
Electronic Thesis and Dissertation Repository

5-16-2018 11:00 AM

Optimal Trading of a Storable Commodity via Forward Markets

Behzad Ghafouri, *The University of Western Ontario*

Supervisor: Davison, Matthew, *The University of Western Ontario*

A thesis submitted in partial fulfillment of the requirements for the Doctor of Philosophy degree in Statistics and Actuarial Sciences

© Behzad Ghafouri 2018

Follow this and additional works at: <https://ir.lib.uwo.ca/etd>



Part of the [Operational Research Commons](#)

Recommended Citation

Ghafouri, Behzad, "Optimal Trading of a Storable Commodity via Forward Markets" (2018). *Electronic Thesis and Dissertation Repository*. 5394.

<https://ir.lib.uwo.ca/etd/5394>

This Dissertation/Thesis is brought to you for free and open access by Scholarship@Western. It has been accepted for inclusion in Electronic Thesis and Dissertation Repository by an authorized administrator of Scholarship@Western. For more information, please contact wlsadmin@uwo.ca.

Abstract

A commodity market participant trading via her inventory has access to both spot and forward markets. To liquidate her inventory, she can sell at the spot price, take a short forward position, or do a combination of both. A trade is proposed in which there is always a hedging forward contract, which can be considered a *dynamic* cash and carry arbitrage. The trader can adjust the maturity of the forward contract dynamically until the inventory is depleted or a time constraint is reached.

In the first setup, the storage contract (to carry the inventory) is assumed to have a constant cost and a flexible duration. The risk and return characteristics of an Approximate Dynamic Programming (ADP) and a Forward Dynamic Optimization solution are compared. The trade is contrasted with an optimal spot sale among other alternative liquidation strategies. Independent from the underlying stochastic forward price model, it is proved and verified numerically that a partial sale strategy is not optimal. The optimally selected forward maturities are limited to the subset comprising the immediate, next, and last timesteps.

Under a more realistic storage contract, which assumes a stochastic cost and a fixed duration, a new ADP approach is developed. The optimal policy shows the tanker rent decision is accompanied by a buy order since the loss from an empty tanker is more than the gain of renting it cheaply yet early. Given the nonadjustable duration of the rent contract, a longer contract generates a higher value by benefiting from a tanker refill option.

Keywords: Real Options, Cash and Carry Arbitrage, Oil Storage, Forward Trading, Markov Decision Process, Approximate Dynamic Programming, Least Squares Monte Carlo

Acknowledgments

My first and sincere appreciation goes to my advisor Prof. Matt Davison for all I have learned from him and for his continuous inspiration and support in all stages of my PhD. This thesis would not have been possible without him. I would also like to thank him for being open to ideas and for helping me to shape my own interest and ideas. His attitude to research inspired me to start this adventure, and he is the most understanding and committed advisor one can hope to have.

I gratefully acknowledge the Ontario Government and Western University for their financial assistance through Ontario Graduate Scholarships (OGS) and QEII Graduate Scholarships in Science and Technology (QEII-GSST). I also acknowledge the Department of Statistical and Actuarial Sciences, and the School of Graduate and Postdoctoral Studies for their financial support.

I wish to extend my sincere thanks to my proposal defence committee, Prof. Marcos Escobar-Anel, Prof. Rogemar Mamon and Prof. Lars Stentoft, and to my thesis examination committee, Prof. Walid Busaba, Prof. Marcos Escobar-Anel, Prof. David Stanford and Prof. Antony Ware, for their insightful feedback and constructive discussions.

I would like to thank Andrea Liguori for his invaluable help and interest during the initial stages of this research. I also wish to express my deep gratitude to Prof. Somayeh Moazeni for her advice and insight on dynamic programming and for all I learned from her.

I have furthermore to thank all teachers who taught me at different stages of my life specially at the Department of Statistical and Actuarial Sciences.

I am eternally grateful to my family for their unconditional love and support; Mitra, Abbas, Maryam, and Behnam.

Dedication

This dissertation is dedicated to Mitra, Abbas, Maryam, and Behnam who supported and motivated me at every step of the way. In particular, this dissertation is dedicated to my lovely wife; it was her patience, trust, and encouragement that made this journey possible.

Table of Contents

Abstract	i
Acknowledgments.....	ii
Dedication	iii
List of Tables	x
List of Figures	xii
List of Acronyms and Symbols	xx
Chapter 1	1
1 Introduction	2
1.1 An Introduction to Commodity Markets.....	2
1.2 Motivation of the Present Research	6
1.3 A Simple Two-Period Model	9
1.4 Existing Literature.....	12
1.4.1 Commodity Storage Valuation and Optimal Trading.....	12
1.4.2 Price Models	19
1.5 Structure of the Thesis.....	20
1.6 Summary	22
Chapter 2.....	25
2 The Price Model and Simulation	25

2.1	The Model	25
2.2	Parameter Estimates and Simulation.....	28
2.3	Dynamics of the Forward Curve	29
2.4	Summary	31
Chapter 3.....		33
3	Forward Dynamic Optimization.....	33
3.1	Model	33
3.2	Theoretical Interpretation and Comparison.....	38
3.3	Parameters	41
3.4	Computational Results	42
3.4.1	Decision-making analysis at the path level.....	43
3.4.2	Impact of the Initial Condition.....	51
3.4.3	Impact of the Number of Periods.....	52
3.4.4	Impact of the Storage Cost.....	54
3.5	Summary	57
Chapter 4.....		60
4	Optimal Solution with Dynamic Programming.....	60
4.1	Introduction	60
4.2	The Model	60

4.3	Theoretical Results	65
4.4	Algorithmic Solutions	67
4.4.1	Exact Dynamic Programming.....	67
4.4.2	Approximate Dynamic Programming (ADP)	69
4.4.3	Forward Dynamic Optimization (FDO).....	70
4.5	Parameters	71
4.6	Computational Results	72
4.6.1	The Optimal Value.....	72
4.6.2	The Optimal Policy	75
4.6.3	Comparison of ADP and FDO.....	79
4.6.4	Mapping the Decisions	82
4.7	Comparative Characteristics of the Present Trade	85
4.8	Sensitivity Analysis.....	90
4.9	Summary	93
Chapter 5	95
5	A Trading Model Considering Stochastic Storage Costs	95
5.1	Introduction	96
5.2	Oil Tanker Vessels Background.....	98
5.3	Literature on Tanker Freight Rates	101

5.4	Stochastic Storage Cost Modeling	105
5.5	Markov Decision Process (MDP) Model	111
5.6	Theoretical Proposition	117
5.7	Algorithmic Solution (ADP)	118
5.8	Parameters and Computational Setting	120
5.8.1	Benchmark Parameters	120
5.8.2	Basis Functions	122
5.8.3	Number of Paths	123
5.8.4	Inventory Discretization.....	124
5.9	Computational Results	124
5.9.1	Optimal Value	125
5.9.2	Optimal Policy	126
5.9.3	Impact of oil factors initial condition.....	134
5.9.4	Impact of Storage Cost Factors Initial Condition on the Initial Storage Cost	138
5.9.5	Impact of storage cost factors initial condition while ‘keeping the initial cost constant’	141
5.9.6	Impact of storage cost factors initial condition.....	146
5.9.7	Impact of the level of pumping costs	150
5.9.8	Impact of the time horizons	152

5.9.9	Sensitivity Analysis	156
5.10	Summary	160
Chapter 6	162
6	Conclusions and Future Research.....	162
6.1	Conclusions	162
6.2	Principal Contributions	166
6.3	Future Work	166
7	References	168
8	Appendix A.....	178
8.1	Proof of Lemma 4.I:	178
8.2	Proof of Lemma 4.II:	180
8.3	Proof of Proposition 4.I:	185
8.4	Assumption 4.I:	186
8.5	Proof of Proposition 4.II:	189
8.6	Proof of Proposition 5.I.....	197
9	Appendix B.....	202
9.1	Function Simulating Oil and Storage Cost Prices	202
9.2	Function Computing Oil Futures Prices.....	209
9.3	Function Computing the Endogenous State Space	209

9.4	Function Computing the Action Space	211
9.5	Function Representing the Transition Function	212
9.6	Function Formulating the Reward Function	213
9.7	Auxiliary Function – Array Index Mapping	214
9.8	Auxiliary Function – Polynomial Feature Mapping	214
9.9	Function Computing the Optimal Action Given Stochastic Storage Costs	216
9.10	Specify the Parameters and Simulate the Prices	218
9.11	Backward Induction	221
9.12	Out-of-Sample Performance	224
	Curriculum Vitae	233

List of Tables

Table 1.1 All possible decision paths at t_0 and t_1	10
Table 1.2 All possible decision paths expressed explicitly at t_0	11
Table 2.1. Parameters estimated by Hahn et al (2014)	28
Table 3.1. Problem parameters as specified by the base case for FDO analysis.	42
Table 4.1. Comparison of assumptions and MDP formulation between Chapter 3 and 4 problems.	65
Table 4.2. Pseudocode of the exact approach on a $H \times H$ grid of $\chi_{-\xi}$	68
Table 4.3. Pseudocode of ADP (LSM) approach.	70
Table 4.4. Problem parameters defined as Case A and Case B for the exact or ADP analysis....	72
Table 4.5. Optimal value and computational time of the FDO and ADP approaches using Case A of the parameters.	73
Table 4.6. Optimal value and computational time of exact and ADP approaches using Case B of the parameters.	74
Table 4.7. The elements of the MDP describing the covered call and protective put positions...	87
Table 4.8. Comparing the return and risk characteristics of different strategies. The parameters are based on Case B in Table 4.4 ($\Delta R = 1$).	88
Table 4.9. Unfavorable condition for the storage trade; downward-sloping initial forward curve induced by the stochastic factors initial condition set at $W_0 = (\chi_0, \xi_0) = (-0.2, 4.2)$, while the rest of the parameters are based on Case B in Table 4.4 ($\Delta R = 1$). There is not any maturity $T \in (0, \bar{T}]$ to set up a profitable short $F(0, T)$ contract, and thus the static cash and carry and FDO strategies opt out to sell at t_0	89

Table 5.1 Baltic Dirty Tanker Index (BDTI) composition as of November 2008. All vessel must have oil necessary approvals.....	100
Table 5.2. The Mirantes et al. (2012) model parameters estimated by Poblacion (2015) on the TD5 route using the TCE data. Note x^* , x^{**} , and x^{***} show the estimated values are significant at 10%, 5%, and 1%, respectively.	109
Table 5.3. Highlight of the differences between the model used in Chapter 4 and Chapter 5 (the present model).....	117
Table 5.4. Pseudocode for ADP (LSM) approach.	120
Table 5.5. Table of the general problem parameters establishing the benchmark setting.	121
Table 5.6. Impact of the degree of the polynomial basis functions generated using χ_t , ξ_t , χ'_t , ξ'_t , α_t , and α_t^* . The confidence intervals are computed by using 30 repetitions for each degree. The other parameters are as per Table 5.5.	122
Table 5.7. Confidence intervals (95%) of V and V^c resulted from the out-of-sample testing. The parameters are as per Table 5.5.....	125
Table 5.8. Percentage of the paths on which the decision is to rent the tanker at different timesteps. It is assumed the state is that the tanker has not been rented yet, i.e. $x_i = (I_i, R_i, T_i) = (0,0,0)$	128
Table 5.9. Impact of the problem time horizon, \bar{T} (years), and duration of the storage contract, T' (years). All the other parameters are set based on the benchmark values specified in Table 5.5.	153
Table 5.10. Impact of extending the main problem time horizon, \bar{T} , and the duration (length) of the storage contract, T' , on the paths with multiple fill ups.....	156
Table 8.1. Solution to inequalities based on Eq. 8.9.....	180

List of Figures

Fig. 1.1. (a) Historical CME NYMEX WTI Crude Oil Futures (CL#1 to CL#12) showing contango (upward) and backwardation (downward) states in the market at different points in time, (b) Slope of the forward curve at time t_i , Holding Cost (HC), and trade direction. 7

Fig. 1.2. Historical One-year Crude Oil Futures Spread and VLCC Time-Charter Rates; the spread data is based on CME NYMEX WTI contracts, while the VLCC data is from The Baltic and International Maritime Council (BIMCO, 2016) for Time-Charter Rates 1, and Charles Weber Company (“Charles R. Weber Company, Inc.”, 2016) for Time-Charter Rates 2 (per barrel per year). 9

Fig. 1.3 Decision tree of a simple two-period model; actions are only taken at times t_0 and t_1 . 10

Fig. 2.1. Two sample realizations (a and b) of the spot price $S(t)$ (dotted line) and forward curve with monthly increments. At each timestep (month) the one-year forward curve, $F(t, t \leq T \leq t + 1)$ is shown for $t = 0, 1/12, 2/12, \dots, 12/12$ using Table 2.1 parameters. 29

Fig. 2.2. Characteristics of the forward curve based on the parameters of Table 2.1, $T = 1$ year, $\Delta t = 0.25$, $M = 10,000$; (a) Slope versus mean price at $t = 0, 0.25, 0.5$, and 0.75 . (b) Change in the slope versus change in the mean price over each period. Period 1 represents the time from $t = 0$ to 0.25 , Period 2 represents the time from $t = 0.25$ to 0.5 , and Period 3 represents the time from $t = 0.5$ to 0.75 30

Fig. 3.1. New positions to capture any potential gains 36

Fig. 3.2. Histogram of Added Value (\$) over $M = 10K$ paths using base case parameters. 44

Fig. 3.3. A sample forward curve at $t = 0.25$ 44

Fig. 3.4. Payoff and selected contract at $t = 0.25$ as a function of slope at $t = 0$ and $t = 0.25$. The trader holds $F(0,1)$, or advances the maturity one, two, or three periods based on the realized slope at $t = 0.25$. All the parameters are per the base case as specified in Table 2.1 and Table 3.1.... 45

Fig. 3.5. Payoff and selected contract at $t = 0.5$ as a function of slope at $t = 0.25$ and slope at $t = 0.5$. Starting with an existing contract based on the slope at $t = 0.25$, the trader acts by choosing a new contract based on the realized slope at $t = 0.5$. All the parameters are per the base case as specified in Table 2.1 and Table 3.1..... 46

Fig. 3.6. Payoff and selected contract at $t = 0.75$ as a function of slope at $t = 0.5$ and slope at $t = 0.75$. Starting with an existing contract based on the slope at $t = 0.5$, the trader acts by choosing a new contract based on the realized slope at $t = 0.75$. All the parameters are according to the base case as specified in Table 2.1 and Table 3.1..... 48

Fig. 3.7. Payoff levels at (a) $t = 0.25$, (b) $t = 0.5$, and (c) $t = 0.75$ in terms of change in the slope and mean forward prices. Payoffs magnitude is divided into five different levels at each time, and for clarity, two plots are provided in parts (b) and (c). All the parameters are according to the base case as specified in Table 2.1 and Table 3.1..... 50

Fig. 3.8. (a) Added Value (%) as a function of the initial prices; transformation of different (χ_0, ξ_0) to the prices lead to $\$22.20 \leq S_0 \leq \134.29 and long-term price in the range $\$44.7-\181.27 . (b) Certain Added Value (\$) vs Uncertain Added Value (\$) colored by the initial slope; each (χ_0, ξ_0) implies an initial slope and a unique decomposition of Added Value (\$) into certain and uncertain parts. The solid and dashed black lines represent $X + Y = \text{constant}$ and $Y = X$ lines respectively. The parameters (other than the initial conditions) are per the base case as specified in Table 2.1 and Table 3.1..... 51

Fig. 3.9. Impact of N on risk-return characteristics and Sharpe ratio; (a) Added Value (\$) and its standard deviation as a function of N for three different c_H (b) Sharpe ratio as a function of N for three different c_H . All the parameters are according to the base case as specified in Table 2.1 and Table 3.1. 53

Fig. 3.10. Impact of storage cost (c_H) on; (a) Added Value (\$) colored by maturity of the initial contract. (b) Added Value (\$) versus its standard deviation for different values of c_H with labels,

where only certain c_H labels are shown for clarity. All the parameters (other than c_H and N) are per the base case as specified in Table 2.1 and Table 3.1..... 54

Fig. 3.11. Separate contribution of advancing or postponing trades to value, and the impact of Refund Ratio. (a) Uncertain Added Value (\$) generated by Advancing (A) or Postponing (P) trades for different values of Refund Ratio versus storage cost. (b) Added Value (\$) as a function of storage cost (c_H) for different values of Refund Ratio. All the parameters (other than c_H and Refund Ratio) are per the base case as specified in Table 2.1 and Table 3.1..... 56

Fig. 4.1. Histogram of value using the ADP method computed for 150 times. The solid line represents the mean of the histogram. Case A parameters ($\Delta R = 1$) as per Table 4.4 are used. . 75

Fig. 4.2. Optimal decision a_i^R (left) and a_i^T (right) at $t_i = 0.25$ ($i = 4$) and state $x_i = (R_i, T_i) = (1, 0.25)$. The parameters are per Case B in Table 4.4 ($M = 100K, N = 16, \Delta R = 1/3, H = 43$). Different $W_0 = (\chi_0, \xi_0)$ are used with the ADP; $W_0 = (-0.639, 4.637)$ per Case B (top), and $W_0 = (-1.2, 4.2)$ (bottom)..... 77

Fig. 4.3. Optimal decision a_i^R (left) and a_i^T (right) at $t_i = 0.75$ ($i = 12$) and state $x_i = (R_i, T_i) = (1, 0.75)$. The parameters are per Case B in Table 4.4 ($M = 100K, N = 16, \Delta R = 1/3, H = 43$). Different $W_0 = (\chi_0, \xi_0)$ are used with the ADP; $W_0 = (-0.639, 4.637)$ per Case B (top), and $W_0 = (-1.2, 4.2)$ (bottom)..... 78

Fig. 4.4. Results of simulating $M_2 = 100K$ new paths using Case A parameters and comparing ADP vs FDO statistics. 80

Fig. 4.5. Evolution of the spot price and performance comparison on a sample path, where $\$9.35 = V_0^{FDO} < V_0^{ADP} = \11.44 81

Fig. 4.6. Evolution of the spot price and performance on a sample path, where $\$7.36 = V_0^{ADP} < V_0^{FDO} = \9.84 81

Fig. 4.7. Optimal decision (inventory on the left, and maturity on the right) in terms of spot price (S_t) and slope ($(F(t, T) - S_t)/(T - t)$) at times 0.125, 0.375, 0.875, and 0.9375, corresponding to time stages $i = 2, 6, 14,$ and 15 . The red line represents $y = c_H$ line. 84

Fig. 4.8. Histogram of the values obtained on each path using different strategies. The histograms are based on the same set of 10K out-of-sample paths. The parameters are based on Case B in Table 4.4 ($\Delta R = 1$). 88

Fig. 4.9. % change in V_0 versus % change in different parameters; (a) Stochastic factors parameters, (b) time horizon T , storage cost c_H , and initial condition (χ_0, ξ_0) , (c) This panel completes panel ‘b’ by further extending the y-axis further, (d) Number of paths M , and time stages N 91

Fig. 4.10. Two sample realizations of the simulated spot price, $S(t)$, and forward curve, $F(t, T)$, with monthly increments base on parameters in Table 2.1 except $(\sigma_\chi, \sigma_\xi) = (0.70, 0.07)$ in (a), and $(\sigma_\chi, \sigma_\xi) = (0.07, 0.70)$ in (b). 92

Fig. 5.1. Approximate tanker routes; (i) TD3 from 1 (Ras Tanura, Saudi Arabia) to 2 (Chiba, Japan), (ii) TD5 from 3 (Bonny, Nigeria) to 4 (Philadelphia, USA), (iii) TD7 from 5 (Sullom Voe, UK) to 6 (Wilhelmshaven, Germany). 101

Fig. 5.2. Three sample realizations of $F'(t, T)$ (\$/tanker per day) using the parameters of Table 5.2 and initial condition $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*) = (0.3, 9, 0.3, 0.3)$ 110

Fig. 5.3. Maximum time passed since the initiation of a rental contract; (a) if $t_i \leq T'$: it is not possible to have an expired contract. (b) if $T' < t_i$: the contract is expired if it is initiated within the red interval, $[t_0, t_i - T')$ 112

Fig. 5.4. Impact of the number of sample paths (M) and out-of-sample paths (M_2) on V (panel a) and V^c (panel b). The confidence intervals are computed by using 30 repetitions at each (M, M_2) pair. The parameters are as per Table 5.5. 123

Fig. 5.5. Histogram of V (panel a) and V^c (panel b) formed by computing the values 150 times. Values are derived from three different methods; Out-of-sample, In-sample 1 (solve optimally forward in time by implementing the computed Θ), and In-sample 2 (approximate the value function at t_0 by solving Bellman equation). The solid lines represent the mean of the histograms. The parameters are as per Table 5.5. 126

Fig. 5.6. Optimal decision with respect to renting, a_i^I , at $t = t_1 = 1/12$ (a and b) and $t = t_4 = 4/12 = 1/3$ (c and d). State: at both t_1 and t_4 the tanker has not yet been rented, i.e. $x_i = (I_i, R_i, T_i) = (0,0,0)$ 127

Fig. 5.7. Optimal policy at $t = t_7 = 7/12$ with respect to the quantity bought/sold on the spot, a_i^R , is displayed on the left column (a, c, & e), and with respect to the maturity of the forward contract, a_i^T , on the right column (b, d, & f). 130

Fig. 5.8. Three sample paths and the trades executed on them based on the optimal policy. The one-year time-charter rate, $TC(t, t + 1)$, is overlaid with the forward curve, $F(t, T)$, on the left panels, and is overlaid with the forward-spot spread, $F(t, T) - F(t, t)$, on the right panels. The vertical lines specify the decision times to ‘rent’, ‘buy’, or ‘sell’. The forward and spread curves are plotted in colors associated to ‘buy’ (red) and ‘sell’ (blue) at the corresponding times. Tt is the state variable at time t indicating the maturity of the contract held (if any). r_t is the reward (\$) of the action taken at time t . Superscript “c” represents the respective variables under the constant storage cost assumption (V^c) rather than the stochastic storage cost (V). 133

Fig. 5.9. Impact of oil factors initial conditions on the value; V and V^c with respect to χ_0 and ξ_0 for $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*) = (3.39, 8.4, 0.3, 0.4)$, which are the benchmark values specified in Table 5.5, and represent an unfavorable initial condition ($\chi'_0 \gg 0$). 135

Fig. 5.10. Impact of oil factors initial conditions on the value; V and V^c with respect to χ_0 and ξ_0 for $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*) = (-3.39, 8.4, 0.3, 0.4)$, a favorably deviated initial condition ($\chi'_0 \ll 0$)... 137

Fig. 5.11. (a) Histogram of the time step at which the rent decision is made (if any at all) for $(\chi_0, \xi_0) = (-0.2, 4.8)$, unfavorable oil initial condition, and $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*) =$

(-3.39, 8.4, 0.3, 0.4), a favorable storage cost initial condition (b) Mean of the profits or losses made over the 25000 out-of-sample paths at each time step with parameters similar to part (a).
 138

Fig. 5.12. Percentage change in the initial storage cost $TC(0,1)$ as a univariate function of χ'_0 , ξ'_0 , α_0 , or α_0^* while keeping the other three fixed. The change is relative to the benchmark values specified in Table 5.5; $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*) = (3.39, 8.4, 0.3, 0.4)$. The red solid lines show the benchmark values..... 140

Fig. 5.13. Initial storage cost $TC(0,1)$ as a multivariable function of $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$; (a) highlighting changes with respect to χ'_0 and ξ'_0 , (b) highlighting changes with respect to α_0 and α_0^* 141

Fig. 5.14. V as a (univariate) function of (a) χ'_0 , (b) ξ'_0 , (c) α_0 , and (d) α_0^* . The black solid line represents the $V^c = \$5.53$ associated with the constant cost case. The rest of the parameters are set based on the benchmark values specified in Table 5.5. The initial conditions are selected such that their combination leads to a constant $TC(0,1) = \$6.57$ per barrel per year. 142

Fig. 5.15. V as a (bivariate) function of (a) $\chi'_0 - \xi'_0$, where ξ'_0 values are randomly jittered for better demonstration, and (b) $\chi'_0 - \alpha_0$. The rest of the parameters are set based on the benchmark values specified in Table 5.5. The initial conditions are selected such that their combination leads to a constant $TC(0,1) = \$6.57$ per barrel per year..... 143

Fig. 5.16. Comparing storage cost factors $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$ set to $(-3.63, 9.5, -1, 1)$ and $(3.91, 8, 0.5, -0.5)$ in the left and right columns respectively. (a)/(b): histogram of the time step at which the rent decision is made (if any at all), (c)/(d): mean of the profits or losses made over the 25000 out-of-sample paths at each timestep, (e)/(f): sample time-charter rate trajectories. The rest of the parameters are set based on the benchmark values specified in Table 5.5..... 145

Fig. 5.17. Random storage cost initial condition $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*) \in \{-4, \dots, 4\} \times \{8, \dots, 9.5\} \times \{-1, \dots, 1\} \times \{-1, \dots, 1\}$ under a favorable oil initial condition fixed at $(\chi_0, \xi_0) = (-0.6393, 4.6366)$ (benchmark). 147

Fig. 5.18. Random storage cost initial condition $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*) \in \{-4, \dots, 4\} \times \{8, \dots, 9.5\} \times \{-1, \dots, 1\} \times \{-1, \dots, 1\}$ under an unfavorable oil initial condition fixed at $(\chi_0, \xi_0) = (-0.3, 4.3)$.
..... 148

Fig. 5.19. V or V^c as a function of the one-year time-charter rate at t_0 , $TC(0,1)$. Different storage cost factors initial conditions are used to compute $TC(0,1)$, where $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*) \in \{-4, \dots, 4\} \times \{8, \dots, 9.5\} \times \{-1, \dots, 1\} \times \{-1, \dots, 1\}$ are the same as utilized in Fig. 5.17 and Fig. 5.18. Panel (a) shows a favorable oil initial condition; $(\chi_0, \xi_0) = (-0.639, 4.637)$, i.e. the benchmark values specified in Table 5.5, and Panel (b) demonstrates an ‘unfavorable’ oil initial condition $(\chi_0, \xi_0) = (-0.3, 4.3)$.
..... 149

Fig. 5.20. Impact of the pumping costs, c_p^+ (cost when buying) and c_p^- (cost when selling). All the other parameters are set based on the benchmark values specified in Table 5.5.
..... 151

Fig. 5.21. Performance of the algorithms at $(c_p^+, c_p^-) = (3, 3)$ leading to $V = \$4.13$ and $V^c = \$1.63$; (a) histogram of the time step at which the rent decision is made, (b): mean of the profits or losses made over the 25000 out-of-sample paths at each timestep. The rest of the parameters are set based on the benchmark values specified in Table 5.5. 152

Fig. 5.22. The term structure of the initial time-charter rate (panel a), and the same structure when it is multiplied by the rent duration (T) to represent the actual cost of storage in “\$ per barrel” (panel b). The basis for the parameters are Table 2.1 for the oil model, Table 5.2 for the storage cost model, and Table 5.5 for the general assumptions. 153

Fig. 5.23. Impact of extending the main problem time horizon, \bar{T} , and the duration (length) of the storage contract, T' . (\bar{T}, T') is increased from $(2, 1)$ (i.e. the benchmark values specified in Table 5.5) to $(3, 2)$, and the corresponding results are presented on the left and right columns respectively. Panels (a) and (b) show the histogram of the time step at which the rent decision is made (if any at all). Panels (c) and (d) illustrate mean of the profits or losses made over the 25000 out-of-sample paths at each time step. 155

Fig. 5.24. Sensitivity of V (solid lines) and V^c (dashed lines) to risk-premium and drift parameters. The basis for the parameters are Table 2.1 for the oil model, Table 5.2 for the storage cost model, and Table 5.5 for the general assumptions..... 157

Fig. 5.25. Sensitivity of V (solid lines) and V^c (dashed lines) to different parameters. The basis for the parameters are Table 2.1 for the oil model, Table 5.2 for the storage cost model, and Table 5.5 for the general assumptions. 158

Fig. 5.26. Sensitivity of V (solid lines) and V^c (dashed lines) to the oil factors initial conditions, χ_0 and ξ_0 . The parameters are based on Table 5.5, where accordingly the benchmark values are $(\chi_0, \xi_0) = (-0.639, 4.637)$. The impact of χ_0 or ξ_0 can be identified via the symbols specified in the legend. 159

Fig. 8.1. Testing the accuracy of estimating the spread via a linear approximation, i.e. with $a = t_i$ and $b = \bar{T} = 1$, using 10,000 sample paths over t_0, t_1, \dots, t_{N-1} timesteps, shown in different colors. All simulation parameters are based on Table 2.1 and Case B of Table 4.4. 188

Fig. 8.2. Simulated $e^{-r(T-t)}(F(t, T) - c_p)$ curves where $0 \leq t \leq 1$ and $t \leq T \leq 1$ year using Table 2.1 parameters, where temporal discretization is set to $N = 16$ time stages to match that of Case B in Table 4.4. R^2 and standard error of the regression is listed for each curve fitted with a line..... 188

List of Acronyms and Symbols

Acronyms

ADP	Approximate Dynamic Programming
FDO	Forward Dynamic Optimization
LSM	Least Squares Monte Carlo
MDP	Markov Decision Process

Main Symbols

i	Discrete time index ($i \in \{0, 1, 2, \dots, N\}$)
t	time
t_i	Discretized time at timestep i ($t_i = i\Delta t$)
T	Maturity time of a forward (oil or freight rate)
\bar{T}	Problem time horizon by which any inventory must be sold and delivered
T'	Fixed duration of the rent contract (charter period)
Δt	Timestep size for discretization
N	Number of periods of the problem time horizon ($N = T/\Delta t$)
N'	Number of periods of the rent contract ($N' = T'/\Delta t$)
\bar{R}	Maximum storage capacity
ΔR	Inventory discretization size (batch size)
L	Number of batches of the discretized capacity ($L = \bar{R}/\Delta R$)
S_t	Spot price of oil at time t
$F(t, T)$	Forward price of oil at time t with maturity at time T
S'_t	Spot price of the freight rate at time t
$F'(t, T)$	Forward price of the freight rate at time t with maturity at time T
$TC(t, T)$	Time Charter (TC) rate at time t for chartering the vessel from time t to T
V, V^c	Lower-bound estimate of the value function at time zero under the stochastic (V) and constant (V^c) storage cost assumptions
c_P, c_P^+, c_P^-	Pumping costs when buying (c_P^+), selling (c_P^-) inventory, or generic (c_P)
c_H	Holding cost including tanker rent and the cost operating it
δ	One timestep discount factor ($\delta = e^{-r\Delta t}$)
$\Theta(i, x_{i+1})$	Vector of the weights of basis functions in LSM
$\Phi(W_i)$	Vector of the basis functions of W_i in LSM
M, M_2	Number of sample (M) and out-of-sample (M_2) paths

MDP Symbols

I_i	Expiry time of the time-charter contract (tanker rental contract)
R_i	Inventory level

T_i	Forward contract maturity time
x_i	Endogenous state variable (vector); $x_i = (R_i, T_i)$ or (I_i, R_i, T_i)
W_i	Exogenous state variables; $W_i = (\chi_i, \xi_i)$ or $(\chi_i, \xi_i, \chi'_i, \xi'_i, \alpha_i, \alpha_i^*)$
a_i^I	Decision to the rent the tanker (1) or not (0)
a_i^R	Decision on oil quantity to be bought or sold on the spot market
a_i^T	Decision on the maturity of short forward position
a_i	Decision (action) variable; a vector comprising of a_i^T , a_i^R , or a_i^I
\mathcal{X}_i^I	Set of feasible rental contract expiry dates at timestep i ($I_i \in \mathcal{X}_i^I$)
\mathcal{X}_i	Set of feasible states (state space) at timestep i ($x_i \in \mathcal{X}_i$)
$\mathcal{A}_i(x_i)$	Set of feasible actions at timestep i ($a_i \in \mathcal{A}_i(x_i)$)
$V_i(x_i, W_i)$	Value function at timestep i given states (x_i, W_i)

Symbols of Oil Price Model

ξ_t, χ_t	Long-term (ξ_t) and short-term (χ_t) stochastic factors
ξ_0, χ_0	Initial conditions of the corresponding stochastic factors
σ_ξ, σ_χ	Factor volatilities
$\lambda_\xi, \lambda_\chi$	Factor risk premia
$\rho_{\xi\chi}$	Correlation of the Brownian motion increments in the SDEs of factors
μ_ξ	Drift of the long-term factor (ξ_t)
μ_ξ^*	Risk-neutral drift of the long-term factor (ξ_t) defined as $\mu_\xi^* = \mu_\xi - \lambda_\xi$
k	Speed of mean reversion of χ_t

Symbols of Tanker Price Model

$\xi'_t, \chi'_t, \alpha_t, \alpha_t^*$	Long-term (ξ'_t), short-term (χ'_t), and seasonal (α_t and α_t^*) factors
$\xi'_0, \chi'_0, \alpha_0, \alpha_0^*$	Initial conditions of the corresponding stochastic factors
$\sigma_{\xi'}, \sigma_{\chi'}, \sigma_\alpha$	Factor volatilities
$\lambda_{\xi'}, \lambda_{\chi'}, \lambda_\alpha, \lambda_{\alpha^*}$	Factor risk premia
$\rho_{\xi'\chi'}, \rho_{\xi'\alpha}, \rho_{\xi'\alpha^*},$	Correlations of the corresponding Brownian motion increments in the
$\rho_{\chi'\alpha}, \rho_{\chi'\alpha^*}$	SDEs of factors
$\mu_{\xi'}$	Long-term factor (ξ'_t) drift
k'	Speed of mean reversion of χ'_t
φ	Seasonal period

Chapter 1

This thesis introduces and investigates a novel trading strategy using forward contracts, which can be considered a dynamic extension of the cash and carry arbitrage. The standard cash and carry arbitrage is set up once, at which point the profit is known and to be derived from the gap between the forward-spot spread and the costs associated with carrying the inventory. The trader holds a short forward position and can add value by adjusting the maturity of the contract at each timestep. Compared to the storage valuation literature, the stochastic processes are under the physical (rather than risk-neutral) measure here, which makes trading in the forward market a potentially profitable choice. Another important distinction of the present research with the storage valuation literature is that the inventory must be hedged while it is carried through time. The constraint does not permit to speculate on the future (spot or forward) prices by carrying an inventory without securing a buyer through a (forward) contract since initiation. The constraint requires holding a short forward position at all times, the maturity of which is subject to optimization to maximize added value. Commodity owners liquidating an existing inventory or producers hedging an expected production quantity are faced with similar challenges as they utilize financial contracts (e.g. forwards) to reduce risks and increase profits (Bertocchi, Consigli, & Dempster, 2011; Fackler & Livingston, 2002).

This thesis contributes theoretically by proving, in the constant storage cost framework, that independent of the underlying stochastic forward price model, a ‘partial sale’ decision (dividing the inventory sold between the spot and forward markets) is not optimal. Also, it is proved (using certain linearity assumptions) that the optimally selected forward maturity is limited to three choices; the immediate time, the next timestep, or the last timestep. These theoretical results imply that the optimal action set is considerably smaller than the feasible set, incorporating which in the algorithms, allows one to compute the optimal policies much faster.

Different computational solutions are developed and compared; an Approximate Dynamic Programming (ADP) method based on the Least-Squares Monte Carlo (LSM) provides the optimal

policy, while a myopic solution is offered by the Forward Dynamic Optimization (FDO) approach. Despite the near optimal performance of FDO in the valuation literature, also known as the Rolling Intrinsic policy, it does not accomplish comparatively good results here. The underperformance of the FDO compared to the ADP method in terms of the expected profit depends on the market conditions as the FDO policy always takes less risk. Employing the ADP technique, the risk and return characteristics of the dynamic cash and carry trade are compared with alternative strategies to liquidate an inventory. While highly speculative methods such as ‘optimally selling on the spot price’ or ‘protective put’ achieve higher values, the present method demonstrates attractive risk attributes.

Finally, a more realistic framework is considered where both commodity prices and storage costs are stochastic. Determining the trade initiation time and, independently, inventory refill decisions are among the unique contributions of this new setting. In this new context, the numerical results support the previously proven theoretical propositions asserting the nonoptimality of partial sales and optimality of a small subset of the maturities. A new ADP approach is developed to solve this problem having six stochastic drivers (two for the oil and four for the storage cost prices). The novel algorithm exploits the structure of the one-time rent decision and the state variable evolution in continuation function approximation, where the state variable dictates which of the six stochastic factors are used in the LSM regression.

1 Introduction

This chapter provides an introduction to commodity markets in Section 1.1. It is followed by presenting the motivation of the proposed research in Section 1.2 and a simple two-period model illustrating the concept in Section 1.3. The related literature is reviewed Section 1.4. The structure of the thesis and the chapter summary are discussed in sections 1.5 and 1.6 respectively.

1.1 An Introduction to Commodity Markets

Commodity markets are very volatile. For example, during the 2008 financial crisis, in the period July 3, 2008 to December 3, 2008, the Brent crude oil price fell 69%. This will create opportunities as well as challenges for a wide range of market players. The significant economic impact of commodities in the world from production to consumption involving different players such as

hedgers, speculators and arbitrageurs, has been long recognized. Therefore, development and understanding of optimal channels, through which players can transact or manage their risk efficiently has been an ongoing effort.

One important aspect of the commodity markets is that they consist of a physical market and a financial market. In the physical, or ‘spot’ market, players buy or sell the physical commodity for immediate delivery. Although operational constraints often impose a minimum lag before delivery (Hull, 2014). On the other hand, in the financial market, the players take positions in derivative contracts developed based on the spot market. Some of these contracts are futures, forwards, and options, which are essential for risk management and efficient transactions. When the lag between the transaction time and the delivery time is larger than the minimal time mentioned above, the trade becomes a forward agreement. The relation between the physical market and financial market is very important to the participants in the two markets.

A forward contract is an agreement between a buyer and a seller at time t , according to which the seller must deliver, at a fixed future time T , an underlying asset, and the buyer must pay on that date an amount fixed at time t , shown by $F(t, T)$, which is called the time t forward price for date T . Also, let $S(t)$ denote the spot price at time t . If there exists a liquid market for the underlying asset, at maturity of the forward contract, T , it can be argued based on the no-arbitrage principle that $F(T, T) = S(T)$, since an inequality leads to an arbitrage opportunity, in which the cheaper side can be bought in one market and sold in the other market simultaneously. This condition allows the long forward position to realize a profit or loss equal to $F(T, T) - F(t, T) = S(T) - F(t, T)$ when she closes out her position at time T .

In case of storable commodities, the theory of storage tries to connect the spot and forward prices through explaining why a market player holds an inventory. This leads to the notion of convenience yield, which is the benefit deriving from owning the physical asset in the inventory as opposed to a long forward position (Kaldor, 1939; Working, 1948). This benefit is comparable to the dividend that is received by a stock owner but not by the holder of a contract on that stock. From a practical point of view, it is not difficult to see the benefits of having the physical asset at one’s disposal, which mitigates the risks of not having the asset available at the right time due to, for instance,

production disruptions. In the following, some of the main consequences of the theory of storage are listed (Geman, 2005):

- The higher the level of global inventories of a commodity, the lower the commodity price, and vice versa.
- The higher the level of global inventories of a commodity, the lower the volatility of commodity price, and vice versa.
- It can be concluded from the first two points that the price of a commodity and its volatility have a positive correlation, a feature known as inverse leverage effect. This is easy to verify in a condition where a commodity becomes scarce, when both price and the volatility increase. Interestingly, the relation between a stock price and its volatility is the opposite.
- Volatility of $F(t, T)$ usually decreases with T , ceteris paribus, a feature known as the Samuelson effect. This is due to the higher sensitivity of short-term maturities to the arrival of news.
- The convenience yield derived from holding an inventory depends (stochastically) on the inventory levels. Usually when inventory levels are low, the convenience yield is high.

Using no-arbitrage principle arguments¹, it can be easily shown that the relation between the spot price at time t , $S(t)$, and the forward price, $F(t, T)$ for a storable commodity is as in Eq. 1.1.

$$F(t, T) = S(t)e^{(r+c-y_1)(T-t)} = S(t)e^{(r-y)(T-t)} \quad \text{Eq. 1.1}$$

Here r is the continuously compounded interest rate, c is the continuously compounded storage cost, y_1 is the continuously compounded convenience yield, and y is the continuously compounded ‘net’ convenience yield. The underlying assumption is that r and y are constant in the interval $[t, T]$. Also, the fact that the commodity under consideration is storable allows one to assume such a y_1 exists. To highlight the similarities between this case and the dividend-paying stock, $F(t, T)$ is the cost of buying the asset at time t and carrying it to time T , where $e^{(r+c-y_1)(T-t)}$ represent the financing (principle and interest) and storage cost net of the accrued benefit, be it dividends or convenience yield. Briefly speaking, the no-arbitrage principle arguments lead to taking advantage

¹ No-arbitrage arguments result in inequalities in this setting, where the ‘convenience yield’ term is added to allow to generate an equation.

of any mispricing by a ‘cash-and-carry’ arbitrage if $F(t, T) > S(t)e^{(r-y)(T-t)}$, and by a ‘reverse-cash-and-carry’ arbitrage if $F(t, T) < S(t)e^{(r-y)(T-t)}$. It is worth mentioning that when it is not possible to carry a commodity through the interval $[t, T]$, i.e. when dealing with a non-storable commodity such as electricity, the above no-arbitrage arguments no longer work!

The relationship shown between $S(t)$ and $F(t, T)$ is important; it establishes the relation between the spot price and a portfolio of forwards maturing at different T s. In other words, it gives the ‘forward curve’. This is also helpful if there is not a liquid spot market, in which $S(t)$ is known immediately. Since, if there exist two different (liquid) maturities, say $F(t, T_1)$ and $F(t, T_2)$, one can infer the values of $S(t)$ and y .

If $r + c < y_1$, which means an environment with low interest rate, low storage cost, and high benefit of holding the physical commodity, the forward curve will be downward sloping. This condition is known as *backwardation*. For instance, in the case of crude oil, the forward curve has often historically been in backwardation, due to the belief among participants, who value the physical commodity highly in the face of insufficient oil supplies. On the contrary, if $r + c > y_1$, which means an environment where there is no or little value of holding the physical commodity, or with sufficiently large interest rate or storage cost, the forward curve will be upward sloping. This condition is known as *contango*. In the case of oil, usually a bearish sentiment, where there is a weak demand for the spot cargos, is accompanied with an upward-sloping curve (“The Forward Curve for Oil Prices Suddenly Looks Awful for OPEC,” 2017). Therefore, it can be concluded that for storable commodities, the shape of the forward curve, i.e. upward- vs downward-sloping, is directly tied to the magnitude of convenience yield. Forecasting how the shape, or convenience yield, will change is crucial for the market players; it can cause them to switch their positions between long- and short-end maturities.

The futures contracts are similar forward contracts fundamentally, however, there are some differences. They are standardized in terms of maturity, quantity, and quality of the underlying commodity. Futures trade on an exchange to which participants post ‘margin’, greatly reducing any counterparty credit risk. For more details of the way in which margin functions see Hull (2014). Despite these differences, it can be shown that in the absence of stochastic interest rates

and credit risk, the futures and forward on the same underlying commodity and similar maturity are equal, an assumption that will be adapted throughout the present study.

So far, the relation between the spot price, $S(t)$, and the forward price, $F(t, T)$, was explored. However, there is the question of the relation between the expected value of the future spot price, $E[S(T)|\mathcal{F}_t]$, and the forward price, $F(t, T)$. In $E[S(T)|\mathcal{F}_t]$, the expectation is conditional on the information available at time t , \mathcal{F}_t , and under the physical measure. In other words, does forward price help to predict future spot price? The *Rational Expectations Hypothesis* of Keynes and Lucas (Sargent, 1986), postulates that the forward price is a non-biased predictor of the future spot price, i.e. $E[S(T)|\mathcal{F}_t] = F(t, T)$. However, empirical studies showed that the equality does not often hold (at least under the real or physical measure \mathbb{P}). When $E[S(T)|\mathcal{F}_t] < F(t, T)$, it can be contributed to risk-averse players who are willing to pay more to secure a delivery at time T instead of buying it on the spot at time T . When $F(t, T) < E[S(T)|\mathcal{F}_t]$, it may be contributed by a belief by players that there is an oversupply of the commodity to be delivered at time T . The latter case is supported by the *theory of normal backwardation*, which postulates Eq. 1.2.

$$F(t, T) < S(t) < E[S(T)|\mathcal{F}_t] \quad \text{Eq. 1.2}$$

The theory argues that the difference $E[S(T)|\mathcal{F}_t] - F(t, T)$ is a *risk-premium*, which the producers pay to lock in a price at time t , i.e. to hedge, and the speculators earn due to taking the risk. Because this theory assumes that the hedgers are net short as a group, it does not hold in all markets. Ultimately, the existence of a normal backwardation, and thus the sign of the risk-premium, rest on the particular commodity and its inventory levels (Geman, 2005). Another view is to embed the risk-premium in the probability measure, and then compute the conditional expectation of the future spot price, which yields the famous risk-neutral probability measure \mathbb{Q} and the forward price as in Eq. 1.3.

$$F(t, T) = E^{\mathbb{Q}}[S(T)|\mathcal{F}_t] \quad \text{Eq. 1.3}$$

1.2 Motivation of the Present Research

We are motivated by studying and optimizing the off-shore oil storage trade observed in contango markets, i.e. upward-sloping forward curve. In the crude oil context, this trade is also known as

Contango and Carry Trade (Diaz-Rainey et al., 2017), where crude oil is bought cheap on the spot, sold using a forward contract at a higher price, and stored in a tanker until delivery. The trade can be profitable if the gap between the forward and spot price is higher than the costs associated with the storage. Depending on the assumptions made, the problem can take slightly various forms, however, we are mainly interested in understanding, expanding, and optimizing the trading strategies. At each timestep, the trader has the option to sell the oil on the spot or adjust the maturity of her short position, or do both on partial quantities. Undertaking these actions successively until the oil is sold or a deadline is reached generates a sequence of cash flows, the sum of which will define the total profit from this trade.

The storage trade was prevalent during the super contango of late 2008 to early 2009 when oil prices hit a low point. For instance, on Feb 12, 2009, there was a very steep 12-month contango between March 2009 (\$33.98/barrel) and March 2010 (\$55.95/barrel) futures contracts. Some recent historical futures curves based on WTI crude oil contracts are shown in Fig. 1.1.a. This


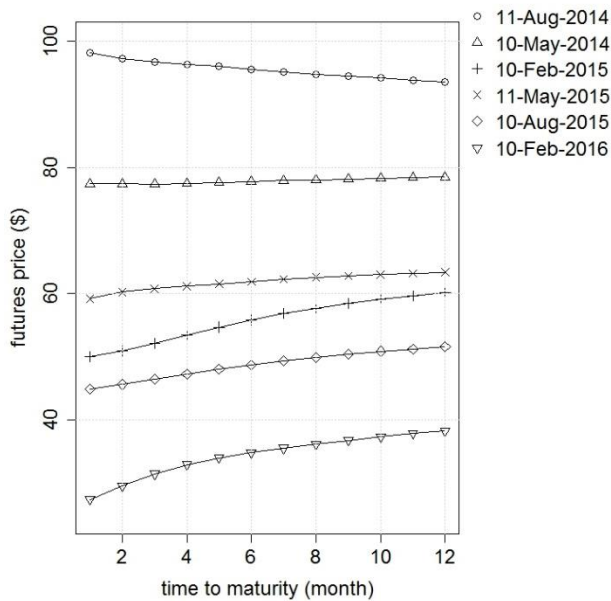


Fig. 1.1. (a) Historical CME NYMEX WTI Crude Oil Futures (CL#1 to CL#12) showing contango (upward) and backwardation (downward) states in the market at different points in time, (b) Slope of the forward curve at time t_i , Holding Cost (HC), and trade direction.

figure highlights various shapes in the one-year futures curve at different points in time, as well as the recent shift towards an upward sloping curve. Fig. 1.1.b shows how the storage trade and the slope of the forward curve are related. The slope of the forward curve, $(F(t_i, t_2) - F(t_i, t_1))/(t_2 - t_1)$, is essentially a measure that captures the change in value between two delivery points in time. Like storage cost, this slope has a unit of dollar per year per barrel, and is closely related to the storage cost. If the slope is greater than the storage cost, there is incentive to long the front-end (buy oil) and short the far-end (contract to sell forward). Indeed, the direction of the trade is reversed if the cost is higher than the slope. So, the spread between, rather than the absolute value of, the forward prices plays a critical role in this argument, where for simplicity, the time value of money is ignored.

Fig. 1.2 illustrates the historical one-year futures spread $F(t, T_2) - F(t, T_1)$ (or the slope since $T_2 - T_1 = 1$) based on WTI crude oil contracts, where the near end of the spread is the front-month contract. The figure shows a positive spread (upward sloping curve) in the period late 2008 to early 2009. When the storage trade is profitable, most of the cheaper onshore storage capacities are exhausted, and traders resort to “floating storage” using offshore oil tankers. Very Large Crude Carriers (VLCC) with capacity of around 2 million barrels are popular for this purpose. The tanker shipping market is very volatile as tanker demand is heavily influenced by the market dynamics of the oil and related products among other factors (Alizadeh et al., 2015). Fig. 1.2 also shows the historical VLCC one-year Time-Charter rates (BIMCO, 2016; “Charles R. Weber Company, Inc.”, 2016) which are annualized on a per barrel basis. The sharp drop in the VLCC charter prices in the late-2008 to 2009 period, combined with large futures spreads, triggered the wide use of offshore storage trades in this period. The slope has stayed positive since the beginning of 2015, which perhaps reflects the new era of abundant crude oil production. Tanker rates declined from \$9.31 to \$5.47 per barrel in the first half of 2016, another positive signal for the storage trade.

Commodity owners or producers intending to liquidate an existing inventory optimally are faced with a similar challenge (Fackler and Livingston, 2002). The standard static cash and carry arbitrage suggests a *set-and-forget* strategy, which means that no further changes are made to the position after the trade is initiated. In this research, the static cash and carry strategy is expanded into a dynamic approach where the trader updates the maturity

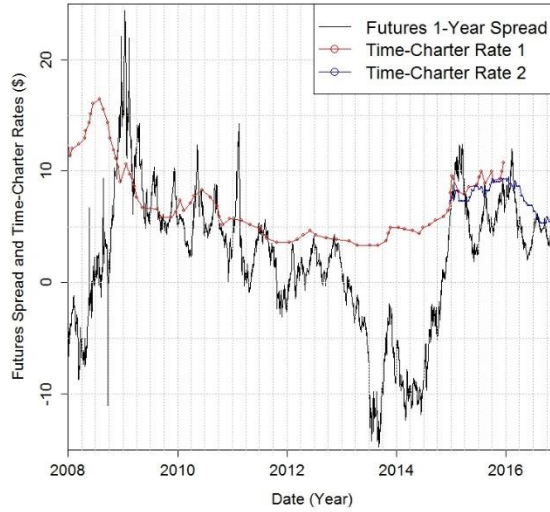


Fig. 1.2. Historical One-year Crude Oil Futures Spread and VLCC Time-Charter Rates; the spread data is based on CME NYMEX WTI contracts, while the VLCC data is from The Baltic and International Maritime Council (BIMCO, 2016) for Time-Charter Rates 1, and Charles Weber Company (“Charles R. Weber Company, Inc.”, 2016) for Time-Charter Rates 2 (per barrel per year).

of the short contract periodically until she delivers the oil. Starting from a full inventory, this can be viewed as a liquidation strategy. Subsequently, the impact of additional flexibilities is also studied, which include a partial sale on the spot and forward markets, optimal buying time rather than starting with a full inventory from the beginning, and an inventory refill option.

1.3 A Simple Two-Period Model

Fig. 1.3 depicts a very simple model of the dynamic cash and carry trading with only two periods. Let R_i and T_i respectively denote the inventory level and the maturity of the short forward position at t_i . Also, a_i denote the action (maturity) chosen at t_i . At t_i , selecting $a_i = t_i$ is equivalent to an immediate sale on the spot market. For simplicity assume the interest rate is zero. Starting from $R_0 = 1$ (full inventory) and $T_0 = 0$ (no initial contract), there are five feasible decision paths as listed in Table 1.1. Here, c_P and c_H denote the pumping cost and the holding cost (per unit of time) respectively.

Starting from the last timestep, there is no decision to be made at t_2 since any inventory must be always sold. Moving backward in time, at t_1 , if the optimization over action a_1 is expressed in




Fig. 1.3 Decision tree of a simple two-period model; actions are only taken at times t_0 and t_1

terms of maximizing the rewards of the corresponding possibilities, one can combine the sub-branch groups, as listed in Table 1.2. Now define $\Lambda(t_i, t_1, t_2) = F(t_i, t_1) - F(t_i, t_2) + c_H(t_2 - t_1)$; this is equivalent to the payoff (spread) that is generated if the trader decides at t_i to switch her contract maturing at t_2 with a contract maturing at t_1 . Similarly, $-\Lambda(t_i, t_1, t_2)$ is the payoff from the reverse action. Finally, moving backward to t_0 yields Eq. 1.4 for V_0 , the optimal value at t_0 . Here, $(x)^+ = \max\{x, 0\}$, and $\mathbb{E}_0[\cdot] = \mathbb{E}[\cdot | \mathcal{F}_{t_0}]$ denotes the expectation under the physical measure with respect to \mathcal{F}_{t_0} , the filtration generated by the stochastic price processes.

Notes	Decisions path	Total reward
Sell the inventory at t_0	①: $a_0 = t_0$	$F(t_0, t_0) - c_P$
Short $F(t_0, t_1)$ at t_0 , and keep the contract at t_1	②: $a_0 = t_1$ and $a_1 = t_1$	$F(t_0, t_1) - c_P - c_H \Delta t$
Short $F(t_0, t_1)$ at t_0 , and postpone the sale by choosing the longer maturity at t_1	③: $a_0 = t_1$ and $a_1 = t_2$	$F(t_0, t_1) - F(t_1, t_1) + F(t_1, t_2) - c_P - 2c_H \Delta t$
Short $F(t_0, t_2)$ at t_0 , and sell immediately on the spot at t_1	④: $a_0 = t_2$ and $a_1 = t_1$	$F(t_0, t_2) - F(t_1, t_2) + F(t_1, t_1) - c_P - c_H \Delta t$
Short $F(t_0, t_2)$ at t_0 , and keep the contract at t_1	⑤: $a_0 = t_2$ and $a_1 = t_2$	$F(t_0, t_2) - c_P - 2c_H \Delta t$

Table 1.1 All possible decision paths at t_0 and t_1 .

Notes	Decisions path	Total reward
Sell the inventory at t_0	①: $a_0 = t_0$	$F(t_0, t_0) - c_P$
Short $F(t_0, t_1)$ at t_0 , and decide whether to keep the contract at t_1 or postpone the sale by choosing the longer maturity at t_1	② & ③: $a_0 = t_1$	$F(t_0, t_1) - c_P - c_H\Delta t$ $+ \max\{-F(t_1, t_1) + F(t_1, t_2) - c_H\Delta t, 0\}$
Short $F(t_0, t_2)$ at t_0 , and decide whether to sell immediately on the spot at t_1 or to keep the contract at t_1	④ & ⑤: $a_0 = t_2$	$F(t_0, t_2) - c_P - 2c_H\Delta t$ $+ \max\{F(t_1, t_1) - F(t_1, t_2) + c_H\Delta t, 0\}$

Table 1.2 All possible decision paths expressed explicitly at t_0 .

Eq. 1.4 compares selling the inventory on the spot market at t_0 with the forward market using either of the two available maturities (t_1 and t_2), with an option to readjust the maturity later.

$$V_0 = F(t_0, t_0) - c_P + \max\{0, -\Lambda(t_0, t_0, t_1) + \mathbb{E}_0 \left[(-\Lambda(t_1, t_1, t_2))^+ \right], -\Lambda(t_0, t_0, t_2) + \mathbb{E}_0 \left[(\Lambda(t_1, t_1, t_2))^+ \right]\} \quad \text{Eq. 1.4}$$

It suggests that the forward positions could potentially result in a higher value than an initial spot sale at $F(t_0, t_0) - c_P$, unless the maximization is always trivially solved to zero. The value $\mathbb{E}_0 \left[(\Lambda(t_1, t_1, t_2))^+ \right]$ is like a call option on the underlying asset price $\Lambda(t, t_1, t_2)$. It provides an option to advance the forward maturity from t_2 to t_1 . Similarly, $\mathbb{E}_0 \left[(-\Lambda(t_1, t_1, t_2))^+ \right]$ is like a put option with the same underlying asset price and provides an option to postpone the forward maturity from t_1 to t_2 . Both options have a strike price of zero. From another perspective, given the form expressed in Eq. 1.5, one may interpret the right-hand side as an exchange option on

$$\mathbb{E}_0 \left[(\Lambda(t_1, t_1, t_2))^+ \right] = c_H\Delta t + \mathbb{E}_0[\max\{F(t_1, t_1) - F(t_1, t_2), -c_H\Delta t\}] \quad \text{Eq. 1.5}$$

$F(t_1, t_1)$ and $F(t_1, t_2)$ with a non-zero strike price of $-c_H\Delta t$. No closed form solution is available for this option (unless $c_H = 0$) although some approximations do exist (Bjerksund and Stensland 2014; Kirk 1995; Margrabe 1978) in the European case of just two periods. When more than two

periods are present, the analogous option starts to combine challenges of moving exercise boundaries and even the approximations fail. To compare the relative performance of the two choices, *suppose* the second and third arguments of the maximum in Eq. 1.4 were equal. After simplification, it would lead to $\mathbb{E}_0[\Lambda(t_1, t_1, t_2)] = \Lambda(t_0, t_1, t_2)$; it would hold if $\Lambda(t, t_1, t_2)$ process is a martingale process. However, it is known that $\Lambda(t, t_1, t_2)$, as a linear combination of forward prices, is a martingale under the risk-neutral measure and not necessarily under the physical measure. Thus, the attempt to show that the last two arguments of the maximum are equal fails. The problem does not seem to have a trivial solution, and Eq. 1.4 should be solved by evaluating the maximum. In this thesis, the above trade is modeled in a general multiperiod setting, and computational solutions are presented.

1.4 Existing Literature

In this study, we extend the static storage trade into a dynamic one by allowing subsequent trading, and investigate the important underlying factors. The problems in the literature relevant to the current research are mainly concerned with the storage asset valuation, where the value of a physical asset is estimated with monetizing the operational flexibilities by making optimal trading decisions. In the following section, the differences and similarities between the existing literature and the present problem will be highlighted. Although the focus of this study is optimal decision making, the decision algorithm requires us to characterize the dynamics of the oil forward term structure. Therefore, the models used to simulate oil prices are reviewed briefly as well.

1.4.1 Commodity Storage Valuation and Optimal Trading

The shape of the futures curves contains important information about the relation between spot and forward prices and the corresponding trading opportunities in the financial markets. At the same time, there is a market for the underlying physical commodity, and these two markets are connected via storage. Particularly, since the recent oil price decline starting in mid-2014, oil inventories have attracted more attention because they reflect fundamental supply-demand factors, which can be exploited in discovering arbitrage opportunities (Ye and Karali, 2016). In this regard, the decision to sell the oil at the spot price, or to store it and use futures contract to deliver at a later point in time is a simple yet important question, especially for the participants in the supply

chain of oil and oil products as well as traders. When the futures curve is in backwardation (downward-sloping), producers may hold the commodity in storage for different reasons; convenience yield as suggested by the classical theory of storage (Kaldor, 1939; Working, 1948), price uncertainty in the framework of a real option to extract oil from a reserve (Litzenberger and Rabinowitz, 1995), or both (Considine and Larson, 2001), where the producer sells the oil by either extracting from the reserve or by pumping from above ground inventory. On the other hand, when the crude oil futures curve is in contango (upward-sloping), it may present a profitable trading opportunity (Ghafouri and Davison, 2017; Jafarizadeh and Bratvold, 2013). Traditionally, the trading strategy involves buying oil on the spot, storing it, and taking a short position in a longer-term futures contract. This means the oil is “bought low” today and a contract to “sell it high” is made to guarantee a later delivery. Of course, the oil must be stored which is not free. But if storage, pumping, and other related costs are outweighed by the difference between the spot and forward prices, this strategy can be profitable.

There is well established evidence in the commodity storage literature about the relation between forward curve shape and storage. By the no-arbitrage principle, as mentioned earlier, the forward price is as given by the following equation.

$$F(0, T) = S_0 \exp[(r + c - y_1)T] \quad \text{Eq. 1.6}$$

Here, S_0 is the spot price, r is the interest rate, c is the storage cost, and y_1 is the convenience yield. In the case of crude oil, Geman and Ohana (2009) verified that the slope of the forward curve can be a good proxy for inventory levels. Using a dataset of inventories and futures prices in the US, they documented the relationship among interest-adjusted spread and inventory levels. They defined interest-adjusted spread between spot and forward prices as $[F(0, T) - S_0 \exp(rT)]/S_0$. When inventory levels are high, convenience yield (y_1) is low because any change in demand can be absorbed through inventories. Thus, $r + c - y_1 > 0$, and the forward curve is upward sloping. Under these conditions, Geman and Ohana (2009) showed that interest-adjusted spread is high, which inspires a cash-and-carry strategy, i.e. long the spot and short the forward. The spot-forward relation has been studied by investigating the cointegration between spot and forward prices. It is found that conclusions about the efficiency of crude oil market is influenced

by the existence of structural breaks (Chen et al., 2014; Chinn et al., 2005; Maslyuk and Smyth, 2009). For instance, Chen et al. (2014) found evidence ‘against’ market efficiency characterized by the no-arbitrage rule for subsamples 1986-2012 and 1986-2004 within WTI data.

The question is how to exploit the potential arbitrage opportunities, beyond the simple cash and carry approach. The storage trade may be regarded as a strategy to monetize the operational flexibilities of storage assets by making optimal decisions and so is best viewed through the real options lens. Similarly, the storage asset *risk-neutral valuation* is based on computing the cash flows resulting from an optimal trading strategy under certain operational constraints, where the expected value of the discounted cash flows under the risk-neutral measure will be the asset value. However, in the present problem, the attempt is made to capture a positive expected profit by optimal trading under the physical measure.

In the context of gas storage valuation (Eydeland and Wolyniec, 2003; Gray and Khandelwal, 2004a, 2004b; Maragos, 2002), there are two heuristic approaches popular among the practitioners; (i) basket of spread options valuation, and (ii) rolling intrinsic valuation, also known as Forward Dynamic Optimization (FDO). *Intrinsic* value is the value that can be secured today by hedging a forward position using the storage facility, which means finding the optimal plan of injections and withdrawals based on the present forward curve. On the other hand, *extrinsic* value is derived from the flexibility to readjust the position in response to forward curve realizations in the future. Each of the two heuristic methods capture the intrinsic and part of the extrinsic value. Re-optimizing at each time step as new information becomes available gives rise to the names ‘rolling’ or ‘dynamic’ in both methods (Lai et al., 2010).

The rolling basket of spread options approach estimates the extrinsic value by the value of a portfolio of calendar-spread options, where the two legs of the spread represent the injection and withdrawal times (Secomandi, 2016). It involves linear programming and spread option valuation. The spread option is on the difference between $F(t_i, t_i)$ and $F(t_i, t_j)$, where t_i and t_j ($t_i < t_j$) represent the injection and withdrawal times respectively. In the case of storage valuation, the strike price, K , will be the marginal costs of injection and withdrawal operations, whereas if the storage facility is rental, one must add the storage cost to those as well. Let $SP_0^{i,j}$ denote the value

of the (spread) option at time t_0 to buy one unit of the commodity at time t_i , and carry it until time t_j to fulfill the short forward position $F(t_i, t_j)$ entered at time t_i .

$$SP_0^{i,j} = e^{-ri\Delta t} E[\max\{e^{-r(j-i)\Delta t} F(t_i, t_j) - F(t_i, t_i) - K, 0\} | \mathcal{F}_0] \quad \text{Eq. 1.7}$$

Here Δt is the discretization time step, and \mathcal{F}_0 is the information available at t_0 . Many such options are possible, formed by considering the combination of all injection times t_i , and withdrawal times t_j . As soon as all option values are computed, the next step is to maximize the value of the portfolio, which is constructed by all the options, by choosing the optimal position quantity in each option subject to the inventory and operational constraints formulated as a linear program (Lai et al., 2010). As mentioned earlier, it is also possible to repeat the optimization at each time step to take advantage of the new information that becomes available with the passage of time. Lai et al (2010) reported that this method provides suboptimal results. This is not surprising because at each time the policy is limited to a series of spread trades, and is devised only based on a deterministic optimization of the intertemporal optionality. Therefore, it does not benefit from the full inherent flexibility. Secomandi (2015) reports that in some cases the rolling basket of spread options significantly underperforms the FDO method. Since this method requires computation of all the spread option values, while it is still suboptimal, only the FDO method among the heuristic approaches is considered for implementation in the present thesis.

Forward Dynamic Optimization (FDO) captures all the intrinsic value as well as part of the extrinsic value. Capturing the full extrinsic value involves a trade-off between risk and reward, where a higher value comes with the risk of larger variations. In comparison to other valuation method, FDO is very intuitive and is favored by practitioners (Secomandi, 2015; Secomandi et al., 2015; Ware, 2013). This is due to transparency, ease of communication with management, and consistency with the methods used by companies to monetize the flexibility of their storage assets.

In the context of crude oil storage trade, to profit from an upward sloping forward curve, Jafarizadeh and Bratvold (2013) explored the Forward Maximization strategy, a static trade where the trader buys oil and simultaneously sells it using the profit-maximizing forward contract. They also attempted to study the dynamic version of this trade, which should lead to the FDO strategy

introduced before; the trader tries to profit from favorable shifts in the forward curve by readjusting her (net) short position at subsequent timesteps. However, there are some issues with their approach, which will be revisited later in this thesis.

In the FDO algorithm, the subsequent adjustments are risk-free, as trades happen only if profitable. Therefore, the trader following FDO is exposed to a very limited downside risk. Between inception and delivery, the portfolio stays neutral by including equal long and short positions in the commodity inventory and futures contract respectively. Although FDO has a low downside risk, it is a myopic strategy which fails to capture the future consequences of present actions.

To avoid the sub-optimality of FDO or basket of spread options, a forward-looking optimal trading algorithm considering the expected impact of the current decisions on the generated value should be used. The literature employing optimal strategies can be divided into two categories based on the price models used; spot vs forward prices. Although the spot models neglect the dynamics of the forward curve, which is indeed very important for the present research, the inherent simplicity (low-dimensionality) has led to their frequent and fruitful use (Bjerksund et al., 2011; Boogert and de Jong, 2008, 2011; Carmona and Ludkovski, 2010; Chen and Forsyth, 2007; Felix and Weber, 2012; Secomandi, 2010; Thompson et al., 2009). Although spot models manage to maintain a low-dimensional optimization problem, the reliability of basing trading decision solely on the spot prices is questionable (Lai et al., 2010). Another reason to use the forward market over the spot market is that transaction costs (bid-ask spread) are lower in the former. Taking advantage of this, Secomandi and Kekre (2014) studied the optimal policy for natural gas procurement to meet a random demand at a fixed future time, T , using the forward market. They assume no storage possibilities. Thus, the procurement can be partially done through either shorting a forward contract at time zero for delivery at time T , or waiting and buying on the spot at time T . In Secomandi and Kekre (2014), while energy procurement is done on the forward market, the forward maturity is fixed at a constant time T . If the forward contract maturity can be chosen by the decision maker, the time gap, between entering the contract and the physical delivery, leads to complications and excessive state space dimensions in the dynamic programming framework, which may easily yield an intractable problem.

Lai et al (2010) considered a high-dimensional forward curve for the natural gas prices using the multivariate Black model, and developed an ADP approach by reducing the problem to a low-dimensional MDP avoiding the intractability of the full information stochastic dynamic program. They derived lower and upper bounds for the storage value to assess heuristic-based approaches such as FDO and basket of spread options. On a set of realistic instances, they concluded that the lower bounds resulting from both heuristics and the ADP approach (reoptimized versions) are all nearly optimal when compared to the upper bound (in the sense of information relaxation and duality approach developed by Brown et al (2010)). In comparison among the three, the FDO approach was found to provide the best compromise between optimality and computational expense. Also, they (artificially) removed the seasonality from the NYMEX natural gas forward curves used, to test their findings in the case of other commodities such as crude oil, which does not exhibit noticeable seasonality. They found that the result does not change structurally, which validates the application of the methods for non-seasonal commodities.

To reduce the problem dimension, Lai et al (2010) made some assumptions about the information and value function approximation, e.g. the forward curve object $F(t_i, t_j; i \leq j)$ is reduced to the spot price, $F(t_i, t_i)$, and next timestep forward price, $F(t_i, t_{i+1})$. Their reward function determining the decision payoff at each time stage is based on the spot price only. In the same formulation, the information contained in the whole forward curve to compute the conditional expectation of the next stage value function is used. However, this proves unwieldy and so, after some simplifications, this is reduced to being conditional only on $F(t_i, t_{i+1})$. Also, Lai et al. (2010) did not consider trading on the forward market since, as Nadarajah et al. (2015) stated, it does not add value if the expectation is under the risk-neutral measure and the reward function is linear with respect to the transacted price. Similarly in other studies, considering term structure of the forward prices is merely for informational purposes in making (spot market trade) decisions, and trading on the forward market is precluded (Nadarajah et al., 2015, 2017). In the absence of forward trading, the one-dimensional action is summarized by $a \in \mathbb{R}$ representing withdraw-and-sell ($a > 0$), buy-and-inject ($a < 0$), or do-nothing ($a = 0$). In the current study, not using a risk-neutral measure prompts trading on the forward market leading to a two-dimensional action corresponding to the spot and forward markets. We allow taking positions for future delivery, which is not

immediately reflected in the inventory level. Thus, the action set involves choosing the contract maturity T , and the reward function depends on $F(t, T)$.

Similar to Lai et al (2010), Nadarajah et al (2015) observe that ADP methods deliver performances that are not much better than FDO (for further theoretical support see Secomandi et al (2015)). Löhndorf and Wozabal (2017) attribute this to “low-dimensional relaxations of the otherwise high-dimensional stochastic optimization problem”. Avoiding these relaxations, the current research sheds light on the relative performance of suboptimal and optimal methods, which may depend on the measure; suboptimal methods may provide near-optimal results under the risk-neutral measure (Lai et al., 2010) and far from it under the historical one (Löhndorf and Wozabal, 2017). Lohndorf and Wozabal (2017) studied natural gas storage contract pricing in an incomplete market setting under a risk measure based on the indifference pricing method. They found that the optimal indifference pricing algorithm results in profits 40% higher than that of FDO (also known as rolling intrinsic valuation), which further highlights the importance of studying optimal approaches. Another feature of the present problem is that the trader is always “flat” oil (no exposure to oil prices by balancing the short contracts with the long ones and the inventory). In other words, the inventory is carried while hedged until the delivery is concluded.

A real option involving multiple decisions by the option holder, such as optimal operation of a storage asset, leads to a potentially intractable MDP due to the curse of dimensionality by (i) the high dimensions of the *exogenous* part of the state space, and (ii) issues associated with estimating expectations of the value function with respect to the exogenous part of the state in the next stage (Powell, 2011, sec. 4.1). The *endogenous* part of the state refers to the variables defining the functional state of the system (or option), which are characterized by the inherent flexibility of the system. These are impacted by operational decisions made and are usually low dimensional. On the other hand, the exogenous part of the state is unaffected by the decisions made, and usually involve a term-structure, e.g. a yield or forward curve, explained by a high dimensional stochastic process. The transition function determines evolution of the endogenous state, whereas the exogenous state evolves according to a stochastic process independent both of the endogenous parts of the states and of the decision that is made. To solve the present problem, an ADP technique based on the Least Square Monte Carlo (LSM) approach (Carriere, 1996; Longstaff and Schwartz,

2001; Tsitsiklis and Roy, 2001) will be used, which (perhaps) allows to overcome the curse of dimensionality when the MDP is expressed as a SDP (Stochastic Dynamic Program). This method offers a cost-effective approach to compute the lower bound of the MDP value, whereas computing the upper bound is more involved and can be computationally expensive (Glasserman, 2003, section 8.7; Nadarajah et al., 2017).

1.4.2 Price Models

To simulate prices, the simplest type of the stochastic models are stochastic differential equations with a single random driver: the so-called one-factor models. Unfortunately, it has been shown that the one-factor models are not able to completely capture the oil price dynamics, and therefore multiple factor models have been introduced (Hilliard and Reis, 1998; Schwartz, 1997). Especially important in the present research, one drawback of one-factor models is that is that they capture neither changes in the forward curve's slope and curvature nor the difference in volatility across various forward maturities.

Hahn et al (2014) studied different approaches to model and forecast oil prices while focusing on stochastic process models. They investigate the suitability of these models after they show that oil prices may exhibit a stationary or non-stationary behavior depending on the time period. They cited this as evidence supporting the use of a two-factor model, where one factor is mean-reverting (stationary) and the other factor is non-stationary. Bhattacharya (1978); Dixit and Pindyck (1994), and Dias (2005) reviewed different mean-reverting processes with the goal of modeling oil prices. One may classify two-factor models into two categories; the first class is essentially a Geometric Brownian Motion (GBM) with a mean-reverting process nested within the drift term of the GBM to model the convenience yield (Gibson and Schwartz, 1990; Ribeiro and Hodges, 2004; Schwartz, 1997). Pindyck (1999) proposes that a more realistic model could include a mean-reverting process added to a stochastically evolving trend line. This suggestion forms the basis for the second class of the two-factor models, where the long-term component is modeled with a GBM process, and is combined with the short-term deviations modeled with a mean-reverting process (Schwartz and Smith, 2000). While Schwartz and Smith (2000) show that the short-term/long-term model can be equivalent to the two-factor model based on the stochastic convenience yield developed in Gibson

and Schwartz (1990), the former is computationally easier and provides a more direct way to separate the two factors.

There are other complex models with more than two factors or with jump processes (Casassus and Collin-Dufresne, 2005; Schwartz, 1997), however, two factor models are the most common ones from a capability versus practicality point of view. For instance, Schwartz (1997) suggests a three-factor model, in which he adds a third factor to the two-factor model of Gibson and Schwartz (1990). The third factor assumes the interest rate follows a Vasicek mean-reverting process (Vasicek, 1977). However, Carmona and Ludkovski (2004) show the addition of the third factor does not improve the model performance greatly because convenience yield volatility is much larger than that of the interest rate.

One consideration for the present research is the computational burden associated with additional numbers of stochastic drivers. The required solution techniques rely on Dynamic Programming (DP), where additional stochastic drivers increase the state space, and the associated computational expense. Although a lower number of stochastic factors is preferred, the developed solution can accommodate (at least in theory) extra stochastic drivers if they are required from a price modeling perspective. Additionally, the simplicity of the two-factor model allows the optimal policy boundary to be nicely presented in the two-dimensional space of the model's two stochastic factors. Finally, as will be discussed in Chapter 5, additional stochastic factors will be incorporated in the problem via introduction of a stochastic storage cost, which is another reason why using a low-dimensional oil forward curve is advantageous. Considering all the above arguments, Schwartz and Smith (2000) short-term/long-term model will be used in this thesis to model the crude oil prices, for which the details will be provided later.

1.5 Structure of the Thesis

In Chapter 2, The Price Model and Simulation, Schwartz and Smith (2000) two-factor model is reviewed in detail. Derivation of futures prices, simulation of prices, as well as the corresponding parameter estimates are presented.

Chapter 3, Forward Dynamic Optimization, is devoted to the sub-optimal yet simple FDO method. The performance and properties of this approach are studied in detail within the most basic problem setting. Here, the trader optimally selects the forward maturity, a_i^T , at each time t_i while from a quantity perspective, the decision is applied to the whole inventory.

Chapter 4, Optimal Solution with Dynamic Programming, focuses on finding the optimal policy within the same problem setting as Chapter 3 with the additional option to sell the inventory on the spot and forward markets partially. However, it is shown theoretically that this partial sale is not optimal, and the structure of the optimal policy is established analytically. The decision variables which the trader selects optimally at each time t_i are the quantity of oil to be sold on the spot market, a_i^R , and the forward maturity, a_i^T . In this chapter, the main approach to solve the problem optimally is ADP, where an exact Dynamic Programming (backward induction) method is also implemented to verify the presented ADP method. Finally, the problem is solved by FDO and the results are compared. The significance of forward curve slope in decision-making is highlighted by illustrating optimal policy with respect to the slope and spot price.

It should be noted that the solution by the exact Dynamic Programming (DP) method used in this thesis is not and should not be mistaken with the *true* (theoretically exact) solution of the presented optimization problems. The distinction between the exact DP and the ADP method is that the former is only based on nested simulations and does not make any assumptions about the structure of the continuation functions (as opposed to the latter). In this sense, the exact DP is also an approximation but with fewer assumptions. In summary, both the exact DP and ADP methods only provide an approximation to the true optimal policy and a low-biased estimate to the true optimal value

In Chapter 5, A Trading Model Considering Stochastic Storage Costs, the problem setting is expanded completely to reflect a more realistic framework. Most notable expansions consist of optimally initiating the trade, inclusion of buy decisions (rather than a sale-only liquidation setup), a storage refill option, and a stochastic storage cost. The decision variables which the trader selects optimally at each time t_i are the quantity of oil to be sold on the spot market (a_i^R), the forward maturity (a_i^T), and the binary decision (a_i^I) of whether or not to initiate the storage rent at t_i . Based

on the structure of the new MDP, the ADP method developed in Chapter 4 is modified to accommodate the new setting. More specifically, the basis functions used in Continuation Function Approximation are shown to be dependent upon the endogenous state (more specifically, whether the storage has been rented or not yet). To highlight the effect the stochastic cost, comparisons with a similar case in which the storage cost is not stochastic are also made.

1.6 Summary

While most of the literature are concerned with risk-neutral pricing of gas storage contracts, the crude oil storage trade has not received much attention. The reason is the former enjoyed widespread use in practice, while in the latter case, the trade has seemed to attract the opportunists since the 2009 financial crisis. One may attribute this to the often downward-sloping shape of the oil forward curve, which is not suitable for the subject trade of this research. However, the fraction of the times when the forward curve is upward-sloping has increased in the last decade, possibly due to structural changes in the market. For instance, the spread between the first and the seventh WTI contracts during the period from 2005 to 2014 indicates a contango condition more than 70 percent of the time, which is contrast to their tendency to be in backwardation in the previous period more than 70 percent of the time (Kemp, 2016).

Fundamentally, demand for natural gas, and consequently the prices, are very seasonal due to weather conditions, while the gas supply is nearly constant. This motivates developing natural gas storage facilities for the underlying engineering and economic reasons. In the case of crude oil, both consumption and production exhibit much less seasonality. Thus, a dedicated oil storage facility would only be built either for strategic reasons, e.g. United States Strategic Petroleum Reserve, or logistics. The economic drivers for storing crude oil apply only a small fraction of the times. Subsequently, an oil tanker provides a way to store crude oil temporarily. We examine the proposed trade first by considering oil tanker rents to be fixed and later, in the final chapter, to be expressed by a stochastic model.

Another feature distinguishing the present research from the existing literature is studying the performance of a trading strategy under the historical measure as opposed to the pricing under the risk-neutral measure. As highlighted in the literature review, the respective performance of different methods, i.e. suboptimal vs optimal, may deviate from what is known and established under the risk-neutral measure. More importantly, it prompts the use of forward contracts in trading in addition to the trades on spot market.

Most of the research in the literature studies natural gas trading on the spot market. Few studies included a simplified version or a complete term structure of the forward prices in the problem, however, it was only for the informational purposes (nested in the conditional expectation of the value function). Nevertheless, the forward maturity is not considered as a decision variable, because the actions are in the form of withdraw-and-sell, buy-and-inject, or do-nothing. There is only one contribution (Löhndorf and Wozabal, 2017) which models the situation in which a trader takes positions and selects the quantity of natural gas in each of the available maturities in the forward market. However, they are concerned with a specific incomplete market setting and pricing framework, which leads to an intractable problem for ADP solutions and requires employing complex techniques (Löhndorf and Wozabal, 2017).

Furthermore, another distinction of the present research with the storage pricing literature is that we are concerned with optimal trading of an asset under a critical constraint. That is the inventory should be carried while hedged until the delivery is concluded. This is a risk-averting assumption rooted in how the cash and carry arbitrage works. Constantly holding a hedging forward position is what leads to the periodic optimization of the contract maturity. The joint optimization of the inventory and a financial contract specification (e.g. maturity) is the salient character of this research.

We contribute to the literature by studying optimal sale (purchase and sale in Chapter 5) of a storable commodity under some risk-reducing assumptions. It also investigates and compares suboptimal but simple approaches. In doing so, it considers the full term structure of the forward curve, both from an informational as well as operational (trading) perspective, where forward maturities can be selected as decision variables. Although the very popular long-term/short-term

price model of Schwartz and Smith (2000) is used in this research, the MDPs developed here are independent from any particular price model. Similarly, the optimization algorithms can easily accommodate other price models. Finally, the introduction of a stochastic storage cost allows to progress toward a more realistic and comprehensive framework, in which most aspects of the trade are incorporated in the MDP as decision variables to be selected optimally by the trader.

Chapter 2

2 The Price Model and Simulation

In Section 2.1, the price model used throughout the thesis is introduced in detail. Section 2.2 reviews the parameter estimates of the model and the simulation process. By simulating the prices, Section 2.3 provides an overview of the dynamics of forward prices and the environment in which the subject trades occur. Section 2.4 offers a summary of the chapter.

2.1 The Model

Throughout this thesis, the Schwartz and Smith (2000) model is used to simulate the evolution of both oil spot prices and the oil forward curve through time. The Schwartz-Smith model has two stochastic factors or state variables. The first state variable χ_t is the short-term factor representing short-term deviations from a long-term trend, which is represented by the second state variable ξ_t , i.e. the long-term factor. The spot price is given in terms of these state variables by Eq. 2.1. Note that we retain, except where minor changes cause no confusion, the original model notation as we review the model.

$$\ln(S_t) = \chi_t + \xi_t \quad \text{Eq. 2.1}$$

Under the physical or \mathbb{P} measure, the two factors are governed by the SDEs Eq. 2.2 and Eq. 2.3.

$$d\chi_t = -k\chi_t dt + \sigma_\chi dW_\chi(t) \quad \text{Eq. 2.2}$$

$$d\xi_t = \mu_\xi dt + \sigma_\xi dW_\xi(t) \quad \text{Eq. 2.3}$$

Where dW_χ and dW_ξ are two correlated processes with correlation coefficient $\rho_{\chi\xi}$, as expressed in Eq. 2.4.

$$dW_\chi dW_\xi = \rho_{\chi\xi} dt \quad \text{Eq. 2.4}$$

Eq. 2.2 describes a (mean reverting) Ornstein-Uhlenbeck process with a mean-reversion level of zero. Eq. 2.3 describes an arithmetic Brownian motion with drift. To compute the forward prices

from this spot price it is, as discussed in Chapter 1, essential to operate under the risk neutral measure \mathbb{Q} . Under this measure the two processes have the following representation.

$$d\chi_t = (-k\chi_t - \lambda_\chi)dt + \sigma_\chi d\tilde{W}_\chi(t) \quad \text{Eq. 2.5}$$

$$d\xi_t = (\mu_\xi - \lambda_\xi)dt + \sigma_\xi d\tilde{W}_\xi(t) \quad \text{Eq. 2.6}$$

$$d\tilde{W}_\chi d\tilde{W}_\xi = \rho_{\chi\xi} dt \quad \text{Eq. 2.7}$$

Here, λ_χ and λ_ξ denote respectively the short- and the long-term market price of risk, or risk premia. Thus, under the risk-neutral measure, the short-term process χ_t reverts to $-\lambda_\chi/k$. Similarly, the drift of the long-term process ξ_t becomes $\mu_\xi^* = \mu_\xi - \lambda_\xi$. It can be shown that given initial values (ξ_0, χ_0) under the risk-neutral measure \mathbb{Q} , (ξ_t, χ_t) follow a bivariate normal distribution with the following mean and covariance structure.

$$E^{\mathbb{Q}}[\xi_t, \chi_t] = \left[\xi_0 + \mu_\xi^* t, e^{-kt} \chi_0 - \frac{\lambda_\chi}{k} (1 - e^{-kt}) \right] \quad \text{Eq. 2.8}$$

$$Cov^{\mathbb{Q}}(\xi_t, \chi_t) = Cov(\xi_t, \chi_t) = \begin{bmatrix} \sigma_\xi^2 t & \frac{\sigma_\chi \sigma_\xi \rho_{\chi\xi}}{k} (1 - e^{-kt}) \\ \frac{\sigma_\chi \sigma_\xi \rho_{\chi\xi}}{k} (1 - e^{-kt}) & \frac{\sigma_\chi^2}{2k} (1 - e^{-2kt}) \end{bmatrix} \quad \text{Eq. 2.9}$$

Therefore,

$$E[\text{Log}(S_t)] = e^{-kt} \chi_0 + \xi_0 + \mu_\xi t \quad \text{Eq. 2.10}$$

$$E^{\mathbb{Q}}[\text{Log}(S_t)] = e^{-kt} \chi_0 + \xi_0 - \frac{\lambda_\chi}{k} (1 - e^{-kt}) + \mu_\xi^* t \quad \text{Eq. 2.11}$$

In the long term ($t \rightarrow \infty$), Eq. 2.10 and Eq. 2.11 representing the two expectations can be considered and compared as two lines in which the slopes (μ_ξ and μ_ξ^*) differ by λ_ξ and the intercepts (ξ_0 and $\xi_0 - \lambda_\chi/k$) differ by λ_χ/k . The risk premia reduce the expected log of the spot price as expressed in Eq. 2.12.

$$E[\text{Log}(S_t)] - E^{\mathbb{Q}}[\text{Log}(S_t)] = \frac{\lambda_\chi}{k} (1 - e^{-kt}) + \lambda_\xi t \quad \text{Eq. 2.12}$$

We can now derive the forward price as the expected value of the spot price under the risk neutral measure by $F(t, T) = E_t^{\mathbb{Q}}[S_T]$. Let $F(t, T)$ denote the forward price at time t for delivery at time T (for the relationship between forward and futures prices please see Chapter 1). Based on the Schwartz-Smith model, $F(t, T)$ can be written as in Eq. 2.13, where ‘ $T - t$ ’ is the *time-to-maturity* of this contract, and $A(T - t)$, expressed in Eq. 2.14, is a deterministic term depending only on the time-to-maturity.

$$F(t, T) = \exp\left[e^{-k(T-t)} \chi_t + \xi_t + A(T - t) \right] \quad \text{Eq. 2.13}$$

$$\begin{aligned} A(T - t) = & (\mu_{\xi} - \lambda_{\xi})(T - t) - \frac{\lambda_{\chi}}{k} (1 - e^{-k(T-t)}) \\ & + \frac{1}{2} \left[\sigma_{\xi}^2 (T - t) + \frac{\sigma_{\chi}^2}{2k} (1 - e^{-2k(T-t)}) \right. \\ & \left. + 2 \frac{\sigma_{\chi} \sigma_{\xi} \rho_{\chi\xi}}{k} (1 - e^{-k(T-t)}) \right] \end{aligned} \quad \text{Eq. 2.14}$$

To examine the forward curve in terms of being upward- or downward-sloping, let us compute the derivative of the forward price with respect to the time-to-maturity, denoted by $z := T - t$, as expressed in Eq. 2.15. If time-to-maturity is very large ($z \rightarrow \infty$), the sign of the derivative matches the sign of $\mu_{\xi}^* + \frac{1}{2} \sigma_{\xi}^2$, which is independent from time t .

$$\begin{aligned} \frac{\partial F(t, t+z)}{\partial z} = & \exp\left[e^{-kz} \chi_t + \xi_t + A(z) \right] \left[\mu_{\xi}^* + \frac{1}{2} \sigma_{\xi}^2 \right. \\ & \left. + (\sigma_{\chi} \sigma_{\xi} \rho_{\chi\xi} - \lambda_{\chi} - k \chi_t) e^{-kz} + \frac{1}{2} \sigma_{\chi}^2 e^{-2kz} \right] \end{aligned} \quad \text{Eq. 2.15}$$

However, for small time-to-maturities z , the sign depends on the long-term as well as the short-term factor parameters, and it is stochastic due to the presence of χ_t . This is why the forward curve can (and does) switch from backwardation (downward) to contango (upward) stochastically. At a fixed time t , the sign can also change as time-to-maturity z varies.

As mentioned earlier, based on the risk-neutral pricing framework, computing $F(t, T)$ is done under the risk-neutral measure. However, the goal of the present research is to evaluate the

performance of an optimal trading strategy and the extent of its profit or loss, which falls in the risk management realm, and thus requires simulating state variables under the physical (historical) measure.

Luckily, the forward curve $F(t, T)$ is fully explained by the above closed-form solution in terms of the only two state variables (χ_t, ξ_t) . When the price has a closed form solution, one can simulate the underlying risk drivers, (χ_t, ξ_t) , under the physical measure, and feed them into $F(t, T)$ to compute future exposures (Schoftner, 2008). In contrast to the present research, simulation under the physical measure has not been considered in previous studies aiming at evaluation of trading strategies such as the suboptimal FDO (Jafarizadeh and Bratvold, 2013), or storage asset valuation (Lai et al., 2010).

2.2 Parameter Estimates and Simulation

In order to simulate the prices, we use the model parameters estimated by Hahn et al (2014) as listed in Table 2.1. The parameters used by Jafarizadeh and Bratvold (2013) were not based on actual estimation of the model and only represented a hypothetical assumption of a favorable condition for this trade. Hahn et al (2014) estimated the parameters in the Schwartz-Smith two-factor model employing the Kalman filter (Kalman, 1960) method using 1990–2013 WTI futures data. This data included several major developments in crude oil markets, compared to Schwartz and Smith (2000) that only included 1990–1996 data. Using the same period data, Hahn et al (2014) found that the results are in good overall agreement with those of Schwartz and Smith (2000).

σ_χ	0.3116	k	1.0880	λ_χ	0.3733	μ_ξ	0.0818
σ_ξ	0.2053	$\rho_{\chi\xi}$	0.0823	λ_ξ	0.1070	$\mu_\xi^* = \mu_\xi - \lambda_\xi$	-0.0252

Table 2.1. Parameters estimated by Hahn et al (2014)

Accordingly, the following discretized version of the SDE's under the physical measure, i.e. Eq. 2.2 to Eq. 2.4, are used to simulate the prices. These discretized equations are the exact forms that can be obtained by integrating the corresponding SDE's (if the exponential functions are replaced by their corresponding first order Taylor expansion, these equations will become equivalent to the Euler approximation of the SDE's).

$$\chi_t = \chi_{t-1}e^{-k\Delta t} + \sigma_\chi \sqrt{\frac{1 - e^{-2k\Delta t}}{2k}} Z_\chi \quad \text{Eq. 2.16}$$

$$\xi_t = \xi_{t-1} + \mu_\xi \Delta t + \sigma_\xi \sqrt{\Delta t} Z_\xi \quad \text{Eq. 2.17}$$

$$Z_\xi \sim N[0,1], \quad Z_\chi \sim N[0,1], \quad \text{Corr}(Z_\chi, Z_\xi) = \rho_{\chi\xi} \quad \text{Eq. 2.18}$$

2.3 Dynamics of the Forward Curve

To provide an overview of the environment in which the trades occur and the optimization algorithms operate, the dynamics of the forward curve is studied. Let a $T = 1$ year time horizon be divided by timesteps of Δt into $N = T/\Delta t$ periods. This N determines the number of timesteps at which trading decisions are made. Although N might not be very large, e.g. $N = 12$ to represent monthly trading, the number of underlying periods used in the simulation is much larger than N to ensure a high-quality simulation of the SDE based on Euler approximation.




Fig. 2.1. Two sample realizations (a and b) of the spot price $S(t)$ (dotted line) and forward curve with monthly increments. At each timestep (month) the one-year forward curve, $F(t, t \leq T \leq t + 1)$ is shown for $t = 0, 1/12, 2/12, \dots, 12/12$ using Table 2.1 parameters, and $(\chi_0, \xi_0) = (-0.639, 4.637)$.

Consistent with the previous notation, let $F(t_1, t_2)$ denote the forward price at time t_1 for delivery at time t_2 . Using $N = 12$, Fig. 2.1 shows two sample realizations of the spot prices and their corresponding forward curves. Here, $(\chi_0, \xi_0) = (-0.639, 4.637)$, which corresponds to a spot price of \$54.45 per barrel and a long-term price of \$103.19 simulating forward prices based on May 2009 market conditions. Various slopes and curvatures are observed among the realizations. Usually, if the spot price is around the lower-end of the range, the forward curve is upward-sloping. However, if the spot price is around the higher-end of the range, it is downward-sloping. This is induced by the mean-reverting nature of the short-term deviation generated by the χ_t factor.

Since it is not possible to illustrate visibly many forward curve realizations on the same plot, and to have a statistically better representation of the forward curve dynamics, a different approach is taken. $M = 10,000$ sample paths are simulated using $N = 4$, i.e. quarterly Δt 's, and the following quantities are computed at times $t = 0, 0.25, 0.5$, and 0.75 ; (i) the mean of the forward prices, and (ii) the slope of the line *fitted* to the forward curve. Fig. 2.2.a shows the mean of forward prices




Fig. 2.2. Characteristics of the forward curve based on the parameters of Table 2.1, $T = 1$ year, $\Delta t = 0.25$, $M = 10,000$, and $(\chi_0, \xi_0) = (-0.639, 4.637)$; (a) Slope versus mean price at $t = 0, 0.25, 0.5$, and 0.75 . (b) Change in the slope versus change in the mean price over each period. Period 1 represents the time from $t = 0$ to 0.25 , Period 2 represents the time from $t = 0.25$ to 0.5 , and Period 3 represents the time from $t = 0.5$ to 0.75 .

versus the slope at all the four timesteps. Both positive and negative slopes are seen, while the mean prices vary mostly between \$40 and \$140. Generally, it can be said that low spot prices relate to high slopes, and high spot prices relate to lower (or even negative) slopes, which has important implications for the storage trade. Moreover, the change in the spot price and the change in the slope are negatively correlated. The likelihood of having a negative slope increases with the passage of time. Percentage of paths with a negative slope is computed as ratio of the number of the paths with a negative slope to the total number of the paths. This ratio increases from 10.4% at $t = 0.25$ to 54% at $t = 0.75$. In other words, the forward curve tends to move toward backwardation (down-ward sloping). However, within the one-year time frame, the forward curve is upward sloping most of the time; only 33.5% of the cases across all paths and all times have a negative slope.

To demonstrate the dynamics of the forward curve in terms of parallel shifts and changes in the slope, the *change* in the mean of forward prices and the *change* in the slope of the forward curve are calculated for two consecutive times, i.e. the change over a period. Period 1 represents the time from $t = 0$ to 0.25, Period 2 represents the time from $t = 0.25$ to 0.5, and Period 3 represents the time from $t = 0.5$ to 0.75. The result is illustrated in Fig. 2.2.b. Regardless of the period, it is seen that usually a decrease (increase) in the slope is accompanied by an increase (decrease) in the mean forward prices. The only impact of period is that the magnitude of the change (dispersion around the origin) increases with the passage of time. Across the three periods, 62.3% of the cases fall in the second quadrant, where a decrease in the slope occurs simultaneously with an increase in the mean price.

2.4 Summary

The Schwartz and Smith (2000) model is reviewed in detail. It is a two-factor model, which includes a long-term factor determining the long-term trend (equilibrium price) and short-term factor capturing the temporary deviations from the long-term trend. Many different forward curve shapes can be replicated with this model. Higher spot prices are accompanied with steeper downward-sloping forward curves, and vice versa. A series of sample paths are simulated using the initial condition $(\chi_0, \xi_0) = (-0.639, 4.637)$, which corresponds to the condition in May 2009.

This is equivalent to the spot price $S_0 = \$54.45$ and a long-term equilibrium price of $\$103.19$. Simulations indicate that as the prices approach the long-term (higher) level, they increase through each period, and the initial upward sloping forward curve tilts downward.

In the next chapter, the first (and simplest) framework for the proposed trade is presented, which is formulated as an MDP model. The Schwartz-Smith model provides the exogenous state variables in the MDP model, which determines the forward prices at each timestep. The proposed trade will result in a profit-maximizing optimization problem, which will be solved by a suboptimal yet simple approach in the next chapter, Chapter 3, and an optimal method in Chapter 4.

Chapter 3

3 Forward Dynamic Optimization

In this chapter, we study the Forward Dynamic Optimization (FDO) method, also known as the Rolling Intrinsic policy. This simple and intuitive strategy is suboptimal, although it may lead to near-optimal results under some circumstances, e.g. gas storage valuation, as discussed in the literature review.

Section 3.1 discusses the assumptions and establishes the MDP formulation of the model. Section 3.2 explains theoretically the intuition behind the FDO policy in terms of the slope of forward curve. Section 3.3 reports the values of parameters used to achieve the numerical results. Sections 3.4 and 3.5 contain the computational results and the chapter summary respectively.

3.1 Model

Assume a trader owns \bar{R} barrels (or \bar{R} units as quantity of oil is described in barrels) of crude oil stored in a tanker. All \bar{R} barrels of oil must be delivered by time \bar{T} , perhaps because the tanker must be returned to its owner by that date. The oil could be delivered earlier however. If the market is in contango and the forward price is higher than the spot price, it would be profitable to buy the oil at $F(t_0, t_0) = S(t_0)$ and immediately sell it forward for $F(t_0, \bar{T})$, thereby locking in a profit of $\bar{R}[F(t_0, \bar{T}) - S(t_0)]$. Assuming the oil is already owned at time t_0 , this strategy involves a short position in a single contract (for \bar{R} barrels) with delivery at time \bar{T} . There may, depending on the evolution of the forward curve, be better opportunities. For instance, if at time $s > t_0$ and delivery time $T_1 < \bar{T}$, $F(s, T_1) > F(s, \bar{T})$, it would be easy for the investor to buy back their time \bar{T} delivery contract and short a T_1 delivery contract. This would add a profit per barrel of $F(s, T_1) - F(s, \bar{T})$ with no risk. Forward Dynamic Optimization is the idea of taking all such immediately profitable contract modifications to profit from fluctuations in the forward market.

We need to introduce some notation before studying this problem in detail. First, we assume (in this chapter) that physical delivery is irrevocable, in the sense that refilling the inventory is not

permitted. However, any pair of long/short contracts (with maturity before or at \bar{T}) may be added to the trading strategy at any time, provided that the net exposure of all contracts is short one unit of the asset, to balance the long position in physical oil stored in the tanker. Therefore, the inventory will remain constantly hedged until delivery. Second, we assume that the time horizon is discretized into N equal periods by timesteps $t_0 = 0, t_1, t_2, \dots, t_N = \bar{T}$. Subsequently, let $\delta = \exp(-r(t_{i+1} - t_i)) = \exp(-r\Delta t)$ denote the constant time-discount factor for one period. We need to account for two types of operational costs. First set of operational costs, denoted by c_P , lumps together the cost of physically transferring the oil from the tanker to the delivery point (“pumping costs”) and any location discount to WTI futures. It is because physical delivery will take place at a port rather than Cushing, Oklahoma, which is the pricing point of the WTI contracts. The second set of operational costs combines the cost of renting the tanker with the cost of crewing/operating it. This “holding cost” is denoted by c_H . We assume that pumping costs must be paid at delivery time of the oil, while holding costs must be paid up front at t_0 . We also assume that if the trader delivers the oil earlier than the end of the tanker rental agreement she will be refunded the unused portion of the holding costs.

At $t_0 = 0$, the trader faces the following set of alternative decisions. One is to sell the oil on the spot market and receive a payoff of $S_0 - c_P$. This will lead to an immediate termination of the decision-making process and zero added value. However, there might exist more profitable alternatives by selling the oil using a forward contract for delivery at a later date. The trader must choose an optimal contract from those available, subject to the time constraints. Of course, the sale at t_0 will be the optimal choice if none of the forward contracts can provide additional value relative to the sale at t_0 . To formalize the problem, we first consider a discrete-time dynamic optimization framework based on the following components:

1. The State Variables (x_i and W_i): At any stage i , the trader owns an inventory level of \bar{R} , which is sold forward with a contract that has a maturity of T_i . In other words, the portfolio of the trader consists of a long asset position with a quantity of \bar{R} , and a short forward position with a maturity of T_i and the same quantity. Let x_i denote the endogenous component of the state variable at time t_i . So, the state can be defined by $x_i = (R_i, T_i)$, which determines the amount of oil in the

inventory, R_i , and the promised delivery date T_i . Also, the stage- i forward curve, $F(t_i, T)$, where $T \in \{t_i, t_{i+1}, \dots, t_N\}$, is fully specified by $W_i = (\chi_i, \xi_i)$, which are treated as exogenous state variables. This indicates that the distribution of W_{i+1} is not affected by the x_i , or the decision made. So, the current state is fully explained by $(x_i, W_i) \in \mathcal{X}_i \times \mathbb{R}^2$, where \mathcal{X}_i is the state space defined as in Eq. 3.1.

$$\mathcal{X}_i = \{\{0\} \times \{0\}\} \cup \{\{\bar{R}\} \times \{t_i, t_{i+1}, t_{i+2}, \dots, t_N\}\} \quad \text{Eq. 3.1}$$

The state $x_i = (0,0)$ is an absorbing state corresponding to an empty inventory condition, from which there will be no further decision making. At t_0 , the initial state is $x_0 = (\bar{R}, 0)$ and $W_0 = (-0.6393, 4.6366)$. This indicates $R_0 = \bar{R}$ and $T_0 = 0$ meaning that the trader starts with \bar{R} units in the tanker, and no forward contract at hand (one that matures today).

2. The Decisions (actions) (a_i): At stage i and state (x_i, W_i) , the decisions to be made is to select a_i , which is the maturity of the short forward position.

$$a_i \in \mathcal{A}_i(x_i) \text{ for } i \in \{0, 1, 2, \dots, N-1\} \quad \text{Eq. 3.2}$$

$\mathcal{A}_i(x_i)$ is the feasible set, which depends on the current state and defined by Eq. 3.3.

$$\mathcal{A}_i(x_i) = \begin{cases} \{0\} & \text{if } R_i = 0 \\ \{t_i, t_{i+1}, t_{i+2}, \dots, t_N\} & \text{if } R_i > 0 \end{cases} \quad \text{Eq. 3.3}$$

If $R_i > 0$, the trader can choose a new maturity from t_i (corresponding to sale on the spot) to t_N (latest possible delivery date). The new maturity a_i can be $a_i > T_i$ or $a_i \leq T_i$, which represents postponing or advancing the current contract maturity respectively.

3. State Transition Function $f_i(x_i, a_i)$: Given the current state (x_i, W_i) and the action a_i , the endogenous part of the next state $x_{i+1} = f_i(x_i, a_i)$ follows (in a deterministic fashion) using the state transition function as expressed in Eq. 3.4.

$$x_{i+1} = (R_{i+1}, T_{i+1}) = f_i(x_i, a_i) = \begin{cases} (0,0) & \text{if } a_i = t_i \text{ or } R_i = 0 \\ (R_i, a_i) & \text{if } a_i > t_i \end{cases} \quad \text{Eq. 3.4}$$

The exogenous part of the state, W_i , evolves based on the stochastic processes of Eq. 2.2 and Eq. 2.3 independently from x_i and a_i .

4. Reward Function $r_i(a_i, x_i, W_i)$: Given the state $x_i = (R_i, T_i)$ and the action a_i at stages $i \in \{0, 1, 2, \dots, N - 1\}$, there will be a reward generated by going short R_i barrels through choosing the contract with maturity a_i , i.e. shorting $F(t_i, a_i)$ contract. Assume that the trader enters time t_i , at which time she holds a short position in a contract that matures at time T_i ($t_i \leq T_i$). According to the assumptions, the tanker rent has already been paid up to time T_i .

Three elements of the payoff are as the following. Firstly, the current contract held must be offset by going long the $F(t_i, T_i)$ contract, which results in a cash outflow of $-e^{-r(T_i-t_i)}(F(t_i, T_i) - c_p)$. Secondly, there is the payoff due to entering the newly chosen short position with maturity a_i , which results in the cash inflow of $e^{-r(a_i-t_i)}(F(t_i, a_i) - c_p)$. Finally, there is the payoff from the rental time adjustment, which is $c_H(a_i - T_i)$. If $a_i > T_i$, the trader must pay extra rent for the additional time beyond T_i . If $a_i < T_i$, the trader will receive a refund equal to the rent for the unused portion. Fig. 3.1 highlights the new long and short contracts being considered if $T_i < a_i$. All the three payoff elements are combined in Eq. 3.5.




Fig. 3.1. New positions to capture any potential gains

$$r_i(a_i, x_i, W_i) = R_i \left[e^{-r(a_i-t_i)}(F(t_i, a_i) - c_p) - e^{-r(T_i-t_i)}(F(t_i, T_i) - c_p) - c_H(a_i - T_i) \right] \quad \text{Eq. 3.5}$$

Note that $a_i = T_i$ is equivalent to a ‘hold’ decision with a zero payoff. At the terminal time, $t_N = \bar{T}$, there is no decision making and $r_N(x_N, W_N) = r_0 = 0$, because by $t_N = \bar{T}$ either the oil has

been already sold, or an existing short contract with $T_N = \bar{T}$ will be fulfilled by delivering the inventory.

The optimization problem, which gives the real option value at t_0 , is expressed by Eq. 3.6. The optimization is over the class Π , where Π is the set of all feasible policies π . A policy π is defined as the set of decision functions $\{A_0^\pi, A_1^\pi, A_2^\pi, \dots, A_{N-1}^\pi\}$, where $A_i^\pi(x_i, W_i): \mathcal{X}_i \times \mathbb{R}^2 \rightarrow \mathcal{A}_i(x_i)$ for $\forall i \in \{0, 1, 2, \dots, N-1\}$.

$$V_0(x_0, W_0) = \max_{\pi \in \Pi} E \left[\sum_{i=0}^{N-1} \delta^i r_i(A_i^\pi(x_i^\pi, W_i), x_i, W_i) \mid (x_0, W_0) \right] \quad \text{Eq. 3.6}$$

The expectation is under the physical measure as the goal is to capture the performance of the trading strategy. Also, x_i^π denotes the random endogenous part of the state at stage i when policy π is implemented. The Bellman equation associated with the above problem is expressed by Eq. 3.7, where $V_i(x_i, W_i)$ is the value function at timestep i and state (x_i, W_i) .

$$V_i(x_i, W_i) = \max_{a \in \mathcal{A}_i(x_i)} \{r_i(a, x_i, W_i) + \delta E[V_{i+1}(f_i(x_i, a), W_{i+1}) \mid W_i]\},$$

$$\forall (x_i, W_i) \in \mathcal{X}_i \times \mathbb{R}^2, \forall i \in \{0, 1, 2, \dots, N-1\},$$

$$V_N(x_N, W_N) = r_N(x_N, W_N) = 0, \quad \forall (x_N, W_N) \in \mathcal{X}_N \times \mathbb{R}^2$$

The myopic approach to solve Eq. 3.7 advised by the FDO method is to ignore $\delta E[V_{i+1}(f_i(x_i, a), W_{i+1}) \mid W_i]$ and to maximize $r_i(a, x_i, W_i)$ by searching for the optimal action $a \in \mathcal{A}_i(x_i)$. Accordingly, the FDO algorithm prescribes the following policy; Eq. 3.8 to be followed sequentially at $i = 0, 1, 2, \dots, N-1$, which leads to stage-wise deterministic payoffs of Eq. 3.9. According to this policy, the trader only maximizes the immediate profit from any present opportunity, where $v_i(x_i, W_i)$ represents the total value generated by such trade at t_i . This policy does not take into account the impact of an action today on the value potentially harvested in the future. We can argue that $v_i(x_i, W_i) \geq 0, \forall i, x_i, W_i$ because we know that in the maximization of Eq. 3.9 one can always choose to hold the same maturity, i.e. $a_i = T_i$, leading to a zero payoff. The way that the above algorithm performs explains why it is called FDO; at $t_0, t_1, t_2, \dots, t_{N-1}$, the trader performs a Forward Dynamic Optimization, which is to readjust her position given the most

recently realized forward prices. This readjustment captures any favorable shift in the forward curve, which can only add non-negative value.

$$A_i^{FDO}(x_i, W_i) = \operatorname{argmax}_{a_i \in \mathcal{A}_i(x_i)} r_i(a_i, x_i, W_i) \text{ for } i \in \{0, 1, 2, \dots, N-1\} \quad \text{Eq. 3.8}$$

$$v_i(x_i, W_i) = \max_{a_i \in \mathcal{A}_i(x_i)} r_i(a_i, x_i, W_i) \text{ for } i \in \{0, 1, 2, \dots, N-1\} \quad \text{Eq. 3.9}$$

Total value achieved by following the FDO policy can be computed based on in Eq. 3.10.

$$\begin{aligned} V_0^{FDO}(x_0, W_0) &= E \left[\sum_{i=0}^{N-1} \delta^i \max_{a_i \in \mathcal{A}_i(x_i)} r_i(a_i, x_i, W_i) \mid (x_0, W_0) \right] \\ &= E \left[\sum_{i=0}^{N-1} \delta^i v_i(x_i, W_i) \mid (x_0, W_0) \right] \end{aligned} \quad \text{Eq. 3.10}$$

Eq. 3.10 can be written more explicitly as Eq. 3.11. The expectations in Eq. 3.10 or Eq. 3.11 can be calculated using Monte Carlo simulations by computing the payoffs along each path in a forward-moving fashion.

$$\begin{aligned} V_0^{FDO}(x_0, W_0) &= E_{t_0} \left[\sum_{i=0}^{N-1} \delta^i \max_{a_i \in \mathcal{A}_i(x_i)} R_i \left[e^{-r(a_i - t_i)} (F(t_i, a_i) - c_P) \right. \right. \\ &\quad \left. \left. - e^{-r(T_i - t_i)} (F(t_i, T_i) - c_P) - c_H(a_i - T_i) \right] \right] \end{aligned} \quad \text{Eq. 3.11}$$

3.2 Theoretical Interpretation and Comparison

To shed light on the intuition behind the FDO policy, let us re-write Eq. 3.9 as Eq. 3.14, where the interim steps are expressed explicitly.

$$\begin{aligned} v_i^*(x_i, W_i) &= \max_{a_i \in \mathcal{A}_i(x_i)} r_i(a_i, x_i, W_i) = \\ &R_i \max_{a_i \in \{t_i, t_{i+1}, t_{i+2}, \dots, t_N\}} \left[e^{-r(a_i - t_i)} (F(t_i, a_i) - c_P) - e^{-r(T_i - t_i)} (F(t_i, T_i) - c_P) \right. \\ &\quad \left. - c_H(a_i - T_i) \right] \end{aligned} \quad \text{Eq. 3.12}$$

For simplicity in exposition, assume that $r = 0$; this constraint can easily be removed at the cost of more complicated notation.

$$v_i^*(x_i, W_i) = R_i \max_{a_i \in \{t_i, t_{i+1}, t_{i+2}, \dots, t_N\}} [F(t_i, a_i) - c_P - F(t_i, T_i) + c_P - c_H(a_i - T_i)] \quad \text{Eq. 3.13}$$

$$= R_i \max_{a_i \in \{t_i, t_{i+1}, t_{i+2}, \dots, t_N\}} [F(t_i, a_i) - F(t_i, T_i) - c_H(a_i - T_i)]$$

$$= R_i \max_{a_i \in \{t_i, t_{i+1}, t_{i+2}, \dots, t_N\}} (a_i - T_i) \left[\frac{F(t_i, a_i) - F(t_i, T_i)}{a_i - T_i} - c_H \right]$$

$$= R_i \max \begin{cases} \max_{t_i \leq a_i < T_i} (a_i - T_i) \left[\frac{F(t_i, a_i) - F(t_i, T_i)}{a_i - T_i} - c_H \right] & a_i < T_i \\ 0 & a_i = T_i \\ \max_{T_i < a_i \leq t_N} (a_i - T_i) \left[\frac{F(t_i, a_i) - F(t_i, T_i)}{a_i - T_i} - c_H \right] & T_i < a_i \end{cases}$$

$$= R_i \max \begin{cases} \max_{t_i \leq a_i < T_i} (T_i - a_i) \left[c_H - \frac{F(t_i, a_i) - F(t_i, T_i)}{a_i - T_i} \right] & a_i < T_i \\ 0 & a_i = T_i \\ \max_{T_i < a_i \leq t_N} (a_i - T_i) \left[\frac{F(t_i, a_i) - F(t_i, T_i)}{a_i - T_i} - c_H \right] & T_i < a_i \end{cases}$$

This can be summarized as Eq. 3.14. The term $(F(t_i, a_i) - F(t_i, T_i))/(a_i - T_i)$ is the slope of the forward curve across the two legs of the spread. The trader can always choose $a_i = T_i$, i.e. ‘hold’, which results in the zero payoff, i.e. $v_i^*(x_i, W_i) = 0$. So, for any choice $a_i \neq T_i$ to be optimal, it must generate a positive payoff. Thus, the third line of Eq. 3.14 indicates that the trader will search

$$v_i^*(x_i, W_i) = \quad \text{Eq. 3.14}$$

$$R_i \max \begin{cases} \max_{t_i \leq a_i < T_i} (T_i - a_i) \left[c_H - \frac{F(t_i, a_i) - F(t_i, T_i)}{a_i - T_i} \right] & a_i < T_i \text{ (advancing)} \\ 0 & a_i = T_i \text{ (holding)} \\ \max_{T_i < a_i \leq t_N} (a_i - T_i) \left[\frac{F(t_i, a_i) - F(t_i, T_i)}{a_i - T_i} - c_H \right] & a_i > T_i \text{ (postponing)} \end{cases}$$

for any maturity $a_i > T_i$ such that it maximizes the slope of the forward curve above c_H . In the first line in Eq. 3.14, the trader will search for any maturity $a_i < T_i$, which minimizes (maximizes the negative of) the slope below c_H . This is equivalent to searching over the *postponing* ($T_i < a_i$) opportunity set and the *advancing* ($a_i < T_i$) opportunity set, respectively, compared to the existing maturity T_i .

In the following, the FDO algorithm presented in this chapter is compared with the literature (Jafarizadeh and Bratvold, 2013); Jafarizadeh and Bratvold (2013) defined *Forward Maximization* as the present value (as of time t_i) of the profit from selling the oil by entering the most profitable forward contract considering pumping and storage costs, as stated in Eq. 3.15. It is assumed that the cost of storage has been paid up to t_i , and $\bar{F}(t_i)$ denotes the forward curve vector at t_i . We are no longer assuming that $r = 0$.

$$FM(\bar{F}(t_i), c_H, c_P, t_i) = \max_{t_i \leq t_j \leq t_N} \{e^{-r(t_j-t_i)}(F(t_i, t_j) - c_P) - c_H(t_j - t_i), 0\} \quad \text{Eq. 3.15}$$

According to Jafarizadeh and Bratvold (2013), which followed the formulation given by Eydeland and Wolyniec (2003) for storage valuation, the FDO algorithm is formulated as in Eq. 3.16.

$$V_0^{FDO,2} = FM(\bar{F}(t_0), c_H, c_P, t_0) + E_{t_0} \left[\sum_{i=1}^N \max\{FM(\bar{F}(t_i), c_H, c_P, t_i) - FM(\bar{F}(t_{i-1}), c_H, c_P, t_{i-1}), 0\} \right] \quad \text{Eq. 3.16}$$

One issue with the above formula is that $FM(\bar{F}(t_i), c_H, c_P, t_i)$ and $FM(\bar{F}(t_{i-1}), c_H, c_P, t_{i-1})$ are net present values associated with times t_i and t_{i-1} , respectively, and thus cannot be subtracted without appropriate time discounting. The other issue is that each element of the summation must be discounted back to time $t = t_0$ to be consistent with $FM(\bar{F}(t_0), c_H, c_P, t_0)$. Finally and most importantly, the difference $FM(\bar{F}(t_i), c_H, c_P, t_i) - FM(\bar{F}(t_{i-1}), c_H, c_P, t_{i-1})$ does not correspond to any possible explicit trade in forward contracts to capture the favorable shifts in forward curves.

This formulation is essentially an impossible trade. Consider the following simplified case, where $N = 2$, i.e. a two-period problem with timesteps $t_0 = 0$, $t_1 = 0.5$, and $t_2 = 1$. Furthermore, assume that in Eq. 3.16, the expression in the maximum operator is positive at t_1 and t_2 , which is equivalent to the existence of a profitable trade at both t_1 and t_2 . According to Eq. 3.16, all but one term will be cancelled out in the telescoping sum, and the resulting expression will be $V_0^{FDO,2} = E_{t_0}[FM(\bar{F}(t_2), c_H, c_P, t_2)]$. This outcome is equivalent to selling the oil at the spot price at time t_2 , $F(t_2, t_2)$, because $t_2 = 1$ is the problem time constraint for delivery of the oil. Selling the oil at the spot price at the problem time horizon requires to carry the long oil position unhedged through time, which is not allowed in the present setup. If the trader were to enter a position at $t_0 = 0$, then she could only “re-adjust” her position at timesteps t_1 and t_2 , which would only generate incremental gains not necessarily summing to a sale price equivalent to $F(t_2, t_2)$. In contrast, the formulation presented in the current research in Eq. 3.10, there is a clear correspondence to trades in the forward contracts.

Another fundamental difference between the present approach and the previous studies, as highlighted in Eq. 3.11 and Eq. 3.16, is as follows; at any time t_i , the decision-making only depends on the time t_i information, W_i or $\bar{F}(t_i)$, as seen in Eq. 3.11. However, in Eq. 3.16, the optimal decision at time t_i incorporates $\bar{F}(t_i)$ as well as $\bar{F}(t_{i-1})$. The method of Eq. 3.16 computes the incremental gain at t_i incorrectly, because it subtracts the old prices as opposed to new prices to account for entering a new long position to offset the existing short position.

3.3 Parameters

We generated $M = 10000$ paths by simulating the state variables with $\Delta t = 1/10080$ year using Eq. 2.16, Eq. 2.17, and Eq. 2.18. The number of partitions in a year ($10,080 = 2^5 \times 3^2 \times 5 \times 7$) is chosen such that it allows perfect divisibility to many larger Δt 's: e.g. it allows to discretize a year into 2, 3, 4, 5, 6, 7, ..., etc. periods, using which corresponding prices from the underlying fine discretization can be extracted. For consistency, the same $\Delta t = 1/10080$ year is used for simulating prices in all cases in this chapter.

The value of trading frequency is set to $\Delta t = 0.25$ (quarterly), which is equivalent to $N = 4$ periods based on the $\bar{T} = 1$ year time horizon. The case with $N = 4$ time periods, $M = 10,000$ simulations, storage cost of $c_H = \$6.57$ per barrel per year, and initial condition of $(\chi_0, \xi_0) = (-0.639, 4.637)$ will be referred to as the base case, and will serve as a basis in all of the investigations, as listed in Table 3.1.

Description	Parameter	Value
Number of periods	N	4
Storage cost	c_H	\$6.57/barrel.year
Pumping cost	c_P	\$3.75/barrel
Number of simulated path	M	10,000
Initial condition of state variables	$W_0 = (\chi_0, \xi_0)$	(-0.639, 4.637)
Initial Spot Price	$S_0 = \exp(\chi_0 + \xi_0)$	\$54.45
Time Horizon (constraint)	\bar{T}	1 year

Table 3.1. Problem parameters as specified by the base case for FDO analysis.

The chosen initial condition and storage cost represents the corresponding values on May-3-2009, as an example of a favorable period of time for this type of trade. The initial condition, $(\chi_0, \xi_0) = (-0.639, 4.637)$, can be translated into $S_0 = \$54.45/\text{barrel}$ and a long-term equilibrium price of $\$103.19/\text{barrel}$ based on the price model. In regard with the storage and pumping cost, Jafarizadeh and Bratvold (2013) assumed a tanker rent (c_H) of $\$0.7\text{M}/\text{year}$ with a pumping cost (c_P) of $\$1\text{M}$. Their assumed tanker rent is equivalent to $\$1,918/\text{day}$ for the whole tanker, which is very low compared to actual prices. Based on actual time charter prices of a 2 million barrel VLCC around May 2009, as shown in Fig. 1.2, we assume the tanker rent to be about $\$36,000/\text{day}$, which is equivalent to the chosen $c_H = (\$36,000 \times 365) / 2,000,000 = \6.57 per barrel per year, which is much higher than the $\$3.5$ per barrel per year assumed by Jafarizadeh and Bratvold (2013). Also, we consider a pumping cost of $c_P = \$3.75/\text{barrel}$. Occasionally, we focus on the impact of a particular parameter and an alternative value relative to the base case is employed, where in such cases, the specific range of the parameter will be provided accordingly.

3.4 Computational Results

Different aspects of the results are studied using the following metrics of value; Added Value (\$) of the FDO strategy is the gain relative to selling the oil on the spot price at time zero, and is equal to $V_0^{FDO}(x_0, W_0)$, or V_0^{FDO} in short. Added Value (%) is useful when comparing the added value under two different initial spot prices or initial conditions.

$$\text{Added Value (\$)} := V_0^{FDO}, \quad \text{Added Value (\%)} := \frac{\text{Added Value (\$)}}{S_0} \quad \text{Eq. 3.17}$$

In the base case, the 95% confidence interval for Added Value (\$) is \$8.94-\$9.05. For a two-million-barrel capacity VLCC tanker, this is equivalent to \$17.88-\$18.10 million. Fig. 3.2 shows the histogram of payoffs across $M = 10K$ paths; the strategy guarantees \$6.19 since it is the minimum achieved value. It is worth noting that this \$6.19 is generated by the spread between $F(0,1)$ and $F(0,0) = S(0)$ (after subtracting the storage cost). The payoffs are significantly skewed to the right. The Added Value (\$) of \$8.99 can be broken down into two parts; (i) \$6.19 generated by selling the oil forward using $F(0,1)$, and (ii) \$2.80 obtained by all the subsequent trades. At $t = 0$, part (i) is known (certain), whereas part (ii) is uncertain. Although most of the value is captured by the initial trade $F(0,1) - S(0)$, the subsequent trades are necessary to capture the remaining 31% of the Added Value (\$). The contribution of the subsequent trades to Added Value (\$) will increase if the initial forward curve is less steep, i.e. in a less favorable environment to start the trading. So, a trader still has an incentive to start although the initial part of the added value is not very high.

3.4.1 Decision-making analysis at the path level

In the analysis of the following section, the slope of forward curve will often be used. To provide intuition about the slope, a symbolic forward curve at $t = 0.25$, $F(0.25, T)$, is shown in Fig. 3.3. Recall in Fig. 3.1, the long and short positions, i.e. the two legs, involved in a trade were explained. Fig. 3.3 demonstrates how the slope between the selected legs of the trade depends on the maturities considered. An approximate ‘slope’ of the forward curve can be estimated using the regression line.

To focus on the impact of the forward curve slope, payoffs are presented in terms of slopes at the current and the preceding timesteps. Realizations of slope and its consequences on the decisions




Fig. 3.2. Histogram of payoffs (\$) (i.e. generated cashflow relative to selling on the spot price along each path) over $M = 10K$ paths using base case parameters.

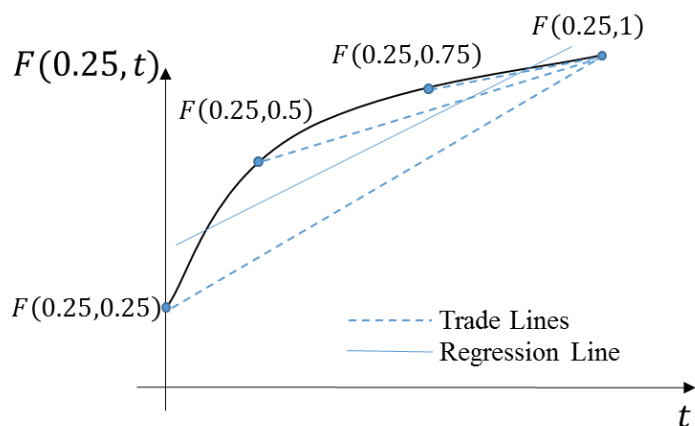


Fig. 3.3. A sample forward curve at $t=0.25$.

and payoffs at each timestep are tracked. Fig. 3.4 shows the payoff at $t = 0.25$, where all paths have the same history because the algorithm starts by shorting $F(0,1)$ at $t = 0$ in the base case. As listed in Fig. 3.4, there are four selected contracts, which correspond to four different decision regions labeled A through D, depending on the slope realized at $t = 0.25$. Region A is where the

realized slope is too high, and $F(0,1)$ is held as the optimal contract because there is no profitable choice. Region B is where the slope is slightly lower, and $F(0.25,0.75)$ is selected accordingly. As the slope falls further, $F(0.25,0.5)$ is selected in regions C. If the slope decreases to the lowest part of the range, region D, $F(0.25,0.25)$ is chosen as optimal. As the slope of forward curve decreases more, the maturity date of the new contract is moved from $t = 1$ (hold existing contract) to $t = 0.25$ (sell on the spot), to generate a larger payoff.




Fig. 3.4. Payoff and selected contract at $t = 0.25$ as a function of slope at $t = 0$ and $t = 0.25$. The trader holds $F(0,1)$, or advances the maturity one, two, or three periods based on the realized slope at $t = 0.25$. All the parameters are per the base case as specified in Table 2.1 and Table 3.1.

Note that the intervals identifying the regions have a small overlap since the slope of the forward curve is computed as the slope of the fitted line approximating the forward curve. Although this approximation works very well, it is not completely accurate due to the inherent curvature of the forward curve. When the algorithm searches for a potential contract to short, it considers each spread trade (long one and short another contract) individually, leading to a slightly different slope

from the fitted line. The dashed lines, which illustrate the concept by dividing the regions, are based on the midpoint of the overlap area.

Subsequently, the contracts holding which the trader may start trading at $t = 0.5$ fall into one of the three types; $F(0,1)$, $F(0.25,0.5)$, and $F(0.25,0.75)$ (listed under the existing contracts in Fig. 3.5). Fig. 3.5 shows the payoff at $t = 0.5$ as a function of the slope at $t = 0.25$ and $t = 0.5$, where seven regions are identified, and labeled A through G. For each region, the existing contracts are determined based on the corresponding slope at $t = 0.25$. Conditional on an existing contract, the realized slope at $t = 0.5$ will determine the new optimal contract.

Region	Slope at $t=0.25$ (x)	Existing contract when entering $t=0.5$	Slope at $t=0.5$ (y)	Action / New contract chosen at $t=0.5$
A	$9.46 \leq x$	$F(0,1)$	(7.95,31.95)	Hold / $F(0,1)$
B	$9.46 \leq x$	$F(0,1)$	(-38.70,8.45)	Advance / $F(0.5,0.75)$ or $F(0.5,0.5)$
C	$7.13 \leq x < 9.46$	$F(0.25,0.75)$	(8.09,22.90)	Postpone / $F(0.5,1)$
D	$7.13 \leq x < 9.46$	$F(0.25,0.75)$	(5.88,8.52)	Hold / $F(0.25,0.75)$
E	$7.13 \leq x < 9.46$	$F(0.25,0.75)$	(-55.79,5.85)	Advance / $F(0.5,0.5)$
F	$5 \leq x < 7.13$	$F(0.25,0.5)$	(5.91,22.89)	Postpone / $F(0.5,0.75)$ or $F(0.5,1)$
G	$5 \leq x < 7.13$	$F(0.25,0.5)$	(-38.79,5.85)	Hold / NA (terminates)

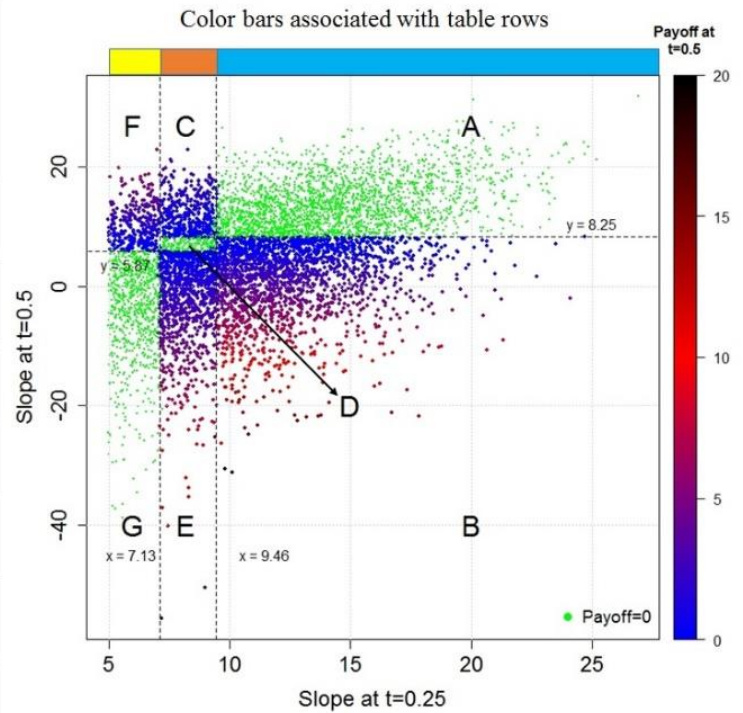


Fig. 3.5. Payoff and selected contract at $t = 0.5$ as a function of slope at $t = 0.25$ and slope at $t = 0.5$. Starting with an existing contract based on the slope at $t = 0.25$, the trader acts by choosing a new contract based on the realized slope at $t = 0.5$. All the parameters are per the base case as specified in Table 2.1 and Table 3.1.

In regions A and B, $F(0,1)$ is the existing contract. If the slope realized at $t = 0.5$ is very high, $F(0,1)$ will be held (region A). Otherwise, $F(0.5,0.75)$ or $F(0.5,0.5)$ will be chosen (region B) depending on how much the slope decreases; if the slope decreases to a large extent, $F(0.5,0.5)$ is preferred to $F(0.5,0.75)$.

In regions C, D and E, the existing contract is $F(0.25,0.75)$. If there is a sufficiently large increase in the slope (region C), the new contract will be $F(0.5,1)$. If there is a sufficiently large decrease in the slope (region E), the new contract will be $F(0.5,0.5)$. Alternatively, if the change in the slope is neither large- nor small-enough (region D), the existing contract, $F(0.25,0.75)$, will be held.

In regions F and G, $F(0.25,0.5)$ is the existing contract. If the slope does not increase enough, the existing contract, $F(0.25,0.5)$, will be held (region G), and the trade will conclude with delivery of the oil. If the slope increases largely, the net short position will be postponed by choosing $F(0.5,0.75)$ or $F(0.5,1)$ (region F) depending on how much the slope increases, where $F(0.5,1)$ is preferred to $F(0.5,0.75)$ at the larger end of the slopes.

If the paths with a zero payoff in region G are considered, the selected optimal maturity at $t = 0.25$ was $t = 0.5$, which is the next time step. So, regardless of how much the slope falls by $t = 0.5$, it is not possible to advance the trade any further. Thus, a small gain at $t = 0.25$ deprived these paths from a larger potential gain at $t = 0.5$. A similar argument for the paths in region A with an existing $F(0,1)$ contract can be given; regardless of how large the increase in the slope at $t = 0.5$ is, it is not possible to postpone this maturity anymore. To conclude, it might have been better to make less profit (or incur some loss) at $t = 0.25$ by avoiding the $F(0.25,0.5)$ or $F(0,1)$ contracts. This observation highlights the short-sighted nature of FDO trading policy.

At $t = 0.75$, existing contracts are comprised of four different types; $F(0.25,0.75)$, $F(0.5,0.75)$, $F(0.5,1)$, and $F(0,1)$. They can be categorized into two maturity dates, $t = 0.75$ and $t = 1$, which are sufficient for tracking the subsequent actions. Fig. 3.5 shows that if the slope at $t = 0.5$ exceeds 8.25, the selected maturity time is 1, and it is 0.75 otherwise, which is reflected in Fig. 3.6.

In Fig. 3.6, the maturity of the existing contracts is $t = 1$ in regions A and B. If the slope realized at $t = 0.75$ is greater than a threshold, the existing contract, $F(0.5,1)$ or $F(0,1)$, will be held (region A). However, if it is less than the threshold (6.97), $F(0.75,0.75)$ will be chosen (region B).

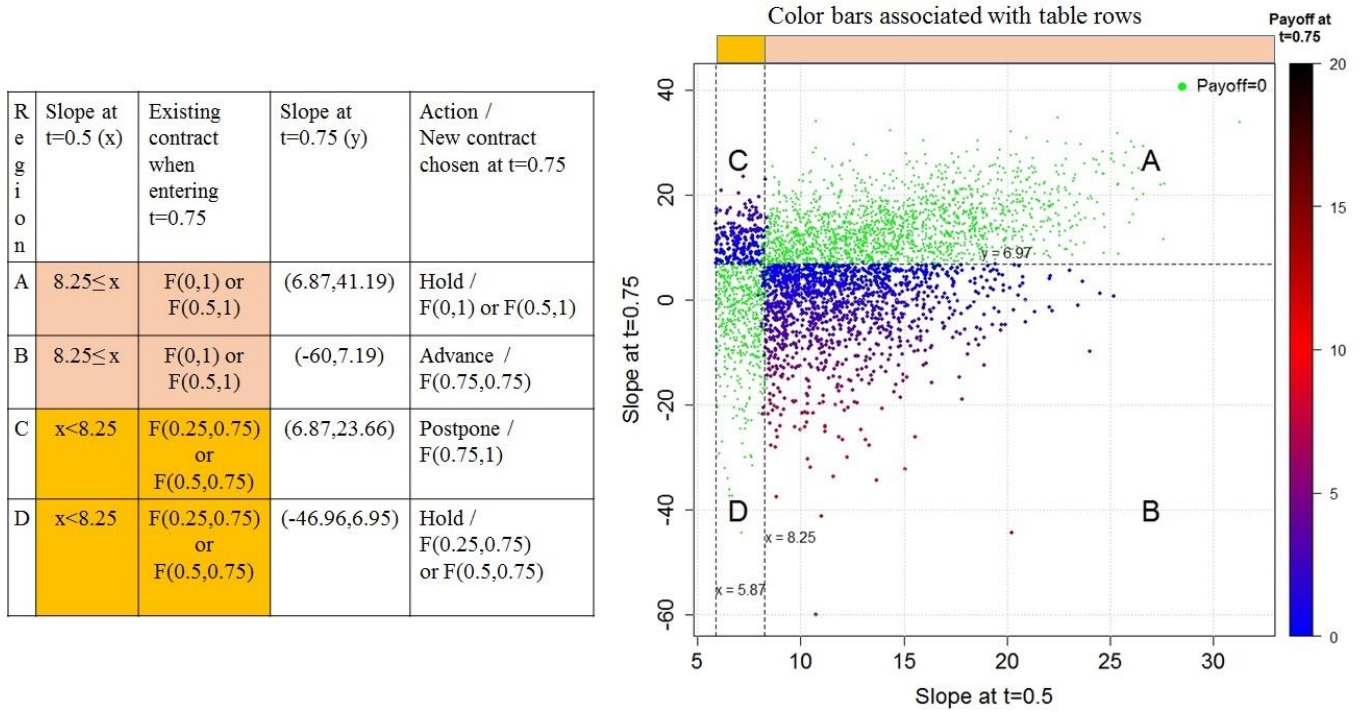


Fig. 3.6. Payoff and selected contract at $t = 0.75$ as a function of slope at $t = 0.5$ and slope at $t = 0.75$. Starting with an existing contract based on the slope at $t = 0.5$, the trader acts by choosing a new contract based on the realized slope at $t = 0.75$. All the parameters are according to the base case as specified in Table 2.1 and Table 3.1.

In regions C and D, the existing maturity date is $t = 0.75$. If the slope realized at $t = 0.75$ is greater than 6.97, $F(0.75,1)$ will be chosen (region C) to effectively postpone the net short position, while a smaller slope at $t = 0.75$ implies the existing contract, $F(0.25,0.75)$ or $F(0.5,0.75)$, will be held (region D), and the trade will terminate without a profit.

Fig. 3.4, Fig. 3.5, and Fig. 3.6 signify the critical role of the forward curve slope in explaining the rationale for selecting new contracts; the realized slope is constantly compared against the storage cost to detect any sufficiently large spread between the two, which may trigger a

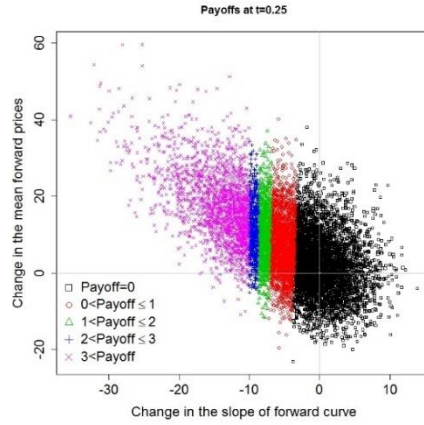
postponing/advancing decision if the currently held maturity permits. In the following, we discuss how there is a threshold of around $c_H = \$6.57/\text{barrel}\cdot\text{year}$, which governs the decisions.

In Fig. 3.4, Fig. 3.5, and Fig. 3.6, the slope marked by the dashed line between regions A and B, i.e. hold $F(0,1)$ or advance respectively, is 9.46, 8.25, and 6.97. Although there is variation among the three values, they get closer to $c_H = \$6.57$, which is due to the slight difference between the slope based on which the trades are made and the slope of the line fitted to the forward curve. With the passage of time, the length of the forward curve decreases due to a lower number of available contracts to a point. At $t = 0.75$, the forward curve is comprised of only two contracts, and the curve coincide with the fitted line. Therefore, the values, 9.46, 8.25, and 6.97, get closer to $c_H = \$6.57$. The remaining discrepancy is because the slope considers neither the time value of pumping cost nor discounting the forward prices both of which are, in contrast, reflected in the selection of the optimal decision-making.


In Fig. 3.7, the realized payoffs at $t = 0.25, 0.5$, and 0.75 are overlaid on a plot whose axes show the *change* in the slope of forward curve and the *change* in mean forward prices through each period. A decrease (increase) in the slope is usually accompanied by an increase (decrease) in the mean forward prices such that over the three timesteps, 62% of the cases falls in the second quadrant; low spot prices relate to high slopes, and high spot prices to low (or even negative) slopes. The magnitude of the change (dispersion around origin) increases with the passage of time.

Fig. 3.7.a shows at $t = 0.25$ the larger the change in the slope, the higher the payoff, whereas the change in mean forward prices has no effect on the result. Payoff levels at $t = 0.5$ and 0.75 , exhibits a weaker dependence on the change in the slope than $t = 0.25$. At $t = 0.5$ and $t = 0.75$, unlike $t = 0.25$, both positive and negative changes in the slope may lead to non-zero payoffs.


In panels (b) and (c) of Fig. 3.7, positive-payoff points are spread around the Y-axis asymmetrically. The cluster of the points on the right-hand side generate value by postponing, and the one on the left by advancing. The left cluster is more populated than the right one, indicating the higher likelihood of advancing since the trade started with a $F(0,1)$ position.



(a) $t = 0.25$



(b) $t = 0.5$




(a) $t = 0.75$

Fig. 3.7. Payoff levels at (a) $t = 0.25$, (b) $t = 0.5$, and (c) $t = 0.75$ in terms of change in the slope and mean forward prices. Payoffs magnitude is divided into five different levels at each time, and for clarity, two plots are provided in parts (b) and (c). All the parameters are according to the base case as specified in Table 2.1 and Table 3.1.

3.4.2 Impact of the Initial Condition

In this section, (χ_0, ξ_0) pairs are chosen based on a 21×21 uniform grid of a χ_0 - ξ_0 domain, where the domain is the rectangle defined by $-0.7 \leq \chi_0 \leq -0.3$ and $3.8 \leq \xi_0 \leq 5.2$. This generates 441 different initial conditions. The results are shown in Fig. 3.8.a, where the (χ_0, ξ_0) rectangular region is mapped into the corresponding “price” region in terms of $S_0 = \exp(\chi_0 + \xi_0)$ and long-term price = $\exp(\xi_0)$ transformations. Some of the historical values of (χ_0, ξ_0) from the period between August 2008 and October 2011 are extracted from Figure 4 in Hahn et al (2014), which shows the estimated evolution of the oil prices in terms of (χ_t, ξ_t) . These are the 21 black squares shown on Fig. 3.8.a, 15 of which lead to a positive Added Value (%). The figure shows that for instance, at $S_0 = \$60$, the long-term price should be greater than \$90 to have a steep-enough initial slope, and thus a non-zero profit.



(a) Added Value (%)

(b) Certain vs Uncertain Added Value (\$)

Fig. 3.8. (a) Added Value (%) as a function of the initial prices; transformation of different (χ_0, ξ_0) to the prices lead to $\$22.20 \leq S_0 \leq \134.29 and long-term price in the range $\$44.7$ - $\$181.27$. (b) Certain Added Value (\$) vs Uncertain Added Value (\$) colored by the initial slope; each (χ_0, ξ_0) implies an initial slope and a unique decomposition of Added Value (\$) into certain and uncertain parts. The solid and dashed black lines represent $X + Y = \text{constant}$ (i.e. total value) and $Y = X$ lines respectively. The parameters (other than the initial conditions) are per the base case as specified in Table 2.1 and Table 3.1.

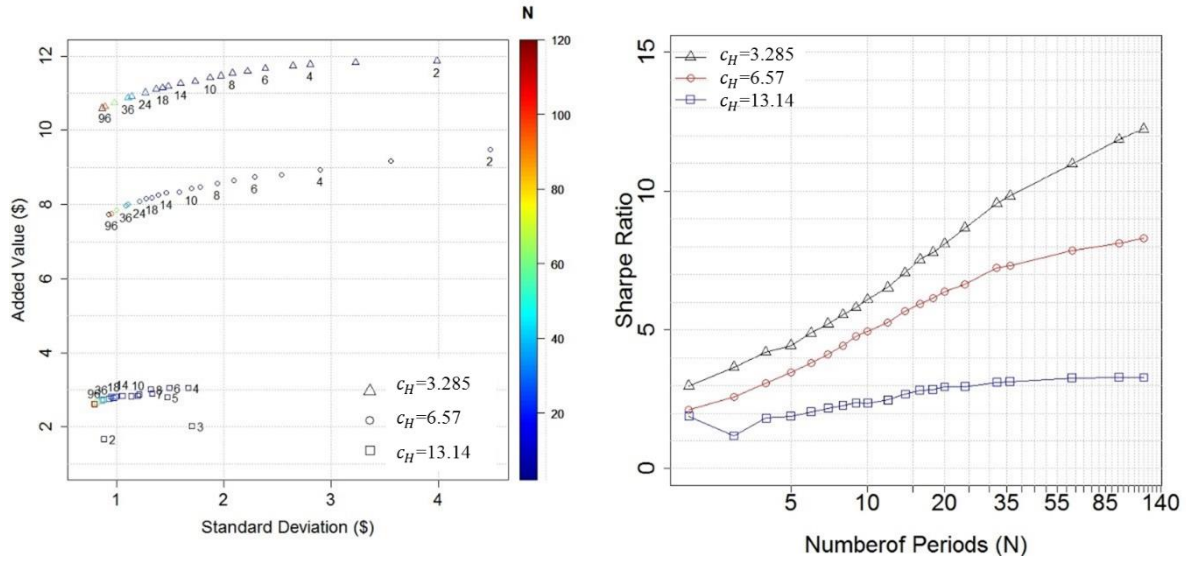
Fig. 3.8.b shows the relation between Certain Added Value (\$) and Uncertain Added Value (\$), where the former is generated by the initial forward maximization, and the latter by all the subsequent trades. In Fig. 3.8.b, both values are computed for each initial condition and are shown on the x- and y-axis. Also, the figure is colored by the initial slope of the forward curve implied by each of the initial conditions. A decrease in χ_0 , i.e. a larger initial deviation in the spot price from the long-term value, or an increase in ξ_0 , i.e. higher long-term price, increases the initial slope of the forward curve. As seen on the plot, the initial slope is a good determinant of Certain Added Value (\$). The solid black lines represent $X + Y = \text{constant}$, where the constant is Added Value (\$) equal to \$5, \$10, \$15 and \$20. The dashed black line is $Y = X$ line. As the initial condition becomes profitable, the uncertain portion is higher than the certain portion, in the area under the $Y = X$ line. However, the situation reverses as we move into the higher Added Value (\$) region. Under very unfavorable conditions the algorithm does not start due to lack of a profitable trade at $t = 0$. However, as soon as an opportunity for the initial trade presents itself, the trading process will begin, where the uncertain portion of profits plays a prominent role. It suggests the trader may consider breaking even on the very first trade, while being more concerned about the opportunities that may arise through future trades by letting the game to initiate.

3.4.3 Impact of the Number of Periods

Different values of N , selected according to the divisibility of the number of timesteps simulated, are chosen to study the impact of the frequency of trading on the results; 2 through 10, 12, 14, 16, 18, 20, 24, 32, 36, 63, 96, and 120. In addition to the base case storage cost, the analysis is repeated for two other values of storage cost; double the base case at $c_H = 6.57 \times 2 = 13.14$, and half of the base case at $c_H = 6.57/2 = 3.285$.

Fig. 3.9.a represents return (reward) and risk by Added Value (\$) and its standard deviation respectively; both decreasing as N increases, although at a diminishing rate. However, the trend of the $c_H = 13.14$ is different from $c_H = 6.57$ or 3.28 , where $N = 2$ or 3 are inferior to higher choices of N from a risk-return perspective. Fig. 3.9.b illustrates the Sharpe ratio defined as Added Value (\$) divided by its standard deviation. Observing that Sharpe ratio increases with N , it might be said that the decrease in the risk more than justifies the decrease in the returns. Also, an increase

in N benefits Sharpe ratio more at a lower storage cost. Ultimately, the choice of N will depend on the risk-return preferences of the trader.




(a) Added Value (\$) vs standard deviation (b) Sharpe ratio vs number of periods

Fig. 3.9. Impact of N on risk-return characteristics and Sharpe ratio; (a) Added Value (\$) and its standard deviation as a function of N for three different c_H (b) Sharpe ratio as a function of N for three different c_H . All the parameters are according to the base case as specified in Table 2.1 and Table 3.1.

The figure suggests that extra trading hurts Added Value (\$), unless the costs are very high, in which case the changes are small (except for $N = 2$ and 3). The reason is frequent trading causes earlier termination of the paths (delivery of oil), while waiting would have been more profitable on average because the likelihood of having a negative slope increases with the passage of time. Because the simulation started with $S_0 = \$54.45/\text{barrel}$ and a long-term equilibrium price of $\$103.19/\text{barrel}$, the prices tend to increase toward the long-term level, and the initial upward sloping forward curve tilts flat or downward. Therefore, a sample path would capture larger gains in the future if it remains in existence longer. It can be concluded that the algorithm performs shortsightedly by opting out for smaller gains sooner in the process. However, frequent trading has tremendous risk reduction effect, as may already be seen even by going from $N = 2$ to $N = 4$.

3.4.4 Impact of the Storage Cost

The impact of c_H is studied by considering different values of c_H from \$0 to \$20, and the analysis is repeated for $N = 2, 4, 8,$ and 16 . Fig. 3.10.a shows the Added Value (\$) versus c_H colored by maturity of the initial forward position. It shows Added Value (\$) decreases as c_H increases for all values of N . As N increases, the number of jumps increases but their size shrinks, to the point at which the graphs seem almost smooth at $N = 16$. The lower the N , the sooner Added Value (\$) touches zero due to lack a suitable maturity at t_0 . For $N = 2$, three regimes are evident in Fig. 3.10.a; (i) $c_H < 10$, (ii) $10 \leq c_H < 16$, and (iii) $16 \leq c_H$. They correspond to an initial short position in $F(0,1)$, $F(0,0.5)$, and $F(0,0)$, respectively. As c_H increases, the longer-term maturity will become unprofitable; it is seen $F(0,1)$ contract is switched to $F(0,0.5)$, where the quick drop at c_H around 10 occurs. Under regime (iii), the problem finishes trivially by selling the oil at t_0 .



(a) Added Value (\$) vs storage cost

(b) Added Value (\$) vs its standard deviation

Fig. 3.10. Impact of storage cost (c_H) on; (a) Added Value (\$) colored by maturity of the initial contract. (b) Added Value (\$) versus its standard deviation for different values of c_H , where only certain c_H labels are shown for clarity. All the parameters (other than c_H and N) are per the base case as specified in Table 2.1 and Table 3.1.

In Fig. 3.10.b, Added Value (\$) is graphed versus its standard deviation, labeled with c_H values. The general trend for $N = 4, 8,$ and 16 is that both return and risk increase as storage cost


decreases. However, for $N = 2$, the graph has two disjoint segments; in the right segment corresponding to $0 \leq c_H < 10$, as return increases, risk decreases. The left segment corresponding to $10 \leq c_H \leq 20$ has a behavior consistent with the general notion that a higher risk is accompanied with a higher reward. A change in c_H from 9.5 to 10 has a profound impact on the performance of $N = 2$ case, and causes Added Values (\$) to suddenly drop well below the other N s. Further investigations revealed that it is due to a sharp decline in Uncertain Added Value (\$); as c_H increases above 9.5, the initially selected maturity becomes 6 months shorter. So, the period for which a potential refund of c_H exists shortens considerably, which translates into a substantial loss for advancing trades, whereas the gains of postponing trades do not increase sufficiently to offset the loss.

In Fig. 3.10.a, most of the difference in results across different values of N originates from the difference in Uncertain Added Value (\$) as the variation in the Certain Added Value is limited ($< \$0.50$). Under a favorable storage cost ($c_H < 10$), a contract $F(0, T_1)$ is selected initially such that T_1 is equal or very close to one. Under these conditions, the profit is mainly made by advancing the trade when the forward curve becomes less steep or downward sloping. The payoff comprises of a “refund” of storage cost and a (positive or negative) forward spread such that the sum is positive. The price dynamics moves in general towards a better spread, and that is why waiting is on average rewarding. Thus, a smaller N leads to a higher profit due to waiting for a later time (not trading myopically soon). Under very unfavorable conditions ($c_H \geq 15$), the trading process starts by shorting $F(0, T_1)$, where $T_1 \ll 1$. In this case, most of the value is generated by the trades triggered by positive forward spreads, and this is when a larger N is beneficial because it increases the opportunity of subsequent trading. For moderate values of c_H ($10 \leq c_H < 15$), the results are mixed as the two forces compete.

To shed further light on the subject, Fig. 3.11.a show the contribution of Advancing and Postponing trades to Uncertain Added Value (\$) separately at different values of Refund Ratio. This ratio is defined as the proportion of the storage cost paid to the trader as a refund if she selects an advancing trade, i.e. she decides that she no longer needs the storage. The refund ratio introduces an asymmetric friction by penalizing the refund transaction. The ratio is set to five

values in the range of $[0,1]$, where 0 corresponds to no-refund, and 1 to full-refund. The four regimes based on the initially selected forward contract, $F(0,1)$ to $F(0,0.25)$, can be seen visibly in Fig. 3.11.a. Regardless of the Refund Ratio, advancing trades (red) contribute more than the postponing ones (blue) when $c_H \leq 10.5$. When $c_H \geq 11$, the contributions of postponing trades begin to overcome the advancing ones with a complete dominance in the $c_H \geq 14.5$ range.

If the Refund Ratio <1 , the trader receives only a partial refund. Based on the shortsighted FDO policy, the trader does not consider this fact when initially deciding to postpone the position, which may have a detrimental effect on the value. Fig. 3.11.b demonstrates the impact of storage cost (c_H) on Added Value (\$) colored by the Uncertain Added Value (\$). It illustrates that as c_H increases Added Value (\$) decreases at all refund ratios. Added Value (\$) decreases as refund



(a) Uncertain Added Value (\$) vs storage cost (b) Added Value (\$) vs storage cost

Fig. 3.11. Separate contribution of advancing or postponing trades to value, and the impact of Refund Ratio. (a) Uncertain Added Value (\$) generated by Advancing (A) or Postponing (P) trades for different values of Refund Ratio versus storage cost. (b) Added Value (\$) as a function of storage cost (c_H) for different values of Refund Ratio. All the parameters (other than c_H and Refund Ratio) are per the base case as specified in Table 2.1 and Table 3.1.

ratio drops from 1 to 0 at all c_H levels, where the difference is most pronounced in the mid-range c_H values.

Focusing on the zero-refund case in Fig. 3.11.a, any trade is solely triggered by a profitable spread. Given the initial condition and the environment dynamics, the spread is usually in favor of the advancing trades. In the first regime, the initial contract is $F(0,1)$ and most profits are due to advancing trades. The contribution of advancing trades under a zero refund (red triangle) is almost constant with respect to c_H , but it becomes increasing as Refund Ratio increases to one. Because the $F(0,1)$ position cannot be postponed further, an opportunity for postponing trades can only occur after an advancing trade has been executed. This is reflected in the trend of the postponing trades, which resembles the increasing trend of the advancing trades at any Refund Ratio.

The second regime starts with $F(0,0.75)$, where one-period postponing opportunities are possible relative to the previous $F(0,1)$ regime, while advancing opportunities are reduced by one period. This results in the step-like change in the trends of both types. Within a regime, postponing is increasingly more expensive as c_H rises, and thus a declining trend is observed. There is a link between postponing and advancing trades; as the number of one type of the trades increases (decreases), the number of the other type increases (decreases) as well because each type shifts the net short position in the opposite direction of the other type through time, which creates potential trading opportunities. However, a large enough refund ratio provides advancing trades with a big profit advantage. Accordingly, advancing trades exhibit a declining trend for low refund ratios and an increasing one for high refund ratios.

In the third regime, the conditions are like the second one, except that postponing trades are more important than before due to the shorter initial maturity ($F(0,0.5)$ instead of $F(0,0.75)$). In the fourth and last regime, postponing trade is the dominant type because the initial contract maturity is only 0.25. This regime represents an unfavorable environment, where a small number of trades happens at marginally small profits due to extremely high storage costs. Across all refund cases, even under the full refund, both types of contributions decrease as c_H rises.

3.5 Summary

We explore the trading opportunities that arise from an upward sloping forward curve using the FDO algorithm. The transparency and simplicity of FDO makes it attractive for risk-averse sellers, however, the myopic nature of the policy makes it vulnerable to underperformance depending on

the circumstances. The deficiencies in the existing FDO policy formulations are addressed and an explicit trading method using forwards is suggested. Given the pivotal role of forward curve dynamics in the storage trade, the results are interpreted in terms of the changes in mean forward prices (shifting), and forward curve slope (tilting). The impact of the curvature (bending) is not directly illustrated since within short time horizons the effect should be insignificant. The slope of the forward curve, approximated with a line, seems to explain the decisions reasonably well (there was only a narrow overlap at the decision boundaries), and thus illustrating the decisions as a function of the curvatures is not considered. Nevertheless, since all the maturities are considered in the optimization, any potential impact of the curvature is captured in decision-making. Mean forward prices does not contribute in explaining the payoff level or selected maturity because the trades are based on the spreads rather than the absolute values. The change in the realized slope through each period explains the behavior of the algorithm well. The decision-making process at each timestep can be summarized based on the slope realized at that time and the slope at the preceding time step. The impact of all previous slopes is embedded in the maturity of the contract currently held, shaping possible future choices.

Given the initial conditions, the forward curve evolution through time is mainly characterized by an increase in price levels and a decrease in the slope. Initial condition of the state variables is found to have a significant effect on the added value. We decompose the total added value into two parts, certain value (due to the trade at t_0) and uncertain value (due to the trades after t_0), where it is verified that the initial slope can approximate the certain part very well. The Added Value (\$) is maximized when the gap between a high long-term price and a low spot price increases. However, the 'Uncertain' Added Value (\$) is maximized at the highest long-term price but a moderate deviation of spot from it because a moderately-sloped forward curve can offer more future trading opportunities.

The impact of number of trading periods, N , on the added value is intertwined with the storage cost, c_H , regime. If the storage cost is low enough such that the longest maturity is selected initially, additional trading hurts; it is due to the suboptimality of FDO policy where opting for small gains sooner is preferred to waiting for larger gains in the future. If the storage cost is extremely high, a

larger N performs better because it offers a better ability in extracting uncertain added value when the initial maturity is short. However, for c_H in the mid-range, there is not a clear best performer and the difference among the results is smaller generally. It is found that changes in N strongly influences standard deviation, seen a risk measure. In conclusion, to choose an appropriate N , one should consider the regime based on the prevailing storage cost and risk-return preferences.

The contribution of advancing and postponing trades to Uncertain Added Value (\$) is studied under different Refund Ratios. Generally, the dominant type of trades influences the other one through the interlink between advancing and postponing opportunity creation. Which trade type is dominant is determined based on the regime, i.e. which initial maturity (1, 0.75, 0.5 or 0.25) is resulted by the prevailing c_H . As c_H increases, the initial maturity becomes shorter, which provides less advancing opportunities and more postponing opportunities. Therefore, moving from regime one to four, the overall level of postponing trades is increasing, while the overall level of advancing trades exhibits a decreasing trend. To measure the impact of Refund Ratio on the total Added Value (\$), the difference between full- and no-refund cases is studied; it is found to be around \$1.4 at maximum, which occurs at $c_H = 10.5$.

In Chapter 3, a myopic solution is provided by the FDO approach to solve the optimization problem resulted from the proposed trading methodology. In the next chapter, Chapter 4, we will study the optimal solution to this problem. The optimal solution will be obtained using an ADP and an exact dynamic programming technique (backward induction with nested simulations). In Chapter 4, the problem framework is the same Chapter 3, except that the trader is allowed to sell part of her inventory on the spot market while the rest of her inventory can be sold on the forward market. In addition, tests with and without the partial sale feature are conducted and compared. In fact, it will be shown computationally and theoretically that permitting the partial sale does not add value.

Chapter 4

4 Optimal Solution with Dynamic Programming

This chapter is structured as follows; the main objectives of this chapter are introduced in Section 4.1. In Section 4.2, the framework is formulated as a Markov Decision Process (MDP). Section 4.3 presents the theoretical findings about the structure of the optimal policy and value function. Section 4.4 reviews the algorithmic solutions and Section 4.5 reports the parameter values used in this chapter. In Section 4.6, the computational results, including the optimal value and policy, are discussed. Section 4.7 compares the risk and reward of the proposed cash and carry trade with those of other strategies such as a sale on the spot or a covered call position. Sensitivity of the results to the parameter estimates are studied in Section 4.8. The chapter summary is included in Section 4.9.

4.1 Introduction

In this chapter, we study the dynamic programming approach, which can provide the optimal solution to the optimization problem introduced in the previous chapter (Eq. 3.6). In this chapter, the problem framework is the same as in the previous chapter. The only difference is that the previous framework is expanded by allowing to sell the oil partially on the spot market and partially on the forward market. This dynamic cash and carry (or *contango and carry* in the context of oil trading) problem will be solved optimally using dynamic programming, and sub-optimally using FDO for comparison.

4.2 The Model

The assumptions, which are the same as in the problem set up in the previous chapter, are reviewed;

- The asset must be sold on the spot or must be hedged by a short forward position.
- Refilling the inventory is not permitted.
- Buying contracts for speculation is not permitted; long contracts may only be purchased to offset existing short contracts. Thus, if the inventory is not empty, the net exposure of all contracts entered must always be “short” with a quantity equal to the existing inventory.

- If the trader delivers the oil before the end of the term of the current rental agreement of the tanker, she will be reimbursed the unused portion of the term. At any time, it is also possible to buy additional rental time if the trader wishes to do so. So, early termination or extending the rental term can be done without any friction or penalty throughout the problem time horizon at a fixed rental rate.

Assume the trader has \bar{R} units of a commodity in storage. This inventory must be sold either via the spot market or the forward market, by time \bar{T} . The problem time horizon, $[0, \bar{T}]$, is discretized into N equidistant stages by $t_i = i\Delta t$ for $i \in \mathcal{J} = \{0, 1, 2, \dots, N\}$, where $\Delta t = \bar{T}/N$. At time t_i , the maturities $T \in \{t_i, t_{i+1}, \dots, t_N\}$ are available for the forward contract $F(t_i, T)$. At any time t_i , the portfolio of the trader consists of a long inventory position with a quantity of R_i , and a short forward contract position with a maturity of T_i and a quantity equal to the long inventory position. The inventory level can be the range of $[0, \bar{R}]$, which can be discretized uniformly into L levels by $\Delta R = \bar{R}/L$, which determines the selling batches allowed as $0, \Delta R, 2\Delta R, \dots, \bar{R}$. We consider a discrete-time dynamic optimization framework based on the following components:

1. The State Variables (x_i and W_i): The endogenous component of the state variable is $x_i = (R_i, T_i)$, expressing the inventory level and the maturity of the short forward contract (contracted delivery date). The endogenous variables only depend on operational decisions made. Stochastic factors specifying the forward curve constitute the exogenous component of the state variable $W_i = (\chi_i, \xi_i)$, which is unaffected by the decisions made. The current state is fully explained by $(x_i, W_i) \in \mathcal{X}_i \times \mathbb{R}^2$, in which \mathcal{X}_i is the state space defined as in Eq. 4.1.

$$\mathcal{X}_i = \{\{0\} \times \{0\}\} \cup \{(0, \bar{R}) \times \{t_i, t_{i+1}, t_{i+2}, \dots, t_N\}\} \quad \text{Eq. 4.1}$$

The state $x_i = (R_i = 0, T_i = 0)$ is an absorbing state in the present MDP. Practically, this corresponds to an empty inventory condition, and $T_i = 0$ is set for notational convenience because contract maturity is meaningless at an empty-tanker state. The initial state is specified by $x_0 = (R_0 = \bar{R}, T_0 = 0)$ and $W_0 = (\chi_0, \xi_0)$. Respectively, $R_0 = \bar{R}$ and $T_0 = 0$ indicate that the trader starts with a full tanker and no forward contract at hand (immediate maturity).

2. The actions (\mathbf{a}_i): At any stage $i \in \mathcal{J} \setminus \{N\}$ and state (x_i, W_i) , the decision is $a_i = (a_i^R, a_i^T) \in \mathcal{A}_i(x_i) \subseteq \mathbb{R}^2$, where a_i^R denotes the quantity of commodity to be sold in the spot market, and a_i^T refers to the new maturity of the forward contract to short the remaining inventory after transacting on the spot ($R_i - a_i^R$). Here, $\mathcal{A}_i(x_i)$ is the feasible set given by Eq. 4.2.

$$(a_i^R, a_i^T) \in \mathcal{A}_i(x_i) = \begin{cases} \{0\} \times \{0\} & \text{if } R_i = 0 \\ \{[0, R_i) \times \{t_{i+1}, t_{i+2}, t_{i+3}, \dots, t_N\}\} \cup \{R_i\} \times \{0\} \} & \text{if } R_i > 0 \end{cases} \quad \text{Eq. 4.2}$$

If $R_i = 0$, the only feasible action is $(0,0)$. If $R_i > 0$, the trader can choose to sell a quantity between $a_i^R = 0$ (do not sell) and $a_i^R = R_i$ (sell the entire inventory). The action $a_i^R = R_i$ excludes the possibility for selling a forward contract; hence $a_i = (R_i, 0)$ is feasible, as captured by the $\{R_i\} \times \{0\}$ term. For actions $a_i^R < R_i$, the trader can update the maturity of her contract to a new one chosen from $\{t_{i+1}, t_{i+2}, t_{i+3}, \dots, t_N\}$.

3. State Transition Function $f_i(x_i, \mathbf{a}_i)$: Given the current state (x_i, W_i) and an action a_i , the next endogenous state $x_{i+1} = (R_{i+1}, T_{i+1})$ is determined by the state transition function $x_{i+1} = f_i(x_i, a_i)$ defined in Eq. 4.3 to Eq. 4.5.

$$x_{i+1} = f_i(x_i, a_i) = (R_{i+1}, T_{i+1}) \quad \text{Eq. 4.3}$$

$$R_{i+1} = \begin{cases} 0 & \text{if } R_i = 0 \\ R_i - a_i^R & \text{if } R_i > 0 \end{cases} \quad \text{Eq. 4.4}$$

$$T_{i+1} = \begin{cases} 0 & \text{if } R_i = 0 \\ a_i^T & \text{if } R_i > 0 \end{cases} \quad \text{Eq. 4.5}$$

The exogenous part of the state, W_i , evolves based on the stochastic processes of Eq. 2.16 and Eq. 2.17 independently from x_i and a_i .

4. Reward Function $r_i(\mathbf{a}_i, x_i, W_i)$: Given an action $a_i = (a_i^R, a_i^T)$ at time step i , $i < N$, the stage reward includes three components generated by: (i) selling a_i^R barrels of oil on the spot, (ii) selling $R_i - a_i^R$ barrels through shorting the $F(t_i, a_i^T)$ contract, and (iii) offsetting the current contract held

by taking a long position in $F(t_i, T_i)$. Eq. 4.6 combines all these elements to express the reward function. It is assumed that $\partial r_i / \partial a_i^T$ and $\partial r_i / \partial a_i^R$ are not always zero (r_i is not a constant).

$$r_i(a_i, x_i, W_i) = a_i^R [F(t_i, t_i) - c_P - c_H(t_i - T_i)] + (R_i - a_i^R) \left[e^{-r(a_i^T - t_i)} (F(t_i, a_i^T) - c_P) - c_H(a_i^T - T_i) \right] - R_i e^{-r(T_i - t_i)} (F(t_i, T_i) - c_P) \text{ for } i \in \mathcal{J} \setminus \{N\} \quad \text{Eq. 4.6}$$

$$\text{and } r_i(x_i, W_i) = r_0 = 0, \text{ for } i = N$$

Here, c_P denotes the deterministic costs due at delivery of the oil, such as pumping cost and any location discount to WTI futures. The holding cost, denoted by a deterministic constant c_H , summarizes all the costs associated with operating the tanker including the rent, assumed payable at the start of each rental period. The payoffs from rental time adjustment assume that the rental cost is charged on a per barrel per year basis. Note that, given W_i , all the forward prices and thus $r_i(a_i, x_i, W_i)$ are deterministic. The terminal reward at the final stage $i = N$ is $r_N(x_N, W_N) = 0$. By $t_N = \bar{T}$ either the oil has been already sold, or an existing short contract with $T_N = \bar{T}$ will be fulfilled by delivering all the remaining inventory.

Lemma 4.1: The reward function r_i is neither concave nor convex with respect to $a_i = (a_i^R, a_i^T)$. The proof is provided in Appendix A.

The dynamic optimization problem of the trader to maximize the total expected reward given the initial state (x_0, W_0) is expressed by Eq. 4.7.

$$V_0(x_0, W_0) = \max_{\pi \in \Pi} E \left[\sum_{i=0}^{N-1} \delta^i r_i(A_i^\pi(x_i^\pi, W_i), x_i, W_i) \mid (x_0, W_0) \right] \quad \text{Eq. 4.7}$$

The optimization is over the class Π of all feasible policies π . A policy π is a set of decision rules $\{A_0^\pi, A_1^\pi, \dots, A_{N-1}^\pi\}$, where $A_i^\pi(x_i, W_i): \mathcal{X}_i \times \mathbb{R}^2 \rightarrow \mathcal{A}_i(x_i)$ for $\forall i \in \mathcal{J} \setminus \{N\}$. Here, $\delta = e^{-r\Delta t}$ denotes the discount factor per stage, and x_i^π refers to the random (endogenous) state at stage i when policy π is implemented. The expectation is taken with respect to the physical measure. Let $V_i(x_i, W_i)$ denote the optimal value function starting from state (x_i, W_i) . The Bellman equation associated with the problem is expressed by Eq. 4.8.

$$V_i(x_i, W_i) = \max_{a \in \mathcal{A}_i(x_i)} \{r_i(a, x_i, W_i) + \delta E[V_{i+1}(f_i(x_i, a), W_{i+1}) | W_i]\},$$

$$\forall (x_i, W_i) \in \mathcal{X}_i \times \mathbb{R}^2, \forall i \in \mathcal{J} \setminus \{N\},$$
Eq. 4.8

$$V_N(x_N, W_N) = r_N(x_N, W_N) = 0, \quad \forall (x_N, W_N) \in \mathcal{X}_N \times \mathbb{R}^2$$

The goal is to find the optimal policy, π^* that maximizes Eq. 4.7. To do so, Eq. 4.8 can be directly employed to compute all optimal value functions at all possible states by discretizing the two-dimensional domain of the continuous random variable W_i ; it requires stepping backward in time, looping through all possible states (x_i, W_i) , and searching for the optimal action. At each t_i and (x_i, W_i) , an optimal action is obtained by Eq. 4.9. This approach will lead to the Exact Dynamic Programming, which will be examined in the Algorithmic Solutions section.

$$a_i^* \in \operatorname{argmax}_{a \in \mathcal{A}_i(x_i)} \{r_i(a, x_i, W_i) + \delta E[V_{i+1}(f_i(x_i, a), W_{i+1}) | W_i]\}$$
Eq. 4.9

The high dimensionality of the state space and the required large number of conditional expectation estimation and optimizations (one per each state and each time step) renders this approach computationally prohibitive. This phenomenon is known as the curse of dimensionality for dynamic programming (Bertsekas, 2012; Powell, 2011). To overcome this phenomenon, an alternative approach, known as Approximate Dynamic Programming (ADP), will be presented in Algorithmic Solutions section.

Before finishing this section, it might be helpful to compare the problem setups between the present and the previous chapter. The following table compares the assumptions and the corresponding MDP formulations of the problem presented above and those of Chapter 3. As summarized in Table 4.1, the framework is the same except for allowing to sell the inventory partially on the spot and partially on the forward markets. The impact of this change is reflected in the MDP formulations.

Model Features	Chapter 3	Chapter 4
Initiation time	Fixed at t_0	Same
Starting inventory level	Full (\bar{R} filled at t_0)	Same
Inventory refill option	Not allowed	Same
Rental contract duration	Optimally chosen (matching the selected forward maturity)	Same
Extension, or early termination of the rental contract	Allowed	Same
Refund of storage cost if contract terminated early	Allowed	Same
Quantity basis for charging the rent	Per inventory level (R_i)	Same
Storage cost value	Known constant	Same
Partial sale on spot/forward	Not allowed	Allowed
Endogenous state variables	1. Inventory level (R_i) 2. Forward maturity (T_i)	1. Inventory level (R_i) 2. Forward maturity (T_i)
State Space, \mathcal{X}_i	$\{\{0\} \times \{0\}\} \cup \{\{\bar{R}\} \times \{t_i, t_{i+1}, t_{i+2}, \dots, t_N\}\}$	$\{\{0\} \times \{0\}\} \cup \{(0, \bar{R}) \times \{t_i, t_{i+1}, t_{i+2}, \dots, t_N\}\}$
Decision variables, a_i	1. Forward maturity (a_i)	1. Quantity sold on the spot (a_i^R) 2. Forward maturity (a_i^T)
Feasible action space, $\mathcal{A}_i(x_i)$	If $R_i = 0$: $\{0\}$ If $R_i > 0$: $\{t_i, t_{i+1}, t_{i+2}, t_{i+3}, \dots, t_N\}$	If $R_i = 0$: $\{0\} \times \{0\}$ If $R_i > 0$: $\{[0, R_i) \times \{t_{i+1}, t_{i+2}, t_{i+3}, \dots, t_N\}\} \cup \{\{R_i\} \times \{0\}\}$

Table 4.1. Comparison of assumptions and MDP formulation between Chapter 3 and 4 problems.

4.3 Theoretical Results

The following propositions and Lemma 4.II (used to prove Proposition 4.I) summarizes the structural results. The proofs are provided in Appendix A.

Lemma 4.II: The lemma has two parts;

(i) the value function can be written in form of $V_i(x_i, W_i) = R_i v_i(T_i, W_i)$ (a multiple of R_i), $\forall i \in \mathcal{J} \setminus \{N\}$, and

(ii) if $V_{i+1}(x_{i+1}, W_{i+1}) = R_{i+1}v_{i+1}(T_{i+1}, W_{i+1})$, then at stage i any action $(a_i^R, a_i^T) = (0 < a_i^R < R_i, a_i^T)$ is dominated by either $(0, a_i^T)$ or $(R_i, 0)$.

Proposition 4.I: In the SDP problem set out by Eq. 4.7 (subject to Eq. 4.1 to Eq. 4.6), partial sale of the inventory is never optimal. That is for $\forall i \in \mathcal{J} \setminus \{N\}$, any action $(a_i^R, a_i^T) = (0 < a_i^R < R_i, a_i^T) \in \mathcal{A}_i(x_i)$ is dominated by the action $(0, a_i^T)$ or $(R_i, 0)$.

Proposition 4.II: Assume that the difference between the ‘adjusted’ forward prices can be written as expressed in Eq. 4.10 (the validity of this assumption, denoted as Assumption 4.I, is examined in Appendix A). Then, the value function and optimal actions structure for $\forall i \in \mathcal{J} \setminus \{N\}$ is expressed by Eq. 4.11. Here, $E_i[\cdot]$ denotes $E[\cdot | W_i]$.

$$e^{-r(t_1-t_i)}[F(t_i, t_1) - c_P] - e^{-r(t_2-t_i)}[F(t_i, t_2) - c_P] \approx m_i(t_1 - t_2) \quad \text{Eq. 4.10}$$

$$V_i(x_i, W_i) = R_i(m_i - c_H)(t_i - T_i) + \Delta t R_i u_i,$$

$$u_i = \max\{0, A_i, B_i\} \quad \text{Eq. 4.11}$$

$$(a_i^R, a_i^T) = \begin{cases} (R_i, 0) & \text{If } A_i < 0 \text{ and } B_i < 0 \\ (0, t_N) & \text{If } B_i > A_i \text{ and } B_i > 0 \\ (0, t_{i+1}) & \text{If } A_i > B_i \text{ and } A_i > 0 \end{cases}$$

$$A_i = m_i - c_H + \delta E_i[u_{i+1}],$$

$$B_i = (N - i)(m_i - c_H) + \delta E_i[u_{i+1} - (N - i - 1)(m_{i+1} - c_H)]$$

The implication of Eq. 4.10 is that at each t_i , the expression on the left-hand side can be expressed by a line with a slope of m_i , regardless of the two maturities t_1 and t_2 . Note that for $i = N - 1$, $(0, t_N) = (0, t_{i+1})$, and $A_i = B_i$; that is the second and third arguments merge, and the maximization reduces to $\max\{0, A_i\}$. To shed some light on the intuition behind the above results, let us focus on $V_i(x_i, W_i)$ expression in Eq. 4.11; the term $R_i(m_i - c_H)(t_i - T_i)$ represents offsetting the current short contract and selling the inventory on the spot as well as the associated storage cost adjustment, i.e. the action pair $(R_i, 0)$. In the term, $\Delta t R_i u_i$, u_i weighs two other actions against $(R_i, 0)$. The three arguments of the maximum operator correspond respectively to actions

$(R_i, 0)$, $(0, t_N)$, or $(0, t_{i+1})$, which reduce the possible optimal actions to a much smaller subset of the feasible action set $\mathcal{A}_i(x_i)$.

The maximization in Eq. 4.11 states that $(0, t_{i+1})$ is chosen over $(R_i, 0)$ if $A_i > 0$; that is postponing the sale to t_{i+1} is preferred to a sell out at t_i if the total value generated from postponing the sale one period, $1 \times (m_i - c_H)$, plus the value (always non-negative) from keeping the option alive, $\delta E_i[u_{i+1}]$, is greater than zero, i.e. $m_i - c_H + \delta E_i[u_{i+1}] = A_i > 0$.

Similarly, $(0, t_N)$ is chosen over $(R_i, 0)$ if $B_i > 0$; it means postponing the sale from t_i to t_N is preferred to a sell out at t_i if the total value generated from postponing the sale $N - i$ periods, $(N - i)(m_i - c_H)$, minus the “foregone optionality” to do a similar postponing next period, $(N - i - 1)E_i[m_{i+1} - c_H]$, plus the gain from keeping the option alive, $\delta E_i[u_{i+1}]$, is greater than zero. The expression $-(N - i - 1)E_i[m_{i+1} - c_H]$ can also be thought of as the expected value from advancing the next period sale from t_N to t_{i+1} .

Finally, $(0, t_{i+1})$ is preferred to $(0, t_N)$ if $A_i > B_i$, which can be simplified to Eq. 4.12. This equation states that if it is expected that the slope differential will increase in the next stage, $(0, t_{i+1})$ is preferred to $(0, t_N)$ since the payoff from postponing the sale is proportional to the slope differential. In this case, shorting t_{i+1} maturity today places the trader in a better position tomorrow, which allows her to take advantage of the more significant long-term postponing opportunity. In other words, choosing t_{i+1} enables the trader to keep the option alive by ‘minimally reducing’ the future upside potential.

$$m_i - c_H < \delta E_i[m_{i+1} - c_H] \tag{Eq. 4.12}$$

4.4 Algorithmic Solutions

4.4.1 Exact Dynamic Programming

Table 4.2 shows the pseudocode of the exact approach assuming that the χ - ξ domain is discretized into an $H \times H$ grid. The inventory state variable R_i can be discretized to approximate the conditional expectation by simulation and sample averaging. The initial inventory \bar{R} is discretized

equally into L levels by $\Delta R = \bar{R}/L$, which determines the permissible batches to be sold as $0, \Delta R, 2\Delta R, \dots, \bar{R}$.

To achieve a computationally tractable approach, particularly for higher dimensional forward curve models, alternative approaches are considered. Even in the current setting, as will be shown in the following sections, the computational time of the exact method can approach 30 hours. Two possible avenues considered in the following are the ADP and FDO methods. The ADP approach is the focus of this chapter since FDO techniques are studied in the previous chapter and in Ghafouri and Davison (2017).

<p>1. Initialize $\hat{V}_N(x_N, W_N) = 0, \forall x_N, W_N$.</p> <p>2. For $i = (N - 1), (N - 2), (N - 3), \dots, 0$</p> <p style="padding-left: 20px;">3. For each $x_i = (R_i, T_i) \in \{\Delta R, 2\Delta R, \dots, L\Delta R\} \times \{t_i, t_{i+1}, t_{i+2}, \dots, t_N\}$</p> <p style="padding-left: 40px;">4. For each $W_i^h \in \{\chi_1, \chi_2, \dots, \chi_H\} \times \{\xi_1, \xi_2, \dots, \xi_H\}$</p> <p style="padding-left: 60px;">5.I. For each $x_{i+1} = (R_{i+1}, T_{i+1}) \in \{\Delta R, 2\Delta R, \dots, L\Delta R\} \times \{t_{i+1}, t_{i+2}, \dots, t_N\}$</p> <p style="padding-left: 80px;">5.I.a. Simulate $W_{i+1}^b, b = 1, \dots, B$, all initiated from W_i^h</p> <p style="padding-left: 80px;">5.I.b. Compute $\hat{V}_{i+1}(x_{i+1}, W_{i+1}^b)$ by bilinear or nearest neighbor interpolation of the existing $\hat{V}_{i+1}(x_{i+1}, W_{i+1})$ on χ-ξ grid</p> <p style="padding-left: 80px;">5.I.c. Estimate</p> <p style="padding-left: 100px;">$\hat{E}[\hat{V}_{i+1}(x_{i+1}, W_{i+1}) W_i^h] = \sum_{b=1}^B \hat{V}_{i+1}(x_{i+1}, W_{i+1}^b) / B$</p> <p style="padding-left: 60px;">End</p> <p style="padding-left: 40px;">5.II. Compute the optimal value function and actions by</p> <p style="padding-left: 60px;">$\hat{V}_i(x_i, W_i^h) = \max_{a \in \mathcal{A}_i(x_i)} \{r_i(a, x_i, W_i^h) + \delta \hat{E}[\hat{V}_{i+1}(f_i(x_i, a), W_{i+1}) W_i^h]\}$</p> <p style="padding-left: 60px;">$\hat{A}_i^\pi(x_i, W_i^h) = \operatorname{argmax}_{a \in \mathcal{A}_i(x_i)} \{r_i(a, x_i, W_i^h) + \delta \hat{E}[\hat{V}_{i+1}(f_i(x_i, a), W_{i+1}) W_i^h]\}$</p> <p style="padding-left: 40px;">End</p> <p style="padding-left: 20px;">End</p> <p>End</p>

Table 4.2. Pseudocode of the exact approach on a $H \times H$ grid of χ - ξ .

4.4.2 Approximate Dynamic Programming (ADP)

Approximate Dynamic Programming (ADP) is a broad group of algorithmic strategies and modelling techniques that offers several methodologies for tackling the curses of dimensionality in large, multiperiod, stochastic (or deterministic) optimization problems (Powell, 2011). An ADP technique based on the Least Square Monte Carlo (LSM) approach (Carriere, 1996; Longstaff and Schwartz, 2001; Tsitsiklis and Roy, 2001) is used here. This method offers a cost-effective approach to compute the lower bound of $V_0(x_0, W_0)$ (Glasserman, 2003, sec. 8.7; Nadarajah et al., 2017). Approximation architectures are employed to estimate the expectation of the value function, also known as the *continuation function*, as used in Nadarajah et al. (2017) consistent with the pioneering LSM works mentioned above. For each x_{i+1} possible, a continuation value approximation is adopted. The approximation assumes that the continuation value is a linear combination of K basis functions of the exogenous part of the state at stage i , i.e. W_i , as expressed by Eq. 4.13. Here, A' denotes the transpose of A .

$$E[V_{i+1}(x_{i+1}, W_{i+1}) | W_i] = \Theta'(i, x_{i+1})\Phi(W_i) = \sum_{k=1}^K \theta_k(i, x_{i+1}) \varphi_k(W_i), \quad \forall i, x_{i+1} \quad \text{Eq. 4.13}$$

Polynomials are very common in linear approximation architectures (Lagoudakis and Parr, 2003; Longstaff and Schwartz, 2001). Given the two-dimensional domain of W_i , considering the polynomials of degree three leads to $K = 10$ basis functions according to Eq. 4.13. These ten basis functions are summarized in the vector $\Phi(W_i)$ defined by Eq. 4.14.

$$\Phi'(W_i) = [1 \quad \chi_i \quad \xi_i \quad \chi_i^2 \quad \xi_i^2 \quad \chi_i \xi_i \quad \chi_i^3 \quad \xi_i^3 \quad \chi_i^2 \xi_i \quad \chi_i \xi_i^2], \quad \forall i \quad \text{Eq. 4.14}$$

Determining the continuation value is reduced to estimating the vector of weights $\Theta(i, x_{i+1})$, which is found by the least squares regression. Table 4.3 shows the pseudocode of the ADP approach. In step **1**, M price paths are simulated. In step **2**, initialization is done using the fact that, at stage N , the value function is zero for all states (by the deterministic reward function). Step **3** includes the loop moving stage-wise backward in time. In step **4**, $\hat{E}[\hat{V}_{i+1}(x_{i+1}, W_{i+1}) | W_i^m]$ is computed for all possible states in the next stage, i.e. x_{i+1} . The continuation value for any x_i and a_i is computed based on $x_{i+1} = f_i(x_i, a_i)$. This comprised of solving for the regression coefficients in step **5.I**,

and computing the estimated continuation value (fitted value of the regression) in step **5.II**. Finally, in step **6**, the optimal action for each state x_i and path m is computed.

<p>1. Simulate M sample paths of the W_i process for $i = 0, 1, 2, \dots, N$; denoted by $\{W_i^m\}_{m=1}^M$.</p> <p>2. Initialize $\hat{V}_N(x_N, W_N) = 0, \forall x_N, W_N$.</p> <p>3. For $i = (N - 1), (N - 2), (N - 3), \dots, 1, 0^*$</p> <p style="padding-left: 2em;">4. For each $x_{i+1} = (R_{i+1}, T_{i+1}) \in \{\Delta R, 2\Delta R, \dots, L\Delta R\} \times \{t_{i+1}, t_{i+2}, \dots, t_N\}$</p> <p style="padding-left: 4em;">5.a Compute the regression coefficients, $\hat{\Theta}(i, x_{i+1})$, using the M sample paths</p> $\hat{V}_{i+1}(x_{i+1}, W_{i+1}^m) \sim \sum_{k=1}^K \theta_k(i, x_{i+1}) \varphi_k(W_i^m), \quad m = 1, \dots, M$ <p style="padding-left: 4em;">5.b. Compute the CFA (as the fitted value of the regression) using $\hat{\Theta}(i, x_{i+1})$</p> $\hat{E}[\hat{V}_{i+1}(x_{i+1}, W_{i+1}) W_i^m] = \sum_{k=1}^K \hat{\theta}_k(i, x_{i+1}) \varphi_k(W_i^m)$ <p>End</p> <p>6. For each $x_i = (R_i, T_i) \in \{\Delta R, 2\Delta R, \dots, L\Delta R\} \times \{t_i, t_{i+1}, t_{i+2}, \dots, t_N\}$, compute the optimal value function and actions by</p> $\hat{V}_i(x_i, W_i^m) = \max_{a \in \mathcal{A}_i(x_i)} \{r_i(a, x_i, W_i^m) + \delta \hat{E}[\hat{V}_{i+1}(f_i(x_i, a), W_{i+1}) W_i^m]\}$ $\hat{A}_i^\pi(x_i, W_i^m) = \operatorname{argmax}_{a \in \mathcal{A}_i(x_i)} \{r_i(a, x_i, W_i^m) + \delta \hat{E}[\hat{V}_{i+1}(f_i(x_i, a), W_{i+1}) W_i^m]\}$ <p>End</p> <p>End</p> <p>*At $i = 0$, the regression will be replaced with a sample average due to the absence of multiple sample paths at $t = 0$, which means $\hat{E}[\hat{V}_1(x_1, W_1) W_0] = \frac{1}{M} \sum_{m=1}^M \hat{V}_1(x_1, W_1^m)$.</p>

Table 4.3. Pseudocode of ADP (LSM) approach.

4.4.3 Forward Dynamic Optimization (FDO)

As discussed in detail in the previous chapter, the myopic decision rule adopted by FDO just maximizes the immediate reward but ignores any corresponding change in the continuation value. The treatment of the problem in Eq. 4.7 by the FDO strategy is reviewed as a reminder since Chapter 3 is devoted to this method. FDO policy is expressed by its decision rule in Eq. 4.15. The

value is generated by following these actions as expressed by Eq. 4.16. FDO presents a (suboptimal) solution by sequentially maximizing the reward at each timestep while moving forward in time.

$$A_i^{FDO}(x_i, W_i) = \operatorname{argmax}_{a_i \in \mathcal{A}_i(x_i)} r_i(a_i, x_i, W_i) \quad \text{Eq. 4.15}$$

$$V_0^{FDO}(x_0, W_0) = E \left[\sum_{i=0}^{N-1} \delta^i \max_{a_i \in \mathcal{A}_i(x_i)} r_i(a_i, x_i, W_i) \mid (x_0, W_0) \right] \quad \text{Eq. 4.16}$$

4.5 Parameters

In the following numerical simulation, the price model parameters are exactly as before, provided in Table 2.1. We generated $M = 100000$ (50K + 50K antithetic) price paths in the ADP method by simulating the state variables with $\Delta t = 1/480$ year using the Schwartz and Smith (2000) model, Eq. 2.16, Eq. 2.17, and Eq. 2.18. This allows price series discretized on different time intervals (Δt s) to be extracted.

The initial conditions and time horizon also match those defined in the base case in Table 3.1. Recall that the specified W_0 corresponds to a spot price of \$54.45 and a long-term price of \$103.19 simulating forward prices based on May 2009 market conditions, which was a favorable period for this type of trade (Diaz-Rainey et al., 2017; Kemp, 2016). Based on actual time-charter rates of a 2 million barrel VLCC around May 2009 (Ghafouri and Davison, 2017), the tanker rent is assumed to be about \$36,000/day, equivalent to the chosen $c_H = \$6.57$ per barrel per year, and higher than the \$3.5 per barrel per year assumed by Jafarizadeh and Bratvold (2013). The Pumping Cost is assumed to be $c_P = \$3.75$ /barrel, which results in a total cost of \$10.32/barrel, and higher than the \$8.5/barrel cost assumed by Jafarizadeh and Bratvold (2013). The parameters for optimal policy analysis using the ADP algorithm, as well as a comparison between the ADP and the FDO are set based on Case A of Table 4.4. In addition, Case B is defined to provide a basis for comparison between the exact and the ADP algorithms, as well as studying the impact of Propositions 4.I and 4.II. The reason behind introducing Case B ($N = 16$) is that using Case A ($N = 60$) in the exact dynamic programming approach will be very computationally expensive.

4.6 Computational Results

In the following subsections, the numerical results studying several aspects of the problem are presented. Unless stated otherwise, all the results are based on the out-of-sample lower-bound (or more accurately a downward-biased) estimate of the value function at time zero, denoted by V_0 . It should be noted that all the computations in this thesis are performed on a desktop computer with an i7-6700@3.41GHz CPU.

Description	Parameter	Case A	Case B
Initial inventory (barrels)	\bar{R}	1	same as A
Time Horizon (constraint)	\bar{T}	1 year	same as A
Number of time stages ($\Delta t = \bar{T}/N$)	N	60	16
Storage discretization increment	ΔR	1	1, and 1/3
Storage cost	c_H	\$6.57	same as A
Pumping cost	c_P	\$3.75	same as A
Initial condition of the exogenous state variables	$W_0 = (\chi_0, \xi_0)$	(-0.639, 4.637)	same as A
Initial condition of the endogenous state variables	$x_0 = (R_0, T_0)$	(1, 0)	same as A
Total number of simulated antithetic paths in the ADP algorithm (Table 4.3)	M	100,000	same as A
Total number of simulated antithetic paths for out-of-sample estimation	M_2	10,000	same as A
Total number of simulated antithetic W_{i+1}^b in the exact algorithm (Table 4.2)	B	N/A	200
Number of grids points in χ - ξ domain for the exact algorithm (Table 4.2)	H	N/A	43

Table 4.4. Problem parameters defined as Case A and Case B for the exact or ADP analysis.

4.6.1 The Optimal Value

The lower-bound estimate of the optimal value is investigated in this section. The impact of considering Propositions 4.I and 4.II is examined by limiting the partial sale and feasible action set respectively. If one believes that the partial sale is not optimal, ΔR can be set to $\bar{R} = 1$, which avoids any inventory discretization, and thus partial sale. For comparison, a $\Delta R = 1/3$ is also tested. Proposition 4.II is implemented by limiting the feasible actions to the following subset of $\mathcal{A}_i(x_i)$.

$$\mathcal{A}_i^{II}(x_i) := \{(R_i, 0), (0, t_N), (0, t_{i+1})\} \quad \text{Eq. 4.17}$$

In Table 4.5, the performance of ADP algorithm is compared to that of FDO using Case A of the parameters. The value generated by ADP is $V_0 = \$10.70$ (per barrel) with a standard deviation of \$4.58 (standard error of \$0.014), which is equivalent to \$21.4 million for the (2 million barrels) VLCC tanker. The FDO algorithm generates $V_0^{FDO} = \$7.85$. In the ADP case, using $M_2 = 10K$ instead of $M_2 = 100K$ changes the estimate less than 2%, while reducing the computational time by 50%. In FDO case, it is even more computationally beneficial to use $M_2 = 10K$. The V_0 or V_0^{FDO} is the ‘added value’ which the trader captures by following the corresponding policy ‘relative to’ selling her inventory on the spot at t_0 . While the ADP algorithm generates 36% more value than FDO, its computation time is 17 (1.12/0.066) times longer. Employing $\mathcal{A}_i^{II}(x_i)$ can reduce the ADP computation time by about 80% to only 3.4 (0.23/0.066) times longer than that of FDO.

Method	Number of Paths (M)	Number of Out-of-Sample Paths (M_2)	Mean (\$)	Standard Deviation (\$)	Computation Time (hours)
ADP	100K	100K	10.70	4.58	2.08
ADP	100K	10K	10.68	4.56	1.12
ADP using $\mathcal{A}_i^{II}(x_i)$	100K	10K	10.68	4.56	0.23
FDO	100K	N/A	7.854	0.989	0.69
FDO	10K	N/A	7.848	0.986	0.066

Table 4.5. Optimal value and computational time of the FDO and ADP approaches using Case A of the parameters.

Table 4.6 shows the optimal values and computational times of the exact and ADP algorithms using Case B of the parameters. The number of time stages, N , of Case B is smaller than that of Case A, which allows to examine the computationally expensive exact algorithms, and the impact of partial sales. The optimal value generated by the exact and ADP methods differ only about 1%, which validates the results. However, the fastest exact case still takes a longer time than the slowest ADP one. Comparing the two exact variants, it is observed that the (slightly slower) Bilinear variant leads to a slightly higher value estimate than the nearest neighbor approach.

Although there is a small difference among all the computed mean values in Table 4.6, as they fall in the range \$10.60-\$10.70, there are significant differences among the computation times; the computationally slowest case is about 940 (30.10/0.032) times slower than the fastest one. The

computation time variation is due to three factors; (i) the algorithm; ADP vs exact (two variations), (ii) partial sale consideration ($\Delta R = 1$ or $1/3$), and (iii) limiting the feasible action set to $\mathcal{A}_i^{II}(x_i)$. Focusing on one factor and keeping all other factors the same, the ADP algorithm takes between 39-170 times (1.21/0.032-30.10/0.17) less computation time than the exact method. Also, factor (ii) can decrease the computation time between 2-6 times (0.063/0.032-28.2/4.64), while factor (iii) can decrease it between 1.7-10 times (0.055/0.032-28.2/2.75).

Method	Interpolation Technique	Mean (\$)	Standard Deviation (\$)	Storage Discretization ΔR (Proposition 4.I)	Feasible Set Limited to $\mathcal{A}_i^{II}(x_i)$? (Proposition 4.II)	Computation Time (hours)
Exact	Nearest Neighbor	10.63	4.33	1/3	NO	28.20
Exact	Nearest Neighbor	10.60	4.32	1	NO	4.64
Exact	Nearest Neighbor	10.68	4.40	1/3	YES	2.75
Exact	Nearest Neighbor	10.70	4.43	1	YES	1.21
Exact	Bilinear	10.66	4.39	1/3	NO	30.10
Exact	Bilinear	10.65	4.41	1	NO	5.12
Exact	Bilinear	10.70	4.46	1/3	YES	3.16
Exact	Bilinear	10.70	4.47	1	YES	1.24
	ADP	10.64	4.39	1/3	NO	0.17
	ADP	10.64	4.39	1	NO	0.055
	ADP	10.63	4.38	1/3	YES	0.063
	ADP	10.64	4.39	1	YES	0.032

Table 4.6. Optimal value and computational time of exact and ADP approaches using Case B of the parameters.

These numerical experiments corroborate Proposition 4.I's result that the opportunity of partial sale creates no additional value for this problem over either selling no inventory or selling all inventory. In other words, the optimal decision with respect to the maturity, a_i^T , is applied to all the inventory. The reason might be that the whole inventory can be emptied during one timestep without being limited by any constraints, outside those considered here, such as a maximum pumping rate or an illiquid market. Also, this numerical experiment verified Proposition 4.II (and

Assumption 4.I by extension) by showing that limiting the feasible set to $\mathcal{A}_i^{II}(x_i)$ does not change the values.

To estimate a confidence interval for V_0 using the ADP method, the computations are repeated for 150 times using different sample and out-of-sample paths. Fig. 4.1 shows the histogram of such computed values. The mean from 150 repeated simulations is \$10.678 and the 95% confidence intervals is [\\$10.674, \$10.682], which confirms that the values found earlier are within the confidence bounds.




Fig. 4.1. Histogram of value using the ADP method computed for 150 times. The solid line represents the mean of the histogram. Case A parameters ($\Delta R = 1$) as per Table 4.4 are used.

4.6.2 The Optimal Policy

The optimal policy is shown on Fig. 4.2 and Fig. 4.3 by illustrating the evolution of the (χ_i, ξ_i) state variables through time on the $\chi_i-\xi_i$ plane, and the associated optimal decisions $a_i = (a_i^R, a_i^T)$ overlying on that plane. Fig. 4.2 and Fig. 4.3 show the optimal actions a_i obtained by the two methods respectively at $t = 0.25$ and $t = 0.75$. For comparison purposes, the optimal decisions for both the exact and ADP methods are overlaid; the ADP results appears in the dense cluster of simulated paths, which becomes more dispersed as time passes, while the exact method results are reflected on the domain rectangular grid. In addition to the initial condition set per Case B, a

different initial condition, $W_0 = (-1.2, 4.2)$, is also used with the ADP method, which causes the cluster of the simulated paths to move toward the bottom-left of the cluster generated by Case B. This allows further exploration of the ADP policy's responsiveness.

Comparing optimal values and computational times of ADP and the exact approach indicates that the ADP method gives the best and fastest result. The downside is the optimal policy and value function are “locally” calculated, i.e. centered around the evolution of the initial condition. If another initial condition is to be considered, the computation must be repeated in the ADP method, whereas the exact approach has already solved the program for all initial conditions within the χ - ξ domain selected from the beginning.

Comparing Fig. 4.2 to Fig. 4.3, as time passes, the exercise boundary (red-blue boundary within the cluster of simulated points) moves from right to left, which corresponds to a flattening forward curve, as will be seen later. The realizations on the right side of the line indicate a sale on the spot decision, i.e. $(a_i^R, a_i^T) = (R_i, 0)$, and most of the ones on the left show a hold decision, i.e. $(a_i^R, a_i^T) = (0, a_i^T)$, where $a_i^T = 1$ across most of the realized domain, and in $a_i^T = t_{i+1}$ in a small region for the ADP. The points for which $a_i^T = t_{i+1}$, are in a small minority relative to the paths simulated in the ADP method (<1% of the 100K). The results of the exact method show that there is a (different) region in the χ - ξ plane, where this apparently rare policy suggested by the ADP method is optimal.

The interim policies suggested by the ADP and the exact method should not be compared globally with one another since the ADP provides a solution based on the initial condition, which is accurate in the area around the evolution of the stochastic factors. Fig. 4.2 and Fig. 4.3 show that the optimal policies suggested by the two methods often agree. However, as one deviates from the densely-populated areas and gets closer to the extremities, the regression results become weak and thus the obtained policies are not dependable. For instance, consider the minority region (colored in orange) representing an optimal choice $a_i^T = 0.8125 (= 0.75 + \Delta t)$ in Fig. 4.3. A discrepancy is observed between the location of the a_i^T policy prescribed by the two methods; the optimal choice of $a_i^T = 0.8125$ indicated by the exact method is seen in a region situated left of the area suggested by ADP. We believe that this is the result of poor regression approximation in an area far from most




Fig. 4.2. Optimal decision a_i^R (left) and a_i^T (right) at $t_i = 0.25$ ($i = 4$) and state $x_i = (R_i, T_i) = (1, 0.25)$. The parameters are per Case B in Table 4.4 ($M = 100K$, $N = 16$, $\Delta R = 1/3$, $H = 43$). Different $W_0 = (\chi_0, \xi_0)$ are used with the ADP; $W_0 = (-0.639, 4.637)$ per Case B (top), and $W_0 = (-1.2, 4.2)$ (bottom).

data points used to build the regression coefficients (the results are best where they matter the most). Also, the magnitude of the difference in payoffs between choices $a_i = (0, 1)$ and $(1, 0)$ considered over all paths is much larger (~ 10 times) than that of choices $a_i = (0, 1)$ and $(0, 0.8125)$. We believe this is another reason that when searching for the optimal action $a \in$




Fig. 4.3. Optimal decision a_i^R (left) and a_i^T (right) at $t_i = 0.75$ ($i = 12$) and state $x_i = (R_i, T_i) = (1, 0.75)$. The parameters are per Case B in Table 4.4 ($M = 100K$, $N = 16$, $\Delta R = 1/3$, $H = 43$). Different $W_0 = (\chi_0, \xi_0)$ are used with the ADP; $W_0 = (-0.639, 4.637)$ per Case B (top), and $W_0 = (-1.2, 4.2)$ (bottom).

$\mathcal{A}_i(x_i)$ in the ADP algorithm (Step 6 in Table 4.3), a small error can lead to an incorrect decision. However, because the proportion of these points are very small (<1%) they don't impact the results significantly, as confirmed by the similar values in Table 4.6. Given the initial W_0 , the general theme of the optimal policy is to short $F(0,1)$ at stage $i = 0$, corresponding to the decision $a_i =$

$(a_i^R, a_i^T) = (0,1)$. Then the decision is to hold this contract until all inventory is sold on the spot at some point depending on the (χ_i, ξ_i) realization. In other words, on most simulated paths the optimal decision, (a_i^R, a_i^T) , changes from $(0,1)$ to $(1,0)$. As highlighted in the theoretical discussion of the optimal policy in Section 4.3, the optimally selected maturity, a_i^T , is not always equal to 1 or 0. This can be seen in Fig. 4.2 and Fig. 4.3, respectively depicted in light blue and orange. This is consistent with the theoretical results of the Proposition 4.II, where at time stage i , $a_i^T = t_{i+1}$ may be optimal under certain conditions.

4.6.3 Comparison of ADP and FDO

Fig. 4.4 shows the histogram of the value generated using the same set of sample paths for the ADP and FDO methods. The optimal ADP policy generates $V_0 = \$10.71$ on average, which can be broken down into two parts; (i) \$6.19 generated by selling the oil forward using $F(0,1)$, and (ii) \$4.52 obtained by all of the subsequent trades. At t_0 , part (i) is known (certain), whereas part (ii) is uncertain. Although most of the value is captured by the first part, the subsequent trades are necessary to capture the remaining 42% of the value. The contribution of the subsequent trades to value will increase if the initial forward curve is less steep, which is a less favorable environment to start the trading. Thus, a trader still has an incentive to start although the certain part of the added value is not very high. Also, Fig. 4.4 shows that the lower- and upper-end of the outcomes generated by the ADP algorithm are respectively lower and higher than those given by the FDO; a range of \$1.44-\$34.13 vs \$6.19-\$16.38 is observed for ADP and FDO respectively. So, the higher value of the ADP algorithm comes at the price of higher volatility and risk (i.e., standard deviation).

Fig. 4.4.b displays the histogram of the time stage at which the algorithm reaches the absorbing state $x_i = (R_i = 0, T_i = 0)$, where the inventory is sold and trading terminates. It is seen that ADP is more patient than FDO; while there is a peak in the first half on the FDO histogram, most ADP paths indicate that deliveries occur during the second half (at the end) of the time horizon.

To compare the detailed performance of ADP and FDO on the same path, two representative paths are studied here; Fig. 4.5 depicts a sample path in which ADP performs better than FDO, whereas

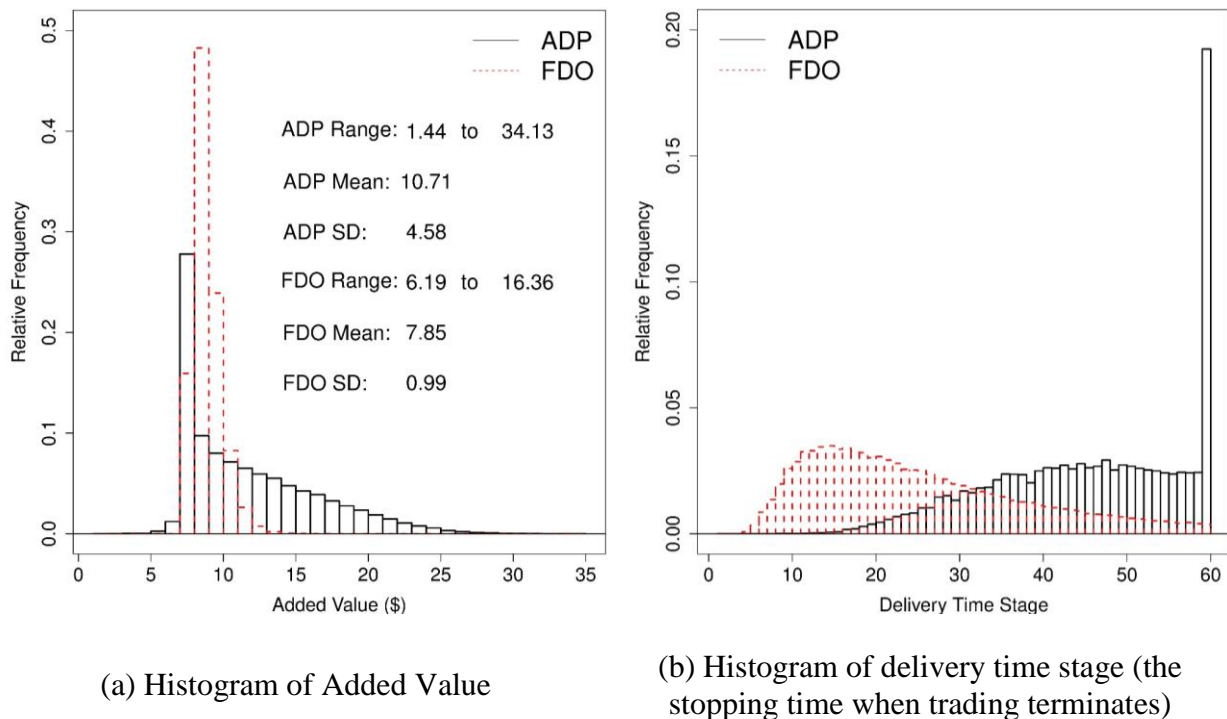
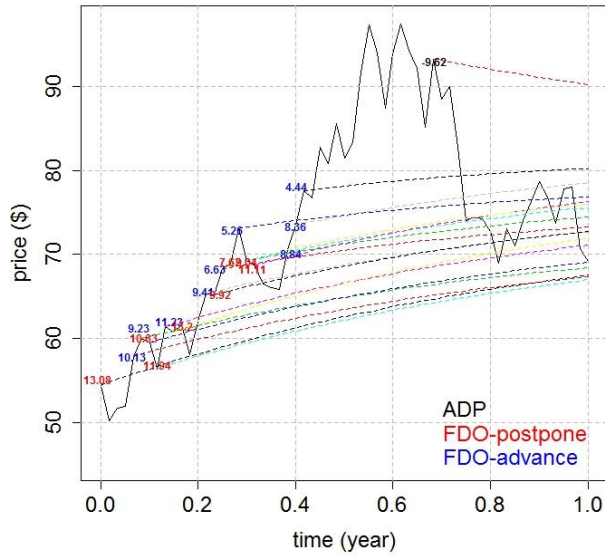
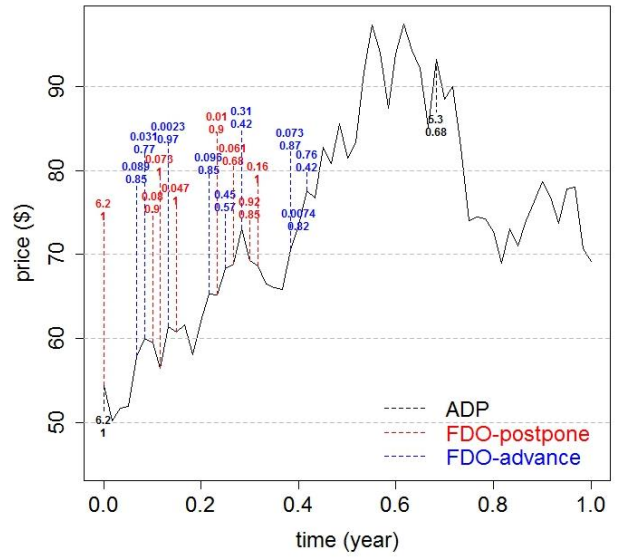


Fig. 4.4. Results of simulating $M_2 = 100K$ new paths using Case A parameters and comparing ADP vs FDO statistics.

Fig. 4.6 exhibits another sample path on which FDO performs better. In both figures, panel (a) shows the forward curve including its slope, when a trade happens, while panel (b) shows the incremental gain and the maturity of the newly taken short position. In Fig. 4.5, ADP waits until $t = 0.68$ year, at which time there is a sufficiently negative slope of -9.62 , and it cashes out the position with an incremental payoff of $\$5.3$. However, in Fig. 4.6, ADP can only collect $\$1.2$ by cashing out at $t = 0.9$ year, whereas the incremental gains of FDO sums up to $\$3.65$ by the time it terminates trading a $t = 0.3$ year. It should be noted that for both paths, the two strategies generate $\$6.20$ from the initial $F(0,1)$ contract and the differences arise from the subsequent trading decisions.

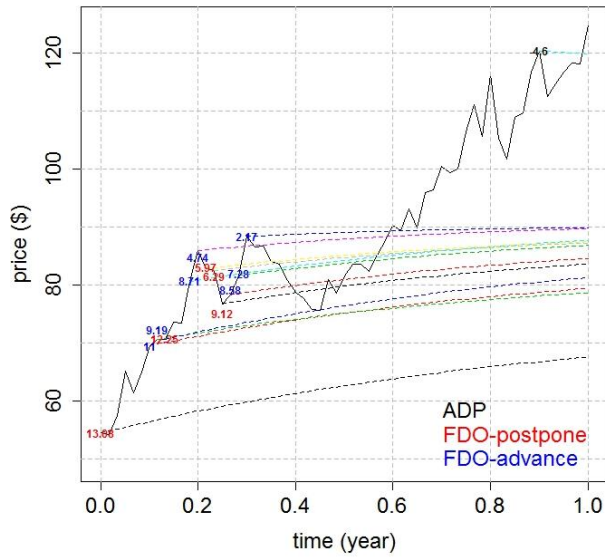


(a) Forward curve and its slope (label) at each trading time

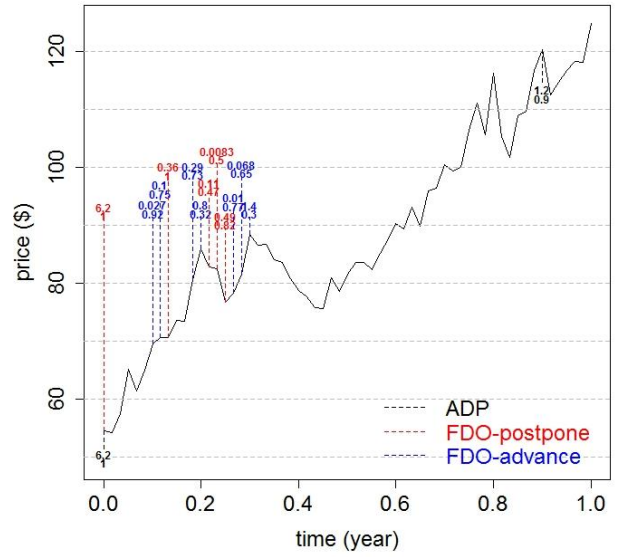


(b) Generated value (top label) and chosen maturity (bottom label) at each trading time

Fig. 4.5. Evolution of the spot price and performance comparison on a sample path, where $V_0^{FDO} < V_0^{ADP} = \11.44



(a) forward curve and its slope (label) at each trading time



(b) Generated value (top label) and chosen maturity (bottom label) at each trading time

Fig. 4.6. Evolution of the spot price and performance comparison on a sample path, where $V_0^{ADP} < V_0^{FDO} = \9.84 .


4.6.4 Mapping the Decisions

To explain the financial intuition behind the optimal decisions, the estimated optimal decisions on the $\chi_t\text{-}\xi_t$ plane are mapped into the $S_t\text{-}b_t$ plane, where S_t is the spot price, and b_t is the slope defined as $(F(t, \bar{T}) - S_t)/(\bar{T} - t)$. The results in this section are based on the ADP algorithm using Case B of parameters. The results are shown in Fig. 4.7 at $t_i = 0.125, 0.375, 0.875,$ and 0.9375 , corresponding to $i = 2, 6, 14,$ and 15 . The red line represents the $y = c_H (= 6.75)$ line. For any given S_t , the optimal decision most often prescribes a sale on the spot when the slope is below (a typically negative) threshold. To analyze the subject more thoroughly, three sources of value affecting the decision are identified as follows.


The first element of the value, V_I , is generated by the immediate reward from the trade, in which the oil is sold at S_t and bought at $F(t, \bar{T})$ to offset the existing contract, i.e. $-(F(t, \bar{T}) - S_t) = -b_t(\bar{T} - t)$. Because V_I is generated by a trade, the payoff can be positive (profit) or negative (loss). For instance, a negative slope of b_t translates into a profit of $V_I = -b_t(\bar{T} - t)$, which decreases as 't' increases assuming b_t remains constant. The correlation between S_t and b_t is usually negative, as low spot prices tend to be occurring with steep forward curves, and vice versa. The second element, $V_{II} \geq 0$, is the refund of the storage cost for the time remaining, $\bar{T} - t$. As the oil is sold on the spot, this amounts to $c_H(\bar{T} - t)$ dollars ($c_H = \$6.57$), which also decreases with time. The third element, $V_{III} \geq 0$, is the continuation value foregone as the trader terminates the position, which is indeed decreasing with time. Due to the geometric nature of the price process, the option value is decreasing as the spot price decreases. It is worth noting that the FDO approach only includes V_I since it is a myopic method.

The trader is indifferent between selling the inventory completely, $a_i^R = 1$, or holding the inventory, $a_i^R = 0$, right at the (red-blue) boundary in Fig. 4.7. Quantitatively speaking, the continuation value, V_{III} , is equal to the sum of the immediate profits or losses, V_I , and storage cost refund, V_{II} . This is demonstrated on panel (b) of Fig. 4.7 equivalently as $V_{III}/(1 - t) = V_I/(1 - t) + V_{II}/(1 - t)$, which can be reformulated as $V_{III}/(1 - t) = -b_t + c_H$. Now, two interesting questions can be answered. The first question is why the boundary is decreasing in S_t . Note that


the refund of storage cost does not change with S_t . However, when S_t is very small, the continuation value is very low, i.e. small $V_{III}/(1-t)$. Thus, the trader can afford to exercise at ‘small’ negative slopes, i.e. small $-b_t > 0$. She can even exercise at small positive slopes (i.e. at a losing trade) if she still makes a profit by receiving the storage cost refund, i.e. small enough $-b_t < 0$ such that $-b_t + c_H > 0$. As S_t increases, the continuation value increases. Therefore, the trader requires a larger immediate reward to exercise at a higher S_t , which in turn implies a larger (negative) slope.




(a) Optimal policy a_i^R (left) and a_i^T (right) at timestep $i = 2$



(b) Optimal policy a_i^R (left) and a_i^T (right) at timestep $i = 6$



(c) Optimal policy a_i^R (left) and a_i^T (right) at timestep $i = 14$



(d) Optimal policy a_i^R (left) and a_i^T (right) at timestep $i = 15$

Fig. 4.7. Optimal decision (inventory on the left, and maturity on the right) in terms of spot price (S_t) and slope $((F(t, T) - S_t)/(T - t))$ at times 0.125, 0.375, 0.875, and 0.9375, corresponding to time stages $i = 2, 6, 14,$ and 15 . The red line represents $y = c_H$ line.

The second question is why the boundary moves upward with the passage of time. As time passes, V_{III} and even $V_{III}/(1 - t)$ decreases, and thus, all else being equal, the boundary moves up since

$V_{III}/(1 - t) = -b_t + c_H$. Since V_{III} is larger on the right side compared to the left (due to a larger S_t), the right side moves up at a faster rate than the left. As shown in panel (d), at the last timestep at which a decision can still be made, $t_i = 0.9375$ ($i = 15$), there is no continuation value ($V_{III} = 0$). Therefore, a positive slope b_t (a loss) is acceptable if $-b_t + c_H = 0$ holds. This means the optimal boundary is the line $y = c_H$, which is well approximated by the numerical simulation.

4.7 Comparative Characteristics of the Present Trade

The performance of the ADP approach is compared with several other strategies in this section. Note that the FDO method provides a myopic solution to the same liquidation problem that is solved optimally by the ADP approach, where both techniques provide policies to hedge and liquidate the inventory using forward contracts. To put the characteristics of trading in the forward markets into perspective, the total value derived from FDO and ADP is compared with other alternatives; (i) sell the inventory on the spot market at t_0 , (ii) the static cash and carry arbitrage, which is to sell the inventory forward using the most profitable forward contract at t_0 , in this case $F(0,1)$, (iii) a *covered call* strategy, (iv) a *protective put* strategy, and (v) a strategy based on selling the inventory optimally on the spot market.

Note that each of these strategies, except for selling forward using $F(0,1)$, does not provide a constant hedge of the inventory (as required by the cash and carry arbitrage) similar to what is offered by FDO or ADP. While covered call and protective put strategies could potentially reduce the risk, the optimal sale on the spot strategy does not benefit from any risk reduction. To set up the covered call strategy, the trader (inventory owner) receives the premium by shorting a call option, and delivers the inventory if the call owner decides to exercise. If the call is not exercised at all, the trader liquidates the inventory at the terminal time. To set up the protective put strategy, the trader buys a put option, and she sells the inventory by selling it on the spot market directly or by exercising the put, whichever that is more profitable.

Both the call and put options are assumed to have a strike price of $K = S_0 = \$54.45$ and an expiry of one year. Also, both options are assumed to be American style, which allows the owner to exercise at any time during $[0, \bar{T}]$, which is similar to the timeframe during which the trader is

allowed to take action in the original problem (based on the forward trading). The price of the call and the put are respectively computed to be \$16.51 and \$3.71 using a least-squares Monte Carlo (by simulating the prices under the risk-neutral measure). In all strategies, there is the pumping cost due at delivery, and the storage cost for one year paid at t_0 . If the inventory is delivered at t_i , there will be a refund of the storage cost, $(\bar{T} - t_i)c_H$.

In the covered call position, the call owner decides when to exercise the option, and assuming this occurs at t_i , the payoff to the trader (inventory owner) is $K - c_P + (\bar{T} - t_i)c_H$. If he does not exercise the call before it expires at \bar{T} , the trader liquidates the inventory at \bar{T} , which generates $S_{\bar{T}} - c_P$.

In the protective put position, the trader owns both the inventory and the option. She decides when to sell her inventory and whether she exercises her put. Specifically, selling the inventory at t_i yields $\max\{K, S_{t_i}\} - c_P + (\bar{T} - t_i)c_H$, where the first argument (K) represents a sale by exercising the put, and the second argument (S_{t_i}) characterizes a sale of the inventory directly on the spot market. Both the covered call and protective put positions are solved by employing an approximate dynamic programming method similar to the ADP approach used in the original problem.

In what follows the results and intuition of the problem, cast as a Markov Decision Process (MDP), are reviewed. The corresponding elements of the MDPs are summarized in Table 4.7. The MDP formulation allows the strategies to be evaluated using an ADP method based on the least-square Monte Carlo. The parameters are based on Case B in Table 4.4 ($\Delta R = 1$).

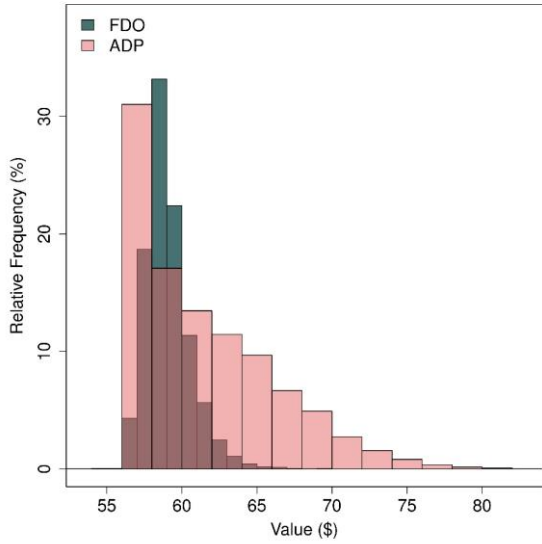
Fig. 4.8 shows the histogram of values resulted from implementing different strategies. Table 4.8 summarizes the histogram information and compares the return and risk features of the above methods. For consistency, all the results are computed using the same paths; the same sample set to build the policy and the same out-of-sample set to test the performance. The difference between the price dynamics under the physical and risk-neutral measure as captured by $\mathbb{E}_0^P[S_{\bar{T}}] = \94.61 and $\mathbb{E}_0^Q[S_{\bar{T}}] = e^{r\bar{T}}F(0, \bar{T}) = \67.88 is clearly evident. There is a large upside in spot prices from

MDP Element	Covered Call	Protective Put
Endogenous state variables	$x_i = R_i$ (inventory level)	$x_i = R_i$ (inventory level)
State Space ($x_i \in \mathcal{X}_i$)	$\mathcal{X}_i = \{0, \bar{R}\}$	$\mathcal{X}_i = \{0, \bar{R}\}$
Exogenous state variables	$W_i = (\chi_i, \xi_i) \in \mathbb{R}^2$	$W_i = (\chi_i, \xi_i) \in \mathbb{R}^2$
Decision variable	Call owner decides to exercise ($a_i = 1$) or not to exercise ($a_i = 0$)	Inventory (and put) owner decides to sell the inventory ($a_i = 1$) or not to sell ($a_i = 0$). She always chooses the higher payoff between exercising the put and selling directly on the spot market.
Feasible action space	$\mathcal{A}_i(x_i) = \begin{cases} \{0\} & \text{If } R_i = 0 \\ \{0, 1\} & \text{If } R_i = \bar{R} \end{cases}$	$\mathcal{A}_i(x_i) = \begin{cases} \{0\} & \text{If } R_i = 0 \\ \{0, 1\} & \text{If } R_i = \bar{R} \end{cases}$
State transition function $x_{i+1} = f_i(x_i, a_i)$	$R_{i+1} = \begin{cases} 0 & \text{if } R_i = 0 \\ R_i \mathbb{I}(a_i = 0) & \text{if } R_i = \bar{R} \end{cases}$	$R_{i+1} = \begin{cases} 0 & \text{if } R_i = 0 \\ R_i \mathbb{I}(a_i = 0) & \text{if } R_i = \bar{R} \end{cases}$
Reward function (for computing the optimal policy of the decision maker)	$r_i(a_i, x_i, W_i) = a_i R_i \max\{\exp(\chi_i + \xi_i) - K, 0\}$	$r_i(a_i, x_i, W_i) = a_i R_i [\max\{K, \exp(\chi_i + \xi_i)\} - c_p + (\bar{T} - t_i)c_H]$
Auxiliary reward function to compute the payoff to the covered call position which does not have any influence on decision-making	$r_i^{\text{aux}}(a_i, x_i, W_i) = a_i R_i [K - c_p + (\bar{T} - t_i)c_H]$ for $i = 0, 1, 2, \dots, N - 1$ $r_i^{\text{aux}}(a_i, x_i, W_i) = \mathbb{I}(\exp(\chi_i + \xi_i) < K) R_i [\exp(\chi_i + \xi_i) - c_p]$ for $i = N$	Not Applicable


Table 4.7. The elements of the MDP describing the covered call and protective put positions.

which to benefit if the inventory was carried unhedged into the future; this is reflected in the histograms and high values achieved under the protective put and spot sale strategies. While the protective put position provides an insurance against price falls, it can gain from the upside potential fully. It is seen that the cost of the put (\$3.71) was not entirely recovered when the value of the protective put is compared to the sale at S_t since adverse events did not occur frequently. On the other hand, in the covered call position, the call is exercised very often which leads to the

spike in the histogram of the covered call position.



(a) Histogram of FDO and ADP



(a) Histogram of Spot Sale, Protective Put, and Covered Call

Fig. 4.8. Histogram of the values obtained on each path using different strategies. The histograms are based on the same set of 10K out-of-sample paths. The parameters are based on Case B in Table 4.4 ($\Delta R = 1$).

Method	Sell on the spot at S_0	Sell forward using $F(0,1)$	Forward Hedging		Option Hedging		No Hedge; Sell optimally at S_t
			FDO	ADP	Covered Call	Protective Put	
Total value (\$)	50.70	56.89	59.06	61.50	60.11	80.53	83.92
Extra value relative to selling at S_0	0.00	6.19	8.36	10.80	9.41	29.83	33.22
99% VaR	50.70	56.89	56.89	56.89	50.07	40.16	33.53
95% VaR	50.70	56.89	57.08	56.89	60.38	41.44	44.44
Range of values	50.70	56.89	[56.8, 69.46]	[53.48, 83.14]	[30.10, 64.64]	[40.16, 231.08]	[14.63, 234.80]

Table 4.8. Comparing the return and risk characteristics of different strategies. The parameters are based on Case B in Table 4.4 ($\Delta R = 1$).

In ADP and FDO, the trader has a short forward position and, in the covered call strategy, she has a short option position. The ADP method generates a higher value relative to the covered call. It could be due to the dynamic updating of the short forward position in the ADP approach compared to maintaining the same short call position in the latter strategy. From an uncertainty perspective, focusing on the 99% value-at-Risk (VaR), ADP ties with FDO, both having 99% VaR of \$56.89. The covered call strategy achieves the second best 99% VaR with a value of \$50.07. It is not surprising that the spot sale strategy has the lowest 99% VaR.

To test the performance of the strategies under a different and unfavorable initial condition, i.e. a downward-sloping initial forward curve, the computations are repeated for $W_0 = (\chi_0, \xi_0) = (-0.2, 4.2)$ corresponding to $S_0 = \$54.60$ and a long-term price of \$66.69. The results are summarized in Table 4.9. In this case, selling forward statically at $F(0, T)$, $\forall T \in (0, \bar{T}]$, is not profitable; it can be seen in the table that selling at $F(0, 1)$ leads to a lower value than selling at S_0 . For the same reason, the FDO strategy does not start trading and it sells the inventory at t_0 . Among the four remaining methods, the ADP ranks first in terms of risk, while selling optimally on the spot price generates the highest value. In summary, if the price starts from a negative χ_0 , i.e. it is initially deviated below the long-term trend, there might be an incentive to engage into a trade

Method	Sell on the spot at S_0	Sell forward using $F(0, 1)$	Forward Hedging		Option Hedging		No Hedge; Sell optimally at S_t
			FDO (Rolling Intrinsic)	ADP	Covered Call	Protective Put	
Total value (\$)	50.85	40.05	50.85	51.44	48.29	54.15	61.32
99% VaR	50.85	40.05	50.85	40.15	28.82	34.74	22.81
95% VaR	50.85	40.05	50.85	42.74	36.69	34.74	30.84

Table 4.9. Unfavorable condition for the storage trade; downward-sloping initial forward curve induced by the stochastic factors initial condition set at $W_0 = (\chi_0, \xi_0) = (-0.2, 4.2)$, while the rest of the parameters are based on Case B in Table 4.4 ($\Delta R = 1$). There is not any maturity $T \in (0, \bar{T}]$ to set up a profitable short $F(0, T)$ contract, and thus the static cash and carry and FDO strategies opt out to sell at t_0 .


postponing the sale of the inventory. However, the degree of risk and the expected profit depends both on the strategy and on the initial condition χ_0 . Note that $S_0 = \$54.49$ and $\$54.60$ in Table 4.8 and Table 4.9 respectively. However, ADP achieves $\$61.50$ and $\$51.44$ under the two scenarios respectively although the spot prices are very close. Comparing $(\chi_0, \xi_0) = (-0.639, 4.637)$ and $(\chi_0, \xi_0) = (-0.2, 4.20)$, the difference in values can be attributed to a deeper initial deviation, χ_0 , and a higher long-term price, ξ_0 , in the former case.

4.8 Sensitivity Analysis


Fig. 4.9 shows sensitivity of the estimated V_0 by illustrating the percent change in V_0 in response to the percent change in a parameter, i.e. $100(x - x_0)/x_0$ with x_0 denoting the initial value of the parameter x . ADP method and Case B of parameters in Table 4.4 with $\Delta R = 1$ is used as the benchmark for the relative assessment.

Among the stochastic factors in Fig. 4.9.a, the result is most sensitive to k and λ_χ , and least sensitive to μ_ξ and $\rho_{\chi\xi}$. Generally, the sensitivity to the short-term factor (χ) parameters is higher than to the long-term factor (ξ), with the exception of the initial condition. Moreover, the only two parameters with a decreasing trend are λ_χ and λ_ξ ; this results from their negative sign as part of the drift term of the risk-neutral SDE's, as seen in Eq. 2.5 and Eq. 2.6. Fig. 4.9.b exhibits increasing and decreasing response to changes in \bar{T} and c_H , respectively, which are consistent with the optionality of value with respect to time horizon and cost. Fig. 4.9.d indicates that the effect of the number of Monte Carlo paths M is insignificant as long as it is greater than about 75K. The number of time stages, N , influence is mainly limited to its lower range, and V_0 does not change drastically for $N > 16$ (the default value in Case B).


Fig. 4.9, panels (b) and (c) indicate that the value is more sensitive to changes in ξ_0 than χ_0 since the former sets the long-term price level exponentially whereas the latter is the (temporary and mean-reverting) short-term deviation from this long-term price. From $F(t, T)$ as expressed in Eq. 2.13, the slope of the curve at time t between maturities T_1 and T_2 is computed by Eq. 4.18. This explains the observed sensitivity of value with respect to ξ_0 ; which occurs because the (initial) slope increases exponentially with ξ_0 .




(a)



(b)



(c)



(d)


Fig. 4.9. % change in V_0 versus % change in different parameters; (a) stochastic factors parameters, (b) time horizon \bar{T} , storage cost c_H , and initial condition (χ_0, ξ_0) , (c) this panel completes panel 'b' by further extending the y-axis further, (d) number of paths M , and time stages N .

$$\frac{F(t, T_2) - F(t, T_1)}{T_2 - T_1} =$$

Eq. 4.18

$$\frac{e^{\xi t}}{T_2 - T_1} \left\{ e^{\chi t \exp[-k(T_2-t)] + A(T_2-t)} - e^{\chi t \exp[-k(T_1-t)] + A(T_1-t)} \right\}$$

Although the impact of ξ_0 is more significant than that of χ_0 , from a factor volatility standpoint, σ_χ is more important than σ_ξ for generating value. Further investigation via numerical simulation shows that a high σ_ξ /low σ_χ translates into a much lower volatility in the slope and curvature of the forward curve compared to a low σ_ξ /high σ_χ . The reason is volatility in the spot price is more



(a) $(\sigma_\chi, \sigma_\xi) = (0.70, 0.07)$

(b) $(\sigma_\chi, \sigma_\xi) = (0.07, 0.70)$

Fig. 4.10. Two sample realizations of the simulated spot price, $S(t)$, and forward curve, $F(t, T)$, with monthly increments base on parameters in Table 2.1 except $(\sigma_\chi, \sigma_\xi) = (0.70, 0.07)$ in (a), and $(\sigma_\chi, \sigma_\xi) = (0.07, 0.70)$ in (b).

important than volatility in the long-term price in driving changes in forward curve slope and/or curvature. Because σ_χ affects the near-end of the curve, its impact is immediately seen in a forward curve switching between contango and backwardation through time. However, σ_ξ affects the price at a distant future point, the impact of which is not significant on the near-end of the curve. The

concept is illustrated in Fig. 4.10, in which two realizations with $\sigma_\xi = 0.07/\sigma_\chi = 0.70$ (left), and $\sigma_\xi = 0.70/\sigma_\chi = 0.07$ (right) are compared.

4.9 Summary

A profit-maximizing dynamic cash and carry arbitrage problem is formulated as a Markov Decision Process (MDP) (Puterman, 2005). The trader decides on a two-dimensional action (a^R, a^T) ; a^R reflects the ‘quantity’ of oil to be sold on the spot, and a^T denotes the ‘maturity’ of the short contract hedging the remaining inventory. Unlike many studies limited to the liquidation on the spot, the second decision variable allows the trader to benefit from the forward market while optimizing over the contract maturity. Simultaneous optimization of the storage management and the financial contract specification, i.e. maturity, is the salient characteristic of this problem.

To investigate the optimal liquidation of a storable commodity, the full term-structure of the forward curve is utilized in the MDP both informationally and as a trading instrument. It is shown that the stage reward function is neither convex nor concave. It is also proved that a partial sale, splitting the quantity between the spot and forward markets, is not optimal. This result does not depend on the underlying stochastic model of forward prices. Moreover, under certain assumptions, it is established that optimal actions are restricted to a small subset of the feasible set: (i) spot sale, (ii) short the forward maturing at the next timestep, or (iii) short the forward maturing at the end of time horizon. Subsequently, the above theoretical propositions are verified using algorithmic solutions, and their significant impact on the computation times are demonstrated.

An approximate dynamic programming (ADP) approach based on Continuation Function Approximation (CFA) with the Least-squares Monte Carlo (LSM) is developed. This approach is benchmarked by the exact optimal solution computed from discretizing the exogenous state space. Both the estimated optimal values and the associated optimal policy obtained using the two methods are in good agreement, except in a limited region in which they differ only with respect to the maturity decision. The errors are believed to be due to the approximation introduced by CFA and the regression process. However, the discrepancy does not greatly impact the estimated optimal value since it only occurs among some minority outliers (<1% of the paths) in the ADP

approach. The computation times show that ADP is at least 39 times faster than the exact method. Moreover, characterizing the decision boundary in the domain of the slope of forward curve and the spot price highlights the critical role of the slope in the optimal action selection.

Furthermore, the ADP method is compared with the simple suboptimal Forward Dynamic Optimization (FDO) approach. While ADP values found to be 36% higher than that of the FDO, it is at least 3.4 times slower. From a risk perspective, the FDO method offers a lower standard deviation, a narrower range, and a guaranteed limited down-side. The histogram of the stopping times shows that the ADP method performs more patiently than FDO in liquidation. To sum up, the added expected profit of the ADP-generated strategies over the myopic FDO strategies comes at a cost of increased risk, whether that risk is measured via standard deviation, range of outcomes or downside. One may consider both methods in an initial analysis, as both can be easily computed using the provided algorithms, since the result will also depend on market details and the way it is parameterized.

Chapter 5

5 A Trading Model Considering Stochastic Storage Costs

In this chapter, a more realistic framework relative to Chapter 4 is introduced. Although the same forward curve model as Chapter 4 is used, a model for stochastic tanker rental costs (independent of the forward model, as will be justified) is added. Nevertheless, taking a comprehensive view in Chapter 5, it contains the solutions based on both stochastic and constant storage cost (the constant cost is similar to Chapter 4 but within the updated framework).

The corresponding optimization problems are solved by a similar methodology. However, a new proposition improving the continuation value estimation (Proposition 5.I) is proved and investigations using the ADP method are performed. It is found that – for the same parameters as used before – while the stochastic storage cost may add value relative to the constant storage cost, the new framework leads to a less value compared to Chapter 4 model due to the non-refundability of storage cost. The changes in the framework and the storage cost does indeed alter the optimal behavior relative to the Chapter 4 model.

The rest of this chapter is organized as follows; Section 5.1 reviews the key assumptions of Chapter 4 framework and presents the suggested modifications. Section 5.2 provides a general introduction on the oil tanker vessels. Section 5.3 reviews the literature on the tanker freight rates. Section 5.4 presents the details of the model used here for stochastic storage costs. Section 5.5 covers formulation of the new framework as a Markov Decision Process (MDP). Section 5.6 states Proposition 5.I, which lays the foundation necessary to build the algorithmic solution (ADP) presented in Section 5.7. General parameters of the problem and experiments establishing the benchmark computational parameters (e.g. number of simulated paths) are reported in Section 5.8. The computational results are studied in Section 5.9, followed by a chapter summary in Section 5.10.

5.1 Introduction

A set of assumptions has been used in the previous chapter. Each assumption and suggestions for improvement are reviewed in the following. This will lead to the new model that will be utilized in this chapter, which reflects a more realistic framework.

► *The trader has a full tanker at t_0 , which implies that she has already started the cash and carry trade, and the problem focus is how to liquidate optimally.*

The problem has been structured based on the assumption that the trade is initiated at some fixed point in time with known initial conditions. Then we determine the optimal policy and calculate the value generated given the initial condition. However, a very important practical question is when to start the trade. Considering the “trade initiation” is more consistent with the practice, where the trader monitors the oil and the tanker rental market and decides optimally when to initiate the trade. This indeed requires a model for stochastic storage costs which leads to the next point.

► *The storage cost throughout the problem time horizon is fixed (not varying with time).*

In a simple setup, in which a trader starts the trade at a fixed point in time, assuming a fixed storage cost would be reasonable if there is no need to refer to the tanker rates in the future. However, tanker rates do fluctuate, and incorporating a variable one allows to study many dynamic aspects of the problem.

► *If the trader decides to return the tanker early, she will receive a pro-rata refund based on the remaining time. The terms of the rental contract may only allow for a fraction of the pro-rata refund.*

This assumption allows the trader to terminate the rental contract and to recover her costs if she decides to sell the inventory earlier than the end of the rental contract. However, it can be imagined that the tanker owner may not accept to pay a (full) refund based on the original rate because, for instance, the tanker rates or demand for them have decreased. This can be (deterministically)

addressed through refunding only a fraction of the pro-rata amount, where the fraction can be 0% representing a worst-case scenario. Although considering a variable tanker rate can improve the issue from a modeling standpoint, availability of the refund may not be still a good assumption from a realistic standpoint. For instance, the SEC (Securities and Exchange Commission) filings of an oil shipping company with a fleet of VLCC tankers state the circumstances under which charterers may terminate charter contracts early: “The events or occurrences that will cause a charter [contract] to terminate or give the charterer the option to terminate the charter generally include a total or constructive total loss of the related vessel, the requisition¹ for hire of the related vessel, the vessel becoming subject to seizure for more than a specified number of days or the failure of the related vessel to meet specified performance criteria.” (“SEC Amendment No.5 to Form F-1,” n.d.). Therefore, it may be more accurate to exclude any refunds from an early termination of the contract.

► *In case of a partial sale of the inventory, the refund of the storage cost will be based on the quantity sold, which implies that the storage cost is charged per barrel per year.*

Even if the refund provision is excluded, it is more realistic to compute the storage cost based on a per tanker per year setting, rather than per barrel per year. This means that as long as there exists a physical or forward position with a non-zero quantity (even smaller than the tanker capacity), the trader has to rent the ‘whole’ tanker because this type of storage facility cannot be shared simultaneously among several users.

In summary, the above adjustments lead to the new set of assumptions defining a novel framework; the trader monitors the oil prices and (stochastic) tanker rates, and decides when to initiate the trade subject to a time horizon. The duration of the rental contract is fixed, and there is no refund if the contract is terminated early. As soon as the tanker is rented, she can select to take a position

¹ Compulsory acquisition of the vessel by states during wartime, where a state forces the owner to charter the vessel to the state at a dictated hire rate.

in the oil market, or wait for another time while she continues to monitor the oil term structure. Therefore, the times to rent a tanker and to initiate the oil trading are chosen optimally and independently (the latter time must be greater than or equal to the former one). Furthermore, the trader has the option to refill the tanker multiple times within the contract duration. Therefore, one key question is how to model the storage costs.

5.2 Oil Tanker Vessels Background

Before developing a stochastic model, we provide a short introduction to tanker markets. Freight rate is the cost of transporting one barrel of oil from port A to port B. One important type of oil tanker vessels used for this purpose is the VLCC (Very Large Crude Carrier) with a capacity of 200,000 to 320,000 DWT (Dead Weight Tonnage). DWT is the maximum permitted weight of the sum of the weights of the vessel, cargo, fuel, fresh water, crew, etc. in tons. Another important class is the Suezmax; it is named after the Suez Canal and characterizes the largest allowable tanker to pass via the canal. The Suezmax vessels, as mid-sized tankers, have a capacity in the range of 120,000 to 200,000 DWT. The most standard capacities in barrels of oil are about 2 and 1 million barrels for VLCC (260000 DWT) and Suezmax (130000 DWT) tankers respectively (the conversion depends slightly on the specific gravity of the oil). There exist the smaller Panamax vessels with a capacity in the range of 60,000 to 80,000 DWT, which are consistent with the size limits of the Panama Canal.

There are two main types of agreements to employ a tanker; Time Charter (TC) and voyage charter contracts. Based on a TC contract the vessel is hired for a specific period, while the vessel owner manages the vessel. However, the charterer directs the vessel which ports to visit and in what order. The charterer is responsible for fuel costs consumed by the vessel, port charges, a daily hire of the vessel crew, and commissions. TC agreements exist for 3, 6, and 9 months as well as 1, 3, 5, and 7 years. Of these the 1-year contract is the most commonly traded one. Usually, TC agreements are traded about one month before chartering commences. TC contracts are also used by tanker owners as risk management tools because when a tanker is chartered for a term, the owners are not faced with the fluctuations in the spot market.

Although TC agreements seem like the most suitable type of arrangement for the present problem, they are not the most frequently used ones. The most common type of contracts involves a ‘voyage’ because tankers are mainly used for transportation (rather than as floating storage). In a voyage charter the charterer pays the vessel owner on a lump-sum, or on a per-ton (or per barrel) basis, known as the freight rate. In return, the vessel owner is responsible for the voyage between a loading and a discharging port, as well as the port, fuel and crew costs. This freight rate is the basis for the spot prices in the tanker market. It is not surprising that most of the research in the marine transportation literature is concerned with the spot freight rate rather than TC rate. The two main methods to negotiate/quote the spot freight rate are WS and TCE as explained in the following.

The Worldwide Tanker Normal Freight Scale, also known as ‘Worldscale’ (WS) is a baseline rate, which provides a convenient way to negotiate the freight across many routes. The ‘Worldscale’ index or WS 100 is updated annually based on the preceding year costs, and serves as a basis to compute tanker spot freight rates. WS 100, quoted in \$/(metric ton), is the cost of transporting a metric ton of cargo using the standard vessel on a round-trip voyage on each tanker route. In a negotiation, the charterers and tanker owners agree on some percentage of WS. For instance, if a charterer and an owner agree on WS 80, then 80% of the published WS 100 on the corresponding route will be the contracted price. This payment includes all related costs like fuel, crew, and port costs, however, the cost of insurance is not included.

The other measure for quoting the spot prices is Time Charter Equivalent (TCE) rate. This rate, quoted in \$/day, is a measure showing the operating performance of a vessel in terms of daily revenues, and is computed by the gross revenue of the tanker minus the expenses (port, canal, and fuel costs) all based on a particular round-trip voyage divided by the round-trip duration in days. So, the main distinction of TCE and WS is that TCE does not include voyage expenses, which are collected by the ship owner, but are merely payments for fuel and port expenses made on behalf of the charterer. Also, it can be said that if the vessel does not incur any depreciation (capital expenses), TCE will be the break-even hire rate. TCE, like WS, is a spot rate measure. However, one can compare TCE rate with Time Charter (TC) rates, which reflects expectations about the future spot freight rates (prices). For instance, if TCE rates are high and the tanker owner expects

TCE rates to increase, she would prefer to charter out the tanker on the spot market rather than using a TC agreement.

One of the important indices in the tanker market is the Baltic Dirty Tanker Index (BDTI), which represents the freight rates of vessels carrying mainly crude oil and other lower distillates of oil such as fuel oil, and is reported daily on the Baltic Exchange based on the settled voyage charter agreements. The routes include tankers sized from VLCC and Suezmax to other smaller sizes. Table 5.1 lists the 18 routes (sub-indices) underlying the BDTI index. In addition to the departure and destination ports, each route specifies the tanker class and capacity (in DWT) to standardize each sub-index. Considering a specific sub-index, e.g. TD5, it is interesting to note that the tanker class, e.g. Suezmax, does not necessarily mean that the route passes through the geographical region implied by the class, e.g. Suez Canal.

Code	Cargo	From	To	Size (DWT)	Class
TD1	Crude	Persian Gulf	US Gulf	280,000	VLCC
TD2	Crude	Persian Gulf	Singapore	260,000	VLCC
TD3	Crude	Persian Gulf	Japan	250,000	VLCC
TD4	Crude	West Africa	US Gulf	260,000	VLCC
TD5	Crude and/or Dirty Products Heat 135F	West Africa	US Atlantic Coast	130,000	Suezmax
TD6	Crude	Black Sea	Mediterranean	135,000	Suezmax
TD7	Crude	North Sea	Continental Europe	80,000	Aframax
TD8	Crude	Kuwait	Singapore	80,000	Aframax
TD9	Crude	Caribbean	US Gulf	70,000	Panamax
TD10D	Fuel Oil (double hull)	Caribbean	US Atlantic Coast	50,000	Panamax
TD11	Crude	Mediterranean	Mediterranean	80,000	Aframax
TD12	Fuel Oil	Antwerp	US Gulf	55,000	Panamax
TD14	No-heat Crude	South East Asia	East Coast Australia	80,000	Aframax
TD15	No-heat Crude	West Africa	China	260,000	VLCC
TD16	Fuel Oil Heat 135F	Black Sea	Mediterranean	30,000	Handysize
TD17	Crude	Baltic	UK or Continental Europe	100,000	Aframax
TD18	Crude	Baltic	UK or Continental Europe	30,000	Handysize

Table 5.1 Baltic Dirty Tanker Index (BDTI) composition as of November 2008. All vessel must have oil necessary approvals.

From a physical trade point of view, the most notable routes are TD3, TD5, and TD7 (Alizadeh and Nomikos, 2009), which are shown on Fig. 5.1. The TD3 sub-index corresponds to the price of a voyage from Persian Gulf (Ras Tanura, Saudi Arabia) to Japan using a VLCC (250,000 DWT) tanker. The TD5 sub-index corresponds to a Suezmax (130,000 DWT) tanker class route from Bonny (Nigeria) to Philadelphia, and the TD7 sub-index represents an Aframax (80,000 DWT) class route from Sullom Voe (North Sea, UK) to Wilhelmshaven (Germany).




Fig. 5.1. Approximate tanker routes; (i) TD3 from 1 (Ras Tanura, Saudi Arabia) to 2 (Chiba, Japan), (ii) TD5 from 3 (Bonny, Nigeria) to 4 (Philadelphia, USA), (iii) TD7 from 5 (Sullom Voe, UK) to 6 (Wilhelmshaven, Germany).

5.3 Literature on Tanker Freight Rates

It is intuitively reasonable that the shape of the forward curve would be related to the cost of tanker rental, as discussed earlier. A first step to developing a stochastic storage cost model is to review the existing literature relating the crude oil market to the freight rate. It is then followed by the freight rate modeling leading to a variable time charter rate for the tanker.

Alizadeh and Nomikos (2004) explore the causality and arbitrage opportunity between WTI futures contracts and physical crude oil in both the Brent and Bonny physical markets. They use

the fact that six different types of imported crude oil can be used in delivery against the WTI futures contracts. To make this cross-market delivery, they consider the freight rates on the corresponding routes between the physical market and Cushing OK, which is the delivery point for WTI futures contracts. They find that WTI futures Granger¹ cause freight rates. Also, the level of freight rates does not impact WTI futures, Brent, or Bonny spot prices. More importantly, they find that the spread between WTI futures and either Brent or Bonny spot price “does not affect” the freight rates, which is the key to their proposed cash-and-carry arbitrage. They explain that cash-and-carry arbitrage is not impacted by the convenience yield. That is because this type of arbitrage occurs when the futures spread is more than the full cost of carry, which is the cost of buying on the spot at time t_1 and delivering it at time $t_2 > t_1$. This condition automatically implies the absence of convenience yield. Alizadeh and Nomikos (2004) provide many instances of evidence where the arbitrage trade exists using back-testing the historical data. The type of trade they consider is to buy Brent or Bonny low on the spot physically, and short WTI futures at a higher price with a maturity of 3 to 5 weeks, during which the oil is transported from its origin to Cushing, Oklahoma. To put this trade into perspective within the strategies discussed in the present research, this is a one-time forward maximization trade without any maturity adjustment or quantity flexibility. Also, the proposed trades are studied through backtesting, where unlike the present work, there is no stochastic model for simulation.

Many studies try to investigate the relation between the spot freight rate and underlying factors such as the oil market. Poulakidas and Joutz (2009) studied the relation between tanker spot prices (West Africa to US Gulf route), WTI spot, WTI Nymex 3-month futures contract, and the amount of oil in inventories, for the period 1998-2006. They find that a high 3-month futures spread, or a low inventory put upward pressure on the tanker spot prices. The correlation between the spot freight rates and oil prices can be ambiguously positive or negative, which is due to demand and supply of oil (Glen and Martin, 2004). If the demand for oil increases, both oil prices and demand

¹ A time series U is said to Granger cause V if lagged values of U contain information that helps predict V beyond the information contained in lagged values of V alone (Granger, 1969).

for tankers increase, which creates a positive correlation. However, if the supply of oil decreases, oil prices increase, while demand for oil transportation decreases. This in turn generates a downward pressure on expected spot freight rates and hence a negative correlation with oil prices is created. Sun et al (2014) investigate the relation between spot freight rates, represented by Baltic Dirty Tanker Index (BDTI), and oil prices. Using Ensemble Empirical Model Decomposition, they decompose each time series into a high-frequency, a low-frequency, and a residual term characterizing the long-term trend. They notice that the decomposition is necessary to explain the correlations of multiscale components of freight rates and oil prices, and they find strong correlation between the two in medium and long term when the relevance structure is taken into account.

Yang et al (2015) study the spillover effect from the crude oil market on the tanker market using data from 2006 to 2014. They find that volatility in the Brent market has a more pronounced impact on the tanker market than the volatility in WTI market. Also, they find that they can classify their sample into two subsamples; June 2006 – April 2009, and April 2009 – April 2014, where the impact of the oil markets on the tanker market is stronger in the former than the latter subsample. They explain the difference via unexpected demand for oil and scarcity of tankers in the first subsample, which prevented a new tanker market equilibrium from being reached rapidly. In the second subsample, excess capacity and intense competition characterizes the tanker market and freight rates.

Shi et al (2013) studied the effect of crude oil price on the tanker market using a structural vector autoregressive (VAR) model. They used monthly time series data from 2002 to 2011, which include oil production in the world (barrel/day), averaged price of WTI and Brent as a proxy for crude oil prices, and Baltic Dirty Tanker Index (BDTI) reflecting the crude oil tanker spot levels. Using impulse response analysis, they investigate the effect of different shocks to the tanker market; as for contemporaneous relationship, while crude oil supply shocks have significant effect on the tanker market, non-supply shocks impact is insignificant. This finding is in agreement with Alizadeh and Nomikos (2004), which finds no evidence indicating that the tanker freight rates are related to the spread between WTI futures and Brent or Bonny physical spot prices, which was the

key to the cash-and-carry arbitrage. They also mention that the effect of crude oil price shocks on the tanker freight market is limited.

Adland et al (2016) study the impact of charterer and owner heterogeneity as well as charterer-owner match effects using 2863 VLCC fixtures (transaction contracts) between 2011 and 2014 reported by brokerage houses. Their empirical method is based on the estimation of fixed effect models and variance decomposition. They find that although market conditions and routes are the two main influential factors on the spot freight rates, characteristics of charterers, owner, and of their matches are also crucial. This is important since it highlights that the spot freight rates are influenced by many determinants and the market conditions and routes alone are not sufficient to explain them.

In conclusion from the above literature, it can be stated that although there is a relation between the freight rates and oil prices, it is heavily impacted by many other factors such as the route (freight index), the geographical location of the oil market, (world) oil inventory levels, supply and demand forces, short-term versus long-term perspective, heterogeneity in the contracts (transaction-specificity), and time period considerations. Since considering all these components explicitly in a model is impractical, representing the freight prices in a standalone model, i.e. one that incorporates the effect of all underlying determinants implicitly, might be justified. The model can still benefit from a multifactor stochastic structure.

There is several research to model the shipping spot freight rates in a continuous-time framework. Usually, the continuous-time models imply a simpler freight price dynamics compared to the empirical discrete-time models (Alizadeh and Nomikos, 2009). For instance, geometric Brownian motion (Koekebakker et al., 2007), and Ornstein-Uhlenbeck (Jørgensen and de, 2010; Sødal et al., 2008, 2009) processes have been used to model the freight spot prices. These have been used in vessel valuation in a real option context (Sødal et al., 2008; Tvedt, 1997) as well as pricing derivatives based on the freight rate (Jørgensen and de, 2010; Koekebakker et al., 2007).

More complicated models incorporating stochastic volatility model of Barndorff-Nielsen and Shephard (2001), or even using Levy process based dynamics are suggested by others (Benth et

al., 2015). It should be mentioned that the freight service as an asset is non-storable, and thus a simple cost-of-carry calculation cannot be applied to value the associated forward contracts, which makes it necessary to use other valuation methods. Also, the non-storability of the asset makes valuing and trading forward contracts crucial because of the need to manage delivery of the asset at a future time efficiently. Taib (2016) derives the price of forward freight contracts using different stochastic models for the underlying freight spot rates (prices) by employing the spot-forward risk-neutral relationship, where she considered the freight spot models studied by Benth et al. (2015) as the underlying models.

Prokopczuk (2011) tests the one-factor models of Black (1976) and Schwartz (1997), and the two-factor models of Schwartz and Smith (2000) and Korn (2005), as well as the three-factor model of Cortazar and Naranjo (2006) in the dry bulk freight market using the futures contracts data from Imarex¹ Freight Futures Market, and finds that the Schwartz and Smith (2000) model provides the best results.

5.4 Stochastic Storage Cost Modeling

In line with the Schwartz and Smith (2000), Mirantes et al. (2012) developed a four-factor stochastic model for natural gas, where the logarithm of the freight spot price is equal to the sum of a long-term factor, a short-term one, and a seasonal factor, where the seasonal factor is complemented by a fourth factor. Poblacion (2015) estimates the four-factor model of Mirantes et al. (2012) using two sets of maritime data, and shows that there is stochastic seasonality in the freight rate dynamics. For the present research, the four-factor stochastic model of Mirantes et al. (2012) will be adapted because of the evidence provided in the literature for its performance, and the availability of the estimated parameters in a relevant setting; both VLCC and Suezmax tankers can be used as floating storage, and the latter tanker is studied by Poblacion (2015). More specifically, the parameters of the four-factor model are estimated using the Time Charter

¹ Imarex (International Maritime Exchange) is an exchange based in Oslo for trading shipping derivatives.

Equivalent (TCE) rates by Poblacion (2015) for the Suezmax tanker on the TD5 route, which is a suitable choice as floating storage.

The Mirantes et al. (2012) four-factor stochastic model, similar to the Schwartz and Smith (2000) model, expresses the logarithm of the freight spot price as the sum of a long-term factor, ξ'_t , a short-term factor, χ'_t , and a seasonal one, α_t . The SDE of the seasonal factor, α_t , is coupled by a fourth factor, α_t^* , since they associate with the real and imaginary parts of a complex factor. It should be noted that this model is originally presented in terms of a notation with factors similar to those of Schwartz and Smith (2000), i.e. χ_t and ξ_t . However, to avoid any confusions, the stochastic factors of Mirantes et al. (2012) model are shown with χ'_t and ξ'_t , denoting the short- and long-term factors respectively. For the same reason, the coefficients notation are also altered in the SDEs with the addition of a “prime” in Mirantes et al. (2012) model. Here, the logarithm of the freight spot price is denoted by S'_t .

$$\ln(S'_t) = \xi'_t + \chi'_t + \alpha_t \quad \text{Eq. 5.1}$$

$$d\xi'_t = \mu_{\xi'} dt + \sigma_{\xi'} dW_{\xi'} \quad \text{Eq. 5.2}$$

$$d\chi'_t = -k' \chi'_t dt + \sigma_{\chi'} dW_{\chi'} \quad \text{Eq. 5.3}$$

$$d\alpha_t = 2\pi\varphi\alpha_t^* dt + \sigma_{\alpha} dW_{\alpha} \quad \text{Eq. 5.4}$$

$$d\alpha_t^* = -2\pi\varphi\alpha_t dt + \sigma_{\alpha} dW_{\alpha^*} \quad \text{Eq. 5.5}$$

Here the long-term drift, $\mu_{\xi'}$, the speed of mean reversion, k' , the seasonal period, φ , factor volatilities, $\sigma_{\xi'}$, $\sigma_{\chi'}$, and σ_{α} are constants. Also $dW_{\xi'}$, $dW_{\chi'}$, dW_{α} , and dW_{α^*} are correlated Brownian motion increments, except dW_{α} and dW_{α^*} , which must be uncorrelated (so, there are $C(4, 2) - 1 = 5$ pairs of correlated increments). dW_{α} and dW_{α^*} are uncorrelated since the last two real SDEs with the same variance are resulted from the SDE of a complex trigonometric component, say $a_t = \alpha_t + i\alpha_t^*$, in the form of $da_t = -i2\pi\varphi a_t dt + \sigma_{\alpha} dW_{\alpha}$. The risk-neutral version of the Mirantes et al. (2012) model, which is needed for derivative pricing, is as follows.

$$d\xi'_t = (\mu_{\xi'} - \lambda_{\xi'}) dt + \sigma_{\xi'} dW_{\xi'}^{\circ} \quad \text{Eq. 5.6}$$

$$d\chi'_t = (-k'\chi'_t - \lambda_{\chi'})dt + \sigma_{\chi'}dW_{\chi'}^\circ \quad \text{Eq. 5.7}$$

$$d\alpha_t = (2\pi\varphi\alpha_t^* - \lambda_\alpha)dt + \sigma_\alpha dW_\alpha^\circ \quad \text{Eq. 5.8}$$

$$d\alpha_t^* = (-2\pi\varphi\alpha_t - \lambda_{\alpha^*})dt + \sigma_\alpha dW_{\alpha^*}^\circ \quad \text{Eq. 5.9}$$

Here $\lambda_{\xi'}$, $\lambda_{\chi'}$, λ_α , and λ_{α^*} are the risk premiums, and $dW_{\xi'}^\circ$, $dW_{\chi'}^\circ$, dW_α° , and $dW_{\alpha^*}^\circ$ are the increments of the Brownian motion under the risk-neutral measure. The price of freight forward contracts at time t with a time-to-maturity of $(T - t)$ is derived from the expected value of the freight spot price at time T under the risk-neutral measure Q . Mirantes et al. (2012) derived the freight forward prices as expressed by Eq. 5.10 and Eq. 5.11.

$$F'(t, T) = \quad \text{Eq. 5.10}$$

$$\exp(\xi'_t + e^{-k'(T-t)}\chi'_t + \cos(2\pi\varphi(T-t))\alpha_t + \sin(2\pi\varphi(T-t))\alpha_t^* + A_4(T-t))$$

$$A_4(T-t) = (\mu_{\xi'} - \lambda_{\xi'} + 0.5\sigma_{\xi'}^2 + 0.5\sigma_\alpha^2)(T-t)$$

$$- (\lambda_{\chi'} - \sigma_{\xi'}\sigma_{\chi'}\rho_{\xi'\chi'}) (1 - e^{-k'(T-t)})/k'$$

$$+ 0.25\sigma_{\chi'}^2 (1 - e^{-2k'(T-t)})/k'$$

$$- \frac{\lambda_{\alpha^*} + \lambda_\alpha \sin(2\pi\varphi(T-t)) - \lambda_{\alpha^*} \cos(2\pi\varphi(T-t))}{2\pi\varphi}$$

$$\text{Eq. 5.11}$$

$$+ \frac{\sigma_{\xi'}\sigma_\alpha\rho_{\xi'\alpha}(1 - \cos(2\pi\varphi(T-t))) + \sigma_{\xi'}\sigma_{\alpha^*}\rho_{\xi'\alpha^*}\sin(2\pi\varphi(T-t))}{2\pi\varphi}$$

$$+ \frac{\sigma_{\chi'}\sigma_\alpha\rho_{\chi'\alpha}}{k'^2 + (2\pi\varphi)^2} (k' - e^{-k(T-t)}(k'\cos(2\pi\varphi(T-t)) + 2\pi\varphi\sin(2\pi\varphi(T-t))))$$

$$+ \frac{\sigma_{\chi'}\sigma_{\alpha^*}\rho_{\chi'\alpha^*}}{k'^2 + (2\pi\varphi)^2} (2\pi\varphi + e^{-k(T-t)}(k'\sin(2\pi\varphi(T-t)) - 2\pi\varphi\cos(2\pi\varphi(T-t))))$$

To simulate the stochastic factors, Eq. 5.2 to Eq. 5.5 are discretized, which results in the following equations.

$$\chi'_t = \chi'_{t-1}e^{-k'\Delta t} + \sigma_{\chi'} \sqrt{\frac{1 - e^{-2k'\Delta t}}{2k'}} Z_{\chi'} \quad \text{Eq. 5.12}$$

$$\xi'_t = \xi'_{t-1} + \mu_{\xi'}\Delta t + \sigma_{\xi'}\sqrt{\Delta t}Z_{\xi'} \quad \text{Eq. 5.13}$$

$$\alpha_t = \alpha_{t-1} + (2\pi\varphi\Delta t)\alpha_t^* + \sigma_{\alpha}\sqrt{\Delta t}Z_{\alpha} \quad \text{Eq. 5.14}$$

$$\alpha_t^* = \alpha_{t-1}^* + (-2\pi\varphi\Delta t)\alpha_t + \sigma_{\alpha}\sqrt{\Delta t}Z_{\alpha^*} \quad \text{Eq. 5.15}$$

To have the form suitable for simulating the stochastic factors, the last two equations, Eq. 5.14 and Eq. 5.15, must be solved jointly to obtain α_t and α_t^* in terms of α_{t-1} and α_{t-1}^* , The result is as follows.

$$\alpha_t = \{\alpha_{t-1} + (2\pi\varphi\Delta t)\alpha_{t-1}^* + \sigma_{\alpha}\sqrt{\Delta t}[(2\pi\varphi\Delta t)Z_{\alpha^*} + Z_{\alpha}]\}/\{1 + (2\pi\varphi\Delta t)^2\} \quad \text{Eq. 5.16}$$

$$\alpha_t^* = \{\alpha_{t-1}^* + (-2\pi\varphi\Delta t)\alpha_{t-1} + \sigma_{\alpha}\sqrt{\Delta t}[(-2\pi\varphi\Delta t)Z_{\alpha} + Z_{\alpha^*}]\}/\{1 + (2\pi\varphi\Delta t)^2\} \quad \text{Eq. 5.17}$$

Poblacion (2015) estimates the four-factor model of Mirantes et al. (2012) using two sets of data based on Word Scale (WS) and Time Charter Equivalent (TCE) as proxies for freight rate on different routes employing the Kalman filter method. The data includes weekly observations of freight spot and forward TCE and WS prices from Jan 2009 to Feb 2014 across five routes. The freight forward prices are the current month forward, and 1st, 2nd, 3rd, 4th, and 5th months after the current month, in addition to the 3rd, 4th, and 5th quarters after the current quarter. The routes are TC2, TC14, TC6, TD5, and TD16 as defined by the Baltic Exchange. It should be reminded that TC (Tanker-Clean) refers to sub-indices that characterize carrying light or middle distillates such as gasoline or naphtha, and therefore are irrelevant to the crude storage trade. On the other hand, TD (Tanker-Dirty) specifies the sub-indices for heavy condensates such as crude oil. However, between TD5 and TD16, only routes TD5 has a sufficient capacity to be used in the crude oil storage; TD5 characterizes a 1 million-barrel Suezmax tanker, whereas the capacity on the TD16 route is specified as 0.25 million barrels (please see Table 5.1 for dirty tanker sub-index capacities). Table 5.2 lists the parameters estimated by Poblacion (2015). It shows Mirantes et al. (2012) model parameter estimates on the TD5 route using the TCE data. Fig. 5.2 illustrates three

different sample realizations of the freight spot and forward prices using the parameter estimates in Table 5.2.

Parameter	Estimate	Standard Error
$\mu_{\xi'}$	1.1049**	0.4071
k'	4.8854***	0.2584
φ	0.8426***	0.0000
$\sigma_{\xi'}$	0.4134***	0.0605
$\sigma_{\chi'}$	1.9123***	0.1141
σ_{α}	0.2285***	0.0229
$\rho_{\xi'\chi'}$	-0.3849***	0.1071
$\rho_{\xi'\alpha}$	-0.2769	0.1745
$\rho_{\xi'\alpha^*}$	0.5984***	0.1098
$\rho_{\chi'\alpha}$	-0.2453**	0.0967
$\rho_{\chi'\alpha^*}$	-0.7586***	0.0554
$\lambda_{\xi'}$	1.2014**	0.4273
$\lambda_{\chi'}$	-1.4337	1.3257
λ_{α}	-0.0291	0.1939
λ_{α^*}	0.5587**	0.1942

Table 5.2. The Mirantes et al. (2012) model parameters estimated by Poblacion (2015) on the TD5 route using the TCE data. Note x^* , x^{**} , and x^{***} show the estimated values are significant at 10%, 5%, and 1%, respectively.

Koekebakker and Adland (2004) study the forward freight rate dynamics; in addition to the freight spot price, Time Charter (TC) rates of maturities 6, 12, and 36 months are used to construct a smooth term structure on a daily basis. They derive the TC rates as a function of $F'(t, T)$, where $F'(t, T)$ is a model for the continuous term-structure of the freight rates. Let $TC(t, T)$ denote the TC rate at time t for chartering the vessel from time t to T . Note that TC is a constant rate that must be paid to the tanker owner on each instant in $[t, T]$, i.e. per unit of time. They argue that because there is no initial cost for entering a forward contract, the value generated by the difference between the instantaneous freight spot rate and the constant TC must be zero under the risk-neutral measure, as expressed by Eq. 5.18.




Fig. 5.2. Three sample realizations of $F'(t, T)$ (\$/tanker per day) using the parameters of Table 5.2 and initial condition $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*) = (0.3, 9, 0.3, 0.3)$.

$$E_t^Q \left[\frac{\int_t^T e^{-r(u-t)} (F'(u, u) - TC(t, T)) du}{T - t} \right] = 0 \quad \text{Eq. 5.18}$$

Based on the above risk-neutral argument, Koekebakker and Adland (2004) compute $TC(t, T)$ as the following equation as a function of $F'(t, T)$.

$$TC(t, T) \approx \frac{\int_t^T F'(t, u) du}{T - t} \quad \text{Eq. 5.19}$$

Using the $TC(t, T)$ equation, Eq. 5.19, and the Mirantes et al. (2012) closed form formula of $F'(t, T)$, Eq. 5.10, we can compute the integral to calculate $TC(t, T)$. Due to the difficult nature of the integrand, the integral is computed numerically given any t and T . Therefore, it can be assumed that for any given $(\xi'_t, \chi'_t, \alpha_t, \alpha_t^*)$, $F'(t, T)$, and subsequently $TC(t, T)$ can be treated as known (deterministic).

5.5 Markov Decision Process (MDP) Model


The problem time horizon is $[0, \bar{T}]$, and is discretized into N equidistant stages by $t_i = i\Delta t$ for $i \in \mathcal{J} = \{0, 1, 2, \dots, N\}$, where $\Delta t = \bar{T}/N$. The trader is allowed to rent a tanker only once. The rent contract covers a fixed charter period of $T' (< \bar{T})$, and the trader can initiate a tanker rent contract in the period $[0, \bar{T} - T']$, which ensures that the charter period does not exceed the problem time horizon. Assuming that $N' = T'/\Delta t$, i.e. the length of the rent contract is N' periods, the trader may decide to start the rent contract at any $t_i = i\Delta t$ for $i \in \mathcal{J}' = \{0, 1, 2, \dots, N - N'\}$. At time t_i , the tanker time charter (TC) rate for renting the vessel from t_i to $t_i + T'$ is denoted by $TC(t_i, t_i + T')$.

The oil forward contract at time t_i , $F(t_i, T)$, can have the following maturities $T \in \{t_i, t_{i+1}, \dots, t_N\}$. The inventory level can be the range of $[0, \bar{R}]$, which is discretized uniformly into L levels by $\Delta R = \bar{R}/L$, which determines the selling batches allowed as $0, \Delta R, 2\Delta R, \dots, \bar{R}$. We consider a discrete-time dynamic optimization framework based on the following components.


1. The State Variables (x_i and W_i): At any stage i , the trader owns an inventory level of $R_i \geq 0$, and if $R_i > 0$, she also has a short position in a forward contract maturing at time T_i . The quantity of the short position is equal to the quantity of the long inventory position. There is a state variable I_i , which denotes the time when the tanker time charter contract expires. The expiry date of a valid (non-expired) contract, I_i must satisfy $t_i \leq I_i \leq \bar{T}$, i.e. it can be neither in the past nor exceeding the problem time horizon \bar{T} . If we are at time t_i , the rent may have been started at $\{t_0, t_1, t_2, \dots, t_{i-1}\}$, which corresponds to expiry times $\{t_{N'}, t_{N'+1}, t_{N'+2}, \dots, t_{N'+i-1}\}$ respectively. As shown in Fig. 5.3, if $t_i \leq T'$, it is not possible to have an expired contract since sufficient time has not elapsed since any contract initiation time. However, if $t_i > T'$, it is possible to have an “expired” state. The ‘0’ indicates the tanker has not been rented yet, and the sign ‘ \emptyset ’ indicates that it has been rented but the contract has expired (rental period ended). Eq. 5.20 summarizes the above arguments by formulating the set of feasible expiry dates, \mathcal{X}_i^I , at time t_i .

$$I_i \in \mathcal{X}_i^I = \begin{cases} \{0, t_{N'}, t_{N'+1}, \dots, \min(t_{i-1+N'}, \bar{T})\} & \text{if } 0 < t_i \leq T' \\ \{0, \emptyset, t_i, t_{i+1}, \dots, \min(t_{i-1+N'}, \bar{T})\} & \text{if } t_i > T' \end{cases} \quad \text{Eq. 5.20}$$

The endogenous component of the state variable will be $x_i = (I_i, R_i, T_i)$. The stage- i forward curve



(a) An expired contract cannot exist if time elapsed is less than T'



(b) An expired contract may exist if time elapsed is more than T'

Fig. 5.3. Maximum time elapsed since the initiation of a rental contract; (a) if $t_i \leq T'$: it is not possible to have an expired contract. (b) if $T' < t_i$: the contract is expired if it is initiated within the red interval, $[t_0, t_i - T']$.

$F(t_i, T)$, is fully specified by (χ_i, ξ_i) , and the time-charter (TC) term structure, $TC(t_i, t_i + T')$, is fully specified by $(\chi'_i, \xi'_i, \alpha_i, \alpha_i^*)$. Let $W_i = (\chi_i, \xi_i, \chi'_i, \xi'_i, \alpha_i, \alpha_i^*)$ denote the exogenous component of the state variable. Therefore, the current state is fully explained by (x_i, W_i) , where $(x_i, W_i) \in \mathcal{X}_i \times \mathbb{R}^6$, in which \mathcal{X}_i is the state space defined as in Eq. 5.21.

$$x_i = (I_i, R_i, T_i) \in \begin{cases} (I_i, 0, 0) & \text{if } I_i = 0 \text{ or } I_i = \emptyset \\ (I_i, 0, 0) \cup \{I_i \times (0, \bar{R}] \times \{t_i, t_{i+1}, \dots, I_i\}\} & \text{otherwise} \end{cases} \quad \text{Eq. 5.21}$$

\mathcal{X}_i ensures that the trader must have a valid tanker charter contract, i.e. $I_i \neq \emptyset$ and $I_i \neq 0$, to be able to either have a non-zero inventory and/or hold a forward position. Also, the expression $I_i \times (0, \bar{R}] \times \{t_i, t_{i+1}, \dots, I_i\}$ ensures that the maturity of the forward contract held does not exceed the expiry of the tanker charter contract, i.e. $T_i \leq I_i$ when $I_i \neq \emptyset, 0$.

The state $(0,0,0)$ refers to the case where the tanker has never been rented. The state $(\emptyset, 0,0)$ refers to the case where the tanker was rented but the rent contract ended. The initial state is defined by $x_0 = (I_0 = 0, R_0 = 0, T_0 = 0)$ and $W_0 = (\chi_0, \xi_0, \chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$, which indicates that the trader starts without any tanker charter contract, and subsequently zero inventory and no forward positions.

2. The Decisions (actions) (a_i): At stage i and state (x_i, W_i) , the decisions to be made consist of $a_i = (a_i^I, a_i^R, a_i^T)$. Here, a_i^I represents the decision to rent the tanker. If the tanker has not been rented yet, i.e. $I_i = 0$, the trader can decide to rent it by choosing $a_i^I = 1$, or decide to wait by choosing $a_i^I = 0$. However, if $I_i \neq 0$, meaning the tanker has been already rented (either still valid or expired), the only feasible decision is $a_i^I = 0$ showing a decision not to be able to rent again.

In addition, a_i^R denotes the quantity of oil to be bought ($a_i^R > 0$) or sold ($a_i^R < 0$) on the spot, where $a_i^R \in [-R_i, \bar{R} - R_i]$; $-R_i$ is the maximum amount that could be sold, limited by the available inventory, and $\bar{R} - R_i$ is the maximum amount that can be bought, limited by the maximum capacity. After selecting a_i^R , the decision maker chooses a_i^T which shows the maturity of the forward contract to short, which hedges the new inventory $R_i + a_i^R$. Of course, if $R_i + a_i^R =$

0, selecting the forward contract maturity a_i^T is meaningless because the inventory is empty, and thus, the only feasible a_i^T is set to 0 for tractability. Note that decisions are made only at stages $i \in \mathcal{J} \setminus \{N\} = \{0, 1, 2, \dots, N - 1\}$. The actions and the feasible set, $\mathcal{A}_i(x_i)$, can be formalized as follows.

$$a_i = (a_i^I, a_i^R, a_i^T) \in \mathcal{A}_i(x_i) = \begin{cases} (0,0,0) & \text{if } I_i \in \{\emptyset, t_i\} \\ \{\{1\} \times (0, \bar{R}] \times \{t_{i+1}, t_{i+2}, \dots, t_{i+N'}\}\} \cup (1,0,0) \cup (0,0,0) & \text{if } I_i = 0 \text{ and } 0 \leq t_i \leq \bar{T} - T \\ (0,0,0) & \text{if } I_i = 0 \text{ and } \bar{T} - T' < t_i \leq \bar{T} \\ \{\{0\} \times (-R_i, \bar{R} - R_i] \times \{t_{i+1}, t_{i+2}, \dots, I_i\}\} \cup (0, -R_i, 0) & \text{if } I_i \in \{t_{i+1}, t_{i+2}, \dots, t_{i-1+N'}\} \end{cases} \quad \text{Eq. 5.22}$$

In Eq. 5.22, the first line refers to the occasion, where a storage contract has been already expired or expires today, and there will not be any feasible (nontrivial) action. The second line refers to the case where the tanker has not been rented yet, but the trader has still the option to do so. If she decides to rent, she can also take position in the oil market. The third line represents the instance, in which a tanker has not been rented yet, and it is now too late to do so. Thus, it is not possible to take any (nontrivial) action. The fourth line characterizes the circumstances, where the trader holds a valid rent contract and, therefore, she has the possibility of adjusting her position in the oil market.

3. State Transition Function $f_i(x_i, a_i)$: Given the current state (x_i, W_i) and the decisions made a_i , the endogenous part of the next state, x_{i+1} , will be determined by the function $f_i(x_i, a_i)$ as defined in Eq. 5.23. It reflects the transition of the inventory through the immediate action, a_i^R , as well as substitution of the newly selected forward position, a_i^T . If the next stage inventory, R_{i+1} , is zero, the maturity choice will automatically be set to $a_i^T = 0$ according to the feasible set described earlier, and so $T_{i+1} = 0$. Given that $I_0 = 0$, the evolution of I_i via the specified transition function guarantees that I_i changes from zero as soon as the decision maker selects $a_i^I = 1$, and it stays the same for N' periods until it becomes \emptyset for the rest of the time. The exogenous part of the state, W_i , evolves based on the stochastic processes of Eq. 2.16 and Eq. 2.17 independently from x_i and a_i .

$$x_{i+1} = f_i(x_i, a_i) = (I_{i+1}, R_{i+1}, T_{i+1}),$$

where

$$(i) I_{i+1} = \begin{cases} \emptyset & \text{if } I_i \in \{\emptyset, t_i\} \\ I_i & \text{if } I_i \in \{t_{i+1}, t_{i+2}, \dots, t_{i-1+N'}\} \\ t_{i+N'} & \text{if } I_i = 0 \text{ and } a_i^I = 1 \\ I_i & \text{if } I_i = 0 \text{ and } a_i^I = 0 \end{cases} \quad \text{Eq. 5.23}$$

$$(ii) R_{i+1} = R_i + a_i^R$$

$$(iii) T_{i+1} = a_i^T$$

4. Reward Function $r_i(a_i, x_i, W_i)$: Given the action $a_i = (a_i^I, a_i^R, a_i^T)$ at stages $i \in \{0, 1, 2, \dots, N-1\}$, there will be a reward as shown by Eq. 5.24, which is the sum of four parts. The first component is $-a_i^R [F(t_i, t_i) - c_P^- I(a_i^R < 0) + c_P^+ I(a_i^R > 0)]$, which is the payoff generated by buying ($a_i^R > 0$) or selling ($a_i^R < 0$) oil on the spot market. Here, $I(*)$ denotes the indicator function. Note that the pumping costs associated with buying and selling are respectively set to c_P^+ and c_P^- . The second component, $(R_i + a_i^R)e^{-r(a_i^T - t_i)}F(t_i, a_i^T)$, is the payoff generated by short selling $R_i + a_i^R$ barrels through $F(t_i, a_i^T)$ contract. The third component, $-R_i e^{-r(T_i - t_i)}F(t_i, T_i)$, is due to offsetting the current short contract held (maturing at T_i) by going long the $F(t_i, T_i)$ contract.

$$r_i(a_i, x_i, W_i) = \quad \text{Eq. 5.24}$$

$$\begin{cases} -a_i^R [F(t_i, t_i) - c_P^- I(a_i^R < 0) + c_P^+ I(a_i^R > 0)] + (R_i + a_i^R)e^{-r(a_i^T - t_i)}F(t_i, a_i^T) - \\ R_i e^{-r(T_i - t_i)}F(t_i, T_i) - \bar{R}T'TC(t_i, t_{i+N'})I(I_i = 0 \wedge a_i^I = 1) & \text{if } I_i \neq t_i \\ -R_i c_P^- & \text{if } I_i = t_i \end{cases}$$

$$\forall i \in \{0, 1, 2, \dots, N-1\}$$

$$\text{and, } r_N(x_N, W_N) = -R_N c_P^-$$

The fourth and last component is $-\bar{R} \times T' \times TC(t_i, t_{i+N'}) \times I(I_i = 0 \wedge a_i^I = 1)$, which is the one-time tanker rent based on the assumption that the time-charter cost, $TC(t_i, t_{i+N'})$, is charged ‘once’ for the ‘whole’ vessel (\bar{R}) covering the period $[t_i, t_{i+N'}]$. The trader decides whether to rent from t_i to $t_{i+N'} = t_i + T'$ with no refund possibility. The term $I(I_i = 0 \wedge a_i^I = 1)$, where “ \wedge ” denotes logical “and”, guarantees the rental charge is computed based on time t_i if $I_i = 0$, i.e. tanker has not been rented yet, and $a_i^I = 1$, the decision to charter the tanker has been just made.

It is worth mentioning that in the above formulation it is assumed that the existing short contract, $F(t_i, T_i)$, is always offset through the market. This is true even if the maturity indicates a current period delivery, i.e. $T_i = t_i$. In this case, if the trader intends to fulfill a physical delivery, she selects $a_i^R = -R_i$ subsequently. This argument is the reason why in the reward function the only term with the pumping cost is the first term, which associates with the transaction in the physical market. The pumping cost represents all costs related to the physical delivery of the oil, such as pumping cost and any location discount to WTI futures.

When $I_i = t_i$, there is no decision making, remember $(a_i^I, a_i^R, a_i^T) = (0,0,0)$ in this case, which is the reason why $r_i(a_i, x_i, W_i) = -R_i c_{\bar{p}}$. Since by t_i (end of rent contract) an existing short contract with $T_i = t_i$ will be fulfilled by delivering all the remaining inventory, a pumping charge is triggered. If the oil has been already delivered, i.e. $R_i = 0$, the reward is simply zero.

A policy π is set of decision functions $\{A_0^\pi, A_1^\pi, A_2^\pi, \dots, A_{N-1}^\pi\}$, where $A_i^\pi(x_i, W_i): \mathcal{X}_i \times \mathbb{R}^6 \rightarrow \mathcal{A}_i(x_i)$ for $\forall i \in \{0, 1, 2, \dots, N-1\}$. The optimization is over the class Π , which is the set of all feasible policies π . The dynamic optimization problem, which gives the real option value at t_0 , will be given by Eq. 5.25. Here $\delta \in (0,1]$ denotes the constant time-discount factor for one stage, i.e. $\delta = e^{-r\Delta t}$. The expectation is under the physical measure as the goal is to capture the performance of the trading strategy. Also, x_i^π denotes the random endogenous part of the state at stage i when policy π is implemented.

$$V_0(x_0, W_0) = \max_{\pi \in \Pi} E \left[-\delta^N R_N c_P^- + \sum_{i=0}^{N-1} \delta^i r_i(A_i^\pi(x_i^\pi, W_i), x_i, W_i) \mid (x_0, W_0) \right] \quad \text{Eq. 5.25}$$

Model Features	Chapter 4	Chapter 5
Initiation time	Fixed at t_0	Optimally determined
Starting inventory level	Full (filled at t_0)	Empty
Inventory refill option	Not allowed	Allowed
Rental contract duration	Optimally chosen	Fixed (e.g. one year)
Extension, or early termination of the rental contract	Allowed	Not allowed
Refund of storage cost if contract terminated early	Allowed	Not applicable since early termination not allowed
Rent charge quantity basis	Per inventory level	Per whole tanker
Rent charge value basis	Known constant	Stochastically evolving
Endogenous state variables	<ul style="list-style-type: none"> • Inventory level (R_i) • Forward maturity (T_i) 	<ul style="list-style-type: none"> • Inventory level (R_i) • Forward maturity (T_i) • Rental contract expiry (I_i)
Decision variables	<ul style="list-style-type: none"> • Quantity sold on the spot (a_i^R) • Forward maturity (a_i^T) 	<ul style="list-style-type: none"> • Quantity bought or sold on the spot (a_i^R) • Forward maturity (a_i^T) • Rent the tanker or not (a_i^I)

Table 5.3. Highlight of the differences between the model used in Chapter 4 and Chapter 5 (the present model).

To enhance understanding of the framework, it will be helpful to compare the present model with the model introduced in Chapter 4. As summarized in Table 5.3, the comparison displays the richer and more realistic nature of the new model.

5.6 Theoretical Proposition

An Approximate Dynamic Programming (ADP) technique will be employed to solve the above problem. At first, it may seem that the new algorithm is a simple expansion of the ADP algorithm introduced in the previous chapter, where the exogenous state variable, W_i , has a higher dimension

due to the addition of the group of stochastic factors corresponding to the storage cost. Such an algorithm will not be correct because of the intricacies involving the continuation function approximations. Before presenting the new algorithm, let us state the following proposition, which will provide the basis required for the necessary changes in the new algorithm. A proof is presented in the appendix (please see Proof of Proposition 5.I section).

Proposition 5.I: Let $I_i^{-(0)}$ denote any state $I_i \in \mathcal{X}_i^I$ such that $I_i \neq 0$, and $x_i = (I_i^{-(0)}, *, *)$ denote any state $x_i = (I_i, R_i, T_i) \in \mathcal{X}_i$ in which $I_i \neq 0$. $i \in \{0, 1, 2, \dots, N - 1\}$, $V_i\left(\left(I_i^{-(0)}, *, *\right), W_i\right)$ is not a function of the storage cost stochastic factors, $(\chi'_i, \xi'_i, \alpha_i, \alpha_i^*)$, while $V_i((0,0,0), W_i)$ is a function of them.

To explain the intuition behind the proposition, note that $V_i\left(\left(I_i^{-(0)}, *, *\right), W_i\right)$ reflects the value function in a state where the tanker has been already rented. Thus, the determinants of the storage cost, $(\chi'_i, \xi'_i, \alpha_i, \alpha_i^*)$, do not impact the value function at this time, and vice versa for $V_i((0,0,0), W_i)$.

5.7 Algorithmic Solution (ADP)

Similar to the previous chapter, the ADP is developed based on Continuation Function Approximation (CFA) by estimating a vector of weights corresponding to a linear combination of the basis function (see Eq. 5.30). The vector of weights $\Theta(i, x_{i+1})$ is computed by least squares regression. The K basis functions are shown with $\Phi(W_i)$. Let us assume that the basis functions are formed by the polynomials up to degree 3 (the procedure is the same for other degrees). Since $W_i = (\chi_i, \xi_i, \chi'_i, \xi'_i, \alpha_i, \alpha_i^*)$, all the basis functions will be as explained in Eq. 5.26.

$$\bar{\Phi}'(W_i) = \text{Eq. 5.26}$$

$$\left[1 \quad \chi_i \quad \xi_i \quad \chi'_i \quad \xi'_i \quad \alpha_i \quad \alpha_i^* \quad \chi_i^2 \quad \dots \quad \chi_i \xi_i \quad \dots \quad \alpha_i^{*2} \quad \chi_i^3 \quad \dots \quad \chi_i^2 \xi_i \quad \dots \quad \alpha_i \alpha_i^{*2} \quad \alpha_i^{*3}\right]$$

Let $\Phi'_1(W_i)$ denote all the basis functions exclusively with respect to the oil factors, (χ_i, ξ_i) , as expressed by Eq. 5.27.

$$\Phi'_1(W_i) = [1 \quad \chi_i \quad \xi_i \quad \chi_i^2 \quad \xi_i^2 \quad \chi_i \xi_i \quad \chi_i^3 \quad \xi_i^3 \quad \chi_i^2 \xi_i \quad \chi_i \xi_i^2] \quad \text{Eq. 5.27}$$

Let Φ'_2 denote all the basis function included in $\bar{\Phi}'$ but not included in Φ'_1 (i.e. $\{\Phi'_2\} = \{\bar{\Phi}'\} \setminus \{\Phi'_1\}$ in a set subtraction notation). Now, let us reorganize vector $\bar{\Phi}'$ and write it as a block vector such that the first block is Φ'_1 and the second block is Φ'_2 . Let Φ denote this reorganized vector as expressed below. The vector of coefficients $\Theta'(i, x_{i+1})$ is expressed in terms of $\Theta'_1(i, x_{i+1})$ and $\Theta'_2(i, x_{i+1})$, which represent the coefficients corresponding to Φ_1 and Φ_2 respectively.

$$\Phi(W_i) = \begin{bmatrix} \Phi_1(W_i) \\ \Phi_2(W_i) \end{bmatrix} \quad \text{Eq. 5.28}$$

$$\Theta'(i, x_{i+1}) = [\Theta'_1(i, x_{i+1}) \quad \Theta'_2(i, x_{i+1})] \quad \text{Eq. 5.29}$$

Now, the CFA can be written in terms of Φ_1 and Φ_2 and the corresponding coefficients as expressed in Eq. 5.30.

$$\begin{aligned} E[\hat{V}_{i+1}(x_{i+1}, W_{i+1}) | W_i] &= \Theta'(i, x_{i+1}) \Phi(W_i) \\ &= \Theta'_1(i, x_{i+1}) \Phi_1(W_i) + \Theta'_2(i, x_{i+1}) \Phi_2(W_i) \end{aligned} \quad \text{Eq. 5.30}$$

The reason for separating the basis functions into two groups can be explained using the Proposition 5.I; depending on the state x_{i+1} , the expectation $E[\hat{V}_{i+1}(x_{i+1}, W_{i+1}) | W_i]$ can be a function of $(\chi_i, \xi_i, \chi'_i, \xi'_i, \alpha_i, \alpha_i^*)$ or just (χ_i, ξ_i) , which leads to using different basis functions as regressors. As proved in the Proposition 5.I, the expectation is only a function of (χ_i, ξ_i) if $x_{i+1} = (I_{i+1}^{-(0)}, *, *)$, according to which one can rewrite Eq. 5.30 in a state-dependent form.

$$E[\hat{V}_{i+1}(x_{i+1}, W_{i+1}) | W_i] = \begin{cases} \Theta'_1(i, x_{i+1}) \Phi_1(W_i) & \text{if } x_{i+1} = (I_{i+1}^{-(0)}, *, *) \\ \Theta'(i, x_{i+1}) \Phi(W_i) & \text{if } x_{i+1} = (0, 0, 0) \end{cases} \quad \text{Eq. 5.31}$$

The form expressed by Eq. 5.31 is suitable for performing the least squares regression using the relevant basis functions only. Obviously, Eq. 5.31 is a specific case of Eq. 5.30 if one sets the coefficients of the irrelevant basis functions to zero, as expressed by Eq. 5.32. Employing the above CFA methodology, it is now possible to develop the ADP algorithm as presented in Table 5.4.

$$\Theta'_2(i, x_{i+1}) = 0, \quad \forall i, x_{i+1} = (I_{i+1}^{-(0)}, *, *) \quad \text{Eq. 5.32}$$

1. Simulate M sample paths of the W_i process for $i = 0, 1, 2, \dots, N$; denoted by $\{W_i^m\}_{m=1}^M$.
2. Initialize $\Theta(i, x_{i+1}) = 0, \forall i, x_{i+1}$ and $\hat{V}_N(x_N, W_N) = -R_N c_P^-, \forall x_N, W_N$.
3. For $i = (N - 1), (N - 2), (N - 3), \dots, 1, 0^*$

4. For $\forall x_i = (I_i, R_i, T_i) \in \mathcal{X}_i$

Solve for the optimal action based on

$$\hat{V}_i(x_i, W_i^m) = \max_{a \in \mathcal{A}_i(x_i)} \{r_i(a, x_i, W_i^m) + \delta \hat{E}[\hat{V}_{i+1}(f_i(x_i, a), W_{i+1}) | W_i^m]\}$$

$$\hat{A}_i^\pi(x_i, W_i^m) = \operatorname{argmax}_{a \in \mathcal{A}_i(x_i)} \{r_i(a, x_i, W_i^m) + \delta \hat{E}[\hat{V}_{i+1}(f_i(x_i, a), W_{i+1}) | W_i^m]\}$$

where CFA is computed using the previously calculated $\hat{\Theta}(i, x_{i+1})$ as (except for \hat{V}_N)

$$\hat{E}[\hat{V}_{i+1}(x_{i+1}, W_{i+1}) | W_i^m] = \Theta'(i, x_{i+1}) \Phi(W_i^m)$$

End

5. Solve for the prior stage regression coefficients, $\hat{\Theta}(i - 1, x_i)$, by

$$\begin{cases} \hat{V}_i(x_i, W_i^m) \sim \Theta'(i - 1, x_i) \Phi(W_{i-1}^m) & \text{If } x_{i+1} = (0, 0, 0) \\ \hat{V}_i(x_i, W_i^m) \sim \Theta'_1(i - 1, x_i) \Phi_1(W_{i-1}^m) \text{ and } \Theta'_2(i - 1, x_i) = 0 & \text{If } x_i = (I_i^{-(0)}, *, *) \end{cases}$$

End

*At $i = 0$, the regression will be replaced with sample average due to the absence of multiple sample paths at $t = 0$, which means $\hat{E}[\hat{V}_1(x_1, W_1) | W_0] = \sum_{m=1}^M \hat{V}_1(x_1, W_1^m) / M$.

Table 5.4. Pseudocode for ADP (LSM) approach.

5.8 Parameters and Computational Setting

In Section 5.8.1, the benchmark values of parameters used in this chapter are introduced. Sections 5.8.2, 5.8.3 and 5.8.4 respectively study the accuracy of the choices made for the basis functions, the number of simulated paths, and the inventory discretization size.

5.8.1 Benchmark Parameters

In the following computational results, $M = 100000$ (50K + 50K antithetic) sample paths are generated by simulating the (exogenous) state variables with a timestep of 1/480 year. This

represents a fine time-scale discretization from which the required (coarser) samples are later extracted. For simulating the oil price factors, (χ_i, ξ_i) , Eq. 2.12, Eq. 2.13, and Eq. 2.14 together with the parameters provided in Table 2.1 are used, which is the same setting as in the previous chapters. For simulating the storage cost stochastic factors, $(\chi'_i, \xi'_i, \alpha_i, \alpha_i^*)$, Eq. 5.12, Eq. 5.13, Eq. 5.16, Eq. 5.17, and the parameters listed in Table 5.2 are used.

The initial conditions of the oil forward curve match those defined in the base case in previous chapters; $(\chi_0, \xi_0) = (-0.639, 4.637)$, which corresponds to a spot price of \$54.45 per barrel and a long-term price of \$103.19 simulating forward prices based on May 2009 market conditions. The time horizon is set to $\bar{T} = 2$ years, and the duration of the tanker rent contract is set to $T' = 1$ year. All the parameters used to obtain the computational results of this chapter are summarized under in Table 5.5.

Description	Parameter	Value
Time horizon (constraint)	\bar{T}	2
Duration (length) of the storage rent contract	T'	1
Timestep size ($\Delta t = \bar{T}/N = T'/N'$)	Δt	1/12
Inventory capacity (barrels)	\bar{R}	1
Storage discretization increment	ΔR	1
Pumping cost	(c_P^+, c_P^-)	(1.875, 1.875)
Initial condition of the oil state variables	(χ_0, ξ_0)	(-0.639, 4.637)
Initial condition of the storage cost state variables	$(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$	(3.39, 8.4, 0.3, 0.4)
Initial condition of the endogenous state variables	$x_0 = (I_0, R_0, T_0)$	(0, 0, 0)
Total number of simulated antithetic sample paths	M	100,000
Total number of simulated antithetic out-of-sample paths	M_2	25,000
Degree of polynomial basis functions	p	3

Table 5.5. Table of the general problem parameters establishing the benchmark setting.

The initial conditions of the stochastic storage cost factors are selected such that the one-year time charter rate at t_0 is $TC(0,1) = \$8.23$ per barrel per year, which is higher than $\$6.57$ used in the previous chapters (although $\$6.57$ is also tested in Section 5.9.5). The main purpose of such a choice is to establish an unfavorable cost condition at t_0 to encourage the algorithm to search for optimal starting times $t > t_0$. In the three following subsections, we run numerical experiments to test what “computational” parameters, e.g. M , M_2 , etc., are needed to get sufficiently accurate numerical results.

5.8.2 Basis Functions

Polynomial basis functions are used in the ADP algorithm. In the following, the polynomial degree (p) is set to 2, 3, or 4, and confidence intervals of V and V^c are computed by repeating each simulation 30 times. Note that V is the estimate of the value function at time zero assuming the stochastic storage cost, while V^c represents the value achieved under the assumption that the storage cost will remain ‘constant’ and equal to its initial value at t_0 . It is seen that the values increase with p slightly and at a diminishing rate. This is not surprising given that it is the result of obtaining a better fit to the continuation value. However, the computational time almost doubles

Polynomial degree (p)	2	3	4
Total number of basis functions (Φ elements)	28	84	210
Number of basis functions containing oil factors only (Φ_1 elements in Eq. 5.30)	6	10	15
Computational time (minutes)	13.65	24.33	45.55
Confidence Interval of V	[6.138,6.158]	[6.177,6.196]	[6.188,6.206]
V (mean)	6.148	6.186	6.197
% change in V relative to the smaller p	NA	+0.62%	+0.18%
Confidence Interval of V^c	[3.854,3.866]	[3.869,3.881]	[3.873,3.885]
V^c (mean)	3.860	3.875	3.879
% change in V^c relative the smaller p	NA	+0.39%	+0.10%

Table 5.6. Impact of the degree of the polynomial basis functions generated using $(\chi_t, \xi_t, \chi'_t, \xi'_t, \alpha_t, \alpha_t^*)$. The confidence intervals are computed by using 30 repetitions for each degree. The other parameters are as per Table 5.5.

with each increase in the polynomial power. As a result, $p = 3$ has been used in the benchmark case, as specified in Table 5.5, which is also consistent with the degree used in the previous chapter.

5.8.3 Number of Paths

The number of sample paths, $M = 100K$, and out-of-sample paths, $M_2 = 25K$, are specified together with other parameters in Table 5.5. By testing other choices of M and M_2 , it is shown that these values provide relatively good estimates. Fig. 5.4 shows V (panel a) and V^c (panel b) resulting from the out-of-sample and in-sample (solve the problem forward in time after computing the estimate of Θ backward in time) simulations, where values of $M \in \{25K, 50K, 100K, 200K\}$ and $M_2 \in \{0.05M, 0.12M, 0.25M\}$ are employed. The plots clearly show the convergence behavior in the estimates. Setting M_2 equal to $0.25M$ reduces the variability and leads to narrower confidence intervals. Also, $M = 200K$ does not provide greatly superior results as compared to $M = 100K$ given that the computational time doubles (24 vs 59 minutes). The confidence intervals are computed by repeating the simulations 30 times for all combinations of M and M_2 . However,




Fig. 5.4. Impact of the number of sample paths (M) and out-of-sample paths (M_2) on V (panel a) and V^c (panel b). The confidence intervals are computed by using 30 repetitions at each (M, M_2) pair. The parameters are as per Table 5.5.

there exist 90 repeated simulations for in-sample calculations (since the sample paths are already simulated for the out-of-sample tests), which explains why the in-sample confidence intervals are relatively narrow.

5.8.4 Inventory Discretization

According to the parameters in Table 5.5, $\Delta R = 1$, which means that partial sale is not considered inevitably. The tests in this section show that using a smaller value, which enables partial sales, does not lead to higher values. For instance, setting $\Delta R = 1/10$ results in $V = \$6.185$ and $V^c = \$3.881$, which matches values generated by $\Delta R = 1$. It is noteworthy that the computation times for $\Delta R = 1$ and $1/10$ are respectively about 24 and 327 minutes. Therefore, $\Delta R = 1$ is used in the benchmark case as specified in Table 5.5. It suffices to explore $\Delta R = 1$ and not consider more fine-grained tanker discretization. This is similar to the previous chapter (Chapter 4) problem, where it was shown theoretically that a partial sale is not optimal. We believe that the partial sale is not optimal in the Chapter 5 framework although it is not proved in the wider setting here. But experimental approaches suggest it makes no difference and offers strong support for this conjecture.

5.9 Computational Results

In the following subsections, the computational results exploring various aspects of the problem are presented; the optimal value and policy are respectively examined in sections 5.9.1 and 5.9.2. Impact of the initial condition of stochastic factors on the option value and the optimal policy are studied next; Section 5.9.3 discusses the impact of (χ_0, ξ_0) , and Section 5.9.4 investigates the impact of $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$ on $TC(0,1)$ (the one year storage cost at t_0). Section 5.9.5 presents the impact of $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$ on the option value while maintaining $TC(0,1)$ constant, and finally Section 5.9.6 does the same without maintaining $TC(0,1)$ constant. The effects from the level of pumping costs and time horizons are reviewed in sections 5.9.7 and 5.9.8 respectively. The computational results conclude with the sensitivity analysis in Section 5.9.9.

Unless stated otherwise, all the results are based on the out-of-sample (lower-bound) estimation of the value function at time zero, denoted by V . In addition to V , V^c is computed, which denotes the

value achieved under the assumption that the storage cost will remain ‘constant’ and equal to its initial value at t_0 . V^c is computed using the non-stochastic $TC(0,1)$ and the exact same oil forward curve realizations. Computation of V^c is done via an ADP algorithm similar to V only with stochastic oil factors.

5.9.1 Optimal Value

To calculate the optimal value and its confidence intervals, both V and V^c are computed 150 times using different (independent) samples and out-of-sample paths. The histograms of the obtained values are illustrated in Fig. 5.5. In each plot, three histograms are shown; one is based on the “Out-of-sample” computations, which leads to the regularly reported lower-bounds. The remaining two histograms are based on in-sample computations; “In-sample 1” indicates the case in which, the sample paths and the resulting Θ (the basis functions coefficients) are utilized in conjunction with each other in a time-forward optimal decision making on each path. This leads to computing the average payoffs across all paths, which is similar to the out-of-sample computations (except the same set of paths are used again). “In-sample 2” is an approach in which, the sample paths are used to compute Θ , after which the value is obtained by solving Bellman equation at the first timestep. “In-sample 2” approach leads to value being directly dependent upon Θ through the continuation value approximation, and thus more prone to in-sample bias, whereas in “In-sample 1”, the bias can only penetrate via the incorporated policy and the payoffs remain intact based on actual cash flows. The mean values computed (and shown by the solid line on the plots) from each approach confirm the above arguments.

In Table 5.7, the confidence intervals of the estimated V and V^c are reported. It is seen that V is much larger than V^c ; it is because the initial condition of the storage cost is unfavorable (it is too expensive at t_0 but it will most likely decline later). This point will be investigated in great detail in the subsequent sections.

V	[6.181, 6.189]	V^c	[3.878, 3.883]
-----	----------------	-------	----------------

Table 5.7. Confidence intervals (95%) of V and V^c resulted from the out-of-sample testing. The parameters are as per Table 5.5.


(a) V (b) V^c

Fig. 5.5. Histogram of V (panel a) and V^c (panel b) formed by computing the values 150 times. Values are derived from three different methods; Out-of-sample, In-sample 1 (solve optimally forward in time by implementing the computed Θ), and In-sample 2 (approximate the value function at t_0 by solving Bellman equation). The solid lines represent the mean of the histograms. The parameters are as per Table 5.5.

5.9.2 Optimal Policy

The optimal policy determines, at each timestep i and state (x_i, W_i) , what action $a_i = (a_i^l, a_i^R, a_i^T)$ should be taken. Recall that $a_i^l = 1$ represents the decision to rent the tanker (if not already rented), and $a_i^l = 0$ indicates the decision not to rent it. Additionally, $a_i^R \in [-R_i, \bar{R} - R_i]$ denotes the quantity of oil to be bought ($a_i^R > 0$) or sold ($a_i^R < 0$) on the spot. Furthermore, a_i^T shows the maturity of the forward contract to short if the new inventory is not empty, i.e. $R_i + a_i^R \neq 0$.

In the following, the optimal policy at some timesteps i and states x_i is presented in the domain of $W_i = (\chi_i, \xi_i, \chi_i', \xi_i', \alpha_i, \alpha_i^*)$ across the sample paths. It is difficult to illustrate the policy with respect to $W_i \in \mathbb{R}^6$ in a two-dimensional space. However, it will be seen that each of the six elements of W_i does not bear the same degree of influence on the policy. Let us consider the state $x_i = (I_i, R_i, T_i) = (0,0,0)$, i.e. the tanker has not been rented yet, as the benchmark state in Fig. 5.6;




Fig. 5.6. Optimal decision with respect to renting, a_i^I , at $t = t_1 = 1/12$ (a and b) and $t = t_4 = 4/12 = 1/3$ (c and d). State: at both t_1 and t_4 the tanker has not yet been rented, i.e. $x_i = (I_i, R_i, T_i) = (0, 0, 0)$.

panel (a) shows the a_i^I decision at $t = t_1 = 1/12$ in the $\chi_i - \xi_i$ plane, while panel (b) displays the same decision in the $\chi'_i - \chi_i$ plane, where it is seen that the latter pair explains the decision more


distinctively (in separate regions). In Fig. 5.6, panel (c) and (d), the same information is repeated except at $t = t_4 = 4/12$. With the passage of time from $t_1 = 1/12$ to $t_4 = 4/12$, $\chi_i - \xi_i$ has become better in explaining the a_i^l decision while $\chi'_i - \chi_i$ has not changed much. One possible explanation might be as follows; the optimal policy is built on and demonstrated via the evolving Monte Carlo sample paths from the initial condition. While the initial condition (at t_0) of oil forward curve is favorable, the storage cost initial condition is unfavorable. With the passage of time, e.g. at t_1 , a better storage cost condition is realized on some of the paths. Comparing Fig. 5.6.a and b, this is the reason why χ_i and χ'_i (rather than ξ_i) are needed to exhibit the decision boundary at t_1 . After more time has elapsed, at t_4 , it seems that χ_i is the most important factor in generating a clear boundary between the “rent” or “don’t rent” paths. Further investigations reveal that this trend continues in subsequent timesteps, which is consistent with other findings (in the following subsections) suggesting that the oil forward dynamics is more impactful than the storage cost. Table 5.8 summarizes the percentage of the paths on which the decision is to rent the tanker, i.e. $a_i^l = 1$, at different timesteps. At all timesteps, it is assumed that the state is $x_i = (I_i, R_i, T_i) = (0,0,0)$, i.e. the tanker has not yet been rented. This highlights the time varying nature of the optimal policy with respect the evolving $W_i = (\chi_i, \xi_i, \chi'_i, \xi'_i, \alpha_i, \alpha_i^*)$.

t_i	t_1	t_2	t_3	t_4	t_5	t_6	t_7	t_8	t_9	t_{10}	t_{11}	t_{12}
%	44.69	71.61	76.10	70.39	60.1	48.6	39.12	31.54	25.61	20.99	17.56	19.07


Table 5.8. Percentage of the paths on which the decision is to rent the tanker at different timesteps. It is assumed the state is that the tanker has not been rented yet, i.e. $x_i = (I_i, R_i, T_i) = (0,0,0)$.

In Fig. 5.7, the goal is to examine the optimal policy given different states at a single point in time, $t_7 = 7/12$. Three different states are assumed; the state at panels (a) and (b) is that the tanker has not been rented yet, i.e. $x_7 = (I_7, R_7, T_7) = (0,0,0)$. The state at panels (c) and (d) is that the tanker was rented at the previous timestep, i.e. $t_6 = 6/12$, but it is currently empty, i.e. $x_7 = (I_7, R_7, T_7) = (1.5,0,0)$. The state at panels (e) and (f) is that the tanker was rented at the previous timestep, i.e. $t_6 = 6/12$, it is currently full, and the inventory was hedged with a forward contract maturing today, i.e. $x_7 = (I_7, R_7, T_7) = (1.5,1, 7/12)$.


The optimal policies with respect to the quantity bought/sold, a_i^R , and the maturity, a_i^T , are shown respectively on the left and right columns of Fig. 5.7. Although the general structure of the policies corresponding to the three different states are relatively similar, it is interesting to clarify their distinguishing features given these states. In particular, the main decision boundary seems to move to the right as the state changes from $(0,0,0)$ to $(1.5,0,0)$, and finally to $(1.5,1,7/12)$. Accordingly, it is seen that the ‘buy’ decision in $(0,0,0)$ (red dots in Fig. 5.7.a) requires a more attractive condition (smaller χ_i ’s) compared to $(1.5,0,0)$ (red dots in Fig. 5.7.c) since there is not a tanker still rented in the former while there is one already rented in the latter. Similarly, it is seen that the ‘buy’ decision at $(1.5,0,0)$ (red dots in Fig. 5.7.c) to fill up the tanker, requires a more attractive condition compared to a ‘hold’ decision at $(1.5,1,7/12)$ (green dots in Fig. 5.7.e) to just keep the tanker full. Additionally, the less sharp (noisy) boundary line in the state $(0,0,0)$, compared to either of $(1.5,0,0)$ and $(1.5,1,7/12)$ is the result of the decision being dependent on $W_i = (\chi_i, \xi_i, \chi'_i, \xi'_i, \alpha_i, \alpha_i^*)$ in state $(0,0,0)$ versus just (χ_i, ξ_i) in states $(1.5,0,0)$ and $(1.5,1,7/12)$, as discussed theoretically in the earlier sections.




(a) Decision on quantity bought/sold on the spot, a_i^R . State: Tanker not yet rented, $(I_7, R_7, T_7) = (0,0,0)$.




(b) Decision on forward maturity, a_i^T . State: Tanker not yet rented, $(I_7, R_7, T_7) = (0,0,0)$.




(c) Decision on quantity bought/sold on the spot, a_i^R . State: Tanker rented but empty, $(I_7, R_7, T_7) = (1.5, 0, 0)$.



(d) Decision on forward maturity, a_i^T . State: Tanker rented but empty, $(I_7, R_7, T_7) = (1.5, 0, 0)$.



(e) Decision on quantity bought/sold on the spot, a_i^R . State: Tanker rented and full, $(I_7, R_7, T_7) = (1.5, 1, 7/12)$.



(f) Decision on forward maturity, a_i^T . State: Tanker rented and full, $(I_7, R_7, T_7) = (1.5, 1, 7/12)$.

Fig. 5.7. Optimal policy at $t = t_7 = 7/12$ with respect to the quantity bought/sold on the spot, a_i^R , is displayed on the left column (a, c, & e), and with respect to the maturity of the forward contract, a_i^T , on the right column (b, d, & f).


One important point observed in both Fig. 5.6 and Fig. 5.7 is that the optimal policy regarding the forward maturity, a_i^T , has a similar behavior to that of the problem solved in Chapter 4, where it was shown in Proposition 4.II that the optimal policy belongs to a small subset of the feasible set, i.e. $\{(R_i, 0), (0, t_N), (0, t_{i+1})\}$. Although this has not been proved here, the computational results do match this proposition.


So far, the optimal policy has been presented based on the sample path realizations in the W_i domain. In the following, the actual trades resulted from adopting the optimal policy on three out-of-sample paths are presented in Fig. 5.8. Both stochastic and constant storage cost assumptions are tested, which, for brevity, will be referred to as “problem V ” and “problem V^c ” respectively. Accordingly, any variable with a superscript “c” refers to the problem V^c . The trades executed on each of the three paths are shown on the three rows of Fig. 5.8. The panels on the left demonstrate the evolution of the forward curve, $F(t, T)$, overlaid with the one-year time-charter rate, $TC(t, t + 1)$, on a second y-axis. The panels on the right present the same information in a different fashion; on a single y-axis, they display the evolution of the one-year time-charter rate, $TC(t, t + 1)$, overlaid with the forward-spot spread, $F(t, T) - F(t, t)$, which has a direct association with the underlying trades.

The first set of panels, Fig. 5.8.a and b, show that in problems V and V^c rent (and buy oil) action occurs at t_0 and t_2 respectively, since the former anticipates a lower storage cost, while the latter is faced with a constant cost. However, they both sell the inventory at t_{11} when there is a favorable (for cashing out) downward-sloping forward curve. The sum of the rewards is respectively \$8.67 (4.78+3.89) and \$2.87 (2.64+0.23) in problems V and V^c . In comparison to problem V , problem V^c 's value suffered from two things; the obvious one is the more expensive storage cost at the trade initiation. The other factor hurting the value is the shorter time-to-maturity at the sell-out. More precisely, the maturity state variables at the time of sale, t_{11} , are $T_{11}^c = t_{12} = 1$ and $T_{11} = t_{14} = 1.17$ in problems V^c and V respectively. Therefore, the time-to-maturity of the new long forward position (offsetting the existing short position when cashing out) is one month in problem V^c ($T_{11}^c - t_{11} = t_{12} - t_{11} = 1/12$) and three months in problem V ($T_{11} - t_{11} = t_{14} - t_{11} = 3/12$). Subsequently, when the one- and three-month time spreads are translated into spot-forward


price spreads ($F(t, t) - F(t, T)$) using the same forward curve slope, the latter generates a larger profit since essentially the price spread (rise) is equal to the slope times the time spread (run). This difference is because problem V^c starts trading 2 periods sooner than problem V .

The second set of panels, Fig. 5.8.c and d, show that the algorithm buys and sells twice in problem V . The tanker is rented at $t_3 = 3/12$, which sets the expiry of the rent contract at $t_{15} = t_3 + 1 = 15/12$. The first purchase and sale occurs at $t_3 = 3/12$ and $t_9 = 9/12$ respectively. The second purchase and sale occurs at $t_{11} = 11/12$ and $t_{14} = 14/12$, one period before the expiry of rental contract. The total reward from the first buy and sell is \$16.83, while the second one results in \$0.04. It is noteworthy that during both trades, an optimal maturity of $T_t = t_{15}$ is preferred, i.e. the maximum maturity permitted. Similar to the previous path, the trade starts at t_0 in problem V^c . In fact, across all paths, trading starts at t_0 in V^c problem since the initial oil forward curve is favorable and the storage cost will be fixed throughout the problem.







(c) Path #2, $F(t, T)$ view



(d) Path #2, $F(t, T) - F(t, t)$ view



(e) Path #3, $F(t, T)$ view



(f) Path #3, $F(t, T) - F(t, t)$ view

Fig. 5.8. Three sample paths and the trades executed on them based on the optimal policy. The one-year time-charter rate, $TC(t, t + 1)$, is overlaid with the forward curve, $F(t, T)$, on the left panels, and is overlaid with the forward-spot spread, $F(t, T) - F(t, t)$, on the right panels. The vertical lines specify the decision times to ‘rent’, ‘buy’, or ‘sell’. The forward and spread curves are plotted in colors associated to ‘buy’ (red) and ‘sell’ (blue) at the corresponding times. T_t is the state variable at time t indicating the maturity of the contract held (if any). r_t is the reward (\$) of the action taken at time t . Superscript “c” represents the respective variables under the constant storage cost assumption (V^c) rather than the stochastic storage cost (V).

The third set of panels, Fig. 5.8.e and f, show that the algorithm in V problem starts trading at the very last permissible time, t_{12} . It seems that, until that point, no sufficiently upward-sloping forward curve was observed at the same time as a sufficiently low storage cost. Even at the rent/buy decision time, the slope is very small leading to a payoff $r_{12} = -5.44$. However, a very steep downward sloping curve compensate this initial loss very well with a payoff of $r_{16} = 14.85$.

The policy can accommodate a decision to rent the tanker but not to buy oil immediately, i.e. $a_i^I = 1$ and $a_i^R = 0$ at some i . However, it is observed that such a decision never happens; when the algorithm decides to rent the tanker ($a_i^I = 1$), it is invariably accompanied by a buy order ($a_i^R > 0$). This is because it rents the tanker when there is a trading opportunity, i.e. a steep forward curve, to maximize the time during which the tanker is utilized, rather than to wait for an opportunity while holding an empty tanker (even if it is rented at a relatively cheaper cost).

5.9.3 Impact of oil factors initial condition

The initial conditions for the oil factors, χ_0 and ξ_0 , significantly influence option values, V and V^c , by setting the initial price, $\exp(\chi_0 + \xi_0)$, and the long-term price, $\exp(\xi_0)$. In the following, the impact of χ_0 and ξ_0 is studied by computing V and V^c for a range of $(\chi_0, \xi_0) \in \{-0.7, -0.6, -0.5, \dots, -0.1\} \times \{3.8, 4, 4.2, \dots, 5.2\}$. For instance, $(\chi_0, \xi_0) = (-0.1, 4)$ corresponding to spot and long-term prices of \$49.40 and \$54.60 represents an unfavorable initial condition. While $(\chi_0, \xi_0) = (-0.7, 4.4)$ corresponding to spot and long-term prices of \$40.45 and \$81.45 respectively, represents a very favorable initial condition.

In Fig. 5.9 panels (a) and (b), by moving towards the top-left corner, a favorable initial condition (highly upward-sloping forward curve) unsurprisingly increases the option values from zero on the bottom-right corner to $V = \$19.02$ and $V^c = \$18.16$ on the top-left corner. Fig. 5.9.(d) illustrates the difference $V - V^c$, which, unlike V and V^c , does not change uniformly with an improving initial condition. It can be explained by observing three regimes; there are not many profitable opportunities in the ‘very bad’ initial condition region and thus $V - V^c$ is minute because both





Fig. 5.9. Impact of oil factors initial conditions on the value; V and V^c with respect to χ_0 and ξ_0 for $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*) = (3.39, 8.4, 0.3, 0.4)$, which are the benchmark values specified in Table 5.5, and represent an unfavorable initial condition ($\chi'_0 \gg 0$).


values are close to zero. On the other hand, in the ‘very good’ initial condition region, both optimal policies imply a buying decision. Here, entering into the storage contract almost immediately leading to large V and V^c . In this regime, the small $V - V^c$ is attributed to occasional trades (when the storage cost is stochastic) in which the trader rents the tanker at some $t > t_0$ at a lower cost than the initial cost.

In the third (middle) region, the initial condition is only moderately attractive. Under the stochastic storage cost case, better results, i.e. $V - V^c \gg 0$, are achieved by waiting and renting a tanker at a lower cost level, which frequently occurs due to the unfavorable initial condition of storage cost ($\chi'_0 \gg 0$), as seen in panel (c) of Fig. 5.9. However, under the constant storage cost, this expensive initial storage cost persists. For instance, detailed investigations for $(\chi_0, \xi_0) = (-0.6, 5.1)$, where a large difference between V and V^c ($V - V^c = \$1.27$) exists, shows that when stochastic storage cost is considered, the storage is rented at $t = t_1$ on 78% of the paths, and at $t > t_1$ on 22% of the paths. However, under the fixed storage cost assumption the storage is rented at $t = t_0$ on all paths.


The above arguments indicate that the relative behavior of V and V^c as quantified by $V - V^c$ in response to changes in (χ_0, ξ_0) should be interpreted considering the initial condition of storage cost factors, $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$. Therefore, in the following, the above experiment is repeated under a favorable storage cost initial condition. In panels (a) and (b) of Fig. 5.10, by moving towards the top-left corner, a favorable initial condition (highly upward-sloping forward curve) causes both values to increase significantly from zero to $V = V^c = \$25.29$. This is similar to the earlier observations made about Fig. 5.9. However, $V - V^c$ in Fig. 5.10.d shows a different pattern from Fig. 5.9.d due to the favorable storage cost at t_0 , which can be explained as follows. On the top-left corner of Fig. 5.10.d, where the (oil) initial condition is favorable, the optimal policies corresponding to V and V^c coincide because they both indicate a rent decision at t_0 . However, when the (oil) initial condition is unfavorable, there is no immediate rent decision (at $t = t_0$), and the two cases behave differently. To further clarify this, let us continue with a detailed investigation using $(\chi_0, \xi_0) = (-0.2, 4.8)$, as an example of an unfavorable initial condition. Under the stochastic cost assumption, because the storage cost will mostly likely increase in the future, the rent decision may be made in the early stages despite the unfavorable oil forward curve. The trader may or may not wait for a suitable condition to ‘buy oil’. If she does, this will reduce the effective time during which the storage is utilized. Alternatively, the rent can be made at a higher price in the subsequent timesteps when a profitable oil forward curve presents itself, when the condition is suitable for filling up the tanker. All scenarios will hurt the value under the stochastic storage cost. However, under the constant cost assumption, the trader can rent the tanker at the constant (low)




(a) V



(b) V^c



(c) Sample Time-Charter Rate Trajectories



(d) $V - V^c$

Fig. 5.10. Impact of oil factors initial conditions on the value; V and V^c with respect to χ_0 and ξ_0 for $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*) = (-3.39, 8.4, 0.3, 0.4)$, a favorably deviated initial condition ($\chi'_0 \ll 0$).

cost at any future time when the oil forward curve is suitable for the trade.

Fig. 5.11.a shows the histogram of the timestep at which the decision to rent the tanker is made. First, it is seen that the number of paths with a rent decision falls from 52% (13103) of the 25K total paths under the constant cost to 45% (11238) under the stochastic storage cost. Second, it is seen that, under the stochastic cost, more than half (6620 of 11238) of the paths have rent decisions

made at the early stages (t_1 and t_2). However, under the constant cost the frequency of immediate rentals is relatively much smaller (just 3866 paths at t_1 and t_2 from 13103 paths on which renting occurred), with many (2275 paths) of the decisions to rent being made at the last possible timestep, t_{12} .




Fig. 5.11. (a) Histogram of the timestep at which the rent decision is made (if any at all) for $(\chi_0, \xi_0) = (-0.2, 4.8)$, an unfavorable oil initial condition, and $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*) = (-3.39, 8.4, 0.3, 0.4)$, a favorable storage cost initial condition (b) Mean of the profits or losses (reward) made over the 25000 out-of-sample paths at each timestep with parameters similar to part (a).

5.9.4 Impact of Storage Cost Factors Initial Condition on the Initial Storage Cost


The storage cost, $TC(t, t + T)$, is the time-charter rate for renting the tanker at time t for a period of T years. The process $TC(t, t + T)$ is driven by the four stochastic factors $(\chi'_t, \xi'_t, \alpha_t, \alpha_t^*)$ of the storage cost model introduced earlier. The initial storage cost, $TC(t_0, t_0 + T')$ or in short $TC(0, 1)$, is the time-charter rate for renting the tanker at $t_0 = 0$ for a period of $T' = 1$ year. Before studying how changes in the initial condition of storage cost factors, $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$, impacts the option value, V and V^c , it would be informative to know how $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$ influences $TC(0, 1)$ since it is expected that at least part of the impact of $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$ on option value is transmitted via

$TC(0,1)$. Fig. 5.12 shows $TC(0,1)$ as a univariate function of each factor initial condition, $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$, while keeping the other three constant. It is seen that $TC(0,1)$ is an increasing function of all the factors. Also, $TC(0,1)$ is very sensitive to ξ'_0 since an increase in the long-term factor increase the area under the storage cost forward curve significantly (recall that $TC(t_0, t_0 + T')$ is the integral of the storage cost forward from t_0 to $t_0 + T'$ divided by T').


To study $TC(0,1)$ as a multivariable function of $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$, many possible combinations of $\chi'_0 \in [-4,4]$, $\xi'_0 \in [8,9.5]$, $\alpha_0 \in [-1,1]$, and $\alpha_0^* \in [-1,1]$ are selected by discretizing these intervals and considering all the combinations $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*) \in \{-4, -3.5, -3, \dots, 4\} \times \{8, 8.25, 8.5, \dots, 9.5\} \times \{-1, -0.75, -0.5, \dots, 1\} \times \{-1, -0.75, -0.5, \dots, 1\}$. Since the quantity $\exp(\chi'_0 + \xi'_0 + \alpha_0 + \alpha_0^*)$ represents the tanker spot prices, to be in the range of reasonable values, we limit $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$ combinations by requiring them to satisfy $2000 < \exp(\chi'_0 + \xi'_0 + \alpha_0 + \alpha_0^*) < 60000$. The result is about 4000 qualified combinations, which subsequently leads to the corresponding initial storage costs computed in the range of $\$0.95 \leq TC(0,1) \leq \15.46 per barrel per day.

To illustrate $TC(0,1)$ as a function of four variables in two dimensions, two variables, e.g. χ'_0 and ξ'_0 as in Fig. 5.13.a, are selected for X and Y axes. Focusing on a fixed (χ'_0, ξ'_0) point, all the $TC(0,1)$ values resulted from changes in the remaining variables, i.e. α_0 and α_0^* in Fig. 5.13.a, are reflected around the corresponding (χ'_0, ξ'_0) point by jittering the dots to avoid overlaps¹. A similar process is repeated in Fig. 5.13.b by choosing α_0 and α_0^* for the X and Y axes respectively. Fig. 5.13 (a) and (b) show that changes in the χ'_0 and ξ'_0 impacts $TC(0,1)$ more significantly than α_0 and α_0^* within the studied range. Fig. 5.13.a illustrates that $TC(0,1)$ is an increasing function of χ'_0 such that it will increase faster at higher levels of ξ'_0 .


¹ Jittering is achieved by replacing (x, y) coordinate of a point with $(x + u, y + v)$, where u and v are uniform random variables from $u \sim U[-a, a]$ and $v \sim U[-b, b]$. Here, $(a, b) = (0.18, 0.04)$ in Fig. 5.13.a and $(a, b) = (0.08, 0.08)$ in Fig. 5.13.b. The values for (a, b) are chosen such that it leads to an illustration in which the points are sufficiently dispersed within each cluster but the clusters do not overlap with each other.




(a) χ_0'



(b) ξ_0'



(c) α_0



(d) α_0^*

Fig. 5.12. Percentage change (decimal notation) in the initial storage cost $TC(0,1)$ as a univariate function of χ_0' , ξ_0' , α_0 , or α_0^* while keeping the other three fixed. The change is relative to the benchmark values specified in Table 5.5; $(\chi_0', \xi_0', \alpha_0, \alpha_0^*) = (3.39, 8.4, 0.3, 0.4)$. The red solid lines show the benchmark values.




Fig. 5.13. Initial storage cost $TC(0,1)$ as a multivariable function of $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$; (a) highlighting changes with respect to χ'_0 and ξ'_0 , (b) highlighting changes with respect to α_0 and α_0^* .

5.9.5 Impact of storage cost factors initial condition while ‘keeping the initial cost constant’

Impact of the storage cost initial conditions on the option values V is studied using different $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$. However, $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$ are selected such that the initial one-year time-charter rate $TC(0,1) = \$6.57$ in all cases, which is equal to the non-stochastic value used in the previous chapters. To do so, $(\xi'_0, \alpha_0, \alpha_0^*)$ is chosen from $\{8, 8.25, 8.5, \dots, 9.5\} \times \{-1, -0.75, -0.5, \dots, 1\} \times \{-1, -0.75, -0.5, \dots, 1\}$, and χ'_0 is then computed by numerically solving $TC(0,1) = 6.57$ with respect to χ'_0 , as expressed by Eq. 5.19. Since V^c represents the value under a constant storage cost equal to the initial value, V^c will be the same for all $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$ combinations and equal to $V^c = \$5.53$. Recall that the option value under the constant storage cost computed in the previous chapter was higher, \$10.80, since the framework assumptions were different (more relaxed).

The plots in Fig. 5.14 show that although for all the $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$ combinations $TC(0,1) = 6.57$, the resulting values can vary from \$5.53 to \$7.79. Larger V 's are usually achieved when a large χ'_0 is paired with a small ξ'_0 . This represents a temporarily expensive storage cost condition that




Fig. 5.14. V as a (univariate) function of (a) χ'_0 , (b) ξ'_0 , (c) α_0 , and (d) α_0^* . The black solid line represents the $V^c = \$5.53$ associated with the constant cost case. The rest of the parameters are set based on the benchmark values specified in Table 5.5. The initial conditions, $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$, are selected such that their combination leads to a constant $TC(0,1) = \$6.57$ per barrel per year.

will revert to its lower normal levels in subsequent periods, leading to capturing a higher value. Thus, given a constant $TC(0,1)$ at t_0 , V is relatively higher if this $TC(0,1)$ is the result of a short-

term (expensive) deviation from the long-term (lower) trend, rather than a short-lived cheap period reverting to higher storage costs.

The univariate relationships illustrated in Fig. 5.14 are combined and demonstrated jointly in Fig. 5.15, where V is shown on $\chi'_0-\xi'_0$ plane (panel a) and, $\chi'_0-\alpha_0$ plane (panel b). Fig. 5.15.a clearly shows a higher V is generated when a large χ'_0 is paired with a small ξ'_0 , i.e. a short-term expensive period reverting to a cheap long-term trend. Interestingly, a similar observation can be made in Fig. 5.15.b in a seasonality context; Fig. 5.15.b shows by fixing χ'_0 at for example 4, increasing α_0 increases V . It means that given a fixed short-term deviation (χ'_0), if the currently prevailing storage cost is the result of a seasonal increase in prices, there are better chances of achieving a higher value as the prices reverts to seasonally lower levels later. Similarly, if one accounts for seasonality by fixing α_0 , a larger χ'_0 increases value.





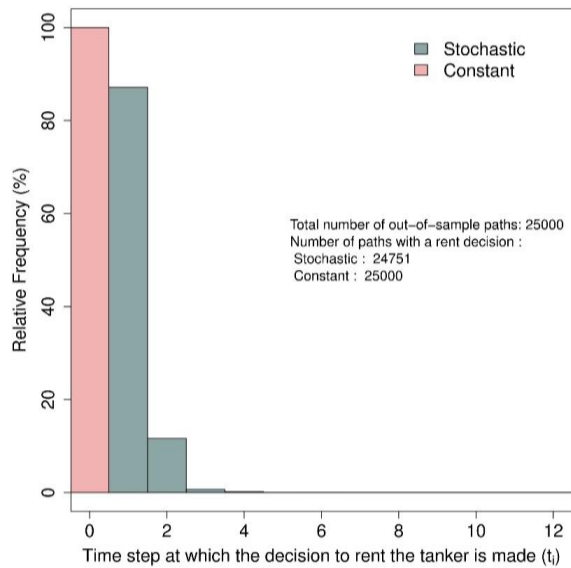
Fig. 5.15. V as a (bivariate) function of (a) $\chi'_0-\xi'_0$, where ξ'_0 values are randomly jittered for better demonstration, and (b) $\chi'_0-\alpha_0$. The rest of the parameters are set based on the benchmark values specified in Table 5.5. The initial conditions, $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$, are selected such that their combination leads to a constant $TC(0,1) = \$6.57$ per barrel per year.

Before making a comparison between the observed V and V^c values, it should be reminded that the oil forward curve initial condition is favorable in the above cases (as it is set based on the benchmark values specified in Table 5.5). The optimal policy under the constant storage cost

assumption (V^c) is to start the trade immediately since the oil forward curve condition is favorable and storage cost is fixed. Under the stochastic storage cost, when χ'_0 is large enough, waiting for a few timesteps before renting the tanker is optimal because the savings from a lower cost of storage is more than the potential loss from an unfavorable oil forward curve, which results in a $V > V^c = \$5.53$, as seen in Fig. 5.14.a. On the other hand, if χ'_0 is small (indicating the storage cost will most probably increase in the future), the optimal decision will imply to rent the storage and start the trade early, a policy similar to what is advised under the constant storage cost assumption. This explains the observation that $V = V^c = \$5.53$ in those cases. Fig. 5.16 verifies the above arguments by detail examination of two extreme ($\chi'_0, \xi'_0, \alpha_0, \alpha_0^*$) cases selected from those tested in Fig. 5.14 and Fig. 5.15. Recall that all ($\chi'_0, \xi'_0, \alpha_0, \alpha_0^*$) combinations result in $TC(0,1) = \$6.57$. The first set of values is $(-3.63, 9.5, -1, 1)$ leading to $V = V^c = \$5.53$, and the second set is $(3.91, 8, 0.5, -0.5)$ resulting in a $V = \$7.79 > V^c = \5.53 . The former set is used in the left panels (a, c and e), while the latter is used in the right panels (b, d, and f). Switching from the first set to the second set, it is seen that how most of renting decisions shifts from t_0 to t_1 (and t_2 to some extent). In Fig. 5.16 panels (c) and (d), mean of profits and losses at each timestep shows that the performance of V and V^c cases match on the left column. On the right column however, the stochastic case performs better by waiting for one period and taking advantage of a lower storage



(a)



(b)




Fig. 5.16. Comparing storage cost factors $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$ set to $(-3.63, 9.5, -1, 1)$ and $(3.91, 8, 0.5, -0.5)$ in the left and right columns respectively. (a)/(b): histogram of the time step at which the rent decision is made (if any at all), (c)/(d): mean of the profits or losses (reward) made over the 25000 out-of-sample paths at each timestep, (e)/(f): sample time-charter rate trajectories. The rest of the parameters are set based on the benchmark values specified in Table 5.5.

cost at t_1 ; the average profit increases from about \$4.40 (at t_0) in the constant case to \$5.50 (at t_1) in the stochastic case. Panels (e) and (f) show some sample time-charter rate paths, which illustrates how the storage cost in the subsequent periods increases on the left and decreases on the right.

5.9.6 Impact of storage cost factors initial condition

In this section, we generate 100 random storage cost initial conditions by selecting the values independently in the following intervals; $\chi'_0 \in [-4,4]$, $\xi'_0 \in [8,9.5]$, $\alpha_0 \in [-1,1]$, and $\alpha_0^* \in [-1,1]$. Subsequently, the one-year time-charter rate at t_0 , $TC(0,1)$, is computed using the selected $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$. Fig. 5.17 shows V , V^c , and $V - V^c$ using the above process under a favorable oil initial condition fixed at $(\chi_0, \xi_0) = (-0.6393, 4.6366)$ (the benchmark values in Table 5.5), while Fig. 5.18 shows the results under an ‘unfavorable’ oil initial condition fixed at $(\chi_0, \xi_0) = (-0.3, 4.3)$.

In both Fig. 5.17 and Fig. 5.18, as seen on panels (a.i) and (b.i), it is evident that both V and V^c mostly depend on χ'_0 - ξ'_0 pair; a low ξ'_0 coupled with a low χ'_0 provides the cheapest storage cost and highest values. This is the result of cheap long-term costs combined with a short-term deviation (to even cheaper costs) from the long-term trend. In Fig. 5.19, V and V^c are plotted with respect to the initial one-year time-charter rates $TC(0,1)$ implied by the same $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$ combinations. This figure confirms that both V and V^c decrease with $TC(0,1)$, which can be mainly attributed to an increase in χ'_0 .

The relative behavior of V and V^c in Fig. 5.19 indicate different characteristics comparing panels (a) and (b). The same observation can be made by comparing $V - V^c$ in panels (c) of Fig. 5.17 and Fig. 5.18. Focusing on a ‘favorable’ oil initial condition in Fig. 5.17.c and Fig. 5.19.a, if the initial storage cost (or equivalently χ'_0) is not too large, the optimal decision will be to rent the storage and start trading early, a policy similar to what is advised under the constant storage cost assumption, which inevitably results in $V = V^c$. However, when the initial storage cost (or similarly χ'_0) is large enough, waiting for a few timesteps before renting the tanker is optimal in problem V , because the savings from a lower storage cost is more than a potential loss from an unfavorable oil forward curve. Under the constant storage cost (problem V^c), as χ'_0 increases, the trader must rent the tanker at an increasingly more expensive initial cost without a possibility to revert to cheaper levels unlike in problem V . These result in a divergence between V and V^c ($V > V^c$) at higher storage costs. In summary, the gap $V - V^c$ starts from zero at cheaper storage cost initial conditions, and will increase as $TC(0,1)$ rises (or χ'_0 increases). Note that in this case $V, V^c \in$




Fig. 5.17. Random storage cost initial condition $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*) \in [-4, 4] \times [8, 9.5] \times [-1, 1] \times [-1, 1]$ under a favorable oil initial condition fixed at $(\chi_0, \xi_0) = (-0.6393, 4.6366)$ (benchmark).






Fig. 5.18. Random storage cost initial condition $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*) \in [-4,4] \times [8,9.5] \times [-1,1] \times [-1,1]$ under an unfavorable oil initial condition fixed at $(\chi_0, \xi_0) = (-0.3, 4.3)$.

Fig. 5.19. V or V^c as a function of the one-year time-charter rate at t_0 , $TC(0,1)$. Different storage cost factors initial conditions $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$ are used to compute $TC(0,1)$, where $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*) \in [-4,4] \times [8,9.5] \times [-1,1] \times [-1,1]$ are the same as utilized in Fig. 5.17 and Fig. 5.18. Panel (a) shows a favorable oil initial condition; $(\chi_0, \xi_0) = (-0.639, 4.637)$, i.e. the benchmark values specified in Table 5.5, and Panel (b) demonstrates an ‘unfavorable’ oil initial condition, $(\chi_0, \xi_0) = (-0.3, 4.3)$.



(a) favorable initial oil forward curve



(b) unfavorable initial oil forward curve

$[0,11)$, and $0 \leq V - V^c < 5$.

Now, turning the focus to an ‘unfavorable’ oil initial condition in Fig. 5.18.c and Fig. 5.19.b, waiting for a few timesteps before renting the tanker is optimal since it allows to exploit a potentially favorable oil forward curve realization. If the initial storage cost is unfavorably high, it has a chance to revert to lower levels with the passage of time under the stochastic cost scenario, while it stays fixed (expensive) under the constant storage cost assumption, and therefore $V > V^c$ in this case.

On the other hand, if the initial storage cost is favorable (low), it may revert to higher levels with the passage of time under the stochastic cost scenario, but the trader cannot take advantage of this initially low storage cost due to the initially unfavorable oil forward curve. However, under the constant storage cost assumption, the trader has access to the same (low) storage cost later. Therefore, it is seen in the range $TC(0,1) < 6$ that $V^c > V$.

In summary, increasing the initial storage cost (or χ'_0) lowers V^c more than it does V , and decreasing it boosts V^c more than it does V . In other words, when waiting is optimal from an oil perspective, having access to a (fixed) cheap storage leads to $V < V^c$, while a locked-in expensive storage results in $V > V^c$. Note that in this case $V, V^c \in [0,1.4)$ and $-0.7 < V - V^c < 0.4$.

5.9.7 Impact of the level of pumping costs

In Fig. 5.20, the impact of the pumping costs, c_p^+ (when buying oil) and c_p^- (when selling oil), are studied. Panels (a) and (b) confirm that as the pumping costs increase, there is indeed a decrease in both V and V^c , which is symmetric with respect to c_p^+ and c_p^- . However, V and V^c vary over different ranges, specifically $1.49 \leq V \leq 9.86$ and $0.24 \leq V^c \leq 7.69$. In both cases, the maximum value is attained when $c_p^+ = c_p^- = 0$ and the minimum when $c_p^+ = c_p^- = 5$. In panel (c), $V - V^c$ is always positive and varies somewhat nonuniformly across the domain. The relative difference $(V - V^c)/V$ illustrated in panel (d), indicates that V outperforms V^c at all (c_p^+, c_p^-) points and the measure increases as the pumping costs increases. In other words, under an ideal condition, i.e. $c_p^+ = c_p^- = 0$, V outperforms V^c by only 22% since both V and V^c lead to relatively

high values. Whereas under the severely adverse conditions, i.e. $c_p^+ = c_p^- = 5$, V outperforms V^c by 84%.




Fig. 5.20. Impact of the pumping costs, c_p^+ (cost when buying) and c_p^- (cost when selling). All the other parameters are set based on the benchmark values specified in Table 5.5.

Fig. 5.21 shows the detailed analysis of the trades for $(c_p^+, c_p^-) = (3, 3)$, where $V = \$4.13$ and $V^c = \$1.63$, i.e. a 61% difference. It indicates that V takes advantage of a cheaper rental cost by initiating rental agreements mainly at t_1, t_2 , and t_3 , compared to V^c initiating at t_0 . That is why some losses are observed in Fig. 5.21.b at times t_{13}, t_{14} , and t_{15} for V , and at t_{12} for V^c ; these losses stem from the rent contract nearing to its end and forcing the trader to make unattractive inventory sales. It is interesting to note that for V valuation problem, a decision to rent is made on a smaller number of paths compared to V^c (20865 vs 25000 respectively), however, Fig. 5.21.b shows that the average profit and loss over all paths are higher generally for V .




Fig. 5.21. Performance of the algorithms at $(c_p^+, c_p^-) = (3, 3)$ leading to $V = \$4.13$ and $V^c = \$1.63$; (a) histogram of the time step at which the rent decision is made, (b): mean of the profits or losses (reward) made over the 25000 out-of-sample paths at each timestep. The rest of the parameters are set based on the benchmark values specified in Table 5.5.


5.9.8 Impact of the time horizons

Table 5.9 summarizes the impact of changing the main problem time horizon \bar{T} , and the duration (length) of the storage contract, T' on both V and V^c . As expected, it is seen that expanding the time horizon \bar{T} increases the option value generally. Interestingly, extending the length of the rent contract, T' , increases the value significantly. Part of this substantial increase can be explained

	V					V^c				
$T' \setminus \bar{T}$	1.5	2	2.5	3	4	1.5	2	2.5	3	4
1	6.123	6.194	6.218	6.240	6.228	3.865	3.882	3.882	3.868	3.856
1.5	-	11.634	11.668	11.632	11.676	-	10.123	10.127	10.120	10.105
2	-	-	18.408	18.451	18.503	-	-	17.122	17.164	17.199

Table 5.9. Impact of the problem time horizon, \bar{T} (years), and duration of the storage contract, T' (years). All the other parameters are set based on the benchmark values specified in Table 5.5.

by Fig. 5.22; panel (a) presents the term structure of the initial time-charter rate, $TC(0, T)$, which is *dollar per barrel per year* units, and panel (b) provides the same information in units of *dollar per barrel*, by factoring in the contract term T and computing the total cost. The benefit of the latter form is that it allows to study the additional rental cost per barrel if the rental contract duration is increased. More specifically, the graph shows there is an additional cost of \$1.62/barrel



(a) Time-charter rate at $t = 0$, $TC(0, T)$, for different rental contract durations (maturities) T quoted in “\$ per barrel per year”.

(b) Time-charter rate at $t = 0$ times the contract duration, $TC(0, T) \times T$, for different maturities T , quoted in “\$ per barrel”.


Fig. 5.22. The term structure of the initial time-charter rate (panel a), and the same structure when it is multiplied by the rent duration (T) to represent the actual cost of storage in “\$ per barrel” (panel b). The basis for the parameters are Table 5.2 for the storage cost model and Table 5.5 for the general assumptions.

if the contract duration is increased from 1 to 1.5 years, and an additional \$0.90/barrel if it is increased from 1.5 to 2 years. On the other hand, the initial oil forward prices are $F(0,1) = \$67.53$, $F(0,1.5) = \$70.44$, and $F(0,2) = \$72.09$, which result in respective spreads of \$2.91 and \$1.65. Therefore, at $t = 0$, comparing the additional value generated from a longer forward maturity versus the additional cost from extending the rental contract duration indicates that the value increases noticeably. Note that analysis ignores all the extrinsic value generated from future trading possibilities because of having the tanker available for a longer time.

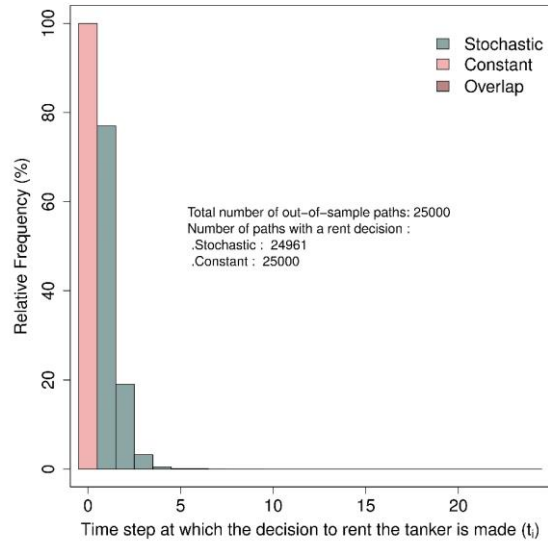
To shed further light on the impact of extending the time horizon, \bar{T} , and contract length, T' , a detailed comparison between two cases are provided in Fig. 5.23; $(\bar{T}, T') = (2, 1)$, i.e. the benchmark values specified in Table 5.5, and $(\bar{T}, T') = (3, 2)$. Under the stochastic cost assumption, it is found that the number of paths on which a rent decision is made increases from 23,415 to 24,961 (out of the 25,000 out-of-sample paths) by extending $(2, 1)$ to $(3, 2)$. Panels (a) and (b) illustrate the histogram of the timestep at which the rent decision is made, which seem to indicate a tendency toward an earlier start when $(\bar{T}, T') = (3, 2)$. Therefore, the more attractive setup, $(\bar{T}, T') = (3, 2)$, leads to trading more often and earlier.

Fig. 5.23, panels (c) and (d), show mean of the profits or losses made over the 25000 out-of-sample paths at each time step. The higher profits (\$4 vs \$2.5) observed at the early times confirm the arguments above that the gain from the initial oil forward curve outweigh the cost of a longer-term storage contract. In addition, the graphs show that the average profit, and the time period over which it is possible have the profits is higher in the case of $(\bar{T}, T') = (3, 2)$.

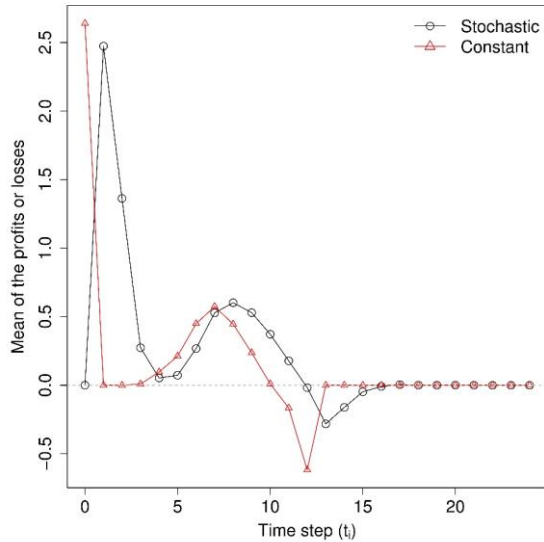
More importantly, detailed investigations revealed that the number of paths on which the tanker is filled more than one time increases from 5 to 340 under the stochastic storage cost, and from 1 to 249 under the constant storage cost assumption when $(\bar{T}, T') = (2, 1)$ is extended to $(3, 2)$. While 340 might not seem to be very many paths (just 1.36% of the total 25K simulated), but this is a “rent-free” trade opportunity on those path after the first trade occurred. Consequently, the average profit on these 340 paths is much higher, at \$39.97, in comparison to rest of the paths (ranging in $[-0.5, 4.4]$). Table 5.10 summarizes the information corresponding to the number paths with




(a) $\bar{T} = 2$ years and $T' = 1$ year



(b) $\bar{T} = 3$ years and $T' = 2$ years



(c) $\bar{T} = 2$ years and $T' = 1$ year



(d) $\bar{T} = 3$ years and $T' = 2$ years

Fig. 5.23. Impact of extending the main problem time horizon, \bar{T} , and the duration (length) of the storage contract, T' . (\bar{T}, T') is increased from $(2, 1)$ (i.e. the benchmark values specified in Table 5.5) to $(3, 2)$, and the corresponding results are presented on the left and right columns respectively. Panels (a) and (b) show the histogram of the timestep at which the rent decision is made (if any at all). Panels (c) and (d) illustrate mean of the profits or losses (reward) made over the 25000 out-of-sample paths at each timestep.

multiple fill ups and the average profit (or loss) on those paths for both stochastic and constant storage cost cases.

	V		V^c	
(\bar{T}, T')	(2, 1)	(3, 2)	(2, 1)	(3, 2)
Number of paths with multiple fill up	5	340	1	249
Average profit or loss over the above paths (\$)	13.31	39.97	14.53	39.08

Table 5.10. Impact of extending the main problem time horizon, \bar{T} , and the duration (length) of the storage contract, T' , on the paths with multiple fill ups.

5.9.9 Sensitivity Analysis

In this section, the goal is to study sensitivity of V and V^c to the parameters of the stochastic differential equations modeling oil prices (Schwartz and Smith (2000)) and storage cost prices (Mirantes et al. (2012)). Sensitivity is computed as the percentage change in V or V^c in response to the percent change in a single parameter. The basis for the parameters are Table 2.1 for the oil model, Table 5.2 for the storage cost model, and Table 5.5 for the general framework assumptions. For brevity, the sensitivity of V and V^c are reported simultaneously, where V results are shown with solid lines and V^c results are illustrated with dashed lines, and the impact of a particular parameter can be identified via the symbol specified in the legend. Fig. 5.24 shows the sensitivity with respect to the risk-premium and drift rate parameters. Firstly, it seems that the impact of oil-related parameters, namely λ_χ , λ_ξ , and μ_ξ , is more pronounced than that of the storage cost parameters. The same observation is made across all parameters including the initial condition $(\chi_0, \xi_0, \chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$. Secondly, V^c is more sensitive than V across all parameters. Recall that V^c is computed under a constant storage cost (equal to the initial value, $TC(0,1)$) in the absence of any future stochasticity. Thus, the only way V^c is affected by the storage model parameters is via $TC(0,1)$, which involves integration of the storage cost forward curve. It might be suggested that V^c is more sensitive than V since when a parameter influence $TC(0,1)$, it remains fixed (certain) for the duration of the problem time horizon in the case of V^c , whereas the stochasticity (uncertainty) in the variable storage cost case relaxes some of that effect through time.






Fig. 5.24. Sensitivity of V (solid lines) and V^c (dashed lines) to risk-premium and drift rate parameters. The basis for the parameters are Table 2.1 for the oil model, Table 5.2 for the storage cost model, and Table 5.5 for the general assumptions.


Fig. 5.25.a shows the sensitivity to the volatility parameters. The previous observations hold here as well; V^c is more sensitive than V , and oil-related parameters are more influential than the storage cost. The graph indicates that a higher σ_χ and σ_ξ increases the value, while a higher $\sigma_{\chi'}$ hurts the value. It might be explained by the fact that the former is the source of value, while the latter is a cost. Fig. 5.25 panel (b) shows the sensitivity to the speed of mean reversion, which indicates that the values are very sensitive to these parameters (still more sensitive to k_χ than $k_{\chi'}$). Panel (c) shows the sensitivity to the correlation parameters; it indicates that the results are relatively not very sensitive to these parameters as the graph ranges between -6% to +7%. The two most influential correlations are $\rho_{\chi\xi}$ and $\rho_{\chi'\xi'}$ with the value being more sensitive to $\rho_{\chi\xi}$ than $\rho_{\chi'\xi'}$. Panel (d) displays the sensitivity to the seasonality period parameter ϕ ; reducing it leads to a decrease in the value which is due to higher storage costs since further investigations reveal that




(a) Volatility parameters



(b) Speed of mean-reversion parameters



(c) Correlation parameters




(d) Seasonality period parameter (ϕ)

Fig. 5.25. Sensitivity of V (solid lines) and V^c (dashed lines) to different parameters. The basis for the parameters are Table 2.1 for the oil model, Table 5.2 for the storage cost model, and Table 5.5 for the general assumptions.

a smaller ϕ means a weaker seasonal pattern (resulting in a faster increase to the long-term higher prices), and higher time-charter rates.

Although in the previous sections the impact of the initial conditions, $(\chi_0, \xi_0, \chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$, on V and V^c was studied, in this part the sensitivity of V and V^c to these parameters are revisited in a fashion consistent with the sensitivity analysis. In Fig. 5.26, percentage change in V (or V^c) in terms of percentage change in one parameter is illustrated. It should be noted that the percentage change in the parameter refers to the absolute value and ignores the sign (important for χ_0). Fig. 5.26.a shows that while both initial conditions are impactful, ξ_0 is more influential than χ_0 , which can be explained by the fact that the latter is a mean-reverting factor. Also, it is seen that V^c more sensitive than V , which may be due to fact that it has a smaller range of values. Similarly, Fig. 5.26.b shows that ξ'_0 affect V and V^c more significantly than χ'_0 does, and both of which have more



(a) Oil factors initial conditions (χ_0, ξ_0)

(b) Storage cost factors initial conditions $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$

Fig. 5.26. Sensitivity of V (solid lines) and V^c (dashed lines) to the initial conditions, χ_0 and ξ_0 . The parameters are based on Table 5.5, where accordingly the benchmark values are $(\chi_0, \xi_0) = (-0.639, 4.637)$ and $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*) = (3.39, 8.4, 0.3, 0.4)$.

dominant impacts compared to α_0 and α_0^* . The direction of change is the opposite comparing (χ_0, ξ_0) and (χ'_0, ξ'_0) since the latter represents a cost.

5.10 Summary

In this chapter a more realistic problem framework is studied. To do so, a stochastically varying storage cost is implemented, which permits to examine the optimal decision making regarding the initiation of a fixed-term time-charter contract, a_i^I , the quantity to buy/sell on the spot, a_i^R , and the forward contract maturity, a_i^T . A modified version of the ADP approach is used to solve the resulting optimization problem. The modification is based on adjusting the basis functions present in the Continuation Function Approximation (CFA) linear structure based on the corresponding endogenous state, x_i , for which the theoretical ground is provided in Proposition 5.I.

To be able to isolate the impact of a stochastic storage cost, the result of a constant storage cost is also provided under the exact same conditions alongside, denoted by V and V^c respectively. In-sample and out-of-sample (lower-bound) values are calculated and compared. The computed optimal values show that given an unfavorable initial storage cost condition, e.g. benchmark values in Table 5.5, the stochastic storage cost algorithm can indeed take advantage of a potential later drop in the cost; the confidence intervals are [6.181, 6.189] and [3.878, 3.883] for V and V^c respectively. Similar to the previous chapter, it is found that a partial sale does not change the optimal value achieved. The optimal policies in regard with (a_i^I, a_i^R, a_i^T) are investigated by presenting them in a two-dimensional domain of exogenous state variables, and by showing out-of-sample trading decisions. Subsequently, the role of the χ_i or equivalently, the forward curve slope in the optimal policy is found to be crucial.

If the storage cost initial condition is unfavorable, i.e. the storage cost is high but will most likely experience a decline in the future, the optimal policy lead to $V^c = V$ when the oil forward curve initial condition is also unfavorable as both yield little profits. However, when the oil forward curve initial condition is favorable $V > V^c$ since V^c suffers from a locked-in expensive storage cost while V benefits from its decline. On the other hand, if the storage cost initial condition is favorable, i.e. the storage cost will most likely increase in the future, the optimal policy lead to

$V^c = V$ when the oil forward curve initial condition is favorable since both start trading early. However, $V^c > V$ when the oil forward curve initial condition is unfavorable since V^c enjoys from a locked-in cheap storage cost while they both wait for a better oil forward realization. In conclusion, the relative impact of the initial conditions on V and V^c significantly depends on the policy, specifically, when a decision to rent is advised.

It is shown that increasing the degree of the polynomial basis function increases the optimal value marginally, however, to avoid excessive computational costs, the third degree is selected as the benchmark. Furthermore, investigating the impact of the pumping costs shows that under severely adverse conditions, i.e. high pumping costs, the algorithm with the stochastic storage cost performs significantly better than the one with constant storage cost. In addition, it is found that the term (duration) of the tanker time-charter contract has a significant effect on the values achieved regardless of the storage cost stochasticity assumption.

Chapter 6

6 Conclusions and Future Research

6.1 Conclusions

The objective of the present research is to explain the off-shore oil storage trade observed in a contango market and aims to present an improved oil trading strategy directly based on transacting in the forward markets. Finding an optimal strategy to exploit a cash and carry arbitrage using both spot and forward markets is an important real options question. The underlying operational flexibilities impact both the problem and its MDP formulation. To study a range of assumptions, two main different frameworks are introduced. In the first framework, studied in Chapter 3 and Chapter 4, the duration of the storage rental contract is flexible, matching the optimally chosen maturity, and the storage cost is constant. In the second framework, examined in Chapter 5, the duration of the storage rental contract is fixed and the storage cost is stochastic (a constant storage cost is also considered in parallel). The additional level of complexity in the second framework allows for a richer and more realistic setup. Furthermore, the above advancements in the framework are accompanied with investigating, comparing, and identifying reliable solution approaches. The three solution methods studied are Forward Dynamic Optimization (Chapter 3), exact Dynamic Programming (Chapter 4), and Approximate Dynamic Programming (Chapter 4 and Chapter 5).

We contribute to the real options literature by studying the optimal decision making with respect to forward contract maturity for delivering an oil inventory, while the seller wishes to always stay hedged with a short position in the forward market. The decision maker can also sell her inventory partially on the spot market, and sell the rest of the inventory forward. There exists extensive literature concerned with natural gas storage valuation. However, a computationally tractable setting like this one in which forward trading decisions do not immediately impact the inventory level has not yet been studied. The present novel setting permits the trader to optimize the maturity of the forward contract for delivering the remaining inventory. This framework does not depend

on a specific price dynamics, and it can accommodate forward curve models with a higher number of stochastic drivers while being tractable.

In Chapter 3, the Forward Dynamic Optimization (FDO) solution method is studied. This heuristic strategy (also known as the Rolling Intrinsic policy) is myopic in the sense that it does not consider the continuation value of a position in evaluating an action. Instead, it just maximizes the immediate reward. Because, as such, FDO considers only those trades with an immediate positive payoff, it cannot initiate any trades under unfavorable initial conditions such as a downward (or insufficiently upward) sloping initial forward curve. It will not consider such trades even if they position the trader for excellent future profits. It should be noted that this apparent disadvantage renders this method to be very low risk (it never accepts any losses) and to have a policy that is completely model independent (no model for what might happen in the future is needed, as the future isn't really considered). To conclude, FDO is simple, transparent, and easy for practitioners to understand. It has attractive risk characteristics. However, it is suboptimal on an expected profit basis.

To mitigate this suboptimality, in Chapter 4, Optimal Solution with Dynamic Programming, optimal solution approaches are studied. The two optimal approaches examined are the exact and Approximate Dynamic Programming (ADP). A comparison with FDO results is also presented. We show that using an ADP strategy, it is possible to increase the added value of FDO strategy by about 36% from \$7.85 to \$10.73. In contrast to FDO, the ADP algorithm can initiate under an adverse initial condition even if this implies an initial loss. Although ADP increases the expected profits relative to FDO, the former leads to a much higher risk with a larger standard deviation (4.6 times) and range (3.2 times). Interestingly, ADP tends to adjust the maturity of the forward contract much less frequently than FDO since it does not trade to capture every (small) profit.

The optimal value of the ADP method is validated against the exact Dynamic Programming method. Compared to the exact method, ADP estimates the option value accurately (less than 1% difference) and 39-170 times faster. The optimal decision is characterized in the two-dimensional domain of χ - ξ using both ADP and exact methods, and a good match is observed. In the region containing most realizations, the decision boundary is nearly vertical. This indicates that χ is the

more important factor in determining the decision. The reason is that χ has a stronger relationship with the forward curve slope. Mapping the optimal decision into the domain of the spot price and the forward curve slope allows to uncover the financial intuition behind the policy. It reveals that the trade is often terminated by selling out the inventory on the spot market when the slope is smaller than a usually negative threshold, i.e. make a profit by receiving the high spot price and paying the low forward price.

It is shown theoretically in Proposition 4.I that a partial sale (splitting the quantity sold between the spot and forward markets) is never optimal. Thus, this flexibility neither adds any value nor changes the optimal policy. The intuition behind this proposition might be explained by the absence of a constraint on the amount of oil that can be sold from the inventory within one period. The proposition is also verified computationally via both the exact and ADP methods. The results of allowing the potential for partial sale ($\Delta R = 1/3$ vs 1) support the proposition, and show a reduction of 2-6 times in computational time. Furthermore, it is proved in Proposition 4.II that the optimal actions are limited to a much smaller subset of the feasible set, $(a_i^R, a_i^T) \in \{(R_i, 0), (0, t_N), (0, t_{i+1})\}$. Accordingly, the corresponding Bellman's equation is expressed in terms of this smaller feasible action set and the forward curve slope. Using Proposition 4.II, the computation time is further reduced between 1.7-10 times. Chapter 4 concludes that the ADP approach can be the workhorse solution technique for the MDP's under consideration, and paves the road for the next chapter.

In Chapter 5, the problem framework is fundamentally expanded, which permits to formulate a richer and more realistic framework. Among the contributions of this chapter are studying a stochastic storage cost, the decision regarding the time of renting the storage, and the decisions with respect to time and quantity for buying/selling the oil (independent from the storage rental time).

The impact of the stochastic storage cost is quantified by computing the value under an assumption that the storage cost remains 'constant' and equal to the initial value at t_0 . In both constant and stochastic storage cost cases, the trader is faced with the exact same oil forward curve realizations. Comparison of the optimal values under the constant and stochastic storage cost, denoted by V and

V^c respectively, leads to the following conclusions; (i) the initial condition of the stochastic storage cost, $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$, impacts the direction in which it evolves dynamically, (ii) the direction in which the storage cost evolves affects the optimal decision of when to rent the storage, and (iii) the opportunities arising in the oil storage and forward markets are considered concurrently in decision-making to balance the tradeoffs between the two sides.

In Chapter 5, many $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$ combinations are selected that result in the same such non-stochastic one-year time-charter rate of Chapter 4, i.e. $TC(0,1) = \$6.57$. It is interesting to find that the option value in Chapter 4 (respectively \$10.73 and \$9.29 with full and zero refund of storage cost) is higher than the value under a constant storage cost in Chapter 5, $V^c = \$5.53$, since the framework assumptions are more relaxed in Chapter 4 (specifically the refund of storage cost at an early discharge). Using these different $(\chi'_0, \xi'_0, \alpha_0, \alpha_0^*)$ combination under the stochastic storage cost result in $\$5.53 \leq V \leq \7.79 in Chapter 5.

Although the ADP approach developed in the previous chapter is the foundation of the solution technique used in the present chapter, Chapter 5 also contributes to the development of an ADP approach via a novel Continuation Function Approximation (CFA) structure, for which the theoretical foundation is established in Proposition 5.I. This proposition exploits the irreversibility of the rent decision in the MDP state/action construct. It establishes that the value function depends on only the oil factors among all the stochastic factors comprising W_i if the state, x_i , indicates an ‘already-rented’ status. Consequently, the basis functions used in the CFA linear structure are adjusted according to x_i , which is an innovation of this research. The impact of the above arguments is also seen on the optimal policies; the decision boundary in χ - ξ plane may be sharp or blurry depending on x_i .

Finally, the sensitivity analysis of V and V^c to different parameters demonstrates that firstly V^c is more sensitive than V in general since any change will be persistent in time (will not evolve through the stochastic cost dynamics). This highlights the importance of including stochastic storage costs to be able to monitor the storage costs in search of an optimal initiation time. Secondly, oil-related parameters are more influential than the storage cost (other than the initial

condition) potentially since the rent decision is made once as opposed to being revisited periodically, i.e. as soon as there is a tanker in place the profit is derived by oil-related trades.

6.2 Principal Contributions

The main contributions of this thesis are summarized as follows:

- I. Extending the spot market trading to both the spot and forward markets in real options involving storage assets, which leads to a joint optimization of the forward maturity and inventory management decisions
- II. Establishing underperformance of the suboptimal heuristic known to perform well in the gas storage valuation context (Chapter 3 and Chapter 4), and offering the optimal trading strategies (Chapter 4 and Chapter 5)
- III. Quantifying the risk and reward characteristics of liquidation under a constantly hedging forward contract compared to other possibilities, e.g. an optimal sale on the spot market (Chapter 4)
- IV. Formulating a realistic framework by introducing stochastic storage costs, which allows studying optimal timing of initiating the trade with storage refill options, and comparing the impact of stochastic versus constant storage costs (Chapter 5)
- V. Characterizing the following properties of the optimal policy; a partial sale of the inventory on the spot and forward market is not optimal, and the optimally selected maturities belong to a small subset of the feasible set (proved in Chapter 4 and numerically observed in Chapter 5)

6.3 Future Work

One avenue for future research is considering a model that incorporates all the stochastic factors corresponding to oil prices and storage costs jointly in a single framework. Such a model can be beneficial in understanding the impact of any correlation between the oil and tanker markets on the trading policy. However, as it was mentioned in the literature review, developing such model would be challenging since the relationship between the oil and tanker market is influenced by many other variables.

It is found that the ADP method provides riskier policies comparing to FDO approach. Accordingly, another perspective to the existing (or similar) problems is through the use of a different objective function. For instance, as opposed to solely maximize the profits in dollar terms,

a risk measure could be minimized, or risk-aversion can be introduced via a utility function of the profits.

Ultimately, considering other commodities would be another promising avenue for further research; for instance in the case of agricultural commodities, Fackler and Livingston (2002) study optimal storage for Illinois soybeans. In particular, a setting where a producer (farmer) anticipates some random (potentially climate-related) crop production level periodically would be interesting. The producer maximizes the profits or minimizes periodic revenue volatilities by optimally selecting the time and quantity of the inventory sold. These actions can impact the prices if the market impact of producers is large such as in an oligopoly. For example, on a collective basis, the efforts and market power of producers aiming for such an outcome has led to the formation of *The Federation of Quebec Maple Syrup Producers*, called “the OPEC of the maple syrup world” by The Economist (“Sticky fingers,” 2013, “The Great Canadian Maple Syrup Heist,” 2013). Their inventory capacity is about 165,000 barrels each valuing US\$1200 approximately.

7 References

- Adland, R., Cariou, P., Wolff, F.-C., 2016. The influence of charterers and owners on bulk shipping freight rates. *Transp. Res. Part E Logist. Transp. Rev.* 86, 69–82. <https://doi.org/10.1016/j.tre.2015.11.014>
- Alizadeh, A.H., Huang, C.-Y., van Dellen, S., 2015. A regime switching approach for hedging tanker shipping freight rates. *Energy Econ.* 49, 44–59. <https://doi.org/10.1016/j.eneco.2015.01.004>
- Alizadeh, A.H., Nomikos, N.K., 2009. Shipping derivatives and risk management, *Shipping Derivatives and Risk Management*. <https://doi.org/10.1057/9780230235809>
- Alizadeh, A.H., Nomikos, N.K., 2004. Cost of carry, causality and arbitrage between oil futures and tanker freight markets. *Transp. Res. Part E Logist. Transp. Rev., Shipping Finance and Port Issues* 40, 297–316. <https://doi.org/10.1016/j.tre.2004.02.002>
- Barndorff-Nielsen, O.E., Shepard, N., 2001. Non-Gaussian Ornstein-Uhlenbeck-based models and some of their uses in financial economics. *J. R. Stat. Soc. Ser. B Stat. Methodol.* 63, 167–207.
- Benth, F.E., Koekebakker, S., Taib, C.M.I.C., 2015. Stochastic dynamical modelling of spot freight rates. *IMA J. Manag. Math.* 26, 273–297. <https://doi.org/10.1093/imaman/dpu001>
- Bertsekas, D.P., 2012. *Dynamic Programming and Optimal Control*, 4th edition. ed. Athena Scientific, Belmont, Massachusetts.
- Bhattacharya, S., 1978. Project Valuation with Mean-Reverting Cash Flow Streams. *J. Finance* 33, 1317–1331. <https://doi.org/10.1111/j.1540-6261.1978.tb03422.x>
- BIMCO, 2016. Tanker shipping: still a strong market as demand stays high, *Market Analysis Report 2016*. Baltic and International Maritime Council (BIMCO), Clarksons.

- Bjerksund, P., Stensland, G., 2014. Closed form spread option valuation. *Quant. Finance* 14, 1785–1794. <https://doi.org/10.1080/14697688.2011.617775>
- Bjerksund, P., Stensland, G., Vagstad, F., 2011. Gas Storage Valuation: Price Modelling v. Optimization Methods. *Energy J.* 32, 203–227.
- Boogert, A., de Jong, C., 2011. Gas storage valuation using a multifactor price process. *J. Energy Mark. Lond.* 4, 29–52. <https://doi.org/10.21314/JEM.2011.067>
- Boogert, A., de Jong, C., 2008. Gas Storage Valuation Using a Monte Carlo Method. *J. Deriv. N. Y.* 15, 81–98. <https://doi.org/10.3905/jod.2008.702507>
- Brown, D.B., Smith, J.E., Sun, P., 2010. Information Relaxations and Duality in Stochastic Dynamic Programs. *Oper. Res.* 58, 785–801. <https://doi.org/10.1287/opre.1090.0796>
- Carmona, R., Ludkovski, M., 2010. Valuation of energy storage: an optimal switching approach. *Quant. Finance* 10, 359–374. <https://doi.org/10.1080/14697680902946514>
- Carmona, R., Ludkovski, M., 2004. Spot convenience yield models for the energy markets. *Contemp. Math.* 351, 65–80.
- Carriere, J.F., 1996. Valuation of the early-exercise price for options using simulations and nonparametric regression. *Insur. Math. Econ.* 19, 19–30. [https://doi.org/10.1016/S0167-6687\(96\)00004-2](https://doi.org/10.1016/S0167-6687(96)00004-2)
- Casassus, J., Collin-Dufresne, P., 2005. Stochastic Convenience Yield Implied from Commodity Futures and Interest Rates. *J. Finance* 60, 2283–2331. <https://doi.org/10.1111/j.1540-6261.2005.00799.x>
- Charles R. Weber Company, Inc. [WWW Document], 2016. URL <http://www.crweber.com/markets1.html> (accessed 11.8.16).

- Chen, P.-F., Lee, C.-C., Zeng, J.-H., 2014. The relationship between spot and futures oil prices: Do structural breaks matter? *Energy Econ.* 43, 206–217. <https://doi.org/10.1016/j.eneco.2014.03.006>
- Chen, Z., Forsyth, P.A., 2007. A Semi-Lagrangian Approach for Natural Gas Storage Valuation and Optimal Operation. *SIAM J. Sci. Comput. Phila.* 30, 30. <http://dx.doi.org.proxy1.lib.uwo.ca/10.1137/060672911>
- Chinn, M.D., LeBlanc, M., Coibion, O., 2005. The Predictive Content of Energy Futures: An Update on Petroleum, Natural Gas, Heating Oil and Gasoline (Working Paper No. 11033). National Bureau of Economic Research.
- Considine, T.J., Larson, D.F., 2001. Uncertainty and the convenience yield in crude oil price backwardations. *Energy Econ.* 23, 533–548. [https://doi.org/10.1016/S0140-9883\(01\)00077-9](https://doi.org/10.1016/S0140-9883(01)00077-9)
- Cortazar, G., Naranjo, L., 2006. An N-factor gaussian model of oil futures prices. *J. Futur. Mark.* 26, 243–266. <https://doi.org/10.1002/fut.20198>
- Dias, M., 2005. Real options with petroleum applications. Doctoral thesis, Industrial engineering department, PUC-Rio, Rio de Janeiro, Brazil.
- Diaz-Rainey, I., Roberts, H., Lont, D.H., 2017. Crude inventory accounting and speculation in the physical oil market. *Energy Econ.* <https://doi.org/10.1016/j.eneco.2017.03.029>
- Dixit, A.K., Pindyck, R.S., 1994. *Investment Under Uncertainty*. Princeton University Press.
- Eydeland, A., Wolyniec, K., 2003. *Energy and Power Risk Management: New Developments in Modeling, Pricing, and Hedging*. John Wiley & Sons.
- Fackler, P.L., Livingston, M.J., 2002. Optimal Storage by Crop Producers. *Am. J. Agric. Econ.* 84, 645–659. <https://doi.org/10.2307/1244842>

- Felix, B.J., Weber, C., 2012. Gas storage valuation applying numerically constructed recombining trees. *Eur. J. Oper. Res.* 216, 178–187. <https://doi.org/10.1016/j.ejor.2011.07.029>
- Geman, H., 2005. *Commodities and Commodity Derivatives: Modelling and Pricing for Agriculturals, Metals and Energy*, 1 edition. ed. Wiley, West Sussex.
- Geman, H., Ohana, S., 2009. Forward curves, scarcity and price volatility in oil and natural gas markets. *Energy Econ.* 31, 576–585. <https://doi.org/10.1016/j.eneco.2009.01.014>
- Ghafari, B., Davison, M., 2017. A Forward Dynamic Optimization Strategy Under Contango Storage Arbitrage with Frictions. *J. Energy Mark.* 10, 59–85. <https://doi.org/10.21314/JEM.2017.166>
- Gibson, R., Schwartz, E.S., 1990. Stochastic Convenience Yield and the Pricing of Oil Contingent Claims. *J. Finance* 45, 959–976. <https://doi.org/10.1111/j.1540-6261.1990.tb05114.x>
- Glasserman, P., 2003. *Monte Carlo Methods in Financial Engineering*, 2003 edition. ed. Springer, New York.
- Glen, D.R., Martin, B.T., 2004. A survey of the modelling of dry bulk and tanker markets. *Res. Transp. Econ.* 12, 19–64. [https://doi.org/10.1016/S0739-8859\(04\)12002-7](https://doi.org/10.1016/S0739-8859(04)12002-7)
- Granger, C.W.J., 1969. Investigating Causal Relations by Econometric Models and Cross-spectral Methods. *Econometrica* 37, 424–438. <https://doi.org/10.2307/1912791>
- Gray, J., Khandelwal, P., 2004a. Towards a realistic gas storage model. *Commod. Now* 1–4.
- Gray, J., Khandelwal, P., 2004b. Realistic gas storage models II: Trading strategies. *Commod. Now* 1–5.
- Hahn, W.J., DiLellio, J.A., Dyer, J.S., 2014. What do market-calibrated stochastic processes indicate about the long-term price of crude oil? *Energy Econ.* 44, 212–221. <https://doi.org/10.1016/j.eneco.2014.04.007>

- Hilliard, J.E., Reis, J., 1998. Valuation of Commodity Futures and Options under Stochastic Convenience Yields, Interest Rates, and Jump Diffusions in the Spot. *J. Financ. Quant. Anal.* 33, 61–86. <https://doi.org/10.2307/2331378>
- Hull, J.C., 2014. *Options, Futures, and Other Derivatives*, 9 edition. ed. Pearson, Boston.
- Jafarizadeh, B., Bratvold, R., 2013. Sell Spot or Sell Forward? Analysis of Oil-Trading Decisions With the Two-Factor Price Model and Simulation. *SPE Econ. Manag.* 5, 80–88. <https://doi.org/10.2118/165581-PA>
- Jørgensen, P.L., de, G.D., 2010. Time charters with purchase options in shipping: Valuation and risk management. *Appl. Math. Finance* 17, 399–430. <https://doi.org/10.1080/13504860903388008>
- Kaldor, N., 1939. Speculation and Economic Stability. *Rev. Econ. Stud.* 7, 1–27. <https://doi.org/10.2307/2967593>
- Kalman, R.E., 1960. A New Approach to Linear Filtering and Prediction Problems. *J. Basic Eng.* 82, 35–45. <https://doi.org/10.1115/1.3662552>
- Kemp, J., 2016. Oil market is back in balance [WWW Document]. Reuters. URL <http://www.reuters.com/article/oil-global-kemp-idUSL8N18Z3HZ>
- Kirk, E., 1995. Correlation in the Energy Markets. *Manag. Energy Price Risk Risk Publications and Enron*, 71–78.
- Koekebakker, S., Adland, R., Sødal, S., 2007. Pricing freight rate options. *Transp. Res. Part E Logist. Transp. Rev.* 43, 535–548. <https://doi.org/10.1016/j.tre.2006.03.005>
- Koekebakker, S., Os, Å., 2004. Modelling forward freight rate dynamics - Empirical evidence from time charter rates. *Marit. Policy Manag.* 31, 319–335. <https://doi.org/10.1080/03088830410001674197>

- Korn, O., 2005. Drift matters: An analysis of commodity derivatives. *J. Futur. Mark.* 25, 211–241. <https://doi.org/10.1002/fut.20139>
- Lagoudakis, M.G., Parr, R., 2003. Least-squares Policy Iteration. *J Mach Learn Res* 4, 1107–1149.
- Lai, G., Margot, F., Secomandi, N., 2010. An Approximate Dynamic Programming Approach to Benchmark Practice-Based Heuristics for Natural Gas Storage Valuation. *Oper. Res.* 58, 564–582. <https://doi.org/10.1287/opre.1090.0768>
- Litzenberger, R.H., Rabinowitz, N., 1995. Backwardation in Oil Futures Markets: Theory and Empirical Evidence. *J. Finance* 50, 1517–1545. <https://doi.org/10.2307/2329325>
- Löhndorf, N., Wozabal, D., 2017. Indifference pricing of natural gas storage contracts. [WWW Document]. URL http://www.optimization-online.org/DB_FILE/2017/02/5863.pdf (accessed 5.11.17).
- Longstaff, F.A., Schwartz, E.S., 2001. Valuing American Options by Simulation: A Simple Least-Squares Approach. *Rev. Financ. Stud.* 14, 113–147. <https://doi.org/10.1093/rfs/14.1.113>
- Maragos, S., 2002. Valuation of the operational flexibility of natural gas storage reservoirs, in: Ronn, E. (Ed.), *Real Options and Energy Management*. Risk Books, London, pp. 431–456.
- Margrabe, W., 1978. The Value of an Option to Exchange One Asset for Another. *J. Finance* 33, 177–186. <https://doi.org/10.1111/j.1540-6261.1978.tb03397.x>
- Maslyuk, S., Smyth, R., 2009. Cointegration between oil spot and future prices of the same and different grades in the presence of structural change. *Energy Policy* 37, 1687–1693. <https://doi.org/10.1016/j.enpol.2009.01.013>
- Mirantes, A.G., Población, J., Serna, G., 2012. The Stochastic Seasonal Behaviour of Natural Gas Prices. *Eur. Financ. Manag.* 18, 410–443. <https://doi.org/10.1111/j.1468-036X.2009.00533.x>

- Nadarajah, S., Margot, F., Secomandi, N., 2017. Comparison of least squares Monte Carlo methods with applications to energy real options. *Eur. J. Oper. Res.* 256, 196–204. <https://doi.org/10.1016/j.ejor.2016.06.020>
- Nadarajah, S., Margot, F., Secomandi, N., 2015. Relaxations of Approximate Linear Programs for the Real Option Management of Commodity Storage. *Manag. Sci.* 61, 3054–3076. <https://doi.org/10.1287/mnsc.2014.2136>
- Pindyck, R.S., 1999. The Long-Run Evolution of Energy Prices. *Energy J.* 20, 1–27.
- Poblacion, J., 2015. The stochastic seasonal behavior of freight rate dynamics. *Marit. Econ. Logist.* 17, 142–162. <https://doi.org/10.1057/mel.2014.37>
- Poulakidas, A., Joutz, F., 2009. Exploring the link between oil prices and tanker rates 36, 215–233.
- Powell, W.B., 2011. *Approximate Dynamic Programming: Solving the Curses of Dimensionality*, 2nd Edition, 2 edition. ed. Wiley, Hoboken, N.J.
- Prokopczuk, M., 2011. Pricing and hedging in the freight futures market. *J. Futur. Mark.* 31, 440–464. <https://doi.org/10.1002/fut.20480>
- Puterman, M.L., 2005. *Markov decision processes: discrete stochastic dynamic programming*. Wiley-Interscience.
- Ribeiro, D., Hodges, S., 2004. A Two-Factor Model for Commodity Prices and Futures Valuation. Presented at the EFMA 2004 Basel Meetings, Social Science Research Network, Basel, Switzerland. <https://doi.org/10.2139/ssrn.498802>
- Sargent, T.J., 1986. *Rational Expectations and Inflation*. Harpercollins College Div, New York.
- Schoftner, R., 2008. On the estimation of credit exposures using regression-based Monte Carlo simulation. *J. Credit Risk* 4, 37–63. <https://doi.org/10.21314/jcr.2008.081>

- Schwartz, E.S., 1997. The Stochastic Behavior of Commodity Prices: Implications for Valuation and Hedging. *J. Finance* 52, 923–973. <https://doi.org/10.1111/j.1540-6261.1997.tb02721.x>
- Schwartz, E.S., Smith, J.E., 2000. Short-Term Variations and Long-Term Dynamics in Commodity Prices. *Manag. Sci.* 46, 893–911. <https://doi.org/10.1287/mnsc.46.7.893.12034>
- SEC Amendment No.5 to Form F-1 [WWW Document], n.d. URL <https://www.sec.gov/Archives/edgar/data/1617049/000119312514405337/d776617df1a.htm> (accessed 3.22.18).
- Secomandi, N., 2016. A tutorial on portfolio-based control algorithms for merchant energy trading operations. *J. Commod. Mark.* 4, 1–13. <https://doi.org/10.1016/j.jcomm.2016.10.003>
- Secomandi, N., 2015. Merchant commodity storage practice revisited. *Oper. Res.* 63, 1131–1143. <https://doi.org/10.1287/opre.2015.1407>
- Secomandi, N., 2010. Optimal commodity trading with a capacitated storage asset. *Manag. Sci.* 56, 449–467. <https://doi.org/10.1287/mnsc.1090.1049>
- Secomandi, N., Kekre, S., 2014. Optimal Energy Procurement in Spot and Forward Markets. *Manuf. Serv. Oper. Manag.* 16, 270–282. <https://doi.org/10.1287/msom.2013.0473>
- Secomandi, N., Lai, G., Margot, F., Scheller-Wolf, A., Seppi, D.J., 2015. Merchant commodity storage and term-structure model error. *Manuf. Serv. Oper. Manag.* 17, 302–320. <https://doi.org/10.1287/msom.2015.0518>
- Shi, W., Yang, Z., Li, K.X., 2013. The impact of crude oil price on the tanker market. *Marit. Policy Manag.* 40, 309–322. <https://doi.org/10.1080/03088839.2013.777981>
- Sødal, S., Koekebakker, S., Aadland, R., 2008. Market switching in shipping - A real option model applied to the valuation of combination carriers. *Rev. Financ. Econ.* 17, 183–203. <https://doi.org/10.1016/j.rfe.2007.04.001>

- Sødal, S., Koekebakker, S., Adland, R., 2009. Value based trading of real assets in shipping under stochastic freight rates. *Appl. Econ.* 41, 2793–2807. <https://doi.org/10.1080/00036840701720853>
- Sticky fingers, 2013. . *The Economist*.
- Sun, X., Tang, L., Yang, Y., Wu, D., Li, J., 2014. Identifying the dynamic relationship between tanker freight rates and oil prices: In the perspective of multiscale relevance. *Econ. Model.* 42, 287–295. <https://doi.org/10.1016/j.econmod.2014.06.019>
- Taib, C.M.I.C., 2016. Forward pricing in the shipping freight market. *Jpn. J. Ind. Appl. Math.* 33, 3–23. <https://doi.org/10.1007/s13160-015-0204-6>
- The Forward Curve for Oil Prices Suddenly Looks Awful for OPEC, 2017. . *Bloomberg.com*.
- The Great Canadian Maple Syrup Heist, 2013. . *Bloomberg.com*.
- Thompson, M., Davison, M., Rasmussen, H., 2009. Natural gas storage valuation and optimization: A real options application. *Nav. Res. Logist. NRL* 56, 226–238. <https://doi.org/10.1002/nav.20327>
- Tsitsiklis, J.N., Roy, B.V., 2001. Regression methods for pricing complex American-style options. *IEEE Trans. Neural Netw.* 12, 694–703. <https://doi.org/10.1109/72.935083>
- Tvedt, J., 1997. Valuation of VLCCs under income uncertainty. *Marit. Policy Manag.* 24, 159–174. <https://doi.org/10.1080/03088839700000067>
- Vasicek, O., 1977. An equilibrium characterization of the term structure. *J. Financ. Econ.* 5, 177–188.
- Ware, A., 2013. Accurate Semi-Lagrangian Time Stepping for Stochastic Optimal Control Problems with Application to the Valuation of Natural Gas Storage. *SIAM J. Financ. Math.* 4, 427–451. <https://doi.org/10.1137/110853546>

- Working, H., 1948. Theory of the Inverse Carrying Charge in Futures Markets. *J. Farm Econ.* 30, 1–28. <https://doi.org/10.2307/1232678>
- Yang, Y., Liu, C., Sun, X., Li, J., 2015. Spillover effect of international crude oil market on tanker market. *Int. J. Glob. Energy Issues* 38, 257–277. <https://doi.org/10.1504/IJGEI.2015.070270>
- Ye, S., Karali, B., 2016. The informational content of inventory announcements: Intraday evidence from crude oil futures market. *Energy Econ.* 59, 349–364. <https://doi.org/10.1016/j.eneco.2016.08.011>

8 Appendix A

This appendix presents the detailed proofs of the lemmas and propositions discussed in the paper.

8.1 Proof of Lemma 4.I:

Lemma 4.I: The reward function $r_i(a_i, x_i, W_i)$, Eq. 8.1, is neither concave nor convex with respect to (a_i^R, a_i^T) .

Proof:

The reward function is expressed by the following equation.

$$\begin{aligned}
 r_i(a_i, x_i, W_i) &= a_i^R [F(t_i, t_i) - c_P - c_H(t_i - T_i)] \\
 &\quad + (R_i - a_i^R) \left[e^{-r(a_i^T - t_i)} [F(t_i, a_i^T) - c_P] - c_H(a_i^T - T_i) \right] \\
 &\quad - R_i e^{-r(T_i - t_i)} [F(t_i, T_i) - c_P]
 \end{aligned} \tag{Eq. 8.1}$$

$$\begin{aligned}
 &= a_i^R \left[F(t_i, t_i) - c_P - e^{-r(a_i^T - t_i)} [F(t_i, a_i^T) - c_P] + c_H(a_i^T - t_i) \right] \\
 &\quad + R_i \left[e^{-r(a_i^T - t_i)} [F(t_i, a_i^T) - c_P] - c_H(a_i^T - T_i) \right] \\
 &\quad - R_i e^{-r(T_i - t_i)} [F(t_i, T_i) - c_P]
 \end{aligned}$$

The partial derivatives of r_i can be computed as follows, where $F' := \partial F(t_i, a_i^T) / \partial a_i^T$ and $F'' := \partial^2 F(t_i, a_i^T) / \partial (a_i^T)^2$.

$$\frac{\partial r_i}{\partial a_i^R} = F(t_i, t_i) - c_P - e^{-r(a_i^T - t_i)} [F(t_i, a_i^T) - c_P] + c_H(a_i^T - t_i) \tag{Eq. 8.2}$$

$$\frac{\partial r_i}{\partial a_i^T} = (R_i - a_i^R) \left[e^{-r(a_i^T - t_i)} \left[-r(F(t_i, a_i^T) - c_P) + F' \right] - c_H \right] \tag{Eq. 8.3}$$

$$\frac{\partial^2 r_i}{\partial (a_i^R)^2} = 0 \tag{Eq. 8.4}$$

$$\frac{\partial^2 r_i}{\partial (a_i^T)^2} = (R_i - a_i^R) e^{-r(a_i^T - t_i)} [r^2(F(t_i, a_i^T) - c_P) - 2rF' + F''] \quad \text{Eq. 8.5}$$

$$\frac{\partial^2 r_i}{\partial a_i^R \partial a_i^T} = \frac{\partial^2 r_i}{\partial a_i^T \partial a_i^R} = -[e^{-r(a_i^T - t_i)} [-r(F(t_i, a_i^T) - c_P) + F'] - c_H] \quad \text{Eq. 8.6}$$

According to Eq. 8.1, r_i is linear with respect to $a_i^R \in [0, R_i]$, and it can be maximized as follows; if $\partial r_i / \partial a_i^R > 0$, then $a_i^R = R_i$ is optimum, and the only feasible choice for the other decision variable is $a_i^T = 0$. On the other hand, if $\partial r_i / \partial a_i^R \leq 0$, then $a_i^R = 0$ is optimum, and a_i^T must be chosen optimally by solving $\partial r_i / \partial a_i^T = 0$, which is the solution of Eq. 8.7.

$$F' - r(F(t_i, a_i^T) - c_P) = e^{r(a_i^T - t_i)} c_H \quad \text{Eq. 8.7}$$

The Hessian matrix at stage i can be written as in Eq. 8.8.

$$H_i = \begin{bmatrix} 0 & b \\ b & d \end{bmatrix} := \begin{bmatrix} 0 & \frac{\partial^2 r_i}{\partial a_i^R \partial a_i^T} \\ \frac{\partial^2 r_i}{\partial a_i^T \partial a_i^R} & \frac{\partial^2 r_i}{\partial (a_i^T)^2} \end{bmatrix} = \quad \text{Eq. 8.8}$$

$$\begin{bmatrix} 0 & e^{-r(a_i^T - t_i)} [r(F(t_i, a_i^T) - c_P) - F'] + c_H \\ e^{-r(a_i^T - t_i)} [r(F(t_i, a_i^T) - c_P) - F'] + c_H & (R_i - a_i^R) e^{-r(a_i^T - t_i)} [r^2(F(t_i, a_i^T) - c_P) - 2rF' + F''] \end{bmatrix}$$

By computing the eigen-values of H_i , Eq. 8.9, it might be shown that $r_i(a_i, x_i, W_i)$ is neither concave or convex by Eq. 8.10.

$$\lambda_1 = \frac{1}{2} [-\sqrt{4b^2 + d^2} + d], \quad \lambda_2 = \frac{1}{2} [\sqrt{4b^2 + d^2} + d] \quad \text{Eq. 8.9}$$

So, $\lambda_1 \lambda_2$ can be expressed by Eq. 8.10.

$$\lambda_1 \lambda_2 = -b^2 \leq 0 \quad \text{Eq. 8.10}$$

Therefore, r_i is neither concave nor convex. Furthermore, Table 8.1 summarizes all possible scenarios regarding the signs of λ_1 and λ_2 , and discusses the corresponding solutions (if any) based on Eq. 8.9.

$\lambda_1 > 0$	No solution
$\lambda_1 = 0$	$b = 0$ and $d \in \mathbb{R}^+$
$\lambda_1 < 0$	$b = 0$ and $d \in \mathbb{R}^-$, or $b \neq 0$ and $d \in \mathbb{R}$
$\lambda_2 > 0$	$b = 0$ and $d \in \mathbb{R}^+$, or $b \neq 0$ and $d \in \mathbb{R}$
$\lambda_2 = 0$	$b = 0$ and $d \in \mathbb{R}^-$
$\lambda_2 < 0$	No solution

Table 8.1. Solution to inequalities based on Eq. 8.9.

From the table, it is deduced that the condition $\lambda_1 \geq 0$ and $\lambda_2 \geq 0$ can hold when

$$b = 0 \text{ and } d \in \mathbb{R}^+ \text{ (resulting in } \lambda_1 = 0 \text{ and } \lambda_2 > 0), \quad \text{Eq. 8.11}$$

and, the condition $\lambda_1 \leq 0$ and $\lambda_2 \leq 0$ can hold when

$$b = 0 \text{ and } d \in \mathbb{R}^- \text{ (resulting in } \lambda_1 < 0 \text{ and } \lambda_2 = 0). \quad \text{Eq. 8.12}$$

Either of the two above cases requires b to be constantly zero. This is not possible because $b = 0$ together with Eq. 8.3 and Eq. 8.6 imply that $\frac{\partial r_i}{\partial a_i^T} = 0$, which is in contrast to the stage reward function assumptions (the reward function is not constant with respect to a_i^R and a_i^T). ■

8.2 Proof of Lemma 4.II:

Lemma 4.II lemma has two parts:

(i) The value function $V_i(x_i, W_i)$ can be written the form of $R_i v_i(T_i, W_i)$, i.e. a multiple of R_i , $\forall i \in \mathcal{J} \setminus \{N\}$.

(ii) If $V_{i+1}(x_{i+1}, W_{i+1}) = R_{i+1} v_{i+1}(T_{i+1}, W_{i+1})$, then at stage i any action $(a_i^R, a_i^T) = (0 < a_i^R < R_i, a_i^T)$ is dominated by either $(0, a_i^T)$ or $(R_i, 0)$.

Proof:

We will prove part (i) of the lemma by backward induction, and as part of this process, we will prove part (ii) as well. Let us consider the Bellman equation, Eq. 4.8, at time stage $i = N - 1$, where the continuation value is zero.

$$\begin{aligned}
V_{N-1}(x_{N-1}, W_{N-1}) &= \max_{a \in \mathcal{A}_{N-1}(x_{N-1})} r_{N-1}(a, x_{N-1}, W_{N-1}) = \\
&\max_{a \in \mathcal{A}_{N-1}(x_{N-1})} a_{N-1}^R [F(t_{N-1}, t_{N-1}) - c_P - c_H(t_{N-1} - T_{N-1})] \\
&\quad + (R_{N-1} - a_{N-1}^R) \left[e^{-r(a_{N-1}^T - t_{N-1})} [F(t_{N-1}, a_{N-1}^T) - c_P] \right. \\
&\quad \left. - c_H(a_{N-1}^T - T_{N-1}) \right] \\
&\quad - R_{N-1} e^{-r(T_{N-1} - t_{N-1})} [F(t_{N-1}, T_{N-1}) - c_P]
\end{aligned} \tag{Eq. 8.13}$$

To solve the maximization, the reward function r_{N-1} can be expressed as the following, which shows that it is linear with respect to a_{N-1}^R .

$$\begin{aligned}
r_{N-1}(a_{N-1}, x_{N-1}, W_{N-1}) &= \\
&a_{N-1}^R \left[F(t_{N-1}, t_{N-1}) - c_P - e^{-r(a_{N-1}^T - t_{N-1})} [F(t_{N-1}, a_{N-1}^T) - c_P] \right. \\
&\quad \left. + c_H(a_{N-1}^T - t_{N-1}) \right] \\
&\quad + R_{N-1} \left[e^{-r(a_{N-1}^T - t_{N-1})} [F(t_{N-1}, a_{N-1}^T) - c_P] \right. \\
&\quad \left. - c_H(a_{N-1}^T - T_{N-1}) \right] \\
&\quad - R_{N-1} e^{-r(T_{N-1} - t_{N-1})} [F(t_{N-1}, T_{N-1}) - c_P]
\end{aligned} \tag{Eq. 8.14}$$

Because r_{N-1} is linear with respect to $a_{N-1}^R \in [0, R_{N-1}]$, $a_{N-1}^R = 0$ if $\partial r_{N-1} / \partial a_{N-1}^R \leq 0$, and $a_{N-1}^R = R_{N-1}$ if $\partial r_{N-1} / \partial a_{N-1}^R > 0$, respectively generating the payoffs p_0 and p_1 defined by Eq. 8.16 and Eq. 8.17. Note that the feasible set is $(a_{N-1}^R, a_{N-1}^T) \in [0, R_{N-1}) \times \{t_N\} \cup \{R_{N-1}\} \times \{0\}$. If $a_{N-1}^R = 0$ is optimal the only feasible a_{N-1}^T will be t_N , whereas if $a_{N-1}^R = R_{N-1}$ is optimal the only feasible a_{N-1}^T will be 0. The condition $\partial r_{N-1} / \partial a_{N-1}^R > 0$ can be written in terms of the more intuitive slope parameter m_{N-1} as expressed by Eq. 8.15.

$$m_{N-1} := \frac{e^{-r(a_i^T - t_i)} [F(t_i, a_i^T) - c_P] - [F(t_i, t_i) - c_P]}{a_i^T - t_i} < c_H \tag{Eq. 8.15}$$

Rewrite the payoff when $m_{N-1} \geq c_H$ as

$$\begin{aligned}
p_0 &:= r_{N-1}((0, t_N), x_{N-1}, W_{N-1}) \\
&= R_{N-1}[e^{-r\Delta t}[F(t_{N-1}, t_N) - c_P] - c_H(t_N - T_{N-1})] \\
&\quad - R_{N-1}e^{-r(T_{N-1}-t_{N-1})}[F(t_{N-1}, T_{N-1}) - c_P] \\
&= R_{N-1}\left[e^{-r\Delta t}[F(t_{N-1}, t_N) - c_P] - c_H(t_N - T_{N-1})\right. \\
&\quad \left.- e^{-r(T_{N-1}-t_{N-1})}[F(t_{N-1}, T_{N-1}) - c_P]\right]
\end{aligned} \tag{Eq. 8.16}$$

and, when $m_{N-1} < c_H$

$$\begin{aligned}
p_1 &:= r_{N-1}((R_{N-1}, 0), x_{N-1}, W_{N-1}) \\
&= R_{N-1}[F(t_{N-1}, t_{N-1}) - c_P - c_H(t_{N-1} - T_{N-1})] \\
&\quad - R_{N-1}e^{-r(T_{N-1}-t_{N-1})}[F(t_{N-1}, T_{N-1}) - c_P] \\
&= R_{N-1}\left[F(t_{N-1}, t_{N-1}) - c_P - c_H(t_{N-1} - T_{N-1}) - e^{-r(T_{N-1}-t_{N-1})}[F(t_{N-1}, T_{N-1}) - c_P]\right].
\end{aligned} \tag{Eq. 8.17}$$

Now, $V_{N-1}(x_{N-1}, W_{N-1})$ can be written in terms of p_0 and p_1 as in Eq. 8.18, which subsequently implies that $V_{N-1}(x_{N-1}, W_{N-1})$ is a multiple of R_{N-1} just like both terms of the sum.

$$V_{N-1}(x_{N-1}, W_{N-1}) = p_0\mathbb{I}(m_{N-1} \geq c_H) + p_1\mathbb{I}(m_{N-1} < c_H) \tag{Eq. 8.18}$$

Now, assume that at stage $i + 1$, the value function can be written in the form of

$$V_{i+1}(x_{i+1}, W_{i+1}) = V_{i+1}((R_{i+1}, T_{i+1}), W_{i+1}) = R_{i+1}v_{i+1}(T_{i+1}, W_{i+1}). \tag{Eq. 8.19}$$

We want to show that the same holds true at stage i ; that is $V_i(x_i, W_i) = R_i v_i(T_i, W_i)$. The Bellman equation at stage i yields

$$V_i(x_i, W_i) = \max_{a \in \mathcal{A}_i(x_i)} \{r_i(a, x_i, W_i) + \delta \mathbb{E}[V_{i+1}(f_i(x_i, a), W_{i+1}) | W_i]\} \tag{Eq. 8.20}$$

$$= \max_{a \in \mathcal{A}_i(x_i)} \left\{ a_i^R [F(t_i, t_i) - c_P - c_H(t_i - T_i)] + (R_i - a_i^R) \left[e^{-r(a_i^T - t_i)} [F(t_i, a_i^T) - c_P] - c_H(a_i^T - T_i) \right] - R_i e^{-r(T_i - t_i)} [F(t_i, T_i) - c_P] + (R_i - a_i^R) \delta \mathbb{E}[v_{i+1}(T_{i+1}, W_{i+1}) | W_i] \right\}.$$

In the following, we argue that that any action $(a_i^R, a_i^T) = (0 < a_i^R < R_i, a_i^T)$ is dominated by either $(0, a_i^T)$ or $(R_i, 0)$.

I. The payoff from $(a_i^R, a_i^T) = (0, a_i^T)$ is

$$p_1(a_i^T) := R_i \left[e^{-r(a_i^T - t_i)} [F(t_i, a_i^T) - c_P] - c_H(a_i^T - T_i) \right] - R_i e^{-r(T_i - t_i)} [F(t_i, T_i) - c_P] + R_i \delta \mathbb{E}[v_{i+1}(a_i^T, W_{i+1}) | W_i]. \quad \text{Eq. 8.21}$$

II. The payoff from $(a_i^R, a_i^T) = (R_i, 0)$ is

$$p_2 := R_i [F(t_i, t_i) - c_P - c_H(t_i - T_i)] - R_i e^{-r(T_i - t_i)} [F(t_i, T_i) - c_P]. \quad \text{Eq. 8.22}$$

III. The payoff from $(a_i^R, a_i^T) = (0 < a_i^R < R_i, a_i^T)$ is

$$p_3(a_i^R, a_i^T) := a_i^R [F(t_i, t_i) - c_P - c_H(t_i - T_i)] + (R_i - a_i^R) \left[e^{-r(a_i^T - t_i)} [F(t_i, a_i^T) - c_P] - c_H(a_i^T - T_i) \right] - R_i e^{-r(T_i - t_i)} [F(t_i, T_i) - c_P] + (R_i - a_i^R) \delta \mathbb{E}[v_{i+1}(a_i^T, W_{i+1}) | W_i]. \quad \text{Eq. 8.23}$$

One can compute the payoff differences as

$$\begin{aligned} p_1(a_i^T) - p_3(a_i^R, a_i^T) &= \\ &= -a_i^R [F(t_i, t_i) - c_P - c_H(t_i - T_i)] + a_i^R \left[e^{-r(a_i^T - t_i)} [F(t_i, a_i^T) - c_P] - c_H(a_i^T - T_i) \right] \\ &\quad + a_i^R \delta \mathbb{E}[v_{i+1}(a_i^T, W_{i+1}) | W_i] = \\ &= a_i^R \left\{ -[F(t_i, t_i) - c_P - c_H(t_i - T_i)] + e^{-r(a_i^T - t_i)} [F(t_i, a_i^T) - c_P] - c_H(a_i^T - T_i) + \right. \\ &\quad \left. \delta \mathbb{E}[v_{i+1}(a_i^T, W_{i+1}) | W_i] \right\} = a_i^R C_i. \end{aligned} \quad \text{Eq. 8.24}$$

Here, C_i is defined as follows, and can be positive or negative depending on a_i^T (and other variables such as W_i).

$$C_i := -[F(t_i, t_i) - c_P - c_H(t_i - T_i)] + e^{-r(a_i^T - t_i)}[F(t_i, a_i^T) - c_P] - c_H(a_i^T - T_i) + \delta \mathbb{E}[v_{i+1}(a_i^T, W_{i+1}) | W_i] \quad \text{Eq. 8.25}$$

The other payoff difference is computed as

$$p_2 - p_3(a_i^R, a_i^T) = (R_i - a_i^R)[F(t_i, t_i) - c_P - c_H(t_i - T_i)] - (R_i - a_i^R) \left[e^{-r(a_i^T - t_i)}[F(t_i, a_i^T) - c_P] - c_H(a_i^T - T_i) \right] - (R_i - a_i^R) \delta \mathbb{E}[v_{i+1}(a_i^T, W_{i+1}) | W_i] = -(R_i - a_i^R)C_i. \quad \text{Eq. 8.26}$$

In summary, for any $(a_i^R, a_i^T) \in [0, R_i) \times \{t_{i+1}, t_{i+2}, t_{i+3}, \dots, t_N\}$, we have

$$p_1(a_i^T) - p_3(a_i^R, a_i^T) = a_i^R C_i, \quad \text{Eq. 8.27}$$

$$p_2 - p_3(a_i^R, a_i^T) = -(R_i - a_i^R)C_i. \quad \text{Eq. 8.28}$$

Thus, action **III**, i.e. $(a_i^R, a_i^T) = (0 < a_i^R < R_i, a_i^T)$, is never optimal because $p_1 \geq p_3 \geq p_2$ if $C_i \geq 0$, and $p_2 > p_3 > p_1$ if $C_i < 0$. Given that the three possible actions **I**, **II**, and **III** partition the feasible set $\mathcal{A}_i(x_i) = [0, R_i) \times \{t_{i+1}, t_{i+2}, t_{i+3}, \dots, t_N\} \cup \{R_i\} \times \{0\}$ into disjoint subsets, we can rewrite Eq. 8.20, i.e. Bellman equation at stage i , as below.

$$V_i(x_i, W_i) = \max_{a_i^T \in \{t_{i+1}, t_{i+2}, t_{i+3}, \dots, t_N\}} p_1(a_i^T) \mathbb{I}(0 \leq C_i(a_i^T)) + p_2 \mathbb{I}(C_i(a_i^T) < 0) = \max_{a_i^T \in \{t_{i+1}, t_{i+2}, t_{i+3}, \dots, t_N\}} \left\{ R_i \left[e^{-r(a_i^T - t_i)}[F(t_i, a_i^T) - c_P] - c_H(a_i^T - T_i) \right] - R_i e^{-r(T_i - t_i)}[F(t_i, T_i) - c_P] + R_i \delta \mathbb{E}[v_{i+1}(a_i^T, W_{i+1}) | W_i] \right\} \mathbb{I}(0 \leq C_i(a_i^T)) + \left\{ R_i [F(t_i, t_i) - c_P - c_H(t_i - T_i)] - R_i e^{-r(T_i - t_i)}[F(t_i, T_i) - c_P] \right\} \mathbb{I}(C_i(a_i^T) < 0) = \quad \text{Eq. 8.29}$$

$$\begin{aligned}
R_i \max_{a_i^T \in \{t_{i+1}, t_{i+2}, t_{i+3}, \dots, t_N\}} & \left\{ e^{-r(a_i^T - t_i)} [F(t_i, a_i^T) - c_P] - c_H(a_i^T - T_i) \right. \\
& \left. - e^{-r(T_i - t_i)} [F(t_i, T_i) - c_P] + \delta \mathbb{E}[v_{i+1}(a_i^T, W_{i+1}) | W_i] \right\} \mathbb{I}(0 \\
& \leq C_i(a_i^T)) \\
& + \{F(t_i, t_i) - c_P - c_H(t_i - T_i) \\
& - e^{-r(T_i - t_i)} [F(t_i, T_i) - c_P]\} \mathbb{I}(C_i(a_i^T) < 0)
\end{aligned}$$

If we define,

$$\begin{aligned}
v_i(T_i, W_i) := & \\
\max_{a_i^T \in \{t_{i+1}, t_{i+2}, t_{i+3}, \dots, t_N\}} & \left\{ e^{-r(a_i^T - t_i)} [F(t_i, a_i^T) - c_P] - c_H(a_i^T - T_i) \right. \\
& \left. - e^{-r(T_i - t_i)} [F(t_i, T_i) - c_P] + \delta \mathbb{E}[v_{i+1}(a_i^T, W_{i+1}) | W_i] \right\} \mathbb{I}(0 \\
& \leq C_i(a_i^T)) \\
& + \{F(t_i, t_i) - c_P - c_H(t_i - T_i) \\
& - e^{-r(T_i - t_i)} [F(t_i, T_i) - c_P]\} \mathbb{I}(C_i(a_i^T) < 0)
\end{aligned} \tag{Eq. 8.30}$$

By substituting Eq. 8.30 into Eq. 8.29, it is proved that

$$V_i(x_i, W_i) = R_i v_i(T_i, W_i). \tag{Eq. 8.31}$$

■

8.3 Proof of Proposition 4.I:

Proposition 4.I: In the SDP problem set out by Eq. 4.7, partial sale of the inventory is never optimal. That is for $\forall i \in \mathcal{J} \setminus \{N\}$, any action $(a_i^R, a_i^T) = (0 < a_i^R < R_i, a_i^T) \in \mathcal{A}_i(x_i)$ is dominated by the action $(0, a_i^T)$ or $(R_i, 0)$.

Proof:

By Lemma 4.II part (i), the value function $V_i(x_i, W_i)$ is a multiple of R_i , $\forall i \in \mathcal{J} \setminus \{N\} = \{0, 1, 2, \dots, N-1\}$. We also know that $V_N(x_N, W_N) = 0$, and thus a multiple of R_N . By part (ii) of Lemma 4.II, if $V_{i+1}(x_{i+1}, W_{i+1})$ is a multiple of R_{i+1} , then at the stage i any action $(a_i^R, a_i^T) = (0 < a_i^R < R_i, a_i^T)$ is dominated by either $(0, a_i^T)$ or $(R_i, 0)$. Combining parts (i) and (ii), it is

concluded that for $\forall i \in \{0, 1, 2, \dots, N - 1\}$ any action $(a_i^R, a_i^T) = (0 < a_i^R < R_i, a_i^T)$ is dominated by either $(0, a_i^T)$ or $(R_i, 0)$. ■

In the following, *Proposition 4.II* is proved, which formulates the structural form of the value function and optimal decisions. To prove this proposition, we assume that Assumption 4.I (stated below) holds.

8.4 Assumption 4.I:

In the following section, *Proposition II* is proved, which formulates the structural form of the value function and optimal decisions. To prove this proposition, Assumption 4.I is used, according to which it is assumed that one can estimate the following spread linearly using a coefficient m_i .

$$e^{-r(a-t_i)}[F(t_i, a) - c_p] - e^{-r(b-t_i)}[F(t_i, b) - c_p] \approx m_i(a - b), \quad a, b \in [0, \bar{T}] \quad \text{Eq. 8.32}$$

Let $Z = T - t$ denote the time-to-maturity and let $f(z) = e^{-rz}[F(t_i, t_i + z) - c_p]$. Using the forward curve model, Eq. 2.13, the first order Taylor expansion of $f(z)$ around $z = 0$, equivalent to small time-to-maturities $T - t \leq 1$, is expressed in Eq. 8.33.

$$f(z) \approx S_t - c_p + \quad \text{Eq. 8.33}$$

$$z \left[S_t \left(-k\chi_t + \frac{\sigma_\chi^2}{2} + 2\sigma_\chi\sigma_\xi\rho_{\chi\xi} + \frac{\sigma_\xi^2}{2} - \lambda_\chi + \mu_\xi - \lambda_\xi - r \right) + rc_p \right] + O(z^2)$$

Using Eq. 8.33, the spread can be approximated by Eq. 8.34. According to this equation, the spread at time t_i is proportional to the time gap, $a - b$, the spot price $S_{t_i} = \exp(\chi_{t_i} + \xi_{t_i})$, and the short-term factor χ_{t_i} .



$$e^{-r(a-t_i)}[F(t_i, a) - c_p] - e^{-r(b-t_i)}[F(t_i, b) - c_p] \approx \quad \text{Eq. 8.34}$$

$$(a - b) \left[S_{t_i} \left(-k\chi_{t_i} + \frac{\sigma_\chi^2}{2} + 2\sigma_\chi\sigma_\xi\rho_{\chi\xi} + \frac{\sigma_\xi^2}{2} - \lambda_\chi + \mu_\xi - \lambda_\xi - r \right) + rc_p \right]$$

To test the accuracy of the above formula in a worst-case scenario, let $a = t_i$ and $b = \bar{T} = 1$; it is the largest possible spread at each time t_i , where a linear approximation causes the maximum error. The actual spread (left-hand side of Eq. 8.34) and the approximated spread (right-hand side of Eq. 8.34) are computed for 10,000 sample paths at t_0, t_1, \dots, t_{N-1} based on the parameters in Table 2.1 and Case B of Table 4.4. Panel (a) in Fig. 8.1 shows the actual versus the approximated spread for each path at each t_i . In the plot, each color represents a timestep, and the solid line is $Y = X$ line. The points are generally close to the $Y = X$ line although there are some deviations when the actual spread is positive. It is noteworthy that as t_i approaches $\bar{T} = 1$, the time gap of the spread shrinks and the approximation becomes usually more accurate.

Panel (b) in Fig. 8.1 shows the histogram of the error (approximate spread minus the actual spread) over all paths and timesteps. It is seen that over the majority of the cases the error is small; it is found that the mean and median of the error is \$1.71 and \$1.33 respectively. Large errors typically occur at large positive spread (down-ward sloping forward curve). It is noteworthy that the approximate spread has a correct sign 87% of the times.

Moreover, to illustrate the degree of linearity, Fig. 8.2 shows two sample realizations of $e^{-r(T-t)}[F(t, T) - c_p]$ versus T at different t . The linear regression line at each t is also computed and overlaid on the corresponding curve. The graphs as well as the regression R^2 and standard error indicate that the linearity assumption can be reasonable. This is particularly true for the later time stages since when t approaches to 1 year (problem time horizon) the length of the curve decreases.



(a) Approximate (left-hand side of Eq. 8.34) versus actual spread (right-hand side of Eq. 8.34) at each path and each timestep. The solid black line represents $Y = X$ line.

(b) Histogram of errors, defined by the approximate spread minus the actual spread, over all paths and all timesteps. The mean (median) of the histogram is \$1.71 (\$1.33).

Fig. 8.1. Testing the accuracy of estimating the spread via a linear approximation, i.e. with $a = t_i$ and $b = \bar{T} = 1$, using 10,000 sample paths over t_0, t_1, \dots, t_{N-1} timesteps, shown in different colors. All simulation parameters are based on Table 2.1 and Case B of Table 4.4.





Fig. 8.2. Simulated $e^{-r(T-t)}[F(t,T) - c_p]$ curves where $0 \leq t \leq 1$ and $t \leq T \leq 1$ year using Table 2.1 parameters, where temporal discretization is set to $N = 16$ time stages to match that of Case B in Table 4.4. R^2 and standard error of the regression is listed for each curve fitted with a line.

8.5 Proof of Proposition 4.II:

Proposition 4.II: If it is assumed that the difference between the ‘adjusted’ forward prices can be written as Eq. 8.35 (i.e. Assumption 4.I holds), then the value function and optimal actions structure for $\forall i \in \mathcal{J} \setminus \{N\}$ are expressed by Eq. 8.36 and Eq. 8.37. It can be seen easily that Eq. 8.37 is a particular case of Eq. 8.36 (for $i = N - 1$) which is written explicitly. Here $\mathbb{E}_i[\cdot]$ denotes $\mathbb{E}[\cdot | W_i]$.

$$e^{-r(t_1-t_i)}[F(t_i, t_1) - c_P] - e^{-r(t_2-t_i)}[F(t_i, t_2) - c_P] = m_i(t_1 - t_2) \quad \text{Eq. 8.35}$$

- **If $0 \leq i \leq N - 2$**

$$V_i(x_i, W_i) = R_i(m_i - c_H)(t_i - T_i) + \Delta t R_i u_i, \quad u_i = \max\{0, A_i, B_i\}$$

$$(a_i^R, a_i^T) = \begin{cases} (R_i, 0) & \text{If } A_i < 0 \text{ and } B_i < 0 \\ (0, t_N) & \text{If } B_i > A_i \text{ and } B_i > 0 \\ (0, t_{i+1}) & \text{If } A_i > B_i \text{ and } A_i > 0 \end{cases} \quad \text{Eq. 8.36}$$

$$A_i := m_i - c_H + \delta \mathbb{E}_i[u_{i+1}]$$

$$B_i := (N - i)(m_i - c_H) + \delta \mathbb{E}_i[u_{i+1} - (N - i - 1)(m_{i+1} - c_H)]$$

- **If $i = N - 1$**

$$V_{N-1}(x_{N-1}, W_{N-1}) = R_{N-1}(m_{N-1} - c_H)(t_{N-1} - T_{N-1}) + \Delta t R_{N-1} u_{N-1}$$

$$u_{N-1} = \max\{0, A_i\} \quad \text{Eq. 8.37}$$

$$(a_{N-1}^R, a_{N-1}^T) = \begin{cases} (R_{N-1}, 0) & \text{If } A_{N-1} < 0 \\ (0, t_N) & \text{If } A_{N-1} > 0 \end{cases}$$

$$A_{N-1} := m_{N-1} - c_H$$

Proof:

We prove Eq. 8.37 individually, and Eq. 8.36 by backward induction. To show Eq. 8.36, we solve the Bellman equation at $i = N - 1$ as follows. At this time stage, the continuation value is zero, and the maximization is performed over the only two available actions $(R_{N-1}, 0)$ and $(0, t_N)$.

$$\begin{aligned}
V_{N-1}(x_{N-1}, W_{N-1}) &= \max_{a \in \mathcal{A}_{N-1}(x_{N-1})} r_{N-1}(x_{N-1}, W_{N-1}, a) = \\
&\max_{a \in \mathcal{A}_{N-1}(x_{N-1})} (m_{N-1} - c_H)[a_{N-1}^R(t_{N-1} - a_{N-1}^T) + R_{N-1}(a_{N-1}^T - T_{N-1})] \\
&= \max\{(m_{N-1} - c_H)R_{N-1}(t_{N-1} - T_{N-1}), (m_{N-1} - c_H)R_{N-1}(t_N - T_{N-1})\} \\
&= R_{N-1} \max\{(m_{N-1} - c_H)(t_{N-1} - T_{N-1}), (m_{N-1} - c_H)(t_N - T_{N-1})\} \\
&= R_{N-1}(m_{N-1} - c_H)(t_{N-1} - T_{N-1}) + R_{N-1} \max\{0, \Delta t(m_{N-1} - c_H)\} \\
&= R_{N-1}(m_{N-1} - c_H)(t_{N-1} - T_{N-1}) + \Delta t R_{N-1} \max\{0, m_{N-1} - c_H\}
\end{aligned} \tag{Eq. 8.38}$$

Let us define

$$u_{N-1} := \max\{0, m_{N-1} - c_H\}. \tag{Eq. 8.39}$$

So, $V_{N-1}(x_{N-1}, W_{N-1})$ can be written as

$$V_{N-1}(x_{N-1}, W_{N-1}) = R_{N-1}(m_{N-1} - c_H)(t_{N-1} - T_{N-1}) + \Delta t R_{N-1} u_{N-1}. \tag{Eq. 8.40}$$

From the order of the arguments in the maximum operator of Eq. 8.39, we deduce the conditions under which each action is optimal, as expressed in Eq. 8.41. This concludes the proof of Eq. 8.36.

$$(a_{N-1}^R, a_{N-1}^T) = \begin{cases} (R_{N-1}, 0) & \text{If } m_{N-1} - c_H < 0 \\ (0, t_N) & \text{If } m_{N-1} - c_H > 0 \end{cases} \tag{Eq. 8.41}$$

To prove Eq. 8.37 using backward induction, we first show that the structure holds for stage $i = N - 2$. Let us consider the Bellman equation, Eq. 17, at time stage $i = N - 2$. For notational brevity, let $E_i[X]$ denote $E[X | W_i]$.

$$\begin{aligned}
V_{N-2}(x_{N-2}, W_{N-2}) &= \\
&\max_{a \in \mathcal{A}_{N-2}(x_{N-2})} r_{N-2}(x_{N-2}, W_{N-2}, a) + \delta \mathbb{E}[V_{N-1}(f_{N-2}(x_{N-2}, a), W_{N-1}) | W_{N-2}] \\
&= \max_{a \in \mathcal{A}_{N-2}(x_{N-2})} (m_{N-2} - c_H)[a_{N-2}^R(t_{N-2} - a_{N-2}^T) + R_{N-2}(a_{N-2}^T - T_{N-2})] \\
&\quad + \delta \mathbb{E}_{N-2}[V_{N-1}(x_{N-1}, W_{N-1})]
\end{aligned} \tag{Eq. 8.42}$$

The candidates for the optimal decision are studied by replacing them into the objective function, and compare the resulting payoffs. Here, we use Proposition I by eliminating the partial sale choices. The candidates can be classified as the following:

0. Choice #0: $(a_{N-2}^R, a_{N-2}^T) = (R_{N-2}, 0)$, i.e. sell everything on the spot and exit, in which case there is not any continuation value.

$$P_0^{N-2} := (m_{N-2} - c_H)R_{N-2}(t_{N-2} - T_{N-2}) \quad \text{Eq. 8.43}$$

1. Choice #1: $(a_{N-2}^R, a_{N-2}^T) = (0, t_N)$, i.e. sell nothing on the spot and short the longest maturity t_N .

$$\begin{aligned} P_1^{N-2} &:= (m_{N-2} - c_H)R_{N-2}(t_N - T_{N-2}) \\ &\quad + \delta \mathbb{E}_{N-2}[R_{N-2}(m_{N-1} - c_H)(t_{N-1} - t_N) + \Delta t R_{N-2} u_{N-1}] \\ &= (m_{N-2} - c_H)R_{N-2}(t_N - T_{N-2}) + \delta \mathbb{E}_{N-2}[-\Delta t R_{N-2}(m_{N-1} - c_H) + \Delta t R_{N-2} u_{N-1}] \\ &= (m_{N-2} - c_H)R_{N-2}(t_N - T_{N-2}) + \delta \Delta t R_{N-2} \mathbb{E}_{N-2}[-(m_{N-1} - c_H) + u_{N-1}] \end{aligned} \quad \text{Eq. 8.44}$$

2. Choice #2: $(a_{N-2}^R, a_{N-2}^T) = (0, t_{N-1})$, i.e. sell nothing on the spot and short t_{N-1} , rather than the longest maturity t_N .

$$\begin{aligned} P_2^{N-2} &:= (m_{N-2} - c_H)R_{N-2}(t_{N-1} - T_{N-2}) \\ &\quad + \delta \mathbb{E}_{N-2}[R_{N-2}(m_{N-1} - c_H)(t_{N-1} - t_{N-1}) + \Delta t R_{N-2} u_{N-1}] \\ &= (m_{N-2} - c_H)R_{N-2}(t_{N-1} - T_{N-2}) + \delta \Delta t R_{N-2} \mathbb{E}_{N-2}[u_{N-1}] \end{aligned} \quad \text{Eq. 8.45}$$

The following summarizes the values generated by taking all potentially optimal actions.

$$P_0^{N-2} = (m_{N-2} - c_H)R_{N-2}(t_{N-2} - T_{N-2}) \quad \text{Eq. 8.46}$$

$$P_1^{N-2} = (m_{N-2} - c_H)R_{N-2}(t_N - T_{N-2}) + \delta \Delta t R_{N-2} \mathbb{E}_{N-2}[-(m_{N-1} - c_H) + u_{N-1}] \quad \text{Eq. 8.47}$$

$$P_2^{N-2} = (m_{N-2} - c_H)R_{N-2}(t_{N-1} - T_{N-2}) + \delta \Delta t R_{N-2} \mathbb{E}_{N-2}[u_{N-1}] \quad \text{Eq. 8.48}$$

Therefore, by substituting Eq. 8.46 to Eq. 8.48 into Eq. 8.42, we get the following.

$$\begin{aligned}
V_{N-2}(x_{N-2}, W_{N-2}) &= \\
&\max_{a \in \mathcal{A}_{N-2}(x_{N-2})} (m_{N-2} - c_H)[a_{N-2}^R(t_{N-2} - a_{N-2}^T) + R_{N-2}(a_{N-2}^T - T_{N-2})] \\
&\quad + \mathbb{E}[V_{N-1}(x_{N-1}, W_{N-1}) | W_{N-2}] \\
&= \max\{P_0^{N-2}, P_1^{N-2}, P_2^{N-2}\} \\
&= R_{N-2} \max\{(m_{N-2} - c_H)(t_{N-2} - T_{N-2}), (m_{N-2} - c_H)(t_N - T_{N-2}) \\
&\quad + \delta \Delta t \mathbb{E}_{N-2}[-(m_{N-1} - c_H) + u_{N-1}], (m_{N-2} - c_H)(t_{N-1} - T_{N-2}) \\
&\quad + \delta \Delta t \mathbb{E}_{N-2}[u_{N-1}]\} \tag{Eq. 8.49} \\
&= R_{N-2}(m_{N-2} - c_H)(t_{N-2} - T_{N-2}) + R_{N-2} \times \\
&\quad \max\{0, 2\Delta t(m_{N-2} - c_H) + \delta \Delta t \mathbb{E}_{N-2}[-(m_{N-1} - c_H) + u_{N-1}], \Delta t(m_{N-2} - c_H) \\
&\quad + \delta \Delta t \mathbb{E}_{N-2}[u_{N-1}]\} \\
&= R_{N-2}(m_{N-2} - c_H)(t_{N-2} - T_{N-2}) + \Delta t R_{N-2} \times \\
&\quad \max\left\{ \begin{array}{l} 0, \\ 2(m_{N-2} - c_H) + \delta \mathbb{E}_{N-2}[-(m_{N-1} - c_H) + u_{N-1}], \\ (m_{N-2} - c_H) + \delta \mathbb{E}_{N-2}[u_{N-1}] \end{array} \right\}
\end{aligned}$$

Let us define

$$u_{N-2} := \max\left\{ \begin{array}{l} 0, \\ 2(m_{N-2} - c_H) + \delta \mathbb{E}_{N-2}[-(m_{N-1} - c_H) + u_{N-1}], \\ (m_{N-2} - c_H) + \delta \mathbb{E}_{N-2}[u_{N-1}] \end{array} \right\} \tag{Eq. 8.50}$$

Thus, by substituting Eq. 8.50 into Eq. 8.49, we have

$$V_{N-2}(x_{N-2}, W_{N-2}) = R_{N-2}(m_{N-2} - c_H)(t_{N-2} - T_{N-2}) + \Delta t R_{N-2} u_{N-2} \tag{Eq. 8.51}$$

From the order of the arguments in the maximum operator, we deduce the conditions under which each action is optimal as specified in Eq. 8.52.

$$(a_{N-2}^R, a_{N-2}^T) = \begin{cases} (R_{N-2}, 0) & \text{If } A_{N-2} < 0 \quad \text{and } B_{N-2} < 0 \\ (0, t_N) & \text{If } B_{N-2} > A_{N-2} \quad \text{and } B_{N-2} > 0 \\ (0, t_{N-1}) & \text{If } A_{N-2} > B_{N-2} \quad \text{and } A_{N-2} > 0 \end{cases} \tag{Eq. 8.52}$$

Here,

$$\begin{aligned}
A_{N-2} &:= (m_{N-2} - c_H) + \delta \mathbb{E}_{N-2}[u_{N-1}] \\
B_{N-2} &:= 2(m_{N-2} - c_H) + \delta \mathbb{E}_{N-2}[-(m_{N-1} - c_H) + u_{N-1}]
\end{aligned}
\tag{Eq. 8.53}$$

The above verifies the value function structure, Eq. 8.37, for time stage $i = N - 2$. Now, we present the remaining part of the backward induction argument.

Suppose

$$\begin{aligned}
V_i(x_i, W_i) &= R_i(m_i - c_H)(t_i - T_i) + \Delta t R_i u_i \\
u_i &= \max \left\{ \begin{array}{l} 0, \\ (N-i)(m_i - c_H) + \delta \mathbb{E}_i[u_{i+1} - (N-i-1)(m_{i+1} - c_H)], \\ m_i - c_H + \delta \mathbb{E}_i[u_{i+1}] \end{array} \right\} \\
(a_i^R, a_i^T) &= \begin{cases} (R_i, 0) & \text{If } A_i < 0 \text{ and } B_i < 0 \\ (0, t_N) & \text{If } B_i > A_i \text{ and } B_i > 0 \\ (0, t_{i+1}) & \text{If } A_i > B_i \text{ and } A_i > 0 \end{cases} \\
A_i &:= m_i - c_H + \delta \mathbb{E}_i[u_{i+1}] \\
B_i &:= (N-i)(m_i - c_H) + \delta \mathbb{E}_i[u_{i+1} - (N-i-1)(m_{i+1} - c_H)]
\end{aligned}
\tag{Eq. 8.54}$$

We need to show that

$$\begin{aligned}
V_{i-1}(x_{i-1}, W_{i-1}) &= R_{i-1}(m_{i-1} - c_H)(t_{i-1} - T_{i-1}) + \Delta t R_{i-1} u_{i-1} \\
u_{i-1} &= \max \left\{ \begin{array}{l} 0, \\ (N-i+1)(m_{i-1} - c_H) + \delta \mathbb{E}_{i-1}[u_i - (N-i)(m_i - c_H)], \\ m_{i-1} - c_H + \delta \mathbb{E}_{i-1}[u_i] \end{array} \right\} \\
(a_{i-1}^R, a_{i-1}^T) &= \begin{cases} (R_{i-1}, 0) & \text{If } A_{i-1} < 0 \text{ and } B_{i-1} < 0 \\ (0, t_N) & \text{If } B_{i-1} > A_{i-1} \text{ and } B_{i-1} > 0 \\ (0, t_i) & \text{If } A_{i-1} > B_{i-1} \text{ and } A_{i-1} > 0 \end{cases} \\
A_{i-1} &:= m_{i-1} - c_H + \delta \mathbb{E}_{i-1}[u_i] \\
B_{i-1} &:= (N-i+1)(m_{i-1} - c_H) + \delta \mathbb{E}_{i-1}[u_i - (N-i)(m_i - c_H)]
\end{aligned}
\tag{Eq. 8.55}$$

We start by writing the Bellman equation at stage $i - 1$, and by substituting $V_i(x_i, W_i)$ with the appropriate expression from the backward induction assumption, i.e. Eq. 8.54, as detailed in Eq. 8.56.

$$\begin{aligned}
& V_{i-1}(x_{i-1}, W_{i-1}) \\
&= \max_{a \in \mathcal{A}_{i-1}(x_{i-1})} (m_{i-1} - c_H) [a_{i-1}^R(t_{i-1} - a_{i-1}^T) + R_{i-1}(a_{i-1}^T - T_{i-1})] \\
&\quad + \delta \mathbb{E}[V_i(x_i, W_i) | W_{i-1}] \\
&= \max_{a \in \mathcal{A}_{i-1}(x_{i-1})} (m_{i-1} - c_H) [a_{i-1}^R(t_{i-1} - a_{i-1}^T) + R_{i-1}(a_{i-1}^T - T_{i-1})] \\
&\quad + \delta \mathbb{E}_{i-1}[R_i(m_i - c_H)(t_i - T_i) + \Delta t R_i u_i]
\end{aligned} \tag{Eq. 8.56}$$

There are $N - i + 2$ candidates for the optimal decision, which are replaced in the objective function to compute the corresponding payoffs. Here, we use Proposition I by eliminating the choices involving a partial sale. The candidates can be classified as the following:

Choice #0: $(a_{i-1}^R, a_{i-1}^T) = (R_{i-1}, 0)$:

$$P_0^{i-1} := R_{i-1}(m_{i-1} - c_H)(t_{i-1} - T_{i-1}) \tag{Eq. 8.57}$$

Choice #1: $(a_{i-1}^R, a_{i-1}^T) = (0, t_N)$:

$$\begin{aligned}
& P_1^{i-1} := (m_{i-1} - c_H) R_{i-1} (t_N - T_{i-1}) \\
&\quad + \delta \mathbb{E}_{i-1}[R_{i-1}(m_i - c_H)(t_i - t_N) + \Delta t R_{i-1} u_i] \\
&= (m_{i-1} - c_H) R_{i-1} (t_N - T_{i-1}) + \delta \mathbb{E}_{i-1}[-(N - i) \Delta t R_{i-1} (m_i - c_H) + \Delta t R_{i-1} u_i] \\
&= R_{i-1} (m_{i-1} - c_H) (t_N - T_{i-1}) + \delta \Delta t R_{i-1} \mathbb{E}_{i-1}[-(N - i)(m_i - c_H) + u_i]
\end{aligned} \tag{Eq. 8.58}$$

Choice #2: $(a_{i-1}^R, a_{i-1}^T) = (0, t_{N-1})$:

$$\begin{aligned}
& P_2^{i-1} := (m_{i-1} - c_H) R_{i-1} (t_{N-1} - T_{i-1}) \\
&\quad + \delta \mathbb{E}_{i-1}[R_{i-1}(m_i - c_H)(t_i - t_{N-1}) + \Delta t R_{i-1} u_i] \\
&= (m_{i-1} - c_H) R_{i-1} (t_{N-1} - T_{i-1}) \\
&\quad + \delta \mathbb{E}_{i-1}[-(N - i - 1) \Delta t R_{i-1} (m_i - c_H) + \Delta t R_{i-1} u_i] \\
&= R_{i-1} (m_{i-1} - c_H) (t_{N-1} - T_{i-1}) + \delta \Delta t R_{i-1} \mathbb{E}_{i-1}[-(N - i - 1)(m_i - c_H) + u_i]
\end{aligned} \tag{Eq. 8.59}$$

...

Choice # $(N - i + 1)$: $(a_{i-1}^R, a_{i-1}^T) = (0, t_i)$ (recall that we are at t_{i-1}):

$$\begin{aligned} P_{N-i+1}^{i-1} &:= (m_{i-1} - c_H)R_{i-1}(t_i - T_{i-1}) \\ &\quad + \delta \mathbb{E}_{i-1}[R_{i-1}(m_i - c_H)(t_i - t_i) + \Delta t R_{i-1} u_i] \\ &= R_{i-1}(m_{i-1} - c_H)(t_i - T_{i-1}) + \delta \Delta t R_{i-1} \mathbb{E}_{i-1}[u_i] \end{aligned} \tag{Eq. 8.60}$$

Substituting the payoffs from different actions in the Bellman equation, Eq. 8.56, yields the following expression.

$$\begin{aligned} V_{i-1}(x_{i-1}, W_{i-1}) &= \max\{P_0^{i-1}, P_1^{i-1}, P_2^{i-1}, \dots, P_{N-i+1}^{i-1}\} \\ &= \max \left\{ \begin{array}{l} R_{i-1}(m_{i-1} - c_H)(t_{i-1} - T_{i-1}), \\ R_{i-1}(m_{i-1} - c_H)(t_N - T_{i-1}) + \delta \Delta t R_{i-1} \mathbb{E}_{i-1}[-(N-i)(m_i - c_H) + u_i], \\ R_{i-1}(m_{i-1} - c_H)(t_{N-1} - T_{i-1}) + \delta \Delta t R_{i-1} \mathbb{E}_{i-1}[-(N-i-1)(m_i - c_H) + u_i], \\ \dots \\ R_{i-1}(m_{i-1} - c_H)(t_i - T_{i-1}) + \delta \Delta t R_{i-1} \mathbb{E}_{i-1}[u_i] \end{array} \right\} \\ &= R_{i-1} \max \left\{ \begin{array}{l} (m_{i-1} - c_H)(t_{i-1} - T_{i-1}), \\ (m_{i-1} - c_H)(t_N - T_{i-1}) + \delta \Delta t \mathbb{E}_{i-1}[-(N-i)(m_i - c_H) + u_i], \\ (m_{i-1} - c_H)(t_{N-1} - T_{i-1}) + \delta \Delta t \mathbb{E}_{i-1}[-(N-i-1)(m_i - c_H) + u_i], \\ \dots \\ (m_{i-1} - c_H)(t_i - T_{i-1}) + \delta \Delta t \mathbb{E}_{i-1}[u_i] \end{array} \right\} \tag{Eq. 8.61} \\ &= R_{i-1}(m_{i-1} - c_H)(t_{i-1} - T_{i-1}) \\ &\quad + R_{i-1} \max \left\{ \begin{array}{l} 0, \\ (m_{i-1} - c_H)(t_N - t_{i-1}) + \delta \Delta t \mathbb{E}_{i-1}[-(N-i)(m_i - c_H) + u_i], \\ (m_{i-1} - c_H)(t_{N-1} - t_{i-1}) + \delta \Delta t \mathbb{E}_{i-1}[-(N-i-1)(m_i - c_H) + u_i], \\ \dots \\ (m_{i-1} - c_H)(t_i - t_{i-1}) + \delta \Delta t \mathbb{E}_{i-1}[u_i] \end{array} \right\} \\ &= R_{i-1}(m_{i-1} - c_H)(t_{i-1} - T_{i-1}) \\ &\quad + R_{i-1} \max \left\{ \begin{array}{l} 0, \\ (m_{i-1} - c_H)(N-i+1)\Delta t + \delta \Delta t \mathbb{E}_{i-1}[-(N-i)(m_i - c_H) + u_i], \\ (m_{i-1} - c_H)(N-i)\Delta t + \delta \Delta t \mathbb{E}_{i-1}[-(N-i-1)(m_i - c_H) + u_i], \\ \dots \\ (m_{i-1} - c_H)\Delta t + \delta \Delta t \mathbb{E}_{i-1}[u_i] \end{array} \right\} \end{aligned}$$

$$\begin{aligned}
&= R_{i-1}(m_{i-1} - c_H)(t_{i-1} - T_{i-1}) \\
&+ \Delta t R_{i-1} \max \left\{ \begin{array}{l} 0, \\ (N-i+1)(m_{i-1} - c_H) + \delta \mathbb{E}_{i-1}[-(N-i)(m_i - c_H) + u_i], \\ (N-i)(m_{i-1} - c_H) + \delta \mathbb{E}_{i-1}[-(N-i-1)(m_i - c_H) + u_i], \\ \dots \\ (m_{i-1} - c_H) + \delta \mathbb{E}_{i-1}[u_i] \end{array} \right\} \\
&= R_{i-1}(m_{i-1} - c_H)(t_{i-1} - T_{i-1}) + \delta \Delta t R_{i-1} \mathbb{E}_{i-1}[u_i] \\
&\quad + \Delta t R_{i-1} \max \left\{ \begin{array}{l} -\delta \mathbb{E}_{i-1}[u_i], \\ (N-i+1)(m_{i-1} - c_H) - (N-i)\delta \mathbb{E}_{i-1}[m_i - c_H], \\ (N-i)(m_{i-1} - c_H) - (N-i-1)\delta \mathbb{E}_{i-1}[m_i - c_H], \\ \dots \\ (m_{i-1} - c_H) \end{array} \right\} \\
&= R_{i-1}(m_{i-1} - c_H)(t_{i-1} - T_{i-1}) + \delta \Delta t R_{i-1} \mathbb{E}_{i-1}[u_i] + \Delta t R_{i-1}(m_{i-1} - c_H) \\
&\quad + \Delta t R_{i-1} \max \left\{ \begin{array}{l} -\delta \mathbb{E}_{i-1}[u_i] - (m_{i-1} - c_H), \\ (N-i)(m_{i-1} - c_H) - (N-i)\delta \mathbb{E}_{i-1}[m_i - c_H], \\ (N-i-1)(m_{i-1} - c_H) - (N-i-1)\delta \mathbb{E}_{i-1}[m_i - c_H], \\ \dots \\ 0 \end{array} \right\} \\
&= R_{i-1}(m_{i-1} - c_H)(t_{i-1} - T_{i-1}) + \delta \Delta t R_{i-1} \mathbb{E}_{i-1}[u_i] + \Delta t R_{i-1}(m_{i-1} - c_H) \\
&\quad + \Delta t R_{i-1} \max \left\{ \begin{array}{l} -\delta \mathbb{E}_{i-1}[u_i] - (m_{i-1} - c_H), \\ (N-i)(m_{i-1} - c_H - \delta \mathbb{E}_{i-1}[m_i - c_H]), \\ (N-i-1)(m_{i-1} - c_H - \delta \mathbb{E}_{i-1}[m_i - c_H]), \\ \dots \\ 0 \end{array} \right\}
\end{aligned}$$

The last $N-i+1$ arguments of the above maximum operator are all multiples of $(m_{i-1} - c_H - \delta \mathbb{E}_{i-1}[m_i - c_H])$. Due to this linear dependence, we can simplify the maximization as follows.

$$\begin{aligned}
V_{i-1}(x_{i-1}, W_{i-1}) &= R_{i-1}(m_{i-1} - c_H)(t_{i-1} - T_{i-1}) + \delta \Delta t R_{i-1} \mathbb{E}_{i-1}[u_i] \\
&+ \Delta t R_{i-1}(m_{i-1} - c_H) \\
&+ \Delta t R_{i-1} \max \left\{ \begin{array}{l} -\delta \mathbb{E}_{i-1}[u_i] - (m_{i-1} - c_H), \\ (N-i)(m_{i-1} - c_H - \delta \mathbb{E}_{i-1}[m_i - c_H]), \\ 0 \end{array} \right\} \\
&= R_{i-1}(m_{i-1} - c_H)(t_{i-1} - T_{i-1}) + \Delta t R_{i-1} u_{i-1}
\end{aligned} \tag{Eq. 8.62}$$

Here,

$$u_{i-1} = \max \left\{ \begin{array}{l} 0, \\ (N-i+1)(m_{i-1} - c_H) + \delta \mathbb{E}_{i-1}[u_i - (N-i)(m_i - c_H)], \\ m_{i-1} - c_H + \delta \mathbb{E}_{i-1}[u_i] \end{array} \right\} \quad \text{Eq. 8.63}$$

The three arguments in the maximum correspond to the (a_{i-1}^R, a_{i-1}^T) decision to be $(R_{i-1}, 0)$, $(0, t_N)$, or $(0, t_i)$, and thus

$$(a_{i-1}^R, a_{i-1}^T) = \begin{cases} (R_{i-1}, 0) & \text{If } A_{i-1} < 0 \quad \text{and } B_{i-1} < 0 \\ (0, t_N) & \text{If } B_{i-1} > A_{i-1} \quad \text{and } B_{i-1} > 0 \\ (0, t_i) & \text{If } A_{i-1} > B_{i-1} \quad \text{and } A_{i-1} > 0 \end{cases} \quad \text{Eq. 8.64}$$

$$A_{i-1} := m_{i-1} - c_H + \delta \mathbb{E}_{i-1}[u_i]$$

$$B_{i-1} := (N-i+1)(m_{i-1} - c_H) + \delta \mathbb{E}_{i-1}[u_i - (N-i)(m_i - c_H)]$$

The above completes the backward induction argument by showing that the value function structure at time stage $i-1$ is indeed in the form of Eq. 8.55. ■

8.6 Proof of Proposition 5.I

Proposition 5.I: Let $I_i^{-(0)}$ denote any $I_i \in \mathcal{X}_i^I$ such that $I_i \neq 0$, and $x_i = (I_i^{-(0)}, *, *)$ denote any $x_i = (I_i, R_i, T_i) \in \mathcal{X}_i$ in which $I_i \neq 0$. $\forall i \in \{0, 1, 2, \dots, N-1\}$, $V_i((I_i^{-(0)}, *, *), W_i)$ is not a function of the storage cost stochastic factors, $(\chi'_i, \xi'_i, \alpha_i, \alpha_i^*)$, while $V_i((0, 0, 0), W_i)$ is a function of them.

Intuitively, $V_i((I_i^{-(0)}, *, *), W_i)$ reflects the value function in a state where the tanker has been already rented. Thus, the determinants of the storage cost, $(\chi'_i, \xi'_i, \alpha_i, \alpha_i^*)$, do not impact the value, however, $(\chi'_i, \xi'_i, \alpha_i, \alpha_i^*)$ do impact $V_i((0, 0, 0), W_i)$ since the tanker has not been rented yet.

Proof:

Suppose $i = N-1$, then

$$I_{N-1} \in \mathcal{X}_{N-1}^I = \{0, \emptyset, t_{N-1}, t_N\} \quad \text{Eq. 8.65}$$

$$x_{N-1} \in \mathcal{X}_{N-1} = \{(0,0,0), (\emptyset, 0,0), (t_{N-1}, 0,0), (t_N, 0,0)\} \quad \text{Eq. 8.66}$$

$$\cup \{t_{N-1}\} \times (0, \bar{R}] \times \{t_{N-1}\} \cup \{t_N\} \times (0, \bar{R}] \times \{t_{N-1}, t_N\}$$

Bellman's equation $\forall x_{N-1} = (I_{N-1}, R_{N-1}, T_{N-1}) \in \mathcal{X}_{N-1}$ will be as per Eq. 8.67 and Eq. 8.68. Given the reward function $r_i(a_i, x_i, W_i)$, the value function depends on the stochastic factors $(\chi'_i, \xi'_i, \alpha_i, \alpha_i^*)$ through the term ' $\bar{R}T' C_i^H(t_i, t_{i+N'}) I(I_i = 0 \wedge \alpha_i^I = 1)$ ', which is only non-zero if $I_i = 0$, and subsequently $x_i = (0,0,0)$. For V_{N-1} in Eq. 8.68, since $I_{N-1} = I_{N-1}^{-(0)}$, the value function does not depend on $(\chi'_{N-1}, \xi'_{N-1}, \alpha_{N-1}, \alpha_{N-1}^*)$. However, V_{N-1} in Eq. 8.67 does depend on $(\chi'_{N-1}, \xi'_{N-1}, \alpha_{N-1}, \alpha_{N-1}^*)$.

1. For $x_{N-1} = (0,0,0)$:

$$\begin{aligned} V_{N-1}((0,0,0), W_{N-1}) &= \max_{a \in \mathcal{A}_{N-1}(x_{N-1})} r_{N-1}(a, (0,0,0), W_{N-1}) + \delta(-R_N c_P^-) \quad \text{Eq. 8.67} \\ &= V_{N-1}((0,0,0), (\chi_{N-1}, \xi_{N-1}, \chi'_{N-1}, \xi'_{N-1}, \alpha_{N-1}, \alpha_{N-1}^*)) \end{aligned}$$

2. For any $x_{N-1} = (I_{N-1}^{-(0)}, *, *)$:

$$\begin{aligned} V_{N-1}\left(\left(I_{N-1}^{-(0)}, *, *\right), W_{N-1}\right) &= \max_{a \in \mathcal{A}_{N-1}(x_{N-1})} r_{N-1}\left(a, \left(I_i^{-(0)}, *, *\right), W_{N-1}\right) + \delta(-R_N c_P^-) \quad \text{Eq. 8.68} \\ &= V_{N-1}\left(\left(I_i^{-(0)}, *, *\right), (\chi_{N-1}, \xi_{N-1})\right) \end{aligned}$$

Before moving to the previous timestep, $i = N - 2$, let us examine the functional dependence of the expectations of V_{N-1} to the stochastic factors since they will appear in the Bellman's equations. The expectation in Eq. 8.69 is obviously a function of all the six stochastic factors included in W_{N-2} . However, the expectation in Eq. 8.70 is only a function of (χ_{N-2}, ξ_{N-2}) since V_{N-1} is only a function of (χ_{N-1}, ξ_{N-1}) , and (χ_i, ξ_i) and $(\chi'_i, \xi'_i, \alpha_i, \alpha_i^*)$ are independent $\forall i \in \mathcal{J}$. Here f_1 and f_2 denote generic functions.

$$E[V_{N-1}((0,0,0), W_{N-1}) | W_{N-2}] = f_1(\chi_{N-2}, \xi_{N-2}, \chi'_{N-2}, \xi'_{N-2}, \alpha_{N-2}, \alpha_{N-2}^*) \quad \text{Eq. 8.69}$$

$$E[V_{N-1} \left(\left(I_{N-1}^{-(0)}, *, * \right), W_{N-1} \right) | W_{N-2}] =$$

$$E[V_{N-1} \left(\left(I_{N-1}^{-(0)}, *, * \right), (\chi_{N-1}, \xi_{N-1}) \right) | (\chi_{N-2}, \xi_{N-2}, \chi'_{N-2}, \xi'_{N-2}, \alpha_{N-2}, \alpha_{N-2}^*)] =$$

$$f_2(\chi_{N-2}, \xi_{N-2}) \quad \text{Eq. 8.70}$$

At $i = N - 2$, \mathcal{X}_{N-2}^I and \mathcal{X}_{N-2} are as follows.

$$I_{N-2} \in \mathcal{X}_{N-2}^I = \{0, \emptyset, t_{N-2}, t_{N-1}, t_N\} \quad \text{Eq. 8.71}$$

$$x_{N-2} \in \mathcal{X}_{N-2} = \{(0,0,0), (\emptyset, 0,0), (t_{N-2}, 0,0), (t_{N-1}, 0,0), (t_N, 0,0)\}$$

$$\cup \{t_{N-2}\} \times (0, \bar{R}] \times \{t_{N-2}\}$$

$$\cup \{t_{N-1}\} \times (0, \bar{R}] \times \{t_{N-2}, t_{N-1}\}$$

$$\cup \{t_N\} \times (0, \bar{R}] \times \{t_{N-2}, t_{N-1}, t_N\} \quad \text{Eq. 8.72}$$

At $i = N - 2$, Bellman's equation $\forall x_{N-2} \in \mathcal{X}_{N-2}$ will be

1. For $x_{N-2} = (0,0,0)$:

$$V_{N-2}((0,0,0), W_{N-2}) =$$

$$\max_{a \in \mathcal{A}_{N-2}(x_{N-2})} r_{N-2}(a, (0,0,0), W_{N-2}) + \delta E[V_{N-1}(x_{N-1}, W_{N-1}) | W_{N-2}] \quad \text{Eq. 8.73}$$

$$= V_{N-2}((0,0,0), (\chi_{N-2}, \xi_{N-2}, \chi'_{N-2}, \xi'_{N-2}, \alpha_{N-2}, \alpha_{N-2}^*))$$

In this case, $r_{N-2}(a, (0,0,0), W_{N-2})$ is a function of $(\chi'_{N-2}, \xi'_{N-2}, \alpha_{N-2}, \alpha_{N-2}^*)$, and $E[V_{N-1}(x_{N-1}, W_{N-1}) | W_{N-2}]$ may or may not be a function of $(\chi'_{N-2}, \xi'_{N-2}, \alpha_{N-2}, \alpha_{N-2}^*)$ depending on the action taken and the resulting x_{N-1} . Therefore, V_{N-2} is a function of all six stochastic factors.

2. For $x_{N-2} = (I_{N-2}^{-(0)}, *, *)$:

$$\begin{aligned}
& V_{N-2} \left(\left(I_{N-2}^{-(0)}, *, * \right), W_{N-2} \right) = \\
& \max_{a \in \mathcal{A}_{N-2}(x_{N-2})} r_{N-2} \left(a, \left(I_{N-2}^{-(0)}, *, * \right), W_{N-2} \right) + \delta E[V_{N-1}(x_{N-1}, W_{N-1}) | W_{N-2}] \\
& = V_{N-2} \left(\left(I_{N-2}^{-(0)}, *, * \right), (\chi_{N-2}, \xi_{N-2}) \right)
\end{aligned} \tag{Eq. 8.74}$$

In Eq. 8.74, $r_{N-2}(a, (1, *, *), W_{N-2})$ is ‘not’ a function of $(\chi'_{N-2}, \xi'_{N-2}, \alpha_{N-2}, \alpha_{N-2}^*)$. Also, $E[V_{N-1}(x_{N-1}, W_{N-1}) | W_{N-2}]$ is ‘not’ a function of $(\chi'_{N-2}, \xi'_{N-2}, \alpha_{N-2}, \alpha_{N-2}^*)$, no matter what action is undertaken, since the resulting x_{N-1} is always in the form of $(I_{N-1}^{-(0)}, *, *)$. Therefore, V_{N-2} is only a function of (χ_{N-2}, ξ_{N-2}) .

Thus, the above arguments prove the proposition for $i = N - 1$ (in Eq. 8.67, Eq. 8.68) and $i = N - 2$ (in Eq. 8.73, and Eq. 8.74). Now suppose the proposition holds for stage i ; that is $V_i \left(\left(I_i^{-(0)}, *, * \right), W_i \right)$ is not a function of the storage cost stochastic factors, $(\chi'_i, \xi'_i, \alpha_i, \alpha_i^*)$, while $V_i((0,0,0), W_i)$ is a function of them. To complete the proof by backward induction, one needs to show the same holds at stage $i - 1$. First, let us express the functionality of the expectations with respect to the stochastic factors at stage $i - 1$, as written in Eq. 8.75 and Eq. 8.76. The fact that (χ_i, ξ_i) and $(\chi'_i, \xi'_i, \alpha_i, \alpha_i^*)$ are independent $\forall i \in \mathcal{J}$ is used in Eq. 8.76. Here, f_3 and f_4 denote generic functions.

$$E[V_i((0,0,0), W_i) | W_{i-1}] = f_3(\chi_{i-1}, \xi_{i-1}, \chi'_{i-1}, \xi'_{i-1}, \alpha_{i-1}, \alpha_{i-1}^*) \tag{Eq. 8.75}$$

$$\begin{aligned}
& E \left[V_i \left(\left(I_i^{-(0)}, *, * \right), W_i \right) | W_{i-1} \right] = \\
& E \left[V_i \left(\left(I_i^{-(0)}, *, * \right), (\chi_i, \xi_i) \right) | (\chi_{i-1}, \xi_{i-1}, \chi'_{i-1}, \xi'_{i-1}, \alpha_{i-1}, \alpha_{i-1}^*) \right] = \\
& = f_4(\chi_{i-1}, \xi_{i-1})
\end{aligned} \tag{Eq. 8.76}$$

Bellman’s equation $\forall x_{i-1} \in \mathcal{X}_{i-1}$ will be

1. For $x_{i-1} = (0,0,0)$:

$$V_{i-1}((0,0,0), W_{i-1}) =$$

$$\max_{a \in \mathcal{A}_{i-1}(x_{i-1})} r_{i-1}(a, (0,0,0), W_{i-1}) + \delta E[V_i(x_i, W_i) | W_{i-1}]$$

Eq. 8.77

$$= V_{i-1}((0,0,0), (\chi_{i-1}, \xi_{i-1}, \chi'_{i-1}, \xi'_{i-1}, \alpha_{i-1}, \alpha^*_{i-1}))$$

In this case, $r_{i-1}(a, (0,0,0), W_{i-1})$ is a function of $(\chi'_{i-1}, \xi'_{i-1}, \alpha_{i-1}, \alpha^*_{i-1})$, and $E[V_i(x_i, W_i) | W_{i-1}]$ may or may not be a function of $(\chi'_{i-1}, \xi'_{i-1}, \alpha_{i-1}, \alpha^*_{i-1})$ depending on the action taken and the resulting x_i . So, $V_{i-1}((0,0,0), W_{i-1})$ is a function of $(\chi_{i-1}, \xi_{i-1}, \chi'_{i-1}, \xi'_{i-1}, \alpha_{i-1}, \alpha^*_{i-1})$.

2. For $x_{i-1} = (I_{i-1}^{-(0)}, *, *)$:

$$V_{i-1}\left(\left(I_{i-1}^{-(0)}, *, *\right), W_{i-1}\right) =$$

$$\max_{a \in \mathcal{A}_{i-1}(x_{i-1})} r_{i-1}\left(a, \left(I_{i-1}^{-(0)}, *, *\right), W_{i-1}\right) + \delta E[V_i(x_i, W_i) | W_{i-1}]$$

Eq. 8.78

$$= V_{i-1}\left(\left(I_{i-1}^{-(0)}, *, *\right), (\chi_{i-1}, \xi_{i-1})\right)$$

In this case, $r_{i-1}\left(a, \left(I_{i-1}^{-(0)}, *, *\right), W_{i-1}\right)$ is ‘not’ a function of $(\chi'_{i-1}, \xi'_{i-1}, \alpha_{i-1}, \alpha^*_{i-1})$. Also, $E[V_i(x_i, W_i) | W_{i-1}]$ is ‘not’ a function of $(\chi'_{i-1}, \xi'_{i-1}, \alpha_{i-1}, \alpha^*_{i-1})$, no matter what action is chosen, since the resulting x_i is always in the form of $\left(I_i^{-(0)}, *, *\right)$. So, $V_{i-1}\left(\left(I_{i-1}^{-(0)}, *, *\right), W_{i-1}\right)$ is just a function of (χ_{i-1}, ξ_{i-1}) . ■

9 Appendix B

This appendix includes the R codes used in Chapter 5, which is built on the stochastic storage cost framework. It also provides the solution assuming that the storage cost is constant. In the following, the code is divided into different parts using subsections with headings for further clarity. Each subsection includes a function or specific pieces of the code that are focused on a narrow task.

9.1 Function Simulating Oil and Storage Cost Prices

```
# This function generates a sample of (chi, ksi, chi_prime, ksi_prime, alpha,
# alpha_star) of length N.out
sample.fcn=function(M, M2, T1, T2, N.out, N.prim, rz_ksi.ksip, rz_chi.chip,
  chi_0, ksi_0, chi2_0, ksi2_0, alp_0, alps_0, r,
  #Oil parameters (Hahn et al)
  mu.ksi,k,lambda.chi,sig.chi,sig.ksi,rho,mu.ksi.star,
  lambda.ksi,
  #storage cost parameters
  mu.ksi2,K2,Phi,sig.ksi2,sig.chi2,sig.alp,lambda.ksi2,
  lambda.chi2,lambda.alp,lambda.alps,
  #correlation matrix of the storage stochastic factors
  r2.chiksi,r2.chialp,r2.chialps,r2.ksialp,r2.ksialps ) {

# Function that simulates two antithetic sample of the evolution of chi-ksi
# from t=0 to t=T, given the initial values
W.process=function(chi_t0,ksi_t0,Z){
  #number of periods: N.out
  #number of partition points: N.out+1
  #problem time horizon: T1
  #problem delta t: dt=T1/N.out

  #there will be N periods and N+1 partition points for SDE discretization
  # so, N is the number of simulations from t to t+1
  N=N.dis / N.out
  #internal dt for SDE discretization
  dt1=dt / N

  #Nt is the total number of steps required
  Nt=N.dis

  #First pair-----
  Z1=Z
  #Generate correlated Brownian increments
  dW.ksi=Z1[1:Nt]
  dW.chi=rho*Z1[1:Nt]+sqrt(1-rho^2)*Z1[(Nt+1):(2*Nt)]

  ksi=rep(ksi_t0,Nt+1)
  chi=rep(chi_t0,Nt+1)

  #SDE UNDER PHYSICAL MEASURE
```



```

#ksi SDE coefficients
a1=1
b1=mu.ksi*dt1
c1=sig.ksi*sqrt(dt1)

#chi SDE coefficients
a2=exp(-k*dt1)
b2=0
c2=sig.chi*sqrt((1-exp(-2*k*dt1))/(2*k))

#loop to calculate ksi and chi
for ( i in 2:(Nt+1)){
  ksi[i]= ksi[i-1]+b1+c1*dW.ksi[i-1] #a1=1 not written to speed up
  chi[i]=a2*chi[i-1]+c2*dW.chi[i-1]
}

res=cbind(chi[seq(1,Nt+1,by=N)],ksi[seq(1,Nt+1,by=N)])

#Second pair, which is the antithetic pair-----
Z1=-Z
dW.ksi=Z1[1:Nt]
dW.chi=rho*Z1[1:Nt]+sqrt(1-rho^2)*Z1[(Nt+1):(2*Nt)]

#SDE's coefficients UNDER PHYSICAL MEASURE are already computed

#loop to calculate ksi and chi
for ( i in 2:(Nt+1)){
  ksi[i]= ksi[i-1]+b1+c1*dW.ksi[i-1] #a1=1 not written
  chi[i]=a2*chi[i-1]+c2*dW.chi[i-1]
}

return( cbind(res, cbind(chi[seq(1,Nt+1,by=N)],ksi[seq(1,Nt+1,by=N)])) )
)

}
#-----

#-----
#Function that simulates two antithetic evolutions of chi, ksi, alpha, and
# alpha_star (storage cost stochastic factors) from t=0 to t=T, given the
# initial values
W2.process=function(chi2_t0,ksi2_t0,alp_t0,alps_t0,Z){
  #number of periods: N.out
  #problem time horizon: T1
  #problem delta t: dt=T1/N.out

  #there will be N periods and N+1 partition points for SDE discretization
  # so N is the number of simulations from t to t+1
  N=N.dis / N.out
  #internal dt for SDE discretization
  dt1 = dt / N

  #total number of steps required

```

```

Nt=N.dis

#First pair-----
dW.chi2=Z[1,]
dW.ksi2=Z[2,]
dW.alp=Z[3,]
dW.alps=Z[4,]

#Initialize the arrays
chi2=rep(chi2_t0,Nt+1)
ksi2=rep(ksi2_t0,Nt+1)
alp=rep(alp_t0,Nt+1)
alps=rep(alps_t0,Nt+1)

#SDE UNDER PHYSICAL MEASURE
#chi2 SDE coefficients
a2=exp(-K2*dt1)
#b2=0
c2=sig.chi2*sqrt((1-exp(-2*K2*dt1))/(2*K2))

#ksi2 SDE coefficients
#a1=1
b1=mu.ksi2*dt1
c1=sig.ksi2*sqrt(dt1)

#alpha and alpha_star SDE coeff
cof=2*pi*Phi*dt1
a3=1+cof^2
c3=sig.alp*sqrt(dt1)

#loop to calculate ksi2 and chi2 and alpha, and alpha_star
for ( i in 2:(Nt+1)){
  chi2[i]=a2*chi2[i-1]+c2*dW.chi2[i-1]
  ksi2[i]= ksi2[i-1]+b1+c1*dW.ksi2[i-1] #a1=1 not written
  alp[i]=( alp[i-1] + cof*alps[i-1] + c3*(cof*dW.alps[i-1]+dW.alp[i-1])
)/a3
  alps[i]=( alps[i-1] - cof*alp[i-1] + c3*(dW.alps[i-1]-cof*dW.alp[i-1])
)/a3
}

res=cbind(chi2[seq(1,Nt+1,by=N)],ksi2[seq(1,Nt+1,by=N)],alp[seq(1,Nt+1,by
=N)],alps[seq(1,Nt+1,by=N)])

#Second pair-----
Z=-1*Z

dW.chi2=Z[1,]
dW.ksi2=Z[2,]
dW.alp=Z[3,]
dW.alps=Z[4,]

#SDEs coefficients UNDER PHYSICAL MEASURE are already computed

#loop to calculate ksi2 and chi2 and alpha, and alpha_star

```

```

    for ( i in 2:(Nt+1)){
      chi2[i]=a2*chi2[i-1]+c2*dW.chi2[i-1]
      ksi2[i]= ksi2[i-1]+b1+c1*dW.ksi2[i-1] #a1=1 not written
      alp[i]=( alp[i-1] + cof*alps[i-1] + c3*(cof*dW.alps[i-1]+dW.alp[i-1])
)/a3
      alps[i]=( alps[i-1] - cof*alp[i-1] + c3*(dW.alps[i-1]-cof*dW.alp[i-1])
)/a3
    }

    return( cbind(res, cbind(chi2[seq(1,Nt+1,by=N)], ksi2[seq(1,Nt+1,by=N)], a
lp[seq(1,Nt+1,by=N)], alps[seq(1,Nt+1,by=N)]) ) )

}
#-----

#-----
#Function that computes storage cost futures price given ttm=T-t (time to m
aturity), chi_t, and ksi_t
fut2.price=function(ttm,chi2_t,ksi2_t,alp_t,alps_t){

#-----
#function A(ttm)
A2=function(ttm,temp0,temp){
  return(

    (mu.ksi2-lambda.ksi2+0.5*sig.ksi2^2+0.5*sig.alp^2)*ttm

    -(1-exp(-K2*ttm))*(lambda.chi2-sig.ksi2*sig.chi2*r2.chiksi)/K2

    +0.25*(1-exp(-2*K2*ttm))*sig.chi2^2/K2

    -(lambda.alps+lambda.alp*sin(temp)-lambda.alps*cos(temp))/temp0

    +sig.ksi2*sig.alp*(r2.ksialp*(1-cos(temp)) + r2.ksialps*sin(temp) )/t
emp0

    +sig.chi2*sig.alp*r2.chialp*(K2-exp(-K2*ttm)*(K2*cos(temp)+temp0*sin(
temp)))/(K2^2+temp0^2)

    +sig.chi2*sig.alp*r2.chialps*(temp0+exp(-K2*ttm)*(K2*sin(temp)-temp0*
cos(temp)))/(K2^2+temp0^2)
  )
}
#-----
temp0=2*pi*Phi
temp=temp0*ttm
fut=ksi2_t +exp(-K2*ttm)*chi2_t + cos(temp)*alp_t + sin(temp)*alps_t + A2
(ttm=ttm,temp0=temp0,temp=temp)
return(exp(fut))

}
#-----

```

```

#-----
#Function that computes Time-Charter rate from time t to T (ttm=T-t; actual
# time) given chi2_t, ksi2_t, etc
TC.fcn=function(ttm,chi2_t,ksi2_t,alp_t,alps_t){

  f.integral=function(ttm,chi2_t,ksi2_t,alp_t,alps_t){
    return(
      integrate(function(x) fut2.price(x,chi2_t=chi2_t,ksi2_t=ksi2_t,alp_t=
alp_t,alps_t=alps_t),lower = 0, upper = ttm)[[1]]/ttm
    )
  }

  # convert whole vessel per day to barrel pey year:
  # TD5 Suezmax capacity is 130,000 Metric Tons or 1,000,000 barrels
  #vectorized computation
  return( (365/1000000)*Vectorize(f.integral)(ttm,chi2_t,ksi2_t,alp_t,alps_
t) )

}
#-----

# Parameters
#####

#end of time horizon
#T1
#rental period length
#T2
#how many timesteps the rental contract length is: N.prim=T2/dt

#Total number of Euler discretization of the period T1
N.dis=(2^5)*(3^2)*5

#make sure there is perfect divisibility
if( (N.dis/N.out)!=round(N.dis/N.out) ) stop()
if( (N.dis/N.prim)!=round(N.dis/N.prim) ) stop()

dt=T1/N.out

D=exp(-r*dt) #one-period discount factor

sig2=array(c(1,r2.chiksi,r2.chialp,r2.chialps,
            r2.chiksi,1,r2.ksialp,r2.ksialps,
            r2.chialp,r2.ksialp,1,0,
            r2.chialps,r2.ksialps,0,1)
          ,dim=c(4,4))

#Construct a block-diagonal matrix by diagonally binding the two covariance
# matrices. Row or columns in the 6x6 matrix are respectively; chi ksi chi'
# ksi' alpha alpha*
library(Matrix) #for block diagonal binding
sig3=bdiag( matrix(c(1,rho,rho,1),ncol=2), sig2 )

```

```

#set a correlation for ksi and ksi'
sig3[2,4]=rz_ksi.ksip
sig3[4,2]=rz_ksi.ksip
#set a correlation for chi and chi'
sig3[1,3]=rz_chi.chip
sig3[3,1]=rz_chi.chip

#Make sure the covariance matrix is positive semi-definite
if( any(eigen(sig3)$values<0) ) stop()

#####

#prepare for parallel computing
#-----
#find the number of cores
n.cores=detectCores()
cl <- makeCluster(n.cores)
# RNG seed will be set to 123
clusterSetRNGStream(cl,123)
clusterExport(cl,c("n.cores", "chi2_0", "ksi2_0", "alp_0", "alps_0", "mu.ks
i2", "K2", "Phi", "sig.ksi2", "sig.chi2", "sig.alp", "lambda.ksi2", "lambda.ch
i2", "lambda.alp", "lambda.alps", "r2.chiksi", "r2.chialp", "r2.chialps", "r2.ks
ialp", "r2.ksialps", "sig3", "T1", "T2", "N.out", "N.prim", "N.dis", "dt", "r", "D", "
M", "M2", "mu.ksi", "k", "lambda.chi", "sig.chi", "sig.ksi", "rho", "mu.ksi.star", "la
mbda.ksi", "fut2.price", "TC.fcn", "W.process", "W2.process", "rnd.fcn"),envir = e
nvironment() )
# register the cluster
registerDoParallel(cl)
#-----

#generate M (sample) + M2 (out-of-sample) antithetic sample paths and compute
# TC rates in parallel
#-----
sample.state3= foreach (m = seq(1,M+M2,by=2), .combine = rbind ) %dopar% {

  #generate all the required random numbers (oil & storage) for path m
  Z=rnd.fcn(N.dis)

  #oil price: first 2 rows of Z
  W2=W.process(chi_t0=chi_0,ksi_t0=ksi_0,Z[1:2,])
  #output:
  sample.core = rbind( W2[,1:2], W2[,3:4])

  #storage cost: second 4 rows of Z
  W2=W2.process(chi2_0,ksi2_0,alp_0,alps_0,Z[3:6,])
  #output:
  sample.core2 = rbind( W2[,1:4], W2[,5:8])

  cbind(sample.core,sample.core2)
}

```

```

#-----

#reshape the sample.state matrix:
# change the dimension to c(N.out+1,2 or 4,M+M2)
#-----
sample.state=array(0,dim = c(N.out+1,2,M+M2))
sample.state.2=array(0,dim = c(N.out+1,4,M+M2))

for (t in 1:(N.out+1)){
  t.ind=seq(t, (M+M2)*(N.out+1), by=(N.out+1) )
  #oil price
  sample.state[t,1,]=sample.state3[t.ind,1] #chi
  sample.state[t,2,]=sample.state3[t.ind,2] #ksi
  #storage cost
  sample.state.2[t,1,]=sample.state3[t.ind,3] #chi2
  sample.state.2[t,2,]=sample.state3[t.ind,4] #ksi2
  sample.state.2[t,3,]=sample.state3[t.ind,5] #alpha
  sample.state.2[t,4,]=sample.state3[t.ind,6] #alpha_star
}
rm(sample.state3)

#Separate out-of-sample paths from the in-sample paths
#out-of-sample
sample.path2 = sample.state[,,(M+1):(M+M2)]
sample.path2.storage = sample.state.2[,,(M+1):(M+M2)]
#in-sample
sample.state = sample.state[,,(1):(M)]
sample.state.2= sample.state.2[,,(1):(M)]

#Pre-compute all the needed Time-Charter Rates (TC, storage costs, or HC)
sample.TC.both = foreach (t = 1:(N.out+1-N.prim), .combine = rbind ) %dopar
% {
  res1=TC.fcn( T2, chi2_t=sample.state.2[t,1,],ksi2_t=sample.state.2[t,2,],
  alp_t=sample.state.2[t,3,],alps_t=sample.state.2[t,4,])
  res2=TC.fcn( T2, chi2_t=sample.path2.storage[t,1,],ksi2_t=sample.path2.st
orage[t,2,],alp_t=sample.path2.storage[t,3,],alps_t=sample.path2.storage[t,4,
])
  matrix(c(res1,res2),ncol=(M+M2),nrow=1)
}
sample.TC=sample.TC.both[,1:M]
sample.TC.2=sample.TC.both[, (M+1):(M+M2)]
#-----

# shut down the cluster
stopCluster(cl)

return(list("a"=sample.state,"b"=sample.state.2,"c"=sample.TC,
          "d"=sample.path2,"e"=sample.path2.storage,"f"=sample.TC.2 ))
}

```

9.2 Function Computing Oil Futures Prices

```
#function that computes futures prices given ttm=T-t (time-to-maturity),
# chi_t, and ksi_t
fut.price=function(ttm,chi_t,ksi_t){

  #-----
  #function A(ttm)
  A=function(ttm){
    return(mu.ksi.star*ttm-(1-exp(-k*ttm))*lambda.chi/k +0.5*((1-exp(-2*k*ttm)
  ))*sig.chi^2/(2*k) + sig.ksi^2*ttm + 2*(1-exp(-k*ttm))*rho*sig.chi*sig.ksi/k
  )
  }
  #-----

  fut=exp(-k*ttm)*chi_t + ksi_t + A(ttm=ttm)
  return(exp(fut))

}
```

9.3 Function Computing the Endogenous State Space

```
#set of Possible States at time t
statespace.fcn=function(t){

  #SPECIAL STATES: no tanker has been rented (0,1,1)
  pos.states=c(0,1,1)

  # the tanker rent contract cannot be expired yet
  if (t<=(N.prim+1)) {
    if (t!=1) {
      # when tanker is empty: R_t=1: (I_t,1,1)
      foreach ( I_t=(N.prim+1):min(t-1+N.prim,N.out+1) )%:%foreach ( R_t=1 )%
:%foreach ( T_t=1 ) %do% {
        pos.states=rbind(pos.states,c(I_t,R_t,T_t))
      }

      # when tanker is Not empty: R_t>1: (I_t,R_t,T_t)
      foreach ( I_t=(N.prim+1):min(t-1+N.prim,N.out+1) )%:%foreach ( R_t=2:(R
.out+1) )%:%foreach ( T_t=t:I_t ) %do% {
        pos.states=rbind(pos.states,c(I_t,R_t,T_t))
      }
    }
  }

  # there is the possibility that tanker was rented but contract expired (123
4,1,1)
  if (t>(N.prim+1)) {
    pos.states=rbind(pos.states,c(1234,1,1))

    # when tanker is empty: R_t=1: (I_t,1,1)
    foreach ( I_t=t:min(t-1+N.prim,N.out+1) )%:%foreach ( R_t=1 )%:%foreach (
T_t=1 ) %do% {
      pos.states=rbind(pos.states,c(I_t,R_t,T_t))
    }
  }
}
```

```

    }

    # when tanker is Not empty: R_t>1: (I_t,R_t,T_t)
    foreach ( I_t=t:min(t-1+N.prim,N.out+1) )%:%foreach ( R_t=2:(R.out+1) )%:
%foreach ( T_t=t:I_t ) %do% {
        pos.states=rbind(pos.states,c(I_t,R_t,T_t))
    }

}

return(pos.states)

}

#set of Possible States at time t
statespace.sub.fcn=function(t){
    #not include the expired state c(1234,1,1) since it is an absorbing state w
ith zero value
    #not include the not-rented-yet state c(0,1,1) if t>N-N'+1 since it is an a
bsorbing state with zero value
    #not include the state c(t,1,1), just expiring with 0 inventory, since it i
s an absorbing state with zero value

    #SPECIAL STATES: tanker not-rented-yet (0,1,1)
if(t<=(N.out+1-N.prim)){
        pos.states=c(0,1,1)
    }else{
        pos.states=NULL
    }

    # the tanker rent contract cannot be expired yet
if (t<(N.prim+1)) {
        if (t!=1) {
            # when tanker is empty: R_t=1: (I_t,1,1)
            foreach ( I_t=(N.prim+1):min(t-1+N.prim,N.out+1) )%:%foreach ( R_t=1 )%
:%foreach ( T_t=1 ) %do% {
                pos.states=rbind(pos.states,c(I_t,R_t,T_t))
            }

            # when tanker is Not empty: R_t>1: (I_t,R_t,T_t)
            foreach ( I_t=(N.prim+1):min(t-1+N.prim,N.out+1) )%:%foreach ( R_t=2:(R
.out+1) )%:%foreach ( T_t=t:I_t ) %do% {
                pos.states=rbind(pos.states,c(I_t,R_t,T_t))
            }
        }

    # the tanker rent contract cannot be expired yet
else if (t==(N.prim+1)) {
        if (t!=1) {
            # when tanker is empty: R_t=1: (I_t,1,1)
            foreach ( I_t=(N.prim+2):min(t-1+N.prim,N.out+1) )%:%foreach ( R_t=1
)%:%foreach ( T_t=1 ) %do% {

```



```

    pos.states=rbind(pos.states,c(I_t,R_t,T_t))
  }

  # when tanker is Not empty: R_t>1: (I_t,R_t,T_t)
  foreach ( I_t=(N.prim+1):min(t-1+N.prim,N.out+1) )%:%foreach ( R_t=2:
(R.out+1) )%:%foreach ( T_t=t:I_t ) %do% {
    pos.states=rbind(pos.states,c(I_t,R_t,T_t))
  }
}

# there is the possibility that tanker was rented but contract expired (123
4,1,1)
}else if (t>(N.prim+1)) {
  # when tanker is empty: R_t=1: (I_t,1,1)
  foreach ( I_t=(t+1):min(t-1+N.prim,N.out+1) )%:%foreach ( R_t=1 )%:%forea
ch ( T_t=1 ) %do% {
    pos.states=rbind(pos.states,c(I_t,R_t,T_t))
  }

  # when tanker is Not empty: R_t>1: (I_t,R_t,T_t)
  foreach ( I_t=t:min(t-1+N.prim,N.out+1) )%:%foreach ( R_t=2:(R.out+1) )%:
%foreach ( T_t=t:I_t ) %do% {
    pos.states=rbind(pos.states,c(I_t,R_t,T_t))
  }
}

}

return(pos.states)
}

```

9.4 Function Computing the Action Space

```

#action set (in integers):
#does not include xT_t=t b/c selling on the spot is done separately via xR_t
action.set=function(t,I_t,R_t){

  #(xI_t, xR_t, xT_t)

  if( I_t==t | I_t==1234 ){
    return(matrix(c(0,-(R_t-1),1),nrow = 1)) #must sell any existing invento
ry, we set xT_t=0 (=1 in integer)

  }else if( I_t==0 & t<=(N.out+1-N.prim) ) { #no tanker but can rent it now (
implies R_t=1)
    res=NULL
    for (xR_t in (-(R_t-1)+1):(R.out+1-R_t) ){ # in continuous form -R_t+dR<=x
R_t<=R1-R_t
      for(xT_t in (t+1):(t+N.prim) ) { # maximum allowed maturity is
limited by the rental contract expiry (t+N.prim)
        res=rbind(c(1,xR_t,xT_t),res)
      }
    }
  }
}

```

```

    }
  }
  return(rbind(res,
              c(1,-(R_t-1),1),#decide to rent the tanker but not buy any o
il, then xT_t is meaningless, and we set xT_t=0 (=1 in integer)
              c(0,-(R_t-1),1) #decide not to rent the tanker at all
              ))

}else if(I_t==0 & t>(N.out+1-N.prim) ) { #no tanker but can not rent anymo
re
  return(matrix(c(0,0,1),nrow = 1)) #we set xT_t=0 (=1 in integer)

}else{ #already a tanker rent contract in place with expiry at I_t. All xI_
t=0
  res=NULL
  for (xR_t in (-(R_t-1)+1):(R.out+1-R_t) ){ # in continous form -R_t+dR<=x
R_t<=R1-R_t
    for(xT_t in (t+1):I_t ){ # maximum allowed maturity is limite
d by the exiting rental contract expiry I_t
      res=rbind(c(0,xR_t,xT_t),res)
    }
  }
  #if decide to sell everything, then xT_t is meaningless, and we set xT_t=
0 (=1 in integer)
  return(rbind(res,c(0,-(R_t-1),1)))
}
}

```

9.5 Function Representing the Transition Function

```

#state transition function (in integers)
#Note about xR_t: R_i=(i-1)d R_j=(j-1)d so R_j-R_i=(j-i)d
S.transition=function(t,I_t,R_t,xI_t,xR_t,xT_t){

  if( I_t==1234 | I_t==t ){ #it means a delivery must conclude the trading
    #I_2=1234
    #R_2 = 1
    #T_2 = 1
    return(cbind(1234,1,1))

  }else if( I_t==0 & xI_t==1 ){ #just rented the tanker
    #I_2=t+N.prim
    #R_2 = R_t + xR_t
    #T_2 = xT_t
    return(cbind( t+N.prim, R_t+xR_t, xT_t ))

  }else if( I_t==0 & xI_t==0 ){ #no tanker and decided not to rent it
    #I_2=0
    #R_2 = R_t + xR_t: we can igoner xR_t b/c it is zero by the action set
    #T_2 = xT_t
  }
}

```

```

    return(cbind(0, R_t, xT_t))

}else{ #there is already an active rent contract in place
  #I_2=I_t
  #R_2 = R_t + xR_t
  #T_2 = xT_t
  return(cbind(I_t, R_t + xR_t, xT_t))
}
}

```

9.6 Function Formulating the Reward Function

```

#payoff function based on the spread (in integers) VECTORIZED with respect to
chi and ksi
reward.fcn=function(t,I_t,R_t,T_t,xI_t,xR_t,xT_t,chi_t,ksi_t,TC_t){
  #excluded ,chi_t,ksi_t,chi2_t,ksi2_t,alp_t,alps_t from the inputs of the re
ward function

  # t current time
  # R_t level of oil in storage
  # T_t initial maturity
  # x_t new maturity chosen
  # trader considers to transfer maturity frim T_t to x_t

  if( I_t==t | I_t==1234 ){
    return(
      -(R_t-1)*dR*PC2
    )

  }else{
    # compute potential storage cost payment for period [t,T1]
    if((I_t==0)&(xI_t==1)){
      storage.cost=-R1*N.prim*dt*TC_t #TC==HC, TC.fcn((N.out+1-t)*dt,chi2_t,k
si2_t,alp_t,alps_t)
    }else{
      storage.cost=0
    }

    return(
      # sell on spot price
      -xR_t*dR * ( fut.price(ttm = 0 ,chi_t,ksi_t)+PC1*(xR_t>0)-PC2*(xR_t<0)
    )

    # short forward
    +(R_t+xR_t-1)*dR*exp( -r*(xT_t-t)*dt ) * fut.price((xT_t-t)*dt, chi_t,
ksi_t)

    # long existing maturity to offset the current contract held
    -(R_t-1)*dR*exp( -r*(T_t-t)*dt ) * fut.price((T_t-t)*dt, chi_t, ksi_t)
    # potential storage cost payment
    +storage.cost
  )
}
}

```

```
}
```

9.7 Auxiliary Function – Array Index Mapping

```
#Map real I_t to an integer index:
It.indexer.fcn=function(I_t){
  #Other than I_t=0 or 1234 ;
  #Minimum real I_t is T2, which is N.prim+1 in discrete (assume entering at
t=0)
  #Maximum real I_t is T1, which is N.out+1 in discrete (assume entering at t
=T1-T2)
  #We consider index 1 for I_t=0, index 2 for I_t=1234,
  # and 3,4,...,3+(N.out-N.prim) for N.prim+1,...,N.out+1.
  #This reduces the size of the Theta array.

  if (I_t==0){
    return(1)

  }else if (I_t==1234)
    return(2)

  else{
    return(I_t-N.prim+2) #Note that I_t will be always greater than or equal
to N.prim+1
  }

}
```

9.8 Auxiliary Function – Polynomial Feature Mapping

```
#polynomial feature mapping function: (NOT including the intercept)
fmap.fcn=function(a,p){
  # a is the array of features; each feature in a colmun
  #p= 2 or 3 #maximum power of the polynomial
  n=ncol(a)

  if(p==2){
    #all power 2 combos
    ind2=cbind(getall(iterpc(n,2,replace = TRUE)),0)
    #all power 1 combos
    ind1=cbind(getall(iterpc(n,1,replace = TRUE)),0,0)
    #all combos
    ind=rbind(ind2,ind1) #ind3,

    n.factor=nrow(ind)

    res=matrix(0,nrow=nrow(a),ncol=n.factor)

    for(i in 1:n.factor){# loop over all combos each generating one (new) fac
tor
      if ( length(which(ind[i,]!=0)) > 1 ){
```

```

        res[,i]=rowProds(a[,ind[i,]]) #,method = "direct"
    }else{
        res[,i]=a[,ind[i,]]
    }
}

}

if(p==3){
    #all power 3 combos
    ind3=getall(iterpc(n,3,replace = TRUE))
    #all power 2 combos
    ind2=cbind(getall(iterpc(n,2,replace = TRUE)),0)
    #all power 1 combos
    ind1=cbind(getall(iterpc(n,1,replace = TRUE)),0,0)
    #all combos
    ind=rbind(ind3,ind2,ind1)

    n.factor=nrow(ind)

    res=matrix(0,nrow=nrow(a),ncol=n.factor)

    for(i in 1:n.factor){# loop over all combos each generating one (new) fac
tor
        if ( length(which(ind[i,]!=0)) > 1 ){
            res[,i]=rowProds(a[,ind[i,]]) #,method = "direct"

        }else{
            res[,i]=a[,ind[i,]]
        }
    }

}

return( res )

}
#polynomial feature mapping function: (NOT including the intercept)
fmap.singlerow.fcn=function(a,p){
    # a is the array of features; each feature in a colmun
    #p= 2 or 3 #maximum power of the polynomial
    n=ncol(a)

    if(p==2){
        #all power 2 combos
        ind2=cbind(getall(iterpc(n,2,replace = TRUE)),0)
        #all power 1 combos
        ind1=cbind(getall(iterpc(n,1,replace = TRUE)),0,0)
        #all combos
        ind=rbind(ind2,ind1) #ind3,

        n.factor=nrow(ind)

        res=matrix(0,nrow=nrow(a),ncol=n.factor)

```

```

    for(i in 1:n.factor){# loop over all combos each generating one (new) fac
tor
    if ( length(which(ind[i,]!=0)) > 1 ){
        res[,i]=prod(a[,ind[i,]])
    }else{
        res[,i]=a[,ind[i,]]
    }
}

}

if(p==3){
#all power 3 combos
ind3=getall(iterpc(n,3,replace = TRUE))
#all power 2 combos
ind2=cbind(getall(iterpc(n,2,replace = TRUE)),0)
#all power 1 combos
ind1=cbind(getall(iterpc(n,1,replace = TRUE)),0,0)
#all combos
ind=rbind(ind3,ind2,ind1)

n.factor=nrow(ind)

res=matrix(0,nrow=nrow(a),ncol=n.factor)

for(i in 1:n.factor){# loop over all combos each generating one (new) fac
tor
    if ( length(which(ind[i,]!=0)) > 1 ){
        res[,i]=prod(a[,ind[i,]])

    }else{
        res[,i]=a[,ind[i,]]
    }
}

}

return( res )

}

```

9.9 Function Computing the Optimal Action Given Stochastic Storage Costs

```

#function that gets a state and finds the optimal action pair that generates
# the maximum value ( OPTimal POLicy finder)
oppo.fcn=function(t,I_t,R_t,T_t,theta,TC_t,mydata.t){

chi_t = mydata.t[,main.col[1]]
ksi_t = mydata.t[,main.col[2]]

# feasible action set
X=action.set(t,I_t,R_t)

```

```

#initialize temporary value function for every possible action
q=array(0, dim = c(nrow(X), length(chi_t) ) ) #M=length(chi_t)

#loop over all actions x in the feasible set of actions
for (j1 in 1:nrow(X)) {

  #choose the j1-th feasible action
  xI_t=X[j1,1]
  xR_t=X[j1,2]
  xT_t=X[j1,3]

  #compute the reward of the chosen action
  C_t = reward.fcn(t,I_t,R_t,T_t,xI_t,xR_t,xT_t,chi_t,ksi_t,TC_t)

  #compute the next state given the current state and the chosen action
  S2 = S.transition(t, I_t,R_t,xI_t,xR_t,xT_t)

  # CV: expected value of V1[t+1,S2] based on S2

  #continuation value: V2=V[t+1, S2[1], S2[2], ] recover from the list cont
.val
  # S2[1]+1: +1 is the adjustment to transform to array index
  if(t==N.out){
    v.hat=-(S2[2]-1)*dR*PC2
  }else{
    #storage cost functionality is considered (I_t2==0) and # CV is not a f
cn of storage cost if I_t2!=0
    v.hat=theta[t, It.indexer.fcn(S2[1]) ,S2[2],S2[3],]%*%t(mydata.t)
  }

  #total value from the j1-th pair of actions: current payoff + discounted
value function
  q[j1,] = C_t + D*v.hat
}

#for any sample path, choose the optimal action pair
return( apply(q,2,max) )
}

#function that receives a sample sate and finds the Optimal Action pair that
generates
# the maximum value ( OPTimal POLicy finder)
oppo.c.fcn=function(t,I_t,R_t,T_t,theta,TC_t,mydata.t){

  chi_t = mydata.t[,main.col.c[1]]
  ksi_t = mydata.t[,main.col.c[2]]

  # feasible action set
  X=action.set(t,I_t,R_t)

  #initialize temporary value function for every possible action
  q=array(0, dim = c(nrow(X), length(chi_t) ) ) #M=length(chi_t)

```

```

#loop over all actions x in the feasible set of actions
for (j1 in 1:nrow(X)) {

  #choose the j1-th feasible action
  xI_t=X[j1,1]
  xR_t=X[j1,2]
  xT_t=X[j1,3]

  #compute the reward of the chosen action
  C_t = reward.fcn(t,I_t,R_t,T_t,xI_t,xR_t,xT_t,chi_t,ksi_t,TC_t)

  #compute the next state given the current state and the chosen action
  S2 = S.transition(t, I_t,R_t,xI_t,xR_t,xT_t)

  # CV: expected value of V1[t+1,S2] based on S2

  #continuation value: V2=V[t+1, S2[1], S2[2], ] recover from the list cont
.val
  # S2[1]+1: +1 is the adjustment to transform to array index
  if(t==N.out){
    v.hat=-(S2[2]-1)*dR*PC2
  }else{
    #storage cost functionality is considered (I_t2==0) and # CV is not a f
cn of storage cost if I_t2!=0
    v.hat=theta[t, It.indexer.fcn(S2[1]) ,S2[2],S2[3],]%*%t(mydata.t)
  }

  #total value from the j1-th pair of actions: current payoff + discounted
value function
  q[j1,] = C_t + D*v.hat
}

#for any sample path, choose the optimal action pair
return( apply(q,2,max) )
}

```

9.10 Specify the Parameters and Simulate the Prices

```

#Global parameters: part one (fixed ones)
#####
####
#Oil parameters (Hahn et al)
mu.ksi=0.0818
k=1.0880 #MAKE SURE YOU DON NOT USE SMALL "k" LETTER
lambda.chi=0.3733
sig.chi=0.3116
sig.ksi=0.2053
rho=0.0823
mu.ksi.star=-0.0252
lambda.ksi=0.107 #based on lambda.ksi=mu.ksi-mu.ksi.star=0.0818-(-0.0252)

```



```

#storage cost parameters
mu.ksi2=1.1049
K2=4.8854
Phi=0.8426 # we expect the seasonality period to be one year
sig.ksi2=0.4134
sig.chi2=1.9123
sig.alp=0.2285 #zero means deterministic seasonality

lambda.ksi2=1.2014
lambda.chi2=-1.4337
lambda.alp=-0.0291
lambda.alps=0.5587

#correlation matrix of the storage stochastic factors
r2.chiksi=-0.3849
r2.chialp=-0.2453
r2.chialps=-0.7586
r2.ksialp=-0.2769
r2.ksialps=0.5984

#IC
chi_0=-0.63932
ksi_0=4.6366

chi2_0=3.39 #log(46000)-(alp_0+alps_0+ksi2_0)
ksi2_0=8.4 #log(36000)-chi2_0-alp_0-alps_0
alp_0=0.3 #if we set to 0, it'll stay so if the variance is 0 too
alps_0=0.4 #if we set to 0, it'll stay so if the variance is 0 too

#end of time horizon
T1=2

#rental period length (T.prime)
T2=1

#total number of trading periods (16 or 12 per year)
N.out=12*T1
dt=T1/N.out

#how many time steps the rental length is:
N.prim=T2/dt

R1=1 #tanker capacity
R.out=1 #Number of tanker discretization intervals: 1<=R_t<=R.out+1
dR=R1/R.out #discretization step size for R (storage capacity)
r=0.005
D=exp(-r*dt) #one-period discount factor

#HC=6.57 $per barrel per year in constant case
PC1=3 #3.75/2 #PC^+ buying #3.75/2 per barrel
PC2=3 #3.75/2 #PC^- selling

p.l=3 #polynomial degree for oil and storage cost stochastic prices

```

```

library(parallel)
library(doParallel)
## Loading required package: foreach
## Loading required package: iterators
library(MASS) #for correlated rnd generation
library(iterpc) #combinatorics for feature mapping
library(matrixStats) #for row Product of fmap
temp1=cbind(1,fmap.fcn(rbind(c(2.1,3.2,4.3,5.4,6.5,7.6),c(2.1,3.2,4.3,5.4,6.5,7.6)), p.1))
temp2=cbind(1,fmap.fcn(rbind(c(2.1,3.2),c(2.1,3.2)), p.1))
temp3=rbind(c(2.1,3.2,4.3,5.4,6.5,7.6),c(2.1,3.2,4.3,5.4,6.5,7.6))
temp4=rbind(c(2.1,3.2),c(2.1,3.2))
n.basis=ncol(temp1) # 83 + 1 (intercept) 27+1 (for 2nd order)
# columns 1,2,7,22,23 are exclusively associated with the oil factors
oil.col=which( temp1[1,] %in% temp2[1,])
n.basis.sub=ncol(temp2) #number of columns of the subset data
# these columns are the pure factors of oil and storage(i.e. power one)
main.col=which( temp1[1,] %in% temp3[1,])
# these columns are the pure factors of oil only (i.e. power one) within thei
r own ploy transformation
main.col.c=which( temp2[1,] %in% temp4[1,])
rm(temp1,temp2,temp3,temp4)

#initial state (non-integer)
I_0=0 # no tanker still rented at time 0 I_0=0
R_0=0 # the tanker is empty R_0=0
T_0=0 # at time zero, the maturity of the contract held is t=0 (equivalent to
i=1)

#Correlations
rz_ksi.ksip=0 #range: +/- 0.7
rz_chi.chip=0 #range: +/- 0.55

#Numbers of sample paths
M=100000
M2=20000
#####
####

#Global parameters: part two (variables ones)
#####
####
initial.time=Sys.time()

mu.ksi.star=mu.ksi-lambda.ksi

#Generate both samples
sample.input=sample.fcn(M, M2, T1,T2,N.out,N.prim,rz_ksi.ksip,rz_chi.chip,chi
_0,ksi_0,chi2_0,ksi2_0,alp_0,alps_0,r,
#storage cost parameters
mu.ksi,k,lambda.chi,sig.chi,sig.ksi,rho,mu.ksi.star,lambda.ksi,
#storage cost parameters

```

```

mu.ksi2,K2,Phi,sig.ksi2,sig.chi2,sig.alp,lambda.ksi2,lambda.chi2,lambda.alp
,lambda.alps,
  #correlation matrix of the storage stochastic factors
  r2.chiksi,r2.chialp,r2.chialps,r2.ksialp,r2.ksialps      )

sample.time=Sys.time()-initial.time
print(c('sample computation time',sample.time))

#import the generated sample
sample.state      =sample.input$a
sample.state.2    =sample.input$b
sample.TC         =sample.input$c
sample.path2      =sample.input$d
sample.path2.storage =sample.input$e
sample.TC.2       =sample.input$f
rm(sample.input)
#####
####

```

9.11 Backward Induction

```

#prepare for parallel computing
n.cores=detectCores() #find the number of cores
cl <- makeCluster(n.cores)
clusterSetRNGStream(cl)
clusterExport(cl,c("n.cores",
                  "chi2_0", "ksi2_0", "alp_0", "alps_0",
                  "mu.ksi2", "K2", "Phi", "sig.ksi2", "sig.chi2", "sig.alp",
                  "lambda.ksi2", "lambda.chi2", "lambda.alp", "lambda.alps",
                  "r2.chiksi","r2.chialp","r2.chialps","r2.ksialp","r2.ksial
ps",
                  "I_0","R_0","T_0","chi_0","ksi_0",
                  "p.1",
                  "PC1",
                  "PC2",
                  "T1","T2",
                  "N.out","N.prim",
                  "oil.col","main.col","main.col.c",
                  "dt",
                  "R1",
                  "R.out",
                  "dR",
                  "r",
                  "D",
                  "n.basis",
                  "n.basis.sub",
                  "M","M2",
                  "mu.ksi",
                  "k",
                  "lambda.chi",
                  "sig.chi",
                  "sig.ksi",
                  "rho",
                  "mu.ksi.star",

```

```

        "lambda.ksi",
        "action.set",
        "fmap.fcn", "fmap.singlerow.fcn",
        "fut.price",
        "oppo.fcn", "oppo.c.fcn",
        "reward.fcn",
        "S.transition",
        "v.op.fcn",
        "It.indexer.fcn",
        "statespace.sub.fcn",
        "statespace.fcn",
        "getall", #from iterpc library
        "iterpc", #from iterpc library
        "rowProds" #from matrixstatistics library
    ),envir = environment() )
registerDoParallel(cl) # register the cluster

#Solve backward in time
#initialize theta to zero: Maximum real I_t=N.out+1
#Assuming Stochastic Storage Cost
Theta = array(0,dim=c(N.out, It.indexer.fcn(N.out+1), R.out+1, N.out+1, n.bas
is    ) )
#Assuming Constant Storage Cost
Theta.c=array(0,dim=c(N.out, It.indexer.fcn(N.out+1), R.out+1, N.out+1, n.bas
is.sub  ) )

#Initialize data of time t; intercept (1) included manually
t=N.out
mydata.t.input=cbind(1, fmap.fcn(cbind(sample.state[t,1,],sample.state[t,2,],
sample.state.2[t,1,],sample.state.2[t,2,],sample.state.2[t,3,],sample.state.2
[t,4,]), p.1) )

# looping backward starting from one-to-last to 2
for( t in (N.out):2 ){
    print(c("t",t))

    #set of Possible States at time t:
    pos.states=statespace.sub.fcn(t)

    #regression data: map chi_t, ksi_t, chi2_t, ksi2_t, alp_t, and alp2_t into
83 polynomial feature vectors
    #intercept (1) included manually
    mydata=cbind(1, fmap.fcn(cbind(sample.state[t-1,1,],sample.state[t-1,2,],sa
mple.state.2[t-1,1,],sample.state.2[t-1,2,],sample.state.2[t-1,3,],sample.sta
te.2[t-1,4,]), p.1) )

    #Find the optimal action pair and value for each possible state (I_t,R_t,T_
t)
    #(in parallel), and for each sample path (in vectorized computation)
    theta.par=foreach ( i=1:nrow(pos.states),.combine = rbind ) %dopar% {

        #Assuming Stochastic Storage Cost
        V = oppo.fcn(t,

```

```

        I_t=pos.states[i,1],
        R_t=pos.states[i,2],
        T_t=pos.states[i,3],
        theta = Theta,
        #TC is not defined beyond N.out+1-N.prim, i.e. T1-T2
        TC_t = sample.TC[min(t,N.out+1-N.prim),],
        mydata.t = mydata.t.input
    )

#Assuming Constant Storage Cost
V.c = oppo.c.fcn(t,
    I_t=pos.states[i,1],
    R_t=pos.states[i,2],
    T_t=pos.states[i,3],
    theta = Theta.c,
    TC_t = sample.TC[1,],#since TC=TC0 is kept constant at a
ll times

    mydata.t = mydata.t.input[,oil.col]
)

#Compute theta(t-1): Continuation Function Approximation
# It gives continuation value at stage t-1 of the value function for the
next state
# to be S(t) as a function of chi(t-1) and ksi(t-1)

if(pos.states[i,1]==0){ #need to use the storage cost data as well

    #Assuming Stochastic Storage Cost
    #regression coefficients
    fit.coef=as.numeric( lm.fit(mydata,V)$coefficients )

    #Assuming Constant Storage Cost
    fit.c.coef=as.numeric( lm.fit(mydata[,oil.col],V.c)$coefficients )

    #export all results
    list( c(fit.coef, fit.c.coef) )

} else{ #regress on the oil price columns only (doesn't need storage cost)

    #Assuming Stochastic Storage Cost
    #regression coefficients
    fit=lm.fit(mydata[,oil.col],V)$coefficients

    #Theta: update the associated columns and adjust the number of coeffici
ents to
    # n.basis to match with the dimension of Theta
    fit.coef=rep(0,n.basis)
    temp=as.numeric(fit)
    fit.coef[oil.col]=temp #add the intercept coefficient to the associated
oil columns

    #Assuming Constant Storage Cost
    fit.c.coef=as.numeric( lm.fit(mydata[,oil.col],V.c)$coefficients )

```

```

        #export all results
        list( c(fit.coef, fit.c.coef) )
    }

}

#non-parallel: update the Theta matrix using the recently computed coefficients
for (i in 1:nrow(pos.states) ){
    #I_t=pos.states[i,1], R_t=pos.states[i,2], T_t=pos.states[i,3]
    #For I_t, a transformation to a compatible integer for array indexing is
    done
    Theta [t-1, It.indexer.fcn(pos.states[i,1]), pos.states[i,2], pos.states[
i,3], ]=unlist(theta.par[i,1])[1:n.basis]
    Theta.c[t-1, It.indexer.fcn(pos.states[i,1]), pos.states[i,2], pos.states[
i,3], ]=unlist(theta.par[i,1])[(n.basis+1):(n.basis+n.basis.sub)]
}

#data(t-1) computed at time 't', will be data(t) at the previous time step,
't-1'
mydata.t.input = mydata
}

#Theta(t=1) has some NA's, replace them with 0
Theta[which(is.na(Theta))]=0
Theta.c[which(is.na(Theta.c))]=0

theta.time=Sys.time()-initial.time
print(c('Theta computation time',theta.time-sample.time))

print(c('Total time so far',theta.time))

#In-sample estimate of the value function at t0
t=1
m=1:2
mydata.t.input=cbind(1, fmap.fcn (cbind(sample.state[t,1,m],sample.state[t,2,
m],sample.state.2[t,1,m],sample.state.2[t,2,m],sample.state.2[t,3,m],sample.s
tate.2[t,4,m]), p.1) )
v.is=oppo.fcn(t=1,I_t=I_0,R_t=R_0/dR+1,T_t=T_0/dt+1,theta=Theta,TC_t=sample.T
C[1,m],mydata.t=mydata.t.input)[1]
v.c.is=oppo.c.fcn(t=1,I_t=I_0,R_t=R_0/dR+1,T_t=T_0/dt+1,theta=Theta.c,TC_t=sa
mple.TC[1,m],mydata.t=mydata.t.input[m,oil.col] ) [1]

```

9.12 Out-of-Sample Performance

```

#Analyze of the detailed out-of-sample performance using the computed Theta a
nd Theta.c
#It will compute the state/action at each timestep for each path

#Assuming Stochastic Storage Cost
v.detail=foreach (m = 1:M2, .combine = cbind)%dopar% {

```

```

#initialize (in integer form)
I_t=I_0
R_t=R_0/dR+1
T_t=T_0/dt+1

#cumulative value for each path
temp=0
#payoff at each time step
V=array(0,dim=c(N.out+1,1))
#state variables at each time step
S=array(0,dim=c(N.out+1,3))
#initialize the state variables
S[1,]=c(I_t,R_t,T_t)

for (t in 1:N.out) {

  #intercept (1) included manually
  mydata.t.input = cbind(1, fmap.singlerow.fcn(cbind(sample.path2[t,1,m],sa
mple.path2[t,2,m],sample.path2.storage[t,1,m],sample.path2.storage[t,2,m],sam
ple.path2.storage[t,3,m],sample.path2.storage[t,4,m]), p.1) )
  TC_t=sample.TC.2[min(t,N.out+1-N.prim),m]#TC is not defined beyond N.out+
1-N.prim

  # feasible action set
  X=action.set(t,I_t,R_t)

  #initialize temporary value function for every possible action
  q=array(0, dim=c(1,nrow(X)))
  C_t=array(0, dim=c(1,nrow(X)))

  #loop over all actions x in the feasible set of actions
  for (j1 in 1:nrow(X)) {

    #choose the j1-th feasible action
    xI_t=X[j1,1]
    xR_t=X[j1,2]
    xT_t=X[j1,3]

    #compute the reward of the chosen action
    C_t[j1]=reward.fcn(t,I_t,R_t,T_t,xI_t,xR_t,xT_t,
                      chi_t=sample.path2[t,1,m],
                      ksi_t=sample.path2[t,2,m],
                      TC_t)

    #compute the next state given the current state and the chosen action
    S2=S.transition(t,I_t,R_t,xI_t,xR_t,xT_t)

    # Continuation Value (CV): expected value of V1[t+1,S2] based on S2
    if(t==N.out){
      v.hat=-(S2[2]-1)*dR*PC2
    }else{
      # CV is a function of the storage cost factors only if I_t2=0, and CV
is not a

```

```

        #function of storage cost factors if I_t2!=0 by setting all the corre
sponding
        #coefficients to zero
        v.hat=Theta[t, It.indexer.fcn(S2[1]), S2[2],S2[3],]%*%t(mydata.t.inpu
t)

    }

    #total value from the j1-th pair of actions:
    # It is equal to the current payoff + discounted value function
    q[j1] = C_t[j1] + D*v.hat # D stage discount factor

}

#find the optimal action
optimal.id=which.max(q)

#the payoff ONLY comes from the immediate reward function
temp = temp + (D^(t-1))*C_t[optimal.id]

#time step t payoff
V[t]=C_t[optimal.id]

#compute the next state using the transition function
S2=S.transition(t, I_t, R_t, X[optimal.id,1], X[optimal.id,2], X[optimal.
id,3])
I_t=S2[1]
R_t=S2[2]
T_t=S2[3]

S[t+1,]=c(I_t,R_t,T_t)

if(I_t==1234){
  S[(t+1):(N.out+1),1]=I_t
  S[(t+1):(N.out+1),2]=R_t
  S[(t+1):(N.out+1),3]=T_t
  break
}

}

#add the last time stage (N.out+1) inventory left-over penalty (reward)
temp=temp-(S2[2]-1)*dR*PC2
V[N.out+1]=-(S2[2]-1)*dR*PC2

#export value
list("value"=temp,"stage.value"=V,"states"=S)

}

#Assuming Constant Storage Cost
v.c.detail=foreach (m = 1:M2, .combine = cbind)%dopar% {

```



```

#initialize ( in integer form)
I_t=I_0
R_t=R_0/dR+1
T_t=T_0/dt+1

#cumulative value for each path
temp=0
#payoff at each time step
V=array(0,dim=c(N.out+1,1))
#state variables at each time step
S=array(0,dim=c(N.out+1,3))
#initialize the state variables
S[1,]=c(I_t,R_t,T_t)

for (t in 1:N.out) {

  #intercept (1) included manually
  mydata.t.input=cbind(1, fmap.singlerow.fcn(cbind(sample.path2[t,1,m], samp
le.path2[t,2,m]), p.1) )
  # The storage cost is constant for all paths and equal to the initial val
ue
  TC_t=sample.TC.2[1,1]

  # feasible action set
  X=action.set(t,I_t,R_t)

  #initialize temporary value function for every possible action
  q=array(0, dim=c(1,nrow(X)))
  C_t=array(0, dim=c(1,nrow(X)))

  #loop over all actions x in the feasible set of actions
  for (j1 in 1:nrow(X)) {

    #choose the j1-th feasible action
    xI_t=X[j1,1]
    xR_t=X[j1,2]
    xT_t=X[j1,3]

    #compute the reward of the chosen action
    C_t[j1]=reward.fcn(t,I_t,R_t,T_t,xI_t,xR_t,xT_t,
                      chi_t=sample.path2[t,1,m],
                      ksi_t=sample.path2[t,2,m],
                      TC_t)

    #compute the next state given the current state and the chosen action
    S2=S.transition(t,I_t,R_t,xI_t,xR_t,xT_t)

    # Continuation Value (CV): expected value of V1[t+1,S2] based on S2
    if(t==N.out){
      v.hat=-(S2[2]-1)*dR*PC2
    }else{
      # CV is a function of the storage cost factors only if I_t2=0, and CV
is not a

```

```

        #function of storage cost factors if I_t2!=0 by setting all the corre
sponding
        #coefficients to zero
        v.hat=Theta.c[t, It.indexer.fcn(S2[1]), S2[2],S2[3], ]*%t(mydata.t.i
nput)

    }

    #total value from the j1-th pair of actions:
    #It is equal to the current payoff + discounted value function
    q[j1] = C_t[j1] + D*v.hat #Dont forgot D discount

}

#find the optimal action
optimal.id=which.max(q)

#the payoff ONLY comes from the immediate reward function
temp = temp + (D^(t-1))*C_t[optimal.id]

#time step t payoff
V[t]=C_t[optimal.id]

#compute the next state using the transition function
S2=S.transition(t, I_t, R_t, X[optimal.id,1], X[optimal.id,2], X[optimal.
id,3])
I_t=S2[1]
R_t=S2[2]
T_t=S2[3]

S[t+1,]=c(I_t,R_t,T_t)

if(I_t==1234){
    S[(t+1):(N.out+1),1]=I_t
    S[(t+1):(N.out+1),2]=R_t
    S[(t+1):(N.out+1),3]=T_t
    break
}

}

#add the last time stage (N.out+1) inventory left-over penalty (reward)
temp=temp-(S2[2]-1)*dR*PC2
V[N.out+1]=-(S2[2]-1)*dR*PC2

#export value
list("value"=temp,"stage.value"=V,"states"=S)

}

#Saving the data in an intuitive way for the stochastic storage cost
total.values=unlist(v.detail[seq(1,3*M2,by=3)]) #total value over each path
#generated payoff at each path and each timestep

```

```

stage.values=matrix(unlist(v.detail[seq(2,3*M2,by=3)]),ncol=M2)
stage.states=matrix(unlist(v.detail[seq(3,3*M2,by=3)]),ncol=3*M2)
S1=stage.states[,seq(1,3*M2,by=3)] #The I_t state at each path and each times
tep
S2=stage.states[,seq(2,3*M2,by=3)] #The R_t state at each path and each times
tep
S3=stage.states[,seq(3,3*M2,by=3)] #The T_t state at each path and each times
tep

#Saving the data in an intuitive way for the stochastic storage cost
total.values.c=unlist(v.c.detail[seq(1,3*M2,by=3)]) #total value over each pa
th
stage.values.c=matrix(unlist(v.c.detail[seq(2,3*M2,by=3)]),ncol=M2)
#generated payoff at each path and each timestep
stage.states.c=matrix(unlist(v.c.detail[seq(3,3*M2,by=3)]),ncol=3*M2)
S1.c=stage.states.c[,seq(1,3*M2,by=3)] #The I_t state at each path and each t
imestep
S2.c=stage.states.c[,seq(2,3*M2,by=3)] #The R_t state at each path and each t
imestep
S3.c=stage.states.c[,seq(3,3*M2,by=3)] #The T_t state at each path and each t
imestep

#find the time step when the decision to rent the tanker is made
myfun=function(X){
  temp=which(X!=0)
  if (length(temp)>0){
    return(min(temp)-1) # -1 is b/c S1 is the state, while the decision is ma
de in the prior timestep
  } else {
    return(NA)
  }
}
my.storage.rent.timestep=apply(S1, 2, myfun)
my.storage.rent.timestep.c=apply(S1.c, 2, myfun)

#find the time step when the decision to buy oil is made
myfun=function(X){
  temp=which( diff(X) > 0 )
  return( temp ) #add min() to return 'inf' if which() returns empty
}
my.fill.timestep=apply(S2, 2, myfun)
my.fill.timestep.c=apply(S2.c, 2, myfun)

#find the time step when the decision to sell oil is made
myfun=function(X){
  temp=which( diff(X) < 0 )

  temp2=which( diff(X) > 0 )

  #if the trader must sell at the very end
  if( length(temp)!=length(temp2) ) temp=c(temp,N.out+1)
  return( temp )
}

```

```

my.sell.timestep=apply(S2, 2, myfun)
my.sell.timestep.c=apply(S2.c, 2, myfun)

#find how many times we have a fill (buy) decision
myfun=function(X) {
  temp=length(unlist(X))
  return(temp)
}
my.fillnumbers=apply(as.array(my.fill.timestep),1, myfun)
my.fillnumbers.c=apply(as.array(my.fill.timestep.c),1, myfun)

#find the paths numbers on which there is more than one fill (buy) decision
which(my.fillnumbers>1)
which(my.fillnumbers.c>1)

outofsample.time=Sys.time()-initial.time
print(c('Out-of-sample computation time',outofsample.time-theta.time))
print(c('Total time',outofsample.time))

# shut down the cluster
stopCluster(c)

```

Curriculum Vitae

NAME

Behzad Ghafouri, PhD

Department of Statistical and Actuarial Sciences, Western University

EDUCATION

- 2014-2018 **Doctor of Philosophy – Financial Modelling**
Department of Statistical and Actuarial Sciences, Western University, London, ON
Thesis: “Optimal Trading of a Storable Commodity via Forward Markets”
- 2010-2011 **Master of Arts – Economics**
Department of Economics, Western University, London, ON
Project: “Effects of Diversity on Economic Outcomes”
- 2000-2004 **Bachelor of Science – Fluid Mechanics**
Mechanical Engineering Department, K. N. Toosi University of Technology, Tehran, Iran
Project: “An Investigation of Fan Coil Units Performance-Testing Standards”

SUMMARY OF ACADEMIC EXPERIENCE

- 2014-2018 **Research and Teaching Assistant** – Department of Statistical and Actuarial Sciences,
Western University, London, ON
- 2010-2011 **Teaching Assistant** – Department of Economics, Western University, London, ON
- 2007-2009 **Research and Teaching Assistant** – Mechanical and Materials Engineering Department,
Western University, London, ON

PROFESSIONAL CERTIFICATES

- 2012-2014 Passed the three levels of the **CFA (Chartered Financial Analyst) Program**

AWARDS AND SCHOLARSHIPS

- 2015 **Queen Elizabeth II Graduate Scholarship in Science and Technology (QEII-GSST)** –
Government of Ontario
- 2016 **Queen Elizabeth II Graduate Scholarship in Science and Technology (QEII-GSST)** –
Government of Ontario
- 2017 **Queen Elizabeth II Graduate Scholarship in Science and Technology (QEII-GSST)** –
Government of Ontario

PUBLICATIONS

Research Articles (Published)

[1] **B. Ghafouri**, M. Davison, A Forward Dynamic Optimization Strategy Under Contango Storage Arbitrage with Frictions, *Journal of Energy Markets*, 2017, Vol. 10-3

Dissertations

[2] **B. Ghafouri**, Optimal Trading of a Storable Commodity via Forward Markets, PhD Dissertation, Western University, 2018

[3] **B. Ghafouri**, Numerical Modeling of the Flow Over Stators and Support Arms of Automotive Cooling Fans, Master's Dissertation, Western University, 2009

Research Articles (In progress)

[4] **B. Ghafouri**, S. Moazeni, M. Davison, Optimal Liquidation of an Inventory via Forward Markets in a Dynamic Cash and Carry Trade, Under Review, 2018

[5] **B. Ghafouri**, M. Davison, Optimal Initiation and Execution of a Contango and Carry Trade, Under Preparation, 2018

Refereed Conferences

[6] **B. Ghafouri**, M. Davison, S. Moazeni, Optimal Dynamic Liquidation of an Inventory via Forward Markets in a Contango-and-Carry Storage Trade, POMS (Production and Operations Management Society) 29th Annual Conference, Huston, May 4-7, 2018

[7] **B. Ghafouri**, M. Davison, Optimal Selling of a Commodity via Forward Markets in a Cash-and-Carry Trade, BAFI (Business Analytics in Finance and Industry) 2018 Conference, Santiago, Chile, January 17-19, 2018

[8] **B. Ghafouri**, M. Davison, S. Moazeni, Optimal Commodity Selling using Spot and Forward Markets, 2017 INFORMS (Institute for Operations Research and the Management Sciences) Annual Meeting, Houston, October 22-25, 2017