NTNU – Trondheim
Norwegian University of
Science and Technology

Tramp shipping in the LNG trade

Halfdan Bondevik
Steinar Halland Fretheim

Industrial Economics and Technology Management
Submission date: June 2014
Supervisor: Rudolf Egging, IØT
Co-supervisor: Henrik Andersson, IØT

Norwegian University of Science and Technology
Department of Industrial Economics and Technology Management

NTNU
Det skapende universitet

MASTERKONTRAKT
- uttak av masteroppgave

1. Studentens personalia

Etternavn, fornavn	Fødselsdato
Bondevik, Halfdan	**10. aug 1987**
E-post	Telefon
halfdan.bondevik@gmail.com	**40220350**

2. Studieopplysninger

Fakultet	
Fakultet for samfunnsvitenskap og teknologiledelse	
Institutt	
Institutt for industriell økonomi og teknologiledelse	
Studieprogram	Hovedprofil
Industriell økonomi og teknologiledelse	**Anvendt økonomi og optimering**

3. Masteroppgave

Oppstartsdato	Innleveringsfrist
15. jan 2014	**11. jun 2014**
Oppgavens (foreløpige) tittel **Tramp shipping in the LNG trade**	
Oppgavetekst/Problembeskrivelse The purpose is to develop a stochastic optimization model that makes optimal movement and trading decisions for LNG vessels that are contracting LNG on a speculative basis. The decision-support model presented gives the user an opportunity to make new decisions each time updated information is revealed. The results from the stochastic model are compared with a deterministic model in order to decide its usefulness. The model run time is also considered. Main contents: 1. Description of the LNG market 2. Literature review including previous work on similar problems 2. Description of the problem 3. Development of mathematical model and methods for solving the problem 4. Implementation of mathematical model and methods 5. Testing of model and methods with relevant data 6. Discussion of the results and the usefulness of the model and methods	
Hovedveileder ved institutt **Førsteamanuensis Rudolf Egging**	Medveileder(e) ved institutt **Henrik Andersson**
Merknader **1 uke ekstra p.g.a påske.**	

4. Underskrift

Student: Jeg erklærer herved at jeg har satt meg inn i gjeldende bestemmelser for mastergradsstudiet og at jeg oppfyller kravene for adgang til å påbegynne oppgaven, herunder eventuelle praksiskrav.

Partene er gjort kjent med avtalens vilkår, samt kapitlene i studiehåndboken om generelle regler og aktuell studieplan for masterstudiet.

TRONDHEIM, 30. april 2014

Sted og dato

Halfdan Bonellie

Student

Hovedveileder

R G EGGING

Originalen lagres i NTNUs elektroniske arkiv. Kopi av avtalen sendes til instituttet og studenten.

NTNU
Det skapende universitet

MASTERKONTRAKT

- uttak av masteroppgave

1. Studentens personalia

Etternavn, fornavn **Fretheim, Steinar Halland**	Fødselsdato **19. jul 1989**
E-post **steinarhf@gmail.com**	Telefon **95228089**

2. Studieopplysninger

Fakultet **Fakultet for samfunnsvitenskap og teknologiledelse**	
Institutt **Institutt for industriell økonomi og teknologiledelse**	
Studieprogram **Industriell økonomi og teknologiledelse**	Hovedprofil **Anvendt økonomi og optimering**

3. Masteroppgave

Oppstartsdato **15. jan 2014**	Innleveringsfrist **11. jun 2014**
Oppgavens (foreløpige) tittel **Tramp shipping in the LNG trade**	

Oppgavetekst/Problembeskrivelse

The purpose is to develop a stochastic optimization model that makes optimal movement and trading decisions for LNG vessels that are contracting LNG on a speculative basis. The decision-support model presented gives the user an opportunity to make new decisions each time updated information is revealed. The results from the stochastic model are compared with a deterministic model in order to decide its usefulness. The model run time is also considered.

Main contents:
1. Description of the LNG market
2. Literature review including previous work on similar problems
2. Description of the problem
3. Development of mathematical model and methods for solving the problem
4. Implementation of mathematical model and methods
5. Testing of model and methods with relevant data
6. Discussion of the results and the usefulness of the model and methods

Hovedveileder ved institutt **Førsteamanuensis Rudolf Egging**	Medveileder(e) ved institutt **Henrik Andersson**
Merknader **1 uke ekstra p.g.a påske.**	

4. Underskrift

Student: Jeg erklærer herved at jeg har satt meg inn i gjeldende bestemmelser for mastergradsstudiet og at jeg oppfyller kravene for adgang til å påbegynne oppgaven, herunder eventuelle praksiskrav.

Partene er gjort kjent med avtalens vilkår, samt kapitlene i studiehåndboken om generelle regler og aktuell studieplan for masterstudiet.

Trondheim 30/4-14

Sted og dato

Stein H. Sefhin

Student

Hovedveileder

R G EGGING

Originalen lagres i NTNUs elektroniske arkiv. Kopi av avtalen sendes til instituttet og studenten.

NTNU
Det skapende universitet

SAMARBEIDSKONTRAKT

1. Studenter i samarbeidsgruppen

Etternavn, fornavn	Fødselsdato
Bondevik, Halfdan	**10. aug 1987**
Etternavn, fornavn	Fødselsdato
Fretheim, Steinar Halland	**19. jul 1989**

2. Hovedveileder

Etternavn, fornavn	Institutt
Egging, Rudolf	**Institutt for industriell økonomi og teknologiledelse**

3. Masteroppgave

Oppgavens (foreløpige) tittel
Tramp shipping in the LNG trade

4. Bedømmelse

Kandidatene skal ha *individuell* bedømmelse ☐
Kandidatene skal ha *felles* bedømmelse ☒

TRONDHEIM, 30. april 2014
Sted og dato

Halfdan Bondevik
Halfdan Bondevik

Hovedveileder

Steinar H. Fretheim
Steinar Halland Fretheim

Originalen oppbevares på instituttet.

Preface

This thesis is written as the concluding part of a master's degree in Managerial Economics and Operations Research at the Department of Industrial Economics and Technology Management at the Norwegian University of Science and Technology (NTNU). The thesis is a continuation of our project work in the fall of 2013.

The thesis is written within the field of operations research. We present two decision-support tools for speculative trading and transporting of Liquefied Natural Gas (LNG).

We would like to thank our supervisors Ruud Egging and Henrik Andersson for interesting discussions and valuable feedback.

Halfdan Bondevik Steinar H. Fretheim

Trondheim, June 4, 2014

Abstract

LNG is to an increasing extent traded through spot markets and short-term contracts. At the same time, gas production is growing and the number of importers and exporters of LNG is increasing. These trends lead to greater opportunities for actors who are looking into speculative trading of LNG.

In this thesis, we develop two stochastic optimization models for buying, transporting and selling LNG in the spot market. We take the perspective of an actor that owns LNG vessels and does speculative trading. The objective is to maximize profit. This is done by making optimal movement and trade decisions. Income is generated by buying and selling LNG. Costs relate to operating the LNG vessel. The models make a trade-off between maximizing revenue and minimizing cost. The price processes in the ports are stochastic. We use scenarios to represent an approximation of the price development process.

The models make use of a dynamic program to estimate the value of potential trade sequences. Two stochastic models are run in combination with the dynamic program in order to make movement and trade decisions. One is a mixed-integer program (MIP) that is run by commercial optimization software (Xpress-Mosel). The other is a heuristic written in Java. We present solutions for both deterministic and stochastic test instances. The stochastic solution takes uncertainty into consideration and presents the decisions that are best hedged against all outcomes of price development.

Our main focus is on comparing the stochastic and deterministic versions of the two models, in order to identify the solution approach that best solves our problem. The stochastic versions are found to provide better solutions than the deterministic ones. This goes for both models. The heuristic solution outperforms the MIP when considering both profit, run time and stability.

With minor adjustments the models can be used as real life decision tools.

Sammendrag

Produksjonen av gass er økende, og Liquefied Natural Gas (LNG) utgjør en stadig større andel av global gasshandel. 31.7% av all gass som ble handlet i 2012 ble solgt som LNG. LNG handles i økende grad gjennom spothandel og korttidskontrakter, og antallet importører og eksportører er voksende. Dette er faktorer som leder til økende muligheter for aktører som bedriver spekulativ handel med LNG.

I denne avhandlingen utvikler vi to stokastiske optimeringsmodeller for kjøp, transport og salg av LNG i spotmarkedet. Vi tar utgangspunkt i en aktør som eier LNG-skip og bedriver spekulativ handel. Målet er å maksimere profitt. Dette oppnås gjennom å gjøre optimale handels- og bevegelsesbeslutninger. Inntekt genereres ved å kjøpe og selge LNG. Kostnader relaterer til å operere LNG-skipet. Modellene gjør en avveiing mellom maksimering av inntekt og minimering av kostander. Prisprosessen i havnene er stokastisk. Vi bruker scenarioer til å approksimere prisprosessen. Vår fokus er på å sammmenligne modellene med hverandre, med hensikt å finne ut hvilken som best løser vårt problem.

Modellene benytter seg av et dynamisk program til å estimere verden av potensielle handelssekvenser, basert på prisprognoser. The dynamiske programmet er implementert i Java. To stokastiske modeller brukes i kombinasjon med det dynamiske programmet for å finne optimale valg for handelstidspunkt og forflyning. Det ene modellen er et mixed-integer program (MIP) som blir kjørt i kommersiell optimeringssoftware (Xpress-Mosel). Den andre er en heuristikk som er skrevet i Java. Vi presenterer løsninger for både deterministiske og stokastiske testinstanser. Den stokastiske løsningen forholder seg til usikkerhet og presenterer de beslutningene som gir best utgangspunkt for alle utfall av prisutviklingen.

Hovedfokuset vårt er på å sammenligne den deterministiske og stokastiske versjonen av de to modellene, for å finne den fremgangsmåten som best løser vårt problem. De stokastiske versjonene gir generelt bedre resultat enn de deterministiske. Dette gjelder begge modellene. Heuristikken er bedre enn MIPen på både profitt, kjøretid og stabilitet. Modellene kan brukes som beslutningsverktøy hvis det gjøres små justeringer.

Contents

1 Introduction **1**

2 Natural Gas Production and Trade of Liquefied Natural Gas **3**

2.1 The LNG value chain . . . 4
2.2 Geographical markets . . . 5
2.3 Supply of natural gas . . . 6
2.4 Demand for natural gas . . . 9
2.5 LNG market characteristics . . . 10
2.5.1 Infrastructure, integration and diversification . . . 10
2.5.2 Contracts . . . 11
2.5.3 Gas market liberalization . . . 13
2.5.4 Pricing . . . 14
2.6 Technological advances in LNG production . . . 18
2.7 LNG shipping . . . 19
2.7.1 The LNG fleet . . . 19
2.7.2 Transportation costs . . . 20

3 Literature Review **23**

3.1 LNG . . . 23
3.2 Tramp shipping . . . 25
3.3 Stochastic modeling . . . 27
3.4 Stochastic dynamic programming . . . 28
3.5 Our work . . . 29

4 Problem Description **31**

5 Models **33**

5.1 Assumptions and modelling choices . . . 33
5.1.1 Grid . . . 33
5.1.2 Time . . . 35

5.1.3 Sailing cost . . . 35
5.1.4 Ports . . . 36
5.1.5 Trade . . . 36
5.2 Dynamic programming to find port values . . . 37
5.2.1 General description . . . 37
5.2.2 Labels . . . 37
5.2.3 Domination . . . 38
5.2.4 Length of trade horizon . . . 40
5.2.5 End-of-horizon values . . . 41
5.3 Mixed Integer Program (MIP) . . . 43
5.3.1 General description . . . 43
5.3.2 Mathematical model . . . 43
5.4 Heuristic . . . 48
5.4.1 General description . . . 48
5.4.2 Calculating the value of possible moves . . . 48
5.4.3 Stochastic vs deterministic . . . 50

6 Solution Approaches **57**
6.1 Interaction between the models . . . 57
6.2 Dynamic Program and Mixed Integer Program . . . 59
6.3 Dynamic Program and heuristic . . . 60

7 Implementation **63**
7.1 Models . . . 63
7.1.1 Dynamic Program . . . 63
7.1.2 Mixed Integer Program . . . 64
7.1.3 Heuristic . . . 65
7.2 Solution approaches . . . 65
7.2.1 Dynamic Program and Mixed Integer Program . . . 65
7.2.2 Dynamic Program and heuristic . . . 66

8 Test Instances **69**
8.1 Map . . . 69
8.2 Ornstein-Uhlenbeck price process . . . 71
8.3 Fixed parameters . . . 74
8.3.1 Parameters used for price forecasting . . . 74
8.3.2 Number of days in total problem period . . . 76
8.3.3 Price development . . . 76
8.3.4 Size of vessel . . . 77

8.3.5 Basis cost and travelling costs 77
8.3.6 Starting conditions 77
8.3.7 Discounting factor 77
8.3.8 Nonanticipativity constraints in the MIP 78
8.3.9 Maximum run time of MIP 78
8.4 Variable parameters 78
8.4.1 Number of scenarios 79
8.4.2 End of horizon values 79
8.4.3 Rolling horizon length 79
8.4.4 Gap in MIP 80
8.4.5 Binary days in MIP 81
8.4.6 Ports included in heuristic move calculation 82
8.5 Stability testing 82
8.5.1 Determining level of stability 83
8.5.2 Scenarios used in the stability testing 84
8.5.3 Stability situations for the vessel 86
8.6 Base case 87

9 Results and Discussion **89**
9.1 Test of Dynamic Program 90
9.2 Stochastic models 93
9.2.1 Test of MIP 93
9.2.2 Test of heuristic 99
9.2.3 Comparison between MIP and heuristic 104
9.3 Deterministic models 107
9.3.1 Test of MIP 107
9.3.2 Test of heuristic 108
9.3.3 Comparison between MIP and heuristic 108
9.4 Comparison between stochastic and deterministic models 110
9.4.1 MIP 110
9.4.2 Heuristic 110
9.4.3 Route considerations 111

10 Concluding Remarks and Further Research **115**
10.1 Conclusion 116
10.2 Further research 116

Bibliography **119**

Appendix A Result Tables **127**

Appendix B Approximate Conversion Table **134**

Appendix C Stand-alone Mathematical Model **135**

List of Tables

2.1 Average operating costs for a 138,000 m^3 LNG vessel 20

5.1 Number of labels generated as a function of the number of trades 39
5.2 Ports considered for each possible new position 51
5.3 Forecasted port values in thousand \$ for port 1 and port 2 in day 4 and 5 51
5.4 Potential value contribution in thousand \$ from port 1 for the different positions 52
5.5 Potential value contribution in thousand \$ from port 2 for the different positions 52
5.6 Position values in thousand \$ 52
5.7 Move valuesin thousand \$ 52
5.8 Action values in thousand \$ for each of the potential moves for stochastic version 54

8.1 Travel time between ports (in days) 71
8.2 Cost of travelling between ports (in million \$) 71
8.3 Values of parameters used in the Ornstein-Uhlenbeck process 75
8.4 Values used for trading in stability testing 84
8.5 Numerical example of stability test 85
8.6 Values for fixed parameters of the base case 87
8.7 Values for variable parameters of the base case 87

9.1 Value of stochastic solutions in million \$ for a varying number of scenarios 112

A.1 Profit in million \$ and run time without EOH-values when varying rolling horizon length in the DP 127
A.2 Profit in million \$ and run time with EOH-values when varying rolling horizon length in the DP 128
A.3 Profit in million \$ and run time when varying rolling horizon length and use of EOH-values in the DP 128

A.4 Profit in million \$ and run time with decreasing horizon when varying the number of scenarios in the stochastic version of the MIP 128
A.5 Profit in million \$ and run time when varying the number of scenarios and type of rolling horizon in the stochastic version of the MIP 129
A.6 Profit in million \$ when varying the number of binary days in the stochastic version of the MIP 129
A.7 Profit in million \$ and run time when varying gap value in the stochastic version of the MIP 129
A.8 Overall in- and out-of-sample values when varying the number of scenarios in the stochastic version of the MIP 129
A.9 In-sample values for different situations when varying the number of scenarios in the stochastic version of the MIP 130
A.10 Out-of-sample values for different situations when varying the number of scenarios in the stochastic version of the MIP 130
A.11 Profit in million \$ and run time when varying the number of scenarios and port inclusion in the stochastic version of the heuristic 130
A.12 Profit and run time with decreasing horizon when varying the number of scenarios in the stochastic version of the heuristic 130
A.13 Profit in million \$ and run time when varying the number of scenarios and type of rolling horizon in the stochastic version of the heuristic 131
A.14 Overall in- and out-of-sample values when varying the number of scenarios in the stochastic version of the heuristic 131
A.15 In-sample values for different situations when varying the number of scenarios in the stochastic version of the heuristic 131
A.16 Out-of-sample values for different situations when varying the number of scenarios in the stochastic version of the heuristic 131
A.17 Comparison of profit and run time between the stochastic versions of the MIP and heuristic when varying the number of scenarios 132
A.18 Comparison of stability between the stochastic versions of the MIP and heuristic when varying the number of scenarios 132
A.19 Profit in million \$ and run time in the deterministic version of the MIP . 132
A.20 Profit in million \$ and run time in the deterministic version of the heuristic 132
A.21 Comparison of profit and run time between the deterministic versions of the MIP and heuristic 133
A.22 Comparison of profit and run time between the stochastic and deterministic versions of the MIP 133
A.23 Comparison of profit and run time between the stochastic and deterministic versions of the heuristic 133

List of Figures

2.1 LNG value chain 4
2.2 Transportation cost per cubic meter of gas as a function of distance for onshore pipeline, offshore pipeline and LNG 5
2.3 Global LNG trade movements in 2012 6
2.4 Projection of global export/import of LNG in different markets/country groupings, 1990-2030 6
2.5 Historical and projected global natural gas production, 2010-2040 7
2.6 Estimate of total global conventional and unconventional natural gas reserves 8
2.7 Historical and projected global shale gas production, 1990-2030 9
2.8 Historical and projected number of liquefaction plants and regasification terminals, 1990-2020 10
2.9 Average number of suppliers per importer and customers per exporter in LNG trade, 1991-2011 11
2.10 Share of short-term vs long-term contracts 12
2.11 Amount of LNG traded through dedicated, flexible and spot contracts in 2008 13
2.12 Historical price development for Brent oil and gas in Europe, 1999-2009 . 15
2.13 Historical and projected average Henry Hub natural gas spot price 16
2.14 Historical LNG prices in different countries 17
2.15 Hypothetical netback to Nigeria from the U.S. Gulf and the U.K. 17
2.16 Hypothetical netback to Qatar from the U.S. Gulf and from Japan 18
2.17 Membrane LNG tanker 19
2.18 Self supporting LNG tanker 20

5.1 Example grid with one vessel (V), one buy port (B) and one sell port (S) 34
5.2 Allowed moves in the grid 34
5.3 Allowed moves in real world 34
5.4 The decision made at time t takes place between t and $t+1$. 35
5.5 Example label from the DP 38
5.6 Example of label extension in the DP 39

5.7 Example of domination of labels in the DP 40
5.8 Values in a port as a function of remaining days in the trade horizon . . 42
5.9 Straight and diagonal neighbors (green) relative to the vessel (orange) . . 44
5.10 Timing of nonanticipativity constraints for position and trade variables . 47
5.11 Overview of which day the position variables close to the vessel (orange) become binary 47
5.12 Possible options (triangles) for the vessel (orange square) 49
5.13 Area where ports are considered when including ports with fewer travel days 49
5.14 Area where ports are considered when including ports with the same or fewer travel days 50
5.15 Example of ports considered for a northeast move in Atlantic Basin-map. Potential ports are dark blue. Non-potential ports are light blue. 50
5.16 Example with two ports and three possible move options 51
5.17 Example with three ports and two possible options of moving 53
5.18 Values of going west or east in different scenarios 54
5.19 Value of going east in stochastic problem 55

6.1 Graphical overview over the rolling horizon method 58
6.2 Example of optimal route for starting trade in port 1 on day 4, found by the DP 59
6.3 Interaction between MIP and DP. The vessel is orange, potential ports are dark blue and non-potential ports are light blue. 61
6.4 Flow diagram for heuristic approach 62

8.1 Ports used in test instance. Buy ports are red and sell ports are blue. . . 70
8.2 Map fitted to grid. Buy ports are red and sell ports are blue. 70
8.3 Correlation coefficient as a function of distance 74
8.4 Price processes with starting price 5 and varying correlations 75
8.5 Price development for all ports used in the test instances 76
8.6 Constraints from the MIP that are affected by the number of binary days 81
8.7 Values used for moves in stability testing 83
8.8 Map used in numerical example of stability test 85
8.9 Overview over scenarios used for in-sample tests and out-of-sample tests . 86

9.1 Profit and run time without EOH-values when varying rolling horizon length in the DP 90
9.2 Profit and run time with EOH-values when varying rolling horizon length in the DP 91
9.3 Profit and run time when varying rolling horizon length and use of EOH-values in the DP 91

9.4 Profit and run time with decreasing horizon when varying the number of scenarios in the stochastic version of the MIP 93
9.5 Profit and run time when varying the number of scenarios and type of rolling horizon in the stochastic version of the MIP 94
9.6 Profit and run time when varying the number of binary days in the stochastic version of the MIP 95
9.7 Profit and run time when varying gap value in the stochastic version of the MIP 96
9.8 Overall in- and out-of-sample values when varying the number of scenarios in the stochastic version of the MIP 97
9.9 In-sample values for different situations when varying the number of scenarios in the stochastic version of the MIP 98
9.10 Out-of-sample values for different situations when varying the number of scenarios in the stochastic version of the MIP 99
9.11 Profit and run time when varying the number of scenarios and port inclusion in the stochastic version of the heuristic 100
9.12 Profit and run time with decreasing horizon when varying the number of scenarios in the stochastic version of the heuristic 100
9.13 Profit and run time when varying number of scenarios and type of rolling horizon in the stochastic version of the heuristic 101
9.14 Overall in- and out-of-sample values when varying the number of scenarios in the stochastic version of the heuristic 102
9.15 In-sample values for different situations when varying the number of scenarios in the stochastic version of the heuristic 103
9.16 Out-of-sample values for different situations when varying the number of scenarios in the stochastic version of the heuristic 104
9.17 Comparison of profit and run time between the stochastic versions of the MIP and heuristic when varying the number of scenarios 105
9.18 Comparison of stability between the stochastic versions of the MIP and heuristic when varying the number of scenarios 106
9.19 Profit and run time in the deterministic version of the MIP 107
9.20 Profit and run time in the deterministic version of the heuristic 108
9.21 Comparison of profit and run time between the deterministic versions of the MIP and heuristic 109
9.22 Comparison of profit and run time between the stochastic and deterministic versions of the MIP 110
9.23 Comparison of profit and run time between the stochastic and deterministic versions of the heuristic 111

9.24 Trades conducted by the vessel sailing the route with highest profit . . . 113

B.1 Conversion table 134

List Of Abbreviations

p.a.	per annum = per year
bcm	billion cubic meters
bcma	billion cubic meters per annum
Btu	British thermal unit
MMBtu	million British thermal units ($\approx$ 28.3 cubic meters of gas)
t	tons
Mt	Million tons
ADP	Annual Delivery Program
IRP	Inventory Routing Problem
VSS	Value of stochastic solution

A table with conversion factors for gas units is found in Appendix B.

Chapter 1

Introduction

Natural gas accounted for 23.9% of global primary energy consumption in 2012 and is one of the fastest growing energy sources in the world (BP, 2013b). Liquefied Natural Gas (LNG) is natural gas that has been cooled down until turning liquid, effectively reducing the volume to 1/600 of its gaseous state. This makes it viable for long-distance transport by sea. Liquefied Natural Gas' (LNG's) share of global gas trade was 31.7% in 2012. The rest of the gas was sold through pipelines (BP, 2013b). Trading through pipelines requires the producer and end consumer to be connected to the same pipeline network. For many markets this is not a viable option, given the high cost of producing long distance pipelines. The flexibility of LNG trade is also a major advantage over pipelines. Pipelines are fixed and have a limited capacity, while LNG vessels can be redirected to where they are currently needed.

LNG markets have changed in recent years, shifting from predominantly long-term dedicated contracts to an increased use of flexible contracts and spot trade. The emerging spot markets opens up opportunities for speculative traders of LNG.

In this thesis we take the perspective of an actor that owns LNG vessels and does speculative trading of LNG in the spot market. This includes buying, transporting and selling LNG. The goal is to maximize profit.

The closest real-life example that we know of is Golar LNG, who contracted tonnage on a speculative basis in 2002. Finding ship employment proved difficult and they ended up converting vessels into storage systems and regasification terminals (Engelen and Dullaert, 2010a). The LNG market has changed a lot since then, and we believe there are greater opportunities to be successful now.

The purpose of this thesis is to present a stochastic optimization model that makes optimal movement and trading decisions for LNG vessels that are contracting LNG on a speculative basis. This is done by dynamically positioning empty vessels close to ports where prices are low or expected to become so, and then moving in to buy when the time is right. The full vessel is then dynamically positioned close to ports where prices are high or expected to become so, before moving in to sell when the time is right. The costs related to operating LNG vessels are also considered. We seek to optimize the positioning and the timing of trade under price uncertainty. The decision-support models presented gives the user an opportunity to make new decisions each time updated price information is revealed.

The paper is organized as follows. First we present the gas industry, looking at production of natural gas and trade of liquefied natural gas. We then present relevant literature, looking into previous optimization work done on LNG, tramp shipping, stochastic modelling and stochastic dynamic programming. Following this we describe our problem, before presenting three models that are used to solve the problem. Two complete solution approaches that makes use of these models are then presented. An outline of how the models are implemented is presented, before looking at the instances used to test the model. Results from running the tests are then discussed, before we conclude and make suggestions for further research.

Chapter 2

Natural Gas Production and Trade of Liquefied Natural Gas

Projections indicate that global gas production will grow by 2% p.a. (per year) running up to 2030. Liquefied Natural Gas (LNG) is predicted to play an increasingly important role, with production growing at 4.3% p.a. in the same period (BP, 2013a).

LNG is natural gas that has been cooled down until turning liquid. This reduces the volume to $\frac{1}{600}$ of the gaseous state, making storage and transportation more convenient. LNG technology is mainly used to transport natural gas over long distances at sea, where pipelines are not cost effective.

The LNG trade has seen rapid growth, diversification and increased flexibility in cargo movements over the last 20 years. Long-term contracts still dominate, but medium-term contracts, short-term contracts and spot trades have taken up an increasing share of the market. At the same time LNG technology is evolving, continuously making LNG trade more cost effective.

In the following sections we take a closer look at the the fundamentals of natural gas production and LNG trade. The LNG value chain is presented, looking at cost distribution and comparing the use of LNG to pipelines. Geographical markets are discussed, as well as some numbers related to supply and demand of natural gas and LNG. Market characteristics are presented, looking at LNG infrastructure, contracts and pricing. We finally make some remarks about technological advances in LNG production and have a look at the basics of LNG shipping costs.

2.1 The LNG value chain

This section describes the value chain of LNG, from reservoir to the end user. The distribution of costs is also discussed.

The LNG value chain, illustrated in Figure 2.1, begins with natural gas being extracted from a reservoir and sent through pipelines to a liquefaction plant. In the liquefaction plant, impurities are removed from the gas and it is cooled down until passing its boiling point at approximately $-160\,°C$. This process of converting the gas into a liquid effectively reduces the volume to $\frac{1}{600}$ of its gaseous state. Volume reduction is what makes LNG valuable, by enabling cost effective long distance transportation of gas. The LNG is then stored in tanks or directly loaded onto ships, where it is kept below its boiling temperature until reaching a regasification terminal. At the regasification terminal the LNG is pumped into a storage tank, where it is kept until being warmed up, transforming the LNG back into gas. It is then sent into the pipeline system for delivery to end users (SLNG, 2010).

Figure 2.1: LNG value chain
(SLNG, 2010)

The distribution of capital costs in a LNG value chain is approximately as follows (Maxwell and Zhu, 2011):

- Exploration and production: 15-20%
- Liquefaction: 30-45%
- Shipping: 10-30%
- Regasification and storage: 15-25%

As indicated by the above numbers, the process of converting the gas to LNG and back constitutes a major cost. Liquefaction, regasification and storage together adds up to 45-70% of the total supply costs. The alternative to LNG is transporting the gas through pipelines all the way to end consumers. A comparison of cost for the two alternatives is shown in Figure 2.2. We see that the use of LNG makes more sense the greater the distance. This is due to the high capital cost of building pipelines. It is also apparent that offshore pipelines are more expensive than onshore pipelines.

The high cost of transporting natural gas over long distances has lead to a large share of

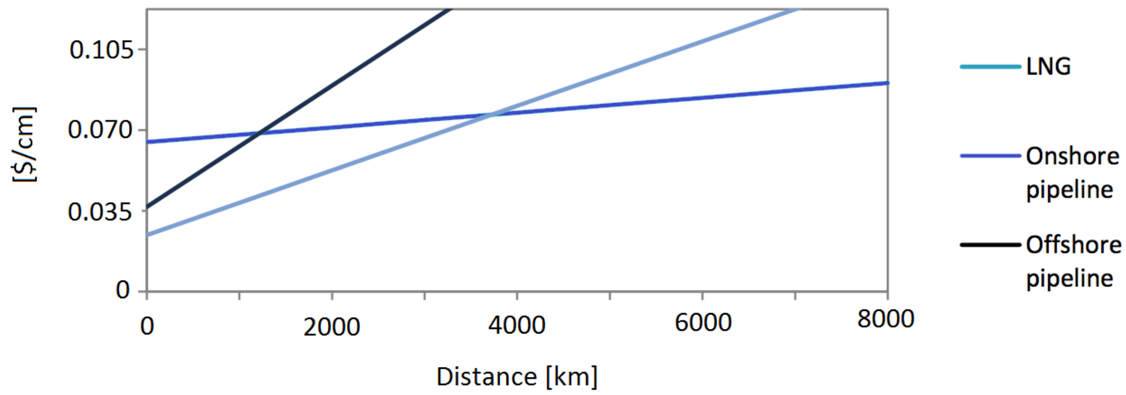

Figure 2.2: Transportation cost per cubic meter of gas as a function of distance for onshore pipeline, offshore pipeline and LNG
(Schwimmbeck, 2008)

trade being done within separate geographical markets. These markets are discussed in the following section.

2.2 Geographical markets

There are three main geographical regions in the LNG trade; the Asia-Pacific Basin, the Atlantic Basin and The Middle East. Japan and South Korea are the largest importers of LNG in the Asia-Pacific market, while Indonesia, Malaysia and Australia are the largest exporters. In the Atlantic Basin, the largest importer of LNG is continental Europe. The largest exporter of LNG is Africa. The countries in the Middle East acts as swing suppliers between the Asia-Pacific and the Atlantic Basin (BP, 2013a). These characteristics can be observed in Figure 2.3, which illustrates worldwide LNG trade in 2012. We also see that some LNG is transported all the way from the Atlantic Basin to the Asia-Pacific Basin.

Figure 2.4(a) shows that all geographical regions are expected to increase their exports of LNG in the coming years. Particularly strong growth is expected in the Atlantic and Pacific Basin. Figure 2.4(b) shows that Europe and the Asian non-OECD countries are responsible for most of the growth in imports.[1]

The increased exports have to be backed up by and increase in production. Predictions for future production are discussed next.

[1] Japan and Korea are OECD countries, while India and China are not (OECD, 2013).

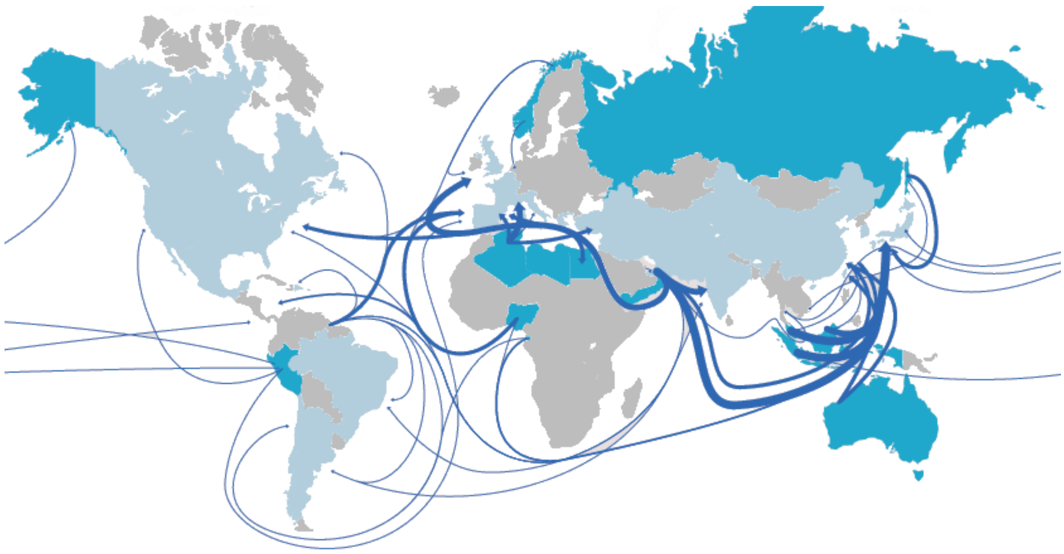

Figure 2.3: Global LNG trade movements in 2012
(GIIGNL, 2013)

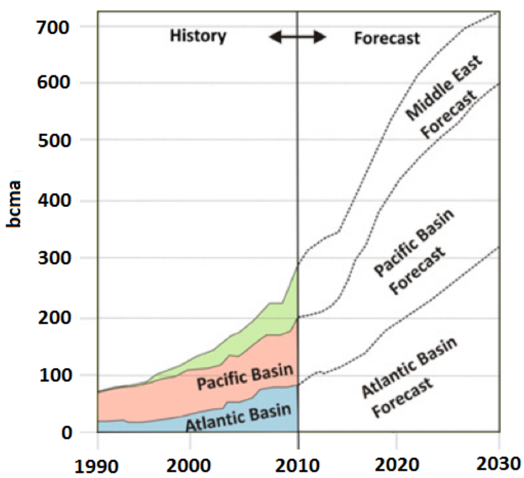

(a) Exports

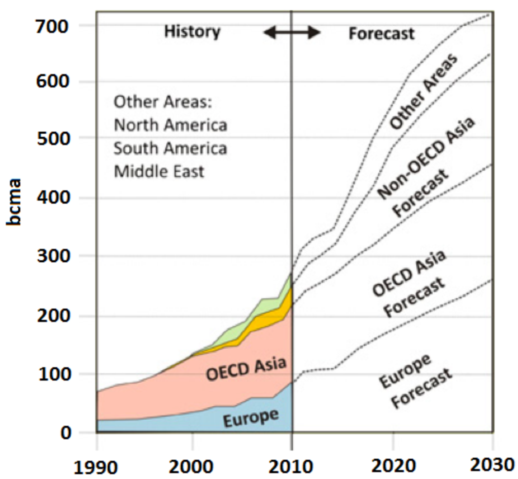

(b) Imports

Figure 2.4: Projection of global export/import of LNG in different markets/country groupings, 1990-2030
(Wood, 2012)

2.3 Supply of natural gas

Natural gas production is expected to see significant growth in the coming years. BP projects that total gas production will grow by 2% p.a. (per year), reaching 4,744 bcma (billion cubic meters per year) by 2030, compared to 3,363 bcma in 2012 (BP, 2013a).

Figure 2.5 shows EIA's (U.S. Energy Information Administration) projection of natural gas production. MENA is the Middle East and North Africa. We see that all country groupings are expected to increase production. The share of production accounted for by each country group remains relatively stable. Comparing LNG export numbers from Figure 2.5 to gas production numbers from Figure 2.4, we find that more than 15% of all gas production in 2030 is projected to be exported as LNG. This view is shared by BP, which predicts that LNG trade will make up 15.5% of global gas trade by 2030 (BP, 2013a).

The largest exporters of LNG in 2012 were Qatar, Malaysia, Australia, Indonesia, Nigeria and Algeria (BP, 2013b). Australia is expected to overtake Qatar as the largest LNG exporter by 2018 and to account for 25% of global LNG production by 2030 (BP, 2013a).

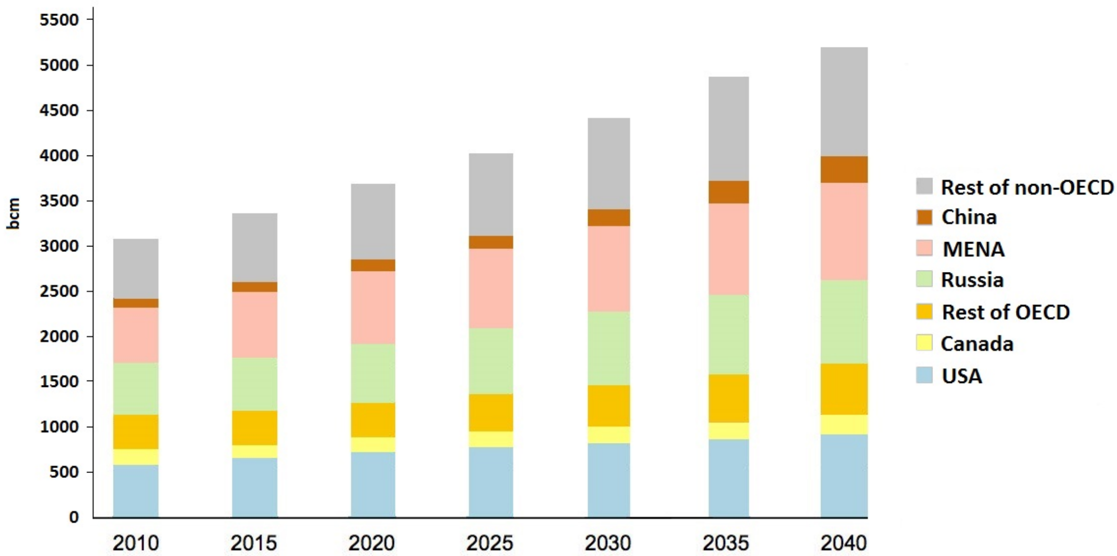

Figure 2.5: Historical and projected global natural gas production, 2010-2040 (EIA, 2013)

The United States is expected to be an especially interesting market in the coming years. Traditionally it has been a large *importer*, but the shale gas revolution is turning this around. Multiple import terminals are now being converted to also handle export of LNG. Planned projects indicate that the U.S. will be a net exporter by 2017 (BP, 2013a). The Energy Information Administration (2013) has predicted that exports will reach 222 bcma by 2040, making up net exports of 12%. In 2011, they had net *imports* of 8% (EIA, 2013). Net imports/exports represent how much gas is imported/exported as a percentage of total consumption. The U.S. will potentially have a key role as an exporter to both Europe and Asia. The distance from the U.S. to these markets is long, meaning that the use of LNG is more likely than building pipelines. This could mean new market opportunities for a speculative trader of LNG.

There are still vast amounts of unused gas resources. Based on current demand and the International Energy Agency's estimate of remaining gas resources, the world has more 200 years of natural gas left (Exxon Mobile, 2013). It is however important to notice that only 23.7% of these resources are actually proven (BP, 2013b). The geographical distribution of total proven and estimated gas resources is shown in Figure 2.6. Unconventional gas resources are less available than the conventional ones, due to lack of technology or high cost of extraction. We see that Russia and the Middle East have a large share of readily available gas resources left.

In the next paragraph we take a closer look at shale gas production, which is going to play an important role in global gas trade in the coming years.

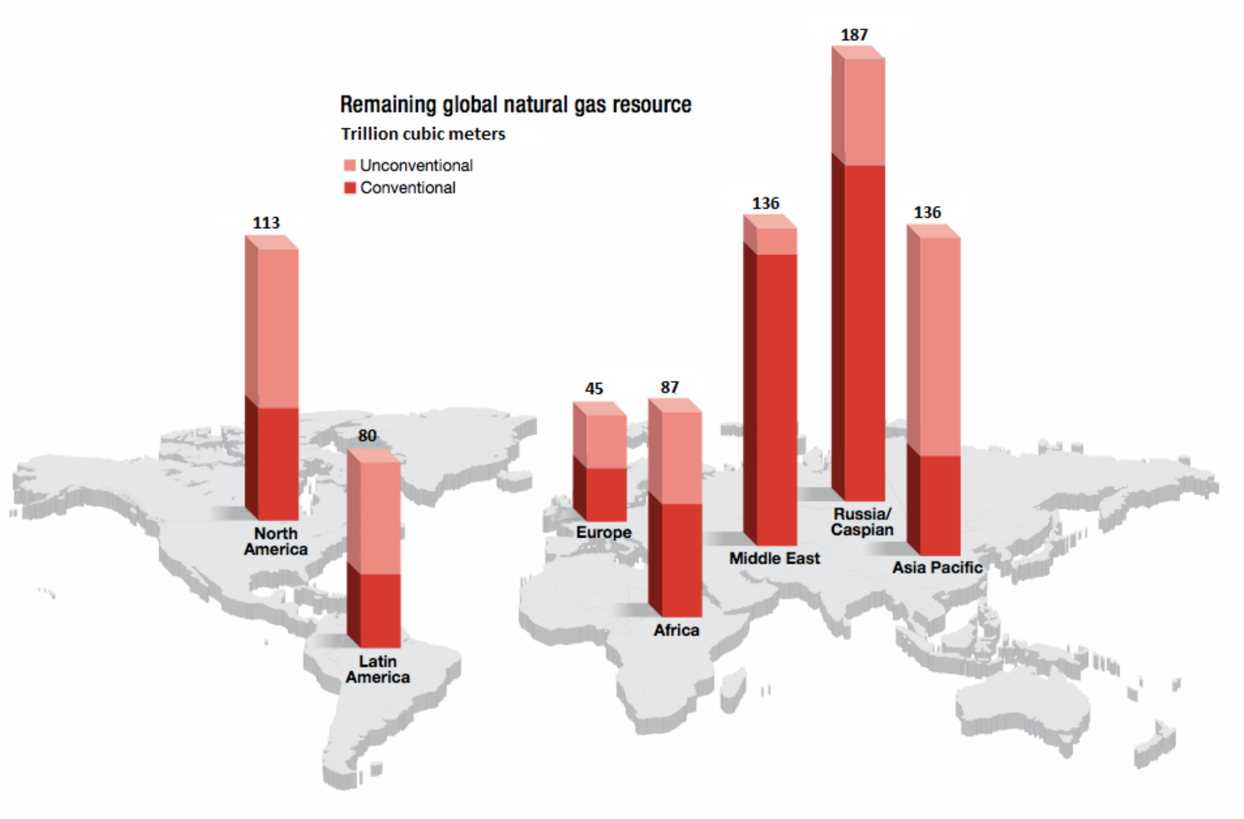

Figure 2.6: Estimate of total global conventional and unconventional natural gas reserves (Exxon Mobile, 2013)

Shale gas. Projected global growth of shale gas production is shown in Figure 2.7. Shale gas is predicted to account for more than 750 bcma by 2030, representing 37% of the expected growth in the world´s natural gas supply, and making up approximately 17% of total gas production in 2030. This is a large increase from 2010, when shale gas only accounted for 3% of total gas production. Shale gas produced in the U.S. and Canada is expected to be responsible for most of the growth, making up 72.8% of shale

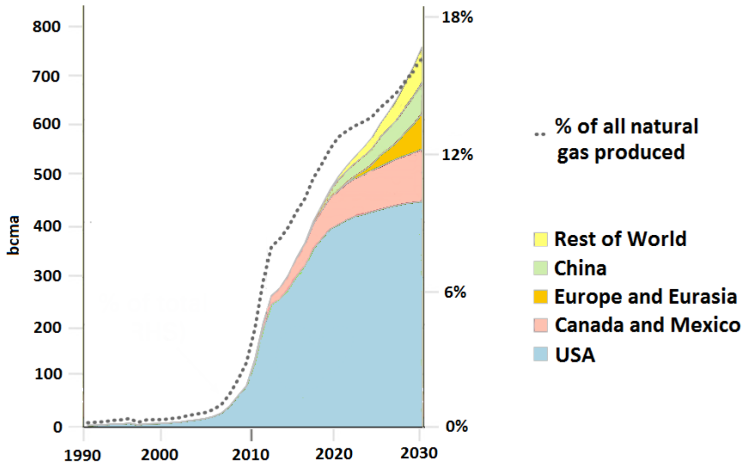

Figure 2.7: Historical and projected global shale gas production, 1990-2030 (BP, 2013a)

gas production in 2030, and almost 10% of total natural gas production. China is the country outside North America that is expected to be most successful in developing shale gas, with estimated production of 62 bcma by 2030. European shale gas production is challenging and is not likely to see any significant growth until after 2030 (BP, 2013a).

2.4 Demand for natural gas

The largest importers of LNG in 2012 were Japan, South Korea, Spain, China and India. Japan alone was responsible for more than 35% of global LNG imports in 2012 (BP, 2013b).

China is likely to experience a rapid increase in imports in the coming years. The growth in shale gas production is not enough to offset the increase in consumption. They will need an import growth of 11% p.a. due to the rapid increase in consumption, reaching 186 bcma by 2030. The EU countries are also not expected to be able to offset their coming decline of conventional gas production, leading to a 48% increase in net imports by 2030, to a total of 413 bcma by 2030 (BP, 2013a). There is also a range of new countries that are seeing gas as a way of diversifying their energy supply (Gkonis and Psaraftis, 2009). One example is the Latin American countries, where Argentina, Chile and Brazil are developing LNG infrastructure. They are likely to become key import countries in the coming years (Wood, 2012).

The market characteristics of the LNG trade are discussed next.

2.5 LNG market characteristics

LNG markets are evolving. New actors are entering the market and more infrastructure is being built. Trade is diversifying and investors are showing increased interest for LNG. Contract terms are evolving, with short-term and flexible contracts making up a larger share of total trade. New sources of gas and new technologies continue to shift the state of the markets. Examples of this can be seen in the U.S. shale gas revolution and in Shell's floating liquefaction plant project. Shale gas has taken the U.S. from being a large importer to likely becoming an exporter by 2017 (BP, 2013a). Shell's floating liquefaction plant is making it possible to liquefy the gas where it is extracted, instead of transporting it by pipeline to a liquefaction plant. This makes it possible to utilize gas resources that have previously been unusable due to the distance from land (Wood, 2012).

2.5.1 Infrastructure, integration and diversification

Figure 2.8 shows historical and projected development of LNG infrastructure. We see that there has been an increase in LNG infrastructure over the last years, and that the the trend is projected to continue in the coming decade. New terminals and plants effectively create new nodes in the LNG trading network.

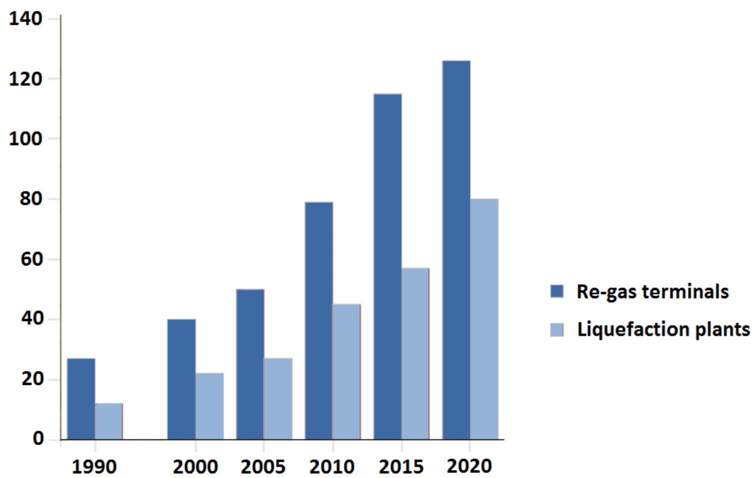

Figure 2.8: Historical and projected number of liquefaction plants and regasification terminals, 1990-2020

(BP, 2013a)

Along with the increase in LNG infrastructure we have seen an increase in diversification of trade over the last decades. One indication is seen in the declining share of LNG accounted for by the largest importer and largest exporter. In 1990, the largest importer accounted for 68% of total imports and the largest exporter accounted for 39% of total exports. In 2012, the numbers were 35% and 32%, respectively (BP, 2013b). Figure 2.9 shows how the number of suppliers per importer and customer per exporter have gone up in this period. The indicates increased competition. Nigeria and Qatar are leading in export diversification, with an average of 20 customers in 2011. Europe and Asia are also expected to further diversify their LNG supply with increased imports from East Africa and the East Mediterranean (Wood, 2012).

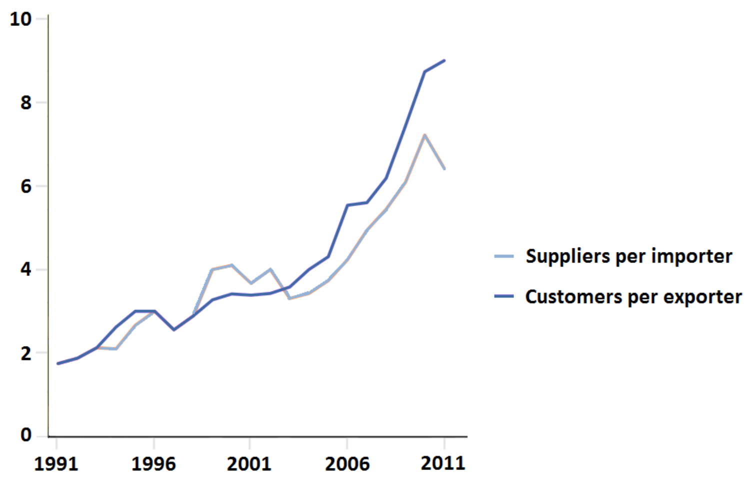

Figure 2.9: Average number of suppliers per importer and customers per exporter in LNG trade, 1991-2011

(BP, 2013a)

2.5.2 Contracts

Contracts in the LNG trade are typically long-term. The main reason for this is the need for risk allocation in the value chain. Sellers face enormous risk in making the multibillion dollar investments inherent in LNG projects. It also takes a long time before revenue is actually generated. Typically there is a delay of more than four years between the final investment decision and project completion (Energy Charter Sec., 2009). This risk is shared with buyers by having contracts that provide long-term off-take agreements (Wood, 2012). Long-term contracts are predicted to make up the major part of trade in the future as well. It is, however, expected that the long-term contracts will continue to become more flexible, allowing cargoes to be traded in the short-term market (Gkonis and Psaraftis, 2009).

The long-term contracts normally include a take-or-pay clause shifting the volume risk to the buyer. The buyer has a volume risk because he has to receive the volumes that are agreed upon in the contract. This can be either more or less than the future demand. If the buyer refuses to receive the agreed upon volume, he still has to pay the whole price (take-or-pay). The seller keeps the price risk through pricing clauses. Pricing clauses usually link the price paid to the price of a substitute product such as oil, but it might also be linked to gas market indicators such as the NBP (Energy Charter Sec., 2009). It is also common that the contracts are dedicated. This means the contracts have destination clauses, preventing the buyer from reselling the LNG. Short-term contracts have traditionally only been used to make up for the imperfect long-term planning (Engelen and Dullaert, 2010b).

There has been an increase in the share of flexible contracts, short-term contracts and spot trade in recent years. Figure 2.10 shows how the share of short-term contracts have increased from 1992 to 2007. This trend has continued, with 25% of all LNG trade in 2012 made through spot or short-term contracts (GIIGNL, 2013). Short-term trading includes contracts of three years or less and balancing trades among long-term contract holders. Spot trades are transactions that are made at once, as opposed to a contract of a future transaction. Some new gas development projects have gone forward with capacity unclaimed, leading to excess volume and potential short-term sales (Rakke et al., 2011).

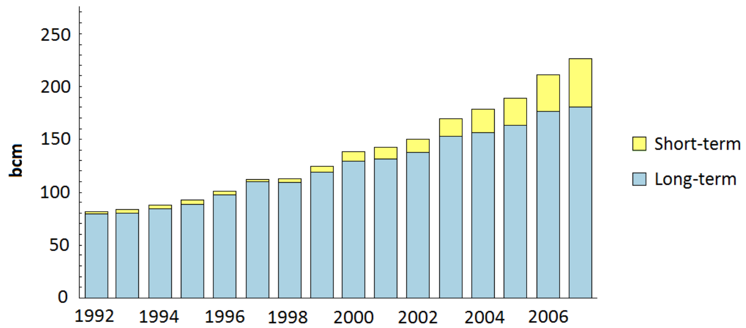

Figure 2.10: Share of short-term vs long-term contracts
(Energy Charter Sec., 2009)

Flexible contracts, as opposed to dedicated contracts, allow the cargoes to be diverted if profitable opportunities emerge. Figure 2.11 shows the share of spot, flexible and dedicated contracts in various markets in 2008. Unfortunately more recent numbers have not been found. We see that the Atlantic Basin had the largest share of flexible contracts, with almost half of the cargoes containing a destination flexibility clause. Trade in the Asia Pacific Basin is more traditional, with almost all of the contracts being dedicated.

Pricing mechanisms and choice of contract type is heavily dependent on the degree of gas market liberalization. The next section discusses factors contributing to gas market

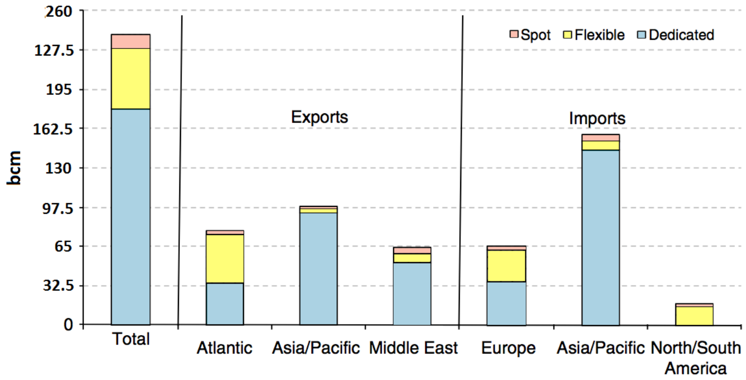

Figure 2.11: Amount of LNG traded through dedicated, flexible and spot contracts in 2008

(Engelen and Dullaert, 2010a)

liberalization, and considers how far different markets have come in the liberalization process.

2.5.3 Gas market liberalization

Successful gas market liberalization creates a liquid market where gas can be traded as a commodity. Energy Charter Secretariat (2009) identifies three preconditions that are necessary for a successful gas industry liberalisation. These are:

- Competitive gas available to market
- Customers free to choose between suppliers
- Open and nondiscriminatory access to transmission system

There are huge differences in how far various countries have come in this liberalization process. One of the main differentiating characteristics of these countries is their dependence on domestic sources relative to imported gas. US and U.K. are two examples of countries where gas supply historically has been mainly domestic. With domestic supply they were able to use government regulation to control the gas trade. Many players were given access on both the demand and the supply side. Security of supply was guaranteed by transparent and liquid markets rather than political protection. Both these markets now have established pricing points that is used for pricing of futures and comparison with other prices. In the North American market the pricing point is the price at the Henry Hub, a major pipeline junction in Louisiana. In the U.K. the price point is theoretical and is known as the National Balancing Point (NBP). In these markets short-term trading has largely replaced long-term contracts, and any long-term contracts being made have price clauses that are linked to gas market indicators (such as NBP) rather than oil prices (Energy Charter Sec., 2009).

Countries that depend on imports have had to negotiate their contracts with exporting countries and rely on long-term contracts, most of which still remain in force. The suppliers of these contracts generally wanted a minimum price or some other kind of guarantee for the entire delivery period. The buyer, however, preferred to have the gas price responsive to the price of substitutes such as oil (Asche et al., 2013). These pricing clauses have prevented gas prices from being established through gas-to-gas price competition. Limited pipeline capacity has also proven a challenge. It is hard to get pipeline capacity for long distance movements of commodity gas, due to capacity constraints in the pipeline grid, most of which is built for predetermined long-term contracts.

The import dependent countries with little or no domestic gas competition have generally not been successful in liberalizing gas trade, and long-term contracts linked to oil-prices still dominate. This is the situation for most of the European Continent and Northeast Asia. The domestic producers that exist are price-takers.

For the import dependent countries that actually have been successful in liberalizing gas trade, competition is on the terms of long-term contracts rather than on prices in a liquid commodity market. There are, however, examples of import-dependent countries that have more competitive trade than others. Two of them are Belgium and the Netherlands, where the increased competition is due to the short distance to the U.K. (Energy Charter Sec., 2009).

2.5.4 Pricing

The global gas market is not liquid. A variety of different mechanisms drive the prices in different regions. As mentioned in the previous section, some markets almost exclusively depend on long-term contracts that are linked to oil prices. Other markets have gradually converted to shorter term contracts, with prices that are based on gas-to-gas competition. In markets where domestic supply predominates over imports, such as in US and U.K., prices often fall below long-term equilibrium levels during domestic surpluses. This is an example of an issue that international pricing of LNG would have to deal with. The ideal global market would be where competition drives equilibrium prices to the long run marginal costs of the supply just necessary to meet demand. This seems unlikely to happen in the near future, with large regional differences and departures from the competitive ideal (Energy Charter Sec., 2009). In regional markets we are likely to see improved moderation of prices as the use of spot contracts increases (Wood, 2012).

The four distinct markets that mainly influence global gas pricing are North America, the U.K., the European Continent and Northeast Asia. As outlined in 2.5.3, North America and the U.K. are liberalized, while the European Continent and Northeast Asia still mainly depend on long-term contracts linked to oil prices. In the next parapgraphs market prices are discussed, before looking at potential arbitrage opportunities for a speculative trader.

Market prices

Asche et al. (2013) have studied the development of the European gas prices shown in Figure 2.12.[2] The figure shows the relationship between prices in three European gas spot markets, namely the National Balancing point (U.K.), Zeebrugge (Belgium) and Title Transfer Facility (the Netherlands), as well as the German long-term contract gas price and the price for Brent oil. We see that the prices in the three spot markets follow each other closely, and also that they seem to be correlated to the oil price. The study done by Asche et al. (2013) did not find any evidence of an independent price determination process in the European gas markets, and conclude that both the contract gas price and the spot gas price is determined by the oil price. EIA found in a 2006 study that natural gas and crude oil prices generally have had a stable relationship, despite some periods where the prices have appeared to decouple (Villar and Joutz, 2006).

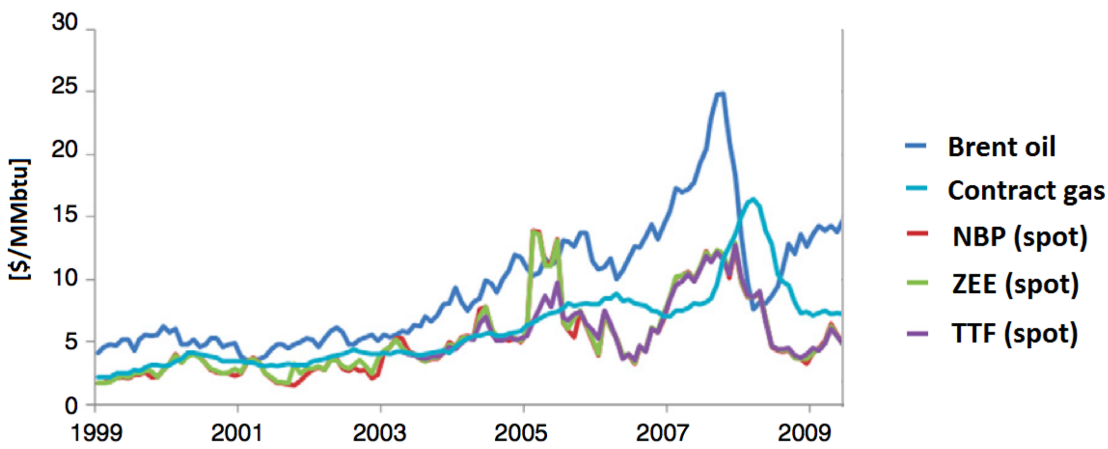

Figure 2.12: Historical price development for Brent oil and gas in Europe, 1999-2009 (Asche et al., 2013)

The Henry Hub in Louisiana has become a price reference point for the U.S. markets. Prices in other parts of the U.S. reflect the transportation cost between the Henry Hub

[2]Prices in the following discussion is listed as \$/MMBtu (million British thermal units). 1 BTU equals 28 cubic meters of natural gas.

and the market in question (Jensen, 2004). The U.S. Energy Information Administration (2013) projects that price in the Henry Hub will rice in the years to come. Their projection of the Henry Hub natural gas spot price in the future can be seen in Figure 2.13. The figure shows that the price is expected to increase steadily the coming years.

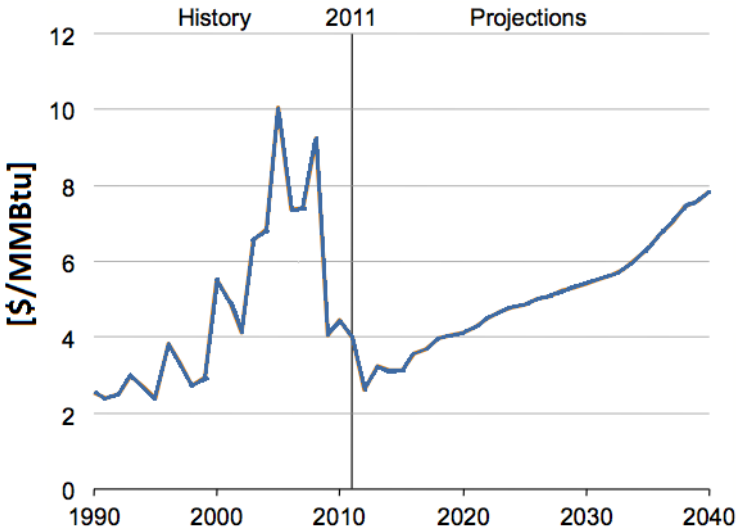

Figure 2.13: Historical and projected average Henry Hub natural gas spot price (EIA, 2013)

Figure 2.14 shows historical gas price development in the U.S.., Germany, U.K. and Japan. The price point for each year is the average gas price for that particular year. A large price difference can be observed between the different countries in recent years. Japan has experienced a steep increase in prices over the last years, especially following the Fukushima disaster of March, 2011, that lead to a shutdown of nuclear power plants. This drove demand for alternative energy supply sources, pushing the prices up. The United States has seen a sharp decline in prices after the shale gas revolution. A large price difference between markets is partially able to sustain because of the high costs of transporting LNG between markets. There are however potential arbitrage opportunities. This is discussed in the next section.

Arbitrage opportunities

Prices being determined independently in each market leads to potential arbitrage opportunities. One example of this can be seen in Figure 2.15. It shows what the market would look like from the perspective of a Nigerian shipper. The oldest numbers are hypothetical, as the U.K. did not have a import terminal before 2005. It is clear that the netbacks would differ dependent on trading with the U.K. or US. This arbitrage potential

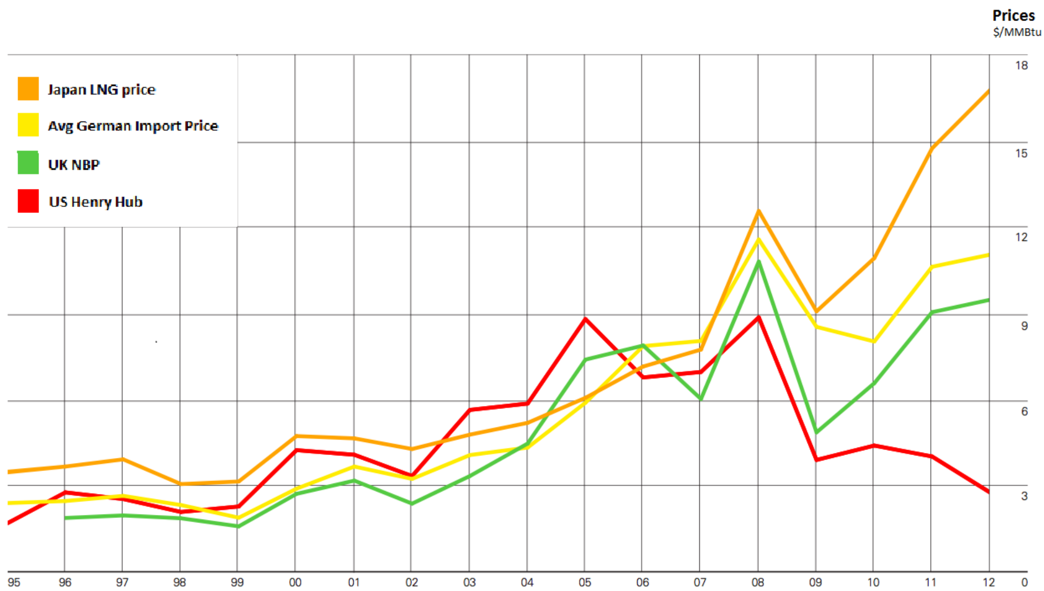

Figure 2.14: Historical LNG prices in different countries
(BP, 2013b)

could be utilized by redirecting flexible cargoes. The numbers also indicate a potential opportunities for a spot trader, which could freely choose the market with the highest return.

Figure 2.16 we see an example of how the Middle East potentially could be used as a source of arbitrage between the Atlantic and Pacific Basins. With its location between the basins, the Middle East has good opportunities to trade in the basin with most favorable price. It can be seen that spot trades between Japan and Qatar have been especially profitable.

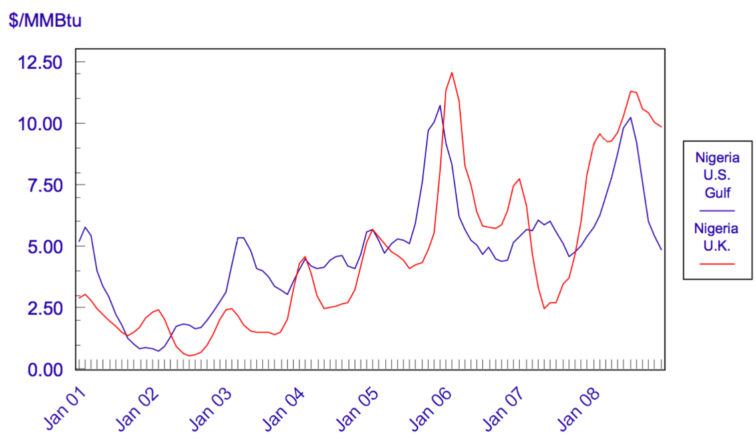

Figure 2.15: Hypothetical netback to Nigeria from the U.S. Gulf and the U.K.
(Energy Charter Sec., 2009)

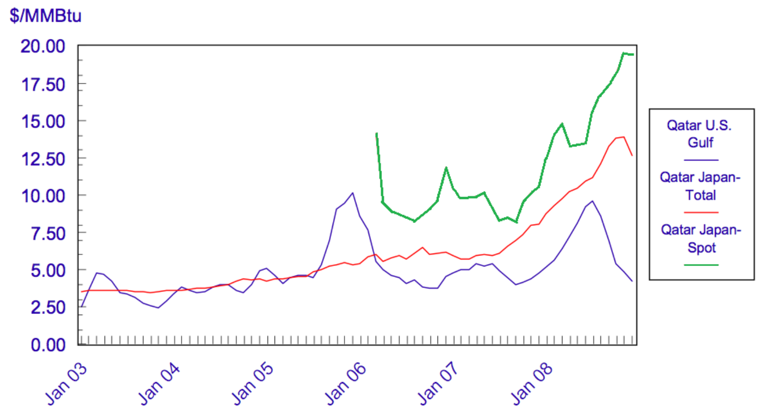

Figure 2.16: Hypothetical netback to Qatar from the U.S. Gulf and from Japan (Energy Charter Sec., 2009)

2.6 Technological advances in LNG production

The enormous profit potential in the LNG industry has helped drive innovation and research on new technical solutions. There has been significant cost reductions in all stages of the LNG chain over the last decades. This has made LNG more competitive compared to other energy sources (Gkonis and Psaraftis, 2009). Logistics costs of liquefaction alone is $\frac{1}{3}$ of what it was in 1980 (Maxwell and Zhu, 2011). The cost of producing certain types of carriers has almost halved over the last 30 years. Processing and storage capacity has increased continuously, with the size of LNG trains going from approximately 1.36 bcma in 1970 to 10.9 bcma in 2010 (GIIGNL, 2013). LNG trains are the facilities in liquefaction plants where the gas is purified and liquefied. The capacity of LNG carriers has increased from 27,500 m^3 in the 1960s to the new 265,000 m^3 Q-max vessels that entered the market in 2008 (Lopac, 2008). In general, we have seen an increase in size of terminals and vessels that have lead to economies of scale (Engelen and Dullaert, 2010a).

Floating liquefaction plants is one of the more important technological innovations evolving in the industry. Shell's *Prelude* project involves a 488 meter long floating liquefaction plant weighing 600,000 tonnes when fully equipped and loaded. The hull was launched in December 2013, but drilling is not expected to begin before 2017. *Prelude* is expected to produce at least 4.9 bcma of LNG (Shell, 2013). Other companies are following closely to see if this is a profitable way of production, and multiple similar projects are under consideration. Onboard regasification and Floating Storage Regasification Units are other technological innovations that have had success in recent years (Wood, 2012).

2.7 LNG shipping

Shipping LNG is particularly expensive due to the need for specially designed vessels that are able to keep storage tank temperatures below −160 °C. Until recently, the main players in LNG transportation have been energy majors and national companies. The increased demand and new conditions in the international energy scene has led to a need of cost reduction and versatility in the market. This has given opportunities for independent vessel owners and other investors to enter the market (Gkonis and Psaraftis, 2009).

2.7.1 The LNG fleet

The world LNG fleet currently consists of 365 vessels with a total capacity of 53,893,000 m^3. 104 new vessels with a total capacity of 9,319,000 m^3 are currently in the order books (Lloyd's, 2013). Two types of containment systems are used for LNG vessels, *membrane system* and *self supporting system*. Examples of the two types are shown in Figure 2.17 and Figure 2.18. The price of LNG carriers shifts with the market. Recent contracts indicate a price of approximately $200 million for a 160,000 m^3 vessel (Reuters, 2013). The prices were as high as $280 million for an equally sized vessel in the early 1990s, but competition and new technology has driven the prices down (Lopac, 2008).

Figure 2.17: Membrane LNG tanker
(Qatargas, 2013)

Figure 2.18: Self supporting LNG tanker
(Global Security, 2013)

2.7.2 Transportation costs

Transporting LNG includes three main types of costs:

- Vessel costs
- Operative costs
- Voyage costs

Vessel costs relate to hiring a vessel. The price for hiring a 160,000 m^3 LNG vessel lies around \$100,000 per day (Golar LNG, 2013a). The rates are highly volatile, going as low as \$30,000 in mid-2010 and reaching \$150,000 in mid-2011 (Golar LNG, 2013b). Short-run factors that impact vessel prices are weather forecasts, energy-policy, initiatives, gas-flow problems and week-to-week changes in drilling activity. Long-run factors include world economic growth, shale extraction, LNG development and carbon policies (Energy, 2010).

Operative costs include different costs related to operating an LNG vessel. An approximation of operative costs for a 138,000 m^3 vessel is shown in Table 2.7.2. They sum up to about \$18,000.

Table 2.1: Average operating costs for a 138,000 m^3 LNG vessel

(Lopac, 2008)

	Insurance	Maintenance	Spare parts	Adm. costs	Crew costs	**Total**
\$/day	5,200	760	1,782	800	9,222	**17,764**

Voyage costs include fuel costs and boil-off, and are heavily dependent on uncertain variables such as weather and fuel price. The boil-off rate is the daily percentage of total cargo capacity that is lost. It is usually in the interval 0.10-0.25% (Lopac, 2008). It is normal to keep some LNG in the tanks also when not currently transporting, to keep the tanks cold. The boil-off contributes to LNG transport having a higher voyage costs than most other types of shipping, making it more expensive for a speculative LNG trader to wait around for market opportunities. Voyace costs sum up to around $100,000 per day (Golar LNG, 2013b).

The total cost of operating an LNG vessel is the sum of vessel, operative and voyage costs. The estimates given above indicate a daily cost of approximately $220,000 for hiring and using a LNG vessel.

Chapter 3

Literature Review

In this thesis we develop a stochastic optimization model for buying, transporting and selling LNG in the spot market. The objective is to maximize profit. Scenarios are used to represent an approximation of the stochastic price development process in the ports. The vessel acts as a tramp ship, not following a predefined schedule, but rather buying and selling loads where and when prices are deemed most favorable. A rolling horizon approach is used to make day-to-day decisions based on updated information. We seek to optimize the positioning of the vessel and the timing of trade under price uncertainty.

We have reviewed literature regarding four main topics that we find relevant to our problem; LNG, tramp shipping, stochastic modeling and stochastic dynamic programming. These topics are presented in the following sections. None of the articles are directly comparable to our thesis, but in total they cover many of the relevant aspects of our study. After presenting the articles we discuss how our work fits in with existing literature.

3.1 LNG

In this section we first present papers regarding the LNG Inventory Routing Problem (IRP). IRPs are problems that take both inventory management and routing into consideration. IRPs span over liquefaction, shipping and regasification in the LNG value chain (Figure 2.1). We move on to describe LNG Annual Delivery Problems (ADPs). An ADP is a complete schedule of every ship's sailing plan for the coming year. Articles combining LNG and uncertainty are then presented, before briefly considering literature related to the LNG production process.

The LNG IRP is first discussed in Grønhaug and Christiansen (2009). They present a supply chain optimization model for the LNG business, looking at transporting and scheduling of LNG ships as well as inventory management at both liquefaction plants and regasification terminals. Both a path-flow and an arc-flow model with pregenerated paths are described. The computational study shows that the path-flow formulation is faster in solving to optimum, but the arc-flow formulation is faster at finding the first integer solution.

Grønhaug et al. (2010) solve the LNG IRP by a branch-and-price method. Inventory management and port capacity constraints are handled in the master problem, while the subproblems generate ship route columns. The model also includes an ad hoc DP-algorithm for solving the longest path subproblem. A variety of acceleration strategies are included in the model as well. The computational study shows that this approach gives much better results (on average more than one magnitude faster) than the model presented in Grønhaug and Christiansen (2009).

Fodstad et al. (2011) and Uggen et al. (2013) study a LNG IRP that includes contract management and trading in the spot market. Goel and Furman (2012) present an arc-flow formulation for the LNG IRP based on the MIP model of Song and Furman (2010). The model optimizes ship schedule decisions together with inventory management at both production and regasification terminals. Construction and improvement heuristics to solve the model efficiently are presented, including several two-ship selection methods. An overview of existing literature in the field of combined routing and inventory management can be found in Christiansen and Fagerholt (2009).

Andersson et al. (2010) study the LNG supply chain, presenting two planning problems that combine transportation planning and inventory management. One is for a producer and the other for a vertically integrated company that controls both the liquefaction and the regasification terminals in addition to transportation. The output of the model is an Annual Delivery Program (ADP).

Rakke et al. (2011) present a rolling horizon heuristic for creating an ADP for a LNG producer. The rolling horizon approach is used to simplify the complex problem into more solvable sub-problems. They minimize the cost of fulfilling long-term contracts while maximizing revenue from selling LNG in the spot market. A similar ADP problem is discussed in Stålhane et al. (2012). Here a multi-start local search heuristic is presented. To improve the heuristic solution they use either a first-descent neighborhood search, branch-and-bound or both. Halvorsen-Weare and Fagerholt (2013) present an alternative model. Their solution method includes decomposition into a routing subproblem and a scheduling master problem. Inner and outer time windows for deliveries are used, with target dates that can be violated at a penalty cost. In Halvorsen-Weare et al. (2013) uncertainties in production rates and sailing times are also considered. Three robustness strategies are tested; adding slack to each sailed round-trip, adding target inventory levels and adding target accumulated berth use. They show that using each of the robustness strategies, and a combination in particular, results in overall lower expected costs. Their results show that there is a significant improvement potential in considering and dealing with the uncertain parameters of LNG vessel routing and scheduling.

Ainouche and Smati (2002) present a stochastic dynamic programming model for reducing production costs in LNG value chains. They take into account the increasingly stochastic nature of LNG development projects resulting from spot market trades replacing long-term contracts. Khalilpour and Karimi (2012) consider contract selection under uncertainty from a LNG buyer's perspective. They present a mixed-integer linear programming model that helps the buyer select the best combination of suppliers and contracts. The sum of purchase and transport costs is minimized.

There are also examples of optimization being used to improve the chemical processes of LNG production. Aspelund et al. (2010) present a optimization-simulation method to minimize the energy requirements of a PRICO LNG process based on Tabu Search and the Nelder-Mead Downhill Simplex. Wahl et al. (2013) use sequential quadratic programming for optimizing a PRICO LNG liquefaction process. Hwang et al. (2013) present a model for optimizing the dual mixed refrigerant (DMR) using the genetic algorithm and sequential quadratic programming.

3.2 Tramp shipping

In this section we start by comparing tramp shipping to other types of shipping. We then present articles about scheduling in tramp shipping, before briefly considering other relevant applications of optimization in tramp shipping.

It is usual to divide shipping operations into three types: liner shipping, industrial shipping and tramp shipping. Liner shipping follows a predefined schedule, similar to a public bus service. Industrial shipping involves operators that are transporting their own cargo between ports. Since the industrial operators own the ship and cargo themselves, their goal is to minimize the cost. If there is any spare capacity on the ship, an industrial operator can transport cargo from the spot market in order to make profits. Unlike liners, there are no given schedules or routes for the industrial operators. Tramp shipping can be compared to the taxi business. Tramp operators might have contracts that oblige them to transport cargo, but they seek the opportunity to pick up available cargoes when the vessel has spare capacity. Their goal is to maximize profit (Hwang et al., 2008).

Christiansen et al. (2004) present a review over ship routing and scheduling. They found that there was an ongoing shift in the market towards more use of tramp shipping, with companies starting to outsource the shipping operations. This leads to more market interaction and increased opportunities for optimization-based tools in decision support. Christiansen et al. (2013) have done a review on papers about ship routing and scheduling from the last decade. They conclude that the volume of research on ship routing and scheduling has more than doubled during this period. LNG shipping is mentioned as one of the fields that has attracted more attention. The review serves as an introduction into various part of ship routing. A review of maritime transportation in general can be found in Christiansen et al. (2007).

Kim and Lee (1997) present a system for ship scheduling for bulk trade. The paper concludes that there are great fluctuations in shipping rates of bulk trades, and that there is great potential for increased profit with proper scheduling.

Brønmo et al. (2007a) present a local search heuristic for short-term tramp shipping scheduling problems, where the objective is to maximize profits. Initial solutions are made by an insertion heuristic, before a local search heuristic is used to improve a given number of the best initial solutions. The paper also states that there has been little attention to the tramp market historically. Korsvik and Fagerholt (2010) use a tabu-search heuristic to solve the same problem. This heuristic allows infeasible solutions in ship-capacity and time windows. The tabu search heuristic perform much better then the multi-start heuristic for large and tightly constrained instances.

Malliappi et al. (2011) present a variable neighborhood search heuristic for solving a routing and scheduling tramp ship problem. The computational results show that this heuristic gives better solutions and faster computation times than the heuristics used by Brønmo et al. (2007a) and Korsvik and Fagerholt (2010). Brønmo et al. (2007b) describe a MP-model of a tramp shipping pickup and delivery problem with time windows, flexible cargoes and multiple ships. Set partitioning is used to solve the problem, with columns generated before the model starts. The objective function is to maximize profits. Brønmo et al. (2010) address the same problem. Instead of generating all columns at the start of the problem, they use dynamic column generation. This solution method can be used with large or loosely restricted instances.

Hwang et al. (2008) use a branch-and-price-and-cut algorithm to present a set-packing model that limits the risk of delivering spot cargoes in tramp routing and scheduling. Due to volatile spot prices there is uncertainty in the spot market, and this model helps ship owners make decisions based on their risk-aversion. Lin and Liu (2011) propose a tramp shipping model that uses a genetic algorithm and simultaneously takes into account the ship allocation, freight assessment and ship routing problem. Fagerholt et al. (2010) present a decision support methodology for strategic planning in tramp and industrial shipping. A combination of optimization and Monte Carlo simulation is used. A rolling horizon principle is applied, where information is revealed to the model as time goes by.

3.3 Stochastic modeling

In this section we start by covering papers that describe stochastic programming in general. We then have a look at reviews on the use of stochastic programming in routing problems. We conclude by presenting papers that have solved strategic routing problems under uncertainty. Stochastic dynamic programming is discussed in the next Section 3.4.

Higle (2005) has written an introductory article to stochastic programming. The article describes different stochastic models as two-stage and multistage, and also proposes solution methods. Higle points out that solving multistage problems is complex, and suggests using decomposition. Flatberg et al. (2007) show the importance of using dynamic and stochastic models, as opposed to deterministic models. This article also describe dynamic and stochastic VRPs. Pillac et al. (2013) classify routing problems from the perspective of information quality and evolution and present a comprehensive review of applications and solution methods for dynamic vehicle routing problems.

Berbeglia et al. (2010) survey dynamic pickup and delivery problems. The article includes basic issues and how the problems can be solved. A pickup and delivery problem with time windows is discussed by Mitrović-Minić (2004). They use a heuristic that considers both the short-term and long-term horizon. Hvattum and Løkketangen (2007) describe a branch-and-regret heuristic for solving stochastic VRPs. The method used in the paper outperforms previous heuristics. Ichoua et al. (2006) introduce probabilistic knowledge about future requests to solve a dynamical real-time vehicle routing and dispatching problem. Savelsbergh and Sol (1998) present a planning model for vehicle routing. The model uses a branch-and-price algorithm and a rolling horizon approach. The vehicle routing problem with time windows is solved by using a multiple scenario approach in Bent and Van Hentenryck (2004). The model includes known requests and a future request based on a probability function.

Christiansen and Fagerholt (2002) solve a shipping problem deterministically. They make results more robust by putting a penalty on solutions that are risky, and thereby handling uncertainty. A risky solution would e.g. be when vessels arrive in ports close to weekends, thus risking having to wait in port until the following Monday. McKinnon and Yu (2011) describe a stochastic ship routing problem with uncertain demand. A branch-and-price algorithm is used to solve the problem. Tirado et al. (2013) consider a dynamic and stochastic maritime routing problem that arises in industrial shipping. Applying customized versions of three well-known heuristics, they show that average yearly cost savings of 2.5% can be achieved by including stochastic information in the model. The savings are found to be substantially larger for instances with partial loads rather than full loads. Shao et al. (2012) present a novel forward dynamic programming method for weather routing that seeks to minimize ship fuel consumption during a voyage.

Ang et al. (2009) use stochastic models when planning container mixes for ships. The aim of the work is to maximize the total expected profits over uncertain scenarios. They use a two-stage stochastic model and solve it with a heuristic algorithm. Shyshou et al. (2010) present a simulation study for a fleet sizing problem, including uncertainty in weather conditions and spot price rates. Alvarez et al. (2011) present an optimization model for fleet sizing and deployment problem. This model is robust and deals with uncertainty in future prices and demand. The model applies the robust optimization technique used by Bertsimas and Sims (2003), which gives the decisions-makers an opportunity to choose their level of risk tolerance.

3.4 Stochastic dynamic programming

In the following section we first present articles about routing and inventory management that uses stochastic dynamic programming. Following this we describe papers solving other problems by applying the same solution approach.

Desai and Lim (2013) use stochastic dynamic programming to determine optimal routing policies in a stochastic dynamic network. They also propose three techniques for pruning stochastic dynamic networks, effectively speeding up the process of attaining optimal routing policies. The techniques includes use of static upper and lower bounds, preprocessing of the network by considering start time and origin of the vehicle, and a mix of the two. Novoa (2009) examines the use of approximate dynamic programming algorithms for the single-vehicle routing problem with stochastic demands from a dynamic or reoptimization perspective. The rollout algorithm is extended by implementing different a priori solutions, look-ahead policies, and pruning schemes. In addition to the direct approaches, Monte Carlo simulation is used.

Azaron and Kianfar (2003) use stochastic dynamic programming to find the dynamic shortest path for source node to sink node in stochastic dynamic networks, where arc lengths are independent random variables with exponential distributions. There is also a environmental variable in each node. This node evolves in accordance with a continuous Markov process and has an impact on the transition time on arcs exiting the node. At each node a decision is made on moving towards the sink node on the best outgoing arc or waiting. The problem is discussed for both full and limited information about environmental variables.

Berman et al. (2001) consider an inventory and routing problem where the amount of product at each customer is a known random process. The objective is to dynamically adjust the amount of product provided to each customer to minimize total expected costs. Costs comprise earliness, lateness, product shortfall and returning non-empty to the depot. The policy is determined by stochastic dynamic programming. Yang and Grothey (2012) use approximate dynamic programming to solve the top-percentile traffic routing problem faced by Internet Service Providers (ISPs). They describe a multistage stochastic optimization problem, where the routing decisions must be made before knowing the amount of traffic to be sent. The integer variables introduced by top-percentile pricing makes it hard to solve exactly. Use of approximate dynamic programming exploits the structure of the problem to construct continuous approximations of the value functions in stochastic dynamic programming.

Boutelier et al. (2000) use dynamic Bayesian networks to represent stochastic actions in Markov decision processes. Dynamic programming algorithms are developed that directly manipulate decision-tree representations of policies and value functions. The method shows significant savings for certain types of domains. Cristobal et al. (2009) outline a stochastic dynamic programming approach where a scenario tree is used in a back-to-front scheme. Multi-period stochastic problems are solved at each given stage of the time horizon. Each subproblem considers the effect of stochasticity of the uncertain parameters from the periods of the given stage, by estimating the expected future value of the objective function. The scheme is applied to a production planning problem and is found to work well for instances on a very large scale.

Kelman et al. (1990) develop a technique called sampling stochastic dynamic programming (SSDP) for reservoir optimization. The technique captures complex structures of the streamflow process by using a large number of sample streamflow sequences. Shapiro (2011) discusses statistical properties and convergence of the SDDP method applied to multistage linear stochastic programming problems. The framework discussed involves generating a random sample from the original distribution and then applying the SDDP algorithm to the constructed Sample Average Approximation problem.

3.5 Our work

We present two models for solving a vehicle routing problem under price uncertainty. The objective is to maximize profit. Stochastic dynamic programming is used to find the value of being in specific ports on given days. The models tries to find optimal movement and trade decisions based on these values. One of the models is based on a mixed integer program and the other on a heuristic. A rolling horizon approach is applied, where new information is revealed to the model as time progresses. The rolling horizon approach has two benefits; it allows us to use updated price information every day, and it splits our problem into smaller subproblems. Our models can be used for any kind of tonnage with only minor adjustments. It is especially well suited for bulk shipping, as we assume a given price per unit of shipped goods.

To our knowledge no previous research has been done on contracting of tonnage on a speculative basis in the shipping industry. Our problem is similar to previous research in the sense that it maximizes profits of a shipping problem under uncertainty. Rakke et al. (2011) use rolling horizon to decrease complexity in a problem that is similar to ours. The difference between our problem and many other models is that we are making move decisions in every time period. This lets us have more flexible movements than traditional routing problems, where movements are dictated by a start port and an end port.

Chapter 4

Problem Description

We look into speculative trading of LNG in the spot market, including the shipping of LNG from loading (buy) ports to unloading (sell) ports. The shipping resembles tramp shipping, where the schedule is dynamically determined by the opportunities that are present in the market.

The problem is considered from the point of view of a ship owner operating a single LNG vessel. The goal of the ship owner is to maximize profit over time by making optimal decisions related to the positioning of the vessel and the timing of trade. This is done by maximizing income from buying and selling, while minimizing the cost of travelling between ports.

Each time period a decision has to be made on whether the vessel should move, and if so, in which direction. When reaching a port, it must be decided if the vessel should trade or wait until a later time period. Each new time period brings an update of prices and a corresponding forecast of future prices. There is uncertainty in the forecasted prices. The ship has to trade full shiploads, and needs to be in a port to make a trade. Using options and futures is not possible.

Chapter 5

Models

In this chapter we state our assumptions and present the models used to solve our problem. There are three main models. The first model described is a Dynamic Program (DP), the second a Mixed Integer Program (MIP) and the third a heuristic. The DP is used in combination with either the MIP or the heuristic when solving the full problem. The role of the DP is to find the best route to travel from each port in the long-term horizon, and assign values to the ports based on this route. The MIP and heuristic then use this information to decide the vessel's short-term actions. The interaction between the models is discussed further in Chapter 6.

5.1 Assumptions and modelling choices

In this section we state our main assumptions and modelling choices. We first present how we have chosen to model geography and time, before looking at specific assumptions regarding ports, vessel and trade. The assumptions stated apply to all models.

5.1.1 Grid

A grid is used to structure the geographical aspect of the problem. An example of a grid is shown in Figure 5.1. Each grid point is a position that can contain a port and/or a vessel. The vessel is only allowed to travel between grid points. Grid points in brown areas are on land and cannot be visited by the vessel.

Figure 5.2 shows how the vessel is allowed to move. It has four straight and four diagonal options. In each time period the vessel can choose between staying put or travelling to one of the neighboring grid points. Ideally the vessel should have the option to travel to

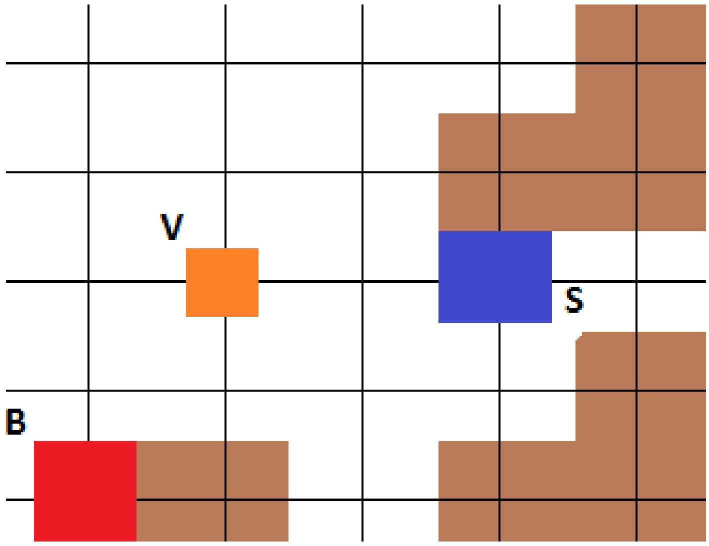

Figure 5.1: Example grid with one vessel (V), one buy port (B) and one sell port (S)

any position within a given area, as marked by the green area in Figure 5.3. Only allowing a limited number of moves is a simplification of the problem. It does however fit well with receiving new prices at predefined time intervals. The run time is also significantly reduced. In addition, the number of grid points can easily be increased or decreased, which makes it easy to adjust the precision of the model.

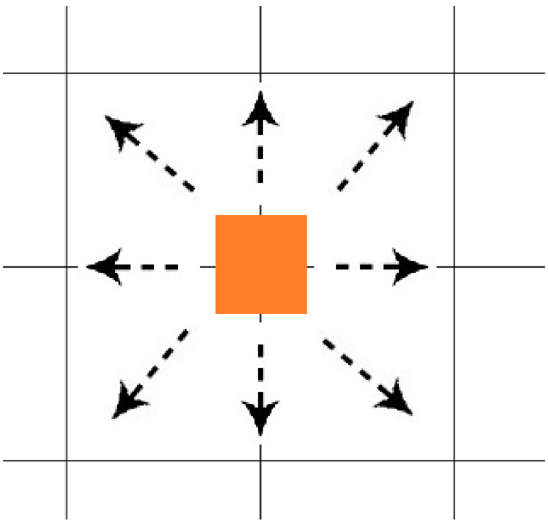

Figure 5.2: Allowed moves in the grid

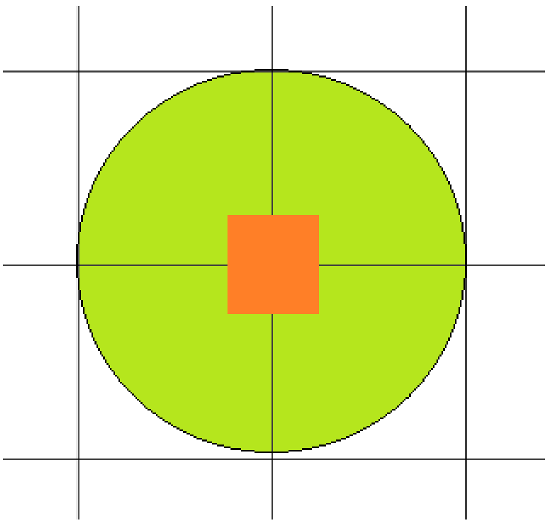

Figure 5.3: Allowed moves in real world

5.1.2 Time

A decision to move, trade or wait is made every time interval. In our test instances one time interval equals one day. It takes one day to move from one grid point to a neighbouring grid point and one day to complete a trade. A decision to wait also lasts one day.

The time interval decides how often the vessel can make a new decision. As long as new price information is not received, there is no need of change the sailing direction of the vessel. Thus, the time interval should be decided based on how often new price information is given. There is a trade-off between a short and a long time interval. Long intervals does not capture all changes in the price. Short time interval would be beneficial in a real-life situation, but it makes the test run times very high when solving for many days. We believe that a time interval of one day both captures the changes in price, and is long enough that the model can be run for a long trade horizon in reasonable time. The choice of time interval is strongly connected to the choice of map. When setting the time interval to one day, we also set the precision of the grid to be one day of travelling between each grid point.

A decision made at time t affects the time interval t to $t+1$. For example, if the vessel decides to make a move at time t, it arrives at the destination at time $t+1$. If the vessel chooses to buy at time t, it is full at $t+1$. Figure 5.4 shows this relationship.

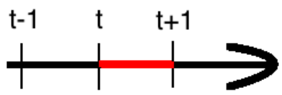

Figure 5.4: The decision made at time t takes place between t and $t+1$.

5.1.3 Sailing cost

The cost of a diagonal move is $\sqrt{2}$ times the cost of a straight move. We have set the cost $\sqrt{2}$ times longer since the distance of a diagonal move is $\sqrt{2}$ longer than a straight move. This is a simplification, as the vessel has to sail faster when moving diagonal, and sailing costs are a cubic function of speed (Norstad et al., 2011). In addition to the cost of sailing, the vessel incurs a basis cost every day. The basis cost relates either to the cost of hiring a vessel, or the alternative cost of owning one.

5.1.4 Ports

The ports are assumed to have infinite supply and demand. This means that the vessel can trade in any port at any time, and that prices in the ports are unaffected by trades done by the vessel. We believe that this is a reasonable assumption, as a single vessel is not likely to impact the market prices alone. An Ornstein-Uhlenbeck price process is used to forecast the prices in the ports. This process is described more in detail in Section 8.2.

5.1.5 Trade

All trades are discounted with a factor of 1% per 30 days. This is because of the time value of money.

5.2 Dynamic programming to find port values

The dynamic programming is used to determine the potential future value of being in a port on a given day. In order to find these values we evaluate potential trade sequences. This is done by generating routes that contain combinations of port visits. Labels are used to keep control of generated routes, and a domination function is applied to limit the number of routes under consideration. In this section we describe the features of the dynamic program.

5.2.1 General description

The dynamic program takes a specific start port, start day and port prices for the trade horizon as input. The trade horizon represents the amount of time the vessel has to complete its trades. Port prices beyond the current day are forecasts based on the method described in Section 8.2. Based on this input, the DP finds the most profitable route from the start ports.

The first thing the vessel has to do is make a trade in the start port. This can be done in the start day, or on one of the subsequent days. All options up to a maximum number of wait days are considered. After the initial trade is made, the vessel moves on to its next trade. If the vessel is in a buy port, the next trade will be in a sell port, and vice versa. Again, it can either move directly to a port and trade, or wait a number of days before making the trade. The program evaluate the impact on future trades. To do this it considers all possible future trade sequences from each of the ports, including all combinations of waitdays. This leads to an enormous number of routes generated. To reduce the number of routes a domination function is applied. It removes routes that are inferior to an other existing route. This domination function is explained in Section 5.2.3.

After the DP has evaluated all potential routes it returns the best route and the value (profit) of this route. This value corresponds to the potential profit that can be earned by moving to the start port on the given start day.

5.2.2 Labels

Labels are used to keep control of routes generated. A label contains current information about port, day, profit, and the route travelled that far. See Figure 5.5 for an illustration. This label indicates that the vessel is in port 2 on day 44 with a profit of $240,932. It started in port 1, and then traded in port 5, 2 and 6, before ending up in port 2.

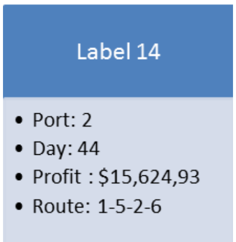

Figure 5.5: Example label from the DP

The first thing done by the DP is to generate labels for the start port. Labels are made for trading straight away, as well as for trading on subsequent days. The initial value of the day parameter corresponds to the number of days the vessel has to travel before reaching the start port, plus the number of wait days before making the first trade. These labels are then extended with all possible options for the next action, creating a new label for each of those possible actions.

The extended labels have updated fields for profit, time spent and route travelled. The change of profit from one label to another is the sum of traveling cost and trade value. The value of the trade can be both negative (when buying) and positive (when selling).

Figure 5.6 shows an illustration of label expansion. The vessel has bought in port 2 on day 44 and is considering its sell port alternatives for the next action. The figure only shows three of the vessel's options, namely selling in port 5 on either day 49, 50 or 51. The vessel also has the option of trading in port 5 on a later day, or trading in another port. We see from the figure that label 16 has lower value than label 15. This is because the price in the port has decreased from day 49 to day 50. In label 17 we however see that the profit is higher than in label 15. This indicates that prices have developed in a beneficial way. It is however important to notice that further expansions of label 15 are likely to get higher values than the expansions of label 17, as it has two more days left for trading. New labels keep on being extended until all labels generated exceed the trade horizon or are dominated by some existing label.

5.2.3 Domination

The number of labels grows exponentially. This is a major issue for model run time. With 3 sell and 3 buy ports we have 3 trade options for each trade. In addition we have the option of waiting up to a certain number of days before making any one of the trades. A maximum of 7 waiting days gives a total of 8 options per port (trading anywhere between day 0 and 8), and a total of 24 options per trade (trading in any of the 3 ports). As each of these options needs to be fully checked, 24^x labels are generated, where x is the number

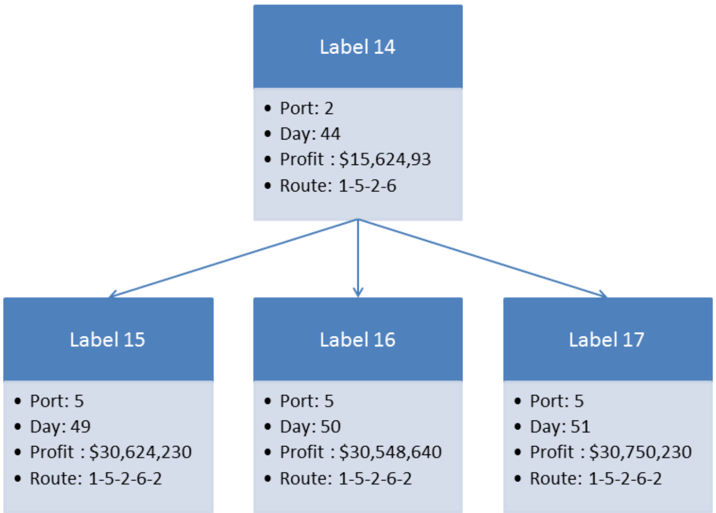

Figure 5.6: Example of label extension in the DP

of trades in the trade horizon. The actual number of trades depends on the length of the trade horizon, how many waiting days are used and the time spent travelling. The number of labels grows rapidly, as shown in Table 5.1.

Table 5.1: Number of labels generated as a function of the number of trades

Number of trades	1	2	3	4	5	6	7
Labels generated	24	576	13,824	331,776	7,962,624	191,102,976	4,686,471,424

A domination function is used to deal with this problem. The domination function effectively removes all labels that are guaranteed to end up with a lower profit than some other existing label. Every new label is compared to all existing labels to see if it is dominated by any of them, or if the new label dominates any of the existing labels. Dominated labels are removed from the list of labels, and are thus not extended further.

Two labels can only dominate each other if they have the same start port and are currently located in the same port. A label is dominated by another label if it has used more or an equal number of days *and* has a lower or equal profit. This is because it is impossible for the dominated label to make higher profits than the dominating one. The dominating label has the same starting point for future trade as the dominated one, but more time left and a higher profit to start off with.

Figure 5.7 shows an example of domination. The two labels are comparable because both started in port 1 and are currently located in port 5. Label 15 dominates label 26 because it has spent fewer days *and* has a higher profit. Note that the route sailed up to the last port does not matter. For the label expansion shown in Figure 5.6 we have that label 15 dominates label 16. Label 17 does *not* dominate label 15, as it has spent more days.

Basis cost are not added to the profit until all labels has been expanded. This prevents labels for a later day to be dominated by an earlier label because of basis cost. In Figure 5.6 this could have happened with label 15 and 16. If basis cost of $100,000 pr day was included for all labels, the real profit of label 15 would be less than the profit of label 16. It would still be dominated by our algorithm.

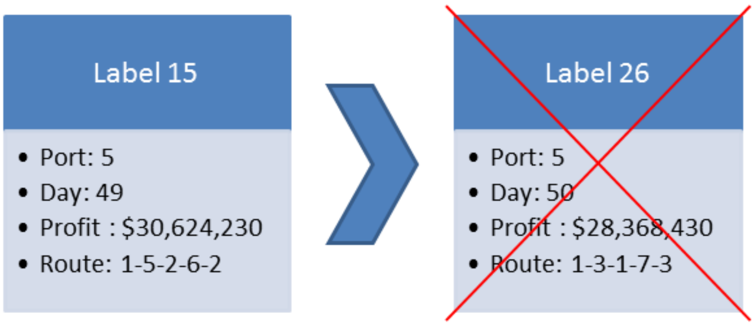

Figure 5.7: Example of domination of labels in the DP

5.2.4 Length of trade horizon

The length of the trade horizon dictates how many days the vessel has at its disposal. Each day can either be used for moving towards a trade port, awaiting better prices in a port or as flexible days at the end of the trade horizon. We define flexible days as days remaining in the end of the trade horizon after finishing the trades early. This is valuable because it gives the vessel a better starting point for trades taking place after the end of the trade horizon. These trades are not considered by the dynamic program, but are relevant when solving a problem in combination with either the MIP or heuristic. In this situation the DP only solves a limited part of the total problem period.

The benefit of finishing trades early is not automatically considered by the DP, as it does not increase the profits generated in the current trade horizon. End-of-horizon values are added to make up for this. They are discussed in Section 5.2.5.

There is a trade-off between making another trade, awaiting better prices and finishing early. The value of making another trade is apparent. It gives profit based on the difference between forecasted buy and sell price. The benefit of waiting in between trades is that it is possible to get better prices in the ports. The estimated value of waiting is decided by the variations in the price forecasts. The value of finishing early is based on an estimated value of future trade, as given by the EOH-values. The model makes a trade-off between the three alternatives based on these value estimations. Normally it performs the maximum number of trades, and then use the remaining days for a combination of awaiting good prices and finishing early.

5.2.5 End-of-horizon values

When run in combination with the MIP or heuristic, the DP only considers a limited part of the total problem period. Trades taking place after the end of this period are not included in the DP. The vessel should still be compensated for finishing its last trade early, as it enables the vessel to start off earlier on a trade taking place after the current trade horizon. This is valid until we approach the end of the total problem period, after which no more trades are made.

Finishing the last trade early does not give direct value in the DP, as it does not increase the profits of the trades made. End-of-horizon values (EOH-values) are added to make up for this, assigning a value to spare days at the end of the trade horizon. The value is only given for unused days in the last 10 days of the trade horizon. This is to prevent the vessel from not making actual trades because of the EOH-values.

Figure 5.8 shows value for a port as a function of remaining days in the trade horizon. Please note that this graph is only meant as an illustration, and does not represent actual numbers. Only the last segment of the trade horizon length is shown. The jump in Figure 5.8(a) is a sale being made. The reason for increases in the graph apart from the trade is the flexible days. They are used to get better prices in the trades prior to the one shown. The graph does not increase steadily. This is because extra days does not necessarily give better prices in the ports. Price development can also be negative, but then the DP will keep the original trade day. That is why the value never decreases.

From 5.8(b) we see that EOH-values reduces the effect of the sale on port value. This is positive, as it reduces the effect of specific trade horizon length on port values. We would like the vessel to travel a route that includes this port, even though it does not have time to make the trade within the current trade horizon. It is important to remember that the trade horizon only considers a limited part of the full problem period. Results from tests with and without EOH-values are presented in Chapter 9.

The vessel is also compensated for buying at the end of the trade horizon. This is done to make buying a valid last action. If no compensation was given for buying a load in the end of the trade horizon, all trade horizons would end with a sale. This is because buying has negative value and decreases the profit. A route ending with a buy would thus never be considered the best route. The buy compensation is also only valid until we approach the end of the total time period, as the vessel should sell in its final trade. The value is

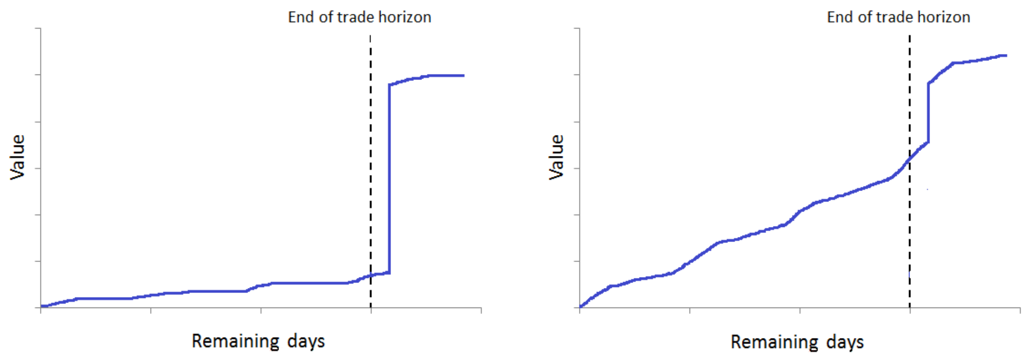

(a) Value graph without EOH-values (b) Value graph with EOH-values

Figure 5.8: Values in a port as a function of remaining days in the trade horizon

given as a multiple of the buy price. The multiple reflects the expected future payback of selling the load that was bought.

5.3 Mixed Integer Program (MIP)

In this section we present a Mixed Integer Program (MIP) used to decide daily vessel actions. It is used in combination with the DP when solving the full problem. The MIP takes port values from the DP as input, and decides which moves to make in the short-term horizon. The interaction between the MIP and DP is described further in Section 6.2. The MIP uses all assumptions stated in Section 5.1.

5.3.1 General description

The MIP is used to decide optimal actions for the vessel. The decision is based on the position and status (empty or full) of the vessel, and port values provided by the DP. Actions considered by the MIP are moving, waiting or trading. A multi-day version of the MIP that can solve the problem without using the DP is shown in Appendix C. The stand-alone model has not been used in testing due to high run time.

5.3.2 Mathematical model

A stochastic model is presented. We start by introducing sets and indices, before looking at parameters and constraints. Finally we present the objective function and constraints.

Sets and indices

$\mathcal{G}$	- Grid points, g
$\mathcal{G}^{\mathcal{N}}(g)$	- Neighboring grid points of g, $\tilde{g}$
$\mathcal{G}^{\mathcal{S}}(g)$	- Straight neighboring grid points of g, $\tilde{g}$
$\mathcal{G}^{\mathcal{D}}(g)$	- Diagonal neighboring grid points of g, $\tilde{g}$
$\mathcal{G}^{\mathcal{P}}$	- Ports, g
$\mathcal{T}$	- Time interval, t
$\mathcal{S}$	- Scenarios, s
$\mathcal{K}_t$	- Index set of scenario subsets at time t, k
Ω_{kt}	- Subset of scenarios at time t, ω

The sets of grid points are divided into multiple subsets. Neighboring grid points of g are the points that can be reached in one time period of traveling. We distinguish between straight neighbors of grid point g, and diagonal neighbors of grid point g. Figure 5.9 shows the placement of straight and diagonal neighbors, relative to the vessel. Straight neighbors are given as $\mathcal{G}^{\mathcal{S}}(g)$, while diagonal neighbors are given as $\mathcal{G}^{\mathcal{D}}(g)$. All neighbors of grid point g are given as the set $\mathcal{G}^{\mathcal{N}}(g)$. $\mathcal{G}^{\mathcal{P}}$ are grid points with ports. Decisions made at time t depict what is done in the time interval t to $t+1$ (as seen in Figure 5.4). s indicates which scenario a variable belongs to. Ω_{kt} are subsets of scenarios. ω is a variable combining time and scenario. ω points to a subset of scenarios at a time t. $\mathcal{K}_t$ is the index set of the scenario subsets at time t. The index sets show which scenarios that belong to each time set in each time interval.

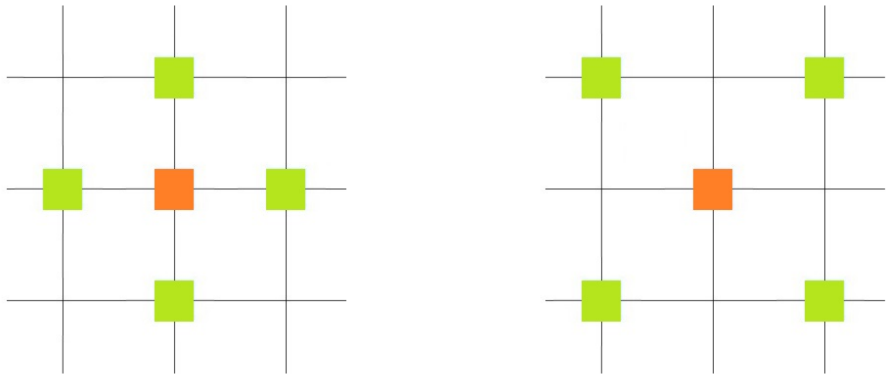

(a) Straight neighbors

(b) Diagonal neighbors

Figure 5.9: Straight and diagonal neighbors (green) relative to the vessel (orange)

Parameters

Q	- Capacity of the vessel
C^B	- Basis cost for the vessel (regardless of moving or not)
C^{MS}	- Extra cost for the vessel if it is moving straight
C^{MD}	- Extra cost for the vessel if it is moving diagonal
V_{gts}	- Value of port g at time t in scenario s

Q is the capacity of the vessel. C^B is the basis cost of the vessel. This is the cost of operating the vessel one day regardless of the vessel action. C^{MS} and C^{MD} are extra costs for the vessel if moving straight or diagonal, respectively. V are port values. The port value is the value the vessel receives when conducting a trade in a port. These value indicate how much the vessel can earn in the future by conducting the first trade in a port at the given time.

Variables

x_{gts} - 1 if vessel is at grid point g at time t in scenario s, 0 otherwise
m^S_{ts} - 1 if vessel moves straight from time t to time $t+1$ in scenario s, 0 otherwise
m^D_{ts} - 1 if vessel moves diagonally from time t to time $t+1$ in scenario s, 0 otherwise
r_{gts} - 1 if vessel trades in grid point g from time t to time $t+1$ in scenario s, 0 otherwise
$x_{g\omega}$ - Variables used as nonanticipativity constraints for position
$r_{g\omega}$ - Variables used as nonanticipativity constraints for trade

Objective function

$$\text{maximize} \sum_{g\in\mathcal{G}^{\mathcal{P}}}\sum_{t\in\mathcal{T}}\sum_{s\in\mathcal{S}} V_{gts}\cdot r_{gts}\cdot Q - \sum_{t\in\mathcal{T}}\sum_{s\in\mathcal{S}}(C^B + C^{MS}\cdot m^S_{ts} + C^{MD}\cdot m^D_{ts})$$

The objective function maximizes profit. The first part of the objective function sums the values from the trades. The second part of the objective function subtracts the basis cost and the costs of the travelling between the ports.

Constraints

$$\sum_{g\in\mathcal{G}} x_{gts} = 1, \qquad t\in\mathcal{T}, s\in\mathcal{S} \tag{5.1}$$

Constraints (5.1) prevent the vessel from being in more than one grid point at a time. Since the position variables (x_{gts}) are binary, only one of them can take the value 1 each time period.

$$x_{g(t+1)s} - \sum_{\tilde{g}\in\mathcal{G}^{\mathcal{N}}(g)} x_{\tilde{g}ts} \leq 0, \qquad g\in\mathcal{G}, t\in 1\ldots\mathcal{T}-1, s\in\mathcal{S} \tag{5.2}$$

Constraints (5.2) ensure that the vessel is only allowed to move to grid points in its neighborhood in the time interval t to $t+1$. In order for the vessel to be in g at time $t+1$, it needs to have been in the neighborhood of g in time t.

$$m^D_{ts} - x_{g(t+1)s} - \sum_{\tilde{g}\in\mathcal{G}^{\mathcal{D}}(g)} x_{\tilde{g}ts} \geq -1, \qquad t\in 1\ldots\mathcal{T}-1, s\in\mathcal{S} \tag{5.3}$$

$$m^S_{ts} - x_{g(t+1)s} - \sum_{\tilde{g}\in\mathcal{G}^{\mathcal{S}}(g)} x_{\tilde{g}ts} \geq -1, \qquad t\in 1\ldots\mathcal{T}-1, s\in\mathcal{S} \tag{5.4}$$

Constraints (5.3) and (5.4) ensure that the move straight or move diagonal variable is set to 1 when the vessel has moved straight or diagonal, respectively, in the time interval t to $t+1$. If the vessel is in grid point g in time $t+1$, and was in the diagonal neighborhood of g in time t, m_{ts}^D is set to 1. The same principle applies for the straight neighborhood and m_{ts}^S.

$$\sum_{g \in \mathcal{G}^{\mathcal{P}}} \sum_{t \in \mathcal{T}} r_{gts} \leq 1, \qquad s \in \mathcal{S} \tag{5.5}$$

Constraints (5.5) ensure that the vessel is only allowed to trade once during the time horizon. For each scenario, the total sum of all trades is less or equal to 1.

$$r_{gts} - x_{gts} \leq 0, \qquad g \in \mathcal{G}^{\mathcal{P}}, t \in \mathcal{T}, s \in \mathcal{S} \tag{5.6}$$

Constraints (5.6) prevent the vessel from trading when it is not present in a port. The trade variable (r_{gts}) of a port has to be lower than the position variable of the same port.

$$m_{ts}^D + m_{ts}^S + \sum_{g \in \mathcal{G}^{\mathcal{P}}} r_{gts} \leq 1, \qquad t \in \mathcal{T}, s \in \mathcal{S} \tag{5.7}$$

Constraints 5.7 ensure that the vessel has to spend one day in a port while trading. If the vessel trades, it can neither move straight nor diagonal.

$$x_{gts} - x_{g\omega} = 0, \qquad t \in \mathcal{T}, k = 1 \ldots \mathcal{K}_t, \omega \in \Omega_{k(t+1)}, \tag{5.8}$$

$$r_{gts} - r_{g\omega} = 0, \qquad t \in \mathcal{T}, k = 1 \ldots \mathcal{K}_t, \omega \in \Omega_{kt}, \tag{5.9}$$

Constraints (5.8)-(5.9) are nonanticipativity constraints (NACs). ω include both scenario and time, and is therefore sufficient for describing which variables should be equal. It is important to notice that the NACs in time t should force the position variable to be the same in time $t+1$ for all scenarios in the same subset. In order for a move in time period t to $t+1$ to be equal between two scenarios, the position of the vessel in the scenarios has to be equal at time $t+1$.

For the trade variables, the NACs are constraining in the same time period as the trade happens. If a trade happens at time t in one of the scenario in the subset, the NACs ensure that the trade happen at time t in the other scenarios in the same subset. The difference in timing of the NACs is shown in Figure 5.10. The times where the NACs are constraining are shown with red circles.

$$r_{gts} \in \{0,1\}, \qquad g \in \mathcal{G}^{\mathcal{P}}, t \in \mathcal{T}, s \in \mathcal{S} \tag{5.10}$$

$$x_{gts} \in \{0,1\}, \qquad g \in \mathcal{G}, t \in \mathcal{T}, s \in \mathcal{S} \tag{5.11}$$

$$m^S_{ts}, m^D_{ts} \in \{0,1\}, \qquad t \in \mathcal{T}, s \in \mathcal{S} \tag{5.12}$$

Constraints (5.10)-(5.12) are binary constraints on the variables.

(a) Position variables (b) Trade variables

Figure 5.10: Timing of nonanticipativity constraints for position and trade variables

Removal of redundant binary constraints

The x-variables that are out of reach for the vessel have their binary constraint removed. By using this approach, the position variables become binary when they need to be binary. This principle is shown in Figure 5.11. The grid points one day away from the vessel are not binary until day two, while the grid points one step further out is not binary until day three, and so on.

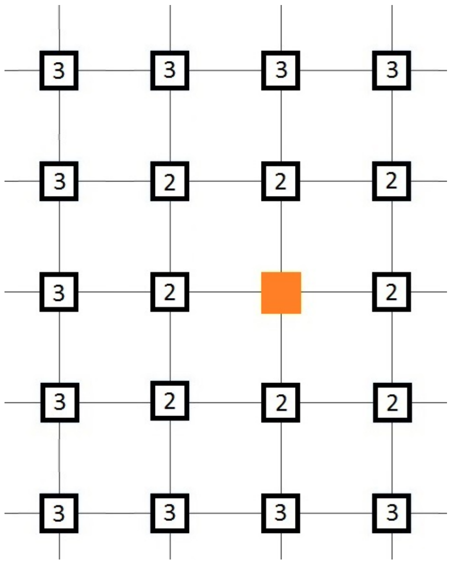

Figure 5.11: Overview of which day the position variables close to the vessel (orange) become binary

5.4 Heuristic

The heuristic presented in this section is an alternative to the MIP for selecting daily vessel actions. It is used in combination with the DP when solving the full problem. The heuristic take port values from the DP as input, and returns a move for the next day. This interaction is described further in Section 6.3. The heuristic uses all assumptions stated in Section 5.1. Both a deterministic and a stochastic version is presented.

5.4.1 General description

In the same way as the MIP, the heuristic makes decisions on whether the vessel should trade, move or stay put. When in a trade port, the vessel has to choose between trading right away or waiting for the next day. When at sea, the vessel has to choose between staying in the same position or moving in one of eight possible directions. It makes the move decision based on the estimated values of travelling to trade ports, as given by the DP. A specific move is more likely if multiple ports with good values are located in that direction.

We first take a look at the movement decisions made, before shortly considering the timing of trade-decision. Examples are included to demonstrate how the heuristic works.

5.4.2 Calculating the value of possible moves

The movement decision made by the vessel reflects where the profit potential is predicted to be highest. This is based on the estimated value of trade ports located in the direction of travel. When choosing its next action, the heuristic calculates values for all eight neighbouring grid points, and the value of staying put. These are the only nine options the vessel has. The points are shown as triangles in Figure 5.12. For each point there are different ports included in the value calculation. This is described in the next paragraphs.

Ports included in value calculation

Value is provided by trading in a port. The benefit of moving in a direction is that it brings you closer to one or more ports. Each of these ports has the potential to be the first trade port in a sequence of trades.

Two alternatives are used when selecting which ports to include in the value calculation of a move. The first alternative is to only include ports that are closer (in travel days)

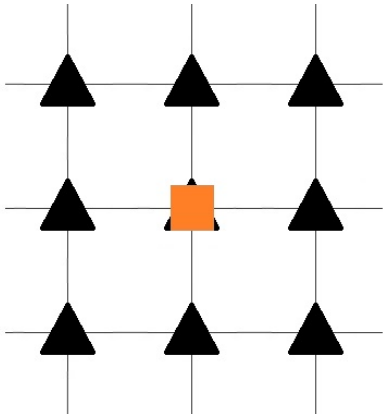

Figure 5.12: Possible options (triangles) for the vessel (orange square)

to the vessel after the move. Figure 5.13 shows the areas included for this alternative. The move considered is marked by a red cross. All ports that are located inside a green field are included in the value calculation. Only a limited grid size is shown, but the area would expand in the same manner for larger grids.

The second alternative is to also include ports that are the same number of days from the vessel after the move. Figure 5.14 shows which grid points are included in this calculation. Ports that you move away from are never included in the value calculation of a move. The values of all ports are included for the option of staying put. Both alternatives are tested for performance in Chapter 9.

Figure 5.15 shows an example of a map with a possible move marked by a red cross. Here the second alternative is used, including all ports that we are not moving away from. The considered ports are marked in blue, while the rest of the ports are marked in yellow.

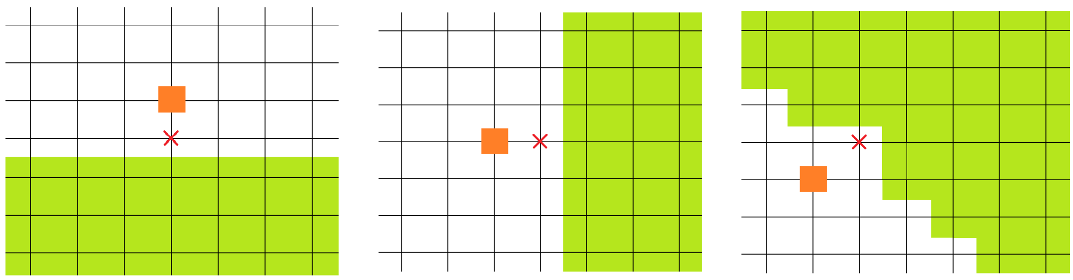

(a) Lower (b) Right (c) Upper right

Figure 5.13: Area where ports are considered when including ports with fewer travel days

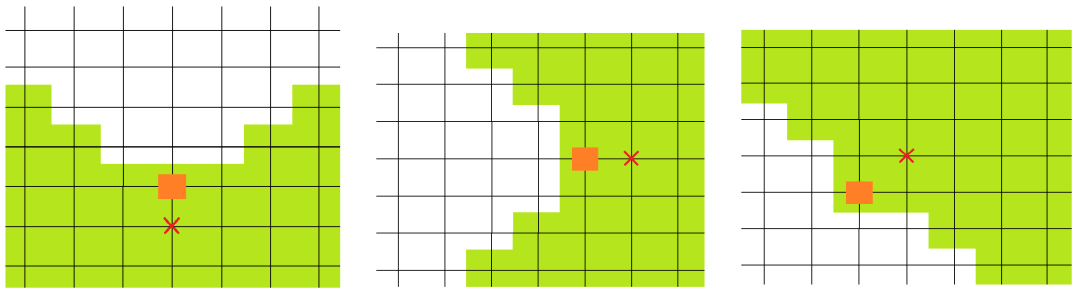

(a) Lower (b) Right (c) Upper right

Figure 5.14: Area where ports are considered when including ports with the same or fewer travel days

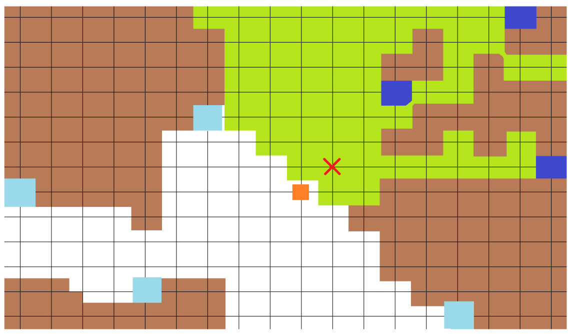

Figure 5.15: Example of ports considered for a northeast move in Atlantic Basin-map. Potential ports are dark blue. Non-potential ports are light blue.

5.4.3 Stochastic vs deterministic

In this section we present two examples that show the difference between the stochastic and deterministic heuristic.

Deterministic heuristic example

In the following example we use the map shown in Figure 5.16. The current position of the vessel is marked by an orange square. The blue squares represent ports. Three potential new positions are shown as green triangles. We assume that the vessel is in the current position on day 1. It can thus be in any of the three new positions on day 2.

Table 5.2 shows the ports that are included when calculating the value of each potential move. It also lists the first day a trade can be made in each of the ports. The second alternative for port inclusion is used, including all ports that the vessel is not moving away from. Port 2 is not included in the value calculation of position 1, as moving there brings the vessel one day further away from this port.

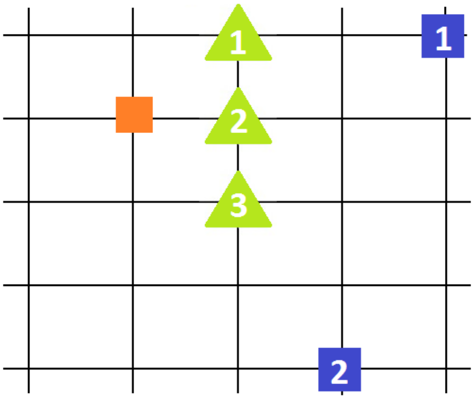

Figure 5.16: Example with two ports and three possible move options

Table 5.2: Ports considered for each possible new position

Position	Ports considered	Earliest trade day port 1	Earliest trade day port 2
1	1	4	6
2	1,2	4	5
3	1,2	4	4
Stay put	1,2	5	5

The values of the ports represent the future value of starting trade in the given port on the given day. Table 5.3 shows the estimated value of starting trade in port 1 and 2 on day 4 and 5.

Table 5.3: Forecasted port values in thousand $ for port 1 and port 2 in day 4 and 5

Port	Day 4	Day 5
1	$10,300'	$10,250'
2	$10,500'	$10,480'

The potential profit of being in a position is decided by the port values, as well as the travel time and the travel cost of moving to the ports. The travel time dictates which of the port values in Table 5.3 can be used, while the travel cost have to be subtracted to find the actual potential profit. Table 5.4 shows the value contribution from port 1 for each of the potential positions. Note that the stay put option is using the port value for day 5. This is because the earliest day the vessel can reach port 1 after a waiting is day 5. Table 5.5 shows the same information for port 2. Position 1 does not have a value for port 2, as moving to position 1 brings the vessel away from port 2. The value given to a position in the deterministic heuristic equals the highest of these port contributions. The final position values are listed in Table 5.6.

Table 5.4: Potential value contribution in thousand $ from port 1 for the different positions

Position	Port value	Travel cost	Position value
1	$10,300'	$200'	$10,100'
2	$10,300'	$240'	$10,060'
3	$10,300'	$280'	$10,020'
Stay put	$10,250'	$340'	$9,910'

Table 5.5: Potential value contribution in thousand $ from port 2 for the different positions

Position	Port value	Travel cost	Position value
1	-	-	-
2	$10,480'	$340'	$10,140'
3	$10,500'	$280'	$10,220'
Stay put	$10,480'	$380'	$10,100'

Table 5.6: Position values in thousand $

Position	Port 1 value	Port 2 value	Position value
1	$10,100'	-	$10,100'
2	$10,060'	$10,220'	$10,220'
3	$10,020'	$10,100'	$10,100'
Stay put	$9,910'	$10,100'	$10,100'

We now have the value of being in each of the potential new positions. What we actually need is however the value of each potential move. This is equal to the position value, less the cost of moving there. The final move values are listed in Table 5.7. We see that moving to position 2 gives the highest value, and is therefore chosen in this example.

Table 5.7: Move valuesin thousand $

Move to position	Position value	Travel cost	Move value
1	$10,100'	$140'	$960'
2	$10,220'	$100'	**$10,120'**
3	$10,100'	$140'	$9,960'
Stay put	$10,100'	$0	$10,100'

Running the deterministic model with multiple price scenarios is done by using the expected value of all price scenarios. The stochastic model handles multiple price scenarios in a different way. This is discussed in the next paragraphs.

Stochastic heuristic example

This section describes how the stochastic version of the heuristic works. An example is used to describe the approach. The example grid used is shown in Figure 5.17. The stochastic example has 21 scenarios. We only consider moving west to position 1 and moving east to position 2. For these two moves it does not matter which alternative for port inclusion is used, as all ports either get closer or further away. Port 1 is included when calculating the value of moving to position 1, while ports 2 and 3 are included when calculating the value of moving to position 2.

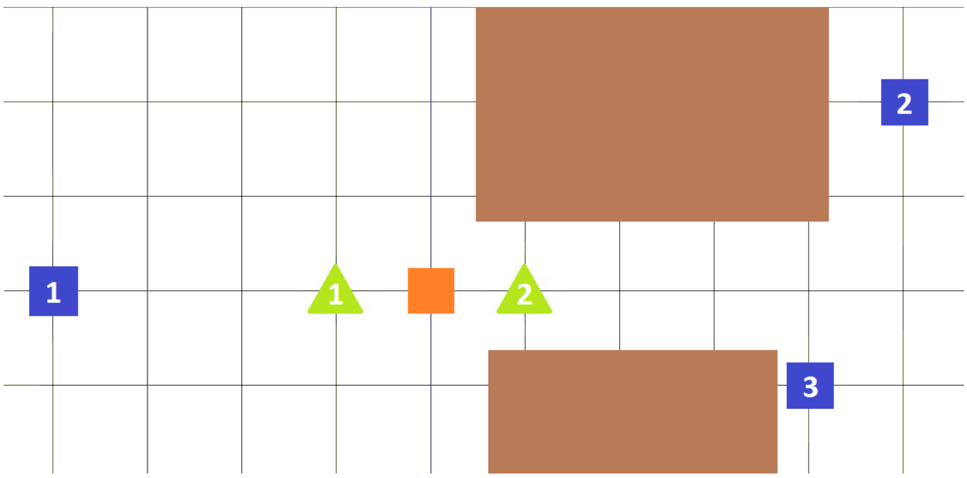

Figure 5.17: Example with three ports and two possible options of moving

Figure 5.18 shows the estimated value of going to the different ports for various scenarios. The values for each port in each scenario are calculated the same way as the grid point values in the previous example. The port values for port 1 and 3 are from day 5 and the port values for port 2 are from day 7. The dotted line shows the expected value of all scenarios. We see that port 1 has the highest expected value.

The stochastic heuristic does not use the expected value, but rather tries to incorporate all available information. This is done by using the best value among the considered ports for each scenario. By doing so we are able to incorporate the positive effect of having multiple ports located in the same direction. Figure 5.19 shows how the graphs of the two ports to the east are combined to evaluate the value of moving east. The combined expected value is higher than for any of the two ports alone. If there are more than two ports pulling in the same direction, the maximum of all these ports are calculated.

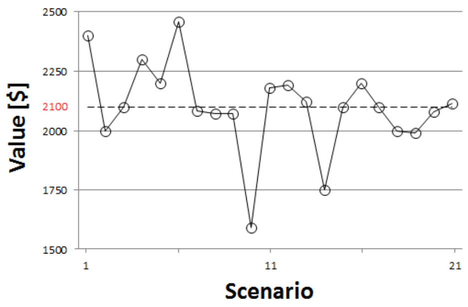

(a) Port 1's attraction of going west

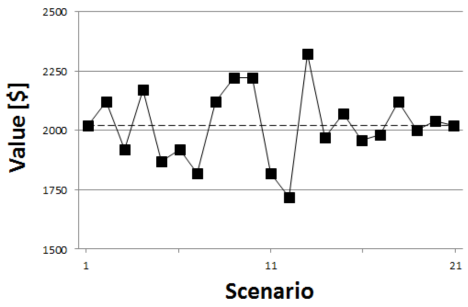

(b) Port 2's attraction of going east

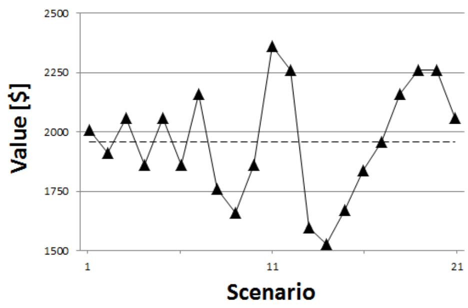

(c) Port 3's attraction of going east

Figure 5.18: Values of going west or east in different scenarios

Table 5.8 shows the value of moving west and moving east. We see that the vessel will travel east in the stochastic version. This is opposite of the deterministic model, which would choose to travel west since port 1 has the highest expected value.

Table 5.8: Action values in thousand $ for each of the potential moves for stochastic version

Move to position	Position value	Travel cost	Action value
West	$2,100'	$100'	$2,000'
East	$2,130'	$100'	**$2,030'**

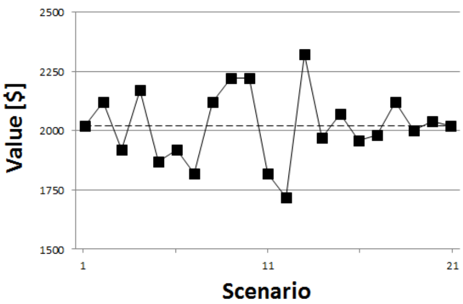

(a) Port 2's attraction of going east

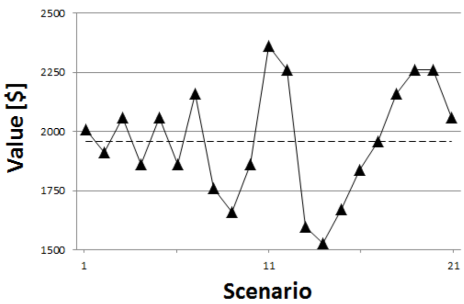

(b) Port 3's attraction of going east

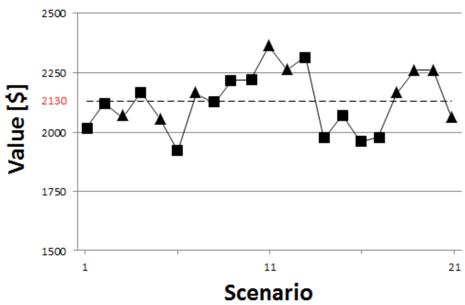

(c) Combined attraction of going east

Figure 5.19: Value of going east in stochastic problem

Chapter 6

Solution Approaches

In this chapter we present two complete solution approaches. Both approaches are based on the models explained in the previous chapter. One combines the DP with the MIP, while the second combines the DP with the heuristic. We first make some general remarks about the interaction between the DP and the two other models. Following this we take a closer look at the combination of the DP and MIP model, before considering the combination of the DP and heuristic. An alternative model which solves the multistage problem using only the MIP is presented in Appendix C.

6.1 Interaction between the models

Using the DP in combination with either the MIP or heuristic model allows us to combine short-term and long-term considerations in an efficient way. The DP is responsible for the long-term aspect. It does not consider the total problem period, but rather finds optimal routes for a limited trade horizon. The period that is considered shifts forward as the problem progresses. We have defined this as a rolling horizon approach. It effectively reduces the complexity of the problem and enables us to solve it faster. We only consider trades d days ahead, even though the problem we are solving is over D days. The principle is shown in Figure 6.1. Each day we move the rolling horizon frame one day ahead. If we are e.g. in day 35 with a rolling horizon length of 60 days, we consider trades from day 35 to 95.

Given a start port and the travel time, the DP finds an optimal route and corresponding profit for the rest of the trade horizon. This is done efficiently by assuming that the vessel moves directly between the ports that look most profitable, given the current price forecast.

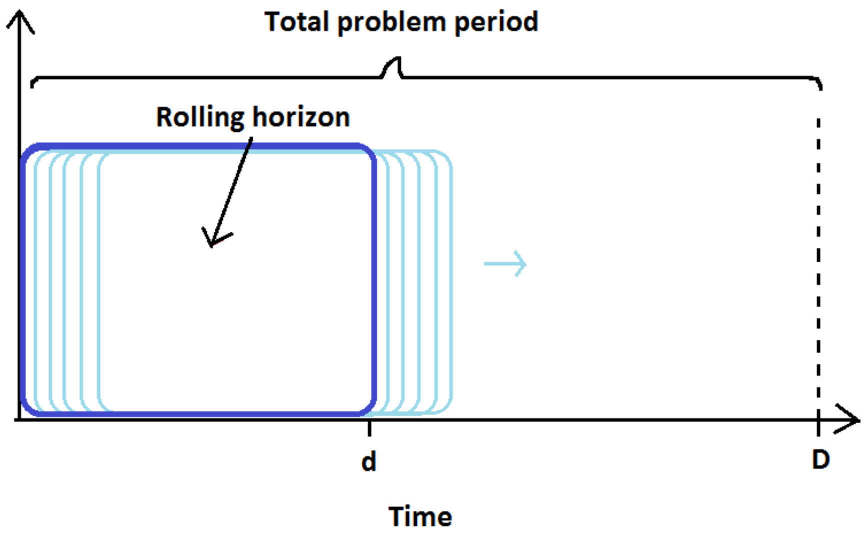

Figure 6.1: Graphical overview over the rolling horizon method

An example of a route calculated by the DP is shown in Figure 6.2. Buy ports are red and sell ports are blue. The vessel is marked by an orange circle. We assume that the vessel is in the current position on day 1. The first trade is made in port 1 on day 4. The dynamic program finds the optimal route to travel after this first trade. The arrows indicate the route of the vessel (1 - 5 - 2 - 6 - 3). The profit of this route corresponds to the value of being in port 1 on day 4.

In a full model run we find optimal routes for initiating trade on subsequent days as well. These routes can potentially have higher profits than for trading straight away, given that prices in the start port develop in a beneficial way. Making the first trade on a later day does however leave fewer days for trading in the rest of the trade horizon, which often leads to lower port values. The DP is run for each potential trade port, calculating routes and corresponding profits for different start days, up to a maximum number of waitdays. A value is assigned to a port for each of these days. This value represents what the long-term profit is expected to be if the vessel conducts the first trade in the port the given day.

The MIP and heuristic make the short-term movement decisions. Based on the port values from the DP, they calculate the best vessel action for the current day. If in a port, it is a trade decision. If at sea, it is a movement decision. The vessel action chosen is returned to the DP, which uses the updated information to calculate new port values for the next day.

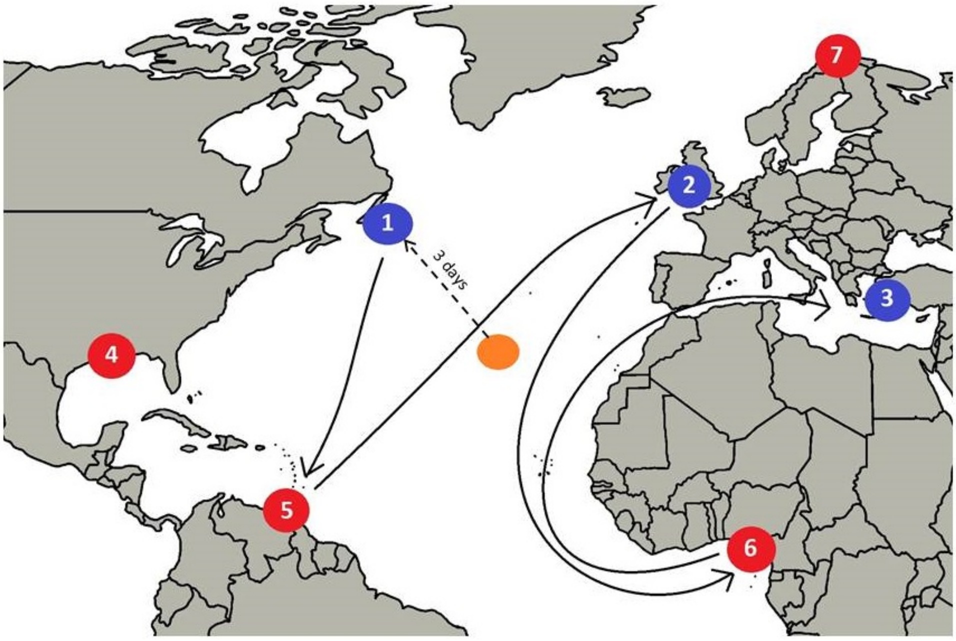

Figure 6.2: Example of optimal route for starting trade in port 1 on day 4, found by the DP

Decisions made the last days of the total problem period

The rolling horizon is made shorter towards the end of the problem period. No trades beyond the problem period are considered. The end of horizon-values are also removed towards the end. This ensures that the vessel sells as its last action.

6.2 Dynamic Program and Mixed Integer Program

The DP takes the position, the state of the vessel (empty or full) as input, and returns values for each potential trade port. If the vessel is full it only returns value for the sell ports. If the vessel is empty, it returns values for the buy ports. Ports with value are called potential ports in the following discussion. Since the dynamic program tells us the value of all future trades after arriving in a port, the MIP only has to consider which port to sail to first. In other words, the DP has simplified the problem from optimizing multiple trades to only optimize the first trade. Based on the values from the DP, the MIP chooses an action for the vessel and sends an updated vessel position and status to the DP, which again updates values for the ports.

There might be situations were the vessel chooses to leave a port before trading in it. This is irrational, but can be explained by how he DP works. In a finite number of days it is only possible to complete a finite number of trades. This is likely to leave some flexible days, as the vessel probably does not need every day in the trade horizon to complete these trades. Sometimes it will be beneficial to spend these flexible days to visit a new port, trade there instead, and still reach the same number of trades as if starting trade in the original port. This leads to the DP returning high values for the other port. The MIP will thus move the vessel towards the other port rather than trading in the port it is currently in. Because of the large capacity of the vessel, small changes in price can tempt the vessel to do this. Inside the rolling horizon it is the correct move to do, but over the total problem period it is almost certain to give a bad result. The problem is partly handled by the EOH-values, which gives the vessel an incentive to finish its trades early. In addition we have added an extra constraint for the vessel. Once in a port, it is not allowed to leave without making a trade. The vessel can stay as many days in a port as it wants, but it has to trade before leaving. This is ensured by giving all other ports a value of 0 until a trade is completed. Including this constraint reduces the complexity of the problem when the vessel is in a port, as it now only has to decide the timing of the trade.

An example of interaction between the MIP and DP is shown in Figure 6.4. The orange square represents the vessel. The dark blue squares represent potential trade ports, while the light blue squares represent ports that are currently *not* potential trade ports. In day 1 we see that port 3 is the only potential trade port. The DP thus calculates values for this port and sends it to the MIP. The MIP then calculates a new position for the vessel based on this port value and returns the new position to the DP. After making the move the vessel is in day 2 and currently located in a potential trade port. The DP then calculates the value of trading the same day and the value of trading subsequent days, and returns these values to the MIP. If the value is highest for trading straight away, the MIP will choose to do so. If not, it will wait until the next day. The next day brings new prices, and the calculation will be made again. In our case the highest value is for trading straight away, and the vessel status is changed. We get new potential trade ports. The DP calculates values for the new potential ports and returns the values to the MIP. The MIP now has to make a new movement decision based on these values. It continues in this matter until the end of the total problem period.

6.3 Dynamic Program and heuristic

A flow diagram for the heuristic is shown in Figure 6.4. The flow diagram is explained in the paragraph below.

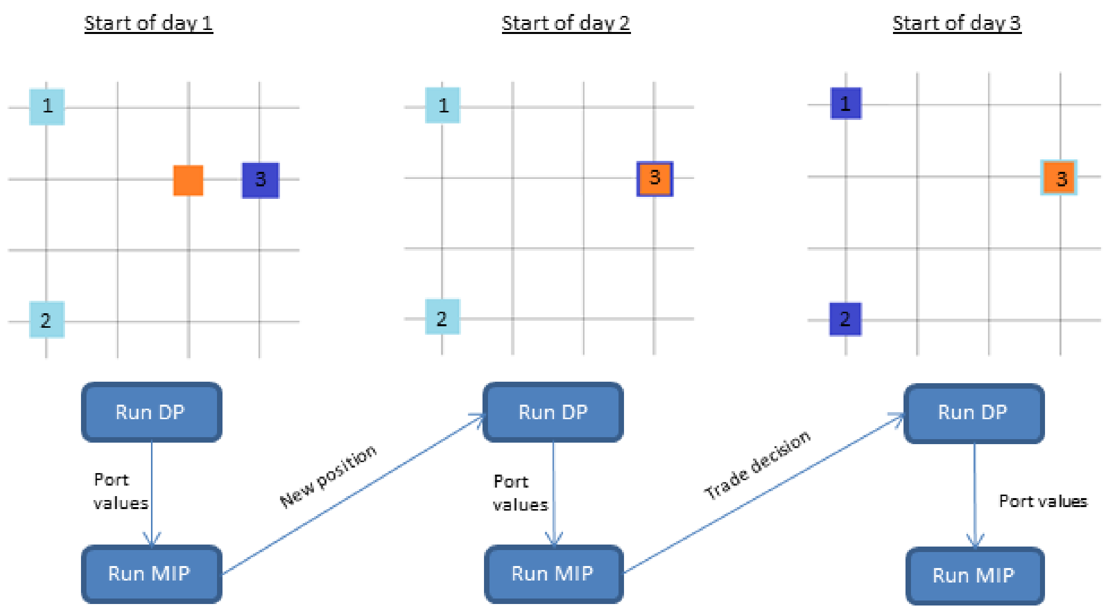

Figure 6.3: Interaction between MIP and DP. The vessel is orange, potential ports are dark blue and non-potential ports are light blue.

The heuristic takes a vessel position and status(full or empty) as input. It then checks whether the vessel is located in a potential trade port, meaning a sell port if it is full and a buy port if it is empty. If the vessel is in a potential trade port, it calculates whether a trade should be made straight away. This calculation is done using the dynamic program. The DP gives the future value of trading in the port on the current day, and the value of trading on subsequent days. This information is returned to the heuristic. If the value of trading on the current day is the highest, the trade will be made. If the maximum profit comes from one the later days, it will wait until the next day and get new price forecasts. The dynamic program then repeats the calculations based on the new price forecasts. As for the MIP, a vessel is not allowed to leave a potential trade port without making a trade. The reason is the same as for the MIP, and is explained in the previous section. If the vessel is not currently in a potential trade port, the heuristic will find a new position for the vessel. This decision is based on port values calculated by the dynamic program, as described in Section 5.4.2. After each decision we move on to the next day. With every new day we get an update of current prices and corresponding price forecasts for each port.

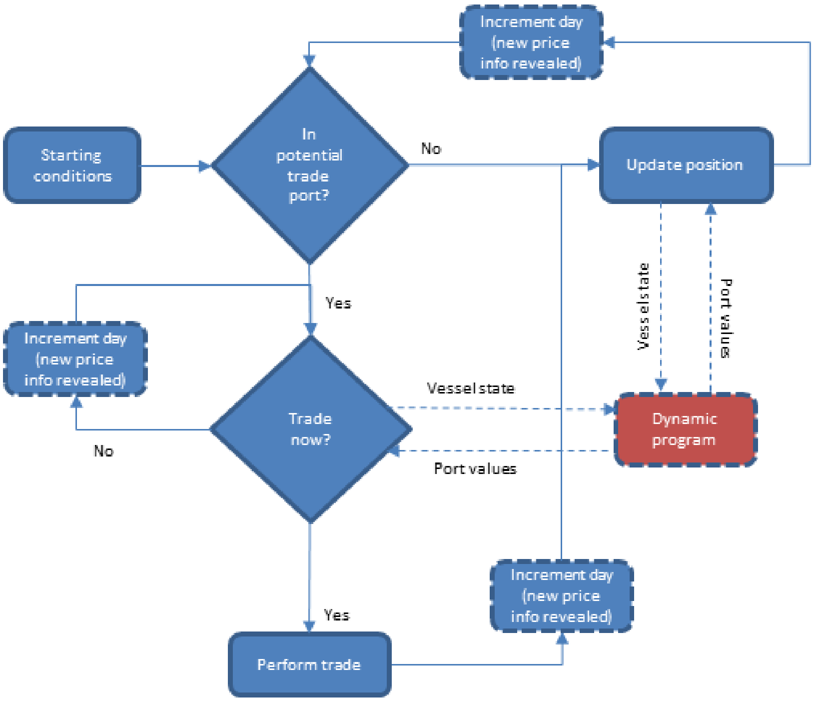

Figure 6.4: Flow diagram for heuristic approach

Chapter 7

Implementation

In this chapter we present a brief outline of the code used to solve our problem. We start with the implementation of models from Chapter 5, before presenting the solution approaches described in Chapter 6.

7.1 Models

In this section we present the implementation of the Dynamic Program, the Mixed Integer Program and the heuristic from Chapter 5.

7.1.1 Dynamic Program

The Dynamic Program (DP) for finding port values is implemented in Java. It takes length of trade horizon, current port prices, price forecasts, start port for trade and travel time to the start port as input. The length of the trade horizon dictates how many days of trade the DP considers. The price forecasts are for all ports for the extent of the time horizon. The start port is the first port where a trade is made.

The output of the dynamic program is the profit of the most profitable trade sequence that initiates trade in the given start port, for each choice of wait days in port. These values represent the potential value of starting trade in the given start port on each of the days.

An overview of the code is shown in Algorithm 1. The code is run for every potential start port of the vessel. Sell ports are potential start ports if the vessel is full, while buy ports are potential start ports if the vessel is empty. A description of how the model works is presented in Section 5.2, while the full code is found on the attached CD. The waiting days dictate how many days the vessel waits before trading, after arriving in a port. Labels contain information about the current state of the vessel. The fields of a label are day, profit and route. nrDays is the number of days in the trade horizon.

```
Input: Trade horizon, current prices, price forecasts, start port and travel time to
       port, b
for all waiting days, w = 0..n do // Waiting days before trading
    create label for start port with day set to b + w; // b is travel time to port
    add label to list of labels;
end
while new labels in list do
    for each new label do
        for all potential trading ports do
            for all waiting days, 0 .. n do
                create new label with updated fields;        // Day, profit, route
                add label to list;
                if label is dominated OR label_day > nrDays then
                    remove label from list;
                end
                else if label dominates otherLabel then
                    remove otherLabel from list;
                end
            end
        end
    end
end
Output: Profit of best route for each wait day option in start port
```

Algorithm 1: Dynamic program

7.1.2 Mixed Integer Program

Xpress-Mosel is used to implement the Mixed Integer Program (MIP). The MIP takes port values from the DP as input. Based on this it finds the most profitable route for the vessel. The output is the first day action of this route. This is a move decision if the vessel is currently at sea, or a decision to trade or not if the vessel is currently in a port. No pseduo-code is presented for the MIP, but the mathematical model is found in Section 5.3, and the full code on the attached CD.

7.1.3 Heuristic

The heuristic is implemented in Java. It takes port values from the DP and current vessel status (empty or full) and position as input and gives an updated vessel status and position as output. If the vessel is in a port to trade, the returned position will be the same as the current position. A short stepwise description is given Algorithm 2. A full description of the model is presented in Section 5.4, and the full code is found on the attached CD. waitToTrade is a boolean variable that says whether it is most profitable to make the trade straight away or to wait for better prices.

```
Input: Port values, vessel status, position
if In potential trade port then
    calculate waitToTrade;
    if waitToTrade == true then
        break;
    end
    else if waitToTrade == false then
        make trade;
        update vessel status;                    // Empty or full
    end
end
else
    calculate grid point values based on port values;
    choose new position based on grid point values;
end
Output: New vessel position and status
```

Algorithm 2: Heuristic

7.2 Solution approaches

7.2.1 Dynamic Program and Mixed Integer Program

The DP and MIP are combined through a BASH server script. This is done because we need to run Java and Xpress interchangeably. An outline of the server script is shown in Algorithm 3. A description of the solution approach is found in Section 6.2. The full code can be found on the attached CD. Price data contains the current prices and a forecast of future prices for all ports and scenarios. The output of the program is the route driven by the port and the corresponding profit, as found in movement.txt. totalProblemPeriod decides how many days the problem is solved for. This is not the same as the trade horizon in the DP. The trade horizon in the DP dictates how many days ahead the DP considers when calculating port values. This period is shorter than the total problem period.

```
Input: Initial status of vessel stored in MIP.txt
currentDay = 1;
while currentDay < totalProblemPeriod do
    import price data for current day;
    run DP (Java) using price data and vessel status from MIP.txt. Store port values
    in verdiData.txt;
    run MIP (Xpress) with port values from verdiData.txt and store updated vessel
    status in MIP.txt;
    update movement.txt with vessel position and profit;
    currentDay ++;
end
Output: Route and profit in movement.txt
```

Algorithm 3: DP and MIP

7.2.2 Dynamic Program and heuristic

The DP and heuristic code are both written in Java and can be run as a combined program. The work flow of the code is described in Algorithm 4. A description of the approach is found in Section 6.3. The code is found on the attached CD. Price data and totalProblemPeriod in the algorithm are the same as for the pseudo code with the DP and MIP.

```
Input: Price data
currentDay = 1;
while currentDay < totalProblemPeriod do
    if in potential trade port then
        run DP to get port values;
        calculate waitToTrade based on port values;
        while waitToTrade == true do
            currentDay++;
            run DP to get port values;
            calculate waitToTrade based on port values;
        end
        make trade and update profit;
    end
    else
        run DP to get port values;
        run heuristic to update position based on port values;
        currentDay++;
    end
end
Output: Route and profit
```

Algorithm 4: DP and heuristic

Chapter 8

Test Instances

In this chapter we describe the test instances used to examine the performance of our models. We start by presenting the map used, including the locations of buy and sell ports. Following this we take a look at the price forecasting method applied. We then discuss the parameters used for testing our models. The parameters are split into two groups; the ones that remain fixed throughout all tests, and those that are changed. Stability tests used are then discussed, before concluding the chapter by presenting the parameter values used in the base case test instance. Results from the tests are presented in Chapter 9.

8.1 Map

We use the Atlantic Basin in our tests. The Atlantic Basin has the highest degree of spot trade, as discussed in Section 2.5.2, and is the market for which we have the most extensive historical price information.

A total of three buy ports and four sell ports are included. The main criteria for selecting ports has been to generate interesting options for the vessel. This is achieved by including ports from all corners of the Atlantic Basin. The ports are shown in Figure 8.1. Figure 8.2 shows the same map fitted to the grid used.

The ports we use in our test instances are Point Fortin (Trinidad and Tobago), Bonny (Nigeria), Hammerfest (Norway), Lake Charles (US), St Johns (Canada), Marmara (Turkey) and Milford Haven (UK). Travelling time between the ports are shown in Table 8.1. Hammerfest and Milford Haven are the two ports that are closest to each other, with a travel time of 5 days. Lake Charles and Marmara are furthest away from each other, with 17 days of travel time.

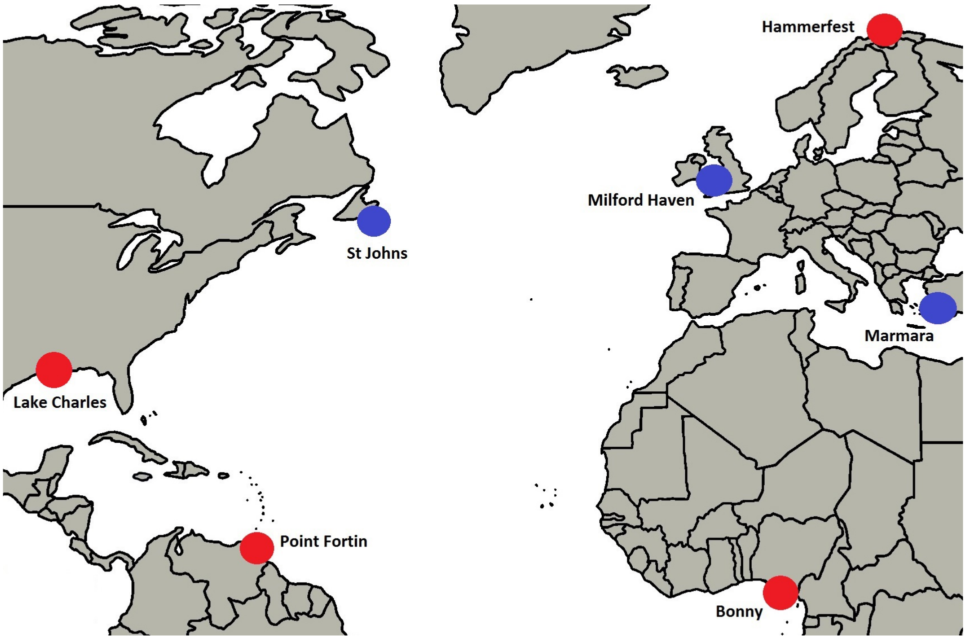

Figure 8.1: Ports used in test instance. Buy ports are red and sell ports are blue. (outline-world map.com)

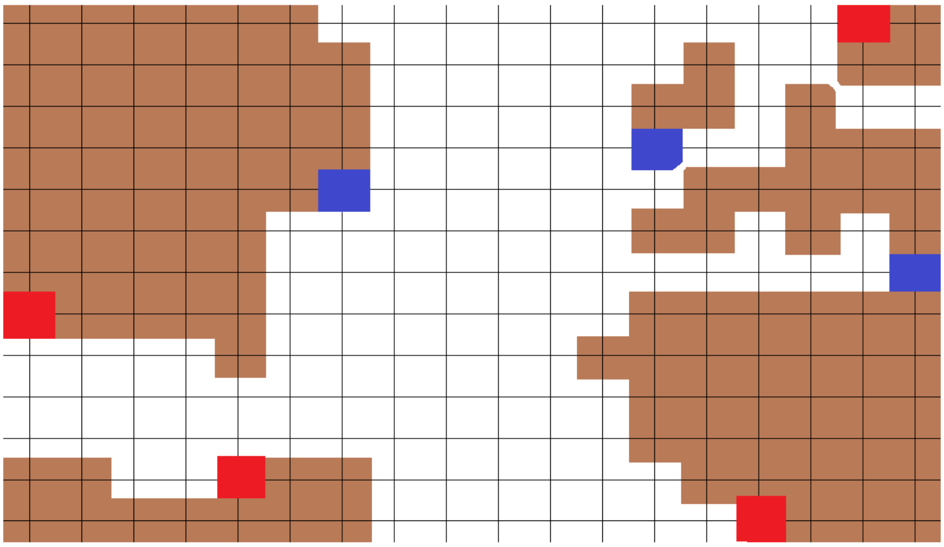

Figure 8.2: Map fitted to grid. Buy ports are red and sell ports are blue.

Table 8.2 shows the travel costs between the ports. Two trips are significantly cheaper than the rest. The trip between Point Fortin and St Johns (cost of M$0.78), and the trip from Hammerfest to Milford Haven (cost of M$0.66). It is important to note that there is *not* a linear relationship between travel time and travel cost. The reason for this is the diagonal moves in the grid. The diagonal moves use the same amount of time as straight moves, but has a higher cost.

Table 8.1: Travel time between ports (in days)

	St Johns (Canada)	Marmara (Turkey)	Milford Haven (UK)
Point Fortin (Trinidad and Tobago)	7	13	8
Bonny (Nigeria)	9	12	10
Hammerfest (Norway)	10	12	**5**
Lake Charles (US)	9	**17**	12

Table 8.2: Cost of travelling between ports (in million $)

	St Johns (Canada)	Marmara (Turkey)	Milford Haven (UK)
Point Fortin (Trinidad and Tobago)	**$0.78”**	$1.50”	$1.12”
Bonny (Nigeria)	$1.18”	$1.40”	$1.20”
Hammerfest (Norway)	$1.16”	$1.40”	**$0.66”**
Lake Charles (USA)	$1.06”	$1.90”	$1.52”

8.2 Ornstein-Uhlenbeck price process

An Ornstein-Uhlenbeck price process is used to forecast future prices in the ports. The process was originally presented by Leonard Ornstein and George Eugene Uhlenbeck (1930). We have made some modifications to the original process in order to make it fit our problem. These are correlation between ports port prices and lower price bounds. We have added correlation between ports because we assume that the price processes are not independent for each port. Due to the correlation we calculate price processes for all ports at the same time. This differs from the original process which calculates one price process at a time. Lower price bounds are used to prevent unrealistically low prices.

The Ornstein-Uhlenbeck price process takes a mean, variance and drift towards the mean as input. The drift towards the mean is an indicator of how much the price tends to return to its long time mean over time. Forecasts of future prices are calculated using the values of these parameters. The parameter values we use are based on on prices in the Henry Hub over the last five years.

Below we present the parameters and the model formulation. The parameters are split into two groups; input parameters and parameters generated randomly by the model. We implement correlation by correlating all the ports relative to one reference port. In the following formulation we have defined port 1 as the reference port.

Input parameters

S_{pt}	- price in port p at time t
C_p	- correlation between port 1 and port p
μ_p	- long-term mean of price in port p
σ	- variance of data, assumed to be equal for all ports
λ	- drift towards mean, assumed to be equal for all ports
t	- length of one time interval
T	- length of the full time period
L_p^{min}	- minimum limit for price at port p

Random parameters

r_{pt}	- random variable for port p in day t
G_{pt}	- random number from Gaussian distribution for port p at time t

Formulation

$$r_{1t} = \sqrt{\frac{1 - e^{-2\lambda \cdot t}}{2 \cdot \lambda}} \cdot G_{1t}, \forall t \in T \tag{8.1}$$

The Ornstein-Uhlenbeck method starts by generating a random number. This is later used to generate prices in the ports. Formula (8.1) generates a vector of random variables to be used for price development in the reference port. The formula uses a number from the Gaussian distribution to achieve randomness. This number is picked randomly for every time t. r_{1t} becomes the reference vector for the other ports. This reference is used when implementing correlation later in the formulation.

$$r_{pt} = C_p \cdot r_{1t} + (1 - |C_p|) \cdot \sqrt{\frac{1 - e^{-2\lambda \cdot t}}{2 \cdot \lambda}} \cdot G_{pt}, p \in 2 \ldots P, \forall t \in T \tag{8.2}$$

Formula (8.2) generates vectors of random numbers to be used for price development in all ports excluding the reference port. These random numbers are affected by the correlation coefficient between each port and port 1. The higher correlation, the higher the chance of getting the same random variable as the reference port. The calculation of correlation coefficients are described in Section 8.2.

$$S_{pt} = S_{p(t-1)} \cdot e^{-\lambda \cdot t} + \mu_p \cdot (1 - e^{-\lambda \cdot t}) + \sigma \cdot r_{pt}, \forall p \in P, t \in 2 \ldots TF \tag{8.3}$$

The random variables r_{pt} from Formula (8.2) are used in Formula (8.3) to decide the port prices in time 2 to T. The prices for $t = 1$ does not have to be calculated, as they are given as input. All other prices are generated by the price process. As seen from the formula, the prices generated depend on the price the day before.

$$\begin{aligned} &\text{if } S_{pt} < L_p^{min} \\ &\text{then } S_{pt} = L_p^{min}, \forall p \in P, \forall t \in 2 \cdot T \end{aligned} \tag{8.4}$$

Formula (8.4) is added to the Ornstein-Uhlenbeck process in order to set a minimum price limit. If the price generated is lower than the lower limit, the price is automatically set to the lower limit. This adjustment is done ongoing for every day, meaning that $L^m in_p$ becomes the starting point for the next day if the price originally was less than $L^m in_p$.

Correlation

This section covers how correlation between port prices is decided, and how it impacts price development.

Historical prices from the Henry Hub in the US and the National Balancing Point (NBP) in the UK are used to decide correlation coefficients. These are the only two locations for which we have good historical price data. We have used this data to calculate the correlation coefficient between the Henry Hub and NBP. By assuming that the correlation between two ports is based on an exponential function where the correlation *decreases* as distance *increases*, we have made a model based which we use to determine correlation between ports. We find it logical that the correlation changes in this way. The further away to ports are from each other, the less of the same factors impact them. This leads the correlation coefficient to decrease. The formula intersects with the calculated correlation coefficient we found by comparing the historical prices from Henry Hub and the NBP.

Formula (8.5) takes distance as input and gives the correlation coefficient as output. The correlation for port 1, the reference port, is defined as 1. d_p is the distance from the reference port to port p. C_p is the correlation between the reference port and the port for which we are generating prices. The function is shown in Figure 8.3. The red lines indicate where the correlation between Henry Hub and the NBP fits into the function.

$$C_p = 0.797 \cdot e^{\frac{d_p}{1000}}, \forall p \in 2 \cdots P \tag{8.5}$$

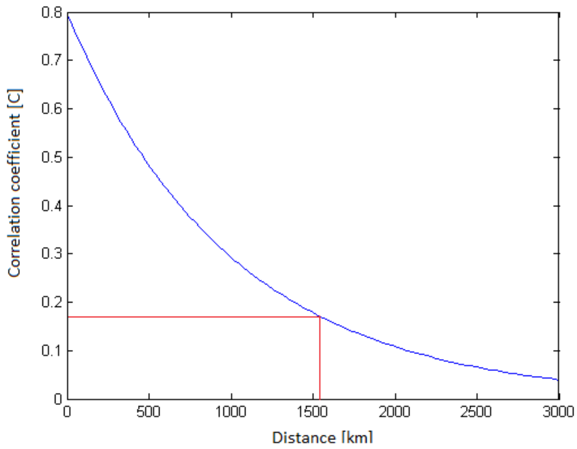

Figure 8.3: Correlation coefficient as a function of distance

Figure 8.4 show how prices vary with correlation. All price paths in the examples have the same starting point. The correlation to the reference price is indicated underneath each figure. The price path from the reference port is included in every figure for comparison. Figure 8.4(a) and 8.4(f) are perfectly negatively correlated, and the price processes is exactly opposite of each other. Figure 8.4(d) is not correlated with any of the others. We have assigned that the correlation coefficients relative to the reference port. This means that the correlation between two of the other ports are not decided by how close they are to each other, but how close each of them are to the reference port.

8.3 Fixed parameters

Some of the parameters are fixed throughout all our tests. These are presented in this section. When testing the model our main focus has been to change parameters that diversifies the performance of our different models. This has decided which parameters we have chosen to fix.

8.3.1 Parameters used for price forecasting

The start prices for each port, lower limits and correlation coefficients used in the Ornstein-Uhlenbeck process are shown in Table 8.3. The start prices are chosen in a manner that should give an incentive to trade in multiple ports. For example, the margin between Hammerfest and Milford Haven is small since those are the two closest port. The minimum price limits are set to \$1.80 for all ports. This is lower than any price registered in the Henry Hub the last five years. The buy ports and the sell ports has the

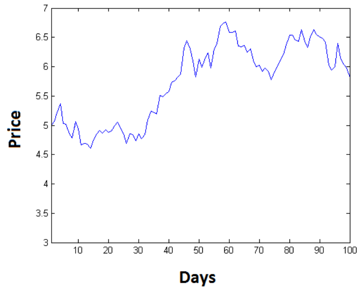

(a) Correlation = 1

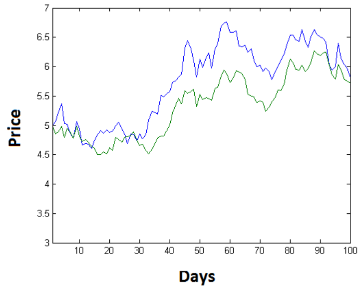

(b) Correlation = 0.5

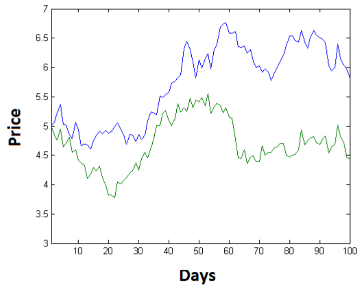

(c) Correlation = 0.2

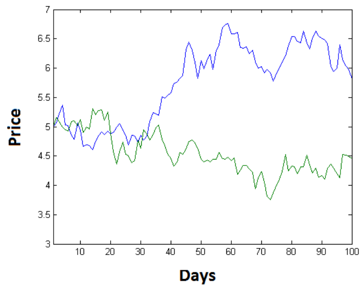

(d) Correlation = 0

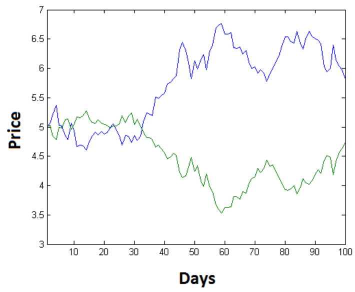

(e) Correlation = -0.7

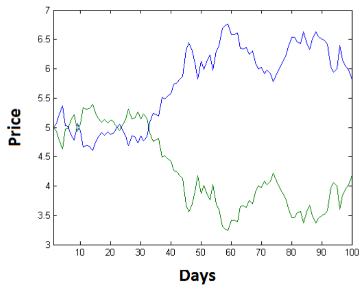

(f) Correlation = -1

Figure 8.4: Price processes with starting price 5 and varying correlations

same lower limit. σ and λ are assumed to be equal for all ports. Their values are 0.163 and 0.0034, respectively. These are calculated from the prices in Henry Hub over the last five years. The long-term mean of a port is set to the same value as the start price. We chose the start prices in order to create many trading options for the vessel, and finds it natural that the prices drift towards these prices to keep the options open.

Table 8.3: Values of parameters used in the Ornstein-Uhlenbeck process

Port	Start price/MMBtu	L_{min}	Correlation
Point Fortin	\$5.3	\$1.8	0.260
Bonny	\$5.2	\$1.8	0.240
Hammerfest	\$5.5	\$1.8	0.474
Lake Charles	\$5.0	\$1.8	0.174
St Johns	\$7.8	\$1.8	0.420
Marmara	\$8.0	\$1.8	0.331
Milford Haven	\$8.1	\$1.8	1.000

8.3.2 Number of days in total problem period

The number of days in the problem period is set to 200. This is based on a trade-off between run time and generating interesting results. With 200 days the model has enough time to see significant differences in the solutions found by the models, without having a too high run times. In a real life situation this concept of problem period is not interesting, as it is the day-to-day decisions that matter.

8.3.3 Price development

Actual price information for various trade ports is highly limited. We have therefore decided to simulate the prices used in our tests. This is done with the Ornstein-Uhlenbeck process, as described above. The price development used is seen in Figure 8.5. For the buy ports, the price in Bonny increases the most, almost reaching a price of \$7/MMBtu. Marmara is the sell port which has the largest decrease in price, down to almost \$6/MMBtu. The price development is decided before the model starts running, but the models does not have access to future prices. They only have price information for their current day, and have to base the actions taken on price forecasts.

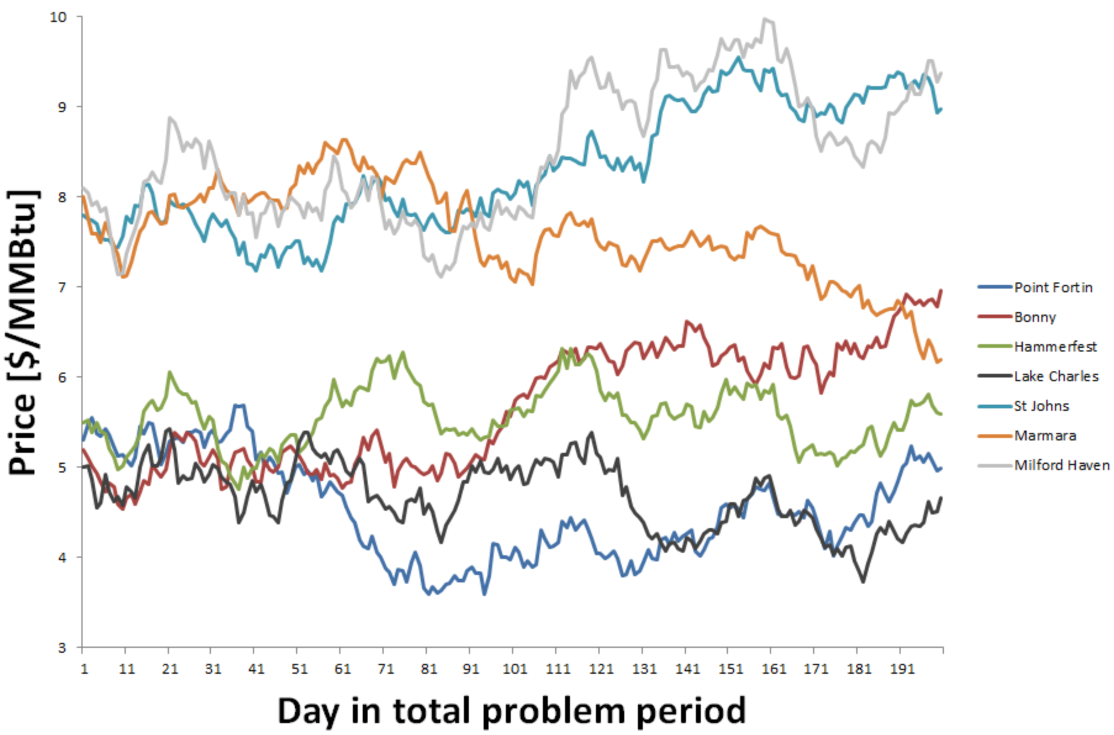

Figure 8.5: Price development for all ports used in the test instances

8.3.4 Size of vessel

The capacity of the vessel is 142,500 m^3, which equals approximately 3,000,000 MMBtu. This is a normal size for an LNG vessel. We have chosen to not vary the size of the vessel. This is because our main focus is comparing the two solution approaches, and we do not see how this could contribute to separate the models from each other.

8.3.5 Basis cost and travelling costs

Each day the vessel incurs a basis cost of $100,000, regardless of moving or not. This cost reflects the cost of hiring a LNG vessel, or the alternative cost of owning one. Making a straight move costs an additional $100,000, while a diagonal move costs $140,000. The diagonal move is more expensive because of the longer distance travelled. The numbers we use are comparable to real numbers, as described in Section 2.7. We have chosen to not test with different values for basis and travel costs, as we do not see how this could contribute to separate the models from each other.

8.3.6 Starting conditions

The vessel starts empty in Milford Haven, which is a sell port. It starts empty because we find it more interesting for the vessel to make complete trades, buying the LNG before selling it. Milford Haven is chosen because it is located in between many ports, and thus is a port that has many interesting trade options to start off with. We have chosen to not test with different start positions. This is because we automatically get to test the model for different positions as the vessel moves throughout the problem period.

8.3.7 Discounting factor

The discounting factor is set to 1% per 30 days. We have considered varying this factor, but have concluded that a change does not lead to a big impact on the model.

The thought behind increasing the discounting factor is to make the vessel get more value from the early trades. It would thus choose routes the routes that performs well early in the trade horizon over routes that perform well late in the trade horizon. In that way the the focus is on the short-term trading, where the price forecasts are more certain. The long-term trading, with more uncertain price forecast, would still be taken into account, but to a smaller extent. This theory is however incorrect, as there is no such thing as routes performing poorly early in the trade horizon when using the dynamic program approach. These routes would be dominated away by better routes before reaching the late trades. There will always be some optimal route that performs well both in the beginning and towards the end.

There is however one factor that might give some impact on a change in the discounting factor. This is the timing of the trades. The best routes will normally perform the maximum number of trades, but the timing of these trades can vary. A high discount factor would push the vessel to make trades early. The discount factor would however have to be high enough to make up for the potentially positive price development from one day to the other.

We have done some small tests with different discount factors, and these tests indicate that the results does not change with a varying discount factor. The value of the discount factor is therefore kept stable throughout our tests.

8.3.8 Nonanticipativity constraints in the MIP

We use nonanticipativity constraints for the first day in the MIP. For each day we are only interested in finding the first action of the vessel, and there is no need to model nonanticipativity constraints for later days. All scenarios that are used in the stochastic model are included in the nonanticipativity subset. This ensures that the vessels in every scenario commit the same action the first day.

8.3.9 Maximum run time of MIP

We have set the maximum run time of the MIP to be 3,600 seconds. The best solution so far is returned if the solver has not a found a solution within the bounds by then. The limit is set to 3,600 seconds because we find it unreasonable to spend more time on an operational level decision.

8.4 Variable parameters

Below we present the parameters that are changed throughout our tests. These are the parameters we find to be most interesting when it comes to impact on model performance. The reason for choosing each of the parameters is described. The parameters values used in the base case are presented in Section 8.6.

8.4.1 Number of scenarios

By varying the number of price scenarios, we try to show that a larger number of scenarios leads to a more correct impression of future prices. Our theory is that this leads to improved vessel actions, which again gives higher profit. We have decided to run tests with five different number of scenarios: 1, 5, 20, 50 and 100. The reason for the increasing intervals is that we believe the increase in profit is diminishing as the number of scenarios increase.

The run time should get higher as we add more scenarios. By doing these tests we try to find a good trade-off between profit and run time.

8.4.2 End of horizon values

We vary the end of horizon values (EOH-values) to examine how much the they change the incentive for the vessel to move. As discussed in Section 5.2.5, the profit loss of not moving gets larger with EOH-values. Our theory is that the vessel moves more frequently with EOH-values, in order to prevent this loss. We have tested the model with and without EOH-values. The EOH-values include $400,000 in value for each flexible day in the end of the rolling horizon, plus a 120 % return on the buy value, if buying as the last trade.

We believe the run time time of the DP will be slightly lower with EOH-values. This is because EOH-values leads to more domination towards the end of the trade horizon. A route that finishes some days before the end of the horizon is likely to dominate some later labels due to its additional EOH-value.

8.4.3 Rolling horizon length

Varying the length of the rolling horizon used in the DP potentially has an impact on the profit. Increasing the length of the rolling horizon makes the dynamic program better approximate the original problem, as the vessel is able to consider trades for a longer period. Run times will however increase. The high uncertainty of future prices also limits how many days ahead it is benefial to consider. By running these test we hope to find the optimal length of the rolling horizon. We have chosen rolling horizon lengths from 20 to 100 days with 10 day intervals.

The run time of the DP should increase with longer trade horizons, as there will be more trades and thus more labels to extend. The run time of the MIP and heuristic should not be effected by a longer rolling horizon. From their point of view the only thing that is changing is the values of the ports in each scenario.

To understand the interaction between the EOH-values and rolling horizon, we have also tested changes in these two parameters together. Both parameters affect the vessel's incentives to move, and we believe that varying them together would make it easier to understand this connection.

Decreasing rolling horizon length vs fixed rolling horizon length

We test the impact of using decreasing rolling horizon length. Decreasing rolling horizon length means that the trade horizon considered by the DP decreases as long as the vessel has not made a trade. This leads to a shorter and shorter trade horizon until the vessel actually performs a trade. After the trade is conducted, the rolling horizon bounces back to its original value. When the trade horizon decreases for each day, the vessel has a greater incentive to move. This is because the vessel has to keep on moving in order to be able to make the trades calculated by the DP. With a fixed-length trade horizon the same opportunities will be there the next day. This makes it less risky to stay put for the vessel, which again leads to less trading and a lower profit.

The decreasing rolling horizon should make the model run slightly faster, as the dynamic program uses more time the longer the horizon is.

8.4.4 Gap in MIP

The gap in the MIP decides the maximum gap between the upper and lower bound of a solution. As soon as this gap is reached, the current solution is accepted. Xpress-Mosel solves an integer problem by closing in from two sides. The model relaxes the problem, finding solutions which become upper bounds. Integer solutions found can be set as lower bounds. The model constantly seeks to decrease the gap between these bounds. A smaller gap should give a better result than a big gap.

The run time of the model should be higher as the gap increases. It requires more work by the solver to close in to a small gap. We have chosen to do tests with gaps of 0.5%, 1%, 2% and 5%.

8.4.5 Binary days in MIP

In this section we discuss some of the constraints from the MIP, which is described in Section 5.3. These constraints are shown in Figure 8.6.

$$m_{ts}^{D} - x_{g(t+1)s} - \sum_{\bar{g} \in \mathcal{G}^{\mathcal{D}}(g)} x_{\bar{g}ts} \geq -1, \qquad t \in 1 \ldots \mathcal{T}-1, s \in \mathcal{S} \tag{5.3}$$

$$m_{ts}^{S} - x_{g(t+1)s} - \sum_{\bar{g} \in \mathcal{G}^{\mathcal{S}}(g)} x_{\bar{g}ts} \geq -1, \qquad t \in 1 \ldots \mathcal{T}-1, s \in \mathcal{S} \tag{5.4}$$

$$r_{gts} - x_{gts} \leq 0, \qquad g \in \mathcal{G}^{\mathcal{P}}, t \in \mathcal{T}, s \in \mathcal{S} \tag{5.6}$$

$$x_{gts} \in \{0,1\}, \qquad g \in \mathcal{G}, t \in \mathcal{T}, s \in \mathcal{S} \tag{5.11}$$

$$m_{ts}^{S}, m_{ts}^{D} \in \{0,1\}, \qquad t \in \mathcal{T}, s \in \mathcal{S} \tag{5.12}$$

Figure 8.6: Constraints from the MIP that are affected by the number of binary days

The MIP can be relaxed by only keeping binary constraints on the position variables the first days. This is possible since we only are interested in the vessel's first day action when we run the MIP together with the DP. When the binary constraints are removed, the vessel is given the opportunity to divide itself into pieces. This is in conflict with Constraints (5.6), which ensures that the vessel has to trade full loads. The vessel can solve this by travelling in parts until it reaches the port, and merge together when arriving. A removal of binary constraints is a relaxation, because it enables the vessel to travel as far as it wants without any cost. This is because Constraints (5.3)-(5.4) are not restricting with fractional variables spread over several grid points. Hence the move variables can be set to 0. The model chooses to set the variables to 0, since the move variables contributes to cost in the objective function. This leads to ports far away becoming more valuable.

The use of binary days can be expressed mathematically, where D^B is used as the number of binary days. Instead of Constraints (5.11) and (5.12) we can use the following formulations:

$$x_{gts} \in \{0,1\}, \qquad g \in \mathcal{G}, t \leq D^B, s \in \mathcal{S}$$
$$x_{gts} \geq 0, \qquad g \in \mathcal{G}, t > D^B, s \in \mathcal{S}$$

$$m_{ts}^{S}, m_{ts}^{D} \in \{0,1\}, \qquad t \leq D^B, s \in \mathcal{S}$$
$$m_{ts}^{S}, m_{ts}^{D} \geq 0, \qquad g \in \mathcal{G}, t > D^B, s \in \mathcal{S}$$

We believe that a low number of integer days gives the vessel an incentive to wait. If the vessel can travel for free in a few days, it could be beneficial for it to wait until then to do the move. Thus, our theory is that the vessel performs better the more integer days are included in the model.

The complexity of the MIP increases with the number of integer days. We therefore expect an increase in run time as well. Tests are run for 2, 3, 4, 5 and 6 binary days to try to find a trade-off between profit and run time.

8.4.6 Ports included in heuristic move calculation

As discussed in Section 5.4.2, there are two alternatives for which ports to include when calculating the value of potential moves. The first alternative is to only include ports that are fewer travel days away after the move. The second alternative inludes all ports that are the same or fewer travel days away after the move. Both alternatives are tested.

The run time for these tests should be the same. This is because the DP, which is the main contributor to run time, solves the same problem for both tests.

8.5 Stability testing

We have included stability testing to see how well the model performs in- and out-of-sample. The basis for performing these tests are stochastic runs of the models. For the in-sample tests we research what decisions would be made by each of the scenarios from the stochastic problem, if they were run independently as deterministic problems. We compare these decisions to the decisions made by the stochastic solution. For the out-of-sample test we generate new scenarios, and research what decisions each these would have made if faced with the same problems as the stochastic model.

The results from the stability tests give us an indication of how well the models are able to exploit the additional information from adding new scenarios. We have divided our stability testing into different situations for the vessel, based on where it is situated relative to the ports, to see if any of the situations seems easier to solve than others. This can later be used to make more effective models, e.g. by using more scenarios for some vessel situations than others.

Our problem differs from others by solving many problems inside a problem period, with no easy way of measuring how good a single move is for the long-term problem. We have not found any literature that tests stability in a similar way, and have therefore created our own method. This method is described below.

8.5.1 Determining level of stability

When deciding the stability of a test, we look at the actions made by the vessel each day. Our model solves a problem over a long time period, and for each day a decision must be made. We try to compare these daily decisions. Decisions made when solving with individual scenarios are compared to the decision made by the stochastic model. A value is given to the deterministic action, depending on how close it is to the action chosen by the stochastic model. There is a big difference between moving in an opposite direction compared to for example staying put. We have assigned a value to all moves, based on how close they are to the stochastic solution. The stability values for a moving vessel are shown in Figure 8.7. The stability values that are used when in a port are shown in Table 8.4.

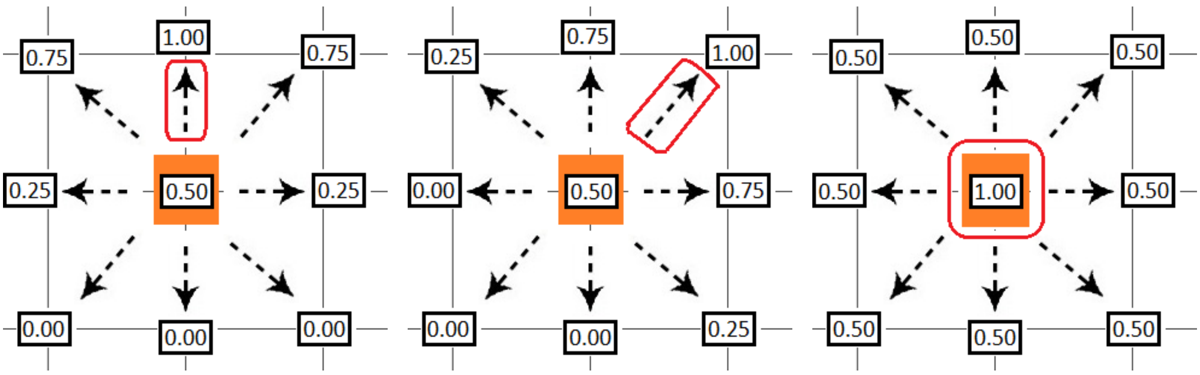

(a) Straight move (b) Diagonal move (c) Staying put

Figure 8.7: Values used for moves in stability testing

We use three different vessel states: (1) the vessel moves, (2) the vessel stays put and (3) the vessel is trading in a port.

The values for state 1 is shown in Figure 8.7(a) and 8.7(b). The value of almost choosing the same move is given as 0.75, while choosing to go in the opposite direction gives a value of 0. If the scenario chooses the same action as the vessel, it gets a value of 1. The reason for these values is that we want to benefit the single scenario for doing almost the correct move. A higher stability value means a higher degree of correlation between the single scenario solution and the stochastic solution.

The values for state 2 is shown in Figure 8.7(c). When the vessel stays put, the single scenario gets value 1 for also staying put. If the single scenario solution chooses to move, it gets value 0.5. The reason for a move giving some value is that moving does not have to be a big mistake. The vessel in the stochastic solution might have been planning to go in that exact same direction the next day. The value of moving cannot be set too high, in case the vessel in the single scenario case is travelling the opposite way of the vessel in the stochastic model.

The values for state 3 are shown in Table 8.4. This state differs from the other states. While states 1 and 2 have nine possible options, state 3 only has two options. The chance of making the right decision with two options is much higher, and we penalize wrong decisions harder. If the stochastic model sells, and the scenario sells, it gets a value of 1. If it does not sell, it gets a value of 0. The same principle goes for not selling. Since state 3 is calculated different from the other states, it is not directly comparable with the stability values calculated for the other states. However, it can be useful to test how the vessels stability in the port differs for a different number of scenarios.

Table 8.4: Values used for trading in stability testing

Stochastic action	Scenario action	
	Sell	Wait
Sell	1	0
Wait	0	1

The absolute values we calculate in the stability testing are of no worth, but the relative difference between tests can decide how the tests perform compared to each other.

Numerical example

We present a numerical example from an in-sample test with five scenarios. The example reflects one of the days in the total time period. In the example, the vessel starts it day in grid point (5,5) and the stochastic solution is to move to point 5,6. Table 8.5 shows the moves chosen by each of the scenarios deterministically. Figure 8.8 shows an overview of the grid points in the example, where indices for every grid point is given. The choice of the stochastic solution is shown with an arrow. The value given for each move is given corresponds to the value in Figure 8.7. The total sample value for a day is calculated as the arithmetic average of the sample values in each scenario.

8.5.2 Scenarios used in the stability testing

An overview over the scenarios used in the in-sample and out-of-sample tests are shown in Figure 8.9. For in-sample stability we consider the same scenarios that were used to

Table 8.5: Numerical example of stability test

Scenario	Move to	Value
1	5,6	1.00
2	5,5	0.50
3	4,6	0.75
4	4,5	0.25
5	5,6	1.00
Total sample value		0.70

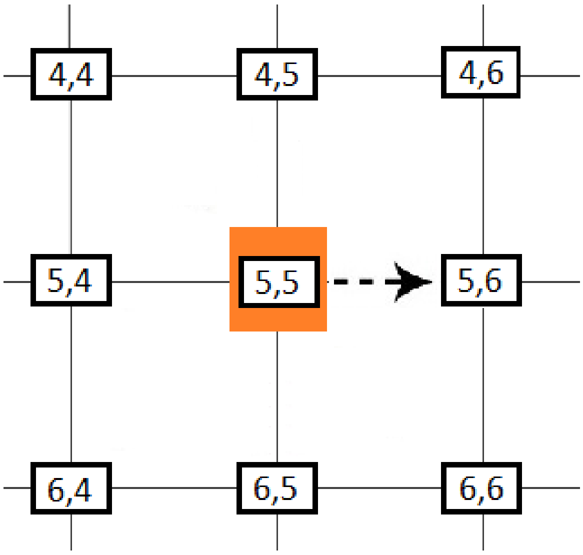

Figure 8.8: Map used in numerical example of stability test

create the stochastic solution. Each day of the problem we find the action that would be performed by the model if run with each of these scenarios, independently. The in-sample stability test gives us the correlation between these deterministic actions and the actions performed in the stochastic solution. This is helpful when selecting how many scenarios to use for solving the stochastic problem. A high degree of correlation indicates that we get a similar result in our model, regardless of which of the in-sample scenarios that is used.

In the out-of-sample testing we test independent scenarios against the solution found by the stochastic model. For each day we generate new scenarios and find the action performed for each of these scenarios. The out-of-sample stability tests lets us research how many scenarios are needed to get a stable solution. The gain from adding even more scenarios is small when we already have produced a solution that fits well with random scenarios. For the out-of-sample tests we test with 50 new scenarios.

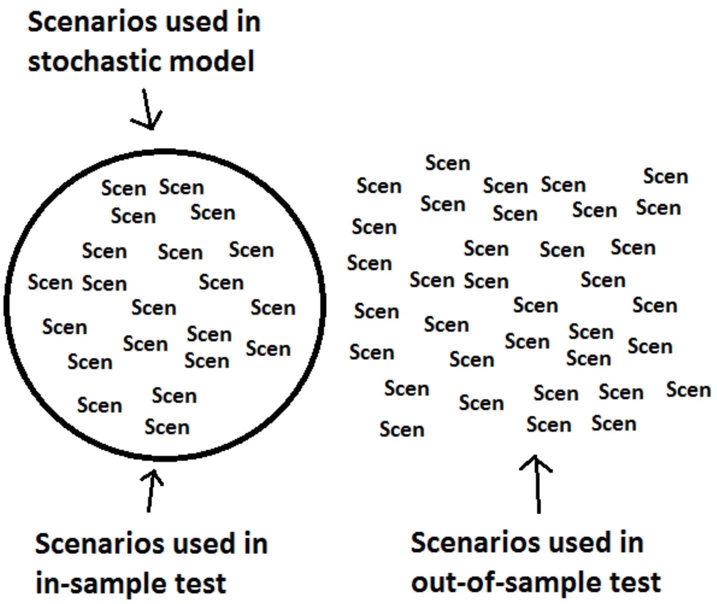

Figure 8.9: Overview over scenarios used for in-sample tests and out-of-sample tests

8.5.3 Stability situations for the vessel

We have divided the stability tests into four vessel situations: (1) on the way from a port, (2) on the way to a port, (3) in a port and (4) in open sea. A vessel is defined on its way from or to a port if it is less than two days of sailing from a port. If the vessel can trade in the port, it is defined on its way to the port. If not, it is on its way from the port. While staying in a port, a vessel is defined as in the port until it trades. After a trade it is defined as travelling from a port, even if it chooses to keep on staying in the port. All situations at least three days from a port is defined as in open sea. We have segmented these results into groups for different number of scenarios. The in port values are not directly comparable to the rest of the values. This is because there are only two options and the stability values are assigned differently when waiting for a trade in a port than for other situations. For the in-sample and out-of-sample tests we have found values of stability for all these four situations.

8.6 Base case

The base case is a test instance with the standard value for all parameters. For the fixed parameters these are the values discussed in the Section 8.3. For the variable parameters we have chosen values we believe to give good trade-off between run time and results. The parameter values used in the base case are listed in Tables 8.6 and 8.7.

Table 8.6: Values for fixed parameters of the base case

Parameter	Value
Vessel size	142,500 m^3
Basis cost, per day	$100,000
Travel straight cost	$100,000
Travel diagonal cost	$140,000
Number of days in total problem period	200
Discounting factor, per 30 days	1%
Days with NACs	1
Maximum run time for MIP	3,600 seconds

Table 8.7: Values for variable parameters of the base case

Parameter	Value
Number of scenarios	20
EOH value, per day	$400,000
EOH value, return on buy	*1.2
Rolling horizon length	60
Decreasing rolling horizon	Yes
Ports included in heuristic move calculation	Fewer travel days
Gap in MIP	1%
Integer days in MIP	3

The base case is used as a starting point when testing with different parameter values. All variable parameters that are not currently being tested are kept at their base case value.

Chapter 9

Results and Discussion

In this chapter we present and discuss the results from our tests. All tests are run as complete instances of the full problem period, solved with a combination of the DP and the MIP or heuristic. In the last two chapters these solution approaches are denoted as MIP and heuristic, even though the DP is included in both solution approaches. The price forecasts for each day are used to decide actions for the vessel, while the real price development is used when calculating the profit. The profit of each test is found after the test has run, by summing the trade values of all trades and subtracting the travel costs. The trade values are negative for buying and positive for selling. The prices used comes from the price development shown in Section 8.3.3.

We start the chapter by considering the dynamic program, looking at the effect of EOH-values and rolling horizon length. Following this the results from stochastic runs of the MIP and heuristic are presented. We discuss the effect on profit and run time when varying parameter values. The difference in performance between the MIP and the heuristic is then discussed, before doing the same tests for the deterministic runs of the MIP and heuristic. We conclude the chapter by comparing the stochastic and deterministic models. All test instances are run with the base case parameters presented in Section 8.6, unless otherwise stated. In this chapter we mainly use figures to present the test results. Tables with values for all diagrams is found in Appendix A.

9.1 Test of Dynamic Program

In this section we examine how the number of days in the rolling horizon and the use of EOH-values impact the performance of the DP. These concepts are described in Section 5.2. The parameters are tested in combination to identify any connection that might exist between the two. The tests are run using the stochastic version of the heuristic. We start the section by describing the results from tests without EOH-values. We then proceed to discuss the tests with EOH-values, before comparing the two tests.

Figure 9.1 shows the profit and run time for decreasing length of rolling horizon, without EOH-values.

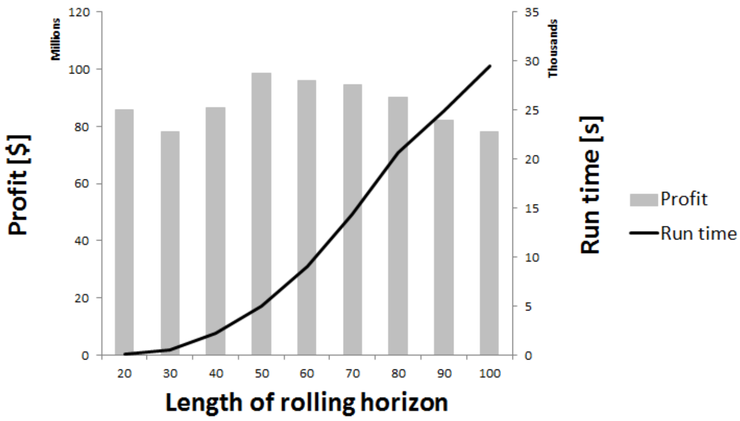

Figure 9.1: Profit and run time without EOH-values when varying rolling horizon length in the DP

The test without EOH-values has a high profit for medium trade horizon lengths. For short trade horizons the vessel gets a low profit because it does not consider the long-term effects when deciding the trades. Long trade horizons also seems to lead to low profits. This is surprising, since a long trade horizon could give the vessel a better perspective over the long-term. We believe that the increase in possible routes for a longer trade horizon is the reason for the low profits. With a high number of routes it is likely that many routes that have almost the same value. This reduces the loss of waiting an extra day, since it is probable that a new route gives about the same profit the next day. This increasing the chance of waiting. By examining the routes of the vessels we can confirm that the vessel hesitates more as the trade horizon increases. A medium trade horizon seems to give a good trade-off between incentives to trade and perspective over the long-term horizon.

Longer trade horizons give higher run time due to an increase in the number of labels generated by the DP.

Figure 9.2 shows profit and run time with EOH-values included in the model. Running the model with EOH-values seems to perform best with a short trade horizon. The EOH-values intensifies trading, giving value for finishing trades early and making buying a valid last action in the DP. This compensates for the fact that few trades are considered. The profit decreases as trade horizon length increases. We believe the reason is the same as for the test without EOH-values, the increased number of options makes the vessel hesitate.

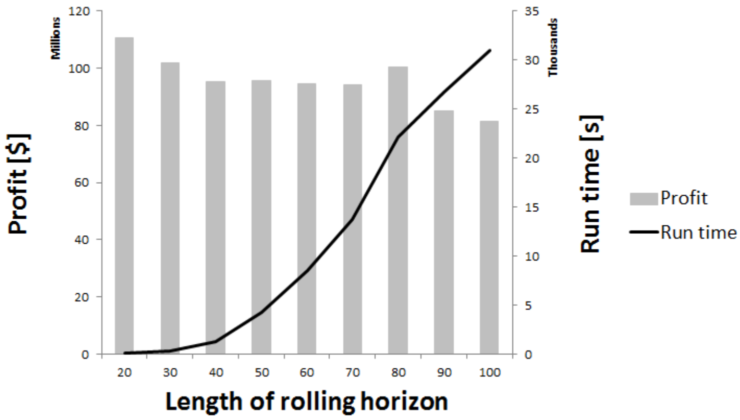

Figure 9.2: Profit and run time with EOH-values when varying rolling horizon length in the DP

There is a growth in run time as the length of the rolling horizon increases. This is because the number of labels created by the DP increases as the trade horizon becomes longer.

Figure 9.3 shows a comparison between the model with EOH-values and the one without EOH-values.

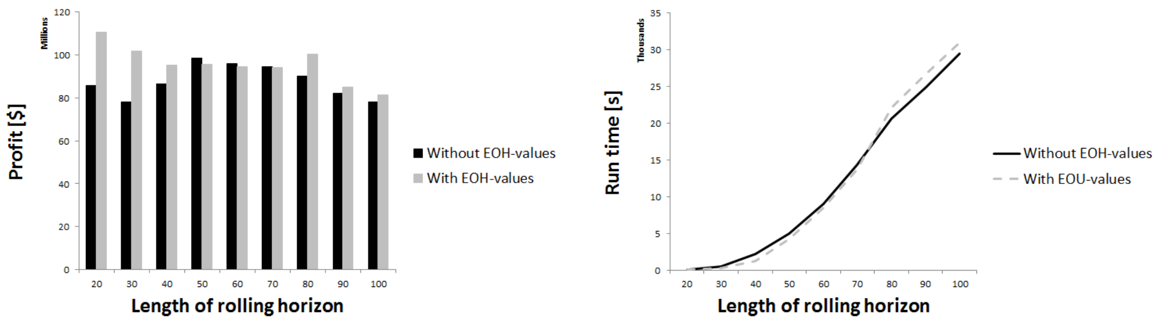

(a) Profit (b) Run time

Figure 9.3: Profit and run time when varying rolling horizon length and use of EOH-values in the DP

The difference between the tests decreases as the number of days in the trade horizon increases. We believe this is because we get more possible route combinations when considering longer trade horizons. It is likely that the vessel is able to use most of its days for trading in some of these route combinations. This reduces the effect of the EOH-values.

The tests with EOH-values are running faster for short rolling horizon lengths. This is because the EOH-values leads to more routes being dominated away towards the end of the trade horizon. This effect diminishes with increase in the trade horizon length, as the relative value of the EOH-values become smaller compared to the total trade profit.

We have done tests showing that the EOH-values have the same influence on the MIP as on the heuristic. The trends are the same as for the heuristic for both tests, and the profit of the tests converges towards each other as the length of the rolling horizon increases.

Conclusion on parameter values

EOH-values leads to increased profit without any significant impact on run time, and should thus be used. Given that EOH-values are used, we get the best results for short trade horizons. Short trade horizons also solve faster. Using EOH-values and a trade horizon of 20 seems to be the best choice of parameters for the dynamic program.

9.2 Stochastic models

In this section we look at the results from the stochastic tests of our models. We first look at the profit, run time and stability of the MIP and heuristic, respectively. Then we conclude the section with a comparison of the two approaches.

9.2.1 Test of MIP

The MIP is tested by varying the number of scenarios, the number of binary day and the size of the gap. These parameters are described in Section 8.4. We also consider the effect of decreasing the size of the rolling horizon until a trade has been made. This is discussed in Section 8.4.3. The rest of the parameters are kept at their base case values, as described in Section 8.6. The stability of the results are tested through in- and out-of-sample stability tests. We conclude by discussing the optimal parameter values for the MIP.

Number of scenarios

Figure 9.4 show how different number of scenarios impact profit and run time.

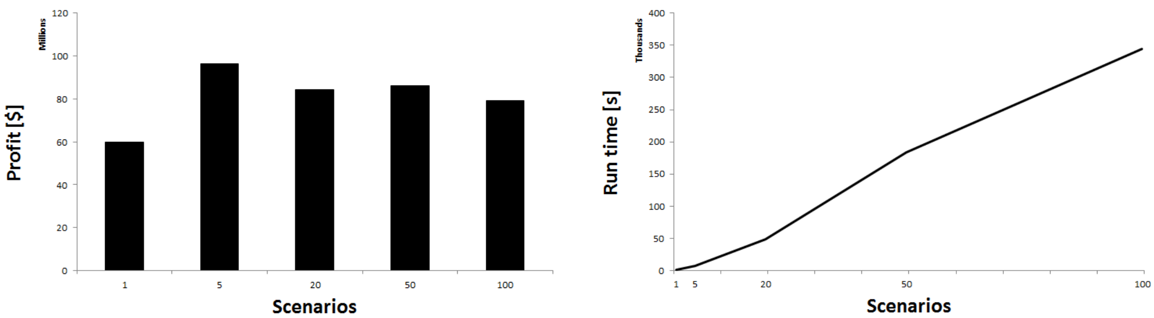

(a) Profit (b) Run time

Figure 9.4: Profit and run time with decreasing horizon when varying the number of scenarios in the stochastic version of the MIP

Using only one scenario gives a particularly poor result. For one scenario the vessel hesitates a lot. The reason for this is that there is no uncertainty from the vessels point of view. Hence, there is no need for the vessel to travel towards a port, unless the vessel has to prevent a decreasing port value for the next day. This might lead the vessel to wait with the trade until the length of the time horizon is exactly long enough to reach the trade. Overall this induce a lot of waiting.

For multiple scenarios the profit seems to decrease with an increase in the number of scenarios. This is not as expected. We thought that a higher number of scenarios would enhance the solutions. We believe that the reason for these results is that the vessel has more options to consider. More options apparently leads the vessel to not sail directly to a port, but rather wait or roam around. The value of better price forecasts is outweighed by the hesitation it leads to. By examining the routes from the results, we see that the number of waiting days increase as the number of scenarios increases. This leads to a smaller number of trades and lower profit.

The run time increases steadily as the number of scenarios increase. This is natural, as the problem becomes more complex with an increasing number of scenarios.

Type of rolling horizon

Figure 9.5 show a comparison of profit and run time between a decreasing rolling horizon and a fixed rolling horizon.

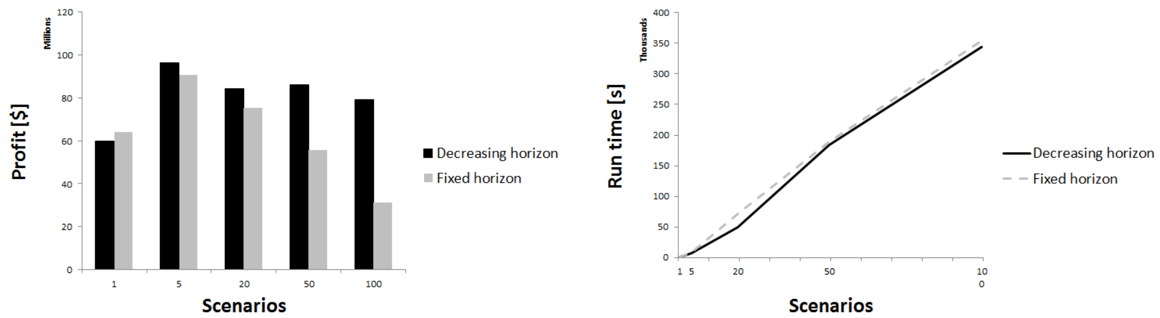

(a) Profit (b) Run time

Figure 9.5: Profit and run time when varying the number of scenarios and type of rolling horizon in the stochastic version of the MIP

It seems clear that using decreasing rolling horizon length leads to higher profit. This is because the decreasing rolling horizon leads to less hesitation by the vessel. It has to move towards ports in order to have the time to fulfill all trades inside the trade horizon, as calculated by the DP. With fixed horizon length the vessel still has all the same options the next day if staying put. This decreases the vessels incentive of moving. By examining the routes chosen by the vessel, we see the same trends as for decreasing horizons. For the fixed horizon the trend of hesitating more as the number of scenarios increase is stronger than for the decreasing horizon.

The run time with decreasing rolling horizon length is a bit shorter than with fixed horizon length. This is because the trade horizons are shorter on average, hence decreasing the run time of the DP.

Binary days

Figure 9.6 shows the impact on profit and run time when varying the number of days with binary position variables. Binary position variables are discussed in Section 8.4.5.

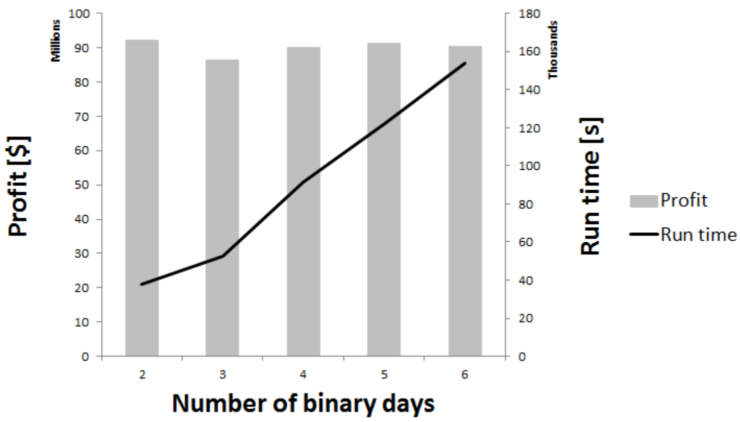

Figure 9.6: Profit and run time when varying the number of binary days in the stochastic version of the MIP

We cannot see a clear trend for the impact on profit. This is surprising, as reducing the number of binary days is a relaxation. When examining the routes for the tests, it seems like the vessel chooses the same route for all tests. The main difference in the routes, is that the tests with more binary days tend to trade earlier in ports than the models without. This should lead to more trades and a higher profit for the tests with many binary days. But it turns out that for some trades, the tests with many binary days waits in a port for the price to increase. These situations give the tests with few binary the chance to catch up. This makes the total number of trades becomes the same overall. The small differences observed in profit, is due to the timing of the trades. We believe that the hesitation of the vessels with few binary days is explained by a wait day giving the vessel an opportunity to travel fir free a later day.

A significant impact on run time can be observed. The test with two binary days runs significantly faster than the test with six binary days, which is natural due to the increased complexity of the problem.

Gap

Figure 9.7 shows the run time and solution value for using different gaps in the MIP model.

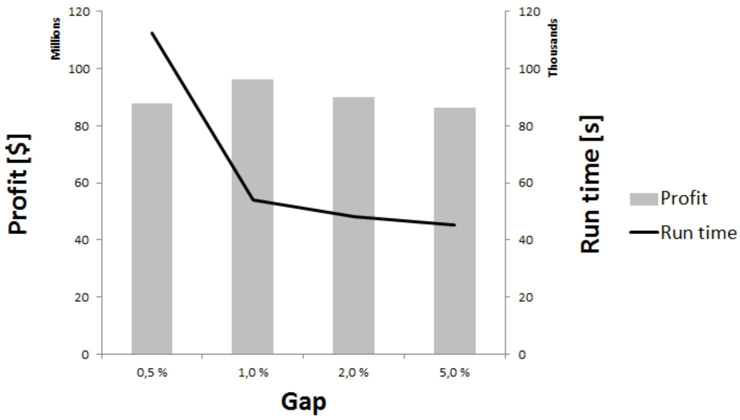

Figure 9.7: Profit and run time when varying gap value in the stochastic version of the MIP

Lower gaps seem to lead to higher profit, with the exception of 0.5%. Studying the route driven shows that this exception is caused by the model waiting longer in ports in order to get optimal prices. This waiting leads to fewer trades completed, and thus lower profit. The 1% test does not make the optimal timing of trade decision, and actually makes mistakes by trading too early. In the end the 1% test is compensated by reaching more trades in the total time period. The 2% and 5% tests makes bad choices both in and between ports.

The run time increases as the gap decreases. The increase is largest between the 0.5% test and 1% test. This can be explained by how the bound converges. The upper bound of the gap decreases slow but steadily as the model is solved. The largest decrease in gap value happens when a new integer solution is found. We have examined results from a sample of tests, which indicate that the 2% and 5% gap is reached for the first or second integer solution. The 1% gap is reached after about five integer solutions, while the 0.5% gap is reached after about ten integer solutions. The decrease in gap by finding a new integer solution diminishes as the number of solutions found increases, which explains the big difference in run time from the 0.5% test to the 1% test.

Stability

The stability tests used are explained in Chapter 8.5. We first discuss overall in- and out-of-sample stability for the stochastic MIP model. Then we discuss the in-sample tests and out-of-sample tests separately for different vessel situations. These situations are described in Section 8.5.3.

Overall stability

Figure 9.8 shows an overview of the results from the stability tests.

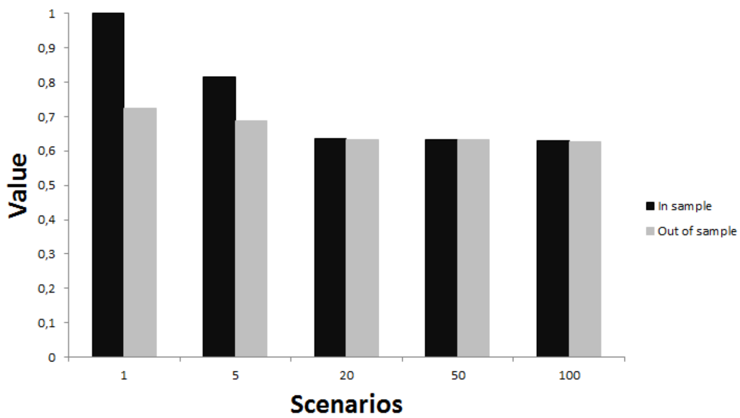

Figure 9.8: Overall in- and out-of-sample values when varying the number of scenarios in the stochastic version of the MIP

For the in-sample test, we neglect the value for one scenario in the discussion since its value is 1 by default. For the rest of the scenarios, the in-sample stability decreases as the number of scenarios gets higher than five. This is surprising, as more scenarios should usually give a better in-sample value. A higher number of scenarios gives more information to the model, and a better opportunity to create a good solution. We believe that the bad results is due to the hesitations of the vessel as the number of scenarios increase. This waiting both weakens the results, and turns out to perform different for the scenarios in the sample.

The out-of-sample values seem to decrease steadily. The explanation is the same as for the in-sample values. It is natural that the out-of-sample values converge with the in-sample values as the number of scenarios increase. With a high number of scenarios, the chance of getting a representative sample of scenarios is big. Hence, the samples of scenarios in the in-sample and out-of-sample tests become more similar to each other as the number of scenarios increase.

In-sample stability

Figure 9.9 shows in-sample stability for different situations.

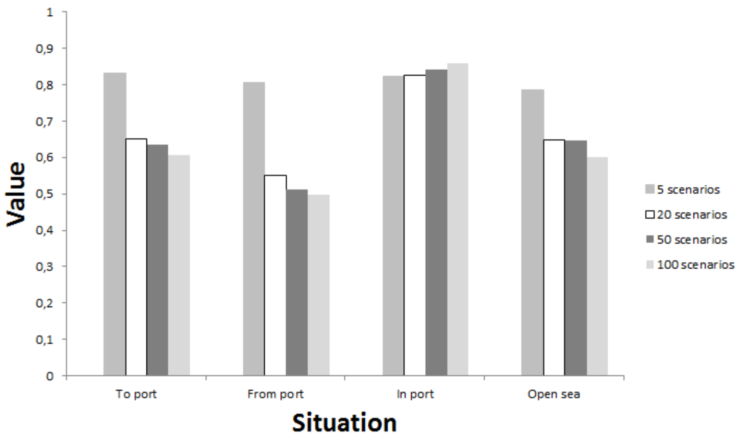

Figure 9.9: In-sample values for different situations when varying the number of scenarios in the stochastic version of the MIP

There seem to be some clear trends for the in-sample stability. Five scenarios perform significantly better than the others in three out of four situations. The trend seems to continue for higher numbers of scenarios as well. It looks to be an advantage with few scenarios in all situations except "in port". We believe this is because the stochastic solution with five scenarios does not hesitate as much as solutions with a higher number of scenarios, as described in the above discussion. "In port" seems to be the only place it is an advantage to have many scenarios. We think this is because there are only two options. With two options the scenarios manages to find the best solution instead of ending up not deciding because of having too many options.

The decrease in stability value "from port" seems to be larger than the other decreases. This might be because the decision "from port" is tougher than the others. Our testing shows that the MIP hesitates more when decisions are tough, thus creating a poor solution.

Out-of-sample stability

Figure 9.10 shows out-of-sample stability for different situations.

The trends in for the out-of-sample stability are similar to the in-sample stability. The stability value decreases for more scenarios in the situations "to port", "from port" and "in open sea", and increases for more scenarios "in port". The explanation for the behaviour is the same as for in-sample stability.

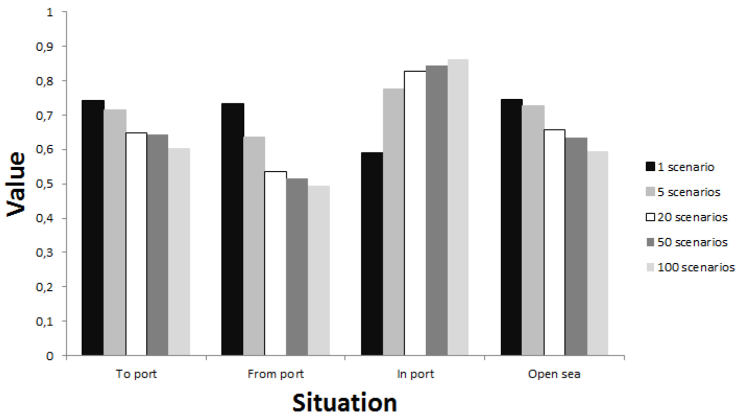

Figure 9.10: Out-of-sample values for different situations when varying the number of scenarios in the stochastic version of the MIP

Conclusion on parameter values

Decreasing length of the rolling horizon gives better results than fixed length. The results from running with different scenario numbers indicate that increasing the number of scenarios beyond 5 decreases the profit. Using 2 binary days seems like a good choice, given the high profit and short run time. 1% seems to be the best choice of gap size, as it is fast and results in high profit.

9.2.2 Test of heuristic

Ports included in value calculation

Figure 9.11 shows how changing the rule for port inclusion impacts profit and run time, respectively. The first alternative is to only include ports that are closer after the move. The second alternative is including all ports that are closer or at an equal distance in travel days after the move. The use of port inclusion is described more thoroughly in Section 5.4.2.

The model performs best when only including ports that are closer after the move. Including ports that do not get closer leads to the vessel being pulled between ports. This makes the vessel roam around rather than sailing to a port to trade. This is confirmed when examining the routes chosen by the vessel. The test with equal and closer ports spends more time waiting in open sea and in ports, thus not reaching as many trades as the test with decreasing horizon.

The run times are basically the same for both alternatives. This is because all the port

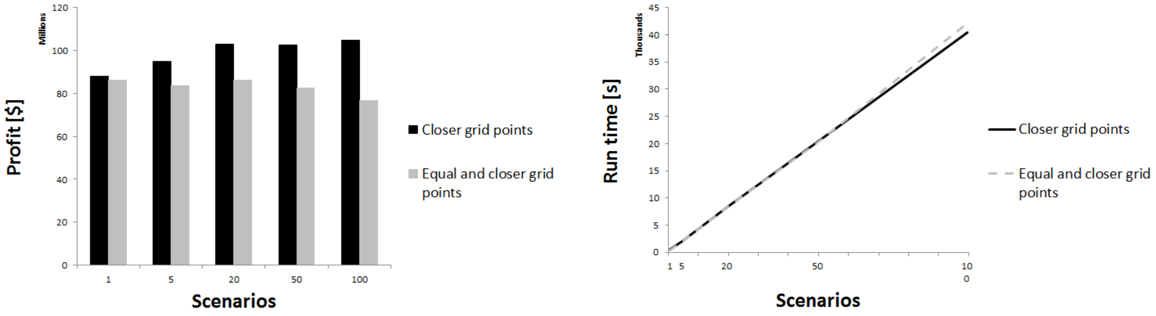

(a) Profit (b) Run time

Figure 9.11: Profit and run time when varying the number of scenarios and port inclusion in the stochastic version of the heuristic

calculations have to be made by the DP in both alternatives. The only difference is which port values are included when considering each of the potential moves.

Number of Scenarios

Figure 9.12 show how the number of scenarios impact the results and run times. respectively.

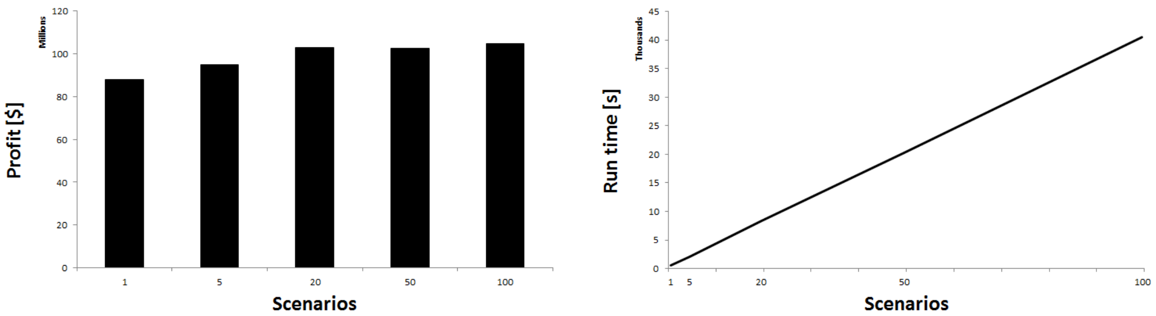

(a) Profit (b) Run time

Figure 9.12: Profit and run time with decreasing horizon when varying the number of scenarios in the stochastic version of the heuristic

Increasing the number of scenarios seems to have a positive effect on the profit. A higher number of scenarios gives a better approximation of future port prices, and thus a better basis for making good decisions. The effect seems to be somewhat diminishing after 20 scenarios. This is probably because 20 scenarios is enough to give a balanced impression of future prices.

When examining the routes we find that the tests with few scenarios hesitates when making big decisions in the tests. This can for example be a decision of crossing the ocean. In the tests with many scenarios, the vessel is confident and crosses the ocean from Europe to America. This is opposite of the MIP, which is indecisive when faced with multiple scenarios. In the tests with few scenarios there is more hesitation, which costs valuable time. It looks like the small number of scenarios is not enough to be confident about a move.

Increasing the number of scenarios leads to a linear growth in run time. This seems natural as the heuristic solves one scenario at a time.

Type of rolling horizon

Figure 9.13 show a comparison between the last test and the same model run with a fixed length of the rolling horizon.

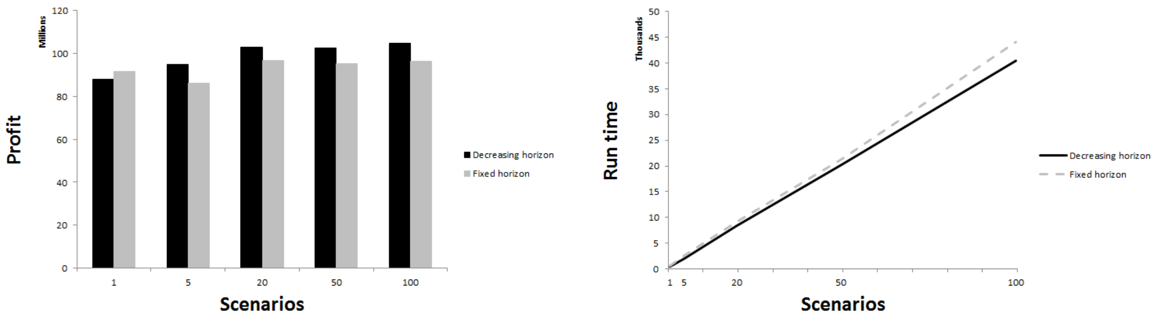

(a) Profit (b) Run time

Figure 9.13: Profit and run time when varying number of scenarios and type of rolling horizon in the stochastic version of the heuristic

The model performs best when using decreasing length of the rolling horizon. This is due to less hesitation by the vessel. A fixed horizon keeps the options open for the vessel, thus not giving any incentives to move. When examining the routes chosen by the different tests, we see that the vessel in the fixed horizon falls behind from the start, by spending more time in the sea and ports.

Solving with decreasing rolling horizon length is a bit faster than with fixed horizon length. This is because the dynamic program on average is run for shorter trade horizons when using the decreasing rolling horizon.

Stability

We start by presenting combined results for the stability tests. We continue by discussing the results from the in-sample test and out-of-sample test for different vessel situations. These situations are described in Section 8.5.3

Overall stability

Figure 9.14 shows total in- and out-of-sample stability for the heuristic.

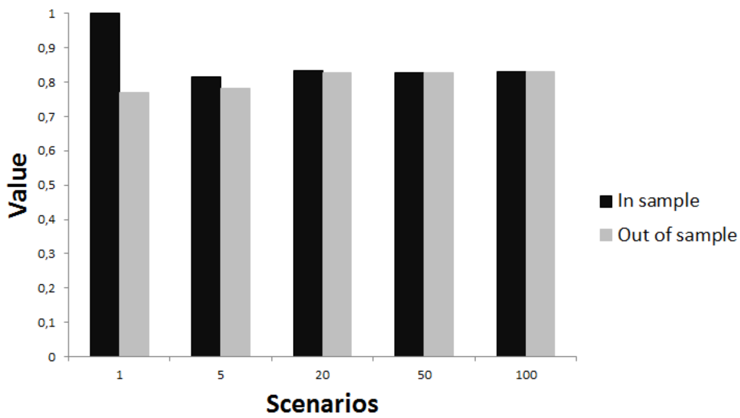

Figure 9.14: Overall in- and out-of-sample values when varying the number of scenarios in the stochastic version of the heuristic

The in-sample value for one scenario is neglected in the discussion, as it is defined as 1.

The in-sample values are higher than the out-of-sample values, as expected. Both the in-sample values and the out-of-sample values increase with the number of scenarios. This is reasonable. A model based on many scenarios is more likely to find a solution that correlates with the independent scenarios. The more scenarios, the less variance and the more correlation. For the out-of-sample a model with many scenarios is likely to be better suited for testing with a range of random scenarios. When the number of scenarios increase, the stability values from the in-sample test and out-of-sample tests converge towards each other. This is logical, since the samples of scenarios in the in-sample and out-of-sample tests become more similar to each other as the number of scenarios increase.

In-sample stability

Figure 9.15 shows the in-sample test values for different vessel situations.

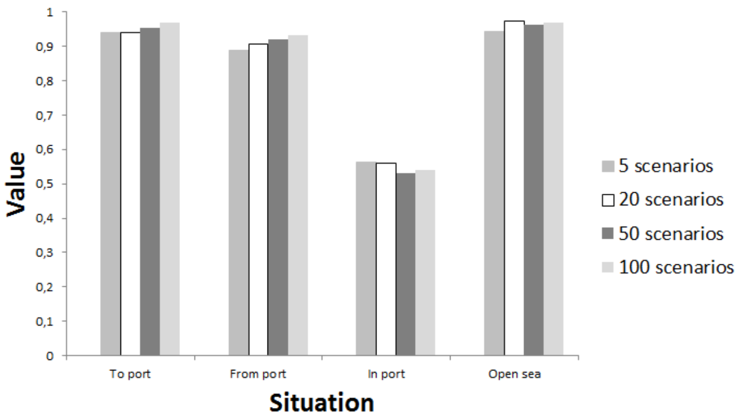

Figure 9.15: In-sample values for different situations when varying the number of scenarios in the stochastic version of the heuristic

We can see two main trends. The first is that the stability values is generally higher for "to port" and "open sea". These are easier decisions to make than "from port", and it is logical that the scenarios correlate more for these situations. "From port" can in theory choose to go to every port, hence the vessel has the possibility of going to many different grid points. For "to port" and in "open sea" the decision of which port to visit is more *settled* (if the vessel is close to a port value of this port becomes larger relative to other ports), hence limiting the choices for the vessel. The second is that the stability values increase as the number of scenarios increase for the "from port" situation. We believe that this is because the "from port" situation is tough, which increases the benefit of having many scenarios.

Out-of-sample stability

Figure 9.16 shows the results of the out-of-sample stability tests for different vessel situations.

We find three trends. The values of "to port" and "open sea" is generally higher than "from port". The reason for this is that "to port" and "open sea" is easier decisions than "from port". The added value of multiple scenarios is apparent both in "to port" and "from port". The effect is strongest when leaving a port. This indicates that the use of many scenarios has most impact when making tough decisions.

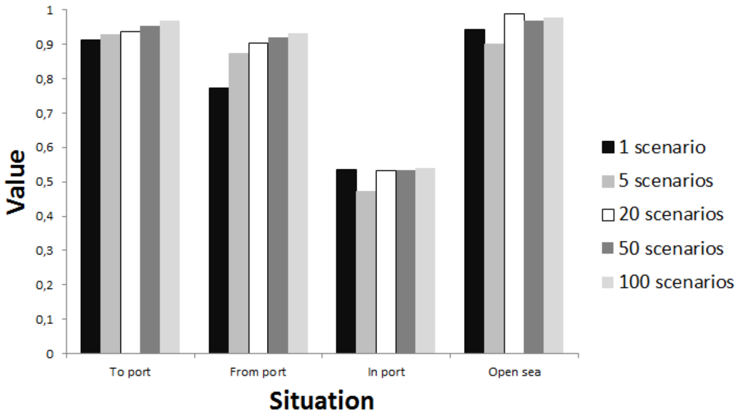

Figure 9.16: Out-of-sample values for different situations when varying the number of scenarios in the stochastic version of the heuristic

Conclusion on parameter values

It is apparent that decreasing rolling horizon length should be used when running heuristic tests. It provides better results in a shorter run time. EOH-values should also be used, since it enhances the performance without increasing the run time. The optimal number of scenarios is a tougher trade off, as more scenarios results in a longer run time. The results indicate a marginal growth in profit when increasing the number of scenarios beyond 20. It does however only take 200 seconds to solve the problem for one day when using 100 scenarios. It would thus be reasonable to use 100 scenarios.

9.2.3 Comparison between MIP and heuristic

Figure 9.17 show a comparison of profit and run time for the stochastic versions of the MIP and heuristic.

Profit

The heuristic generally give better solutions than the MIP model. Adding more than 5 scenarios is positive for the heuristic result, while the impact seems to be opposite for the MIP.

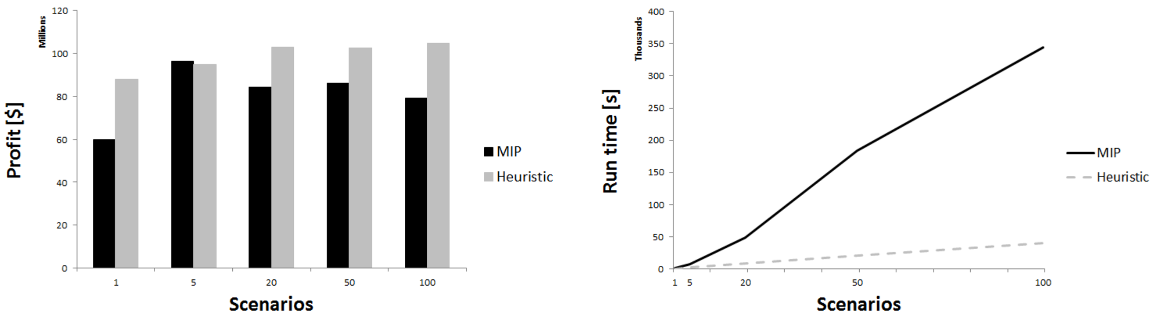

(a) Profit

(b) Run time

Figure 9.17: Comparison of profit and run time between the stochastic versions of the MIP and heuristic when varying the number of scenarios

Run time

The run time clearly favors the heuristic for all tests. The difference increases with the number of scenarios.

Stability

For both in-sample-stability and out-of-sample-stability, the heuristic outperforms the MIP. It seems like the heuristic gets increasingly better stability value as the number of scenarios increase, while the opposite happens for the MIP. The results from the MIP is surprising, as the most natural development is that the stability increases as the number of scenarios increases.

It is hard to find an obvious explanation for why the MIP and the heuristic develop in opposite directions. One hypothesis is the number of possible options for the models. The MIP has to consider all possible routes 20 days ahead, while the heuristic just considers one move. The stability, and solution, seems to decrease with a increasing number of scenarios when there are many options to consider. This is shown both by the port inclusion of the heuristic test and the difference in trends of in sample values for "in port" and "from port" in the MIP. Our theory is that the MIP overall deals with to many options, leading to a worse result and stability as the number of scenarios increases.

There is a correlation between the performance in the stability tests and the profit of the stochastic models. This makes sense, as a low stability value indicates that few scenarios matches the solution, which suggests that the solution is bad.

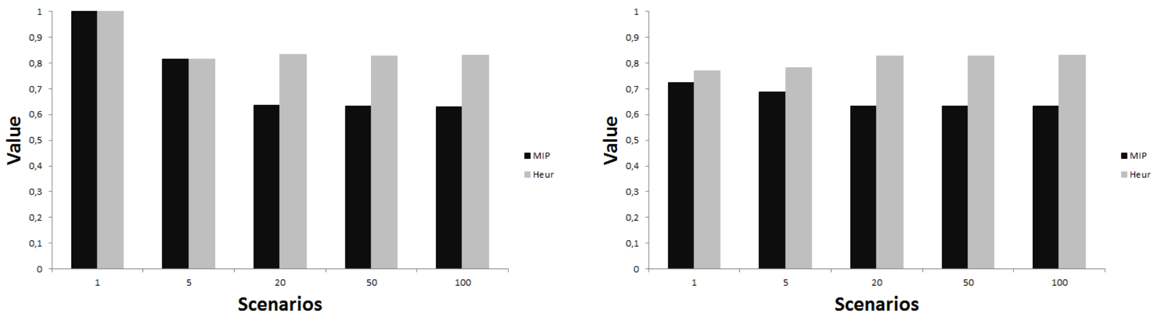

(a) In-sample stability (b) Out-of-sample stability

Figure 9.18: Comparison of stability between the stochastic versions of the MIP and heuristic when varying the number of scenarios

Route considerations

The vessels in the MIP and heuristic basically choose the same route. The best route from all our tests is found is shown in Section 9.3.3. This route is similar to the other routes from our solutions. The difference in profit is made up by the frequency of trade. The heuristic makes the vessel trade more often than the MIP, resulting in higher profit. The MIP hesitates more both in and in between ports. This is caused by the MIP's indecisiveness when faced with many opportunities.

Summary

The heuristic performs better than the MIP when considering both profit, run time and stability. It is the best solution approach for solving the stochastic problem.

9.3 Deterministic models

In this section we consider the results from the deterministic tests of our model. The deterministic model uses one price scenario that contains the expected value of all price scenarios under consideration. The expected values are calculated as the arithmetic mean of all price scenarios. In these tests we have only tested the impact of varying the number of scenarios, as our main goal with the tests is to discover any differences between the stochastic and deterministic models. The rest of the parameter are kept at their base case values, as described in Section 8.6.

9.3.1 Test of MIP

Figure 9.19 shows the impact on run time and profit when varying the number of scenarios used to compute the expected value.

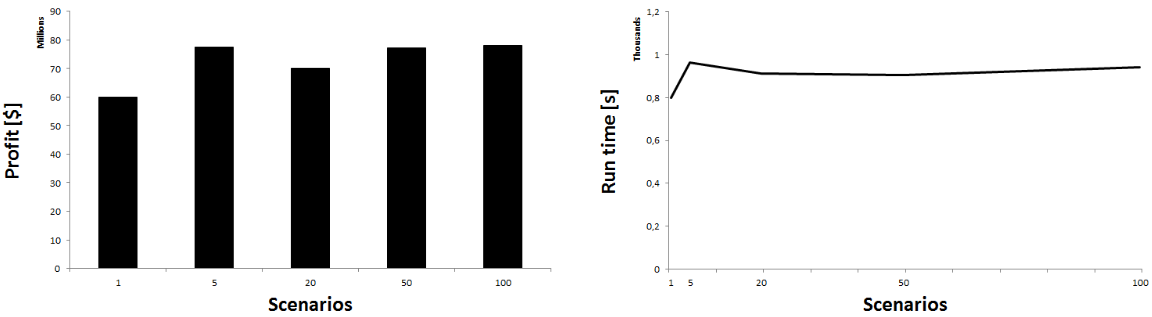

(a) Profit (b) Run time

Figure 9.19: Profit and run time in the deterministic version of the MIP

The results seems to be increasing with the number of scenarios. This is logical, since an expected value of a set of scenarios becomes closer to the true mean when the number of scenarios increase. Since a random walk process is used when simulating prices, the true mean should be a reasonable indication of future prices. The tests with one and five scenarios has more variance than the others, as the expected value computed may happen to correlate well with the actual price development. It looks like the EV of five scenarios happens to fit well with the actual price development.

The run time is equal for all scenarios. This is as expected, since the problem solved basically is the same for all tests - by solving one scenario with an expected value.

9.3.2 Test of heuristic

Figure 9.20 show the impact on profit and run time when varying the number of scenarios used to compute the expected value.

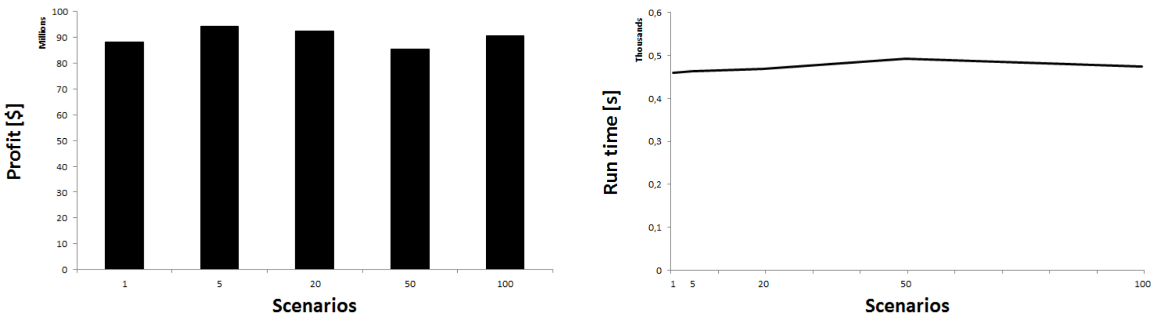

(a) Profit (b) Run time

Figure 9.20: Profit and run time in the deterministic version of the heuristic

There is no clear trend of how the number of scenarios impact the profit of the heuristic. There does not seem to be any gain from using more scenarios when calculating the expected value. This is a bit unexpected, since we believed that an expected value close to the true mean would perform well.

The run times are similar for all tests. This is expected since the same problem is run for each test, only varying the expected valued of the price scenario.

9.3.3 Comparison between MIP and heuristic

Figure 9.21 shows a comparison of the deterministic MIP and heuristic for different number of scenarios.

Profit

The heuristic outperforms the MIP in the deterministic runs. It has higher profit for all test instances.

Run time

The run time is almost double for running the MIP compared to the heuristic. The run times are however negligible, as the daily run time on an operational level would be less than $\frac{1000}{200} = 5$ seconds.

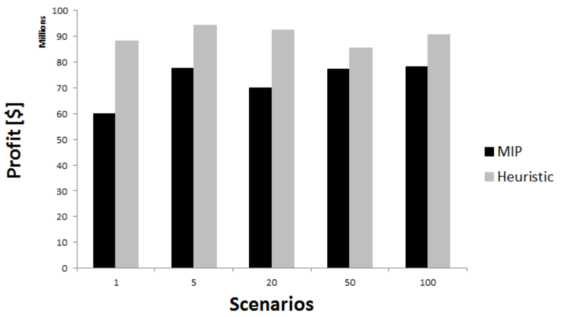

(a) Profit

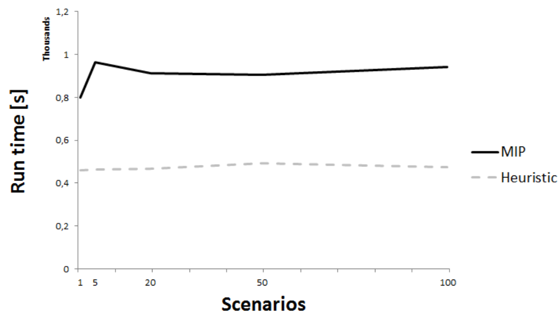

(b) Run time

Figure 9.21: Comparison of profit and run time between the deterministic versions of the MIP and heuristic

Route considerations

The heuristic trades more frequently than the MIP. The vessel in the MIP is indecisive, with many waiting days. The waiting days leads to a smaller number of trades and less profit. This is caused by the vessel being attracted by many ports at the same time, not being able to decide which one it should go for.

Summary

The heuristic gives better results for all test runs. It is the best solution approach for solving the deterministic problem.

9.4 Comparison between stochastic and deterministic models

In this section we discuss the differences between the stochastic and deterministic versions of the models when it comes to profit and run time. We show results for both the MIP and the heuristic.

9.4.1 MIP

Figure 9.22 show a comparison of profit and run time for the stochastic and deterministic tests of the MIP.

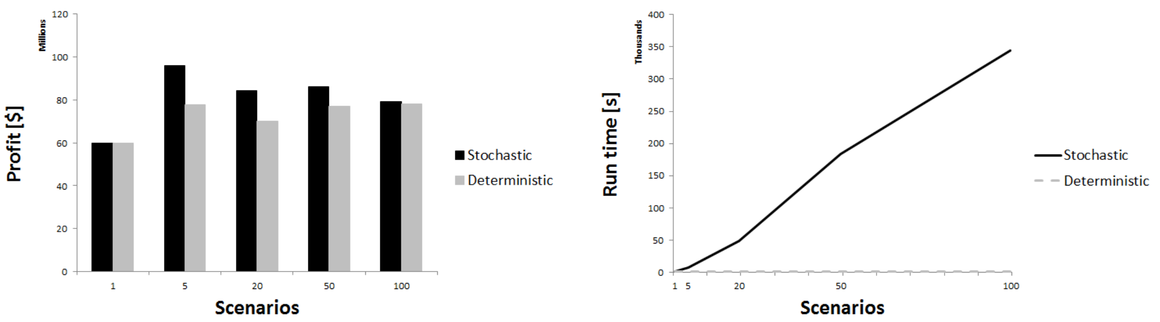

(a) Profit

(b) Run time

Figure 9.22: Comparison of profit and run time between the stochastic and deterministic versions of the MIP

The stochastic version of the MIP performs better than the deterministic version. The difference diminishes as the number of scenarios increase. This is because the vessel hesitates when having to consider many alternatives. As the number of scenarios increases the number of options increase, and the vessel hesitates more. The hesitation outweighs the gain of having more information.

The run time of the stochastic model is significantly higher than for the deterministic model, and increases linearly with the number of scenarios. This is logical since the stochastic problem is far more complex than the deterministic.

9.4.2 Heuristic

Figure 9.23 shows a comparison of profit and run time for the stochastic and deterministic tests of the heuristic.

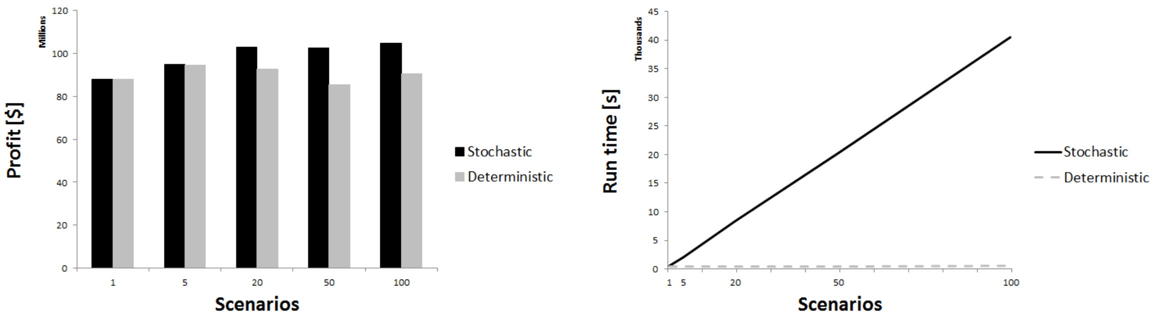

(a) Profit (b) Run time

Figure 9.23: Comparison of profit and run time between the stochastic and deterministic versions of the heuristic

The profit from running the stochastic model increases with the number of scenarios. The same trend is *not* seen for the deterministic version. The reason for a higher profit is that the stochastic model uses more of the information given, thus being able to create a better result.

The run time for the stochastic version is significantly higher than for the deterministic version, and increases linearly with the number of scenarios. This makes sense since the stochastic version solves the same problem as the deterministic problem one time for each scenario.

9.4.3 Route considerations

The vessel hesitates less in the stochastic versions of the models, resulting in more trades. The difference is especially big early in the problem period, indicating that the decreasing horizon approach is exploited better by the stochastic than the deterministic model. The timing of the trades when in ports are also better in the stochastic versions, meaning that the vessel performs more profitable trades.

In Figure 9.24 we have shown the route of the best solution. This is the heuristic with a 20-day horizon using EOH-values, shown in Figure 9.2. The vessel uses three different trade routes between the ports. Trade 1 is between Hammerfest and Milford Haven, trade 2 is between Point Fortin and St Johns while trade 3 is between Point Fortin and Milford Haven.

The vessel starts in Milford Haven, and conducts trade 1 six times. Then the vessel sails to Point Fortin and conducts trade 2 two times. It continues by doing trade 3 once, before doing trade 2 four more times. Most of the routes from our test instances follow the same pattern, but with differences in travelling time between the ports and waiting time in the ports.

Summary

The stochastic versions outperforms the deterministic ones for both the MIP and the heuristic. The value of the stochastic solutions (VSS) for different scenarios are shown in Table 9.1. This table shows that the value of the stochastic solutions seem to increase with an increasing number of scenario for the heuristic, while the opposite happens for the MIP. The VSS is not shown for one scenario, since the stochastic solution of one scenario by definition is the same as the deterministic solution.

Table 9.1: Value of stochastic solutions in million $ for a varying number of scenarios

Scenarios	MIP	Heuristic
5	\$18.6”	\$0.5”
20	\$14.3”	\$10.4”
50	\$9.1”	\$17.0”
100	\$1.2”	\$14.0”

The run time for the stochastic models are significantly higher than for the deterministic models. A trade-off has to be made between profit and run time. Solving the stochastic version of the heuristic with 100 scenarios is still quite fast, solving in less than 5 minutes. This seems like a reasonable amount of time to use when solving a daily problem.

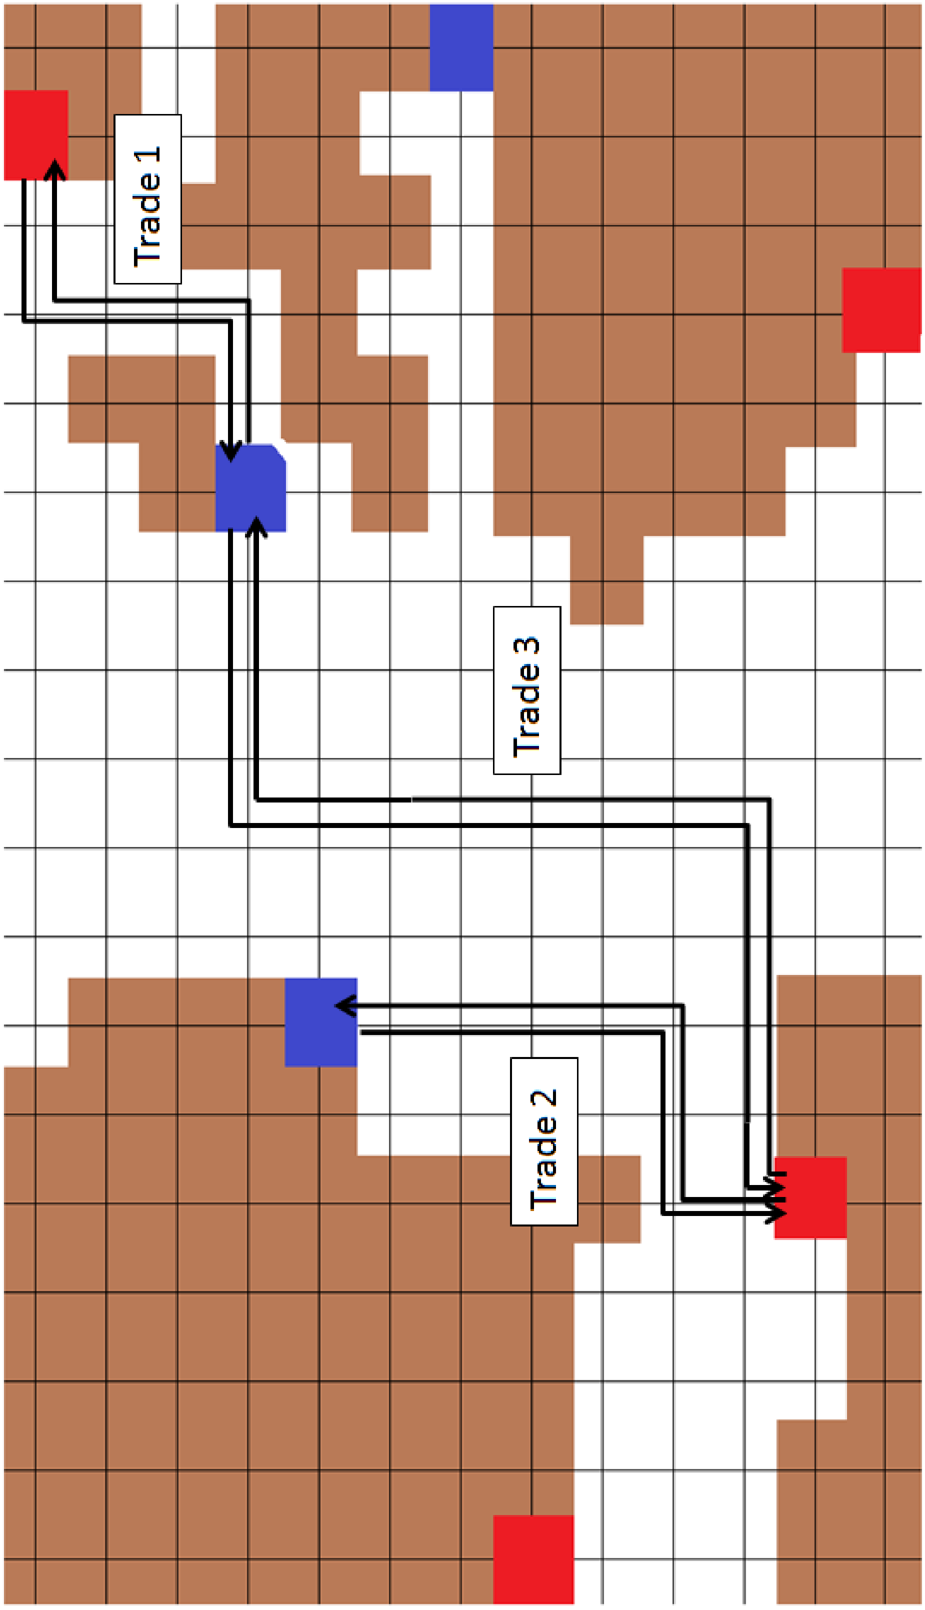

Figure 9.24: Trades conducted by the vessel sailing the route with highest profit

Chapter 10

Concluding Remarks and Further Research

In this thesis we have developed two decision support models for trading LNG on a speculative basis. One model is based on a Mixed Integer Program approach, the other on a heuristic. Given multiple scenarios of spot prices in different ports, the models produce movement and trade decisions for LNG vessels. The goal is to maximize profits by buying cheap and selling expensive, while minimizing cost related to moving between ports. A review of existing literature found no previously published material with the same focus.

Golar LNG tried to do speculative trading of LNG in 2002, but were not able to find attractive tonnage for their vessels. LNG markets have opened up a lot since then, with increased diversification and more LNG available for open trade. The share of flexible contracts, short-term contracts and spot trade have increased. New opportunities arise with this market development, and tools that address these opportunities are needed. On a strategic level it is interesting to research whether speculative trading of LNG is profitable at all. On a tactical level it is interesting to look at how to best structure the use of flexible long-term contracts and spot trade. The models presented in this thesis are relevant on an operational level. They can be used for continuous speculative spot trading decisions, but with small modifications they can potentially also be used to analyze when it is profitable to divert flexible LNG cargoes.

We have tested our models in to determine the one with best capability of being used as a decision support tool, based on profit and run time. Both models have been tested with different values of parameters, in order to see which parameters that has the largest impact. The stochastic versions are compared to the deterministic ones, in order to research whether there is any advantage of using a stochastic approach.

10.1 Conclusion

The heuristic model performs better than the MIP. The main difference between the routes found in the heuristic and MIP is that the vessel hesitates less in the heuristic. This increases the number of trades, and also the profit. It seems like the MIP performs worse because it has several options that it needs to consider, and is not able to decide which one to go for. In the heuristic the value of a given port is only included in a move calculation if the vessel is heading towards that port. The MIP considers all ports when calculating the value a move, and calculates a route of 20 days. When testing the heuristic with more ports included in the move calculation, we get worse results. This is another indication that more options give worse solutions.

The stochastic version of the heuristic performs better than the deterministic version. The stochastic model is able to use the information from the additional scenarios to create a significantly better solution. The value of the stochastic solution increases with the number of scenarios, but so does also the run time. The gain of adding scenarios seems to be highest for the first 20 scenarios.

The models have not been compared to real life trading. It is however usable for real life situations. New ports are easily added and the various parameters can be changed. The model can also be used for any other kind of tonnage, with only minor adjustments. It is especially well suited for bulk shipping, as we assume full shiploads and a given price per unit of shipped goods.

10.2 Further research

We present six areas of further research. These are discussed in the next paragraphs.

The EOH-values can be improved. It is likely that the EOH-values could be improved to generate better results. The current version is simple, with a predefined value given per flexible day at the end of the trade horizon, and a compensation for buying as the last trade. The flexible day value could probably be tweaked further to better approximate the value of finishing the last trade early. The value could e.g. be based on price forecasts. This is hard to implement, as the vessel can choose to travel to any other port, and each of the ports has different price development. The value would have to consider all of the price developments. The compensation for buying could also better approximate the actual value it represents, based on what the return from a future sell is likely to be.

The price forecast can be more realistic. The price forecasts have a large impact on the economical results of our models. In the current model we base the price forecast for all ports on historical data from the Henry Hub. These forecasts could be substantially improved by collecting historical prices and other relevant information for all ports considered. We have not been able to find other historical price data then from Henry Hub and the National Balancing Point. More complete price data could be used to improve the forecasts generated by the Ornstein-Uhlenbeck process, or as a basis for another price forecasting method.

The map representation can be enhanced. Our models uses a grid and movements between grid points to approximate the geographical and time aspects of our problem. This is a highly simplified representation of real life. It should be possible to make finer grid points or even try to make it possible to travel in any direction. This can be done in combination with lowering the length of the time intervals. This should raffine the move that could be chosen to travel, and thus enhance the profit.

Using different number of scenarios for different situations By examining the out-of-sample values from the stability tests, we see that there is a need of more scenarios when tough decisions are made, for example when leaving a port. A future model could use many scenarios when making hard decisions, and few for the easier ones.

Improve the MIP. If the MIP is to be used in further research, it has to be improved. One way to improve the MIP could be by introducing a heuristic that only considers closer grid points. Then the MIP would basically solve the same problem as the heuristic. Including a smaller gap in order to get a better answer is not an option, as run time would increase significantly, and the model could not be used in day-to-day planning.

The results can be compared to alternative uses of the LNG vessel. We do not know how our models perform against using the vessel for other types of shipping. It would be interesting to test our model on real test instances and compare the results to alternative uses of the vessel, such as renting it out for fulfilling long-term contracts.

Bibliography

A. Ainouche and Dr A. Smati. Optimization of LNG chain by stochastic dynamic programming model. In *17th World Petroleum Congress*, 2002.

J. Alvarez, P. Tsilingiris, E. Engebrethsen, and N. Kakalis. Robust fleet sizing and deployment for industrial and independent bulk ocean shipping companies. *Information Systems and Operational Research*, 49(2):93–107, 2011.

H. Andersson, M. Christiansen, and K. Fagerholt. Transportation planning and inventory management in the LNG supply chain. In E. Bjørndal, M. Bjørndal, P. M. Pardalos, and M. Rönnqvist, editors, *Energy, Natural Resources and Environmental Economics*, Energy Systems, pages 427–439. Springer Berlin Heidelberg, 2010.

J. Ang, C. Cao, M. Parlar, and H. Ye. Two-stage stochastic integer programming model for multiperiod sea cargo mix problem in container shipping industry. *IEEE Transactions on Systems, Man, and Cybernetics Part A: Systems and Humans*, 39(2):460–465, 2009.

F. Asche, B. Misund, and M. Sikveland. The relationship between spot and contract gas prices in Europe. *Energy Economics*, 38(0):212–217, 2013.

A. Aspelund, T. Gundersen, J. Myklebust, M.P. Nowak, and A. Tomasgard. An optimization-simulation model for a simple LNG process. *Computers & Chemical Engineering*, 34(10):1606–1617, 2010.

A. Azaron and F. Kianfar. Dynamic shortest path in stochastic dynamic networks: Ship routing problem. *European Journal of Operational Research*, 144(1):138–156, 2003.

R. W. Bent and P. Van Hentenryck. Scenario-based planning for partially dynamic vehicle routing with stochastic customers. *Operations Research*, 52(6):977–987, 2004.

G. Berbeglia, J. F. Cordeau, and G. Laporte. Dynamic pickup and delivery problems. *European Journal of Operational Research*, 202(1):8–15, 2010.

O. Berman and R. C. Larson. Deliveries in an inventory/routing problem using stochastic dynamic programming. *Transportation Science*, 35(2):192–213, 2001.

D. Bertsimas and M. Sim. Robust discrete optimization and network flows. *Mathematical Programming*, 98(1-3):49–71, 2003.

C. Boutilier, R. Dearden, and M. Goldszmidt. Stochastic dynamic programming with factored representations. *Artificial Intelligence*, 121(1):49–107, 2000.

BP. BP World Energy Outlook 2030. January 2013a.

BP. BP Statistical Review of World Energy. June 2013b.

G. Brønmo, M. Christiansen, K. Fagerholt, and B. Nygreen. A multi-start local search heuristic for ship scheduling - a computational study. *Computers & Operations Research*, 34(3):900–917, 2007a.

G. Brønmo, M. Christiansen, and B. Nygreen. Ship routing and scheduling with flexible cargo sizes. *Journal of the Operational Research Society*, 58(9):1167–1177, 2007b.

G. Brønmo, B. Nygreen, and J. Lysgaard. Column generation approaches to ship scheduling with flexible cargo sizes. *European Journal of Operational Research*, 200(1):139–150, 2010.

M. Christiansen and K. Fagerholt. Robust ship scheduling with multiple time windows. *Naval Research Logistics*, 49(6):611–625, 2002.

M. Christiansen and K. Fagerholt. Maritime inventory routing problems. In C. Floudas and P. Pardalos, editors, *Encyclopedia of Optimization*, pages 1947–1955. Springer US, 2009.

M. Christiansen, K. Fagerholt, and D. Ronen. Ship routing and scheduling: Status and perspectives. *Transportation Science*, 38(1):1–18, 2004.

M. Christiansen, K. Fagerholt, B. Nygreen, and D. Ronen. Chapter 4 Maritime Transportation. *Handbooks in Operations Research and Management Science*, 14(C):189–284, 2007.

M. Christiansen, K. Fagerholt, B. Nygreen, and D. Ronen. Ship routing and scheduling in the new millennium. *European Journal of Operational Research*, 228(3):467–483, 2013.

M. P. Cristobal, L. F. Escudero, and J. F. Monge. On stochastic dynamic programming for solving large-scale planning problems under uncertainty. *Computers & Operations Research*, 36(8):2418–2428, 2009.

S. Desai and G. J. Lim. Solution time reduction techniques of a stochastic dynamic programming approach for hazardous material route selection problem. *Computers & Industrial Engineering*, 65(4):634–645, 2013.

EIA. International Energy Outlook 2013. July 2013.

Bergen Energy. Global gas market development and the oil-gas link, 2010. URL `https://www.bergen-energi.com/arch/_img/9272859.pdf`. Accessed: 2014-19-03.

Energy Charter Sec. Developments in LNG Trade and Pricing, 2009. URL `http://www.encharter.org/fileadmin/user_upload/document/LNG_2009_ENG.pdf`. Accessed: 2014-19-03.

S. Engelen and W. Dullaert. Transformations in gas shipping: Market structure and efficiency. *Maritime Economics & Logistics*, 12(3):295–325, 2010a.

S. Engelen and W. Dullaert. Designing a decision support model for the LNG market. *Transportation Planning & Technology*, 33(8):719–731, 2010b.

Exxon Mobile. The Outlook For Energy: A View To 2040. 2013.

K. Fagerholt, M. Christiansen, L. M. Hvattum, T. Johnsen, and T. Vabø. A decision support methodology for strategic planning in maritime transportation. *Omega*, 38(6):465–474, 2010.

T. Flatberg, G. Hasle, O. Kloster, E. J. Nilssen, and A. Riise. Dynamic and stochastic vehicle routing in practice. In V. Zeimpekis, C. D. Tarantilis, G. M. Giaglis, and I. Minis, editors, *Dynamic Fleet Management*, volume 38 of *Operations Research/Computer Science Interfaces Series*, pages 41–63. Springer US, 2007.

M. Fodstad, K. T. Uggen, F. Rømo, A. Lium, G. Stremersch, and S. Hecq. LNGSheduler: A rich model for coordinating vessel routing, inventories and trade in the liquefied natural gas supply chain. *The Journal of Energy Markets*, 3(4):31–64, 2011.

GIIGNL. The LNG Industry in 2012. 2013.

K. G. Gkonis and H. N. Psaraftis. The LNG market: a game theoretic approach to competition in LNG shipping. *Maritime Economics & Logistics*, 11(2):227–246, 2009.

Global Security, 2013. URL `http://www.globalsecurity.org/military/systems/ship/tanker-lng-history.htm`. Accessed: 2013-12-02.

V. Goel, K.C. Furman, J. H. Song, and A.S. El-Bakry. Large neighborhood search for LNG inventory routing. *Journal of Heuristics*, 18(6):821–848, 2012.

Golar LNG. Golar LNG Q3 Results 2013, 2013a. URL `http://hugin.info/133076/R/1746651/588040.pdf`. Accessed: 2013-12-11.

Golar LNG. Golar LNG Q2 Results 2013 Presentation, 2013b. URL `http://www.slideshare.net/TradeWindsnews/golar-lng-q2-2013-results-presentation`. Accessed: 2013-12-11.

R. Grønhaug and M. Christiansen. Supply chain optimization for the liquefied natural gas business. In J. Nunen, M. G. Speranza, and L. Bertazzi, editors, *Innovations in Distribution Logistics*, volume 619 of *Lecture Notes in Economics and Mathematical Systems*, pages 195–218. Springer Berlin Heidelberg, 2009.

R. Grønhaug, M. Christiansen, G. Desaulniers, and J. Desrosiers. A branch-and-price method for a liquefied natural gas inventory routing problem. *Transportation Science*, 44(3):400–415, 2010.

E. E. Halvorsen-Weare and K. Fagerholt. Routing and scheduling in a liquefied natural gas shipping problem with inventory and berth constraints. *Annals of Operations Research*, 203(1):167–186, 2013.

E. E. Halvorsen-Weare, K. Fagerholt, and M. Rönnqvist. Vessel routing and scheduling under uncertainty in the liquefied natural gas business. *Computers & Industrial Engineering*, 64(1):290–301, 2013.

J. L. Higle. Stochastic programming: optimization when uncertainty matters. *Tutorials in Operations Research*, pages 30–53, 2005.

L. M. Hvattum, A. Lokketangen, and G. Laporte. A branch-and-regret heuristic for stochastic and dynamic vehicle routing problems. *Networks*, 49(4):330–340, 2007.

H. S. Hwang, S. Visoldilokpun, and M. J. Rosenberger. A branch-and-price-and-cut method for ship scheduling with limited risk. *Transportation Science*, 42(3):336–351, 2008.

J. H. Hwang, M. I. Roh, and K. Y. Lee. Determination of the optimal operating conditions of the dual mixed refrigerant cycle for the LNG FPSO topside liquefaction process. *Computers & Chemical Engineering*, 49:25–36, 2013.

S. Ichoua, M. Gendreau, and J. Y. Potvin. Exploiting knowledge about future demands for real-time vehicle dispatching. *Transportation Science*, 40(2):211–225, 2006.

J. T. Jensen. The development of a global LNG market. *Oxford Institute for Energy Studies*, 2004.

J. Kelman, J.R. Stedinger, L.A. Cooper, E. Hsu, and S. Yuan. Sampling stochastic dynamic programming applied to reservoir operation. *Water Resources Research*, 26 (3):447–454, 1990.

R. Khalilpour and I.A. Karimi. Contract selection under uncertainty. LNG buyers' perspective. *Computer Aided Chemical Engineering*, 31:1487–1491, 2012.

S. H. Kim and K. K. Lee. An optimization-based decision support system for ship scheduling. *Computers & Industrial Engineering*, 33(3-4):689–692, 1997.

J. E. Korsvik, K. Fagerholt, and G. Laporte. A tabu search heuristic for ship routing and scheduling. *Journal of the Operational Research Society*, 61(4):594–603, 2010.

D. Y. Lin and H. Y. Liu. Combined ship allocation, routing and freight assignment in tramp shipping. *Transportation Research Part E: Logistics & Transportation Review*, 47(4):414–431, 2011.

Lloyd's, 2013. URL `http://www.lloydslistintelligence.com/llint/gas/index.htm`. Accessed: 2013-12-11.

A. A. Lopac. Recent trends in transporting of LNG, liquefied natural gas. In *11th International Conference on Transport Science,„Transport Policy", Słowenia*, pages 28–29, 2008.

F. Malliappi, J.A. Bennell, and C.N. Potts. A variable neighborhood search heuristic for tramp ship scheduling. *Lecture Notes in Computer Science*, 6971 LNCS:273–285, 2011.

D. Maxwell and Z. Zhu. Natural gas prices, LNG transport costs, and the dynamics of LNG imports. *Energy Economics*, 33(2):217–226, 2011.

K. McKinnon and Y. Yu. Solving stochastic ship fleet routing problems with inventory management using branch and price. 2011.

S. Mitrović-Minić, R. Krishnamurti, and G. Laporte. Double-horizon based heuristics for the dynamic pickup and delivery problem with time windows. *Transportation Research Part B: Methodological*, 38(8):669–685, 2004.

I. Norstad, K. Fagerholt, and G. Laporte. Tramp ship routing and scheduling with speed optimization. *Transportation Research Part C: Emerging Technologies*, 19(5):853–865, 2011.

C. Novoa and R. Storer. An approximate dynamic programming approach for the vehicle routing problem with stochastic demands. *European Journal of Operational Research*, 196(2):509–515, 2009.

OECD, 2013. URL `http://www.oecd.org/about/membersandpartners/`. Accessed: 2013-12-13.

outline-world map.com. http://www.outline-world-map.com/. Accessed: 2014-04-16.

V. Pillac, M. Gendreau, C. Guéret, and A. L. Medaglia. A review of dynamic vehicle routing problems. *European Journal of Operational Research*, 225(1):1–11, 2013.

Qatargas, 2013. URL http://www.qatargas.com/English/MediaCenter/news/Pages/Article_120.aspx. Accessed: 2013-12-02.

J. G. Rakke, M. Stålhane, C.R. Moe, M. Christiansen, H. Andersson, K. Fagerholt, and I. Norstad. A rolling horizon heuristic for creating a liquefied natural gas annual delivery program. *Transportation Research Part C: Emerging Technologies*, 19(5):896–911, 2011.

Reuters, 2013. URL http://www.reuters.com/article/2013/08/05/idUSnHUGddv2+70+ONE20130805. Accessed: 2013-12-11.

M. Savelsbergh and M. Sol. DRIVE: Dynamic routing of independent vehicles. *Operations Research*, 46(4):474–490, 1998.

R. G. Schwimmbeck, 2008. URL http://www.pipeline-conference.com/abstracts/pipeline-vs-lng. Date Accessed: 2013-12-06.

W. Shao, P. Zhou, and S. K. Thong. Development of a novel forward dynamic programming method for weather routing. *Journal of Marine Science & Technology (Japan)*, 17(2):239–251, 2012.

A. Shapiro. Analysis of stochastic dual dynamic programming method. *European Journal of Operational Research*, 209(1):63–72, 2011.

Shell, 2013. URL http://www.shell.com/global/aboutshell/major-projects-2/prelude-flng.html. Accessed: 2013-12-08.

A. Shyshou, I. Gribkovskaia, and J. Barceló. A simulation study of the fleet sizing problem arising in offshore anchor handling operations. *European Journal of Operational Research*, 203(1):230–240, 2010.

SLNG, 2010. URL http://www.slng.com.sg/about-lng-value-chain.html. Accessed: 2013-09-17.

J. Song and K. C. Furman. A maritime inventory routing problem: Practical approach. *Computers & Operations Research*, 2010.

M. Stålhane, J. G. Rakke, C. R. Moe, H. Andersson, M. Christiansen, and K. Fagerholt. A construction and improvement heuristic for a liquefied natural gas inventory routing problem. *Computers & Industrial Engineering*, 62(1):245–255, 2012.

G. Tirado, L. M. Hvattum, K. Fagerholt, and J. F. Cordeau. Heuristics for dynamic and stochastic routing in industrial shipping. *Computers & Operations Research*, 40 (1):253–263, 2013.

K. T. Uggen, M. Fodstad, and V. S. Nørstebø. Using and extending fix-and-relax to solve maritime inventory routing problems. *TOP*, 21(2):355–377, 2013.

G. E. Uhlenbeck and L. S. Ornstein. On the Theory of the Brownian Motion. *Physical Review*, 36:823–841, September 1930.

J. A. Villar and F. L. Joutz. The relationship between crude oil and natural gas prices. *Energy Information Administration, Office of Oil and Gas*, 2006.

P. E. Wahl, S. W. Løvseth, and M. J. Mølnvik. Optimization of a simple LNG process using sequential quadratic programming. *Computers & Chemical Engineering*, 56:27–36, 2013.

D. A. Wood. A review and outlook for the global LNG trade. *Journal of Natural Gas Science & Engineering*, 9(0):16–27, 2012.

X. Yang and A. Grothey. Solving the top-percentile traffic routing problem by approximate dynamic programming. *IMA Journal Management Mathematics*, 23(4):413–434, 2012.

Appendix A

Result Tables

All the profits in the tables are rounded of to the closest $100,000.

Table A.1: Profit in million $ and run time without EOH-values when varying rolling horizon length in the DP

Scenarios	Profit	Run time
20	$85.9”	96
30	$78.0”	519
40	$86.7”	2,189
50	$99.0”	5,001
60	$96.1”	9,042
70	$94.5”	14,364
80	$90.3”	20,637
90	$82.1”	24,955
100	$78.2”	29,480

Table A.2: Profit in million $ and run time with EOH-values when varying rolling horizon length in the DP

Scenarios	Profit	Run time
20	$110.7”	95
30	$102”	365
40	$95.4”	1,258
50	$95.6”	4,269
60	$94.6”	8,480
70	$94.2”	13,699
80	$100.4”	22,113
90	$85.1”	26,714
100	$81.6”	31,027

Table A.3: Profit in million $ and run time when varying rolling horizon length and use of EOH-values in the DP

Trade horizon length	Without EOH-values		With EOH-values	
	Profit	Run time	Profit	Run time
20	$85.9”	96	$110.7”	95
30	$78.0”	519	$102”	365
40	$86.7”	2,189	$95.4”	1,258
50	$99.0”	5,001	$95.6”	4,269
60	$96.1”	9,042	$94.6”	8,480
70	$94.5”	14,364	$94.2”	13,699
80	$90.3”	20,637	$100.4”	22,113
90	$82.1”	24,955	$85.1”	26,714
100	$78.2”	29,480	$81.6”	31,027

Table A.4: Profit in million $ and run time with decreasing horizon when varying the number of scenarios in the stochastic version of the MIP

Scenarios	Profit	Run time
1	$60.0”	806
5	$96.2”	7,564
20	$84.3”	49,193
50	$86.3”	184,165
100	$79.3”	344,277

Table A.5: Profit in million $ and run time when varying the number of scenarios and type of rolling horizon in the stochastic version of the MIP

Scenarios	Fixed RH		Decreasing RH	
	Profit	Run time	Profit	Run time
1	$64.0”	991	$60.0”	806
5	$90.6”	8,848	$96.2”	7,564
20	$75.2”	70,790	$84.3”	49,193
50	$55.6”	189,786	$86.3”	184,165
100	$31.2”	354,138	$79.3”	344,277

Table A.6: Profit in million $ when varying the number of binary days in the stochastic version of the MIP

Binary days	Profit	Run time
2	$92.1”	38,017
3	$86.5”	52,408
4	$90.1”	91,510
5	$91.2”	122,249
6	$90.3”	153,675

Table A.7: Profit in million $ and run time when varying gap value in the stochastic version of the MIP

Gap	Profit	Run time
0.5%	$87.6”	112,432
1.0%	$95.8”	54,215
2.0%	$89.9”	48,140
5.0%	$86.2”	45,412

Table A.8: Overall in- and out-of-sample values when varying the number of scenarios in the stochastic version of the MIP

Scenarios	In-sample value	Out-of-sample value
1	1	0.723
5	0.814	0.687
20	0.637	0.632
50	0.634	0.634
100	0.629	0.626

Table A.9: In-sample values for different situations when varying the number of scenarios in the stochastic version of the MIP

Scenarios	To ports	From ports	In ports	Open sea
1	1.000	1.000	1.000	1.000
5	0.832	0.807	0.823	0.787
20	0.650	0.550	0.825	0.650
50	0.634	0.513	0.841	0.645
100	0.607	0.498	0.860	0.602

Table A.10: Out-of-sample values for different situations when varying the number of scenarios in the stochastic version of the MIP

Scenarios	To ports	From ports	In ports	Open sea
1	0.741	0.734	0.591	0.746
5	0.717	0.635	0.777	0.728
20	0.648	0.535	0.829	0.658
50	0.642	0.514	0.843	0.634
100	0.604	0.494	0.862	0.593

Table A.11: Profit in million $ and run time when varying the number of scenarios and port inclusion in the stochastic version of the heuristic

Scenarios	Closer grid points		Equal and closer points	
	Profit	Run time	Profit	Run time
1	$88.1”	494	$86.3”	469
5	$94.9”	2,027	$83.5”	2,034
20	$103.0”	8,369	$86.3”	8,305
50	$102.6”	20,278	$82.4”	20,168
100	$104.6”	40,482	$76.5”	42,177

Table A.12: Profit and run time with decreasing horizon when varying the number of scenarios in the stochastic version of the heuristic

Scenarios	Profit	Run time
1	$88.1”	494
5	$94.9”	2,027
20	$103.0”	8,369
50	$102.6”	20,278
100	$104.6”	40,482

Table A.13: Profit in million $ and run time when varying the number of scenarios and type of rolling horizon in the stochastic version of the heuristic

Scenarios	Fixed RH		Decreasing RH	
	Profit	Run time	Profit	Run time
1	$91.8”	550	$88.1”	494
5	$86.3”	2,564	$94.9”	2,027
20	$96.8”	9,107	$103.0”	8,369
50	$95.4”	21,346	$102.6”	20,278
100	$96.5”	44,174	$104.6”	40,482

Table A.14: Overall in- and out-of-sample values when varying the number of scenarios in the stochastic version of the heuristic

Scenarios	Overall	
	In sample	Out of sample
1	1.000	0.769
5	0.816	0.781
20	0.834	0.826
50	0.828	0.829
100	0.831	0.831

Table A.15: In-sample values for different situations when varying the number of scenarios in the stochastic version of the heuristic

Scenarios	To ports	From ports	In ports	Open sea
1	1.000	1.000	1.000	1.000
5	0.941	0.907	0.561	0.942
20	0.940	0.907	0.561	0.975
50	0.952	0.922	0.532	0.963
100	0.969	0.933	0.540	0.970

Table A.16: Out-of-sample values for different situations when varying the number of scenarios in the stochastic version of the heuristic

Scenarios	To ports	From ports	In ports	Open sea
1	0.914	0.774	0.537	0.944
5	0.929	0.874	0.470	0.900
20	0.938	0.905	0.533	0.988
50	0.954	0.921	0.535	0.969
100	0.968	0.932	0.539	0.979

Table A.17: Comparison of profit and run time between the stochastic versions of the MIP and heuristic when varying the number of scenarios

Scenarios	MIP		Heuristic	
	Profit	Run time	Profit	Run time
1	\$60.0”	806	\$88.1”	494
5	\$96.2”	7,564	\$94.9”	2,027
20	\$84.3”	49,193	\$103.0”	8,369
50	\$86.3”	184,165	\$102.6”	20,278
100	\$79.3”	344,277	\$104.6”	40,482

Table A.18: Comparison of stability between the stochastic versions of the MIP and heuristic when varying the number of scenarios

Scenarios	In sample		Out of sample	
	MIP	Heuristic	MIP	Heuristic
1	1.00	1.00	0.723	0.769
5	0.814	0.816	0.687	0.781
20	0.637	0.834	0.632	0.826
50	0.634	0.828	0.634	0.829
100	0.629	0.831	0.633	0.831

Table A.19: Profit in million \$ and run time in the deterministic version of the MIP

Scenarios	Profit	Run time
1	\$60.0”	799
5	\$77.6”	965
20	\$70.1”	914
50	\$77.2”	907
100	\$78.1”	942

Table A.20: Profit in million \$ and run time in the deterministic version of the heuristic

Scenarios	Profit	Run time
1	\$88.1”	460
5	\$94.4”	463
20	\$92.6”	469
50	\$85.6”	493
100	\$90.6”	474

Table A.21: Comparison of profit and run time between the deterministic versions of the MIP and heuristic

Scenarios	MIP		Heuristic	
	Profit	Run time	Profit	Run time
1	$60.0”	799	$88.1”	460
5	$77.6”	965	$94.4”	463
20	$70.1”	914	$92.6”	469
50	$77.2”	907	$85.6”	493
100	$78.1”	942	$90.6”	474

Table A.22: Comparison of profit and run time between the stochastic and deterministic versions of the MIP

Scenarios	Profit		Run time	
	Stochastic	Deterministic	Stochastic	Determinstic
1	$60.0”	$60.0”	806	799
5	$96.2”	$77.6”	7,564	965
20	$84.3”	$70.0”	49,193	914
50	$86.3”	$77.2”	184,165	907
100	$79.3”	$78.1”	344,277	942

Table A.23: Comparison of profit and run time between the stochastic and deterministic versions of the heuristic

Scenarios	Profit		Run time	
	Stochastic	Deterministic	Stochastic	Deterministic
1	$88.1”	$88.1”	494	478
5	$94.9”	$94.4”	2,027	482
20	$103.0”	$92.6”	8,369	478
50	$102.6”	$85.6”	20,278	482
100	$104.6”	$90.6”	40,482	494

Appendix B

Approximate Conversion Table

The table in Figure B.1 is used for converting different units throughout the thesis.

	1 billion cubic metres NG	1 billion cubic feet NG	1 million tonnes oil equivalent	1 million tonnes LNG	1 trillion British thermal units	1 million barrels oil equivalent
1 billion cubic metres NG	1	35.3	0.9	0.740	35.7	6.6
1 billion cubic feet NG	0.028	1	0.025	0.021	1.01	0.19
1 million tonnes oil equivalent	1.11	39.2	1	0.820	39.7	7.33
1 million tonnes LNG	1.36	48	1.22	1	48.6	8.97
1 trillion British thermal units	0.028	0.99	0.025	0.021	1	0.18
1 million barrels oil equivalent	0.15	5.35	0.14	0.110	5.41	1

Figure B.1: Conversion table
(BP, 2013a)

Appendix C

Stand-alone Mathematical Model

Mathematical model

We begin by introducing the sets and indices of the model. Then we present the parameters and variables, before presenting the objective function and constraints.

Sets and indices

$\mathcal{G}$	- Grid points, g
$\mathcal{G}^{\mathcal{N}}(g)$	- Neighboring grid points of g, $\tilde{g}$
$\mathcal{G}^{\mathcal{S}}(g)$	- Straight neighboring grid points of g, $\tilde{g}$
$\mathcal{G}^{\mathcal{D}}(g)$	- Diagonal neighboring grid points of g, $\tilde{g}$
$\mathcal{G}^{\mathcal{L}}$	- Buy ports, g
$\mathcal{G}^{\mathcal{U}}$	- Sell ports, g
$\mathcal{T}$	- Time interval, t
$\mathcal{S}$	- Scenarios, s
$\mathcal{K}_t$	- Index set of scenario subsets at time t, k
Ω_{kt}	- Subset of scenarios at time t, ω

Parameters

Q	- Capacity of vessel
C^B	- Basis cost for the vessel (not moving)
C^{MS}	- Extra cost for the vessel if it is moving straight
C^{MD}	- Extra cost for the vessel if it is moving diagonal
P_{gts}	- Price at port g at time t in scenario s

Variables

x_{gts} - 1 if the vessel is at grid point g at time t in scenario s, 0 otherwise
m_{ts}^{S} - 1 if the vessel moves straight from time t to time $t+1$ in scenario s, 0 otherwise
m_{ts}^{D} - 1 if the vessel moves diagonally from time t to time $t+1$ in scenario s, 0 otherwise
b_{gts} - 1 if the vessel buys in grid point g at from time t to time $t+1$ in scenario s, 0 otherwise
s_{gts} - 1 if the vessel sells in grid point g from time t to time $t+1$ in scenario s, 0 otherwise
f_{ts} - 1 if the vessel is full from time t to time $t+1$ in scenario s, 0 otherwise
x_{gt} - Variables used as nonanticipativity constraints for position
f_{gt} - Variables used as nonanticipativity constraints for status

Objective function

maximize $\sum_{g\in\mathcal{G}}\sum_{t\in\mathcal{T}}\sum_{s\in\mathcal{S}}(P_{gts}\cdot s_{gts} - P_{gts}\cdot b_{gts})\cdot Q - \sum_{t\in\mathcal{T}}\sum_{s\in\mathcal{S}}(C^{B} + C^{MS}\cdot m_{ts}^{S} + C^{MD}\cdot m_{ts}^{D})$

Constraints

$$\sum_{g\in\mathcal{G}} x_{gts} = 1, \qquad t\in\mathcal{T}, s\in\mathcal{S} \tag{C.1}$$

$$x_{g(t+1)s} - \sum_{\tilde{g}\in\mathcal{G}^{\mathcal{N}}(g)} x_{\tilde{g}ts} \leq 0, \qquad g\in\mathcal{G}, t\in 1\ldots\mathcal{T}-1, s\in\mathcal{S} \tag{C.2}$$

$$m_{ts}^{D} - x_{gts} - \sum_{\tilde{g}\in\mathcal{G}^{\mathcal{D}}(g)} x_{\tilde{g}(t+1)s} \geq -1, \qquad t\in 1\ldots\mathcal{T}-1, s\in\mathcal{S} \tag{C.3}$$

$$m_{ts}^{S} - x_{gts} - \sum_{\tilde{g}\in\mathcal{G}^{\mathcal{S}}(g)} x_{\tilde{g}(t+1)s} \geq -1, \qquad t\in 1\ldots\mathcal{T}-1, s\in\mathcal{S} \tag{C.4}$$

$$\sum_{g\in\mathcal{G}^{\mathcal{U}}} s_{gts} - f_{ts} \leq 0, \qquad t\in\mathcal{T}, s\in\mathcal{S} \tag{C.5}$$

$$\sum_{g\in\mathcal{G}^{\mathcal{L}}} b_{gts} + f_{ts} \leq 1, \qquad t\in\mathcal{T}, s\in\mathcal{S} \tag{C.6}$$

$$f_{(t+1)s} - f_{ts} - \sum_{g\in\mathcal{G}^{\mathcal{L}}} b_{gts} + \sum_{g\in\mathcal{G}^{\mathcal{U}}} s_{gts} = 0, \qquad t = 1\ldots\mathcal{T}-1, s\in\mathcal{S} \tag{C.7}$$

$$m_{ts}^{S} + m_{ts}^{D} + \sum_{g\in\mathcal{G}^{\mathcal{L}}} b_{gts} + \sum_{g\in\mathcal{G}^{\mathcal{U}}} s_{gts} \leq 1, \qquad t\in\mathcal{T}, s\in\mathcal{S} \tag{C.8}$$

$$b_{gts} + s_{gts} - x_{gts} \leq 0, \qquad g \in \mathcal{G}^{\mathcal{L}} \cup \mathcal{G}^{\mathcal{U}}, t \in \mathcal{T}, s \in \mathcal{S} \tag{C.9}$$

$$x_{gts} - x_{g\omega} = 0, \qquad t \in \mathcal{T}, k = 1 \ldots \mathcal{K}_t, \omega \in \Omega_{k(t+1)} \tag{C.10}$$

$$f_{gt} - f_{g\omega} = 0, \qquad t \in \mathcal{T}, k = 1 \ldots \mathcal{K}_t, \omega \in \Omega_{k(t+1)} \tag{C.11}$$

$$b_{gts} \in \{0, 1\}, \qquad g \in \mathcal{G}^{\mathcal{L}}, t \in \mathcal{T}, s \in \mathcal{S} \tag{C.12}$$

$$s_{gts} \in \{0, 1\}, \qquad g \in \mathcal{G}^{\mathcal{U}}, t \in \mathcal{T}, s \in \mathcal{S} \tag{C.13}$$

$$x_{gts} \in \{0, 1\}, \qquad g \in \mathcal{G}, t \in \mathcal{T}, s \in \mathcal{S} \tag{C.14}$$

$$m^S_{ts}, m^D_{ts}, f_{ts} \in \{0, 1\}, \qquad t \in \mathcal{T}, s \in \mathcal{S} \tag{C.15}$$