**, 134–146.
- Behr, P., Guettler, A. & Miebs, F. (2013), 'On portfolio optimization: Imposing the right constraints', *Journal of Banking & Finance* **37**(4), 1232–1242.
- Ben Mabrouk, A. (2020), 'Wavelet-based systematic risk estimation: Application on GCC stock markets: The Saudi Arabia case', *Quantitative Finance and Economics* 4(4), 542–595.
- Bernoulli, D. (1954), 'Exposition of a new theory on the measurement of risk', *Econo-metrica* 22(1), 23–36.
- Bertsimas, D., Lauprete, G. J. & Samarov, A. (2004), 'Shortfall as a risk measure: Properties, optimization and applications', *Journal of Economic Dynamics and Control* **28**(7), 1353–1381.
- Bhat, C. R. & Eluru, N. (2009), 'A copula-based approach to accommodate residential self-selection effects in travel behavior modeling', *Transportation Research Part B: Methodological* **43**(7), 749–765.
- Bjørnland, H. C. (2009), 'Oil price shocks and stock market booms in an oil exporting country', *Scottish journal of political economy* **56**(2), 232–254.
- Blitzstein, J. K. & Hwang, J. (2014), Introduction to Probability, CRC Press.
- Bollerslev, T. (1986), 'Generalized autoregressive conditional heteroskedasticity', *Journal of Econometrics* **31**(3), 307–327.
- Bongini, L., Degli Esposti, M., Giardina, C. & Schianchi, A. (2002), 'Portfolio optimization with short-selling and spin-glass', *The European Physical Journal B – Condensed Matter and Complex Systems* 27(2), 263–272.
- Bozorg-Haddad, O., Solgi, M. & Loáciga, H. A. (2017), Meta-Heuristic and Evolutionary Algorithms for Engineering Optimization, Vol. 294 of Wiley Series in Operations Research and Management Science, John Wiley & Sons.

- Brito, R. P. & Vicente, L. N. (2013), Efficient cardinality / mean-variance portfolios, *in* 'IFIP Conference on System Modeling and Optimization', Springer, pp. 52–73.
- Cesarone, F., Scozzari, A. & Tardella, F. (2013), 'A new method for mean-variance portfolio optimization with cardinality constraints', *Annals of Operations Research* **205**(1), 213–234.
- Chan, J. C. & Kroese, D. P. (2010), 'Efficient estimation of large portfolio loss probabilities in t-copula models', *European Journal of Operational Research* 205(2), 361– 367.
- Chang, T.-J., Meade, N., Beasley, J. E. & Sharaiha, Y. M. (2000), 'Heuristics for cardinality constrained portfolio optimisation', *Computers & Operations Research* 27(13), 1271–1302.
- Charfeddine, L. & Al Refai, H. (2019), 'Political tensions, stock market dependence and volatility spillover: Evidence from the recent intra-GCC crises', *The North American Journal of Economics and Finance* **50**, 101032.
- Chen, A. H., Liang, Y.-C. & Liu, C.-C. (2012), An artificial bee colony algorithm for the cardinality-constrained portfolio optimization problems, *in* '2012 IEEE Congress on Evolutionary Computation', IEEE, pp. 1–8.
- Chen, R. & Xu, J. (2019), 'Forecasting volatility and correlation between oil and gold prices using a novel multivariate gas model', *Energy Economics* **78**, 379–391.
- Cheng, G., Li, P. & Shi, P. (2007), 'A new algorithm based on copulas for var valuation with empirical calculations', *Theoretical Computer Science* **378**(2), 190–197.
- Cherubini, U. & Luciano, E. (2001), 'Value-at-risk trade-off and capital allocation with copulas', *Economic Notes* **30**(2), 235–256.
- Cherubini, U., Luciano, E. & Vecchiato, W. (2004), *Copula methods in finance*, John Wiley & Sons.
- Cheung, Y.-W. & Lai, K. S. (1995), 'Lag order and critical values of the augmented dickey–fuller test', *Journal of Business & Economic Statistics* **13**(3), 277–280.

- Chiodi, L., Mansini, R. & Speranza, M. G. (2003), 'Semi-absolute deviation rule for mutual funds portfolio selection', *Annals of Operations Research* **124**(1), 245–265.
- Chukwudum, Q. (2018), 'Extreme value theory and copulas: Reinsurance in the presence of dependent risks'.
- Concepcion Ausin, M. & Lopes, H. F. (2010), 'Time-varying joint distribution through copulas', *Computational Statistics & Data Analysis* **54**(11), 2383–2399.
- Corne, D. W., Jerram, N. R., Knowles, J. D. & Oates, M. J. (2001), PESA-II: Regionbased selection in evolutionary multiobjective optimization, *in* 'Proceedings of the 3rd Annual Conference on Genetic and Evolutionary Computation', pp. 283–290.
- Coussement, K., Benoit, D. F. & Van den Poel, D. (2010), 'Improved marketing decision making in a customer churn prediction context using generalized additive models', *Expert Systems with Applications* **37**(3), 2132–2143.
- Daníelsson, J. (2011), Financial Risk Forecasting: The Theory and Practice of Forecasting Market Risk with Implementation in R and Matlab, John Wiley & Sons.
- de Melo Mendes, B. V. & de Souza, R. M. (2004), 'Measuring financial risks with copulas', *International Review of Financial Analysis* **13**(1), 27–45.
- Deb, K., Pratap, A., Agarwal, S. & Meyarivan, T. (2002), 'A fast and elitist multiobjective genetic algorithm: NSGA-II', *IEEE Transactions on Evolutionary Computation* 6(2), 182–197.
- Deep, K. & Thakur, M. (2007), 'A new crossover operator for real coded genetic algorithms', *Applied Mathematics and Computation* **188**(1), 895–911.
- Demarta, S. & McNeil, A. J. (2005), 'The t copula and related copulas', *International Statistical Review* **73**(1), 111–129.
- DeMiguel, V., Garlappi, L. & Uppal, R. (2009), 'Optimal versus naive diversification: How inefficient is the 1/N portfolio strategy?', *The Review of Financial Studies* 22(5), 1915–1953.

- Di Gaspero, L., Di Tollo, G., Roli, A. & Schaerf, A. (2007), 'Hybrid local search for constrained financial portfolio selection problems', *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems* pp. 44–58.
- Di Gaspero, L., Di Tollo, G., Roli, A. & Schaerf, A. (2011), 'Hybrid metaheuristics for constrained portfolio selection problems', *Quantitative Finance* **11**(10), 1473–1487.
- Di Tollo, G. & Roli, A. (2008), 'Metaheuristics for the portfolio selection problem', *International Journal of Operations Research* **5**(1), 13–35.
- Drake, A. E. & Marks, R. E. (2002), 'Genetic algorithms in economics and finance: Forecasting stock market prices and foreign exchange - a review', *Genetic algorithms and genetic programming in computational finance* pp. 29–54.
- Embrechts, P., Frey, R. & McNeil, A. (2005), 'Quantitative risk management', *Princeton Series in Finance* **10**(4).
- Engle, R. F. (1982), 'Autoregressive conditional heteroscedasticity with estimates of the variance of United Kingdom inflation', *Econometrica: Journal of the Econometric Society* pp. 987–1007.
- Eshelman, L. J. & Schaffer, J. D. (1993), Real-coded genetic algorithms and intervalschemata, *in* 'Foundations of Genetic Algorithms', Vol. 2, Elsevier, pp. 187–202.
- Fabozzi, F. J., Fuss, R. & Kaiser, D. G. (2008), *The handbook of commodity investing*, John Wiley & Sons.
- Fama, E. F. & French, K. R. (1988), 'Permanent and temporary components of stock prices', *Journal of political Economy* **96**(2), 246–273.
- Farid, J. (2010), 'Value at risk VaR methods', available at https://financetrainingcourse.com/education/2010/03/ master-class-calculating-value-at-risk-var-var-methods/ . Date accessed: 10/04/2017.
- Fernández, A. & Gómez, S. (2007), 'Portfolio selection using neural networks', Computers & Operations Research 34(4), 1177–1191.

- Fieldsend, J. E., Matatko, J. & Peng, M. (2004), Cardinality constrained portfolio optimisation, *in* 'International Conference on Intelligent Data Engineering and Automated Learning', Springer, pp. 788–793.
- Fisher, I. (1906), The Nature of Capital and Income, The Macmillan Company.
- Fortenbery, T. R. & Hauser, R. J. (1990), 'Investment potential of agricultural futures contracts', *American Journal of Agricultural Economics* **72**(3), 721–726.
- Fortin, I. & Kuzmics, C. (2002), 'Tail-dependence in stock-return pairs', Intelligent Systems in Accounting, Finance & Management 11(2), 89–107.
- Frahm, G., Junker, M. & Szimayer, A. (2003), 'Elliptical copulas: Applicability and limitations', *Statistics & Probability Letters* 63(3), 275–286.
- Fryzlewicz, P. (2007), 'Lecture notes: Financial time series, ARCH and GARCH models', *University of Bristol*.
- Garey, M. R. & Johnson, D. S. (1979), *Computers and intractability*, Vol. 174, Freeman San Francisco.
- Goldberg, D. E. (1989), *Genetic Algorithms in Search, Optimization, and Machine Learning*, Addison-Wesley.
- Gorton, G. & Rouwenhorst, K. G. (2006), 'Facts and fantasies about commodity futures', *Financial Analysts Journal* **62**(2), 47–68.
- Graham, B. (2003), The Intelligent Investor, revised edn, Harper Collins.
- Graham, D. I. & Craven, M. J. (2021), 'An exact algorithm for small-cardinality constrained portfolio optimisation', *Journal of the Operational Research Society* 72(6), 1415–1431.
- Guijarro, F. (2018), 'A similarity measure for the cardinality constrained frontier in the mean–variance optimization model', *Journal of the operational research society* 69(6), 928–945.
- Gulpinar, N., An, L. T. H. & Moeini, M. (2010), 'Robust investment strategies with discrete asset choice constraints using dc programming', *Optimization* **59**(1), 45–62.

- Halmos, P. R. (2013), *Measure Theory*, Vol. 18 of *Graduate Texts in Mathematics*, Springer.
- Harper, D. (2017), 'An introduction to value at risk (var)', *available at http://www.investopedia.com/articles/04/092904.asp*. Date accessed: 10/04/2017.
- Hastie, T. & Tibshirani, R. (1990), *Generalized Additive Models*, Vol. 43 of *Monographs* on *Statistics and Applied Probability*, Chapman & Hall, CRC Press, Washington DC.
- Hochreiter, R. (2010), 'A note on evolutionary stochastic portfolio optimization and probabilistic constraints', *arXiv preprint arXiv:1001.5421*.
- Hoe, L. W., Hafizah, J. S., Zaidi, I. et al. (2010), 'An empirical comparison of different risk measures in portfolio optimization', *Business and Economic Horizons* 1(1), 39–45.
- Hooks, J. A. (1989), Inflation, inflation accounting, and real stock returns: An alternative test of the proxy hypothesis, Michigan State University.
- Hu, L. (2006), 'Dependence patterns across financial markets: A mixed copula approach', *Applied Financial Economics* **16**(10), 717–729.
- Huang, J.-J., Lee, K.-J., Liang, H. & Lin, W.-F. (2009), 'Estimating value at risk of portfolio by conditional copula-GARCH method', *Insurance: Mathematics and Economics* 45(3), 315–324.
- Jain, B. J., Pohlheim, H. & Wegener, J. (2001), On termination criteria of evolutionary algorithms, *in* 'Proceedings of the 3rd Annual Conference Companion on Genetic and Evolutionary Computation', Morgan Kaufmann Publishers (San Francisco, CA).
- Jensen, G. R., Johnson, R. R. & Mercer, J. M. (2000), 'Efficient use of commodity futures in diversified portfolios', *Journal of Futures Markets: Futures, Options, and Other Derivative Products* **20**(5), 489–506.
- Jin, Y., Qu, R. & Atkin, J. (2014), A population-based incremental learning method for constrained portfolio optimisation, *in* '2014 16th International Symposium on Symbolic and Numeric Algorithms for Scientific Computing', IEEE, pp. 212–219.

- Jin, Y., Qu, R. & Atkin, J. (2016), Constrained portfolio optimisation: The state-ofthe-art markowitz models, *in* 'Proceedings of the 5th International Conference on Operations Research and Enterprise Systems - Volume 1: ICORES', pp. 388–395.
- Joe, H. (1997), *Multivariate Models and Multivariate Dependence Concepts*, CRC Press.
- Joe, H. & Xu, J. J. (1996), 'The estimation method of inference functions for margins for multivariate models', available at https://open.library.ubc.ca/ soa/cIRcle/collections/facultyresearchandpublications/52383/ items/1.0225985.
- Jondeau, E. & Rockinger, M. (2006), 'The copula-GARCH model of conditional dependencies: An international stock market application', *Journal of International Money and Finance* **25**(5), 827–853.
- Jones, C. P. & Wilson, J. W. (2006), 'The impact of inflation measures on the real returns and risk of us stocks', *Financial Review* **41**(1), 77–94.
- Kakouris, I. & Rustem, B. (2014), 'Robust portfolio optimization with copulas', *European Journal of Operational Research* **235**(1), 28–37.
- Kalayci, C. B., Polat, O. & Akbay, M. A. (2020), 'An efficient hybrid metaheuristic algorithm for cardinality constrained portfolio optimization', *Swarm and Evolutionary Computation* **54**, 100662.
- Kenney, F. & Keeping, E. (1951), *Mathematics of Statistics Part Two*, D. Van Nostrand Company, Inc Princeton,; New Jersey; Toronto; New York; London.
- Kizys, R., Doering, J., Juan, A. A., Polat, O., Calvet, L. & Panadero, J. (2022), 'A simheuristic algorithm for the portfolio optimization problem with random returns and noisy covariances', *Computers & Operations Research* **139**, 105631.
- Kojadinovic, I., Yan, J. et al. (2010), 'Modeling multivariate distributions with continuous margins using the copula r package', *Journal of Statistical Software* **34**(9), 1–20.
- Kolm, P. N., Tütüncü, R. & Fabozzi, F. J. (2014), '60 years of portfolio optimization: Practical challenges and current trends', *European Journal of Operational Research* 234(2), 356–371.
- Konno, H., Shirakawa, H. & Yamazaki, H. (1993), 'A mean-absolute deviationskewness portfolio optimization model', *Annals of Operations Research* 45(1), 205– 220.
- Konno, H. & Suzuki, K.-I. (1995), 'A mean-variance-skewness portfolio optimization model', *Journal of the Operations Research Society of Japan* **38**(2), 173–187.
- Konno, H. & Yamazaki, H. (1991), 'Mean-absolute deviation portfolio optimization model and its applications to tokyo stock market', *Management science* 37(5), 519– 531.
- Krzemienowski, A. & Szymczyk, S. (2016), 'Portfolio optimization with a copulabased extension of conditional value-at-risk', *Annals of Operations Research* 237(1-2), 219–236.
- Li, D., Sun, X. & Wang, J. (2006), 'Optimal lot solution to cardinality constrained mean– variance formulation for portfolio selection', *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics* **16**(1), 83–101.
- Liagkouras, K. (2019), 'A new three-dimensional encoding multiobjective evolutionary algorithm with application to the portfolio optimization problem', *Knowledge-Based Systems* **163**, 186–203.
- Lim, S., Kim, M.-J. & Ahn, C. W. (2020), 'A genetic algorithm (ga) approach to the portfolio design based on market movements and asset valuations', *IEEE Access* 8, 140234–140249.
- Lin, C.-C. & Liu, Y.-T. (2008), 'Genetic algorithms for portfolio selection problems with minimum transaction lots', *European Journal of Operational Research* **185**(1), 393–404.
- Lin, L., Cao, L., Wang, J. & Zhang, C. (2004), 'The applications of genetic algorithms in stock market data mining optimisation', *Management Information Systems*.

- Liu, Y. & Qin, Z. (2012), 'Mean semi-absolute deviation model for uncertain portfolio optimization problem', *Journal of Uncertain Systems* **6**(4), 299–307.
- Looney, B. (2020), 'BP statistical review of world energy', *BP Statistical Review, London UK*.
- Luke, S. (2013), *Essentials of Metaheuristics*, second edn, Lulu. Available for free at http://cs.gmu.edu/~sean/book/metaheuristics/.
- Lwin, K., Qu, R. & Kendall, G. (2014), 'A learning-guided multi-objective evolutionary algorithm for constrained portfolio optimization', *Applied Soft Computing* 24, 757– 772.
- Maghyereh, A. & Al-Kandari, A. (2007), 'Oil prices and stock markets in GCC countries: New evidence from nonlinear cointegration analysis', *Managerial Finance*.
- Maghyereh, A. I., Awartani, B. & Tziogkidis, P. (2017), 'Volatility spillovers and crosshedging between gold, oil and equities: Evidence from the gulf cooperation council countries', *Energy Economics* **68**, 440–453.
- Mai, J.-F. & Scherer, M. (2017), Simulating Copulas: Stochastic Models, Sampling Algorithms, and Applications, Vol. 6 of Series in Quantitative Finance, World Scientific.
- Mansini, R. & Speranza, M. G. (1999), 'Heuristic algorithms for the portfolio selection problem with minimum transaction lots', *European Journal of Operational Research* **114**(2), 219–233.
- Marimoutou, V., Raggad, B. & Trabelsi, A. (2009), 'Extreme value theory and value at risk: Application to oil market', *Energy Economics* **31**(4), 519–530.
- Maringer, D. (2008), 'Heuristic optimization for portfolio management [application notes]', *IEEE Computational Intelligence Magazine* **3**(4), 31–34.
- Maringer, D. & Kellerer, H. (2003), 'Optimization of cardinality constrained portfolios with a hybrid local search algorithm', *Or Spectrum* **25**(4), 481–495.
- Markowitz, H. (1952), 'Portfolio selection', The Journal of Finance 7(1), 77–91.

- Martin, I. (2021), 'On the autocorrelation of the stock market', *Journal of Financial Econometrics* **19**(1), 39–52.
- Mensi, W., Hammoudeh, S. & Kang, S. H. (2015), 'Precious metals, cereal, oil and stock market linkages and portfolio risk management: Evidence from saudi arabia', *Economic Modelling* **51**, 340–358.
- Merton, R. C. (1980), 'On estimating the expected return on the market: An exploratory investigation', *Journal of financial economics* **8**(4), 323–361.
- Messaoud, S. B. & Aloui, C. (2015), 'Measuring risk of portfolio: GARCH-copula model', *Journal of Economic Integration* pp. 172–205.
- Mirović, V., Živkov, D. & Njegić, J. (2017), 'Construction of commodity portfolio and its hedge effectiveness gauging–revisiting dcc models.', *Finance a Uver: Czech Journal* of Economics & Finance 67(5).
- Mitchell, G. G., O'Donoghue, D., Barnes, D. & McCarville, M. (2003), 'Generepair – a repair operator for genetic algorithms', available at http://mural. maynoothuniversity.ie/10351/1/DO-GeneRepair-2003.pdf.
- Mitra, S. & Ji, T. (2010), 'Risk measures in quantitative finance', *International Journal of Business Continuity and Risk Management* **1**(2), 125–135.
- Nelsen, R. B. (2007), An Introduction to Copulas, Springer Science & Business Media.
- Nelson, D. B. (1992), 'Filtering and forecasting with misspecified ARCH models i: Getting the right variance with the wrong model', *Journal of econometrics* 52(1-2), 61– 90.
- Neumann, F. & Witt, C. (2013), Bioinspired computation in combinatorial optimization: Algorithms and their computational complexity, *in* 'Proceedings of the 15th Annual Conference Companion on Genetic and Evolutionary Computation', ACM, pp. 567– 590.
- Oh, K. J., Kim, T. Y., Min, S.-H. & Lee, H. Y. (2006), 'Portfolio algorithm based on portfolio beta using genetic algorithm', *Expert Systems with Applications* **30**(3), 527– 534.

- Ottaway, M. & Hamzawy, A. (2011), *Protest movements and political change in the Arab world*, Vol. 28, Carnegie Endowment for International Peace, Washington, DC.
- Owen, J. & Rabinovitch, R. (1983), 'On the class of elliptical distributions and their applications to the theory of portfolio choice', *The Journal of Finance* **38**(3), 745–752.
- Pai, G. V. & Michel, T. (2009), 'Evolutionary optimization of cconstrained k-means clustered assets for diversification in small portfolios', *IEEE Transactions on Evolutionary Computation* **13**(5), 1030–1053.
- Paquete, L., Schiavinotto, T. & Stützle, T. (2007), 'On local optima in multiobjective combinatorial optimization problems', *Annals of Operations Research* **156**(1), 83.
- Patton, A. J. (2004), 'On the out-of-sample importance of skewness and asymmetric dependence for asset allocation', *Journal of Financial Econometrics* **2**(1), 130–168.
- Puchinger, J., Raidl, G. R. & Pferschy, U. (2010), 'The multidimensional knapsack problem: Structure and algorithms', *INFORMS Journal on Computing* **22**(2), 250–265.
- Quaranta, A. G. & Zaffaroni, A. (2008), 'Robust optimization of conditional value at risk and portfolio selection', *Journal of Banking & Finance* **32**(10), 2046–2056.
- Raquel, C. R. & Naval Jr, P. C. (2005), An effective use of crowding distance in multiobjective particle swarm optimization, *in* 'Proceedings of the 7th Annual Conference on Genetic and Evolutionary Computation', ACM, pp. 257–264.
- Rockafellar, R. T. & Uryasev, S. (2002), 'Conditional value-at-risk for general loss distributions', *Journal of Banking & Finance* **26**(7), 1443–1471.
- Rockafellar, R. T., Uryasev, S. et al. (2000), 'Optimization of conditional value-at-risk', *Journal of risk* **2**, 21–42.
- Rubbaniy, G., Khalid, A. A., Syriopoulos, K. & Samitas, A. (2021), 'Safe-haven properties of soft commodities during times of covid-19', *Journal of Commodity Markets* p. 100223.

- Ruiz-Torrubiano, R. & Suárez, A. (2010), 'Hybrid approaches and dimensionality reduction for portfolio selection with cardinality constraints', *IEEE Computational Intelligence Magazine* 5(2), 92–107.
- Ruppert, D. (2011), *Statistics and Data Analysis for Financial Engineering*, Vol. 13 of *Springer Texts in Statistics*, Springer.
- Sabino da Silva, F. B. & Ziegelman, F. (2017), 'Robust portfolio optimization with multivariate copulas: A worst-case CVaR approach', *available at SSRN: http:* //dx.doi.org/10.2139/ssrn.3076283.
- Salahi, M., Mehrdoust, F. & Piri, F. (2013), 'CVaR robust mean-CVaR portfolio optimization', *International Scholarly Research Notices* **2013**.
- Salcedo-Sanz, S. (2009), 'A survey of repair methods used as constraint handling techniques in evolutionary algorithms', *Computer Science Review* **3**(3), 175–192.
- Samanipour, F. & Jelovica, J. (2020), 'Adaptive repair method for constraint handling in multi-objective genetic algorithm based on relationship between constraints and variables', *Applied Soft Computing* **90**, 106143.
- Satyanarayan, S. & Varangis, P. (1996), 'Diversification benefits of commodity assets in global portfolios', *The Journal of Investing* **5**(1), 69–78.
- Schaerf, A. (2002), 'Local search techniques for constrained portfolio selection problems', *Computational Economics* 20(3), 177–190.
- Scherer, B. & He, L. (2008), 'The diversification benefits of commodity futures indexes: A mean-variance spanning test', *The handbook of commodity investing* pp. 241–265.
- Schmidt, J. & Curry, B. (2022), 'How inflation erodes the value of your money', available
 at https://www.forbes.com/advisor/investing/what-is-inflation/
 . Date accessed: 10/08/2022.
- Shaw, D. X., Liu, S. & Kopman, L. (2008), 'Lagrangian relaxation procedure for cardinality-constrained portfolio optimization', *Optimisation Methods & Software* 23(3), 411–420.

- Shukla, A., Pandey, H. M. & Mehrotra, D. (2015), Comparative review of selection techniques in genetic algorithm, *in* 'Futuristic Trends on Computational Analysis and Knowledge Management (ABLAZE), 2015 International Conference on', IEEE, pp. 515–519.
- Simaan, Y. (1997), 'Estimation risk in portfolio selection: The mean variance model versus the mean absolute deviation model', *Management science* **43**(10), 1437–1446.
- Siriopoulos, C., Tsagkanos, A., Svingou, A. & Daskalopoulos, E. (2021), 'Foreign direct investment in GCC countries: The essential influence of governance and the adoption of ifrs', *Journal of Risk and Financial Management* **14**(6), 264.
- Skolpadungket, P., Dahal, K. & Harnpornchai, N. (2007), Portfolio optimization using multi-objective genetic algorithms, *in* 'Evolutionary Computation, 2007. CEC 2007. IEEE Congress on', IEEE, pp. 516–523.
- Smillie, A. (2008), New copula models in quantitative finance, PhD thesis, Tanaka Business Centre, Imperial College London.
- Statman, M. & Clark, J. (2013), 'End the charade: Replacing the efficient frontier with the efficient range', *Available at SSRN 2230548*.
- Su, C.-W., Sun, T., Ahmad, S. & Mirza, N. (2021), 'Does institutional quality and remittances inflow crowd-in private investment to avoid dutch disease? A case for emerging seven (E7) economies', *Resources Policy* 72, 102111.
- Szegö, G. (2005), 'Measures of risk', *European Journal of Operational Research* **163**(1), 5–19.
- Terpezan Tabără, O. A. (n. d.), 'The importance of value at risk method in the management of banking risk', *available at http://www.asecu.gr/files/ RomaniaProceedings/64.pdf*. Date accessed: 17/03/2017.
- Tobin, J. (1958), 'Liquidity preference as behavior towards risk', *The Review of Economic Studies* **25**(2), 65–86.

- Trivedi, P. K., Zimmer, D. M. et al. (2007), 'Copula modeling: An introduction for practitioners', *Foundations and Trends in Econometrics* **1**(1), 1–111.
- Umar, M., Mirza, N., Rizvi, S. K. A. & Furqan, M. (2021), 'Asymmetric volatility structure of equity returns: Evidence from an emerging market', *The Quarterly Review of Economics and Finance*.
- Walter, M. J., Walker, D. J. & Craven, M. J. (2020), Visualising evolution history in multiand many-objective optimisation, *in* 'International Conference on Parallel Problem Solving from Nature', Springer, pp. 299–312.
- Wang, C.-P., Huang, H.-H. & Jou, D. G. (2011), 'Dynamic portfolio frontier in a mean– variance framework', *Applied Financial Economics* 21(17), 1255–1261.
- Wang, L., Ahmad, F., Luo, G.-I., Umar, M. & Kirikkaleli, D. (2021), 'Portfolio optimization of financial commodities with energy futures', *Annals of Operations Research* pp. 1–39.
- Welsh, W. (1999), 'On the reliability of cross-correlation function lag determinations in active galactic nuclei', *Publications of the Astronomical Society of the Pacific* **111**(765), 1347.
- Werner, F. & Sotskov, Y. N. (2006), *Mathematics of Economics and Business*, Routledge.
- Whitley, D. (2001), 'An overview of evolutionary algorithms: Practical issues and common pitfalls', *Information and Software Technology* **43**(14), 817–831.
- Woodside-Oriakhi, M., Lucas, C. & Beasley, J. E. (2011), 'Heuristic algorithms for the cardinality constrained efficient frontier', *European Journal of Operational Research* 213(3), 538–550.
- Xu, R.-T., Zhang, J., Liu, O. & Huang, R.-Z. (2010), An estimation of distribution algorithm based portfolio selection approach, *in* '2010 International Conference on Technologies and Applications of Artificial Intelligence', IEEE, pp. 305–313.
- Yan, J. et al. (2007), 'Enjoy the joy of copulas: With a package copula', *Journal of Statistical Software* **21**(4), 1–21.

- Young, M. R. (1998), 'A minimax portfolio selection rule with linear programming solution', *Management Science* **44**(5), 673–683.
- Zaremba, A. (2015), 'Portfolio diversification with commodities in times of financialization', *International Journal of Finance & Banking Studies* **4**(1), 18–36.
- Zhu, S. & Fukushima, M. (2009), 'Worst-case conditional value-at-risk with application to robust portfolio management', *Operations research* **57**(5), 1155–1168.
- Zitzler, E., Laumanns, M. & Thiele, L. (2001), 'SPEA2: Improving the strength pareto evolutionary algorithm', *TIK-report* **103**.