



FACULTY OF SCIENCE AND TECHNOLOGY

MASTER'S THESIS

Study programme / specialisation:
Petroleum Geosciences and Engineering

The spring semester, 2023

Open

Author:
Marianne Soledad Yanez Sanchez

Supervisor at UiS:
Aojie Hong

Thesis title:
**Impact of risk attitude on optimal IOR initiation time:
A case study solved in a sequential decision-making framework powered
by machine learning-based non-linear regression.**

Credits (ECTS): 30

Keywords:

Sequential decision-making
Least-squares Monte Carlo
Risk attitude
Machine learning
Non-linear regression

Pages: 152
+ appendix: 161

Stavanger, June 15th, 2023

Abstract

Impact of risk attitude on optimal IOR initiation time: A case study solved in a sequential decision-making framework powered by machine learning-based non-linear regression

BEng. Marianne Soledad Yanez Sanchez

The University of Stavanger, 2023

Supervisor: Aojie Hong

The least-squares Monte Carlo algorithm (LSM) is an efficient approximate dynamic programming algorithm for solving sequential decision-making problems, leveraging regression. Previous studies have showcased the LSM workflow and linear regression in a sequential decision problem for optimizing the improved-oil-recovery (IOR) initiation and termination time, based on expected monetary value maximization as the decision criterion under risk neutrality.

In this work, risk attitude is introduced in the IOR optimization problem to assess the impact on the decisions. Risk behaviours are modelled using utility functions, and the optimal decision strategy is found by maximizing the expected utility. Since the utility functions introduce non-linearity, machine learning non-linear regression techniques are used in the LSM workflow to approximate the expected utilities.

Results suggest that risk-averse decision-makers prefer longer primary recovery lifetime compared to risk-neutral and risk-seeking decision-makers. This behaviour is attributed to the NPV uncertainty related to the CAPEX incurred by switching to secondary recovery. Risk-averse decision-makers prefer shorter secondary recovery lifetime. This behaviour is attributed to the OPEX and production late-stage marginal

cash inflow. The more risk-seeking the decision-maker is, the sooner they prefer to switch to secondary recovery, and the longer they would run the secondary recovery.

The value of the information increases as the decision-maker is more risk-seeking. The differences in the production lifetime decisions with the consideration of future information versus the decisions ignoring future information also increase as the decision-maker is more risk-seeking.

A change in the problem setting to a more marginal and uncertain case shows that risk-averse decision-makers would not run the project. Risk neutral decision-makers would only run the project if future information were incorporated. This reinforces the importance of sequential decision-making, where value is created from information. Risk-seeking decision-makers would run the project with or without information.

The novelties and contributions from the present work include:

- Modelling, demonstration, and discussion of the impact of different risk attitudes on decisions.
- Selection and application of the best machine learning method for non-linear regression in the LSM approach.
- Demonstration of the value of considering future information in solving sequential decision-making problems.

Dedication

To my husband...

... for being there when I could not be there.

... for empowering me when I thought it would be more difficult.

... for being a geoscientist when I wanted an opinion.

... for letting me be ... “me”.

To my daughters, Anna Victoria and Isabella...

... for growing while I grow in some other way.

... for understanding what is difficult to understand.

... for being the main drivers in everything I do.

Acknowledgements

I would like to dedicate some words to my Supervisor, Dr. Aojie Hong, not only for guiding me through the entire research and teaching me so much academically, but also for his individual qualities, passion for his work, willingness to share and setting this task for success. I am thankful for his wise mentoring and look forward to continuing working together in sharing the knowledge we have gained through his work.

I also want to thank my friends from the University of Stavanger, with whom I shared many of the courses of the master's program, for your kindness and support, making my path smoother. Marta Pernalet, Martin Coffey, Deniz Seyfeli, Alina Shashel, Siavash Heydari, Kelvin Wireko, and many others that supported me in one way or another, I want to say, "Thank you", for your openness, willingness to learn and for striving for excellence.

Table of contents

Abstract.....	ii
Dedication.....	iv
Acknowledgements.....	v
List of Figures.....	viii
List of Tables.....	xii
Glossary.....	xiii
Chapter 1: Introduction.....	1
Chapter 2: Decision analysis.....	5
2.1. Introduction to decision analysis.....	5
2.2. Expected utility theory and risk attitude.....	9
Chapter 3: Sequential decision-making.....	13
3.1. Introduction to sequential decision-making.....	13
3.2. Formulation of sequential decision problems.....	14
3.3. Solving sequential decision problems.....	17
3.4. Value of information (VOI) analysis.....	17
Chapter 4: Least-squares Monte Carlo algorithm (LSM).....	20
4.1. Monte Carlo simulation (MCS).....	21
4.2. Regression.....	22
4.3. Least-squares Monte Carlo general workflow.....	24
4.4. Example.....	27
Chapter 5: Machine learning non-linear regression methods.....	36
5.1. K nearest neighbours (KNN).....	36
5.2. Random forest (RF).....	37
5.3. Extreme gradient boost (XGB).....	37
5.4. Support vector regression (SVR).....	38
5.5. Piecewise regression.....	38

Chapter 6: General workflow of the present study	40
Chapter 7: Selection of the non-linear regression method.....	43
7.1. K-fold cross-validation	43
7.2. Scatter and Q-Q plots of regression residuals.....	47
7.3. Comparison of different regression methods.....	54
Chapter 8: Case study – Optimal time to initiate and terminate the IOR project ..	66
8.1. Problem definition	66
8.2. Least square Monte Carlo implementation for the IOR problem	75
8.3. Case study results.....	84
8.4. Change of ooi in the problem setting.....	120
Chapter 9: Discussion	127
Chapter 10: Conclusions and Recommendations	134
References.....	136
Appendix A: Sensitivity analysis of the non-linear regression hyperparameters on the fitting scores.....	139
Appendix B: Bi-dimensional regression 3D fit plots.....	141
Appendix C: LSM Utilities validation.....	144
Appendix D: LSM Python code for the IOR initiation and termination time case study with risk-neutral attitude, using OLS regression.....	146
Appendix E: LSM Python code for the IOR initiation and termination time case study with risk-neutral attitude, using XGB regression.....	146
Appendix F: LSM Python code for the IOR initiation and termination time case study with risk-averse attitude, using XGB regression.....	147
Appendix G: LSM Python code for the IOR initiation and termination time case study with risk-seeking attitude, using XGB regression.....	147

List of Figures

Figure 1. Example of a simple influence diagram. Decision is switch or not to switch to secondary recovery..	7
Figure 2: Example of a simple decision tree.....	8
Figure 3. Illustration of the equivalence rule.	10
Figure 4. Examples of trends of utility functions for different risk attitudes.	12
Figure 5: Schematic representations and decision trees for solving sequential decision-making problems using different approaches	15
Figure 6. Schematic workflow of Monte Carlo simulation	22
Figure 7. Conceptual illustration on a linear regression process, which reduces the unexplained variation in the response variable.	23
Figure 8. Conceptual illustration of the ordinary least-squares regression (OLS)	23
Figure 9. General flowchart for implementing real options values in the LSM algorithm.	26
Figure 10. Oil production rate under primary recovery as a function of time for three realizations..	28
Figure 11. Fully structured decision tree for the simple IOR initiation time exercise for a risk-averse case, with $Rho = \$800$ million..	29
Figure 12. General workflow of the present work.	42
Figure 13. Sensitivity analysis of hyperparameter ‘number of neighbours’ from KNN algorithm on the fitting scores.	46
Figure 14. Sensitivity analysis of hyperparameter ‘maximum depth’ from XGB algorithm on the fitting scores.	47
Figure 15. Scatterplots of a single-dimensional OLS regression of NPVs in a risk-neutral case.....	49
Figure 16. Scatterplots of a single-dimensional OLS regression of utilities in a risk-averse case.	50
Figure 17. Scatter and Q-Q plots of a single-dimensional OLS..	52
Figure 18. Single-dimensional Random Forest regression QC plots for risk-averse utilities.....	53
Figure 19. Single-dimensional Extreme Gradient Boost regression QC plots for risk-averse utilities	54
Figure 20. Single-dimensional regression fits using different algorithms for risk-averse utilities	56
Figure 21. Single-dimensional regression residuals using different algorithms for risk-averse utilities.	57
Figure 22. Multi- dimensional regression fits using different algorithms for risk-averse utilities.	60
Figure 23. Multi-dimensional regression residuals using different algorithms for risk-averse utilities	61
Figure 24. Bi-dimensional regression 3D fit plot, using OLS and KNN.....	62
Figure 25. Bi-dimensional regression 3D fit plot, using RF and XGB.	63
Figure 26. Scatterplots of a single-dimensional XGB regression of NPVs in a risk-neutral case.....	65
Figure 27. Influence diagram of the IOR initiation and termination time case study for risk-averse or risk-seeking profiles.....	74

Figure 28. LSM workflow for the present case study.....	76
Figure 29. Scatterplots of utilities versus oil rates at year 30 for different risk-aversion tolerances..	87
Figure 30. Scatterplots of utilities versus oil rates at year 30 for different risk-seeking tolerances..	88
Figure 31: LSM utilities validation: Risk-averse EU re-converted using Rho= \$50 million.....	91
Figure 32: LSM utilities validation: Risk-averse EU re-converted using Rho= \$100 million.....	91
Figure 33: LSM utilities validation: Risk-averse EU re-converted using Rho= \$200 million.....	92
Figure 34: LSM utilities validation: Risk-seeking EU re-converted using Rho= -\$50 million.....	94
Figure 35: LSM utilities validation: Risk-seeking EU re-converted using Rho= -\$100 million.....	95
Figure 36: LSM utilities validation: Risk-seeking EU re-converted using Rho= -\$200 million.....	95
Figure 37: LSM utilities validation: Risk-seeking EU re-converted using Rho= -\$400 million.....	96
Figure 38: LSM utilities validation: Risk-seeking EU re-converted using Rho= -\$800 million.....	96
Figure 39: Number of years of primary recovery estimated by risk-seeking decision-makers.....	97
Figure 40. Sensitivity analysis of different risk tolerances on the VOI, for risk-averse attitude and neutrality.	98
Figure 41. Sensitivity analysis of different risk tolerances on the VOI, for risk-seeking attitude and neutrality.	99
Figure 42. Boxplot of the final NPVs for different risk tolerances, for risk-averse attitude.....	100
Figure 43. Boxplot of the final NPVs for different risk tolerances, for risk-seeking attitude.....	100
Figure 44. Difference in years of primary recovery lifetime between a risk-neutral and risk-averse decision-maker.	102
Figure 45. Difference in years of secondary recovery lifetime between a risk-neutral and risk-averse decision-maker.....	104
Figure 46. Difference in years of primary recovery lifetime between a risk-neutral and risk-seeking decision-maker.	105
Figure 47. Difference in years of secondary recovery lifetime between a risk-neutral and risk-seeking decision-maker.....	106
Figure 48. Decisions with and without information of primary recovery lifetime for risk-averse and neutral attitudes.....	107
Figure 49. Decisions with and without information of secondary recovery lifetime for risk-averse and neutral attitudes.....	108
Figure 50. Difference in average decision with information and without information of primary recovery lifetime, for risk-averse and neutral attitudes.	109
Figure 51. Difference in average decision with information and without information of secondary recovery lifetime, for risk-averse and neutral attitudes.....	110

Figure 52. Boxplots of difference in decisions with information and without information of primary recovery lifetime, for risk-averse and neutral attitudes.....	111
Figure 53. Boxplots of difference in decisions with information and without information of secondary recovery lifetime, for risk-averse and neutral attitudes....	111
Figure 54. Difference in average decision with information and without information of primary and secondary recovery lifetimes, for risk-averse and neutral attitudes..	112
Figure 55. Decisions with and without information of primary recovery lifetime for risk-neutral and risk-seeking attitudes..	113
Figure 56. Decisions with and without information of secondary recovery lifetime for risk-neutral and risk-seeking attitudes..	114
Figure 57. Difference in average decision with information and without information of primary recovery lifetime, for risk-neutral and risk-seeking attitudes.	115
Figure 58. Difference in average decision with information and without information of secondary recovery lifetime, for risk-neutral and risk-seeking attitudes.....	116
Figure 59. Boxplots of difference in decisions with information and without information of primary recovery lifetime, for risk-neutral and risk-seeking attitudes.	116
Figure 60. Boxplots of difference in decisions with information and without information of secondary recovery lifetime, for risk-neutral and risk-seeking attitudes.	117
Figure 61. Difference in average decision with information and without information of primary and secondary recovery lifetimes, for neutral and risk-seeking attitudes.	118
Figure 62. Difference in average decision with information and without information of primary recovery lifetime, for risk-averse and risk-seeking attitudes.	119
Figure 63. Difference in average decision with information and without information of secondary recovery lifetime, for risk-averse and risk-seeking attitudes.	120
Figure 64. Histograms of the OOIP prior distribution corresponding to the original and new problem setting.	121
Figure 65. Sensitivity analysis on CE over different risk profiles for the IOR initiation and termination time problem with a change on the OOIP distribution.	123
Figure 66. Sensitivity analysis on VOI over different risk profiles for the IOR initiation and termination time problem with a change on the OOIP distribution. ...	124
Figure 67. Sensitivity analysis primary recovery lifetime decision over different risk profiles for the IOR initiation and termination time problem with a change on the OOIP distribution.	125
Figure 68. Sensitivity analysis secondary recovery lifetime decision over different risk profiles for the IOR initiation and termination time problem with a change on the OOIP distribution.	126
Figure 69. Sensitivity analysis of hyperparameter ‘maximum leaf nodes’ from Random Forest algorithm on the fitting scores.....	139
Figure 70. Sensitivity analysis of hyperparameter ‘C: Regularization from SVR algorithm on the fitting scores.	139
Figure 71. Sensitivity analysis of hyperparameter ‘gamma’ from SVR algorithm on the fitting scores.	140
Figure 72. Sensitivity analysis of hyperparameter ‘number of bins’ from Piecewise algorithm (RF estimator) on the fitting scores.	140

Figure 73. Bi-dimensional regression 3D fit plot, using OLS for NPV versus oil rates (risk-neutral case).....	141
Figure 74. Bi-dimensional regression 3D fit plot, using XGB for NPV versus oil rates (risk-neutral case).....	141
Figure 75. Bi-dimensional regression 3D fit plot, using SVR for utilities versus oil rates (risk-averse case).....	142
Figure 76. Bi-dimensional regression 3D fit plot, using Piecewise 10 slopes (OLS regressor) for utilities versus oil rates (risk-averse case).....	142
Figure 77. Bi-dimensional regression 3D fit plot, using Piecewise 4 slopes (XGB regressor) for utilities versus oil rates (risk-averse case).....	143
Figure 78: LSM utilities validation: Risk-averse EU re-converted using different risk tolerances..	144
Figure 79: LSM utilities validation: Risk-seeking EU re-converted using different risk tolerances.	145

List of Tables

Table 1. Likelihood functions for the measured oil production rates under primary recovery.....	28
Table 2. Paths of measured oil rates for simplified LSM example.....	30
Table 3. Utilities for alternatives at time t_2 for simplified LSM example.	31
Table 4. Expected utilities for alternatives at time t_2 for simplified LSM example. ...	31
Table 5. Utilities for alternatives at time t_1 for simplified LSM example.	32
Table 6. Expected utilities for alternatives at time t_1 for simplified LSM example. ...	32
Table 7. Utilities for alternatives at time t_0 for simplified LSM example.	33
Table 8. Table representation of optimal decision policy for simple IOR problem solved using LSM.	33
Table 9. Comparison of VOI estimated by Hong et al. (2019), replicated roll-back decision tree and LSM workflows using OLS and XGB.....	34
Table 10. Comparison of VOI estimated by rolling back the decision tree and LSM workflows using XGB.	35
Table 11. Optimized hyperparameters suggested by 3-Fold K validation using <i>GridSearchCV</i> on multi-dimensional regressions of utilities, for a risk-averse case. .	45
Table 12. R^2 scores for single-dimensional regressions of utilities versus observed oil rates, using different algorithms.	58
Table 13. R^2 scores for multi-dimensional regressions of utilities versus observed oil rates, using different algorithms.	64
Table 14: Means, standard deviations, and boundaries for the truncated normal distributions of the production model parameters.	70
Table 15: Correlation coefficients between the production-model parameters.	71
Table 16: Economic parameters.....	73
Table 17. Comparison of EVWOI, EVWI and VOI estimated by Hong et al. (2019) and the replicated Python workflow in this work.	81
Table 18. Comparison of EVWOI, EVWI and VOI estimated by the MATLAB and Python implementation when using common fixed 100 realizations.	82
Table 19. Comparison of EVWI and VOI when using OLS versus XGB in a neutral risk case.....	84
Table 20. Comparison of CEWOI, CEWI and VOI in a risk-neutral case versus a risk-averse case.	86
Table 21: Maximum LSM utility among all realizations.....	93
Table 22. CEs, VOI, and production lifetime decisions for the IOR problem with the change of OOIP settings, obtained with different risk profiles.	122

Glossary

bbl: Barrels

CAPEX: Capital expenditure

CE: Certain equivalent

CF: Cash flow

CLO: Closed-loop optimization

DWI: Decision with information

DWOI: Decision without information

EMV: Expected monetary value

ENPV: Expected net present value

EU: Expected utility

EUT: Expected utility theory

EV: Expected value

IOR: Improved oil recovery

KNN: K nearest neighbour

LSM: Least-squares Monte Carlo

MCS: Monte Carlo simulation

ML: Machine learning

MM\$: Million dollars

MMbbl: Million barrels

NLR: Non-linear regression

NPV: Net present value

OLS: Ordinary linear regression

OOIP: Original oil in place

OPEX: Operational expenditure

QC: Quality control

RF: Random forest

Rho: Risk tolerance

SDM: Sequential decision-making

SRDM: Sequential reservoir-decision-making

STD: Standard deviation

SVR: Support vector regression

VOI: Value of information

WI: With information

WOI: Without information

XGB: Extreme gradient boost

Chapter 1: Introduction

Good decision-making skills are not “born-with” abilities. Choosing wisely is a trained skill that improves with experience and observations. A “decision” is defined as a “conscious and irrevocable” allocation of resources to achieve desired objectives; it is conscious because it must be voluntary, and it is irrevocable because reconsidering a different decision later would imply loss of the resources already employed (Abbas and Howard, 2015).

Howard (1966) defined decision analysis as a “systematic procedure for transforming opaque decision problems into transparent decision problems by a sequence of transparent steps”. Based on this, Abbas and Howard (2015) indicate that decision analysis lays out four key elements of rational decision-making: objectives, alternatives, information and preferences (including risk attitude). The complexity of the decision problem depends on these elements, and additionally, uncertainty – decisions are related to the future and the future is uncertain because of incomplete or unavailable information. A decision context where successive observations/information are recorded before successive decisions is termed as sequential decision-making (SDM) (Diederich, 2001).

Hydrocarbon reservoir management is “a set of operations and decisions by which a reservoir is identified, measured, produced, developed, monitored and evaluated from its discovery through depletion and final abandonment” (Wiggins and Startzman, 1990). It commonly involves many complex decision problems under many uncertainties. Some examples of sequential decision-making problems in reservoir management are improved-oil-recovery (IOR) initiation timing (Hong et al. 2019), sequential drilling (Ahmadi and Bravtvold, 2023) and geosteering (Kullawan et al., 2018). The essential of solving for the optimal decision strategy of a sequential

decision-making problem is the consideration of future information (i.e., learning over time).

Dynamic programming approaches are designed for modelling and solving SDM problems. However, the wide use of dynamic programming approaches in the oil and gas industry is hindered by their computational intensiveness. The computational intensiveness can be partially mitigated by introducing some approximations in dynamic programming approaches, i.e., approximate dynamic programming. One of such approaches is the least-squares Monte Carlo approach (LSM) developed by Longstaff and Schwartz (2001). The key features of this approach are (1) Monte Carlo simulation for uncertainty modelling, (2) regression for the assessment of conditional expected value given information and (3) solving sequential decision problems backwards in time (i.e., starting from the last decision point). These features will be introduced in detail later.

The LSM algorithm was originally developed for real options valuation in finance and then has been applied in many other fields, such as health, mining and environment (Andréasson and Shevchenko, 2017). The energy sector is not an exception; examples of LSM applications are pricing the abandonment-timing option in an oilfield project (Jafarizadeh and Bratvold, 2012), solving the sell-spot or sell-forward problem for trading a specific volume of oil (Jafarizadeh and Bratvold, 2013) and valuation of carbon capture and storage investments under prices variation (Hongrui et al., 2016).

Hong et al. (2019) and Wui (2019) used the LSM algorithm to determine the optimal IOR initiation and termination time, in a sequential decision-making framework, which incorporates future information of oil-rate measurements for uncertainty updating (or learning) over time. They compared the decision strategy and outcomes solved using LSM with that of the “state-of-the-art” reservoir management approach – closed-loop optimization (CLO) (Wang et al., 2009). Since CLO considers

only currently available information, but not future information, to support the current decision, it can lead to a suboptimal decision strategy, compared to LSM that guarantees the (near-) optimal decision strategy.

Most LSM applications focus on risk-neutral cases, i.e., expected monetary value (EMV) maximization and use linear regression for conditional EMV assessment. The impact of risk-averse and seeking attitudes on sequential decisions have not been extensively investigated and discussed in the literature. When risk-averse and seeking attitudes are involved, utility functions and expected utility maximization should be applied (Von Neumann and Morgenstern, 1947). Risk-averse and seeking utility functions introduce non-linearity. Approximating a non-linear relationship using linear regression poses a number of challenges, including regression fit errors. Andréasson and Shevchenko (2017) reviewed many attempts to resolve this problem, but the solutions either add computational time or introduce bias. They suggested non-linear regression based on neural networks as an alternative to improve the fit to a non-linear relationship. In general, machine learning-based non-linear regression could be used.

The primary goal of this thesis work is to investigate the impact of risk attitudes on sequential decisions, based on the synthetic case study introduced by Hong et al. (2019). The novelties and contributions include:

- Modelling, demonstration and discussion of the impact of different risk attitudes on the IOR initiation and termination time decisions.
- Selection and application of the most suitable machine learning method for non-linear regression in the LSM algorithm.
- Demonstration of the value of considering future information in solving sequential decision-making problems.

This thesis is organized in the following manner. First, decision analysis and risk attitude modelling will be briefly reviewed. Second, the steps of LSM application will be detailed with an example. Third, some commonly used machine learning

methods for non-linear regression will be described. Forth, the results of the thesis work will be presented and discussed, starting with the regression performances of the machine learning methods and then the results of the IOR initiation and termination time decision problem solved with the selected non-linear regression method and LSM. Finally, the thesis work will be concluded, including a discussion and recommendations for future works.

Chapter 2: Decision analysis

In this chapter basic concepts of decision analysis are introduced. Additionally, the concept of risk attitude and expected utility theory are reviewed.

2.1. INTRODUCTION TO DECISION ANALYSIS

The discipline of decision analysis was founded by Ronald A. Howard in 1960s. Decision Analysis is a systematic procedure for transforming opaque decision problems into transparent decision problems by a sequence of transparent steps (Howard, 1966).

A simple definition of decision-making is “choosing the alternative that best fits a set of goals” (Bratvold and Begg, 2010). Three elements define the foundation of decision modelling: objectives, alternatives, and information. Objectives define the goals that a decision-maker desires to reach. Alternatives are acts that the decision-maker can take for reaching the objectives. A decision problem must have at least two alternatives from which for the decision-maker to choose one; otherwise, the decision-maker have no choice and no decision to make if there is only one alternative. Information is used to predict how well or poorly every alternative can reach the objectives (Bratvold and Begg, 2010).

A “good outcome (or result)” is defined as “a future state of the world that we prize relative to other possibilities”, and a “good decision” is defined as “an action we take that is logically consistent with our objectives, the alternatives we perceive, the information we have, and the preferences we feel” (Abbas and Howard, 2015). One type of preference, involved in many decision problems, is our feelings of uncertainty – our desirability of the variety of possible outcomes/consequences of our actions (Jordaan, 2004). The feeling of uncertainty is quantified and modelled in the expected utility theory (EUT) which will be reviewed in the next section.

Probability is the measure or quantification of uncertainty. The elicitation of a probability aims at a quantitative statement of opinion about the uncertainty associated

with the event under consideration (Jordaan, 2004). Many uncertain events have outcomes described by quantitative variables. In the oil and gas industry, some examples are original oil in place (OOIP), hydrocarbon prices, and profit. These outcomes are continuous, meaning they can take on any value between defined bounds. An example of discrete variables is the number of wells. No matter a variable is continuous or discrete, if one doesn't know and cannot control or decide on the value of the variable for sure, the variable is an uncertain variable. OOIP, future prices, and future profit are uncertain variables, whilst number of wells is usually a control/decision variable because one can decide how many wells one wants to drill.

One very common approach to making decisions under uncertainty is to estimate one single value¹ of each uncertain variable and use the estimated values to calculate the payoff of each objective, for example net present value (NPV) for monetary payoffs. This approach is called “deterministic” because it considers only a single “deterministic” assessment of the outcomes of uncertain variables and ignores other possible outcomes. Thus, the deterministic approach often leads to sub-optimal decisions and values. The statistically correct “probabilistic” approach is to consider all possible outcomes of each key uncertain variable by assessing probability distributions of possible outcomes for the key uncertain variables. In this case the decisions are based on the probability distributions of objective payoffs associated with the underlying uncertain variables.

Influence diagrams and decision trees are useful decision analysis tools for framing and solving decision problems. An influence diagram illustrates the decisions, uncertain variables, calculations, values, and connecting arrows that symbolize the influences between them. Uncertain variables and events are typically shown by ellipses; decision nodes are typically shown as rectangles; and value nodes are shown as rounded rectangles (Von Winterfeldt and Edwards, 2007).

¹ The single value may be a base value, a most-likely value, an expected value, etc.

Figure 1 shows an example of a simple influence diagram. Different uncertain events are represented by ellipses. The problem involves a single decision, represented by the rectangle, and the value is represented by the rounded rectangle. As an example, the decision of this problem could be the production strategy: continue with primary recovery or to switch to secondary recovery. Uncertainties are the oil production rate and oil price, and the value is the revenue. The arrow that connects the decision node to the revenue node indicates that the decision directly affects the revenue. The arrows that connect the uncertain events to the revenue node indicate that the uncertain events directly affect the revenue. The arrow from the decision to the oil production indicates that the outcomes of this uncertainty depend on the decision alternatives. If an arrow is drawn from an uncertain event to the decision node, this means that the outcome of the uncertain event is ‘known’ before the decision is made.

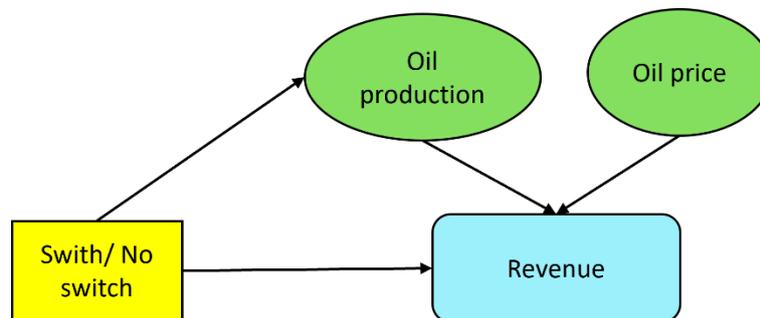


Figure 1. Example of a simple influence diagram. Decision is switch or not to switch to secondary recovery. Uncertainties are oil rates and oil price. Value is the revenue.

Decision trees can be easily used to visualize and communicate the sequence of decisions and uncertain events as well as decision alternatives and uncertain outcomes. A decision node is represented by a square node, with labelled branches to the side representing the alternatives. An uncertainty is represented by a circle node, with branches to the right indicating possible outcomes, also known as “prospects;” the group of possible outcomes or prospects represents a “deal”.

Figure 2 shows the respective decision tree for influence diagram in Figure 1. The alternatives are to continue with primary production or to switch to secondary recovery. The uncertain events are the oil production rate and oil price, from which independent possible outcomes are drawn as branches to the right; these represent the prospects. Numbers on each branch represent the probability of each possible outcome, and the numbers at the right end of each branch represent the payoffs, in this case revenue. The probabilities of the oil prices are independent of the oil production rates and decision alternatives. However, as indicated in the influence diagram in Figure 1, the probabilities of the oil rates are different depending on which alternative from the decision is taken.

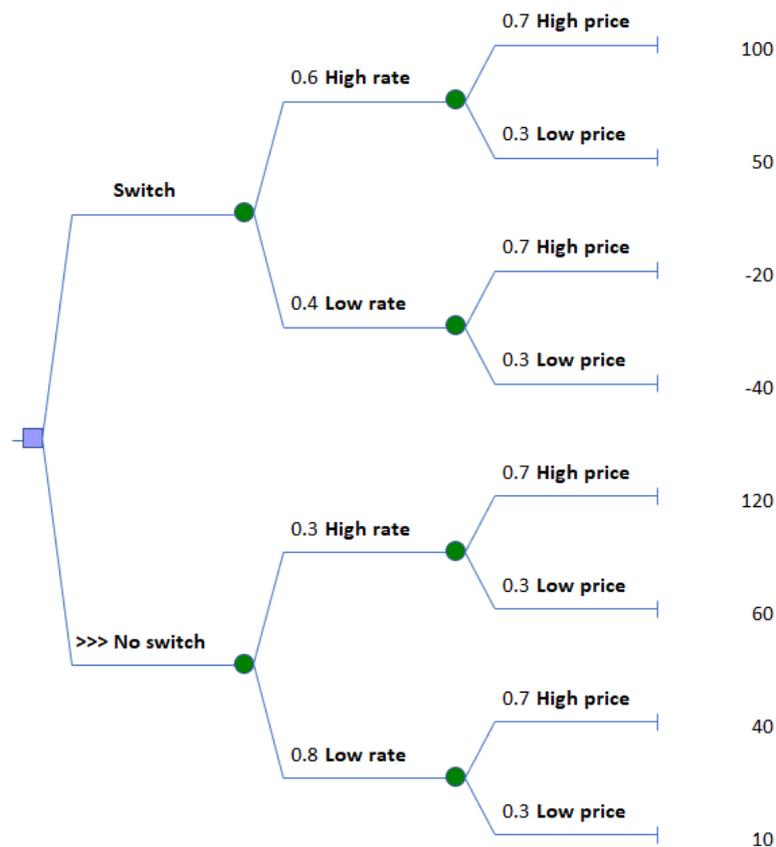


Figure 2: Example of a simple decision tree. Decision is switch or not to switch to secondary recovery. Uncertainties are oil rates and oil price. Value is the revenue.

Later in this thesis, a full decision tree and an influence diagram will be presented to illustrate the specific sequential decision problem studied.

2.2. EXPECTED UTILITY THEORY AND RISK ATTITUDE

The concept of utility can be traced back to Nicolaus Bernoulli in 1700s. In the decision theory, utility is defined as the personal value of a monetary value. The expected utility theory (EUT), developed by Von Neumann and Morgenstern (1947) based on 4 axioms, implies that all rational decision-makers must have utilities that can be numerically expressed, through a utility function that accounts for a decision-maker's attitude towards uncertainty, and leads to consistent and normative decisions. EUT has been practically implemented in various fields such as economy, game theory and decision theory, but few implementations in the oil and gas industry (Hong and Bartvold, 2023).

Abbas and Howard (2015) reformulated the Von Neumann and Morgenstern (1947)'s axioms to the five rules of actional thought:

1. The probability rule: A decision-maker must be able to assess a probability distribution of possible outcomes for an uncertain event.
2. The order rule: the decision-maker must be able to order prospects according to their preferences.
3. The equivalence rule: As illustrated in Figure 3, given prospect A is preferred to prospect B, and prospect B is preferred to prospect C (noted as $A > B > C$), the decision-maker must be able to assign a probability p so that they are indifferent between getting B for sure (no uncertainty) and a deal with a chance p of getting A and $1-p$ of getting B. In such case, the probability p is called the decision-maker's preference probability for

prospect B. The concept of preference probability and its relation to utility will be detailed below.

4. The substitution rule: Provided the indifference between prospect B and the A-C deal as illustrated in Figure 3, the decision-maker must be willing to exchange prospect B with the A-C deal, or vice versa.
5. The choice rule: By following the 4 rules above and with the assigned preference probabilities of all prospects, the decision-maker's preference to a deal can be quantified by a (calculated) preference probability. For different deals, the decision-maker should take the deal that has the greatest preference probability.

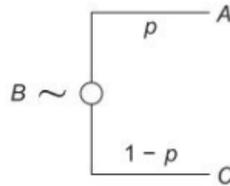


Figure 3. Illustration of the equivalence rule. From Abbas and Howard (2015).

The equivalence rule introduces the concept of preference probability and certain equivalent. The probability p in Figure 3 describes the decision-maker's preference of prospect B on the basis of prospect A (the most preferred) and C (the least preferred), and thus p is called the decision-maker's preference probability for prospect B. Prospect B is called the certain equivalent (CE) of the A-C deal because to the decision-maker, B (a certain deal) is equivalent (or exchangeable) to the A-C deal (an uncertain deal) (Abbas and Howard, 2015). With the concept of certain equivalent, a deal can be represented by a corresponding certain equivalent and for different deals, the decision-maker should take the deal that is equivalent to a certain equivalent that is the most preferred among the certain equivalents of all deals (an alternative expression of the choice rule). The rankings of different deals by certain equivalents and by

preference probabilities should be the same because a more preferred certain equivalent (or prospect) should have been assigned a greater preference probability.

The curve formed by plotting preference probabilities on the y-axis versus monetary amounts on the x-axis is called a utility function (Figure 4). A utility function summarizes all preference information necessary for a decision-maker to make a rational decision; in other words, the decision-maker spontaneously obey the five rules of actional thought if he/she uses a utility function.

For monetary prospects, a utility function is defined for the whole range of wealth, and the only requirement is that a utility function must be non-decreasing as the wealth increases, which reflect more money is preferred to less money in general. A utility function describes a decision-maker's feeling about changes of their wealth relative to the initial state of their wealth and their feelings about uncertainty (Howard, 1966).

For a risk-neutral decision-maker, expected monetary value (EMV) is the consistent decision criterion (i.e., the objective is to maximize EMV). However, when a decision-maker is risk-averse or risk-seeking, EMV is no longer a consistent decision criterion; the consistent decision criterion should be expected utility (EU).

For a risk-neutral decision-maker, the utility function must be linear:

$$u(x) = a + bx \quad (1)$$

where u denotes utility and x denotes monetary gain (or loss if x is a negative value) and a and b ($b > 0$) are coefficients that can be arbitrary values because a and b have no impact on CE and ranking of decision alternatives.

For risk-averse and risk-seeking attitudes, they are commonly modelled using an exponential utility function because of its simplicity – only one parameter, risk tolerance, needs to be specified:

$$u(x) = a - \text{sgn}(\rho) b e^{-\frac{x}{\rho}} \quad (2)$$

where ρ ($\rho \neq 0$) denotes risk tolerance and $\text{sgn}(\rho)$ takes the sign $\text{sgn}(\rho) = 1$ if ρ is positive and $\text{sgn}(\rho) = -1$ if ρ is negative. For risk aversion $\rho > 0$, and for risk-seeking $\rho < 0$. ρ approaching to $+\infty$ represents risk-neutrality: ρ approaching to 0 from the positive direction represents extreme risk aversion, and ρ approaching to 0 from the negative direction represents extreme risk-seeking.

Figure 4 shows an illustration of the shapes of the utility function for different risk attitudes.

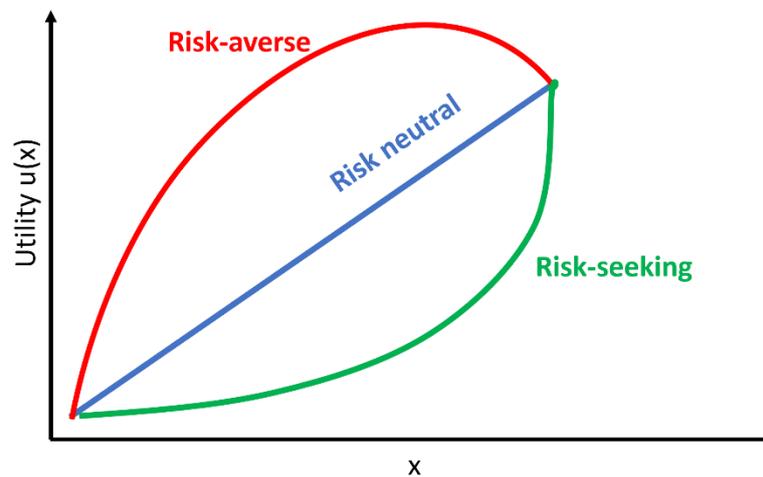


Figure 4. Examples of trends of utility functions for different risk attitudes.

For a risk-averse decision-maker the CE for an uncertain prospect is smaller than the expected value (EV) of the prospect; the CE is equal to the EV under risk-neutrality; for risk-seeking decision-makers the CE is greater than the EV (Howard, 1966).

Chapter 3: Sequential decision-making

In this chapter sequential decision analysis will be introduced, explaining how sequential decision problems can be formulated and solved. A review of dynamic programming and approximate dynamic programming will be provided. After explaining how the information can be used in the different SDM approaches, the general concept of value of information (VOI) will be explained.

3.1. INTRODUCTION TO SEQUENTIAL DECISION-MAKING

Warren Powell, author of the book “Approximate Dynamic Programming: Solving the curses of dimensionality” (Powell, 2011), wrote: “Sequential decision analytics is an umbrella for a vast range of problems that consist of the sequence: decision, information, decision, information, decision, information, ...”. The outcome of each decision (as a part of information) can affect future decisions. The objective of solving an sequential decision-making problem is to find the optimal decision strategy – a plan of actions depending on what has happened earlier and what information has revealed – that optimizes an objective (e.g., maximizing profit, minimizing cost or maximizing NPV) (Bratvold and Begg, 2010).

In a SDM scheme, multiple decisions are ordered chronologically, and the outcome of a previous decision is considered known before a next decision is made. Additional information gathered over time is driven by the decisions taken before the information can be used to update the probabilities of remaining uncertainties, and consequently update the expected objective values of all future decision alternatives. Therefore, the decisions with the consideration of the additional information might be different from that if the additional information is not considered, and the decision policy might also change (Ahmadi and Bratvold, 2023). SDM relies on the integration of additional information gathered and continuous uncertainty updating given the additional information over time, to support current decision(s).

3.2. FORMULATION OF SEQUENTIAL DECISION PROBLEMS

The SDM framework can be represented by a decision tree structure, which is reviewed in more detail in this section. This section also shows the differences between the full formulation of SDM problems from other decision-making approaches that do not fully consider the information gathered over time.

3.2.1. Sequential Decision trees

As introduced in the previous chapter, a decision tree is a diagram that expresses the relationship and time-sequence between the main elements of a decision problem: decisions, uncertainties and payoffs (Bratvold and Begg, 2010). In Chapter 4: LSM algorithm, Section 4.4: Example, a fully structured decision tree example is introduced to explain the least-squares Monte Carlo method. Yet, to understand the different SDM approaches and the relation between their components Figure 5 shows compact versions of decision trees². The decision points are denoted D_i , the prior uncertainty³ is denoted as $U|I_0$ and additional information gathered after decision D_{i-1} and before D_i is denoted as I_i . Hong et al. (2019) and Ahmadi and Bratvold (2023) well describe three different approaches for solving sequential decision-making problems. The first one is represented in Figure 5a, called the myopic approach. In this approach, a single decision point and an initial uncertainty $U|I_0$ are considered; other future decisions and information are disregarded, thus, the current decision D_0 depends only on the prior information I_0 .

² A compact decision tree omits details of alternatives and possible outcomes, and only decision and uncertainty nodes are shown.

³ A prior probability refers to the initial estimate of the probability of an outcome. A posterior probability refers to the updated probability of the outcome, given some other outcome happens; this is usually known as conditional probability.

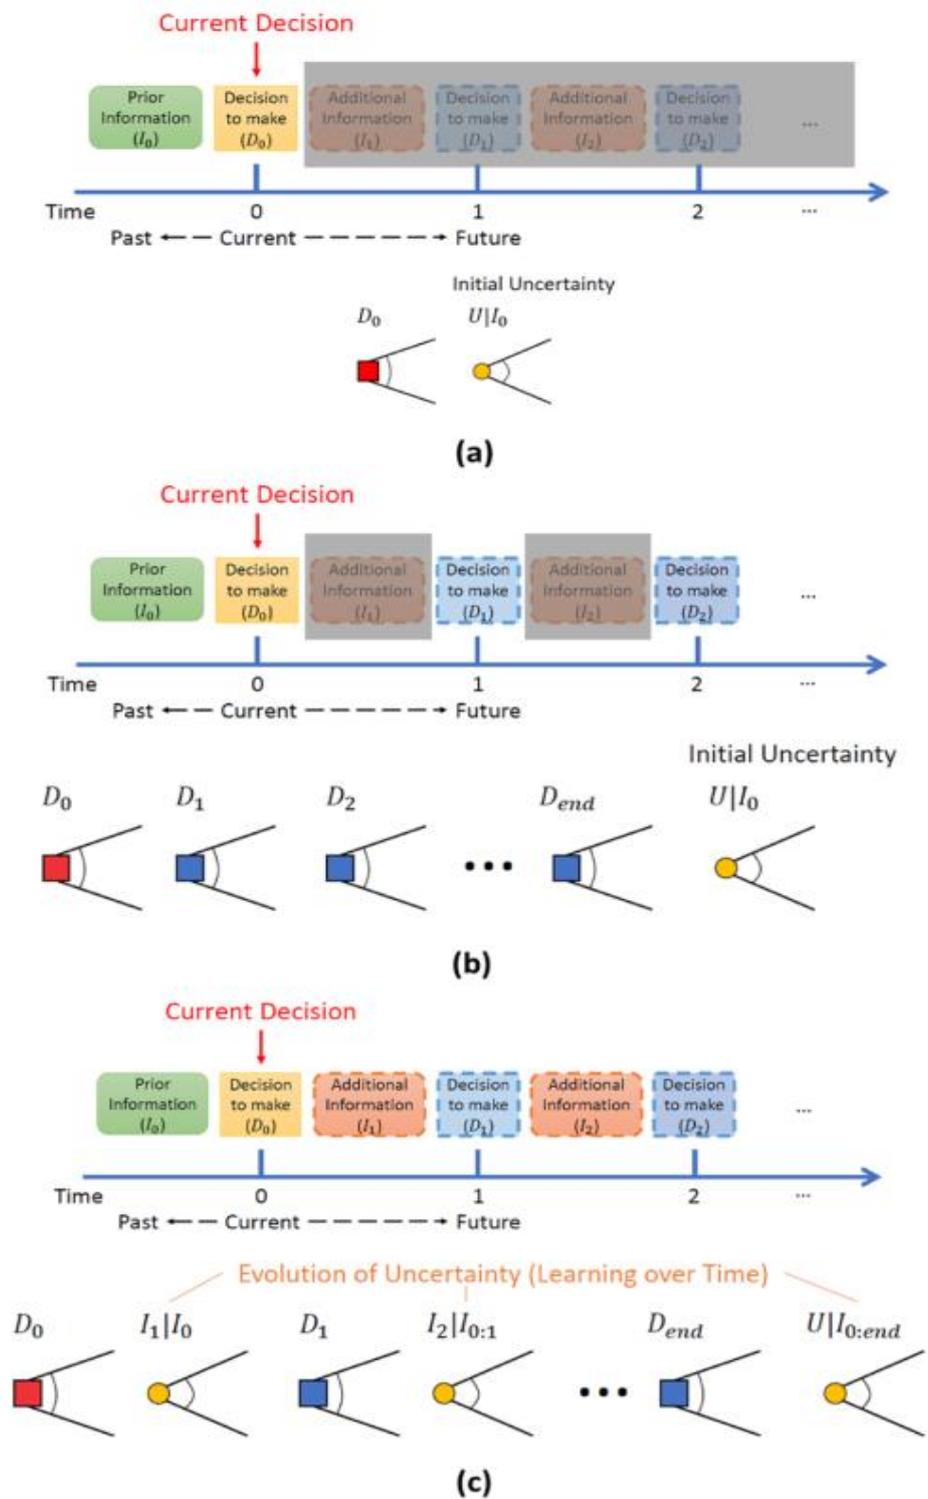


Figure 5: Schematic representations and decision trees for solving sequential decision-making problems using different approaches; (a) myopic approach, (b) naïve approach, and (c) farsighted approach. From Hong et al. (2019), Ahmadi and Bravtvoid (2023).

In the naïve approach, illustrated in Figure 5b, also only prior information is considered, but instead of one decision point, the sequence of decisions is considered; future information is ignored for the decision-making. The farsighted/dynamic programming approach, Figure 5c, is the one that considers the future information obtained over time; additional information will be gathered, and uncertainty probabilities will be updated for supporting a next decision in time. This approach models the full decision-information structure, guaranteeing the optimal decision policy (Hong et al., 2019). The optimal decision policy is solved using the rollback procedure of decision tree, from the last point in time towards time zero, which is described in the sub-section below.

3.2.2. Bellman equation

The mathematical formulation behind the decision tree of the farsighted approach is the Bellman equation (Bellman, 1957), formulated below.

$$V_i(\omega_i) = \max_{a \in A(i)} \left[\sum_{x \in X(a)} P(x|\omega_i) (r_a^x + \delta V_{i+1}(\omega_{i+1}^{a,x})) \right] \quad (3)$$

where i represents a decision point that includes available information ω_i , and $A(i)$ represents a set of feasible alternatives. A single alternative in $A(i)$ is noted as a , with corresponding uncertain outcomes $X(a)$. $P(x|\omega_i)$ represents the conditional probability of an uncertain outcome x given information ω_i . r_a^x represents the immediate payoff from alternative a given outcome x , and $V(\omega_i^{a,x})$ represents the future expected value, estimated by selecting the best alternatives in the future, given alternative a , and outcome x . Hence, $V(\omega_i^{a,x})$ corresponds to the optimal expected value computed at the future decision nodes (D_i for $i > 0$) in Figure 5c. δ is the discount factor.

3.3. SOLVING SEQUENTIAL DECISION PROBLEMS

This addresses how complex SDM problems are solved, and what techniques have been developed to overcome high computational costs.

3.3.1. Curses of dimensionality and dynamic programming

It is acknowledged that the myopic and naïve models might lead to suboptimal results, but it is also understood that the farsighted approach can be computationally costly. The complexity of a full decision tree for sequential decision-making grows quickly with the number of decision points, the number of alternatives at each decision point, the number of uncertainties and the number of possible outcomes; this is known as the “curse of dimensionality” (Powell, 2011).

A large range of SDM problems are solved using dynamic programming (Bellman, 1957). Dynamic programming enables the identification of optimal decision policies of SDM problems under uncertainty. It uses the principle of decomposition to break a general decision problem into a set of sub-problems, making it the most suitable approach for the optimization of sequential decisions and uncertainties (Ahmadi and Bravtvoid, 2023). However, this approach also suffers from the curse of dimensionality and the computational cost increases exponentially when solving large, real-world decision schemes (Bellman, 1957).

Approximate dynamic approaches have been developed to partially overcome the curses of dimensionality. The LSM algorithm used in the present work falls within this umbrella and is based in approximating the expected value by using regression. The details of the LSM algorithm are documented in the next chapter.

3.4. VALUE OF INFORMATION (VOI) ANALYSIS

Decisions can be supported by acquiring and using information. The value associated with collecting the information is called value of information (VOI).

Although the VOI is a concept that applies to general decision analysis (not necessarily sequential decision analysis), it is included in this section after the various approaches for solving SDM problems have been explained in relation to how the information is used.

Abbas and Howard (2015) describe the VOI as the maximum amount that a decision-maker should pay for gathering additional information. The information gathered can be perfect or imperfect, where the first one is associated with the uncertainty being revealed in full (Ahmadi and Bratvold, 2023). When the information is imperfect, the outcomes of the decision problem are uncertain, and this uncertainty is quantified assigning probabilities. In general, the use of information aims to make better decisions, i.e., increasing the chances of obtaining the desired outcomes.

In the case of decision-makers with neutral or constant risk aversion preferences (see Chapter 2, section 2.2: Expected utility theory and risk attitude), VOI is calculated as the difference between the expected value obtained with free additional information and the expected value obtained with no additional information (Bratvold and Begg, 2010):

$$VOI = \max\{0, \Delta\} \quad (4)$$

$$\Delta = EVWI - EVWOI$$

where EVWI is the expected value with additional information and EVWOI is the expected value without additional information. The minimum value of VOI is 0 because if $EVWI < EVWOI$, one can always choose to not gather the additional information (Abbas and Howard, 2015). In the case of risk profiles described with an exponential utility function, the same formulas apply, but replacing the EVs with certain equivalents (CEs).

Since the uncertain variables normally have a high number of dimensions, it is difficult to represent their probability distribution through an analytical form. The probability distributions of these variables are usually approximated using Monte Carlo

sampling (MCS) (see Chapter 4, section 4.1: Monte Carlo simulation). The decision without information (DOWI) is the alternative that maximizes the EV over the prior value:

$$EVWOI = \max_{a \in A} \left[\int_x v(x, a) p(x) dx \right] \approx \max_{a \in A} \left[\frac{1}{B} \sum_{b=1}^B v(x^b, a) \right] \quad (5)$$

where x denotes the underlying uncertain variables of value v . a is a decision alternative from the set A , x represents the distinctions of interests, $v(x, a)$ is the value for an alternative a and a given realization x , and $p(x)$ is the prior probability density of x . x^b represents a sample of x from MCS.

Similarly, the decision with information is an alternative that maximizes the expected value over the posterior value given a realization of observations y :

$$EVWI(y) = \max_{a \in A} \left[\int_{x|y} v(x|y, a) p(x|y) dx|y \right] \approx \max_{a \in A} \left[\frac{1}{B} \sum_{b=1}^B v(x^b|y, a) \right] \quad (6)$$

where $x|y$ denotes x given y and $p(x|y)$ is the posterior⁴ probability density of x given y . The expected value with information is then:

$$EVWI = \left[\int_y EVWI(y) p(y) dy \right] \approx \left[\frac{1}{C} \sum_{c=1}^C EVWI(y^c) \right] \quad (7)$$

where $p(y)$ is the marginal probability density of y , i.e., the preposterior⁵ probability density of y (Hong et al., 2018).

⁴ A posterior probability refers to the updated probability of the outcome, given some other outcome happens; this is usually known as conditional probability.

⁵ Preposterior probability refers to the probability of the evidence/observations.

Chapter 4: Least-squares Monte Carlo algorithm (LSM)

The least-squares Monte Carlo (LSM) algorithm, proposed by Longstaff and Schwartz (2001), is a state-of-the-art approximate dynamic programming method developed for real options valuation. LSM has several advantages compared to other methods for real options valuation: Black-Scholes option pricing model (Black and Scholes, 1973), binomial model (Cox et al., 1979), risk-adjusted decision trees (Brandao et al., 2005), Monte Carlo simulation (MCS) (Boyle, 1977), and hybrid real options models (Neely III and de Neufville, 2001). LSM consists of two main steps: simulation of uncertainties and least-squares regression (Jafarizadeh and Bratvold, 2013); this is followed by solving the sequential decision problem backwards in time. The conditional expectation function is estimated by regressing the values, e.g., NPVs, utilities, as a function of the uncertain variables at each decision point. The uncertain variables are generated via Monte Carlo simulation (MCS) at all time periods. The optimal decision policy is identified by maximizing the expected value or expected utility at each decision point. The regression function to use constitutes an important element of the LSM algorithm (Ahmadi and Bratvold, 2023).

The LSM algorithm allows to efficiently handle multiple uncertainties, mitigating their curse of dimensionality. However, this technique still suffers from the curses of dimensionality of decision points and decision alternatives. Thus, most implementations of LSM are on decision problems with only two or three alternatives at each decision point, such as the “continue” and “exercise” alternatives in real options.

The following two sub-sections explain the main two stages of the LSM algorithm, and a further sub-section explains the general LSM workflow for sequential decision-making.

4.1. MONTE CARLO SIMULATION (MCS)

Monte Carlo simulation (MCS) is a popular tool used for uncertainty analysis. In the context of decision analysis, direct assessment of probability distributions of objective attributes (e.g., NPV and hydrocarbon recovery) is challenging. Instead, these attributes are often calculated based on models that relates the attributes (output variables) to underlying uncertain variables (input variables) whose probability distributions are easier to assess (e.g., reservoir properties, costs and prices). Monte Carlo simulation in this context is used to propagate the uncertainty from the underlying variables to the objective attributes (Bratvold and Begg, 2010).

The MCS procedure is as follows: Given a model, the uncertainty of input variables is modelled with probability distributions. A set of the values of the input variables is sampled from the probability distributions and used in the model to calculate a value of an output variable. The value of the output variable is stored. The sampled values of the input variables and the corresponding value of the output variable is called a “realization”. The procedure of sampling the input variables, calculating and storing the output variable is repeated multiples times, and the statistics (e.g., mean, variance, percentiles) of the output variable can be approximated by the statistics of the calculated and stored values. Figure 6 illustrates the MCS procedure.

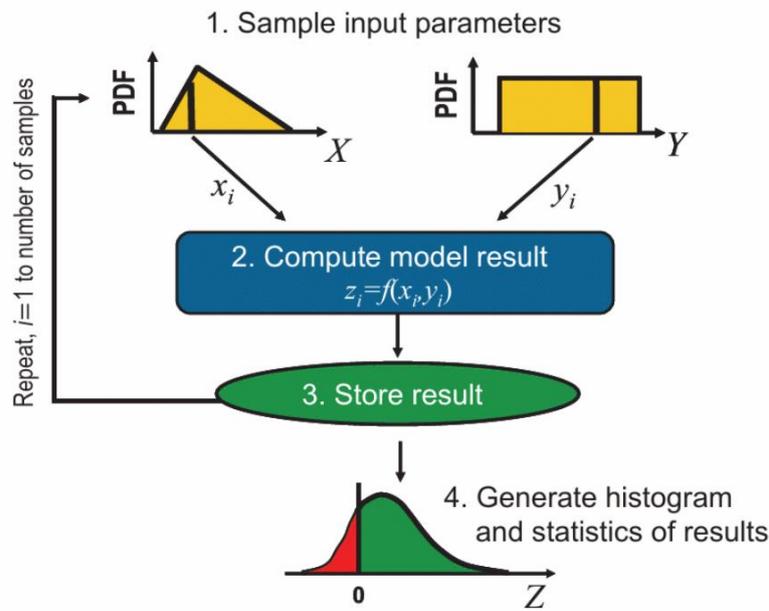


Figure 6. Schematic workflow of Monte Carlo simulation. From Bratvold and Begg (2010).

4.2. REGRESSION

Regression analysis is used to estimate relationships between independent variables, often called predictors, and dependant variables. The objectives of regression analysis are: (1) to identify the “best” model among several candidates in case the physics of the system does not provide an unique mechanistic relationship, and (2) to determine the “best” values of the model parameters (Agami Reddy, 2011).

The linear regression model is widely used, where a linear function is used to approximate the relation between the predictors and dependant variables. The usefulness of linear models is very broad, even to fit non-linear functions when the range of predictors is limited to a specific range (Cook and Weisberg, 1982).

The conceptual illustration of a linear regression process is shown in Figure 7. The main objective of the regression is to minimize the deviations of the points from the prospective regression line. The method of the least-squares is one of the most common methods used, where the ‘best fit’ criteria is defined by a minimization of the sum of the squares of the residuals (Agami Reddy, 2011).

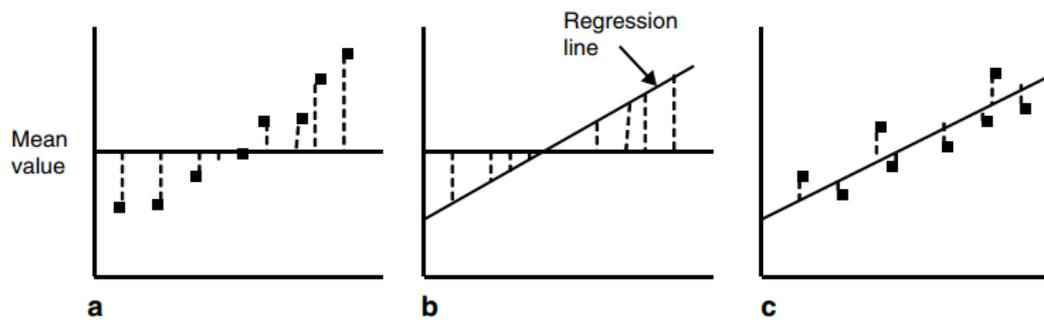


Figure 7. Conceptual illustration on a linear regression process, which reduces the unexplained variation in the response variable. (a) Total variation before regression. (b) Explained variation due to regression. (c) Residual variation after regression. From Agami Reddy (2011).

4.2.1. Ordinary least-squares (OLS)

One of the most commonly used linear prediction methods is the ordinary least-square (OLS) regression. This is illustrated in Figure 8 , where the squared sum of the vertical differences D_i between the line and the observation points are minimized $((\sum_{i=1}^n D_i^2)^{1/2})$.

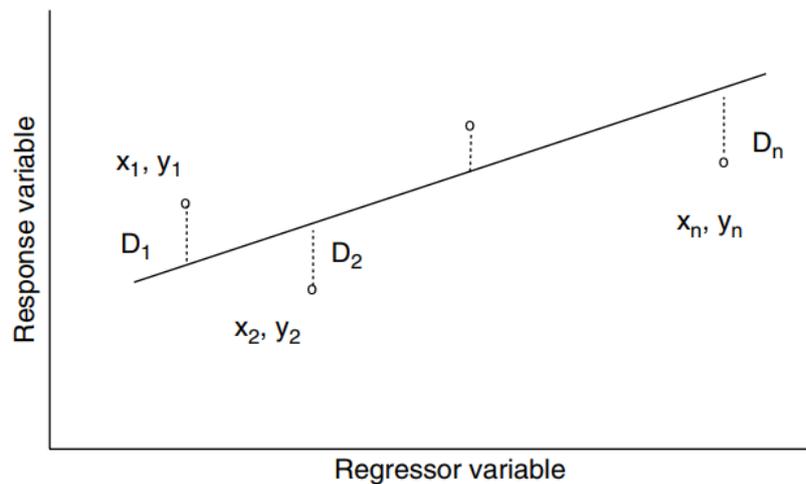


Figure 8. Conceptual illustration of the ordinary least-squares regression (OLS). From Agami Reddy (2011).

A single-dimensional or univariate regression can be represented only with two parameters, a and b , where a represents the intercept and b represents the slope. The

objective of the regression is to determine the values of a and b so they can satisfy equation 8, with a minimum error ϵ .

$$y = a + bx + \epsilon \quad (8)$$

When more than one independent variable influences the dependant variable, a multivariate regression will explain more of the variation and provide a better prediction than a single regression. The mathematical form, known as the “additive multiple linear regression model” is defined as

$$y = \beta_0 + \beta_1x_1 + \dots + \beta_px_p + \epsilon \quad (9)$$

where x_i are the independent variables, β_i are unknown constants (regression model parameters) and ϵ is the regression error (Agami Reddy, 2011). In this case the parameters β_i are to be identified such as the sum of squares function L is minimized: $L = \sum_{i=1}^n \epsilon_i^2 = (Y - X\beta)'(Y - X\beta)$, where X and Y represents matrixes of the multiple variables x and y variables, and β is a vector containing $\beta_0, \beta_1, \beta_2, \dots$

4.3. LEAST-SQUARES MONTE CARLO GENERAL WORKFLOW

Ahmadi and Bravtvoid (2023) recently documented a very structured explanation of the general LSM workflow based on real options, illustrated in Figure 9, and divided it in four main steps. Below is a summary of these general LSM steps:

1. Simulate sufficient realizations of the uncertain variables for all time periods via MCS. This will constitute the multiple “paths.”
2. Calculate the cash flows for all the paths that are functions of the simulated uncertain variables and other deterministic variables. Often the payoffs are added along each path as cumulative cash flows and discounted back to the initial point in time by applying a discount rate, generating the net present value (NPV).

3. Execute the optimal decision policy employing the backward induction strategy (Bellman, 1957), starting from the last decision point back to time zero:
 - 3.1. In the last period, there are no further options in the project, hence the optimal policy is found by comparing the available alternatives' payoffs and exercising the alternative with the highest expected value.
 - 3.2. For previous time periods, the immediate payoff of the exercising options should be compared to the expected value for continuing the project in the next time step. In the LSM algorithm, the conditional expected value is found by least square regression:
 - a. Fit a regression function $Y = f(X)$, where X represents the set of uncertain variables in the current period and Y are the discounted continuation values of the project in the next period
 - b. The optimal decision policy is found by the highest payoff alternative, including the continuation alternative.
 - c. Update the NPV matrix at each time step, by including the continuation values calculated in the previous steps.
 - d. Repeat consecutively for each time step, backwards in time.
4. At time zero (first decision point) calculate the average of the values, which represents the expected value for the project with the optimal decision strategy.

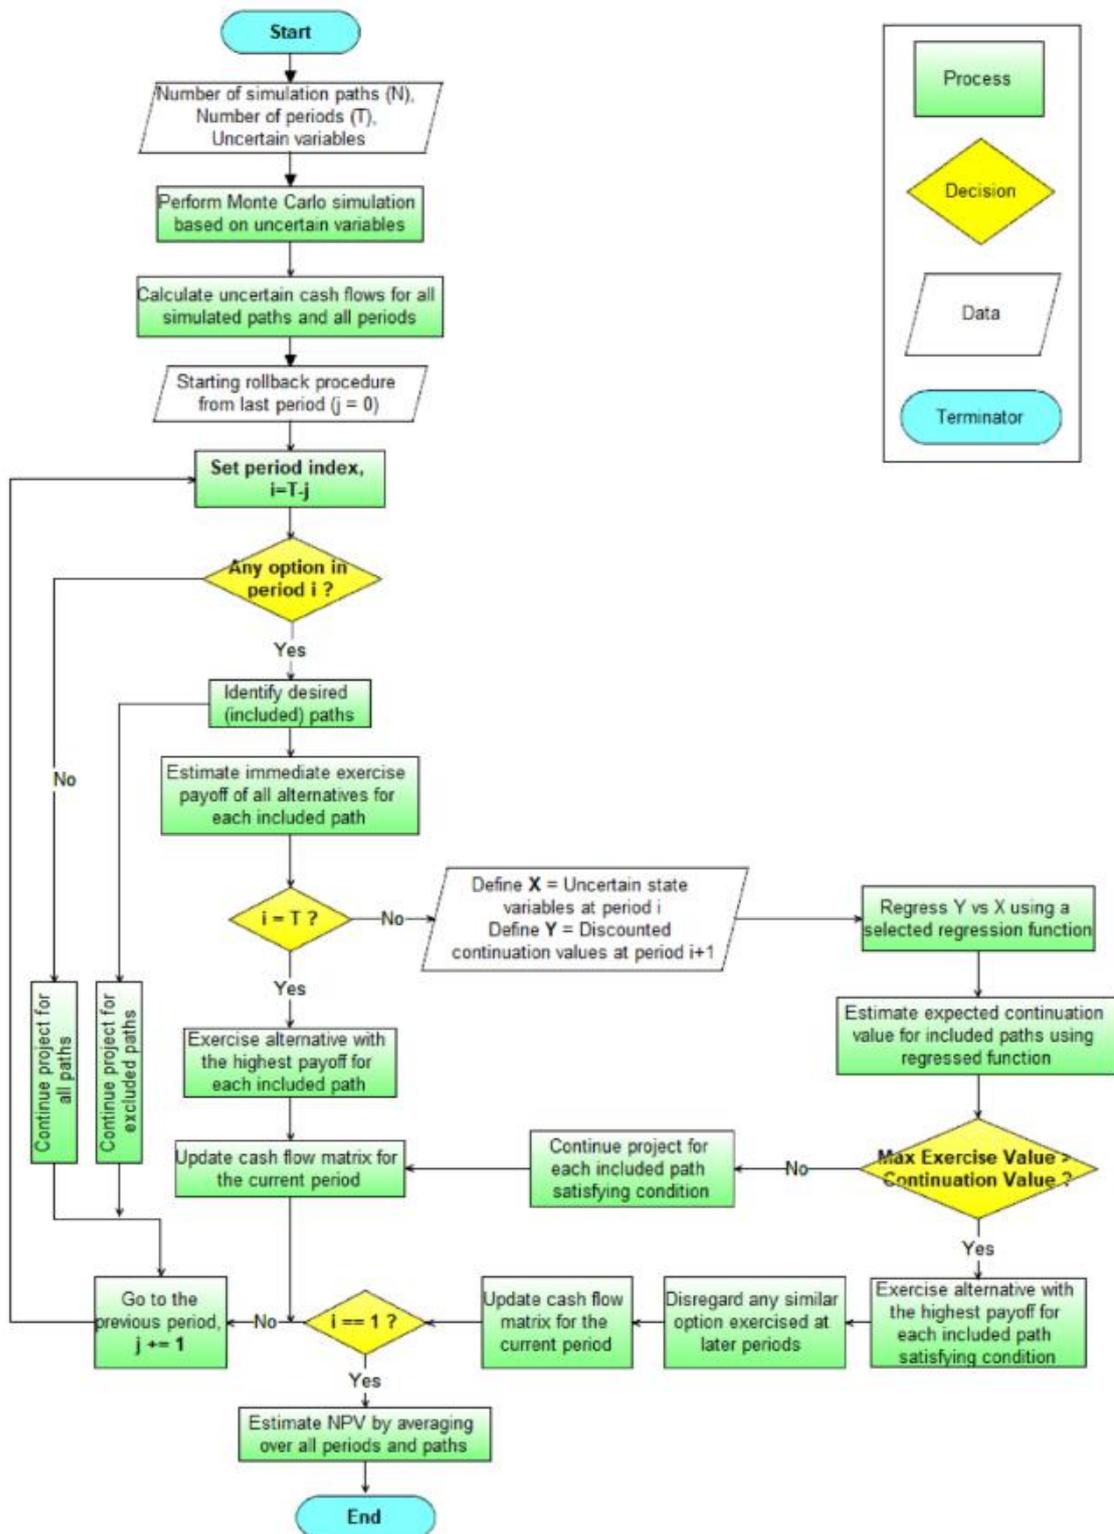


Figure 9. General flowchart for implementing real options values in the LSM algorithm. From Ahmadi and Bravtvoid (2023).

4.4. EXAMPLE

In this sub-section, a simplified version of the case study, that will be detailed in Chapter 8, is presented, showcasing a fully structured SDM problem, which can be solved using the LSM algorithm and aligned with the traditional decision tree roll-back procedure.

4.4.1. Problem setting and fully structured decision tree

This exercise was introduced by Hong et al. (2019) and exemplifies a fully structured representation for sequential decision-making, with only three decision points in time. It is an improved oil recovery (IOR) initiation time problem that defines an oil field with a 15-year production lifetime and decision alternatives of ‘switching’ or ‘not switching’ from primary to secondary recovery at years 0, 5 and 10.

The oil production rates are modelled as a function of time with uncertain parameters of three equiprobable geological realizations, R1, R2 and R3. The oil production rates, over time and under primary recovery, of the three realizations are illustrated in Figure 10, where t_i are the decision points in time. The likelihood functions of oil production rate measurements are summarized in Table 1. For example, if the truth is R1, the probability of the measurement at t_1 indicating “high rate” (i.e., a measured data of 18.1 million bbl/yr) is 0.75, and the probability of the measurement at t_2 indicating “high rate” (i.e., a measured data of 15.2 million bbl/yr) is 0.80.

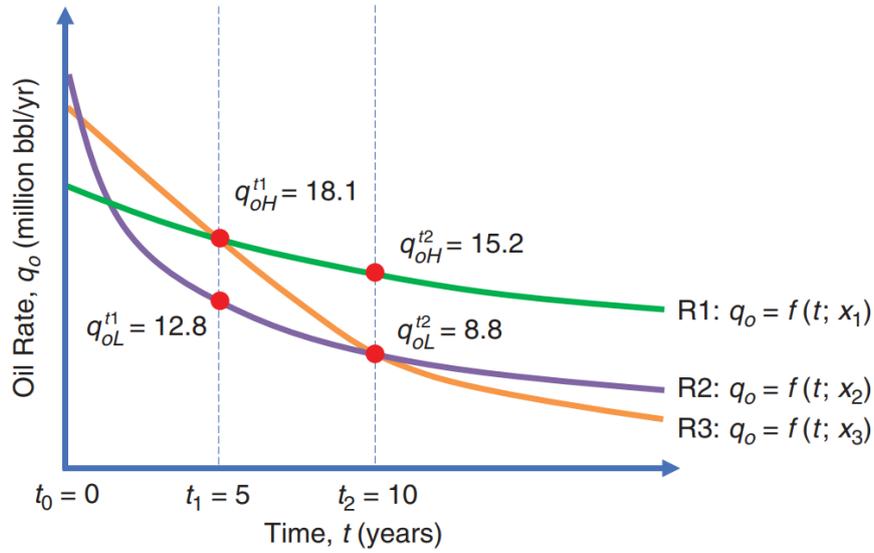


Figure 10. Oil production rate under primary recovery as a function of time for three realizations. From Hong et al. (2019).

		Given the Truth is		
		R1	R2	R3
Measurement at t_1 says	High rate ($q_{oH}^{t_1}$)	0.75	0.25	0.75
	Low rate ($q_{oL}^{t_1}$)	0.25	0.75	0.25
Measurement at t_2 says	High rate ($q_{oH}^{t_2}$)	0.80	0.20	0.20
	Low rate ($q_{oL}^{t_2}$)	0.20	0.80	0.80

Table 1. Likelihood functions for the measured oil production rates under primary recovery. From Hong et al. (2019).

Figure 11 shows the full decision tree for the problem setting, where the NPVs and utilities corresponding to the decision alternatives and geological realizations are listed at the right end of the tree. In this example, an exponential utility function with risk tolerance (Rho) of \$800 million (risk-averse) is used to convert NPVs to utilities.

The decision problem is solved by rolling back the decision tree, resulting in an EU of 0.6 and CE of \$677.19 million.

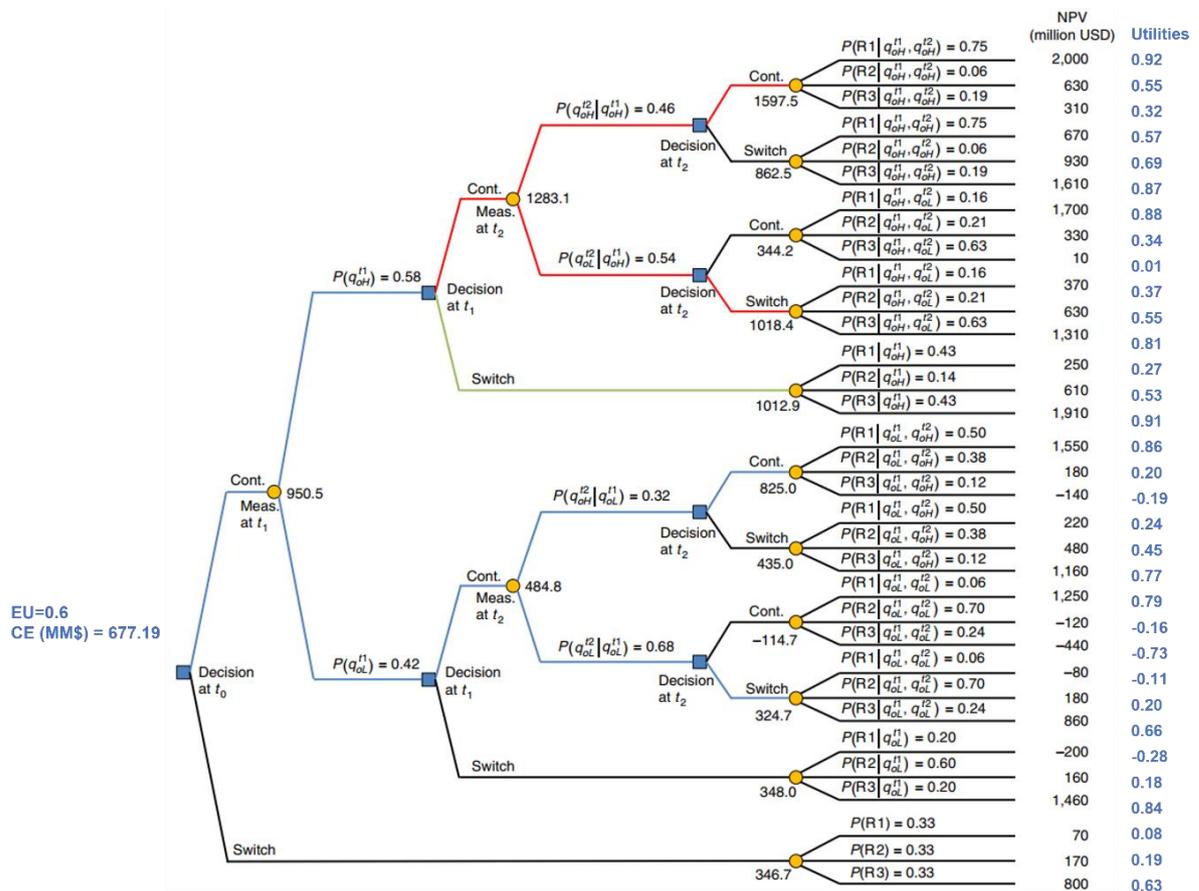


Figure 11. Fully structured decision tree for the simple IOR initiation time exercise for a risk-averse case, with Rho = \$800 million. Modified from Hong et al. (2019).

After manually solving the decision tree, the steps of using LSM to solve the same decision problem will be detailed, with a few paths for illustrative purposes. The first step is to generate six data paths or realizations of the observed oil rates, using MCS based on the prior probabilities (three equiprobable realizations – each realization’s probability is 1/3) and the likelihood probabilities (listed in Table 1). The data paths generated are shown in Table 2.

Path	Geological realization	Data at	
		$t_1 (q_o^{t1})$	$t_2 (q_o^{t2})$
1	1	18.1	15.2
2	3	18.1	15.2
3	3	18.1	15.2
4	1	12.8	15.2
5	3	12.8	8.8
6	2	12.8	8.8

Table 2. Paths of measured oil rates for simplified LSM example.

As rolling back decision trees, LSM starts at the last decision point, t_2 in this example, and then recursively solves the sequential decision problem backwards in time. Table 3 lists the utilities corresponding to the alternatives of “Continue” with primary recovery and “Switch” to secondary recovery, at time step t_2 , for each path. At decision point t_2 , additional information - oil rates at t_1 and t_2 - will have been measured and thus the conditional expected utilities (EU) at t_2 given additional information are conditional on the oil rates measured at both t_1 and t_2 (Hong et al., 2019). Hence, utilities are regressed on oil rates at t_1 and t_2 using a non-linear regression algorithm, XGB in this example (see Chapter 5 for background on non-linear regression methods and Chapter 7 for the selection of regression method for the present case study), to assess the conditional EUs at t_2 given additional information. The regression is done for both alternatives, “Continue” and “Switch,” respectively.

Table 4 lists the conditional EU at t_2 , assessed from regression, for each path and each alternative. The optimal decision at t_2 for each path is identified by choosing the highest EU for each path (underlined value in the table).

		Data at		Utilities	
Path	Geological realization	$t_1 (q_o^{t1})$	$t_2 (q_o^{t2})$	Continue	Switch
1	1	18.1	15.2	0.918	0.567
2	3	18.1	15.2	0.321	0.866
3	3	18.1	15.2	0.321	0.866
4	1	12.8	15.2	0.856	0.240
5	3	12.8	8.8	-0.733	0.659
6	2	12.8	8.8	-0.162	0.201

Table 3. Utilities for alternatives at time t_2 for simplified LSM example.

		Data at		Utilities		Conditional Expected Utilities	
Path	Geological realization	$t_1 (q_o^{t1})$	$t_2 (q_o^{t2})$	Continue	Switch	Continue	Switch
1	1	18.1	15.2	0.918	0.567	0.520	<u>0.767</u>
2	3	18.1	15.2	0.321	0.866	0.520	<u>0.767</u>
3	3	18.1	15.2	0.321	0.866	0.520	<u>0.767</u>
4	1	12.8	15.2	0.856	0.240	<u>0.856</u>	0.240
5	3	12.8	8.8	-0.733	0.659	-0.448	<u>0.430</u>
6	2	12.8	8.8	-0.162	0.201	-0.448	<u>0.430</u>

Table 4. Expected utilities for alternatives at time t_2 for simplified LSM example.

After identifying the decision policy for t_2 , LSM progresses one time step backwards, to t_1 , and the utilities corresponding to the two alternatives “Continue” and “Switch” at t_1 are listed in Table 5. Note that if the decision at t_1 is “Switch,” there will be no more decision at t_2 ; if the decision at t_1 is “Continue,” the decision at t_2 will be made based on the optimal decision policy that has been identified at the previous step – taking Path 1 as an example, its optimal decision at t_2 is “Switch” (expected utility = 0.767 in Table 4) and thus its “Continue” utility at t_1 is its “Switch” utility at t_2 (0.567). At decision point t_1 , only oil rates at t_1 will have been measured as additional information; therefore, the utilities are regressed on t_1 measurements for the assessment of conditional EU given additional information at t_1 . The assessed EUs at t_1 for both

alternatives and each path are listed in Table 6. Once again, the best decision at t_1 is the alternative with the highest EU for each path.

		Data at	Utilities	
Path	Geological realization	$t_1 (q_o^{t1})$	Continue	Switch
1	1	18.1	0.567	0.268
2	3	18.1	0.866	0.908
3	3	18.1	0.866	0.908
4	1	12.8	0.856	-0.284
5	3	12.8	0.659	0.839
6	2	12.8	0.201	0.181

Table 5. Utilities for alternatives at time t_1 for simplified LSM example.

		Data at	Utilities		Conditional Expected Utilities	
Path	Geological realization	$t_1 (q_o^{t1})$	Continue	Switch	Continue	Switch
1	1	18.1	0.567	0.268	<u>0.767</u>	0.695
2	3	18.1	0.866	0.908	<u>0.767</u>	0.695
3	3	18.1	0.866	0.908	<u>0.767</u>	0.695
4	1	12.8	0.856	-0.284	<u>0.572</u>	0.245
5	3	12.8	0.659	0.839	<u>0.572</u>	0.245
6	2	12.8	0.201	0.181	<u>0.572</u>	0.245

Table 6. Expected utilities for alternatives at time t_1 for simplified LSM example.

Finally, LSM moves to t_0 . Since there is no additional information available at t_0 (the information that is available at t_0 is the prior information that has been embedded in the prior probability distribution), the EU for each alternative at t_0 is calculated as the mean of the utilities of each path, listed in Table 7, which results in an EU for “Continue” of 0.669, which corresponds to a CE of \$885.32 million, and an EU for “Switch” of 0.376, which corresponds to a CE of \$377.15 million. The optimal decision policy is summarized in Table 8 for each path.

		Utilities	
Path	Geological realization	Continue	Switch
1	1	0.567	0.084
2	3	0.866	0.632
3	3	0.866	0.632
4	1	0.856	0.084
5	3	0.659	0.632
6	2	0.201	0.191

Table 7. Utilities for alternatives at time t_0 for simplified LSM example.

		Data at		Optimal decision at		
Path	Geological realization	$t_1 (q_o^{t_1})$	$t_2 (q_o^{t_2})$	t_0	t_1	t_2
1	1	18.1	15.2	Continue	Continue	Switch
2	3	18.1	15.2	Continue	Continue	Switch
3	3	18.1	15.2	Continue	Continue	Switch
4	1	12.8	15.2	Continue	Continue	Continue
5	3	12.8	8.8	Continue	Continue	Switch
6	2	12.8	8.8	Continue	Continue	Switch

Table 8. Table representation of optimal decision policy for simple IOR problem solved using LSM.

4.4.2. Validation of VOI assessment using LSM

In this section, the VOI assessed using LSM algorithm will be validated by comparing against the VOI obtained analytically using the decision tree.

The first step was to manually resolve a risk-neutral decision tree using the roll-back procedure, and compare the VOI results obtained by Hong et al. (2019). The VOI is calculated as the difference between the expected value with information (EVWI) and the expected value without information (EVWOI). The tree branch with information (Figure 11) is solved under neutrality, delivering the EVWI. For the EVWOI, a separate branch of the tree is built with the decisions at t_0 , t_1 and t_2 to the

left, because the decisions are taken before the information from the uncertainties are known. The probabilities assigned to the uncertainties for each geological realization are the prior probabilities. The tree is rolled-back to obtain the EVWOI. Then the VOI is calculated as the difference between EVWI and EVWOI, and documented in Table 9.

The VOI was also calculated through the LSM algorithm for validation. OLS and XGB regressions were tested for the same problem setting, using 100,000 realizations that respected the prior and likelihood probabilities. Table 9 shows the VOIs obtained by Hong et al. (2019), the replicated decision tree roll-back procedure and the LSM algorithm. The replicated roll-back approach has delivered the same result from Hong et al. (2019), \$214.7 million, validating the replication of the example. The LSM algorithm using OLS and XGB has delivered the same result, \$214.35 million, being very similar to the estimation by the decision tree. The small difference between the decision tree (i.e., analytical) VOI and the LSM VOI (214.7 vs. 214.4) is attributed to the sampling error inherent in MCS. In this case, the VOI is not different/improved by XGB, compared to OLS, because this simplified problem is discretized into four possible measured values of oil rates (12.8 and 18.1 at t_1 and 8.8 and 15.2 at t_2) and the regression surface of expected NPV (ENPV) (t_1 rate, t_2 rate) is defined by 4 points; thus, the impact of using OLS (linear regression) and XGB (non-linear regression) is not very different. This comparison validates that XGB performs as well as OLS for linear regression, which is as expected.

Estimated value	Hong et al. (2019)	Replicated decision tree	LSM (OLS)	LSM (XGB)
VOI (MM\$)	214.7	214.7	214.4	214.4

Table 9. Comparison of VOI estimated by Hong et al. (2019), replicated roll-back decision tree and LSM workflows using OLS and XGB.

Once the working decision tree was validated, the utility function was introduced in both, the decision tree and LSM workflow, to model a risk-averse attitude. For the LSM workflow, the two types of regressions, OLS and XGB, were tested. Table 10 summarizes the VOI results. It can be observed that the analytical and the LSM algorithms have delivered very similar results, \$93.39 million estimated rolling back the tree, and \$93.16 million using LSM with XGB and OLS regression. The reason why the OLS and XGB have delivered the same results is the discretization of the four possible measured values of oil rates in this simplified problem, as explained before. The similarity of the VOI obtained by rolling back the tree versus the LSM algorithm provides confidence in the use of LSM and non-linear regression for calculation of VOI, procedure that will be extensively used in the thesis case study going forward.

Estimated value	Decision tree	LSM (OLS)	LSM (XGB)
VOI (MM\$)	93.39	93.16	93.16

Table 10. Comparison of VOI estimated by rolling back the decision tree and LSM workflows using XGB.

Chapter 5: Machine learning non-linear regression methods

It has been mentioned that linear functions are very commonly used to approximate the relation between predictors and dependant variables. If the relation between the independent and dependant variables is not linear, linear regression might not be the optimal fitting approach, hence, different methods would be required. In the present case study, the introduction of risk attitude involves the assessment of expected utilities for maximisation during the LSM algorithm. Since risk-averse and risk-seeking utility curves are non-linear functions of monetary values, the resulting relation between the utilities and the observations (oil rates) results non-linear. Using a linear regression approach for these cases may result in sub-optimal decision policies.

This chapter introduces several non-linear regression methods based in machine learning. The theory of the machine learning regression methods is outside the scope of the present thesis; therefore, the methods will be introduced very briefly. These methods are later tested for approximating the expected utilities in the present case study.

5.1. K NEAREST NEIGHBOURS (KNN)

KNN is a machine learning classifier based on identifying the nearest neighbours of a sample and using those neighbours to determine the class of the sample. k number of the nearest neighbours are used in determining the class. When multiple classes are present in the neighbour's sample, simple majority voting or distance weighted voting is used for the resulting classification (Cunningham and Delany, 2021).

KNN method is extremely sensitive to the user-defined parameter k . Small k can lead to overfitting the model. A large k might result in the regression model not capturing the actual behaviour of the dataset (Shahali, 2022). One way to determine the

optimal k is performing K -fold cross-validation. See Chapter 7: Selection of regression method, Section 7.1: K -fold cross-validation.

5.2. RANDOM FOREST (RF)

Random forest is a supervised learning algorithm based on a growing-tree approach. Yang et al. (2016) describes regression tree as a decision tree-based prediction model, where the decision tree is built using the samples of the dataset: Starting from the entire set of samples, a regression tree selects one independent input variable among all and performs binary split into two child sets, under the condition that the two child nodes give increased purity of the data compared with its single parent node. Purity is defined as the deviation of predicting with the mean value of the output variable. The split is performed iteratively until a terminating criterion is satisfied. The nodes that are not further partitioned are called leaves. After growing a large tree, a pruning process is employed to remove the leaves contributing insignificantly to the purity improvement (Breiman et al., 1984).

Random forest grows an ensemble of trees governed by random vectors, and letting them vote for the most popular classification. Breiman (2001) explains that for a k^{th} tree, a random vector Θ_k is generated, independent of the past random vectors $\Theta_1 \dots \Theta_{k-1}$ but with the same distribution. A tree is grown using the training set and k , resulting in a classifier $h(x, k)$ where x is an input vector. After many trees are generated, each tree casts a unit vote for the most popular class.

5.3. EXTREME GRADIENT BOOST (XGB)

XGB is a machine learning method based on regression trees. The algorithm is based on the idea of “boosting”, which combines all the predictions of a set of “weak” learners to develop a “strong” learner, through assistive training strategies. The overall

goal is to minimize over and underfitting, but also reducing computational time (Mokhtar et al., 2021).

To avoid overfitting problems, XGB uses an analytic expression that assess the “goodness” of the model from the original prediction (Mokhtar et al., 2021). This involves a loss function and the number of observations. A regularization parameter indicates the minimum loss needed to further partition a leaf node.

5.4. SUPPORT VECTOR REGRESSION (SVR)

The fundamental working guideline of SVR is to perform the data mapping in some other dot product spaces through non-linear mapping and perform the linear algorithm in the feature space (Shamshirband et al., 2015). SVR tries to find the optimal regression hyperplane so that most training samples lie within an deviation margin ϵ around this hyperplane. (Chen et al., 2018)

The problem can be described as convex optimization. The aim of SVR is to find a function of the input vector that has, at most, ϵ deviation from the total training sample, and at the same time as flat as possible by minimizing the norm of a model parameter (Chen et al., 2018). SVR then minimises two terms in the objective function, the first one is the insensitive loss function, i.e., only sample training error greater than a user-specific threshold is considered in the loss function. The other term is model complexity, which is expressed as the sum of squared regression coefficients (Yang et al., 2016).

5.5. PIECEWISE REGRESSION

The piecewise regression method partitions one input variable into multiple mutually exclusive segments/slopes and fit one local regression function for each segment/slope. The regression function for each segment can be linear or non-linear.

The split into segments (and optimisation of the breakpoints) and the computation of the regression coefficients are performed simultaneously, with an objective of minimising the total training error (Yang et al., 2016).

For any sample, its training error is equal to the absolute deviation between the real output and the predicted output for the segment. A recent approach starts with solving a linear regression (no segmentation) on the complete set of data with least absolute deviation. Then each input variable works as partition feature (a priori a single input variable) and the model is solved for 2 segments. The feature with minimum training error is kept and if its error represents a percentage reduction of more than a user-defined tolerance β from the global linear regression without data partition, the procedure continues with iteratively increasing number of segments until the β training reduction criterion is not satisfied. In this case it is decided that the current split provides no improved fit than previous iteration fit (Yang et al., 2016).

Chapter 6: General workflow of the present study

The general methodology used in the present study is illustrated in the Figure 12. The first step in this study was to replicate the problem setting in a risk-neutral case and LSM workflow with OLS regression used by Hong et al. (2019) and Wui (2019) – the IOR initiation and termination time problem. The replication was achieved by coding in Python. The replicated results were validated against Hong et al.’s and Wui’s results. This is documented in Chapter 8: Case study, Section 8.2: LSM implementation. The codes for the replication were the base for the further work that extended the decision problem to non-risk-neutral cases.

The next phase was to start the non-linear regression (NLR) testing using different machine learning (ML) methods in a relevant dataset to the case study. The dataset for the testing was generated using the replicated code, from which a subset of oil rates and their respective NPVs for different years of production were extracted. A utility function with a preliminary risk tolerance was used to convert the monetary values to utilities; the preliminary risk tolerance was calculated as one third of the expected NPV achieved by the replicated code under risk-neutrality. The oil rates and respective utilities were used to test single-dimensional and multi-dimensional regressions, using different NLR algorithms.

For each NLR regression algorithm, the optimal parameters were determined. To choose the most suitable method for the test data, several quality control measurements were performed, including assessment of the fitting scores, 2D and 3D visual inspections of the fit, residual scatter plots and QQ-plots (See Chapter 7: Selection of regression method).

After choosing the optimal regression method for the test data, the LSM algorithm was implemented using NLR as a fitting method for the expected values and/or utilities. Having seen that the chosen NLR method provided a better fit even in

a linear relation (see Chapter 7: Selection of regression method), a first validation was performed using risk-neutral and using the chosen NLR for the fitting. These results were compared against the original OLS approach, to validate the accuracy of NLR in the LSM algorithm. Following this, risk attitudes were introduced into the case study problem and modelled using exponential utility functions. This is summarized in Chapter 8: Case study, Section 8.2: LSM implementation.

The value of information (VOI) obtained through the LSM algorithm was also validated using a conventional decision tree approach that resembled a simplified version of the current case study. This validation was documented in Chapter 4: LSM algorithm, Section 4.4: Example, Sub-section 4.4.2: Validation of VOI. Once the VOI method was validated, sensitivity analysis of different risk tolerances was conducted. The impacts of the risk attitude and the information on the decision policy were assessed, and the outcomes of the case study were discussed. See Chapter 8, for detailed documentation on all the mentioned assessments, and Chapter 9 for the discussion.

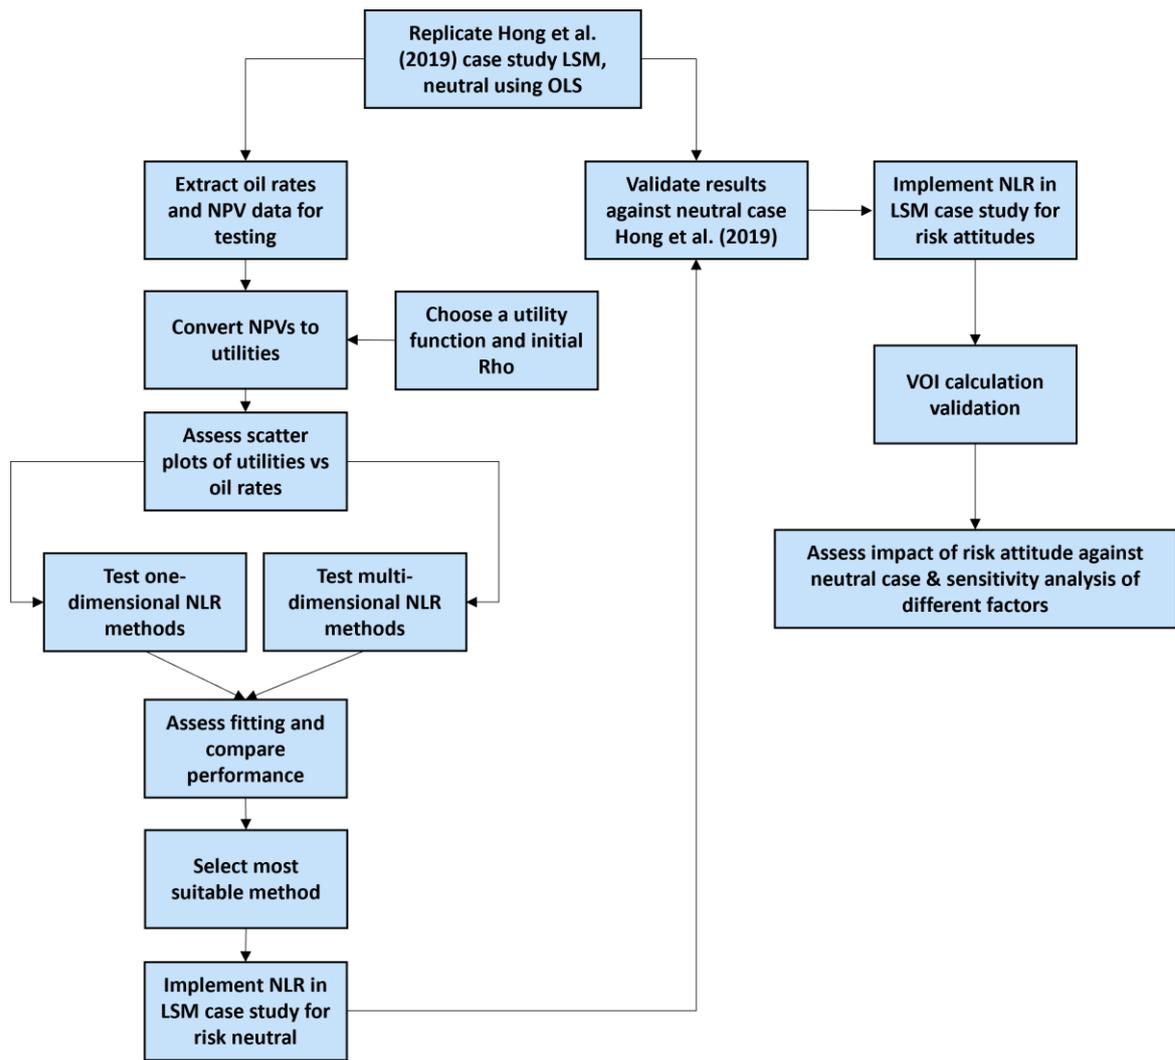


Figure 12. General workflow of the present work.

Chapter 7: Selection of the non-linear regression method

In this chapter a review of the non-linear regression (NLR) testing is performed, including documentation of different quality control displays and measurements, and the final selection of a regression method to be used in the LSM algorithm for resolving the case study.

The data used for testing was generated with the Python workflow created for the case study, where oil rates have been simulated over 50 years of production, using Monte Carlo simulation, and respective NPVs have been calculated, as described in Chapter 8. Therefore, this extraction of data is representative of the problem in question.

The X axis for regressions consisted of sets of 10,000 samples of oil rates for each year of production from year 1 to year 30, which would enable multi-dimensional regression tests. The Y for regressions consisted of 10,000 samples of the corresponding NPVs at year 30. To obtain the NPVs at year 30, the LSM workflow was solved starting at year 50, and ‘rolled-back’ down to year 30, when the NPVs were assessed for the different alternatives, and the optimal NPVs were exported for testing, along the oil rates from year 1 to year 30 (see details about the LSM workflow in Chapter 8: Case study, Section 8.2: LSM implementation). The NPVs were then converted to utilities, using the exponential utility function with a risk tolerance of \$200 million. Tests of single dimensional and multi-dimensional regression were performed.

7.1. K-FOLD CROSS-VALIDATION

Machine learning methods are controlled by ‘hyperparameters’ to avoid under- and overfitting. In underfitting, the model is unable to capture the relation between the independent variables and the target variable. In overfitting the model can fit the training data very well but cannot predict the validation/test data. In order to find the most suitable hyperparameters a k-fold cross-validation method was used. In k-fold

cross-validation, the dataset is randomly split into k groups (fold) and treat each group as a validation sample or test dataset. Since the partition is random, the variance of the accuracy estimates can be large for statistical inference (Wong and Yeh, 2020). Several studies have suggested to perform k-fold cross validation repeatedly to obtain reliable accuracy estimates for statistical comparison (Wang et al., 2017; Witten et al., 2011).

The fitting scores from k-fold cross-validation quantify overfitting or underfitting. In non-linear regression algorithms the fitting scores are usually the squares of the coefficient of multiple correlation or coefficient of determination, denoted as R^2 . In this work, Python module *scikit-learn* is used for the regression algorithms and the determination of their respective R^2 scores. In *scikit-learn* R^2 quantifies the proportion of variance of the real data that has been explained by the independent variables in the model. It provides an indication of goodness of fit and therefore a measure of how well unseen samples are likely to be predicted by the model, through the proportion of explained variance (scikit-learn.org). In scikit-learn R^2 is defined as follows:

$$R^2(y, \hat{y}) = 1 - \frac{\sum_{i=1}^n (y_i - \hat{y}_i)^2}{\sum_{i=1}^n (y_i - \bar{y})^2} \quad (10)$$

where \hat{y}_i is the predicted value of the i-th sample from regression, y_i is the corresponding true value for total n samples and

$$\bar{y} = \frac{1}{n} \sum_{i=1}^n y_i$$

Best possible score is 1.0 and it can be negative because cross-validation is normally performed on random permutations on of train and test data, therefore the model can be arbitrarily worse (scikit-learn.org).

For each NLR algorithm a selection of the most sensitive hyperparameters was made for the cross-validation. First, a 3-fold cross-validation was performed, using the scikit-learn module “*GridSearchCV*”. *GridSearchCV* implements a “fit” and a “score” method. The hyperparameters of the estimator are optimized by cross-validation over combinations within a dictionary of parameters (grid). This enables searching over any given sequence of parameter settings (scikit-learn.org).

Table 11 shows a summary of the hyperparameters selected for each NLR method, and the results suggested by the 3-fold validation with *GridSearchCV*.

Algorithm	Hyperparameter	Single-dimensional regression	Multi-dimensional regression
K nearest neighbour	Lead size	1	1
	Number of neighbours	27	27
	p	1	2
Random forest	Number of estimators	70	100
	Maximum features	0.7	0.2
	Maximum leaf nodes	12	54
Extreme gradient boost	Maximum depth	3	3
	Eta	0.1325	0.1325
	Subsample	1	0.875
Support vector regression	Kernel	poly	rbf
	C	63	15.87
	Gamma	0.223	0.001
Piecewise regression	Number of bins	2	10
	Estimator	RF	XGB

Table 11. Optimized hyperparameters suggested by 3-Fold K validation using *GridSearchCV* on multi-dimensional regressions of utilities, for a risk-averse case.

After the joint optimization of parameters, a confirmation was performed through visual inspection of the R^2 plots, obtained by varying each parameter individually and performing 10-k fold validation. This assessment can be seen as sensitivity analysis on each parameter to assess the effect of the fit on train and test

data. For each parameter assessed, the other parameters were kept fixed. Graphs were generated for each parameter, plotting R^2 against the tested parameter. Figures below show some examples of this sensitivity analysis, for the parameters that generated more variation in KNN and XGB R^2 scores (see Appendix A for the rest of the NLR algorithms and parameters).

Figure 13 shows the sensitivity of parameter ‘number of neighbours’ on the R^2 for train and test data, when using KNN regression. It is observed that the maximum R^2 for the test data is reached at about 30-60 neighbours, however on the train data the optimal number of neighbours would be around 30, where the curve becomes horizontally flat and shortly after starts to linearly decrease. This plot confirms the suggestions from *GridSearchCV*, where around 30 is the optimal number of neighbours.

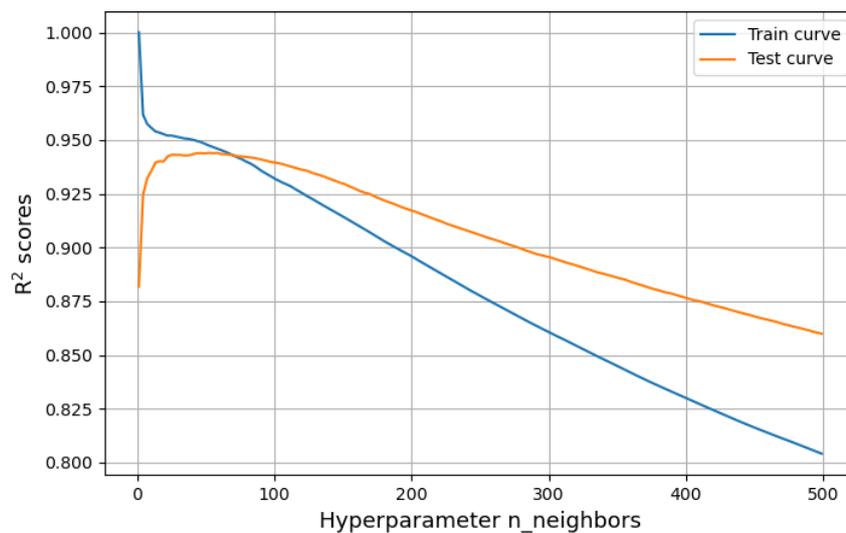


Figure 13. Sensitivity analysis of hyperparameter ‘number of neighbours’ from KNN algorithm on the fitting scores.

Figure 14 shows the sensitivity of parameter ‘maximum depth’ from XGB algorithm on the R^2 scores obtained when using XGB regression. It is noted that both the train and the test data reach R^2 stability at around 2 maximum depth. However, the train curve suggests a higher value of the parameter. This result is also in agreement

with *GridSearchCV*'s suggestion of a maximum depth of 3, where both train and test data have high R^2 scores.

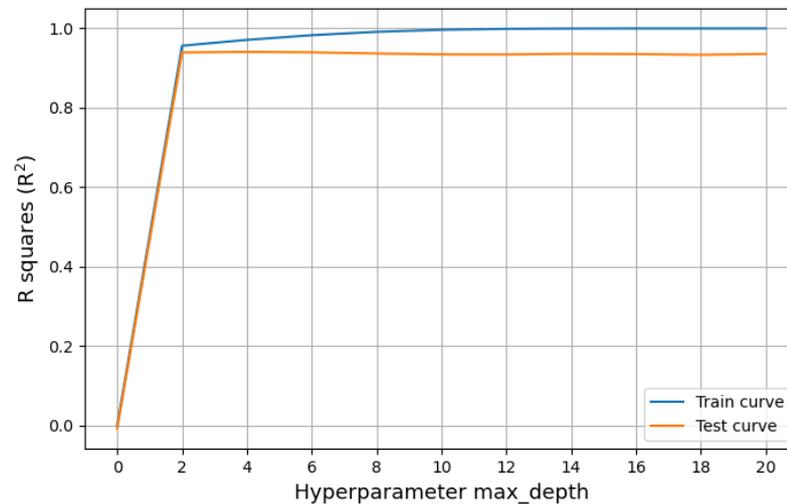


Figure 14. Sensitivity analysis of hyperparameter ‘maximum depth’ from XGB algorithm on the fitting scores.

7.2. SCATTER AND Q-Q PLOTS OF REGRESSION RESIDUALS

One of the quality control (QC) measurements in a regression procedure is to compute the residuals (or errors) ε , between the measured dependant value y_i and the predicted value \hat{y}_i , $\varepsilon = \hat{y}_i - y_i$, and plot them against the independent variables. Inference for the general linear model makes several assumptions, including independence of errors, normality, and homogeneity of variance (homoscedasticity) (Schützenmeister et al., 2012), which would imply that the residuals are equally spread around $\varepsilon = 0$. A scatterplot of these errors provides a visual inspection to assess this condition. For a suitable non-linear regression fitting, minimum bias should still be aimed, and the error is expected to be randomly distributed.

Figure 15 shows the single-dimensional OLS regression of NPVs for a risk-neutral case, and typical 2D QC plots, (a) the regression fit scatter plot and (b) the

residuals scatterplot. Figure 15a shows that the relation between the independent variables (oil rates) and the dependent variables (NPVs) is mostly linear, hence, a straight regression line (blue line) is sensibly representative of the sample trend. There are some inaccuracies below 2 MM bbl, where the cloud formed by the red dots has a slightly different slope. To explain this behaviour, it is worth to make the following observation, reminding that the data used is an output from the LSM workflow execution down to year 30: A 'linear case' means the objective function is a linear function of observables for given decisions and 'a non-linear case' means the objective function is a non-linear function of observables for given decisions. Even for a linear case, the relation between objective values and observed values might not be linear when a decision strategy is involved because the relation between objective values and the decision strategy might not be linear. Such non-linearity is represented by the slight change of slope below 2 MM bbl.

The residuals scatter plot show that the errors are generally unbiased, having an overall random distribution, and centred around the zero-error line. A local average of the residuals was calculated for facilitating the visual analysis: the residuals were first sorted by increasing oil rates; an average of the residuals was calculated for each consecutive group of 100 samples of oil rates and the resulting residual value was plotted at the respective 'y' of the graph, and at the 'x' corresponding to the median value of the 100 oil rates; this was repeated for each consecutive group, calculating average and standard deviation. The local average, also called windowed average, is displayed in green colour in Figure 15b, and the local standard deviation is shown pink; these lines show that between 2 and 3.5 MM bbl the error is closest to zero, indicating unbiasedness. As mentioned previously, there are slightly larger errors towards the lower end of the oil rates where the regression line does not effectively capture the tale of the oil samples. Wider scattering of the predicted values is observed at the highest oil rates in Figure 15a, and consequently larger errors as shown in Figure 15b. The local

standard deviation is quite homogeneous over the oil rates. The R^2 score for this fit resulted in 0.654.

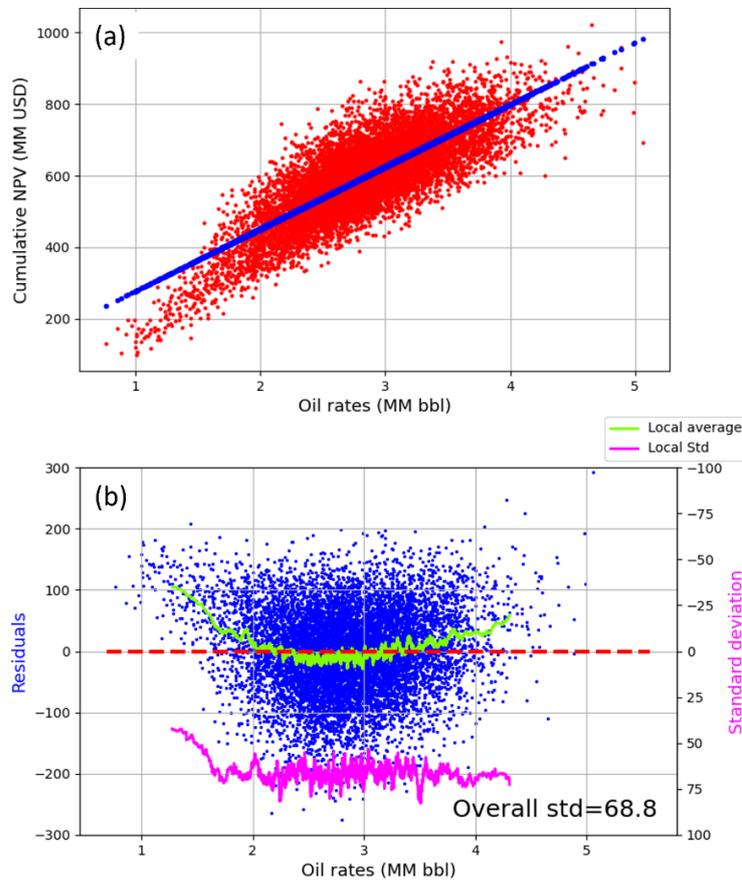


Figure 15. Scatterplots of a single-dimensional OLS regression of NPVs in a risk-neutral case. (a) Scatterplots of NPVs vs oil rates (red) and predicted NPVs (blue). (b) Residuals scatterplot (blue), with windowed average (green) and standard deviation (pink).

Figure 16 shows equivalent displays for a single-dimensional OLS regression of utilities for a risk-averse case. In this case the relation between the oil rates and the utilities is not linear, evidenced by the exponential shape of the samples (the red dots in Figure 16a). If a linear regression is used, the errors are not homogeneously distributed, as seen in Figure 16b. There are two indications of not appropriate fit: first, the local average, represented by the green line is not close to zero for most of the oil

rates range, with much larger residual values at the lower end; second, the local standard deviation, represented by the pink line, is not homogeneous over the oil rates range, showing bias. The R^2 scores for this fit resulted in 0.509, which is significantly lower than the score obtained in the risk-neutral case. This supports that a linear fit would not be suitable for a risk-averse case, as expected.

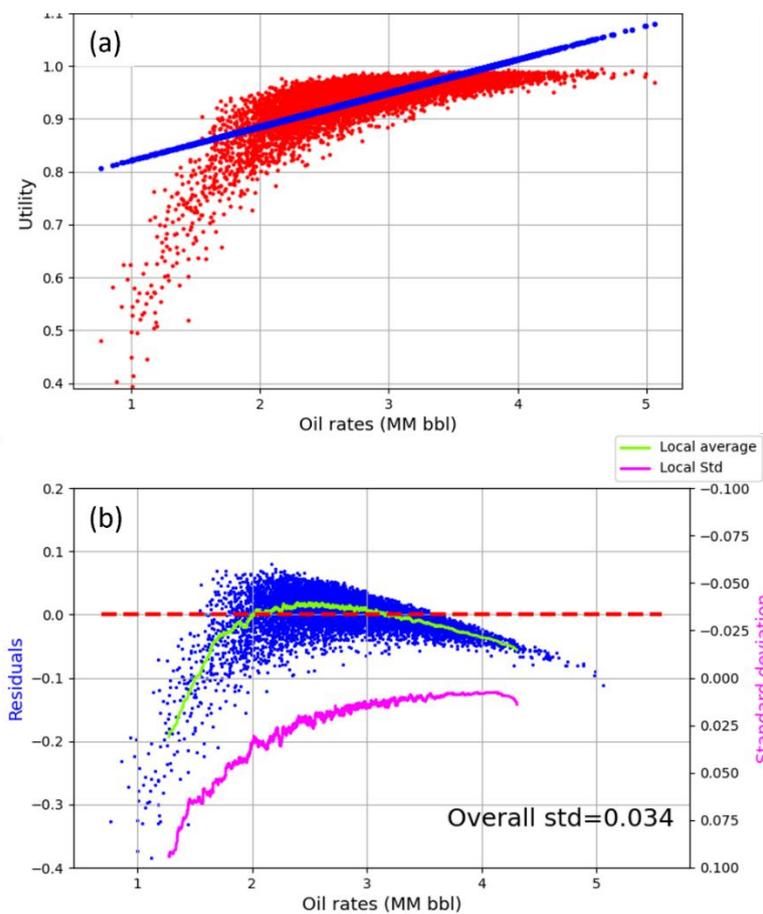


Figure 16. Scatterplots of a single-dimensional OLS regression of utilities in a risk-averse case. (a) Scatterplots of utilities vs oil rates (red) and predicted utilities (blue). (b) Residuals scatterplot (blue), with windowed average (green) and standard deviation (pink).

Another useful way to visually verify the goodness of the fit is by drawing the regression residuals Q-Q plots (Wilk and Gnanadesikan, 1968), under the expectation that regression errors should be normally distributed to respect homoscedasticity. Q-Q

plots are very suitable to evaluate whether a sample is normally distributed. Consider a sample with size n , given by x_1, \dots, x_n , and F_n is the empirical distribution function based on the sample, a Q-Q plot is built by plotting the theoretical quantiles, known as the standard normal variate (normal distribution with mean 0 and standard deviation 1), $F^{-1}(F_n(x_i))$, against the samples quantiles, $x(i)$. If the empirical distribution is consistent with the theoretical distribution, F , the points in the Q-Q plot fall on the line of identity. This applies for distributions within a common location and scale family. Therefore, Q-Q plots can be used for verifying the distribution of a sample set for normality (Loy et al., 2016).

The Q-Q plots presented in this work are generated via Python module *statsmodels*, with a 'standardized' line, where the resulting points in the Q-Q plot are scaled by the standard deviation of the given sample set and the mean is also added (statsmodels.org). Figure 17 shows the residuals Q-Q plots for OLS performed on the neutral and risk-averse cases presented previously in Figure 15 and Figure 16, respectively. For the neutral case in Figure 17a,c, it can be observed that the points of the Q-Q plot overall fall on the standardized line. However, when using OLS fitting in a risk-averse case Figure 17b,d, it is evident that the ends of the curve are falling outside the straight line, indicating that the fit is inadequate especially in the lower end oil rates, as concluded previously.

Going forward with the non-linear regression testing in this work, residual scatter and Q-Q plots were both considered to support the determination of the best fit. This was also accompanied by evaluation of the R^2 scores, and visual inspection of the data fit in 2D and 3D displays.

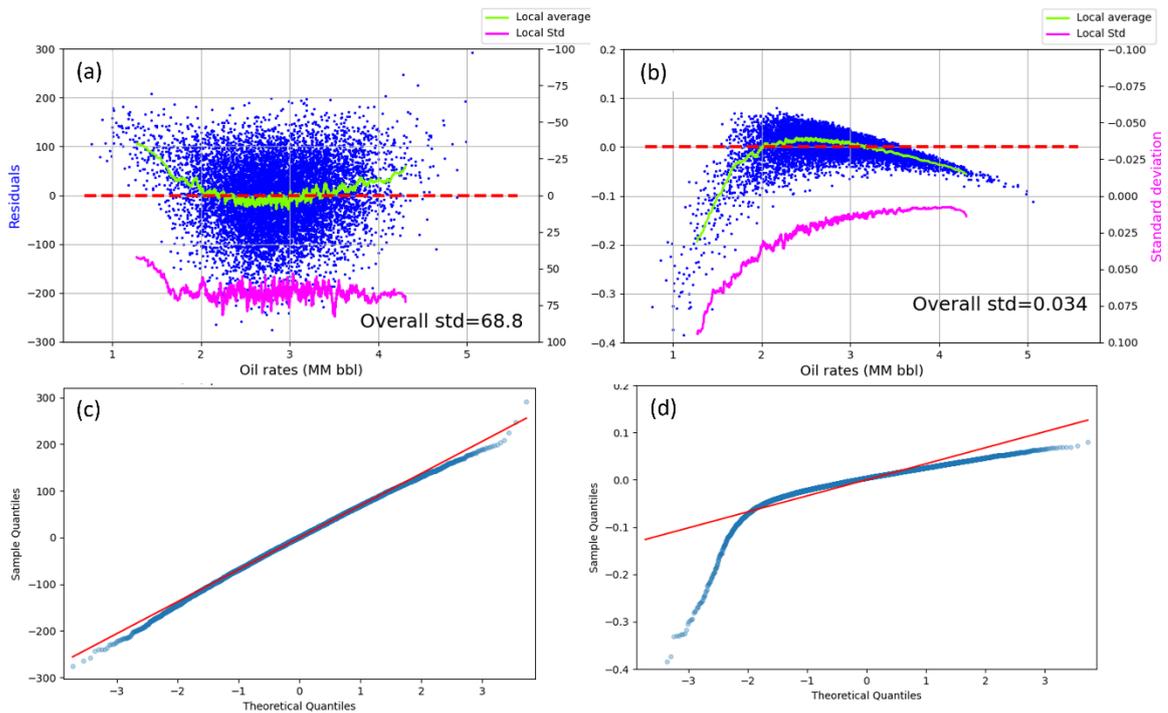


Figure 17. Scatter and Q-Q plots of a single-dimensional OLS. Regression residuals scatterplot for (a) risk-neutral and (b) risk-averse. Residuals Q-Q plot (blue) for (c) risk-neutral and (d) risk-averse; red line is a standardized line.

Figure 18 and Figure 19 show the QC plots for RF and XGB non-linear regressions of utilities on a risk-averse case, before and after the parameters are optimized. Parameter optimization refers to the results of the 3-Fold *GridSearchCV* validation (Table 11), subsequently confirmed by 10-Fold R^2 sensitivity analysis. In a single-dimensional regression, it is expected that the predicted values form a line rather than a cloud of points, where the latter would indicate overfitting. It is also expected that after the parameters are optimized the residuals have larger magnitudes and have an unbiased and larger standard deviation (in case of overfitting the residuals would be closer to the zero line). For both regressors, overfitting is noticed before the parameters are optimized. This is evidenced by the cloud of predicted points instead of a line (Fig. 16a, c and e). After the parameters are optimized, the predicted utilities draw a line, as expected (b,d,f). This is true for the two regressors presented in this example. Another

observation is that after parameters are optimized, the standard deviation of the errors increases, from 0.01 to 0.022 for RF, and from 0.018 to 0.021 for XGB, which indicates an improvement from the overfitted result. Q-Q plots show that the fitting is challenging at the edges of the oil rates, especially the lower end. RF has a slightly worse fit in this case.

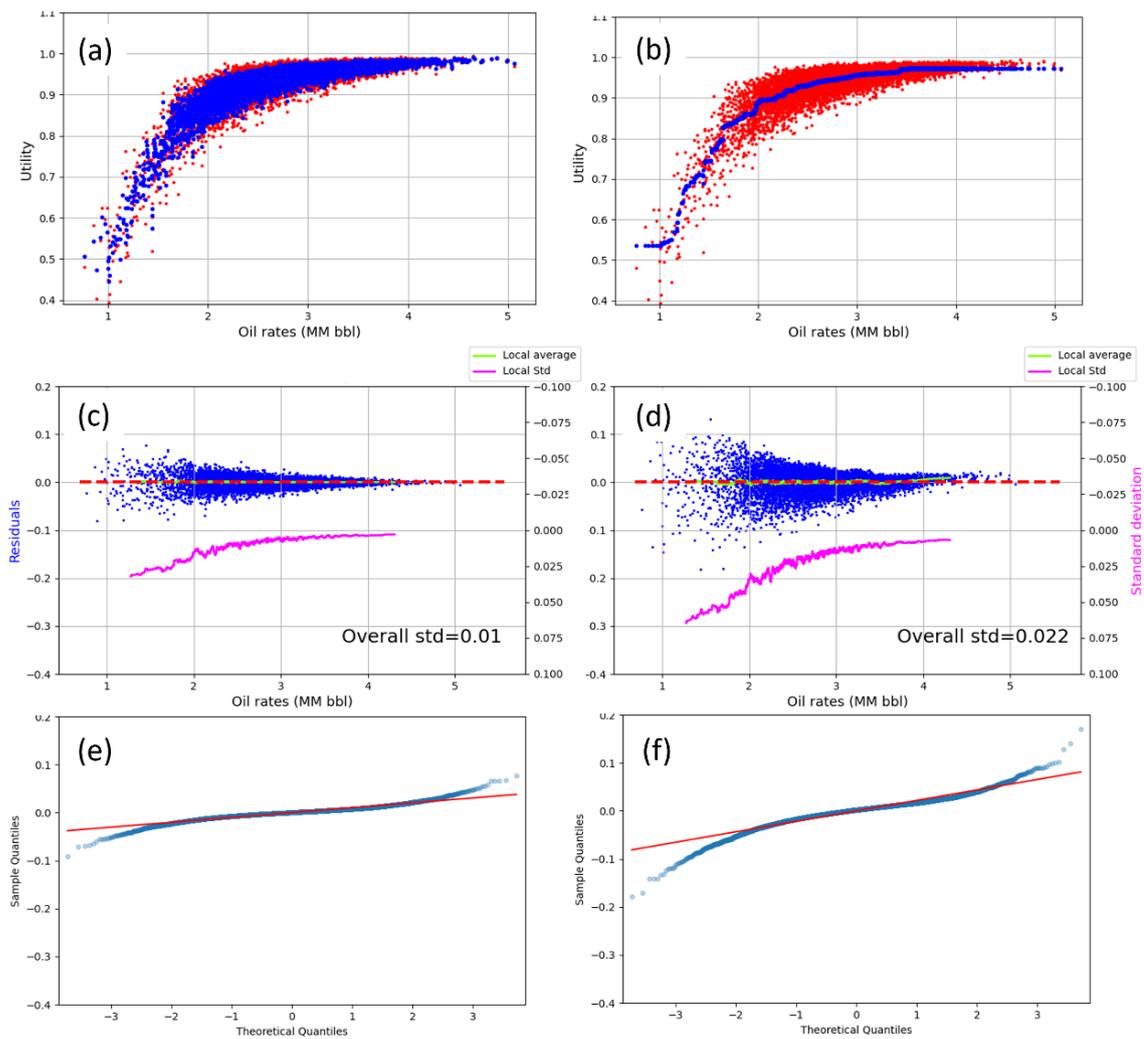


Figure 18. Single-dimensional Random Forest regression QC plots for risk-averse utilities: (a, c, e) before and (b, d, f) after parameter optimization. (a,b): Utilities (red) and predicted utilities (blue). (c,d): Residuals scatterplots. (e,f): Residuals Q-Q plots (blue) and standardized line (red).

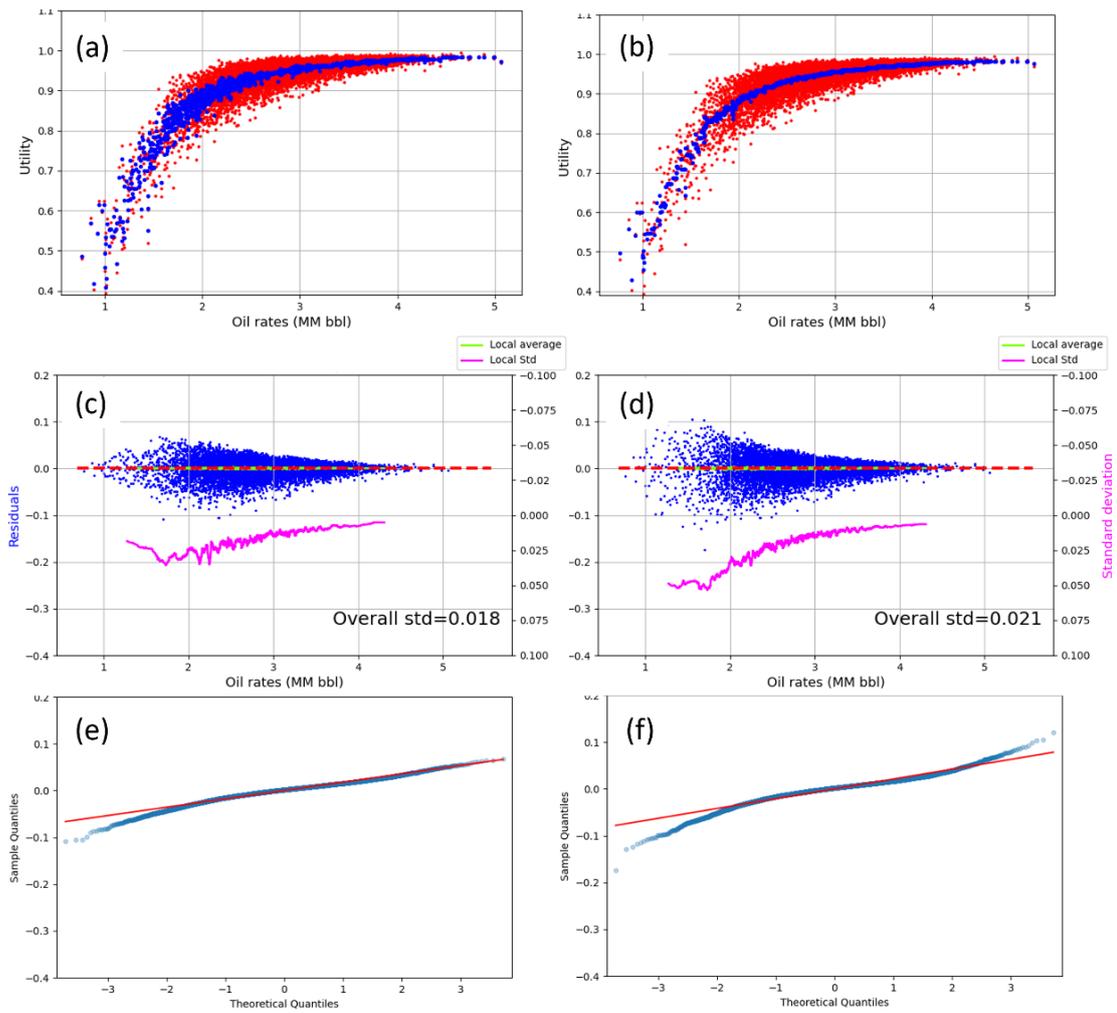


Figure 19. Single-dimensional Extreme Gradient Boost regression QC plots for risk-averse utilities: (a, c, e) before and (b, d, f) after parameter optimization. (a,b): Utilities (red) and predicted utilities (blue). (c,d): Residuals scatterplots. (e,f): Residuals Q-Q plots (blue) and standardized line (red).

7.3. COMPARISON OF DIFFERENT REGRESSION METHODS

In this section, two phases of testing are shown: single and multi-dimensional regressions. Single-dimensional regression was assessed first to facilitate the understating of the scatterplots and expectations of the different algorithms' predictions with respect the exponential behaviour of the data. For each case, the respective optimized parameters shown in Table 11 were used.

7.3.1. Single-dimensional regression

In the case of single-dimensional regression testing only production oil rates at year 1 were used. Figure 20 and Figure 21 show respectively the regression results and residuals using different linear and non-linear methods. As seen before, the red clouds clearly show the non-linearity of the utilities against the oil rates. Linear regression (a) is unable to grasp the trend of the data (Figure 20a); this is evidenced in the biased distribution of the errors and uneven spread (Figure 21a).

The non-linear regressors deliver a much more suitable fit, with some variations. SVR (e) is unable to capture the details in the higher end of the oil rates, and the residuals show strong bias in this place; this is evident in the regression residuals presented in Figure 21e, showing strong bias. KNN (b) and RF (c) produce a good result overall but struggle to grasp the lowest oil rates, below 1 MM bbl, producing a flat tale in this area. A similar observation is noted for Piecewise regression with a 10-slope split using RF regressor (f), where the lowest end is a discontinuous tale; this is also confirmed by the regression residuals shown in Figure 21f, evidencing important bias. XGB (d) offers the most sensible regression, with an overall fit including the edges, and exhibiting the least dipping windowed standard deviation of the errors in the low oil rates (Figure 21d), as other methods exhibit.

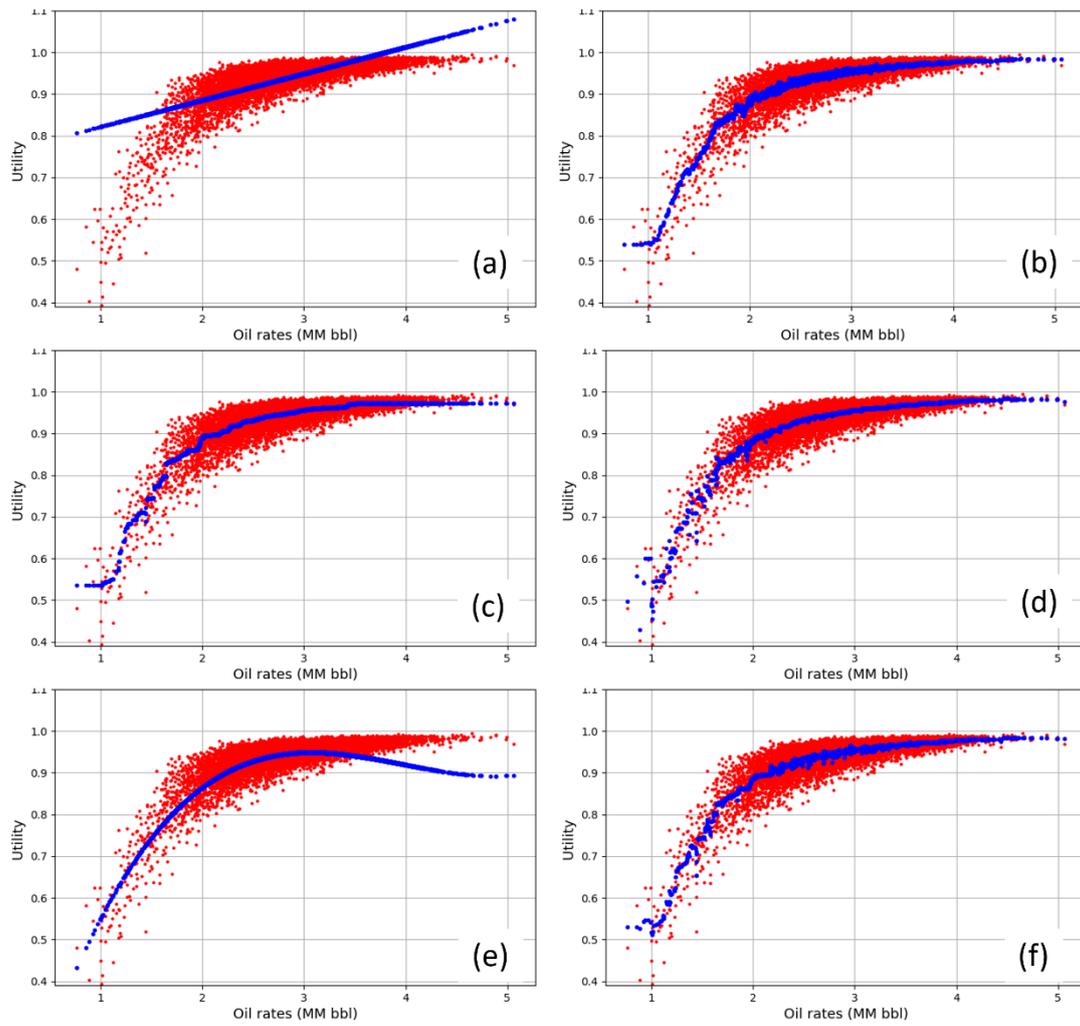


Figure 20. Single-dimensional regression fits (blue) using different algorithms for risk-averse utilities (red): (a) OLS, (b) (KNN), (c) Random Forest, (d) XGB, (e) SVR, (f) Piecewise using RF, 10 slopes.

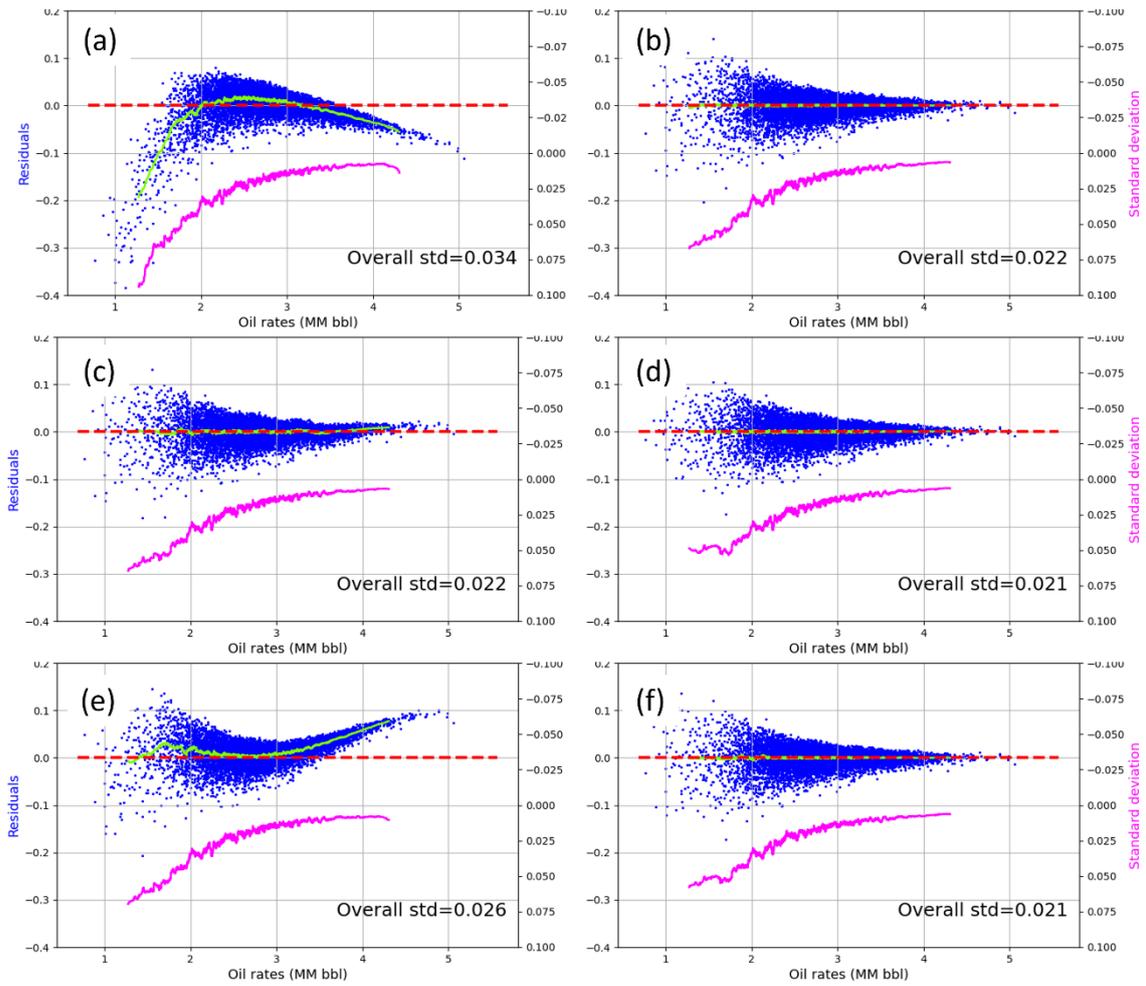


Figure 21. Single-dimensional regression residuals (blue) using different algorithms for risk-averse utilities: (a) OLS, (b) (KNN), (c) Random Forest, (d) XGB, (e) SVR, (f) Piecewise using RF, 10 slopes.

Table 12 shows R^2 scores corresponding to single-dimensional regressions of utilities versus observed oil rates, using the different algorithms. The numbers confirm the visual judgements made previously. OLS and SVR show very poor accuracy with R^2 scores of 0.509 and 0.645 respectively. KNN and RF produce better R^2 scores, both around 0.79, but the highest scores are achieved by XGB and Piecewise, with scores of 0.812 and 0.813 respectively. Despite Piecewise produced technically the highest score, XGB is still preferred, since its fit is more accurate at the

very low oil rates, as observed in the scatterplots from Figure 20d and Figure 21d. Piecewise had shown lack of accuracy at fitting the lower end, as shown in Figure 20f.

Method	OLS	KNN	RF	XGB	SVR	Piecewise (10 slopes RF)
R² score	0.509	0.799	0.797	0.812	0.645	0.813

Table 12. R² scores for single-dimensional regressions of utilities versus observed oil rates, using different algorithms.

From this phase of the testing, it is concluded that XGB is the most appropriate non-linear algorithm for single-dimensional regression of risk-averse utilities in the present case study.

7.3.2. Multi-dimensional regression

For the case of multi-dimensional regression testing, production oil rates from year 1 to year 30 were incorporated. Parameter optimization was performed multi-dimensionally. Figure 22 and Figure 23 show the regression results and residuals with the different linear and non-linear methods. In a multi-dimensional regression, the output regression points conform a 3D ‘surface’; when this surface is plotted in a 2D graph of the dependant versus independent variables, equivalent to the graph in Figure 22, the predicted values (blue dots) are seen as a cloud of points, instead of the line that would result from a single-dimensional regression. For this reason, to facilitate the visual understanding of the regression fits, oil rates scatterplots in Figure 22 show only the rates corresponding to year 1, even when the regression has been done multi-dimensional. The residuals scatterplots in Figure 23 are shown for all the production years from 1 to 30.

Similarly seen in the single-dimensional case, Figure 22 shows that OLS (a) and SVR (e) are not suitable for the fit, evidenced by the residuals local standard deviation, (Figure 23a,e), being asymmetrically distributed. Piecewise (f) produces an apparent better fit (Figure 22f), but the residuals evidence overfitting in the case of multi-dimensional case (errors are all very close to zero, Figure 23f), therefore it is inadequate.

KNN (b) and RF (c) are visually better fits, however they both have some failures: KNN suffers the same difficulty at the lower rates as in the single-dimensional case, producing a flat lower tail (Figure 22b). RF overcomes this issue in the multi-dimensional case, but the predicted trend is slightly dislocated in the upper end, underestimating the utilities (Figure 22c); this creates a slight bias in the residuals distribution in this area (Figure 23c), making it not the most desirable solution. XGB once again offers the most sensible regression, with an overall fit from low to high ends of the oil rates (Figure 22d); the residual plot shows a very stable distribution, and the windowed standard deviation doesn't exhibit bias at any of the edges (Figure 23d).

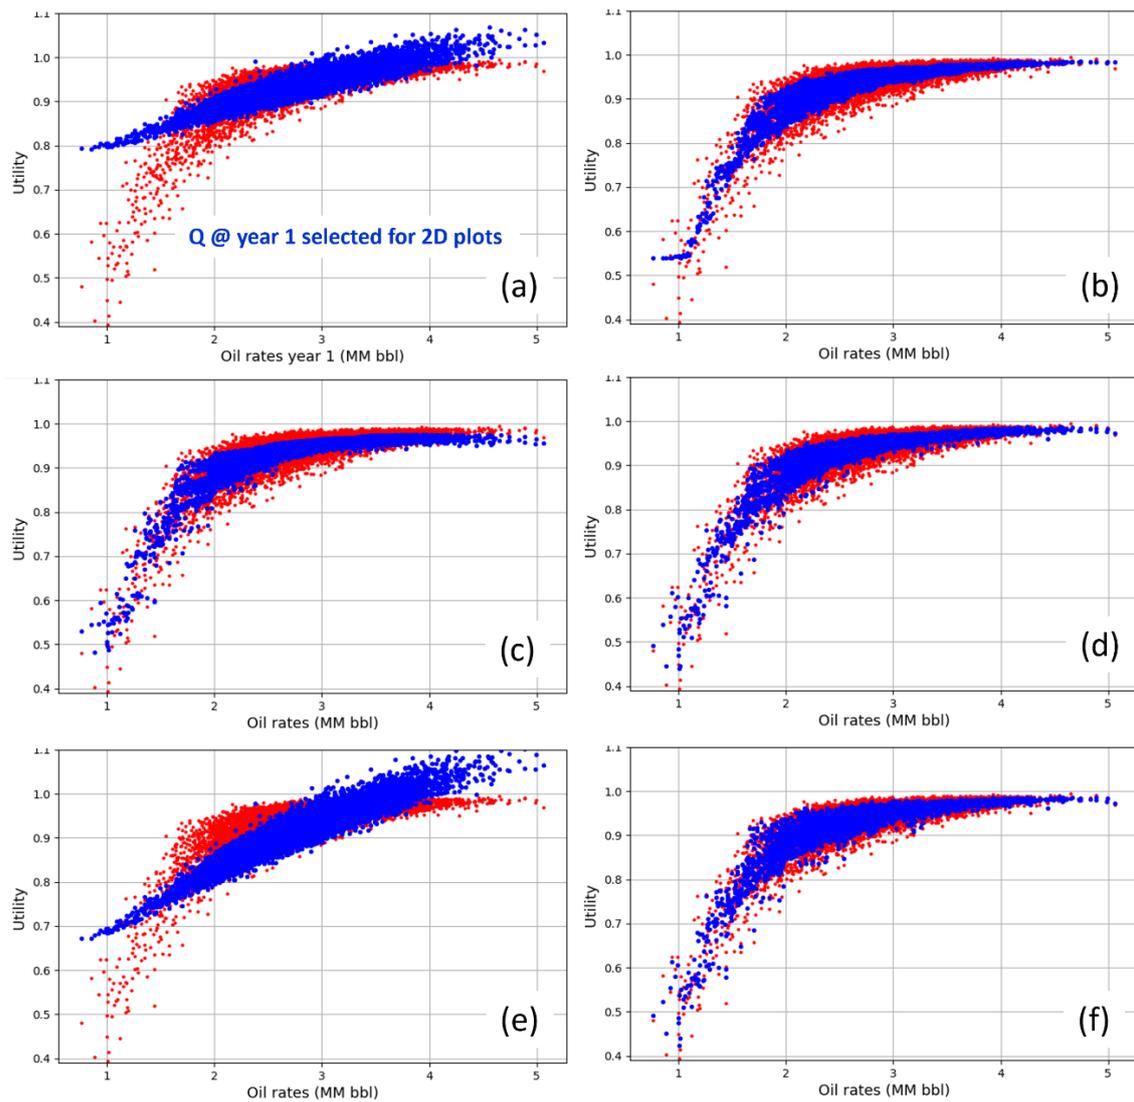


Figure 22. Multi- dimensional regression fits (blue) using different algorithms for risk-averse utilities (red): (a) OLS, (b) (KNN), (c) Random Forest, (d) XGB, (e) SVR, (f) Piecewise using XGB, 4 slopes. For 2D display only oil rates at year 1 are selected.

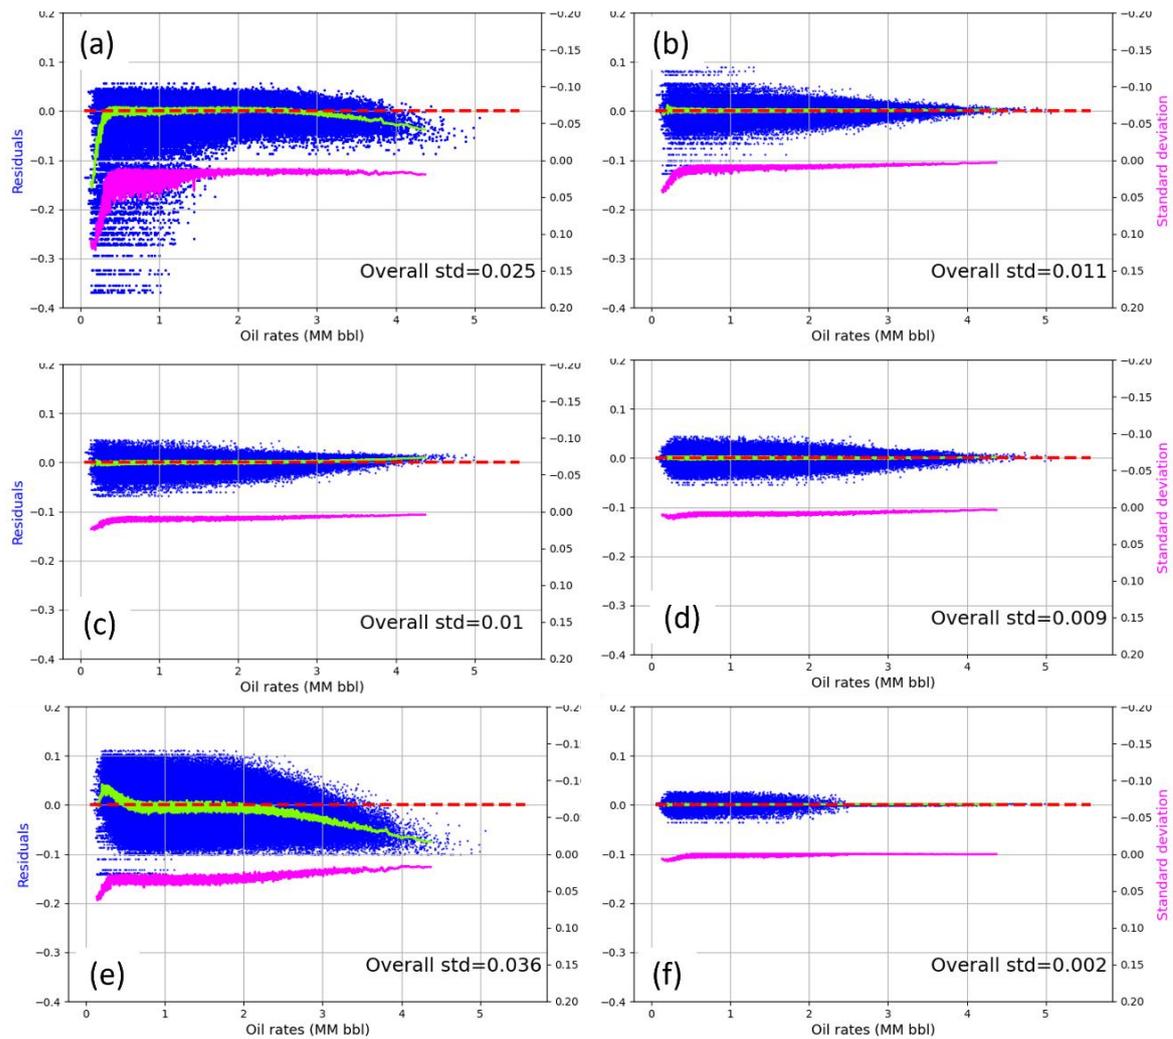


Figure 23. Multi-dimensional regression residuals (blue) using different algorithms for risk-averse utilities: (a) OLS, (b) (KNN), (c) Random Forest, (d) XGB, (e) SVR, (f) Piecewise using XGB, 4 slopes.

To visually assess the accuracy of the fit in a 3D view, a bi-dimensional regression has been performed using oil rates from years 1 and 30. Figure 24 shows the regression results for (a) OLS and (b) KNN. In a 3D perspective the inaccuracy of the OLS (Figure 24a) becomes very evident, especially at the lower rates. KNN (Figure 24b) offers a better solution, but also in the 3D the data point at the lowest rates can be seen left off the regression surface. Figure 25a shows the regression result for RF, which handles much better the lower rates; however one conclusion of this testing, is that the various QC plots need to be assessed in conjunction; the residuals from Figure

23 show that RF creates a slight bias at the highest oil rates, reason for which this method is not chosen for the case study. Figure 25b shows the 3D fit from XGB, where an overall suitable surface fit can be seen; moreover, the XGB residuals are the most homogeneous from all the algorithms tested, with the flattest local standard deviation shown in Figure 23d. For other algorithms 3D QCs, see Appendix B.

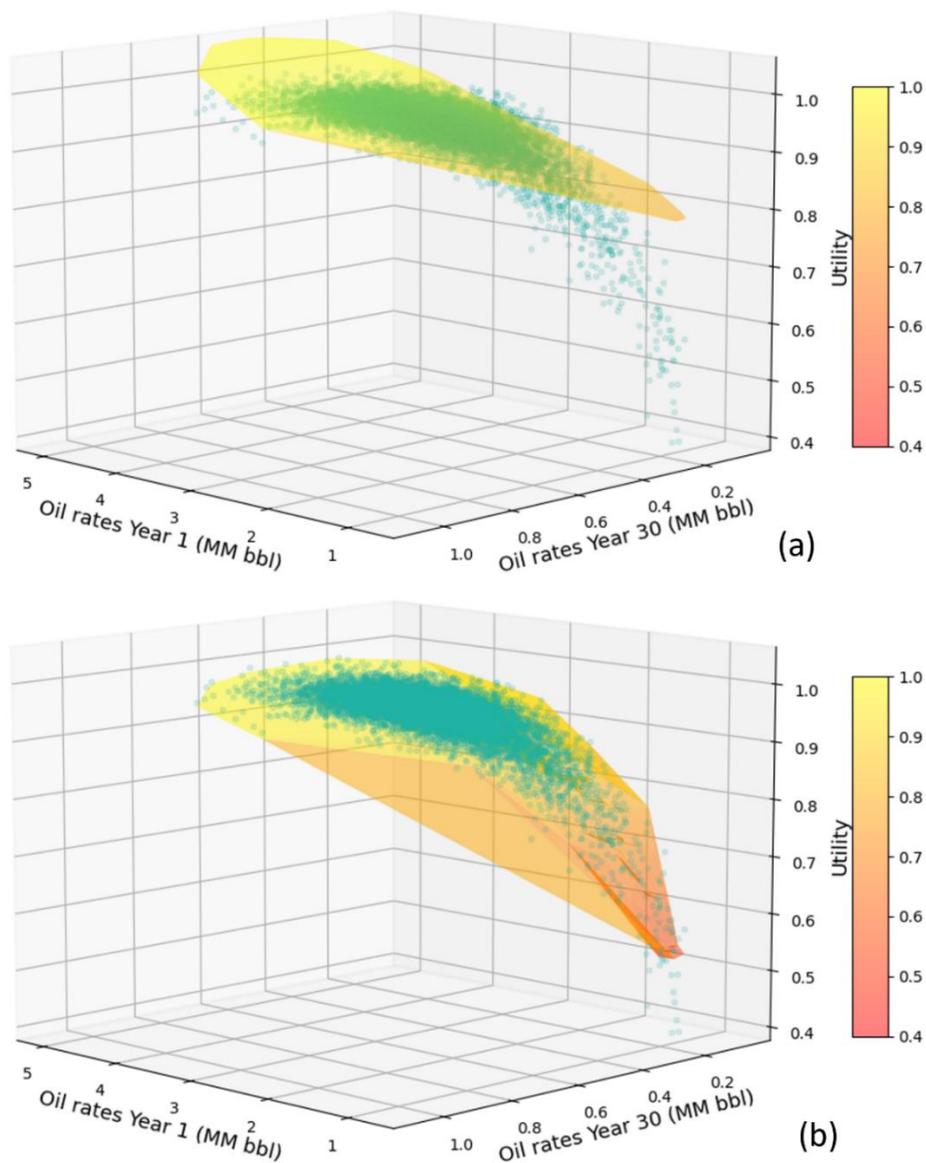


Figure 24. Bi-dimensional regression 3D fit plot, using (a) OLS and (b) KNN. Blue dots: input data. Warmer colors: regression surface.

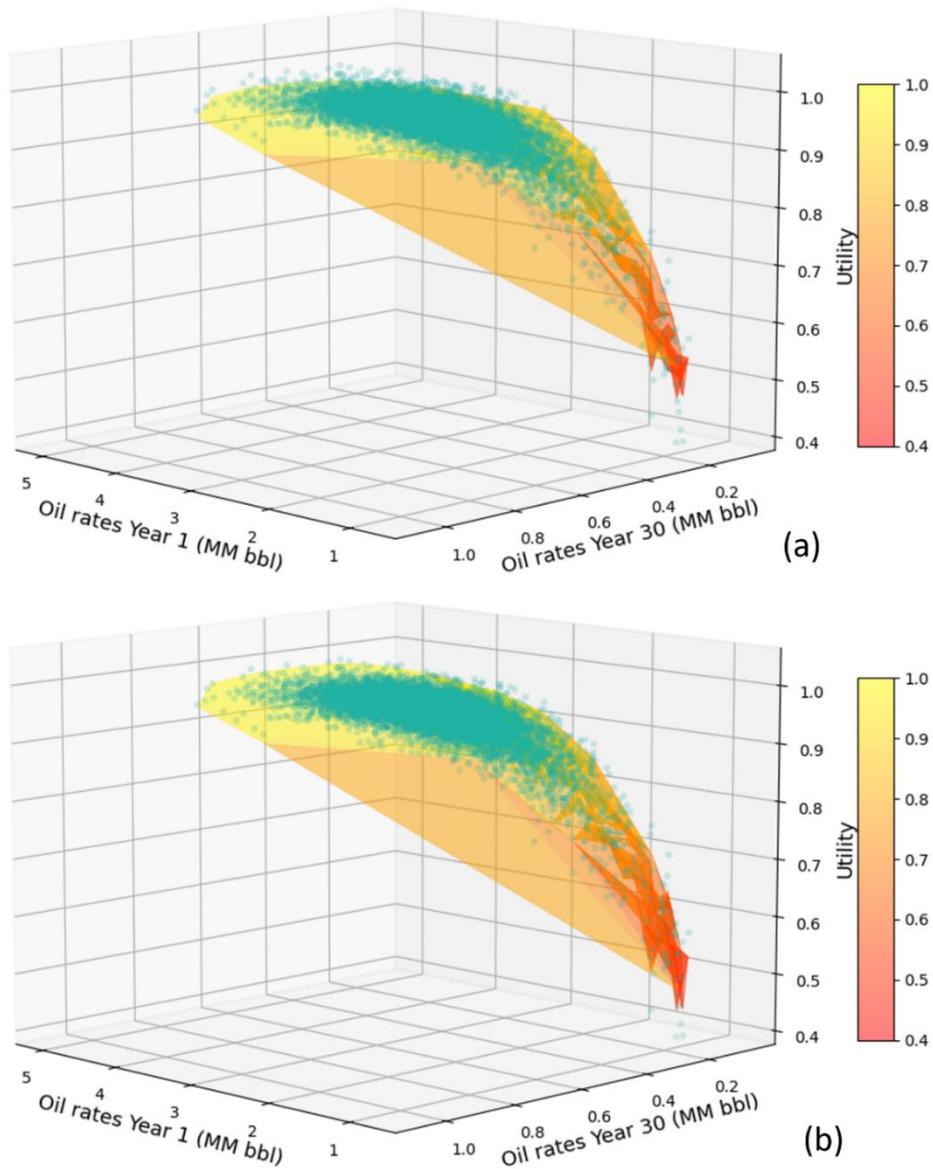


Figure 25. Bi-dimensional regression 3D fit plot, using (a) RF and (b) XGB. Blue dots: input data. Warmer colors: regression surface.

To quantify the results, Table 13 documents the R^2 scores corresponding to multi-dimensional regressions of utilities versus observed oil rates from year 1 to year 30, using different algorithms. In the multi-dimensional case, the scores also support previous inferences made with the 2D and 3D plots. OLS and SVR show very poor

accuracy with R^2 scores of 0.736 and 0.449 respectively. KNN and RF produce better R^2 scores, both around 0.95. XGB achieves a score of 0.963 and Piecewise achieves 0.997. Despite Piecewise produced again the highest R^2 score, its residuals evidenced overfitting, as shown in Figure 23f. Therefore, XGB is considered a more suitable fit.

Method	OLS	KNN	RF	XGB	SVR	Piecewise (10 slopes RF)
R^2 score	0.736	0.951	0.956	0.963	0.449	0.997

Table 13. R^2 scores for multi-dimensional regressions of utilities versus observed oil rates, using different algorithms.

From the multi-dimensional regression testing, it is also concluded that XGB is the most suitable machine learning non-regression method for the specific case study. The selection of XGB is consistent with a recent work which chose the same algorithm as part of a simulation-regression approach to estimate VOI in a case study of CO₂ injection and storage at the Utsira formation. For details on this work please refer to Shahali (2022).

Since the case study involves validation of the neutral case, which implies linear NPV-oil rates relationships, it is tempting to assess the performance of XGB on a linear case and compare with previous results obtained with OLS. Figure 26 shows the single-dimensional regression of NPVs on a risk-neutral case using XGB algorithm. Comparing with the OLS regression results for the same case, shown previously in Figure 15, it can be inferred that the XGB provides a more accurate fit. The slightly different slope observed in lower end of the oil rates (below 2 MM bbl) is better captured, as shown in Figure 26a, and the residuals are more symmetrical and less biased at the edges; this is evidenced by much flatter local average at the tails, as shown in Figure 26b.

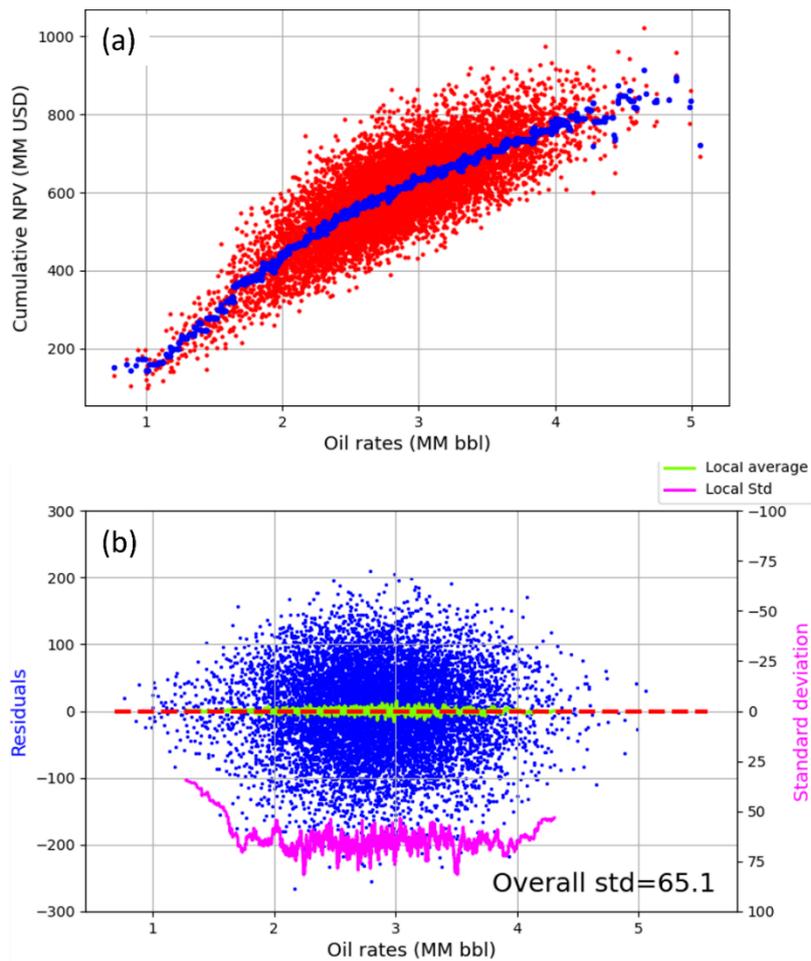


Figure 26. Scatterplots of a single-dimensional XGB regression of NPVs in a risk-neutral case. (a) Scatterplots of NPVs vs oil rates (red) and predicted NPVs (blue). (b) Residuals scatterplot (blue), with windowed average (green) and standard deviation (pink).

The R^2 score for the XGB fit resulted in 0.689. Recalling the OLS exercise, the R^2 score obtained was 0.654. This quantitative improvement in the XGB score supports the conclusion from the visual assessment, that XGB is more suitable even for risk-neutral assessments in the present case study, and therefore will provide more reliable predictions.

Chapter 8: Case study – Optimal time to initiate and terminate the IOR project

The current case study uses hypothetical oil production and economic parameters, to build on a sequential decision-making problem focused on hydrocarbon production optimization. The problem, model and parameters are explained in this section.

8.1. PROBLEM DEFINITION

This case study builds on the work by Hong et al. (2019) and Wui (2019). They considered a hypothetical production field with a maximum lifetime of 50 years and assessed the optimal initiation and termination time of the improved oil recovery (IOR) phase. For this, they used a fully structured sequential reservoir-decision-making (SRDM) scheme, which considered the sequences of uncertainty, learning and decisions. The potential decision alternatives at each year of production were (1) to continue with primary production, (2) to switch to secondary production, or (3) stop production. The methodology featured the LSM workflow with linear regression (OLS) to estimate the conditional expectation at each alternative, at each time step, and the problem was solved recursively executing the decision policy driven by the highest expected value. The case was treated under risk-neutrality; hence, the objective function was to maximize the expected monetary value, namely expected net present value (ENPV).

The novelties of the present work are (1) the introduction of risk attitude to assess the impact on the decisions, and (2) the use of machine learning non-linear regression, instead of OLS regression, to estimate the conditional utilities. The optimal decisions obtained in the new sequential decision-making problem were compared to those obtained while neutral in Hong et al. (2019), and sensitivity analysis was performed on various aspects for the different risk profiles.

The introduction of risk attitudes brings two main differences:

- Utility function: As discussed previously, a risk-averse or risk-seeking decision-maker, has a utility function that can be used to define their preference probabilities versus the monetary values. To comply with the risk attitude, the optimal decision policy is driven by maximizing the expected utilities (EU) (Hong and Bartvold, 2023), differently than in Hong et al. (2019)'s work, where ENPV were used directly

- Non-linearity: The utility function introduces non-linearity between the uncertain variables and the utilities. For a neutral case, the relation between the uncertainties and the monetary values is linear, therefore ordinary linear regression (OLS) was used in Hong et al. (2019) to approximate the expected values (expected NPV) at each decision point during the LSM algorithm. In the present case, to approximate the expected utility, machine learning methods for non-linear regression were evaluated and implemented to find the best fit of the utilities versus the uncertain variables.

As described by Hong et al. (2019), there is uncertainty about the reservoir properties and the effects of recovery mechanisms. Following their work, these uncertainties were quantified by assigning probability distributions to the parameters of a production model. The production model is based on the exponential decline and was used to calculate the potential production oil rates that would be observed at each year of production during a 50-year lifetime. These oil rates, are then incorporated in the sequential decision-making scheme as “information” obtained sequentially, therefore used to support further decisions. The production model used is explained in the next sub-section. For the monetary values, namely cash flows and NPVs, a representative economic model was used, which is also explained later. These monetary amounts were then converted to utilities, using an exponential utility function.

8.1.1. Two-factor production model

In this work, the two-factor production model proposed by Parra-Sanchez (2010) was used to model oil production. This model is useful and very tractable, meaning that enables an efficient analysis of the problem within the time and resources available (Hong et al., 2019).

The two-factor production model is based on an exponential decline and can include multiple phases of production recovery. The exponential decline is defined as follows:

$$E_R(t) = E_R^0 + (E_R^\infty - E_R^0) * (1 - e^{-\frac{t}{\tau}}) \quad (11)$$

The model parameters are described below:

$E_R(t)$: Recovery factor until time t

E_R^0 : Recovery factor at time 0

E_R^∞ : Theoretical ultimate recovery factor

t : Total lifetime of the recovery phases.

τ : Time constant for production

The recovery factor E_R is the fraction of the recovered amount of hydrocarbons against the original oil in place (OOIP) and depends on the reservoir properties and production mechanisms. This production model is called the “two-factor” model due to two parameters that strongly depend on the reservoir properties: (1) the theoretical ultimate recovery factor (E_R^∞), and (2) a time constant (τ) that indicates how fast the increment of the recovery factor is for a recovery mechanism.

Based on the case studies from Hong et al. (2019) and Wui (2019), the present work also focused in two recovery phases only - primary and secondary recovery. The time at which the primary recovery is stopped, and the secondary recovery is started, is called in this work the ‘switch time’, and this switch from the primary recovery to the secondary recovery is assumed to happen only once and to be irreversible.

For the primary recovery phase, the two-factor model can be expressed as follows:

$$E_{R1}(t) = E_{R1}^0 + (E_{R1}^\infty - E_{R1}^0) * (1 - e^{-\frac{t}{\tau_1}}) \quad (12)$$

E_{R1}^0 : Primary recovery factor at time zero

E_{R1}^∞ : Theoretical ultimate recovery factor for primary production

For the secondary recovery phase, the two-factor model the exponential decline can be defined as:

$$E_{R2}(t) = E_{R1}(t_{R1}) + (\Delta E_{R2}^\infty) * (1 - e^{-\frac{-(t-t_{R1})}{\tau_2}}) \quad (13)$$

$E_{R1}(t_{R1})$: Primary recovery factor at the end of the primary production

ΔE_{R2}^∞ : Theoretical ultimate increment of secondary recovery factor when there is a switch from primary to secondary recovery phase.

With the recovery factors the oil production rates are calculated at each time step as follows:

$$q_{o,j} = \frac{N_o * [E_R(t_j) - E_R(t_{j-1})]}{\Delta t_j} \quad (14)$$

$q_{o,j}$: Oil production rate at time step j

N_o : Estimated OOIP

Δt_j : Period between time step j and time step j-1

8.1.1.1. Model parameters

To model the oil production, as per Hong et al. (2019), bounded probability distributions were assigned over the model parameters using prior knowledge to quantify the uncertainties in the reservoir properties and the effects of the primary- and secondary-recovery mechanisms. The uncertain variables in the production model are:

OOIP: Original oil in place

E_{R1}^{∞} : Theoretical ultimate recovery factor for primary production

τ_1 : Time constant for primary production

τ_2 : Time constant for primary production

ΔE_{R2}^{∞} : Theoretical ultimate increment of secondary recovery factor and time constant for secondary production

Table 14 shows the truncated normal distribution parameters for these uncertainties. Table 15 shows the correlation coefficient between them.

Parameter	N_o (MM bbl)	E_{R1}^{∞} (fraction)	τ_1 (years)	ΔE_{R2}^{∞} (fraction)	τ_2 (years)
Mean	240	0.2	16	0.15	7
Standard deviation	35	0.05	2	0.05	1.5
Minimum	10	0.05	1	0.01	1
Maximum	1000	0.5	30	0.31	13

Table 14: Means, standard deviations, and boundaries for the truncated normal distributions of the production model parameters. From Hong et al. (2019); Wui (2019).

	N_o	E_{R1}^{∞}	τ_1	ΔE_{R2}^{∞}	τ_2
N_o	1	-0.8	0.16	0.56	-0.08
E_{R1}^{∞}	-0.8	1	0.2	-0.7	0.1
τ_1	0.16	0.2	1	-0.3	-0.2
ΔE_{R2}^{∞}	0.56	-0.7	-0.3	1	-0.3
τ_2	-0.08	0.1	-0.2	-0.3	1

Table 15: Correlation coefficients between the production-model parameters. From Hong et al. (2019); Wui (2019).

Multi-variate Monte Carlo simulation were used to generate the samples of the model parameters. Wui (2019) demonstrated that a sample size of 10,000 is large enough for sensible estimation of the value of information (VOI), therefore, 10,000 samples of each model parameter were generated for this study. The production oil rates were then calculated using equation 14, for each year of production lifetime.

Imperfect information can be represented by adding a measurement error (i.e., noise) to an information value, before building the regression model (Shahali, 2022). In this study, the measured oil rate Q_j was sampled from a normal distribution with a mean of $q_{o,j}$, and a standard deviation as a resulting of a percentage of $q_{o,j}$:

$$Q_j \sim N\left(q_{o,j}, \frac{\% \text{ noise}}{100} \times q_{o,j}\right) \quad (15)$$

The measurement error considered in this work was 10%. The samples Q_j are the ‘measured/observed’ oil production rates and these are the ones used in regression.

8.1.2. Economic model

In the present study, the objective function is to maximize the expected utility, therefore, monetary values need to be estimated first. A representative economic model was used to calculate the cash flows from each time step, with fixed oil price, capital

expenditure (CAPEX) and operational expenditure (OPEX). The cash flow from time steps $j-1$ to j is calculated as follows:

$$CF_j = q_{o,j} \times P_o \times \Delta t_j - (CAPEX_j + OPEX_j) \quad (16)$$

CF_j : Cash flow at time step j

P_o : Oil price

Δt_j : Time period between time step j and time step $j-1$

$CAPEX_j$: Capital expenditure at time step j

$OPEX_j$: Operational expenditure at time step j

The NPV of the cash flows over the production lifetime is given by the following equation:

$$NPV = \sum_{j=0}^{n_t} \frac{CF_j}{(1+r)^{t_j}} \quad (17)$$

n_t : Number of time steps

CF_j : Cash flow at time step j

r : discount rate

As $q_{o,j}$ is a function of the primary and secondary recovery lifetimes, NPV is also a function of these two lifetimes. It should be noted that the oil rates $q_{o,j}$ are not considered a Markovian process over time, so the conditional ENPV and EU given measured oil rates should be assessed by regressing NPVs and utilities on all oil rates that have been measured over time before a decision.

In this work, the economic parameters are assumed to be invariant over time, so the problem setting does not model the learning of economic uncertainties but only the

learning of reservoir and production uncertainties over time. The economic parameters assumed for this case study, as per Hong et al. (2019), are listed in the Table 16.

Economic parameters	Values	Units
Oil Price	50	\$/bbl
CAPEX (Primary)	50	\$M
CAPEX_2After1 (Secondary)	40	\$M
CAPEX_2No1 (Secondary)	75	\$M
OPEX (Primary)	20	\$MM/year
OPEX (Secondary)	30	\$MM/year
Discount rate	12	%, per year

Table 16: Economic parameters. From Hong et al. (2019); Wui (2019).

CAPEX_2After1: Capital cost of initiating secondary recovery after having primary recovery.

CAPEX_2No1: Capital cost of initiating secondary recovery without having primary recovery.

The CAPEX will be deducted only at the year when the recovery phase is starting. The OPEX will be deducted every year according to the recovery type (i.e.: primary or secondary).

8.1.3. Utility function

The corresponding utilities to the NPV values were calculated using a utility function. An exponential utility function was implemented, which could be used to model risk-aversion and risk-seeking behaviours, according to the sign of the risk tolerance (Hong and Bartvold, 2023). The utility function used is expressed below:

$$u(x) = 1 - \text{sgn}(\rho) e^{-\frac{x}{\rho}} \quad (18)$$

where x represents monetary amounts (NPV): positive values for gain and negative values for loss. ρ denotes the risk tolerance; $\rho > 0$ for risk aversion and $\rho < 0$ for risk-seeking. $\text{sgn}(\rho)$ takes the sign $\text{sgn}(\rho) = 1$ if ρ is positive and $\text{sgn}(\rho) = -1$ if ρ is negative.

8.1.4. Influence diagram

After having described the problem setting, model and economic parameters, the influence diagram of this case study is shown in Figure 27. This diagram is designed for risk attitudes different than neutral, where the objective function is to maximise expected utility, rather than NPV which in this case is drawn as an uncertainty influencing the utility. Utility is represented with the blue rounded rectangle. The yellow rectangle represents the decision to assess at each time step, and the uncertainties (reservoir properties, oil rates, measured oil rates and NPV) are shown in green ellipses.

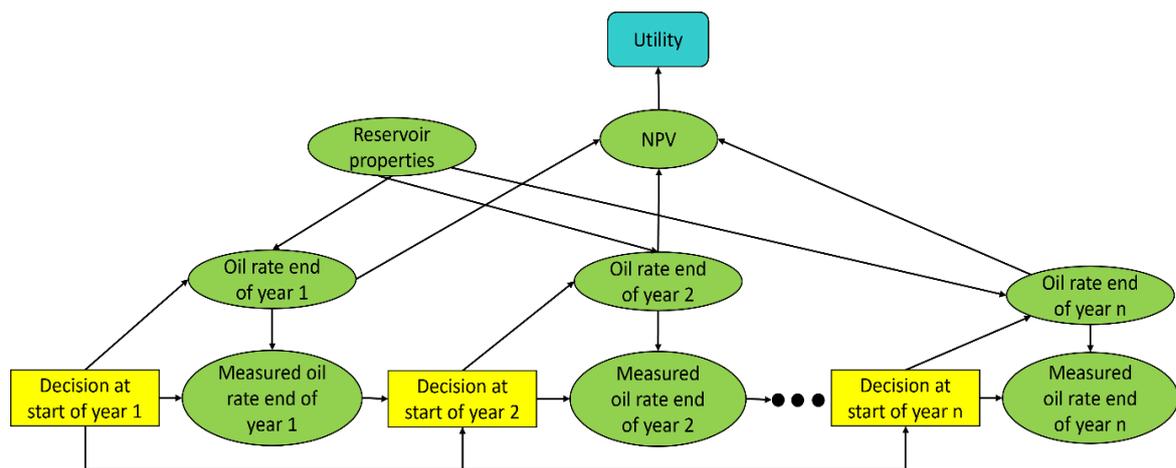


Figure 27. Influence diagram of the IOR initiation and termination time case study for risk-averse or risk-seeking profiles. Arrows should be drawn from each decision node to the 'oil rate' and 'measured oil rate' uncertainty for all years, and to the 'NPV' uncertainty node. For simplicity of the diagram in this picture some of these arrows are omitted.

8.2. LEAST SQUARE MONTE CARLO IMPLEMENTATION FOR THE IOR PROBLEM

The approach used to solve the IOR sequential decision problem in the present work is based on the LSM algorithm, incorporating the observation of oil rates with time in support of the decisions. The workflow was programmed in Python, replicating the MATLAB implementation created by Wui (2019) and Hong et al. (2019), both designed for a neutral risk attitude and using OLS in the regressions. The details of this workflow and its validation are explained in the next sub-sections.

8.2.1. Workflow and modelling in Python

The LSM algorithm implemented in Python for this work is illustrated in Figure 28. This schematic view was put together based on the general LSM methodology well explained by Ahmadi and Bravtvoid (2023) , which is now adapted to the specific case study.

The methodology starts by defining the number of paths (individual realizations) of the uncertain variables, and the number of time steps. Following the recommendation from Wui (2019), 10,000 paths were defined, being a suitable sample size for an accurate assessment of the VOI. The number of time steps was kept the same as that in Wui (2019)'s work, with a maximum production lifetime of 50 years, where each year is an observation and decision point. After this definition the approach continues with the forward simulation of the uncertain variables over all the paths; this consisted in a multi-variate Monte Carlo simulation over five model parameters, as per truncated normal distributions defined in Table 14 and correlation coefficients defined in Table 15:

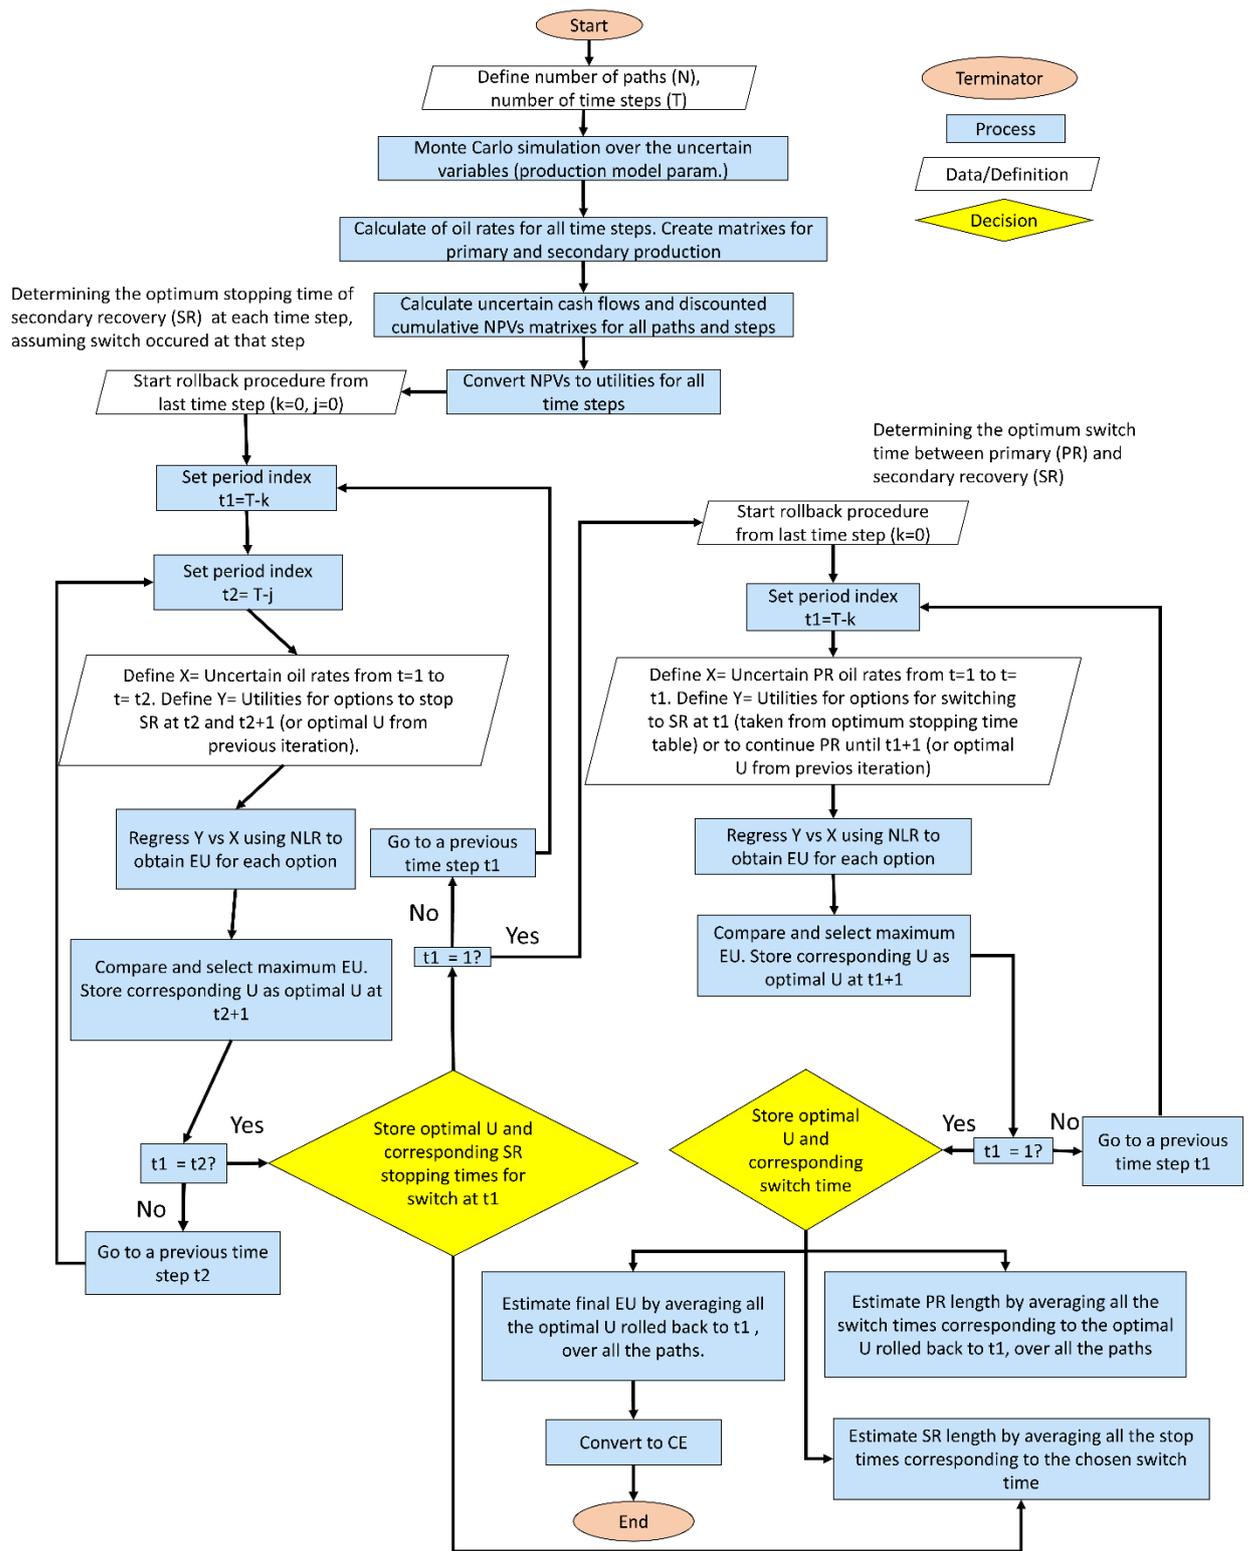


Figure 28. LSM workflow for the present case study. Modified from Ahmadi and Bravtvoild (2023) to reflect the specific IOR initiation and termination problem, and incorporation of risk attitude.

These model parameters enabled the calculation of the oil rates at each time step of the production lifetime, using equation 14. Measurement error was accounted for by performing Monte Carlo simulation of a normal distribution where the means were the calculated oil rates, and then standard deviations were 10% of these calculated values, as per equation 15. This new sample set is considered the “observed” oil rates and the ones used for regression.

The corresponding cash flows were also calculated for each time step and allowing for any possible combination of primary recovery and secondary recovery lifetimes, using equation 16, followed by the NPV calculation (equation 17). Then, using the exponential utility function (equation 18) the NPVs were converted to utilities. A preliminary risk tolerance of \$200 million was used. Sensitivity analysis of risk tolerance will be performed later.

The identification of the optimal decision policy was split into two phases:

1. Determining the optimal stopping time of secondary recovery, given a switch from primary production at each year, for a given path.
2. Determining the optimal switch time from primary to secondary recovery, with secondary recovery’s corresponding optimal stopping time determined in the previous step, for a given path.

To incorporate the ‘learnings over time’, the decision-making process starts at year 50, and recursively progress backwards in time in a rollback procedure. Two timing indexes are created to represent the switch from primary to secondary recovery at each year, and the possible stopping times of the secondary recovery at each of those years. For each decision point, all the oil rates from year 1 to the assessment year are used, and the utilities corresponding to that assessment year. To exercise the optimal decision policy, the expected utilities (EU) from the available options (explained next) are compared at each decision point. As described in the LSM algorithm, the EU are estimated by using non-linear multi-dimensional regression (using the chosen XGB

algorithm) where the X axis is represented by the oil rates, and the Y axis is represented by the EU from each option.

1. Determining the optimal stopping time of secondary recovery, given a switch from primary production at each year:

The workflow starts by defining the optimal secondary recovery stopping time assuming “switch” from primary to secondary recovery happened at the start of year 50. In such case, there are only two alternatives available: “stop secondary recovery at the start of year 50” or “continue secondary recovery and stop at the end of year 50” (maximum lifetime). Here the utilities corresponding to these two alternatives are regressed on the oil rates from year 1 to the start of year 50 (this is the information known until that decision point). The EUs from each option are compared, the largest EU is chosen, and its corresponding utility and decision (stopping time) are stored in a database that will be used later in the second part of the workflow. This stored utility and decision represents the optimal secondary recovery stopping time if the switch to secondary recovery happened at the start of year 50. This concludes the analysis for a potential switch occurring at the start of year 50.

The process continues backwards in time, and now the analysis consists in finding optimal secondary recovery stopping time if the switch from primary recovery happened at the start of year 49. At the start of year 49, there are three alternatives now: “stop secondary recovery at the start of year 49”, “continue secondary recovery and stop at the start of year 50” or “continue secondary recovery and stop at the end of year 50”. First, the alternatives of “continue and stop at the start of year 50” and “continue and stop at the end of year 50” are compared. The utilities corresponding to these two alternatives are regressed on the oil rates from year 1 to the start of year 50 (all information known until the decision point). The EUs from each option are compared, the largest EU is chosen, and its corresponding utility is “temporarily” stored. This

optimal utility is temporarily stored because now it will be compared with the alternative of “stop secondary recovery at the start of year 49”. The regression now considers the rates from year 1 to the start of year 49 (information known until the new decision point). The EUs from each option are compared, the largest EU is chosen, and its corresponding utility and decision (stopping time) are stored. This stored utility and decision represents the optimal secondary recovery stopping time if the switch to secondary recovery happened at the start of year 49. This concludes the analysis for a potential switch occurring at the start of year 49.

This algorithm is repeated consecutively for each potential switch year until year 1 is reached. As the switch time moves backwards by a year the number of decision points increases by 1. For each switch year assessment, the stopping time analysis starts back at year 50. Within the same switch year assessment, as the stopping decision time moves backwards by a year, the number of points for regression decreases by 1. For all years except year 1 the regressions are multi-dimensional. At the end of this loop, a database containing the optimal stopping time for each potential switch year has been recorded for each realization and will be used in the second phase, to be explained next.

2. Determining the optimal switch time, from primary to secondary recovery:

Following the stopping time decision policy, the approach continues with the assessment of optimal IOR initiation time, namely the “switch” time from primary to secondary recovery. The first step is to compute the utilities for a full lifetime of primary production from year 1 until the end of year 50. The workflow now starts by comparing two alternatives at the start of year 50: “continue with primary production until the end of year 50” and a “switch to secondary production at the start of year 50” (the utilities for the “switch at the start of year 50” alternative, with its optimal stopping time, were already computed in the first part of the workflow). In this step the utilities corresponding to these two alternatives are regressed on the oil rates from year 1 to the

start of year 50 (this is the information known until that decision point). The EUs of the two alternatives for each path are compared, the alternative with a largest EU is chosen, and the utility associated with the chosen/optimal alternative is “temporarily” stored for each path. This optimal utility is temporarily stored because the corresponding optimal alternative will be now compared to the alternative of “switch to secondary recovery at the start of year 49” (also already calculated in the first part of the workflow, with its optimal stopping time). The regression now considers the oil rates from year 1 to the start of year 49 (information known until the decision point at the start of year 49). The EUs from each decision alternative are compared, the largest EU is chosen, and its corresponding utility is “temporarily” stored, to be used in the comparison with the alternative of “switch to secondary recovery at the start of year 48” (calculated in the first part of the workflow). The regression now considers the oil rates from year 1 to the start of year 48 are performed to calculate the EUs, and the process continues in the same logic backwards in time.

This is repeated consecutively for each year until reaching year 1. As the time moves backwards by a year the number of decision points increases by 1 and the number of points for regression decreases by 1. At the end of this loop, the optimal utilities correspond to the optimal IOR initiation time decision policy. The optimal lengths of the primary and secondary recovery are stored for each path. Final EU is calculated by averaging the utilities over all the paths. This represents the expected utility with “imperfect information”. This is then converted to certain equivalent (CE) using the inverse utility function (solving for ‘x’ in equation 18). This is considered the “CE with information”.

As mentioned in Chapter 3, section 3.4: Value of information analysis, when the utility function is exponential, the VOI can be calculated as the subtraction of the “CE with information” and the “CE without information.” To estimate the “CE without information,” the expected utility “without information” is first computed by averaging

the utilities of each possible alternative of the problem over their respective multiple realizations, and then choosing the alternative that delivers the largest EU. This EU is then converted to CE using the inverse utility function. After this the VOI is computed.

8.2.1.1. Validation against predecessor LSM workflow, using risk-neutral and OLS

Once the described LSM workflow was implemented in Python, the next immediate step was to validate it by comparing the results against those by Hong et al. (2019) and Wui (2019) for a risk-neutral case and using OLS.

Table 17 shows the expected value without information (EVWOI), expected value with information (EVWI) and value of information (VOI), estimated with the Python implementation for this thesis work and by the predecessor’s implementation by Hong et al. (2019).

Estimated value	Hong et al. (2019)	Python implementation 2023
EVWOI (MMS)	756.20	755.98
EVWI (MMS)	805.90	805.29
VOI (MMS)	49.70	49.31

Table 17. Comparison of EVWOI, EVWI and VOI estimated by Hong et al. (2019) and the replicated Python workflow in this work.

Hong et al. (2019) estimated a VOI of \$49.7 million, while the new implementation in Python delivered a VOI of \$49.31 million. Although the results are very comparable, it is noted decimal precision differences in the VOI with both workflows. Mild differences are expected due to the following reasons:

- Randomness implied by the Monte Carlo sampling of the model parameters.

- Randomness implied by the Monte Carlo sampling of the measurement error added to the observed rates, described in equation 15.

To understand any potential difference not explained by the reasons above, an exercise was run using equal fixed 100 realizations in both implementations, MATLAB predecessor’s code and the replicated Python code. Using a relatively small number of realizations facilitated a 1-to-1 comparison of the results in the steps at which differences were found. Table 18 shows the results of this new test with additional decimal precision. MATLAB code estimated a VOI of \$22.9941 million, while the new implementation in Python delivered a VOI of \$22.9955 million.

Estimated value	MATLAB implementation	Python implementation
EVWOI (MMS)	2274.1557	2274.1557
EVWI (MMS)	2297.1499	2297.1513
VOI (MMS)	22.9941	22.9955

Table 18. Comparison of EVWOI, EVWI and VOI estimated by the MATLAB and Python implementation when using common fixed 100 realizations.

The decimal differences were assessed at specific steps. It was found that each software slightly differs in the decimal precision of the regression results when the input ‘y’ variables (NPV) are very similar. An example of this case was found when assessing the optimal stopping time of secondary recovery if switching to secondary recovery at year 10: At year 16 a regression is performed for two options for stopping time, having respective NPVs of \$732.21 million and \$732.25 million. The regressed NPVs with MATLAB were respectively \$732.2117 million and \$732.1668 million; the regressed NPVs with new Python code were respectively \$732.2117 million, \$732.2141 million. The optimal stopping time is determined by choosing the maximum ENPV, hence, for

MATLAB the optimal option would be first option (\$732.2117 million), and for Python would have been the second option (\$732.2141 million). In this case the mild decimal differences in the regression results would cause a difference of 1 year in the determination of the stopping time.

From this analysis, it can be concluded that another reason contributing to the mild differences in the results from Table 17 is the slightly different handling of decimal precisions by each software, when regressing very similar ‘y’ values. Despite this, it was considered that the final EVWI and VOI results are similar enough, so the Python code is valid for the rest of the assessments.

For further analysis in this work, fixed 10,000 model parameters realizations were created, as well as fixed measurement errors added to the oil rates (using random normal functionality in Python-NumPy), to keep consistency among the different comparisons and sensitivity analysis presented in this work going forward.

8.2.1.2. Validation of non-linear regression in a risk-neutral case

After the new code was validated against the predecessor results, a further estimation was performed using non-linear regression and compared to OLS. In Chapter 7: Selection of the regression method, it has been presented that the XGB has delivered a more accurate fit for NPVs versus oil rates, compared to the fit provided by OLS. Therefore, it was considered that XGB results would be more reliable than the ones provided by OLS. Table 19 shows the LSM solution of the case study with both algorithms.

Estimated value	OLS	XGB
EVWI (MMS)	805.29	807.79
VOI (MMS)	49.31	49.81

Table 19. Comparison of EVWI and VOI when using OLS versus XGB in a neutral risk case.

The EVWOI is \$755.98 million, and the EVWI using OLS and XGB is \$805.29 million and \$807.79 million respectively. This represents a small difference in the VOI, \$49.31 million when using OLS, and \$49.81 million when using XGB. As explained in Chapter 7, the relation between the NPVs and the observables (oil rates) is linear, however, the ENPV may not be linear against the observables due to a combination of optimal alternatives (a mixture of “switch” and “continue” alternatives for future decision points) in regression. This is the reason why a non-linear regression method can lead to a more reliable solution, i.e., more accurate VOI assessment and more optimal decision policy even for a linear case. Having performed this analysis, the VOI using XGB will be the reference for further comparisons. In the same way, risk-neutral cases in this work are treated with XGB going forward.

8.3. CASE STUDY RESULTS

In this section, the results of the case study are presented, where a utility function has been introduced and the LSM algorithm is implemented using non-linear regression, specifically XGB, accounting for exponential relation between the NPVs and the utilities. The new objective function consists of maximizing expected utility (EU), differently to the neutral case where the objective function is to maximize the expected NPV (ENPV). Various risk tolerances are assessed through sensitivity analyses on different aspects.

8.3.1. Introduction of risk attitude and impact on the final monetary value (CE)

The first step of the case study is to introduce risk attitude in the IOR initiation problem. As explained in the workflow, a utility function is used to convert the NPVs to utilities, and the LSM implementation is run using XGB to assess the conditional expected EUs given measured oil rates. The final EU (calculated by averaging the utilities over all the paths - refer to section 8.2: LSM implementation, sub-section 8.2.1: Workflow and modelling in Python) is converted back to a monetary value, certain equivalent (CE), and compared to the expected NPV obtained when the risk attitude is neutral.

For this first assessment, a preliminary risk-averse utility function was used, with a risk tolerance (Rho) of \$200 million. Table 20 shows the CE⁶ results for a risk-neutral and risk-averse case. Suffix WOI in the Table 20 (and going forward in this document) represents “without information, and WI represents “with information”. The risk-averse case is additionally solved using OLS to demonstrate the difference in the results against using XGB regression.

The CEs obtained by a risk-neutral decision-maker differ to the CEs obtained by a risk-averse decision-maker, which is expected due to the nature of the different risk attitudes. The risk-neutral case has a CEWI of MM 807.8, and the risk-averse case has a CEWI of \$699.6 million. The difference of approx. \$107 million can be regarded as the “cost” of being risk-averse, which is defined as “risk premium”. Moreover, the VOI for this risk-averse decision-maker is dramatically different than the one for a neutral person, being \$11.61 million for the risk-averse case and \$49.81 million for the risk-neutral case. A risk-averse person seems to value the use of the information less than a risk-neutral for this specific case study. In the next section, sensitivity analysis of the risk tolerance over the VOI is performed to further understand this behaviour. As for the CEs without information, a risk-neutral would obtain a CE of \$755.98 million

⁶ For risk-neutral cases, the expected NPV (ENPV) is the same as certain equivalent (CE)

and a risk-averse would obtain a CE of \$687.99 million. The different CE results respond to different decisions between both decision-makers and the risk premium (a risk-averse person values an uncertain deal lower than the deal's EV).

On the other hand, opportunity was taken to compare the risk-averse solution solved by OLS and XGB. As per seen in Chapter 7, OLS does not provide an optimal prediction of the EUs over the oil rates when the profile is risk-averse. It can be observed that using a not adequate function for the regression fitting, leads to different results. In this specific test, the VOI is overestimated by the OLS solution, being \$440 thousand larger than the one obtained when solved by XGB. The difference in the VOI estimation by OLS and XGB regressions may lead to different decisions strategies.

Estimated value	Risk-neutral	Risk-averse ($\rho = 200$)	Risk-averse ($\rho = 200$)
	(XGB)	(XGB)	(OLS)
CEWOI (MMS)	755.98	687.99	687.99
CEWI (MMS)	807.79	699.60	700.04
VOI (MMS)	49.81	11.61	12.05

Table 20. Comparison of CEWOI, CEWI and VOI in a risk-neutral case versus a risk-averse case. Risk tolerance = \$200 million.

In all the assessments going forward, XGB regressor is used in the LSM algorithm to solve the IOR problem, given any risk attitude.

8.3.2. Sensitivity analysis of different risk attitudes

In this section, sensitivity analysis of different risk attitudes on various decision aspects is documented. Risk-averse and risk-seeking profiles are considered in the assessments. Figure 29 and Figure 30 show the scatterplots of utilities versus oil rates

at production year 30, for different risk tolerances. Neutral risk attitude is also included as a reference, where the y axis is represented by NPVs.

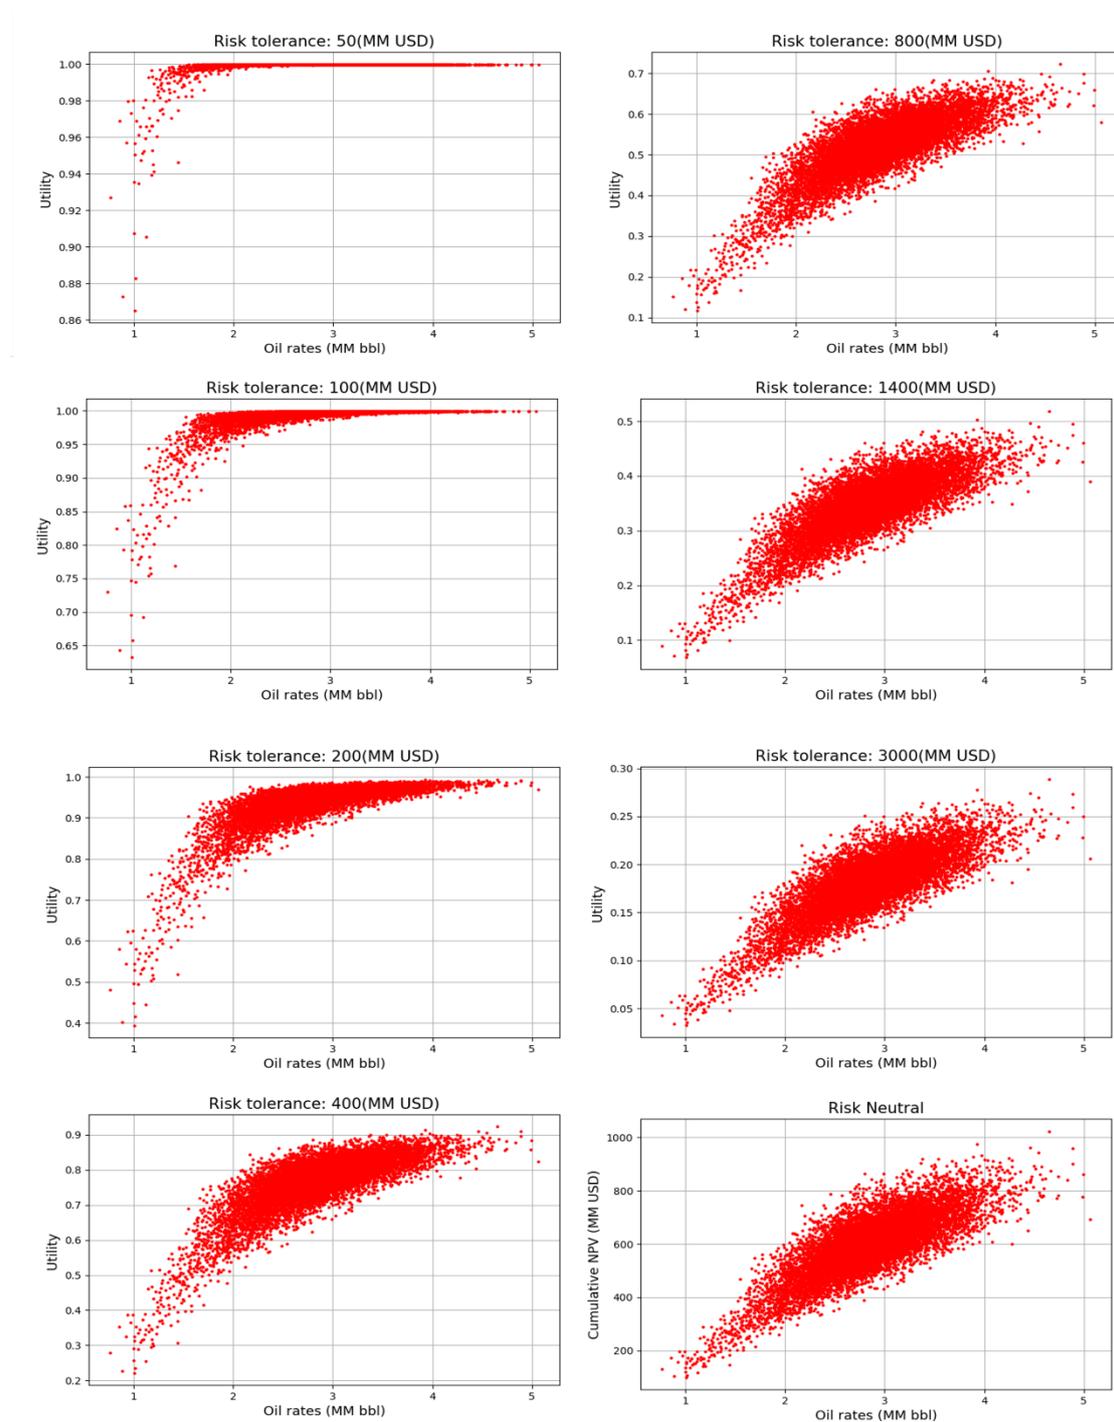


Figure 29. Scatterplots of utilities versus oil rates at year 30 for different risk-aversion tolerances. Risk-neutral NPVs versus oil rates included as reference.

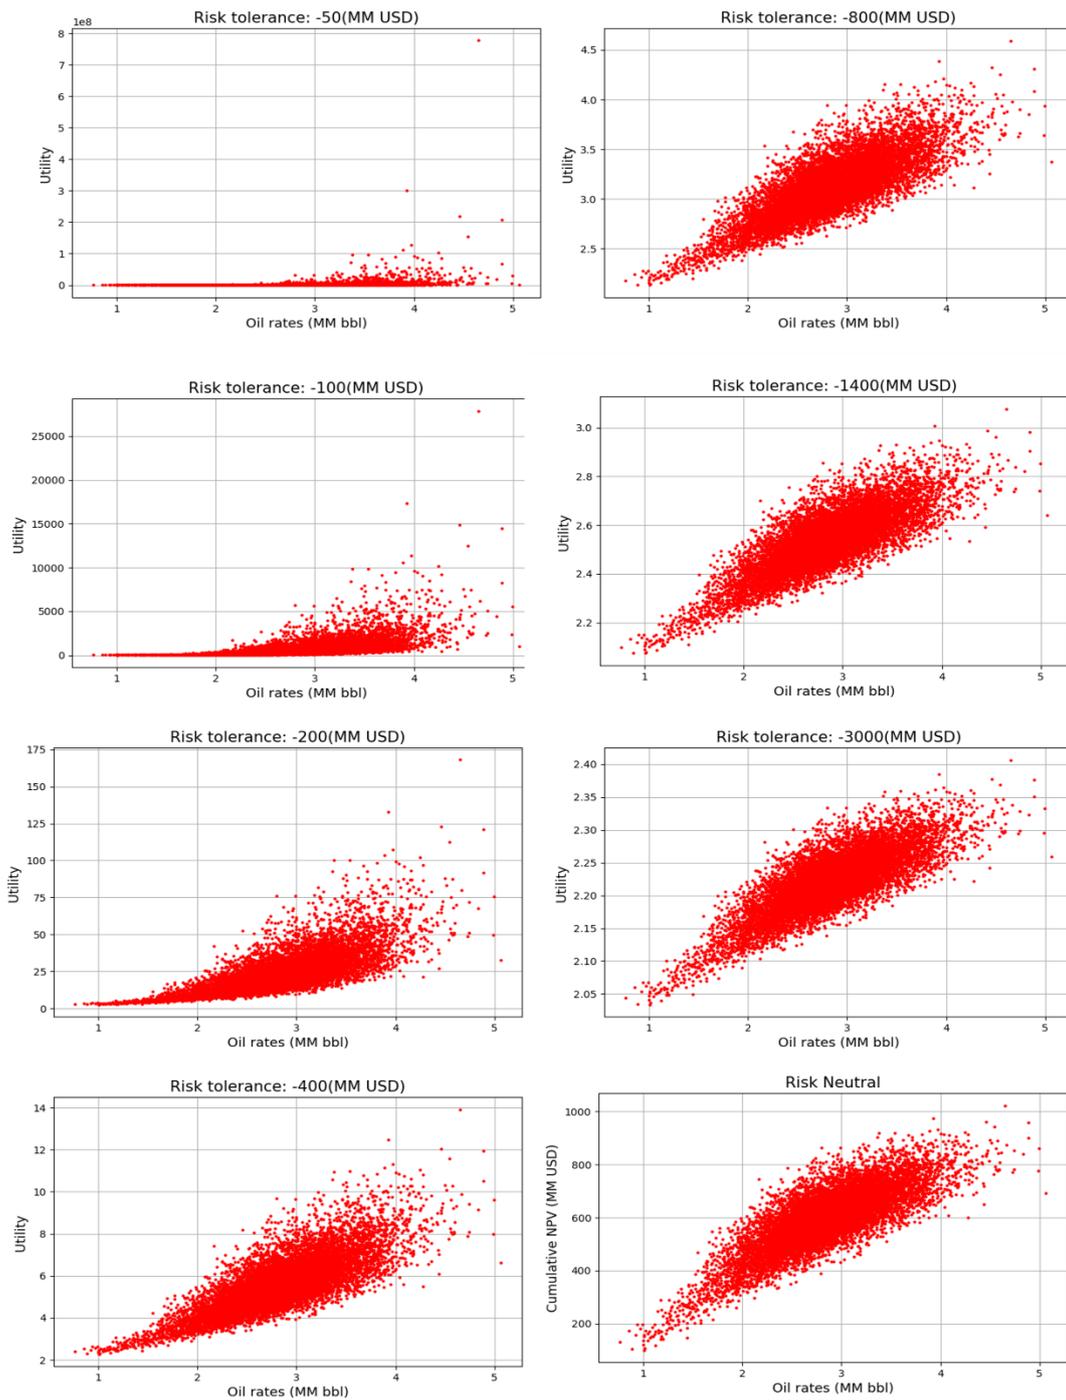


Figure 30. Scatterplots of utilities versus oil rates at year 30 for different risk-seeking tolerances. Risk-neutral NPVs versus oil rates included as reference.

8.3.2.1. LSM utilities validation

Before starting any sensitivity analysis, a validation of the final utilities obtained by the LSM algorithm using the different risk tolerances was performed. Figure 29 shows that for small risk tolerances (more risk-averse) the curve is concave, and progressively becomes more linear as the risk tolerance increases (approaching neutral risk attitude). For the very small risk tolerances such as \$50 million and \$100 million most of the utilities are very close to '1', generating a very narrow spread of the points. This is because for small positive risk tolerances the exponential term of the utility function tends to '0' (refer to utility function in equation 18). These cases might suffer from numerical issues and therefore deliver misleading expected utilities during the LSM algorithm.

Figure 30 shows that for negative risk tolerances closer to zero (more risk-seeking) the curve is convex, and progressively becomes more linear as the risk tolerance decreases (approaching neutral risk attitude). For the negative risk tolerances closest to zero, -\$50 million and -\$100 million, a similar observation to the risk-averse case is noted. In these cases, the exponential argument of the utility function tends to 1, hence, the exponential term of the utility equation (equation 18) results in very large numbers. This might also cause numerical issues during the LSM algorithm.

To validate the cases for which the LSM algorithm is stable and discard the cases of numerical issues, a utility validation was performed following the procedure below, for risk aversion and risk-seeking profiles:

1. For each case of risk tolerance (Rho), convert LSM final utilities to NPVs with their respective Rho, for all realizations.
2. For each case Rho, convert the NPVs back to utilities with an individual Rho.
3. Calculate the expected utility (EU) for each Rho case.
4. Plot the EUs for all Rho cases.

5. Assess the maximum EU from the graph: The maximum EU should be obtained when the Rho use for re-conversion is the same Rho from which the original utilities come from. If this condition is not fulfilled, the solved LSM decision strategy would be sub-optimal for two possible reasons: numerical issue because of small absolute value of Rho or regression error in EU assessment.
6. Repeat steps 2-5 using different Rho for the re-conversion.

Figure 31, Figure 32 and Figure 33 show examples of the results of re-converting the utilities using Rho of \$50 million, \$100 million, and \$200 million respectively. Figure 31 shows that for the case of Rho=50 there is an apparent indication that there is no numerical issue, as the maximum re-converted EU is obtained when using Rho=50, as expected. The same situation is observed for the case of Rho=200 in Figure 33, where the maximum re-converted EU is obtained when using Rho=200. For the case of Rho=100, in Figure 32, the maximum re-converted EU is not obtained when using Rho=100, which indicates there is either a numerical issue or a regression error for this specific sampled realization/path set. This case should be discarded for further analysis, since the LSM solution for rho=100 is sub-optimal for this specific problem and sample setup. For all remaining cases no numerical issue was found (see Appendix C, where all cases are documented).

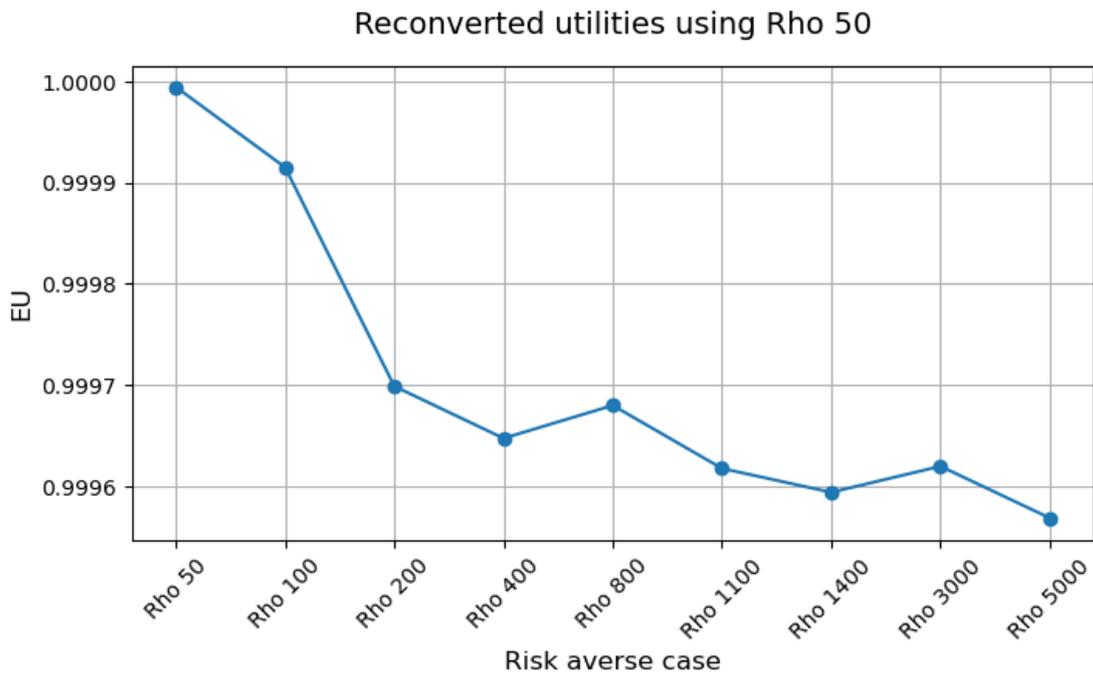


Figure 31: LSM utilities validation: Risk-averse EU re-converted using Rho= \$50 million.

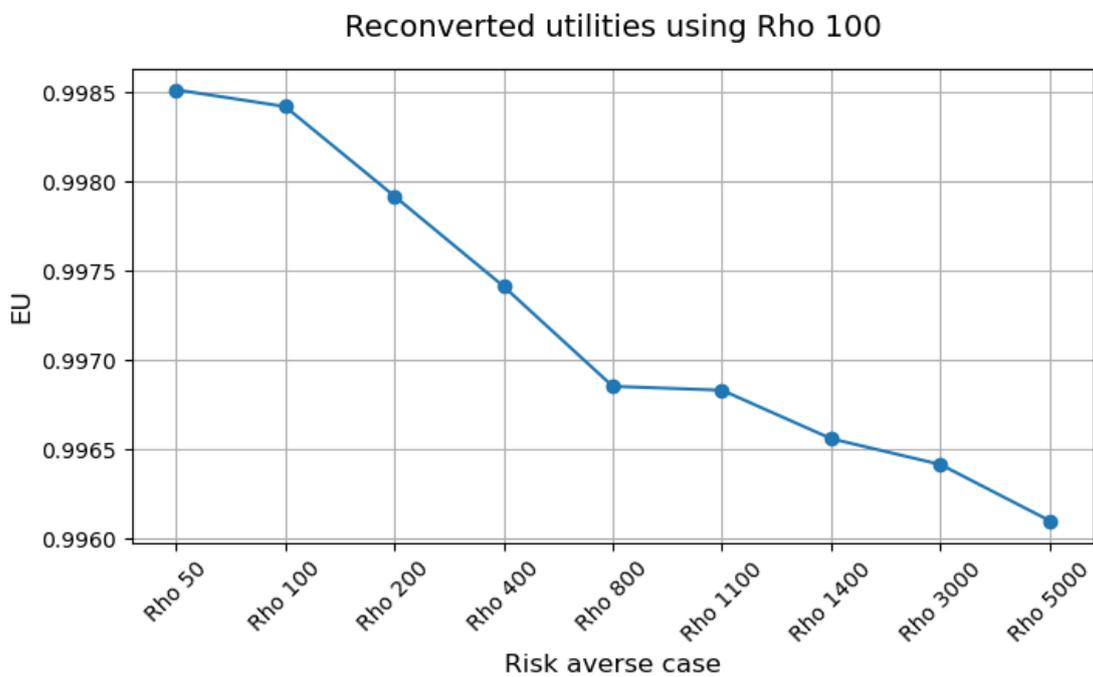


Figure 32: LSM utilities validation: Risk-averse EU re-converted using Rho= \$100 million.

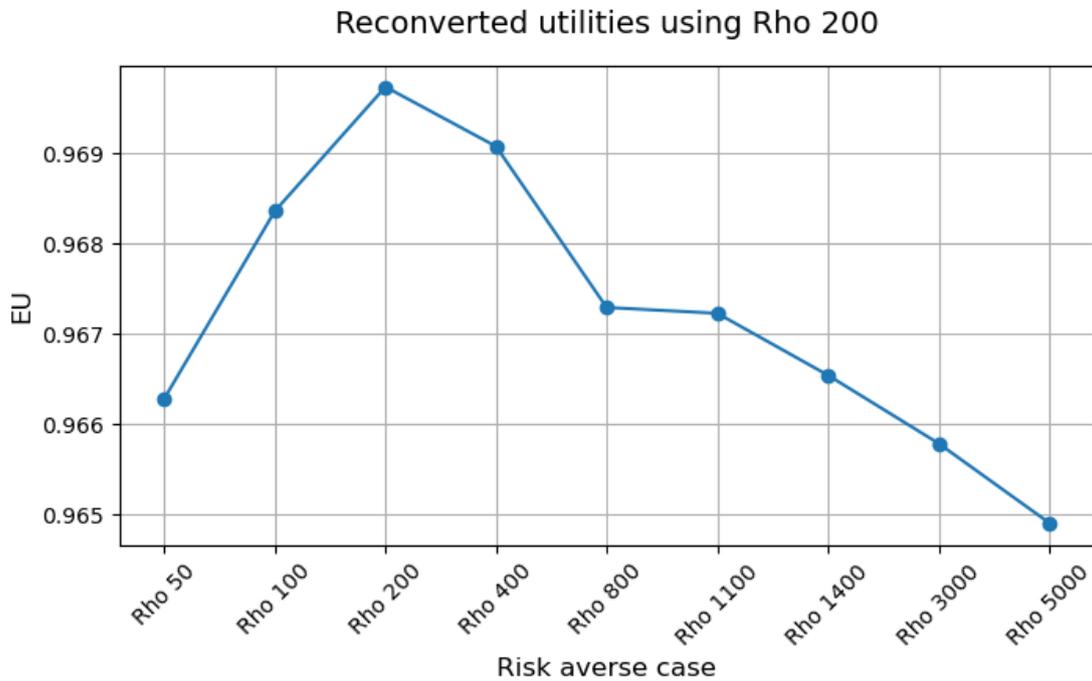


Figure 33: LSM utilities validation: Risk-averse EU re-converted using Rho= \$200 million.

Given that Rho=100 case was considered unstable, there was a suspicion that Rho=50 might be unstable as well, and validation result in Figure 31 might have been a matter of ‘luck’ with the specific set of Monte Carlo realizations used. To confirm this, the maximum LSM utility corresponding to each case where assessed. By definition, the exponential utility function used (equation 18) cannot return a value of ‘1’, which would imply an infinite NPV; hence, if any Rho case happens to have a realization that has a value of ‘1’, it could be considered an unstable case. Table 21 shows the maximum utilities for each case. Risk tolerances of \$50 million and \$100 million contain at least one realization with a final utility of ‘1’. This means these cases are suffering from numerical issue and should be discarded for further analysis. From

this point onwards, risk-averse sensitivity analysis is only done from risk tolerance of \$200 million towards risk-neutrality (Rho=infinity).

Rho	Maximum utility
50	1.000000
100	1.000000
200	0.999990
400	0.999268
800	0.972949
1100	0.908857
1400	0.874431
3000	0.584609
5000	0.440645

Table 21: Maximum LSM utility among all realizations.

An equivalent analysis was performed on the risk-seeking cases. Figure 33 to Figure 38 show some examples for re-conversion of utilities using Rhos of -\$50 million, -\$100 million, -\$200 million, -\$400 million and -\$800 million respectively. For the cases of Rho=-200 and Rho=-400, Figure 36 and Figure 37, the maximum EU is not found when using these values for the re-conversion; hence these cases imply numerical issues and can be discarded. For the case of Rho=-800, Figure 38, the maximum EU is obtained when using the original respective Rho for re-conversion, as expected, and can be considered a reliable case. This is also the case of negative Rhos further from zero, where no numerical issue was found (see Appendix C, where all cases are documented).

For the cases of Rho=-50 and Rho=-100 the graph looks exactly the same, and the results are not conclusive, showing the same EU for the first four Rhos, which could indicate a potential numerical issue. To understand this further, the final decisions for

primary recovery lifetime were plot as histograms for the potential unstable cases: Rhos of -\$50 million, -\$100 million, -\$200 million and -\$400 million respectively; this is shown in Figure 39. The histogram shows that all the decisions from all realizations are ‘0’ years of primary recovery, and for all Rhos. This ‘constant’ behaviour is not expected when using MCS and represents another indication that Rhos closer to zero suffer from numerical issues, delivering unreliable LSM results. From this point onwards, risk-seeking sensitivity analysis is only done from risk-neutrality (Rho=infinity) up to Rho=-800 and discards Rhos closer to zero.

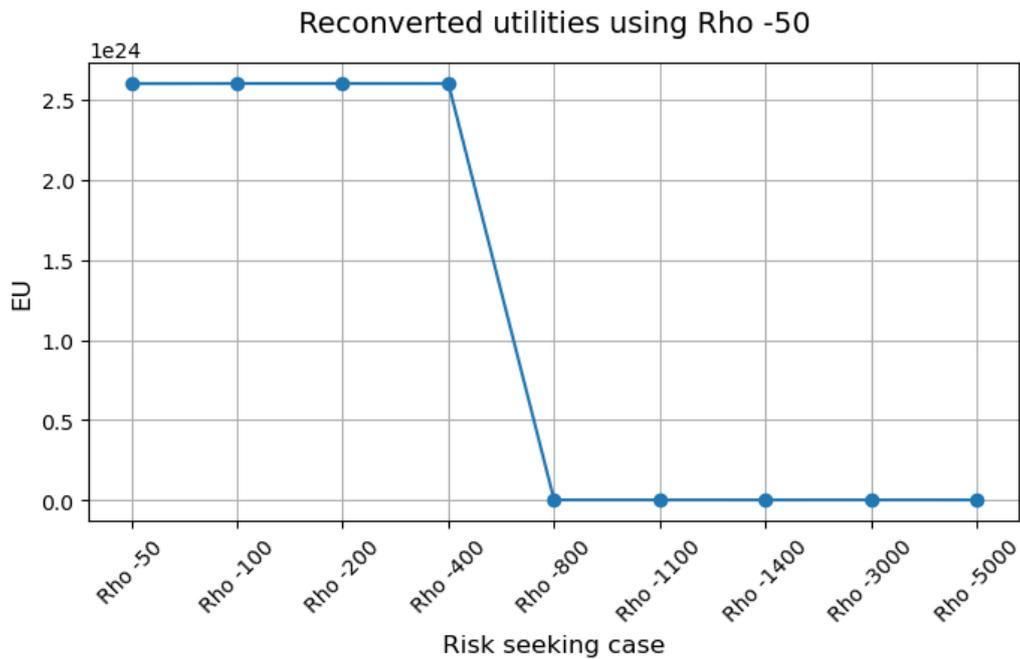


Figure 34: LSM utilities validation: Risk-seeking EU re-converted using Rho= -\$50 million.

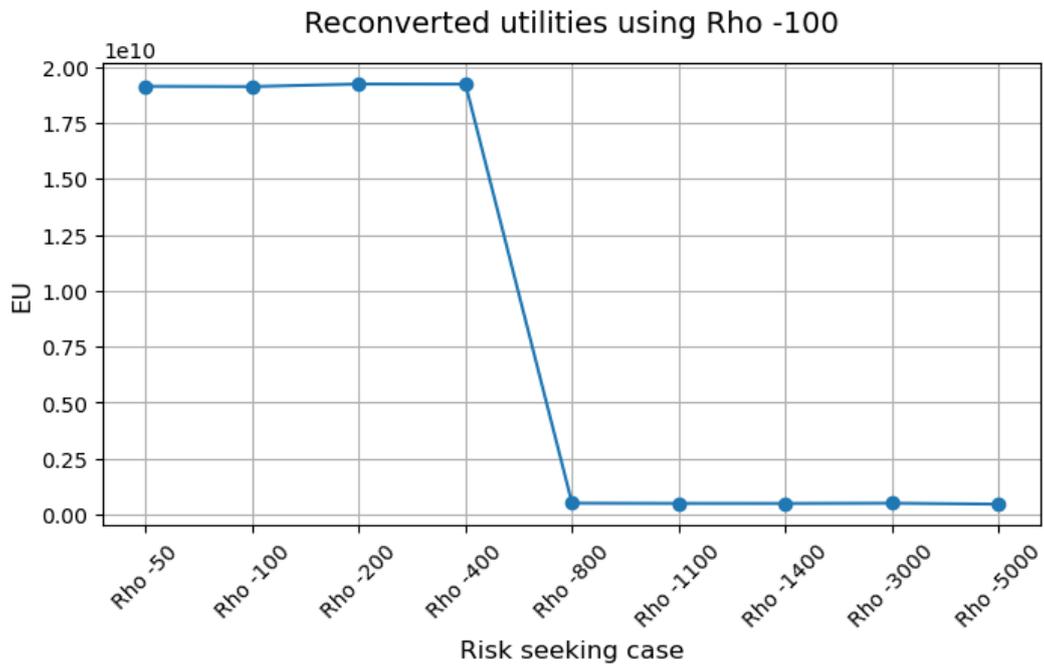


Figure 35: LSM utilities validation: Risk-seeking EU re-converted using Rho= -\$100 million.

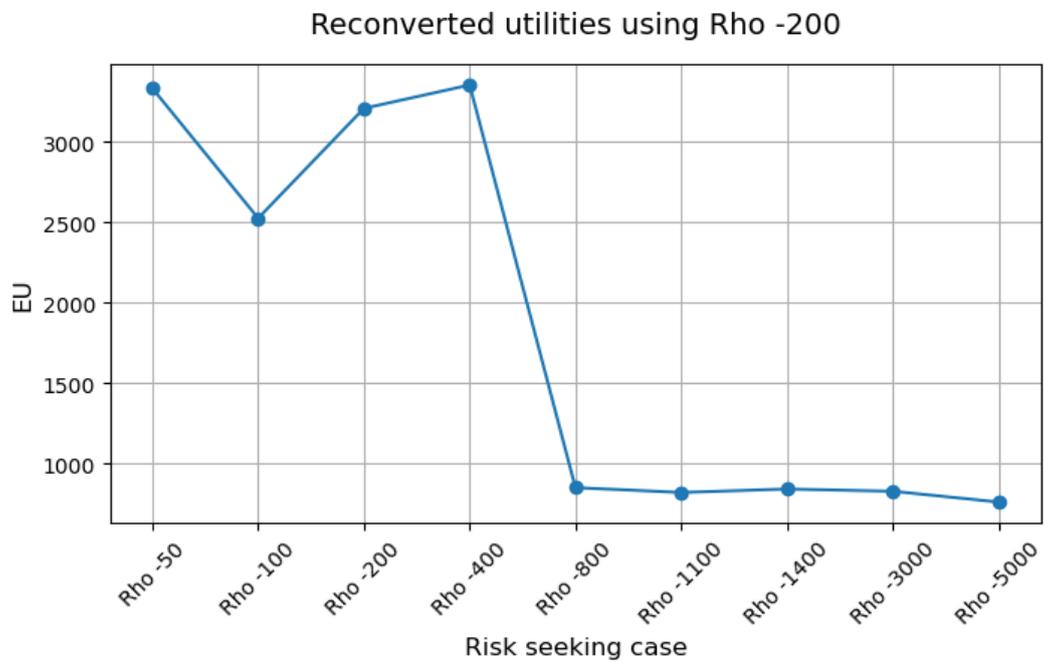


Figure 36: LSM utilities validation: Risk-seeking EU re-converted using Rho= -\$200 million.

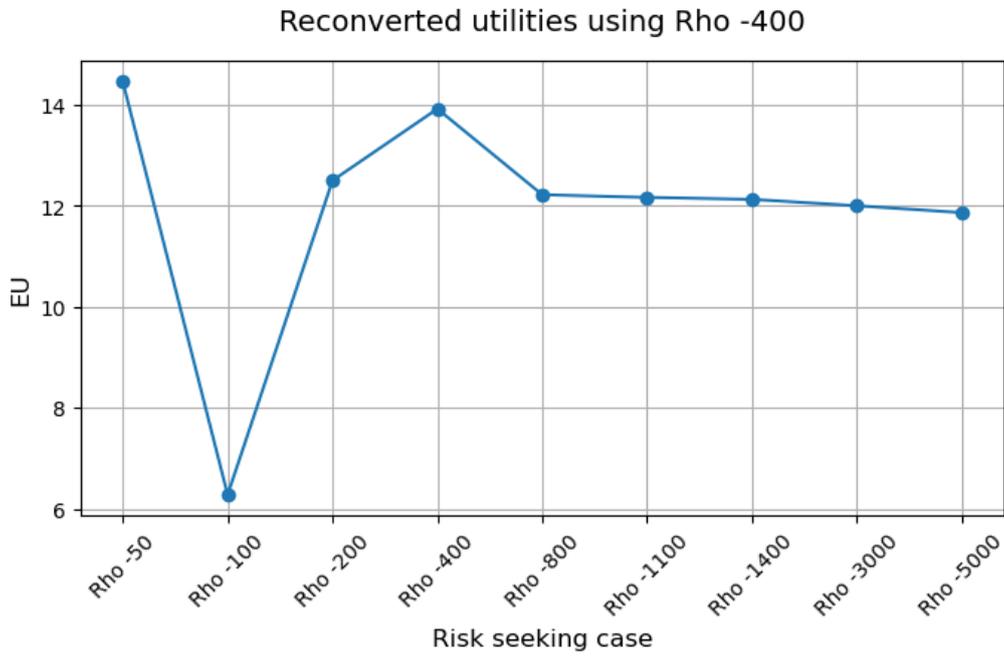


Figure 37: LSM utilities validation: Risk-seeking EU re-converted using Rho= -\$400 million.

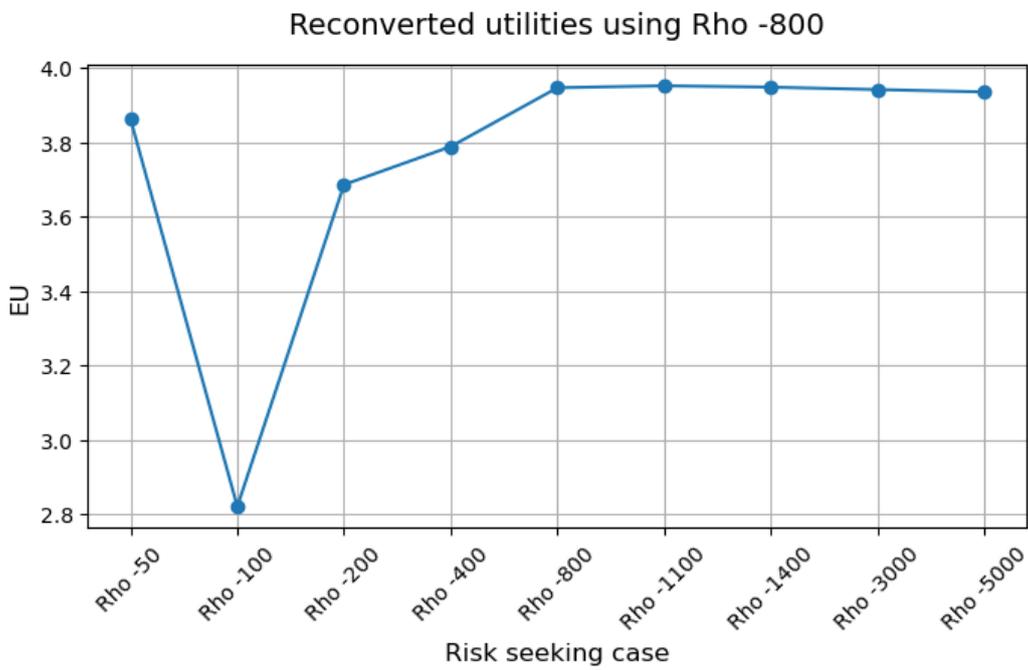


Figure 38: LSM utilities validation: Risk-seeking EU re-converted using Rho= -\$800 million.

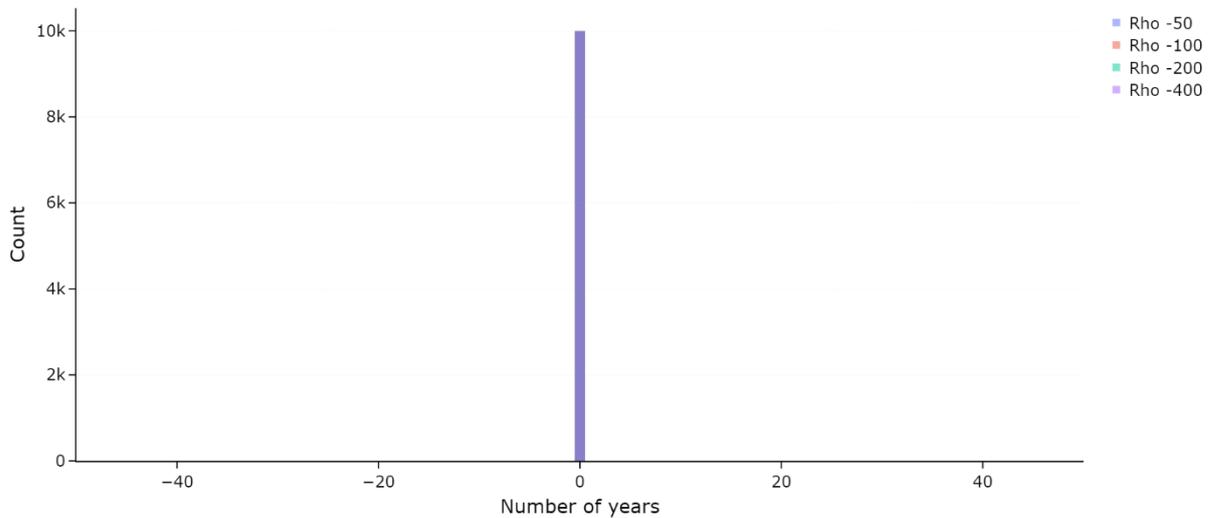


Figure 39: Number of years of primary recovery estimated by risk-seeking decision-makers.

8.3.2.2. Impact of risk attitude on VOI

A sensitivity analysis on the VOI for different risk tolerances was performed and presented in Figure 40 and Figure 41, for risk-averse and risk-seeking profiles respectively. The risk-averse graph indicates that the larger the risk tolerance, the larger the VOI, increasing from \$11.61 million for Rho 200, to \$47.07 million for Rho 5000. Infinite Rho represents a risk-neutral case with VOI of \$47.15 million.

This increasing trend might be interpreted as follows: the less risk-averse a decision-maker is, the more they will value the information, and their decisions might be highly influenced by the information available. This interpretation will be complemented later when sensitivity analysis on the decisions is performed.

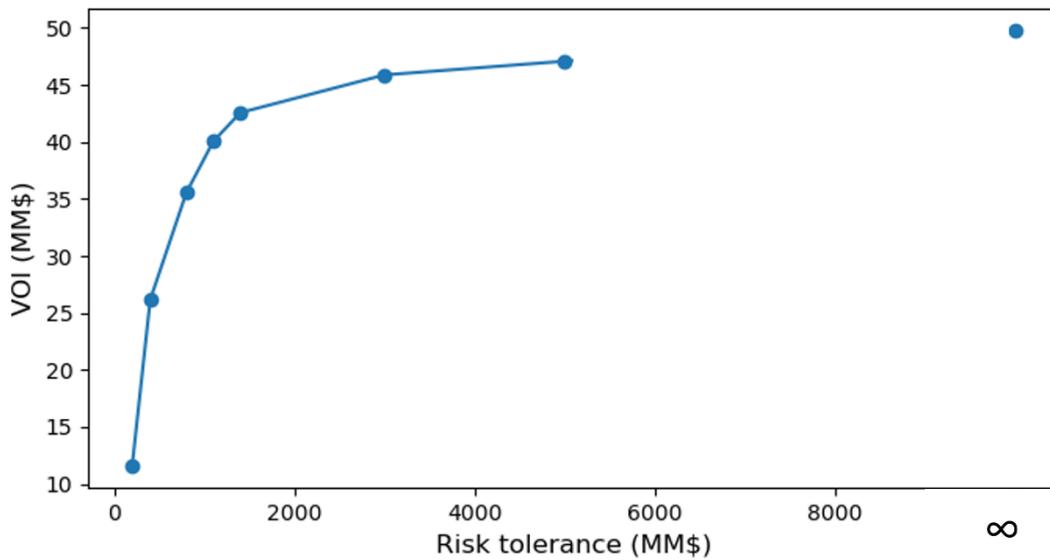


Figure 40. Sensitivity analysis of different risk tolerances on the VOI, for risk-averse attitude and neutrality.

For risk-seeking decision-makers, the value of information seems to increase from Rho -800 to Rho -1400, going from VOI of \$22.97 million to VOI \$53.11 million respectively. From that point to Rho -5000, the VOI reaches around \$50 million with a steadier trend. The VOI settles at neutrality with \$47.04 million. This means that the impact of the risk attitude on the VOI is larger for negative Rhos that are closer to zero, down to Rho around -1500. Between Rho around -3000 towards neutrality the risk attitude has less impact on VOI.

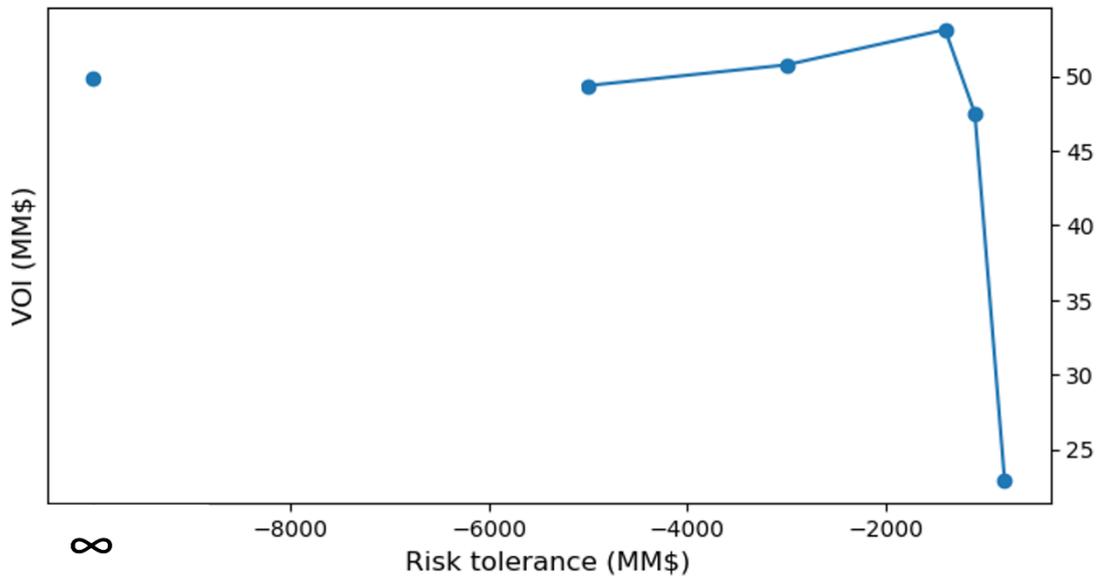


Figure 41. Sensitivity analysis of different risk tolerances on the VOI, for risk-seeking attitude and neutrality.

Figure 42 and Figure 43 show the distribution of the final NPVs for the 10,000 Monte Carlo realizations after resolving the LSM algorithm for different risk tolerances, and expressed through boxplots, to assess the spread. In the risk-averse case, the less risk-averse the decision-maker is, the larger the spread is and the larger expected NPV, which indicates that the use of the utility function accounting for risk-averse behaviour performs as expected – a risk-averse decision-maker prefers less uncertainty and is willing to give up some value (i.e., risk premium) in ENPV for less uncertainty. This is also true for the risk-seeking case, where the more risk-seeking decision-maker, the slightly larger the spread of the NPVs is, although the difference between the different cases is not so large as seen as in the risk aversion cases.

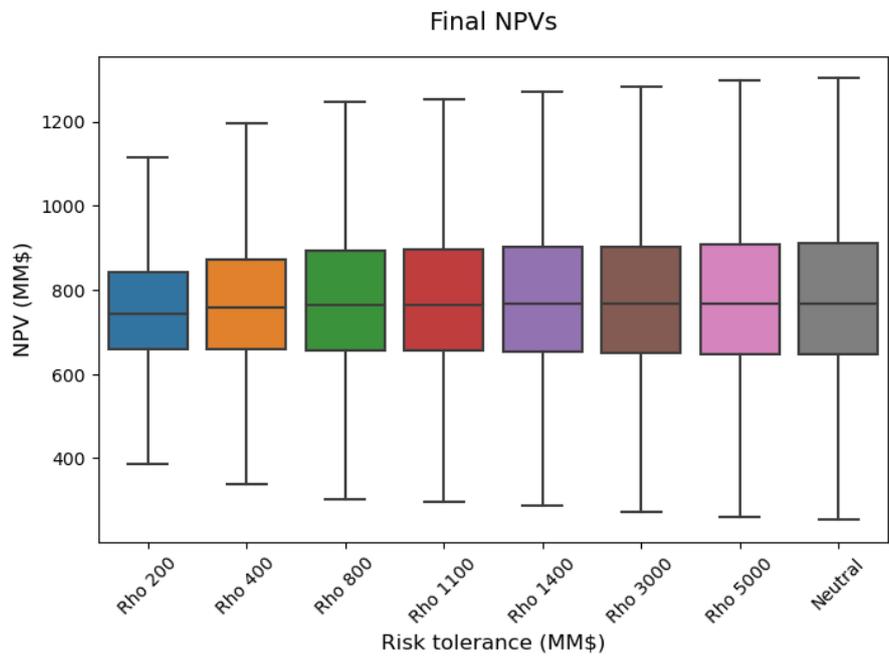


Figure 42. Boxplot of the final NPVs for different risk tolerances, for risk-averse attitude.

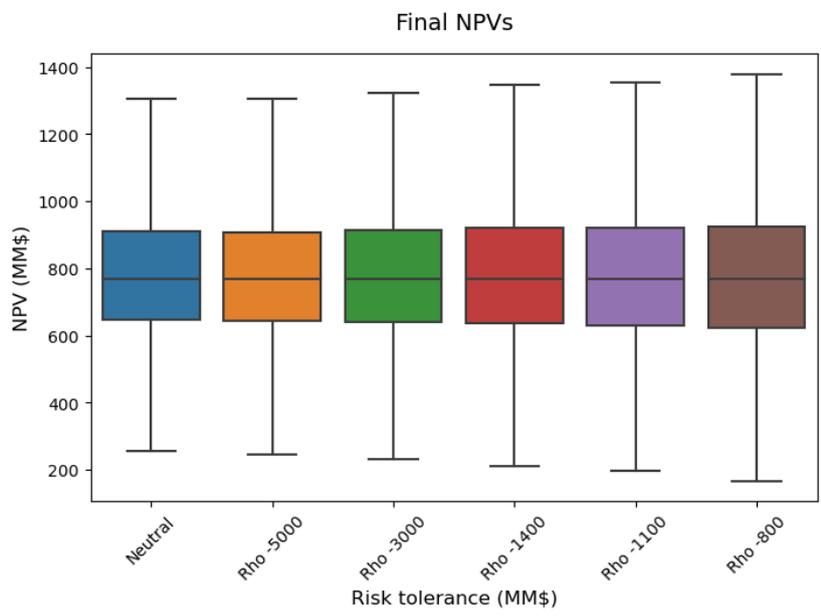


Figure 43. Boxplot of the final NPVs for different risk tolerances, for risk-seeking attitude.

8.3.2.3. Impact of risk attitude on decisions

After assessing the impact of different risk attitudes on the VOI and NPV distribution, the impact on the decision strategy for primary and secondary recovery lifetimes are also assessed and discussed in this section. The lifetime of the primary recovery represents the time at which the IOR (improved oil recovery) phase initiates, hence, this is the switch time between primary and secondary recovery. The addition of the primary and the secondary recovery lengths totalizes the overall lifetime of the field.

Figure 44 shows the difference in years of primary lifetime between a risk-neutral and a risk-averse decision-maker, for different Rhos. The difference is calculated as the primary lifetime of the risk-neutral case minus the primary lifetime of a risk-averse case. The overlapping histograms show a dominance towards the negative differences. Most of them peak between 0 and -2 years of difference against the risk-neutral case, but the more risk-averse, the more asymmetrical the distribution is, with increasing skewness towards the negative differences. $\text{Rho} = 8000$ gives a relatively symmetrical histogram; Rhos between 3000 and 1100 cover some dominance towards -3 years difference against neutral, and Rhos 800 to 200 have dominance down to -8 years of difference. Because the difference is calculated as the primary lifetime of the risk-neutral case minus the primary lifetime of a risk-averse case, a negative difference means that a risk-averse decision-maker prefers more years for primary production than a risk-neutral decision-maker. A risk-averse decision-maker prefers to collect more data (longer primary recovery) to reduce uncertainty before deciding to switch to secondary recovery. The interpretation of the observed behaviour also relies on the capital expenditure (CAPEX), which is incurred when switching from primary to secondary production, as explained in section 8.1.2: Economic model. If the extra recovery from secondary recovery cannot cover the extra CAPEX, there will be extra loss, which makes the NPV distribution extend to lower NPVs. On the other hand, if the extra recovery from secondary recovery can cover the extra CAPEX, there will be extra gain,

which makes the NPV distribution extend to higher NPVs. Accounting for the two scenarios, the spread of the NPV distribution for switching to secondary recovery is larger than the spread of the NPV distribution for continuing with primary recovery; that is, there will be larger uncertainty if switching to secondary recovery, which a risk-averse decision-maker does not prefer unless the increase in expected NPV from primary recovery to secondary recovery can compensate the decision-maker's risk premium. This then, results in a risk-averse decision-maker preferring to delay the switch to secondary recovery in comparison with neutral decision-makers. As shown in Figure 44, the more risk-averse (towards $Rho=200$), the larger negative difference against neutral (histograms displace to the left in the figure).

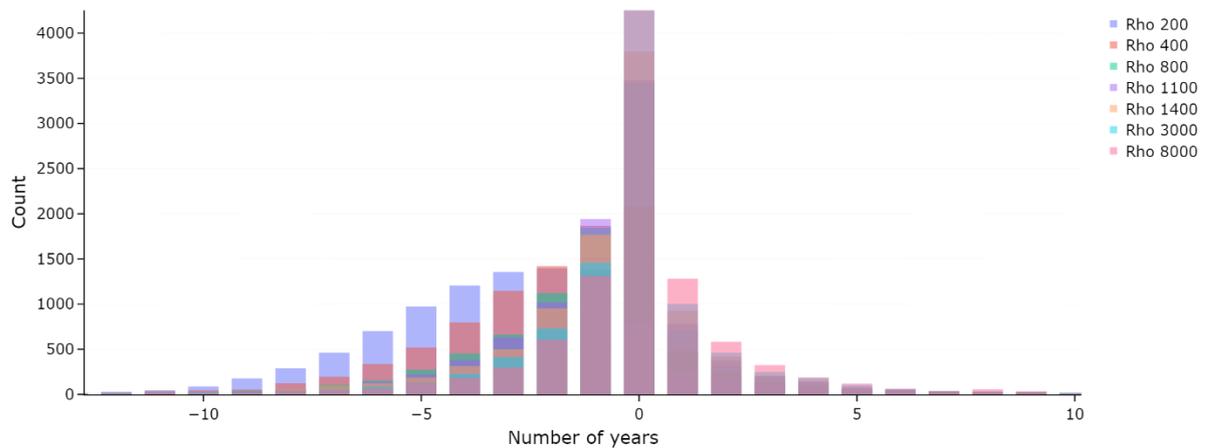


Figure 44. Difference in years of primary recovery lifetime between a risk-neutral and risk-averse decision-maker.

Figure 45 shows the difference in years of secondary lifetime between a risk-neutral and risk-seeking decision-maker, for different Rhos. The difference is calculated as the secondary lifetime of the risk-neutral case minus the secondary lifetime of a risk-averse case. In this case the dominance of the histograms displaces to positive side of the plot. Again, the least risk-averse, $Rho 8000$ has the most

symmetrical decision distribution, but Rhos between 3000 and 1100, have dominances towards 5 years of difference; the histograms for the most risk-averse profiles, between Rho 800 and 200, have dominances towards 9 years of differences against neutral. Being the order of the subtraction 'neutral minus risk-averse', this means that the neutral prefers longer secondary recovery than the risk-averse decision-maker. A risk-averse decision-maker prefers to stop the secondary recovery sooner. As shown in the section 8.1.2: Economic model, the operational expenditure (OPEX) is larger for secondary recovery than primary recovery. If the late-stage secondary recovery is low, the NPV is reduced by the difference between OPEX and cash inflow; however, if the late-stage secondary recovery is high, the NPV is increased by the difference between cash inflow and OPEX. Thus, a risk-averse decision-maker does not prefer the uncertainty associated with the late-stage production even though the EVs of the late-stage cash inflows are still positive. In other words, because risk-averse's $CE < EMV$ (expected monetary value) of net cash flow, the risk-averse's CE of net cash flow can be negative even though the EMV of the net cash flow is positive, which indicates that the risk-averse decision-maker prefers to terminate production whilst a risk-neutral decision-maker prefers to continue production. As observed in Figure 45, the more risk-averse (towards $Rho=200$), the larger positive difference against neutral (histograms displace to the right in the figure).

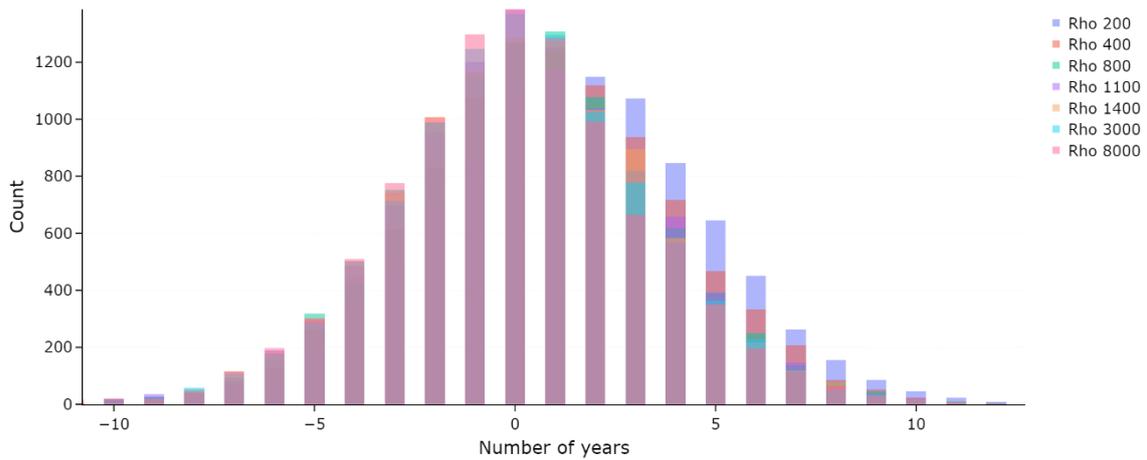


Figure 45. Difference in years of secondary recovery lifetime between a risk-neutral and risk-averse decision-maker.

The equivalent displays are produced now for a neutral decision-maker against a risk-seeking one. Figure 46 shows the difference in years of primary lifetime between a risk-neutral and risk-averse decision-maker. Now that the histogram dominances displaced to the right of the plot. The subtraction is in order ‘risk-neutral minus risk-seeking’, and the dominant results are positive numbers. The interpretation of this is consistent to what observed in Figure 44: the less risk-seeking the decision-maker is, the longer the estimated primary recovery length. Risk-seeking decision-makers prefer to switch to secondary recovery earlier than risk-neutral decision-makers. Risk-seeking decision-makers with Rhos -8000 and -3000 have more dominance between 0 and 2 years of difference against neutral, but Rho -1440 to -800 have dominances that extend up to 6-7 years of difference. The more risk-seeking (towards Rho= -800), the larger positive difference against neutral (histograms displace to the right in the figure).

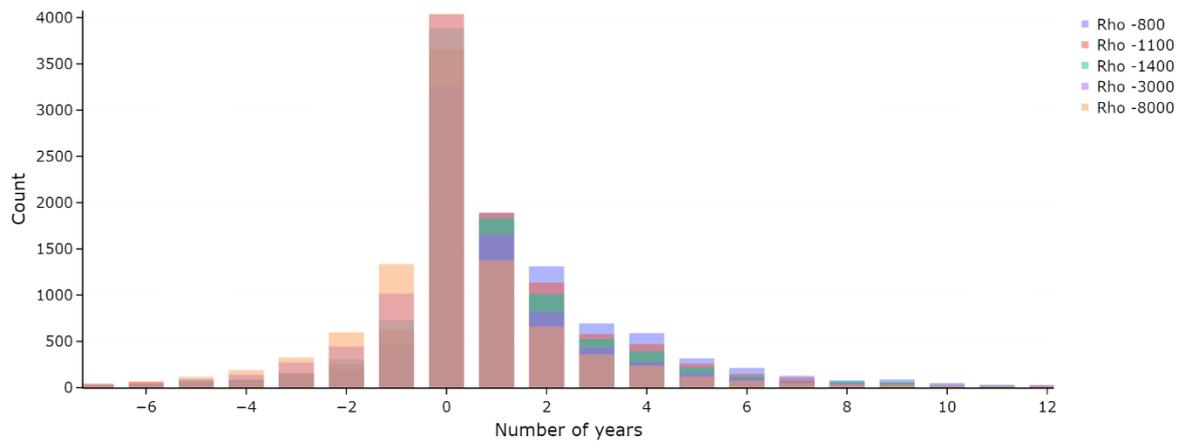


Figure 46. Difference in years of primary recovery lifetime between a risk-neutral and risk-seeking decision-maker.

Histograms in Figure 47 show the differences between risk-neutral and risk-seeking decision-makers for secondary recovery length. There is a slightly stronger dominance on the negative side of the histograms, representing a longer secondary recovery preferred by risk-seeking with respect to risk-neutral decision-makers. Consistent to the analysis of Figure 45, the less risk-seeking a decision-maker is, the earlier they prefer to stop secondary production, attributed to the NPV uncertainty associated with the late-stage secondary production, as explained before. Risk-seeking decision-makers with Rhos -8000 and -3000 have a relatively symmetrical histograms peaking at 0 years difference, although the spread extends ± 10 years. Risk-seeking decision-makers with Rho towards -1100 and -800 have histograms with some dominance down to -5 and -6 years. The more risk-seeking (towards Rho = -800), the larger negative difference against neutral (histograms displace to the left in the figure).

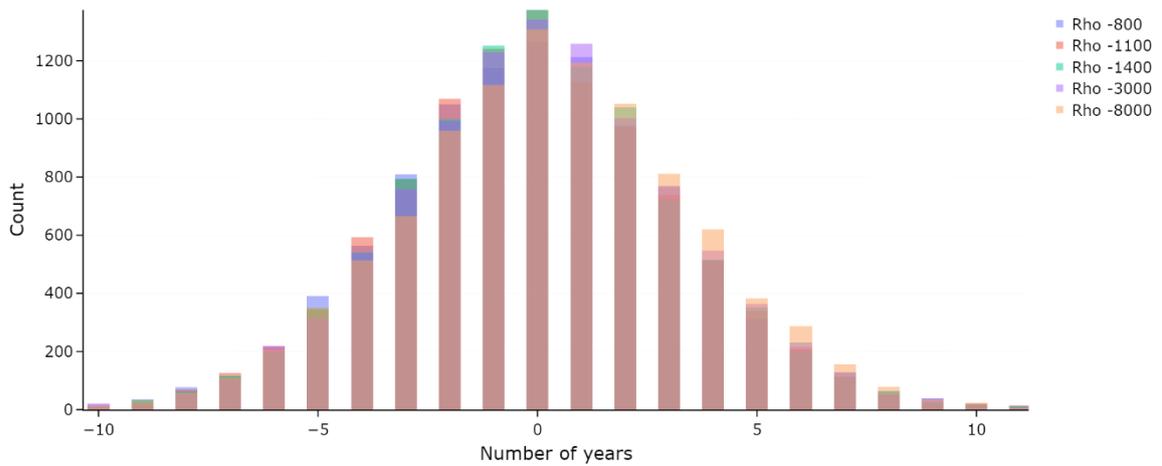


Figure 47. Difference in years of secondary recovery lifetime between a risk-neutral and risk-seeking decision-maker.

8.3.2.4. Impact of information on decisions

After assessing the impact of the risk attitude in the field lifetime decisions, sensitivity analysis was performed to understand the impact of the information on the decisions, for different risk profiles.

Figure 48 and Figure 49 show respectively the histograms of the decisions with information for the primary and secondary production recovery lengths. Rhos of \$200 million and \$800 million are assessed, as well as risk-neutral. The average decision with information is plot as a continuous vertical line, on the same colour as its corresponding histogram. The decision without information is plot as a dashed line.

Figure 48 confirms previous interpretation that a risk-averse decision-maker prefers longer primary recovery compared to a neutral decision-maker. The figure shows that the more risk-averse, the longer the primary recovery is preferred, hence the histogram shifts to the right of the plot. It can be observed that risk-neutral peaks at 1 year of primary recovery; Rho 800 has an irregular distribution with peaks at 6 and 7 years; the most risk-averse decision-maker, Rho 200, peaks at 11 years of primary recovery.

Assessing the difference between the decisions with and without information (the gap between the solid line and dashed line for each Rho value), it seems that the less risk-averse the decision-maker is, the larger this difference is. This impression is later assessed in more detail in Figure 50.

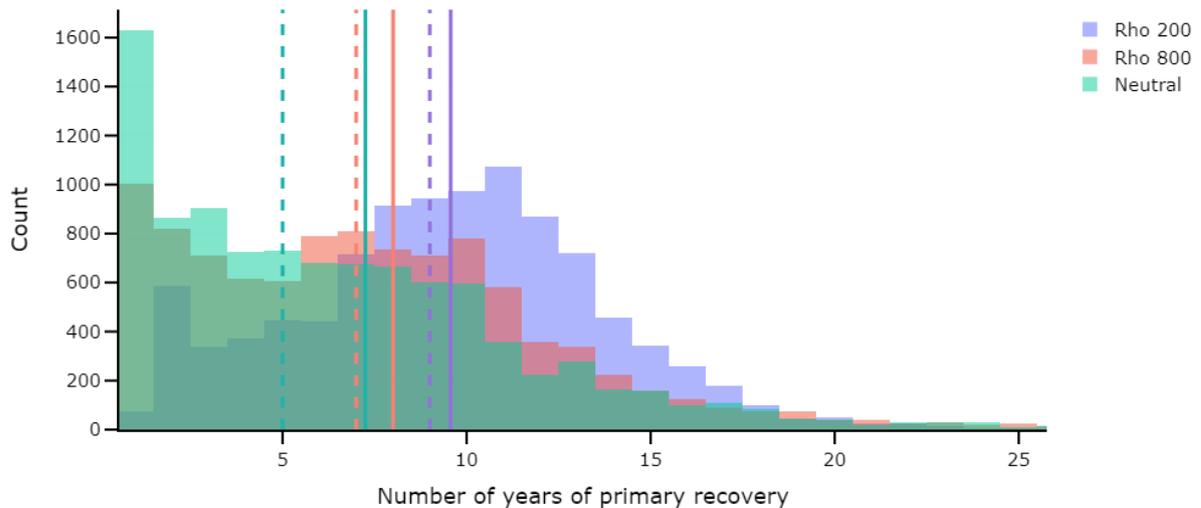


Figure 48. Decisions with and without information of primary recovery lifetime for risk-averse and neutral attitudes. Continues lines: average decision with information. Dotted lines: decision without information.

Figure 49 also agrees with another previous interpretation, that a risk-averse decision-maker will estimate shorter secondary recovery lifetime compared to a risk-neutral decision-maker. The histograms dominances displace to the left as the decision-maker is more risk-averse. Risk-neutral peaks at 11 years of secondary recovery, with relatively symmetrical distribution; Rho 800 also peaks at 10 and 11 but it can be observed that has a tale towards the left compared to risk-neutral; the most risk-averse case, Rho 200, peaks at 9 years, with a longer tale to the left. Note that in this case Rho 200 and Rho 800 have the same decision without information, at 14 years of secondary recovery lifetime (– Rho 800 is overlapping Rho 200 in the graph).

Similar to the primary recovery case, there is an apparent indication that the neutral case's decisions would be more impacted by the information when compared to risk-averse's decisions: the difference between the decision with and without information is larger than the ones for the risk-averse cases. This will be better illustrated in Figure 51.

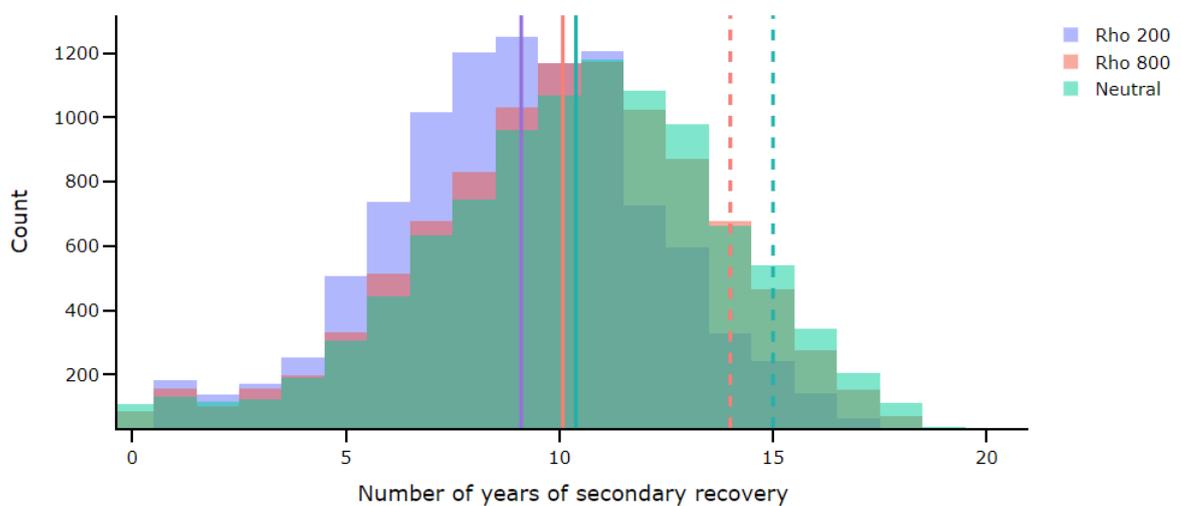


Figure 49. Decisions with and without information of secondary recovery lifetime for risk-averse and neutral attitudes. Continues lines: average decision with information. Dotted lines: decision without information.

Figure 50 and Figure 51 show the difference in years of the average decision with information and the decision without information, respectively for primary and secondary recovery lifetimes, and including all the previously studied risk-averse profiles. Figure 50 shows an overall increasing difference as Rho increases. Starting from 0.6 years difference at Rho 200 to 2.5 years difference at Rh0 3000. This confirms preliminary conclusion drawn in Figure 48. When correlating this observation to the analysis done on Figure 40 - sensitivity analysis on VOI, and Figure 44 - different in

decisions neutral and risk-averse, the interpretation becomes more complete: it has been demonstrated that a very risk-averse decision-maker will delay the switch to secondary recovery, and will extend the primary recovery lifetime; the more risk-averse, the information (i.e., production rates measured during primary recovery) becomes less material (i.e., the information has smaller impact on the switch decisions). This explains why the sensitivity analysis on VOI shows that the more risk-seeking, the more valuable the information becomes. Information is more valuable when it is more material to the switch decisions.

For the secondary recovery case, Figure 51, the trend in the difference in decisions is not monotonic, although the magnitudes are small. At Rho 200 the difference is -4.8 years; then this decreases to -4 years at Rho 800 and -3.9 years at Rho 1100. After this the difference increases towards Rho 1400, with -4.8 years, and -4.5 years difference towards risk-neutral.

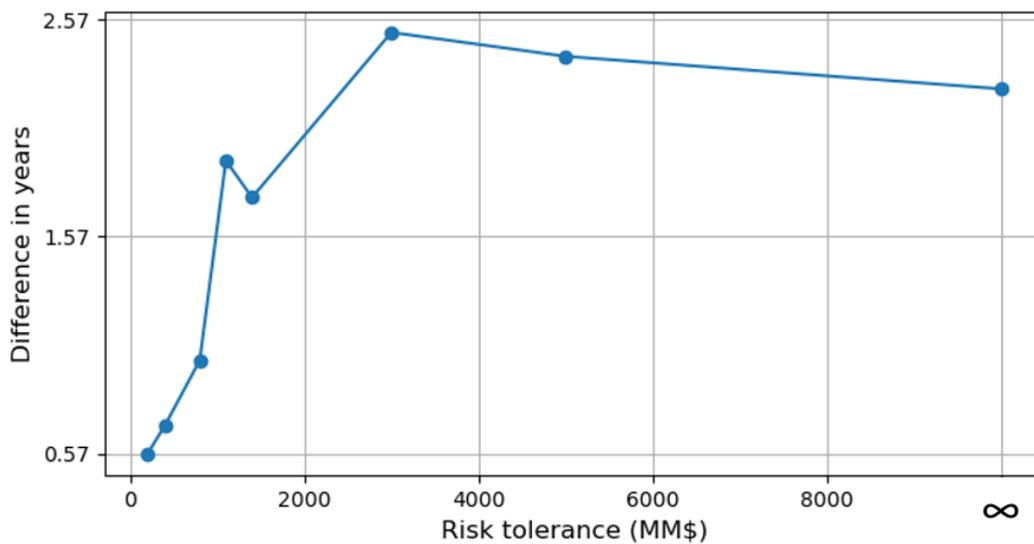


Figure 50. Difference in average decision with information and without information of primary recovery lifetime, for risk-averse and neutral attitudes.

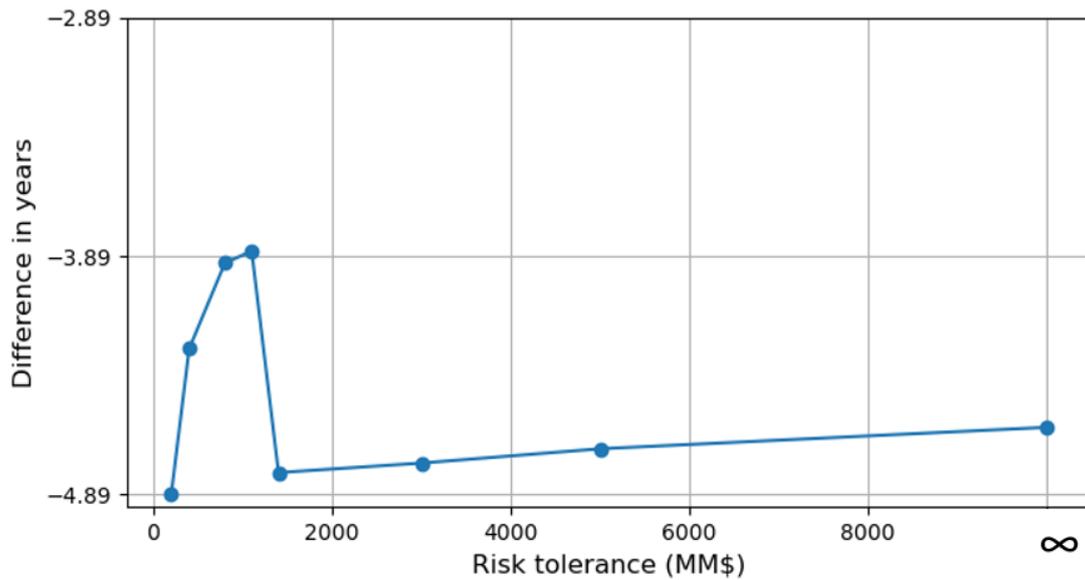


Figure 51. Difference in average decision with information and without information of secondary recovery lifetime, for risk-averse and neutral attitudes.

To understand further this behaviour, boxplots of the difference's spreads were assessed. Figure 52 and Figure 53 show boxplots for the difference in lifetime decisions with information and without information of the respective recovery phases; for each ρ , the spread represents the variation of the lifetime decisions of the 10,000 Monte Carlo samples. The boxplots show that the median (P50) of the differences and the respective spreads do not show any specific trend along the different risk attitudes.

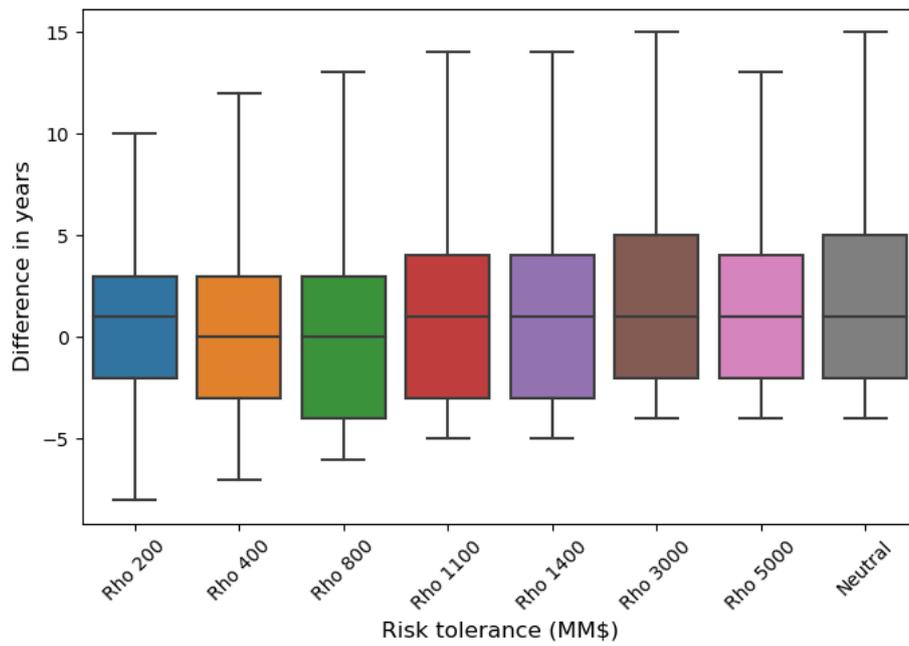


Figure 52. Boxplots of difference in decisions with information and without information of primary recovery lifetime, for risk-averse and neutral attitudes.

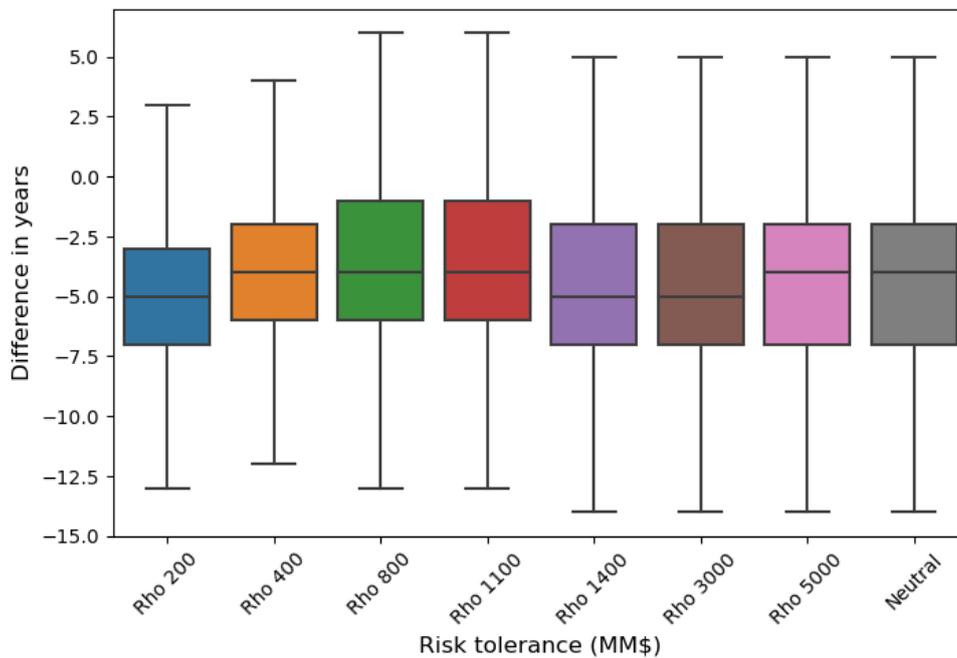


Figure 53. Boxplots of difference in decisions with information and without information of secondary recovery lifetime, for risk-averse and neutral attitudes.

To complement this assessment, a combined bar chart showing the difference in average decision with information and without information for primary and secondary recovery lifetimes is illustrated in Figure 54. It can be observed that this combined chart does not exhibit any evident trend either, even if adding the primary and secondary lengths together (total length of the blue plus the orange bar for each Rho). It is thought that this might respond to the fact that the expected utilities are a combined result of primary and secondary recovery lifetimes. Although not changing monotonically, the differences in the decisions with and without information for the primary recovery are all positive, meaning that the average decision with information is for longer primary recovery phase; as for the secondary recovery the differences are all negative, indicating that the decisions with information are for shorter secondary recovery, for risk-averse cases.

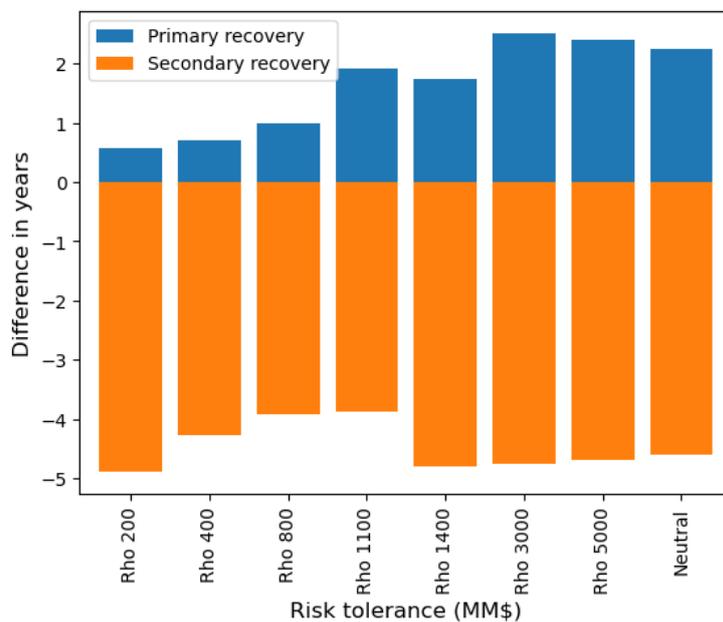


Figure 54. Difference in average decision with information and without information of primary and secondary recovery lifetimes, for risk-averse and neutral attitudes.

Same assessments are now presented for the risk-seeking cases below. Figure 55 and Figure 56 show the histograms of the decisions with information for the primary and secondary production recovery lengths, showcasing two risk-seeking levels (Rhos of -\$800 million and -\$5000 million) and risk-neutral. Average decision with information is plot as a continuous vertical line, and decision without information as a dashed line. For the primary recovery phase, Figure 55, it can be noted that the more risk-seeking the decision-maker is, towards risk tolerance of -800, the histogram displaces to the left, confirming that more risk-seeking decision-makers will chose to have shorter primary recovery phases, consistent with that has been interpreted before. Risk tolerances of -800 and -5000 peak at one year recovery, while risk-neutral has more dominant tale towards the right, suggesting longer primary recovery and delaying the IOR initiation time. Consistent to what concluded in the risk-averse analysis, the primary recovery decisions with and without information appear to differ less as the decision-maker is less risk-seeking (towards risk-neutral).

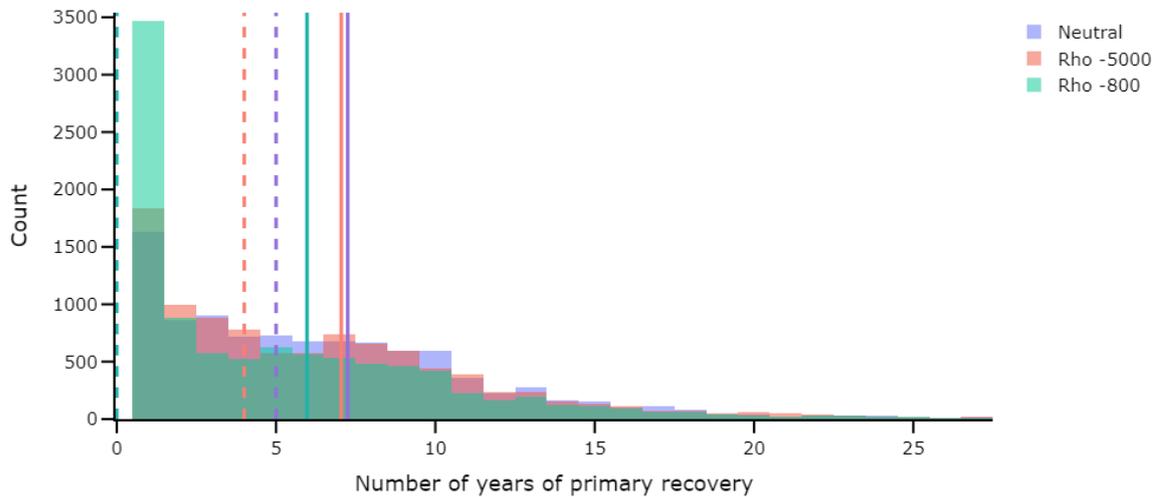


Figure 55. Decisions with and without information of primary recovery lifetime for risk-neutral and risk-seeking attitudes. Continues lines: average decision with information. Dotted lines: decision without information.

For the secondary recovery, shown in Figure 56, the distribution of decisions is slightly more uniform through the three risk profiles presented. Previous analysis have suggested that risk-seeking decision-makers will generally extend the secondary lifetime when compared to neutral and risk-averse cases; an evidence of this can be spotted in the overlapping histograms from Figure 56, where the green colour for $Rho = -800$, the most risk-seeking decision-maker, dominates slightly more towards the right side, at around 14-15 years, while $Rho -5000$ can be seen slightly overriding the green colour on the left side, at around 7 years.

The risk-seeking case with a Rho of -5000 and the risk-neutral case have the same decision without information, at 15 years of secondary recovery lifetime ($Rho -5000$ is overlapping risk-neutral in the graph). The average decision with information for the three risk profiles are very close together, around 10 years of secondary lifetime. Consistent to previous analysis, there is a suggestion that the more risk-seeking the decision-maker is, the greater impact the information has in their decisions., as the difference in decision with and without information becomes larger (the gap between the solid line and dashed line for each Rho value).

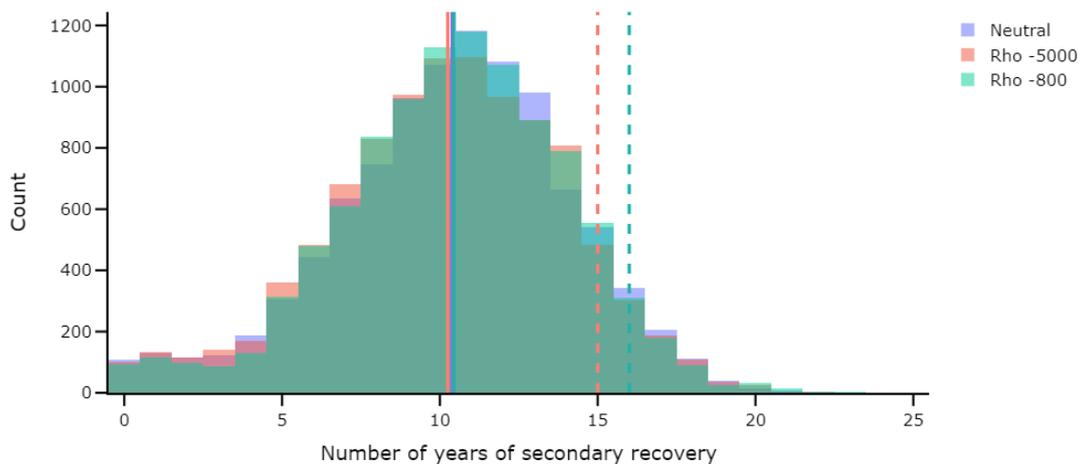


Figure 56. Decisions with and without information of secondary recovery lifetime for risk-neutral and risk-seeking attitudes. Continues lines: average decision with information. Dotted lines: decision without information.

Figure 57 and Figure 58 show the differences in the average decision with information and decision with no information, respectively for primary and secondary recovery lifetimes, for all studied risk-seeking Rhos. For primary recovery, Figure 57, the difference in years increases from 2.25 years at risk-neutral to 6.3 years at Rho -1100, and then a slight decrease to 6 years at Rho-800. This is consistent to what was observed in the risk-averse case: in general, the more risk-seeking, the larger the difference in decisions with and without information. This supports previous observation that the VOI tends to increase as the decision-maker is more risk-seeking.

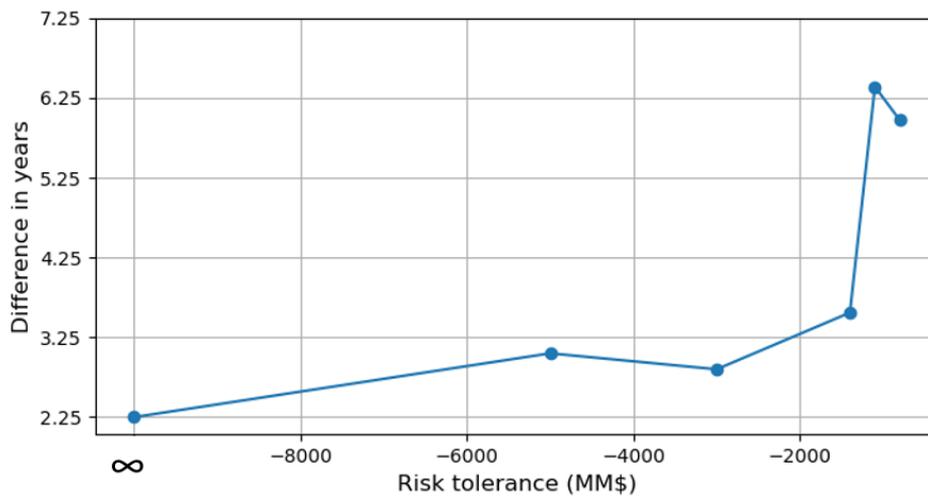


Figure 57. Difference in average decision with information and without information of primary recovery lifetime, for risk-neutral and risk-seeking attitudes.

In the analysis for secondary recovery, Figure 58, there is not a very clear trend. Starting with -4.6 years differences at risk-neutral, then the difference increases slightly to -4.8 at Rho -5000, and then decreasing again to -4.58 years at Rho -1400; then increases abruptly to 5.55 years at Rho -800. Due to this variation, it is considered complementary to assess the spread of the differences, shown in Figure 59 and Figure 60. Figure 59 and Figure 60 show boxplots for the difference in decisions with information and without information of the respective recovery phases lengths; this spread represents the variation of the differences for the 10,000 Monte Carlo samples.

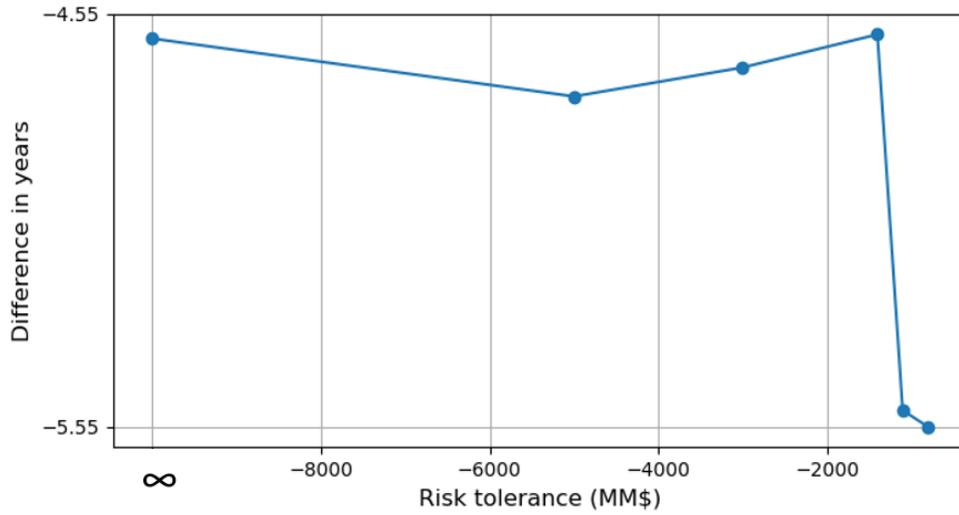


Figure 58. Difference in average decision with information and without information of secondary recovery lifetime, for risk-neutral and risk-seeking attitudes.

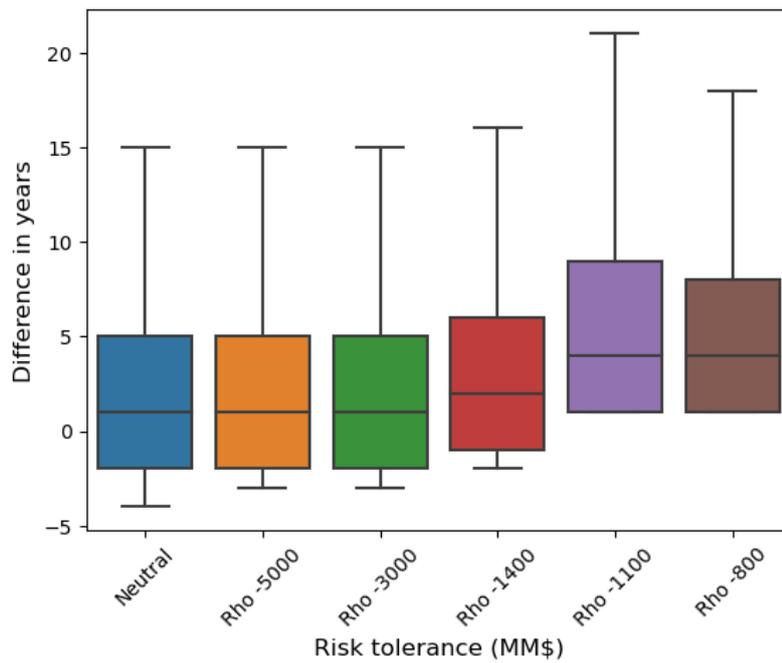


Figure 59. Boxplots of difference in decisions with information and without information of primary recovery lifetime, for risk-neutral and risk-seeking attitudes.

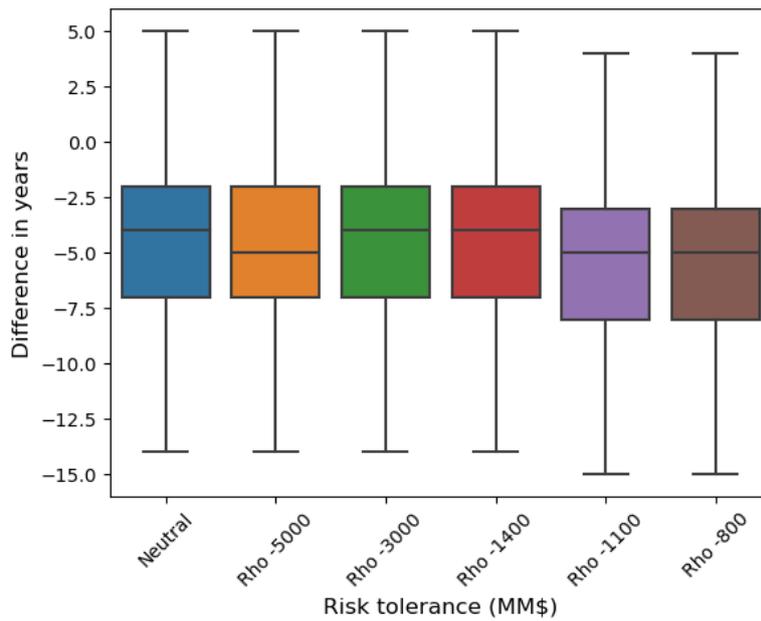


Figure 60. Boxplots of difference in decisions with information and without information of secondary recovery lifetime, for risk-neutral and risk-seeking attitudes.

Like in the risk aversion case, the boxplots from Figure 59 and Figure 60 do not show any monotonic trend for the median (P50) or spread of the differences across the different risk attitudes. A complimentary combined bar chart showing the difference in average decision with information and without information for primary and secondary recovery lifetimes is illustrated in Figure 61. It can be observed that this plot does not exhibit any specific trend either, even for the total lifetime, pointing to the same conclusion drawn before, attributing to the fact that the overall expected utilities result from a combined effect of primary and secondary recovery lifetimes.

Interestingly enough, also for the risk-seeking case there is the same observation about the polarity of the differences for primary and secondary phases: the differences in the decisions with and without information for the primary recovery are all positive, meaning that the average decision with information is for longer primary recovery phase; as for the secondary recovery the differences are all negative, indicating that the decisions with information are for shorter secondary recovery, for risk-seeking cases.

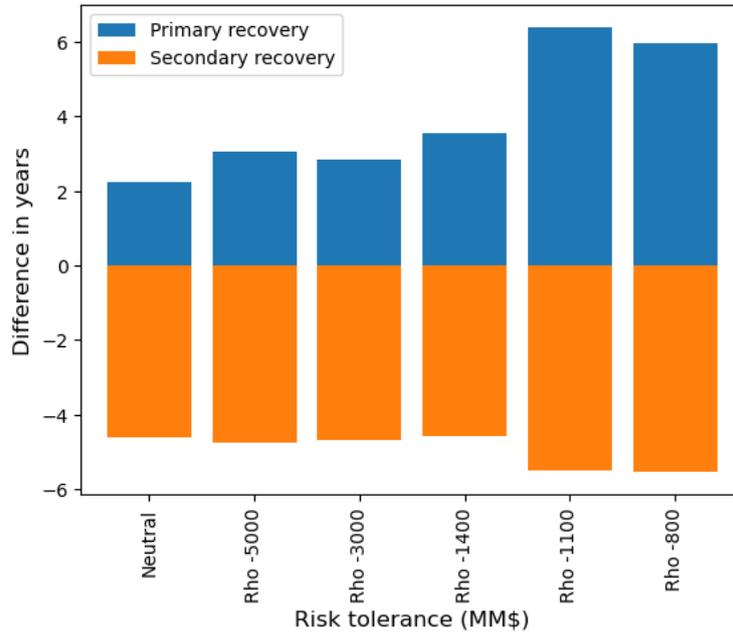


Figure 61. Difference in average decision with information and without information of primary and secondary recovery lifetimes, for neutral and risk-seeking attitudes.

Having assessed the risk-averse and risk-seeking decision-maker attitudes separately, combined graphs were created showing the differences in lifetime decisions with and without information for risk-averse and risk-seeking profiles together. This was performed with the objective of assessing potential overall trends. To combine the different Rhos in a single graph, Gamma, the inverse of Rho ($\text{Gamma} = 1/\text{Rho}$), also called ‘risk aversion coefficient’, is used, where $\text{Gamma} = 0$ models risk-neutral behaviours, positive Gamma further away from 0 models more risk-averse behaviours, and negative Gamma further away from 0 models more risk-seeking behaviours. Figure 62 shows the difference in years of the average decision with information and the decision without information for primary recovery lifetime. The x axis has been inverted to show extreme risk-aversion towards the left of the graph and extreme risk-seek towards the right. As concluded in the individual assessments of the risk attitudes,

there is a clear trend of increasing difference in the decisions with and without information as the decision-maker becomes more risk-seeking.

Figure 63 shows the difference in years of the average decision with information and the decision without information for secondary recovery lifetime. The individual assessments on each risk attitude had not shown any clear trend. However, when combining the risk attitudes in a single plot, an overall increasing trend can be fitted for the absolute differences of the decisions with and without information as the decision-maker becomes more risk-seeking. When comparing to the primary recovery trend, however, the secondary recovery overall trend is much less clear.

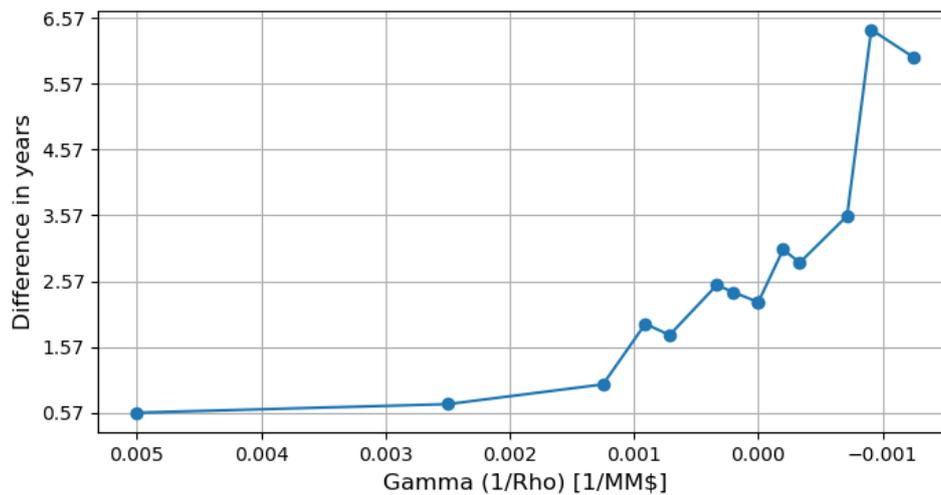


Figure 62. Difference in average decision with information and without information of primary recovery lifetime, for risk-averse and risk-seeking attitudes. Large positive numbers represent extreme risk-averse, small negative numbers represent extreme risk-seeking, zero represent risk-neutral.

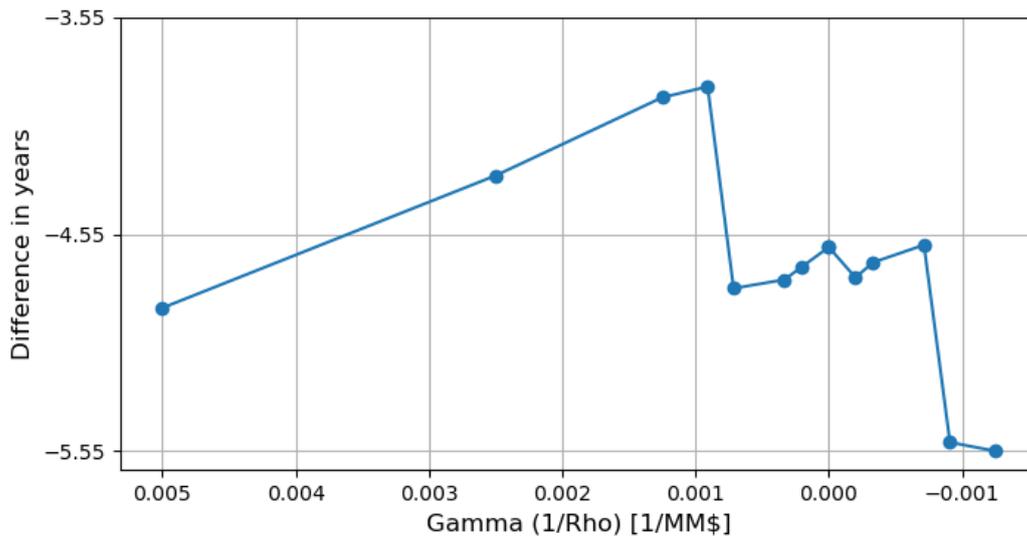


Figure 63. Difference in average decision with information and without information of secondary recovery lifetime, for risk-averse and risk-seeking attitudes. Large positive numbers represent extreme risk-averse, small negative numbers represent extreme risk-seeking, zero represent risk-neutral.

8.4. CHANGE OF OOIP IN THE PROBLEM SETTING

This section will investigate the impact of a change in the original oil in place (OOIP) distribution on the optimal decision strategy. The other parameters of the problem setting will be kept unchanged. The following sub-sections describe the change and discuss the consequent results.

8.4.1. Distribution of model parameter original initial in place (OOIP)

In the original problem, the prior OOIP is modelled using a bounded normal distribution with a mean of \$240 million and standard deviation of \$35 million. In this exercise, the mean is set to \$20 million and standard deviation to \$85 million, with the same distribution bounds. This change implies smaller production profits, as there are less reserves, and a larger uncertainty.

Figure 64 shows the histograms for the OOIP distribution as per the original and the new problem setting. It is noted that in the original problem setting the OOIP

distribution is more symmetrical around the mean. In the new OOIP, the distribution peaks in the lower end, at a range of 10-45 million bbl (OOIP cannot have negative values, hence, peak shows in the lower end), and has a longer tail to the upper end, since the standard deviation is larger than previous one.

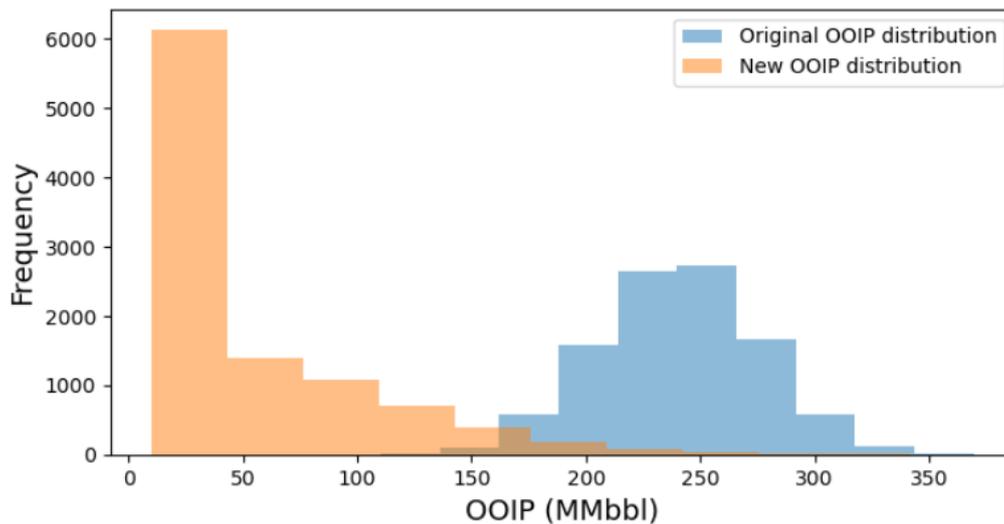


Figure 64. Histograms of the OOIP prior distribution corresponding to the original problem setting, ($N \sim (240, 35)$) and new problem setting ($N \sim (20, 85)$).

8.4.2. Sensitivity analysis of risk attitude

Hong et al. (2019) had also run the IOR initiation and termination case study with this change of OOIP distribution to assess the behaviour on the decisions for a risk-neutral case. Their results demonstrates that a risk-neutral decision-maker, without considering the future information (measurement of oil rates), prefers to not run the project, but the decision-maker prefers to run the project when considering the future information. In the present study, the same conclusion is drawn by solving the problem using the LSM algorithm with XGB for regression. The results confirm that a risk-neutral decision-maker would run the project only when the future information is considered. The risk-neutral behaviour leads to an average primary production lifetime of 1.17 years and an average secondary production lifetime of 3.06 years, when future information is accounted for. Hong et al. (2019) had estimated an EVWI (CEWI) and

VOI of \$55.3 million for the neutral case, which is considered close enough to the present result when using XGB for the LSM workflow, which resulted in \$52.06 million. This result highlights the importance of the sequential decision-making schemes, which allow the incorporation of information over time, enabling the creation of value from future information.

With the risk-neutral case results, it is intuitively presumed that a risk-averse decision-maker would not run the project, and a risk-seeking decision-maker might run the project even without the consideration of future information. This is confirmed by testing different Rho values: Rho = \$200 million (risk-averse), Rho = -\$3000 million (less risk-seeking) and Rho = -\$800 million (more risk-seeking). Table 22 summarizes the results, listing CEs, VOI and primary and secondary production lifetimes for all cases, with (WI) and without (WOI) information. The figures following Table 22 summarize the results in a graphical way. To include all risk profiles in single graphs, Gamma (risk aversion coefficient = $\Gamma = 1/\text{Rho}$) is used. Gamma = 0 models risk-neutral attitudes, positive Gamma further away from 0 models more risk-averse attitudes, and negative Gamma further away from 0 models more risk-seeking attitudes.

	Risk-averse Rho = MM\$ 200	Risk-neutral	Less Risk-seeking Rho = MM-\$3,000	More Risk-seeking Rho = MM -\$800
CE WOI (MM\$)	0	0	4.03	49.67
CE WI (MM\$)	0	52.06	60.07	101.04
VOI (MM\$)	0	52.06	56.04	51.37
Primary lifetime WOI (years)	0	0	0	0
Secondary lifetime WOI (years)	0	0	6	8
Avg. Primary lifetime WI (years)	0	1.17	1.15	0
Avg. Secondary lifetime WI (years)	0	3.06	2.85	3.51

Table 22. CEs, VOI, and production lifetime decisions for the IOR problem with the change of OOIP settings, obtained with different risk profiles.

As originally presumed, a risk-averse decision-maker would not run the project, and a risk-seeking decision-maker would run the project even without considering future information, but in this case, they would proceed straight to secondary recovery. Table 22 and Figure 65 show that a risk-averse decision-maker would assess a CE of 0 with or without information, hence, they would walk away the project. A risk-neutral would assess a CE of \$52.06 million when the future information is considered, and 0 if no information. The two risk-seeking decision-makers would both run the project. The less risk-seeking decision-maker (Rho= -\$3000 million) would obtain a CE of \$4.03 million with no information, and a CE of \$60.07 million with future information. The more risk-seeking decision-maker (Rho= -\$800 million) would assess a CE of \$49.67 million with no information, and a CE of \$101.04 million with future information. The differences between the CEs with and without information in both cases confirms once again the importance of incorporating future information in the decisions, as additional value can be generated. Another important observation is that there is an increasing trend of CE as the decision-maker becomes more risk-seeking, for both cases with and without information.

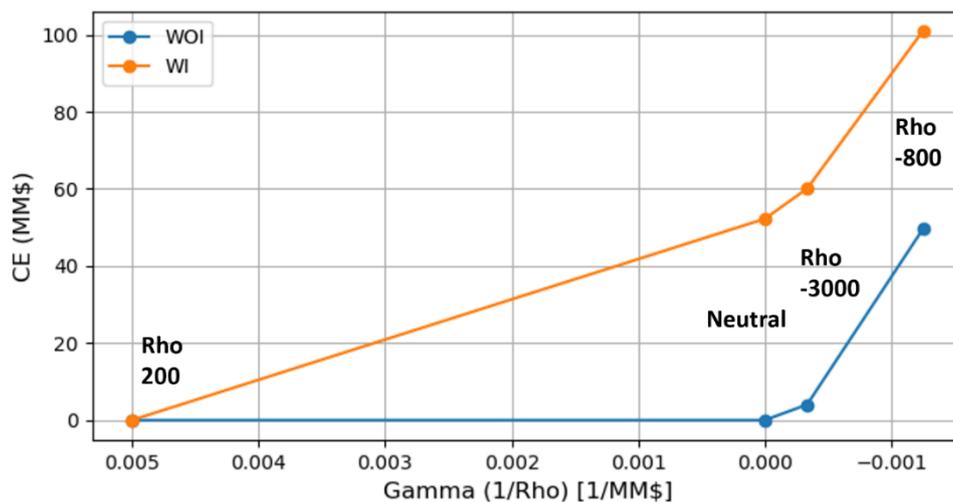


Figure 65. Sensitivity analysis on CE over different risk profiles for the IOR initiation and termination time problem with a change on the OOIP distribution.

Figure 66 shows the sensitivity analysis on VOI for the different risk profiles. VOI would be 0 for the risk-averse case since the project will not be run. For the risk-neutral case the VOI results in \$52.24 million. For the risk-seeking cases, $Rho = -\$3000$ million and $Rho = -800$, the VOI are \$56.04 million and \$51.37 million, respectively. For this problem setting, the differences in VOI between the different risk profiles are not considerable large as long as the project proceeds (for the risk-neutral and seeking cases).

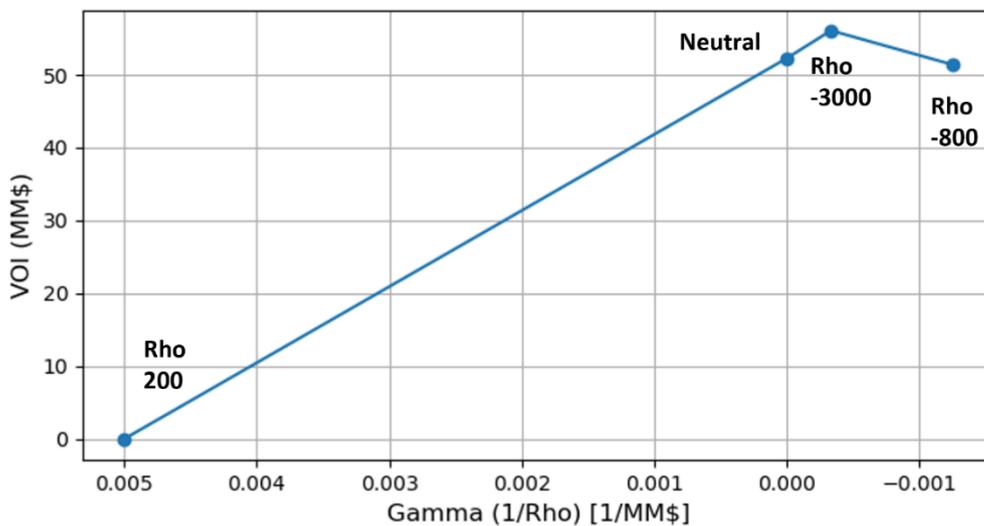


Figure 66. Sensitivity analysis on VOI over different risk profiles for the IOR initiation and termination time problem with a change on the OOIP distribution.

Figure 67 and Figure 68 show the decisions for primary and secondary recovery lifetime respectively, with and without information for the different risk profiles. As noted before, the risk-neutral decision-maker runs the project only with information, with an average primary production lifetime of 1.17 years and an average secondary production lifetime of 3.06 years. If no future information is accounted for, the risk-seeking decision-makers would skip the primary production and proceed straight to secondary production. The less risk-seeking decision-maker ($Rho = -\$3000$ million)

would prefer 6 years of secondary production (also total lifetime) while the more risk-seeking decision-maker ($\text{Rho} = -\$800$ million) would prefer 8 years.

When the future information is incorporated in the decision-making, the less risk-seeking decision-maker would run both, primary and secondary recovery phases, with an average lifetime of 1.15 and 2.85 years, respectively, totalizing 4 years of lifetime in average; this is similar to the result of the risk-neutral decision-maker. However, the more risk-seeking decision-maker would still go straight to secondary production, skipping primary recovery, with an average total lifetime of 3.51 years. This is consistent with the results that have been demonstrated and discussed for the original problem setting: a more risk-seeking decision-maker prefers shorter primary production and longer secondary production, in general.

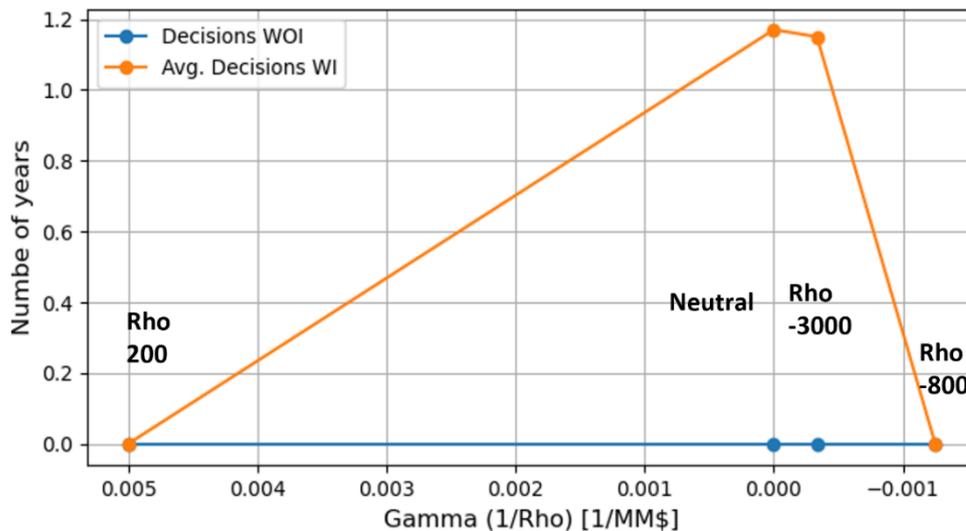


Figure 67. Sensitivity analysis primary recovery lifetime decision over different risk profiles for the IOR initiation and termination time problem with a change on the OOIP distribution.

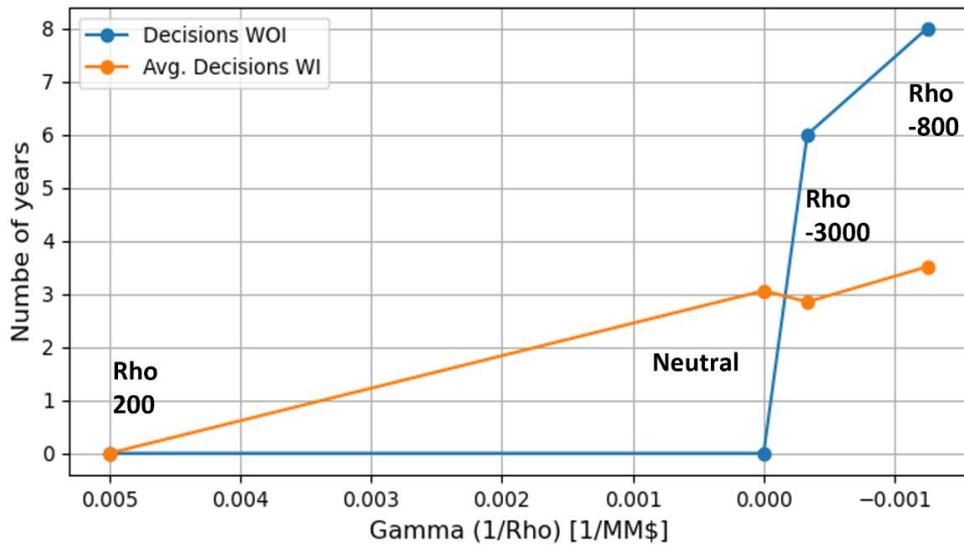


Figure 68. Sensitivity analysis secondary recovery lifetime decision over different risk profiles for the IOR initiation and termination time problem with a change on the OOIP distribution.

Chapter 9: Discussion

Hong et al. (2019) demonstrated the IOR initiation and termination time decision problem in a sequential decision-making framework, which considered information gathering and uncertainty updating (i.e., learning) over time for the optimization of decision strategy, with a numerical case study. Their case study was limited to risk-neutral behaviours with the objective to maximize expected monetary value – NPV. In the present study, exponential utility functions have been introduced to account for different risk attitudes and to investigate their impacts on the IOR initiation and termination time decision. With a utility function, the objective is to maximize expected utility.

Hong et al. (2019) used the LSM algorithm as an approximate dynamic programming approach to solve the sequential decision-making problem. The LSM algorithm is based on Monte Carlo simulation for uncertainty modelling and sampling, and regression for calculating (approximately) conditional expected values given information. Hong et al. used OLS for regression as they argued that the relation between NPVs and the information (measured oil rates) is linear for a given decision alternative. However, a ‘linear case’ means the objective function is a linear function of observables for given decisions and ‘a non-linear case’ means the objective function is a non-linear function of observables for given decisions. Even for a linear case, the relation between objective values and observed values might not be linear when a decision strategy is involved because the relation between objective values and the decision strategy might not be linear. For this reason, and because the inclusion of non-risk-neutral attitude by utility functions introduces non-linearity, machine-based non-linear regression methods have been assessed against OLS in the present study, even for the risk-neutral case.

Several machine learning methods (k-nearest neighbour (KNN), random forest (RF), extreme gradient boost (XGB), support vector regression (SVR) and piecewise regression) have been evaluated for the non-linear regression between utilities and measured oil rates. For each method, the optimal hyperparameters have been determined with a 3-fold cross-validation. Sensible fits have been obtained with KNN, RF and XGB; however, KNN have suffered difficulties to predict at very low oil rates. RF has predicted better at lower oil rates but under-estimated utilities at higher rates. XGB has given the best regression fit with the greatest R^2 score and without overfitting, among the evaluated methods. Therefore, XGB has been chosen and used as the non-linear regression method in the LSM algorithm for solving the IOR initiation and termination time problem in the present work.

XGB regression has been tested against OLS for the risk-neutral case, resulting in better fit and higher R^2 scores. This has validated previous observation that even when the relation between the NPVs and the measured oil rates is linear, when a decision strategy is involved, non-linearity can be introduced. In other words, the relation between NPVs and measured oil rates can be non-linear for mixed decision alternatives (i.e., the optimal future decisions can be a mix of decision alternatives for different realizations, depending on their information paths). Therefore, a non-linear regression method has provided a more accurate prediction.

Sensitivity analysis has been performed over various aspects of the IOR initiation and termination time decision problem for different risk tolerances. The impact of risk tolerance on the VOI has been investigated and discussed. For risk-averse attitudes, the VOI has been found to increase monotonically as the risk tolerance increases, which means that the less risk-averse a decision-maker is, the more they will value the future information. In the case of risk-seeking decision-makers, the VOI trend is slightly more variable; there is a slight VOI increase from risk-neutrality towards a risk tolerance of around \$1500 million, but then the VOI decreases rapidly for risk

tolerances closer to zero. For both risk-averse and risk-seeking decision-makers, the median and the spread of the final NPVs tend to increase as a decision-maker is more risk-seeking. The results are consistent with risk-averse and risk-seeking behaviours and correctly reflect the impact of risk attitude on the decisions.

To assess the impact of the risk attitude on the decisions, sensitivity analysis has been performed for different risk tolerances. One way of quantifying this impact, is to measure the difference between the production lifetime preferred by a risk-neutral decision-maker and the production lifetime preferred by a non-risk-neutral decision-maker. The case study results have demonstrated that risk-averse decision-makers prefer longer primary recovery and risk-seeking decision-makers prefer shorter primary recovery, compared to risk-neutral decision-makers. This behaviour has been attributed to the fact that a risk-averse decision-maker prefers to collect more information to reduce uncertainty before switching to secondary recovery. It has also been attributed to the additional CAPEX incurred for starting the secondary recovery: the extra recovery might or might not cover the CAPEX, and accounting for these two scenarios makes the NPV distribution wider, i.e., uncertainty increases, which is unpreferred by a risk-averse decision-maker.

Regarding the secondary recovery phase, it has been demonstrated that risk-averse decision-makers prefer to stop the secondary recovery sooner than risk-neutral decision-makers, and risk-seeking decision-makers would stop the secondary recovery later than risk-neutral decision-makers. In other words, the more risk-averse a decision-maker is, the shorter secondary recovery they prefer. This has been attributed to the OPEX, especially at the late-stage of secondary recovery, when the expected cash inflow is marginal and might not cover these costs. This is a situation that a risk-averse decision-maker does not prefer – the decision-maker's CE of the net cash flow (= cash inflow – OPEX) distribution can be negative even though the expected net cash flow is positive. For a risk-neutral decision-maker, their CE of the net cash flow equals the

expected net cash flow, by definition; hence, the risk-neutral decision-maker would prefer to continue the secondary recovery as long as the expected net cash flow is positive, resulting in a longer secondary recovery compared to a risk-averse decision-maker. For the same reason, the more risk-seeking a decision-maker is, the longer the secondary recovery they prefer – the risk-seeking decision-maker's CE of the net cash flow distribution can be positive even though the expected net cash flow is negative.

The impact of the future information on the decisions has been assessed via sensitivity analysis for difference risk tolerances. By analysing the difference between the decisions for primary recovery lifetime with and without future information, it has been noted that the more risk-seeking the decision-maker is, the larger this difference is, hence, the more valuable the information becomes. This observation complements previous VOI analysis. For the secondary recovery, a general trend has been observed indicating that the more risk-seeking, the larger difference between the secondary recovery lifetime with information and the secondary recovery lifetime without information. The trend for secondary recovery is not as clear as that for primary recovery – some up and down around a general trend for secondary recovery whilst a clear monotonic trend for primary recovery – because the determination of the secondary recovery lifetime is more complex and depends on both the switch time (the time when the primary recovery is switched to the secondary recovery) and the termination time (the time when the secondary recovery is terminated), compared to the primary recovery lifetime that depends only on the switch time.

The differences in the decisions with and without information for the primary recovery lifetime have been found positive for all risk profiles, indicating that all decision-makers would prefer longer primary recovery when considering future information, in average. This is attributed to the fact that the information during the primary recovery is valuable reducing the uncertainty in secondary recovery and informing the decision-maker to avoid switching to secondary recovery too early. For

the secondary recovery, the differences with and without information have been all found negative, meaning that all decision-makers would prefer shorter secondary recovery when considering future information, in average. The information during the secondary recovery is valuable in reducing the uncertainty in secondary recovery at later stages and informing the decision-maker to avoid terminating the secondary recovery too late.

An additional exercise has been performed by changing the problem setting to a more marginal and more uncertain case, where the OOIP has a lower mean (thus the expected NPV is closer to 0) and a larger standard deviation (thus the spread of NPV distribution is larger). Hong et al. (2019) had run the exercise in a risk-neutral case; they had found that a risk-neutral decision-maker would not run the project without the consideration of future information, but they would proceed to one year of primary recovery and three years of secondary recovery, in average, when accounting for the future information in their sequential decisions. In the present work, Hong et al. (2019)'s risk-neutral case results have been confirmed using XGB regression in the LSM algorithm, obtaining very similar decision strategy and expected NPV. These results reinforce the importance of the sequential decision-making schemes, which allow the incorporation of future information over time, creating value from information.

Other risk attitudes have also been investigated in the changed problem setting. The results indicated that a risk-averse decision-maker would not run the project, with and without considering future information. Risk-seeking decision-makers would run the project even without considering future information. The CEs increase as a decision-maker becomes more risk-seeking, for both cases with and without the consideration of future information. The VOI is not considerably different among the risk-neutral and risk-seeking cases, as long as the project is started.

In the changed problem setting, if no future information is considered, risk-seeking decision-makers would skip the primary production and proceed straight to secondary production. This is still the case for extreme risk-seeking decision-makers when future information is accounted for. However, with the consideration of future information, less risk-seeking decision-makers prefer to start with a short primary recovery and then decide if switch to secondary recovery based on the oil rates measured during the primary recovery, similar to the risk-neutral case.

This work has demonstrated that machine learning-based non-linear regression can be effectively used in LSM. The most important advantage of machine learning is the more accurate utilities prediction compared to OLS and the flexibility to be used with all risk profiles, including risk-neutral. One of the disadvantages is the longer compute time these methods require compared to OLS or ordinary fitting approaches. Although this work was not focused on time efficiency, it could be appreciated that some of the algorithms required much more compute time than others, for example random forest and piecewise regressions compared to k nearest neighbour. For more complex problems, a compromise might need to be made to obtain sufficient prediction quality without incurring into unnecessary costs. The other disadvantage of the machine learning methods is the required efforts during the hyperparameters testing and cross-validation. This testing phase can be exhaustive and might imply increase of both, human-resources, and computer-resources, when compared to ordinary fitting methods.

Another contribution from this work is the validation of the LSM algorithm to solve sequential decision-making problems that involve risk attitudes different than neutral. The LSM has provided with an efficient way to estimate the conditional expected utilities replacing the need of Bayesian updates for conditional probabilities, as used in conventional decision tree approaches. One of the disadvantages is that the methodology delivers an “approximation” of the expected utilities, and not the accurate figures. However, this work has showed that the VOI obtained by LSM is very

comparable to the one obtained through the conventional analytical tree approach, indicating the approximation in this specific problem has been suitable. Nonetheless, this might be validated in other problem settings.

The results from the present study indicate that the risk attitude change the decisions. It also shows that value of information is not the same for all risk profiles. Moreover, the overall results depend on problem setting. The contrasting results over the different risk tolerances also showed how sensitive a project strategy can be to the risk attitude. However, there is not “right or wrong” on a corporation to be risk-averse, risk-neutral, or risk-seeking. The benefit of being risk-averse is that potential losses are avoided when the scenarios are very uncertain and/or marginals. However, a risk-averse corporation might reject project opportunities that might result profitable. On the other hand, a risk-seeking corporation might decide to invest and eventually loose when the scenario is uncertain; but on the other hand, they embrace project opportunities. The important highlight of the present study is that, in all cases, the risk attitude can be modelled, and its impact can be assessed, which enables robust scenario-testing in sequential decision-making problems.

Chapter 10: Conclusions and Recommendations

In this work, a sequential decision-making problem has been solved using the LSM algorithm - an approximate dynamic programming approach. A hypothetical field with maximum 50 years of production has been assumed, and the objective was to assess the optimal lifetime for primary and secondary recovery - IOR initiation and termination time problem. Geological and petrophysical uncertainties have been taken into consideration in a decline-curve-based production model, and certain economic parameters have been used. The novelty of the work has been the introduction of risk attitude and the assessment of the different risk profiles in the decisions.

This thesis work has demonstrated that the risk attitude can be effectively modelled, and its impact on decisions can be assessed. In the IOR initiation and termination time problem, sensitivity analysis on the different risk profiles led to the following remarks:

- Different risk attitude leads to different decisions.
- Impact and value of information depends on risk attitude.
- Results are specific to problem setting.

The results have also demonstrated the importance of incorporating future information in decision modelling, which enables creation of value. This is relevant for all risk attitudes.

Another contribution of the present work is the effective application of machine learning-based non-linear regression for sequential decision modelling problems. The non-linear regression has been used to approximate the expected utilities during the LSM algorithm, to exercise the optimal decision policy. Non-linear regression can be successfully used in any risk attitude case, including risk-neutral cases.

The approach presented in this work is effective, efficient and is not limited to oil and gas problems, but has applicability to many other fields, for example, pharmaceuticals, medicine, agriculture, and new energy supply.

Recommendations for further work include:

- Testing other machine learning-based non-linear regression methods such as neural networks.
- Using other dynamic programming approaches instead of LSM, for example, reinforcement learning.
- Assessing the impact on risk attitude in problems that involve more complex production models (including water and gas production) and/or uncertain economic parameters (e.g., oil and gas prices).
- Application to other decision contexts related to the oil and gas industry, for example, finding the optimal time to drill an infill well.

References

- Abbas, A. E., and R. A. Howard, 2015, *Foundations of decision analysis*, Pearson Education.
- Agami Reddy, T., 2011, *Estimation of Linear Model Parameters Using Least Squares*: Boston, MA, Boston, MA: Springer US, p. 141-182.
- Ahmadi, R., and R. Bratvold, 2023, An exposition of least square Monte Carlo approach for real options valuation: *Journal of Petroleum Science and Engineering*, no. 222, doi: <https://doi.org/10.1016/j.petrol.2022.111230>.
- Andréasson, J., and P. V. Shevchenko, 2017, Bias-corrected Least-Squares Monte Carlo for utility based optimal stochastic control problems: Centre for Financial risk, Faculty of Business and Economics, Macquarie University, v. Working paper 17-08, p. 1-24.
- Bellman, R., 1957, *Dynamic programming*: Princeton, N.J, Princeton University Press.
- Black, F., and M. Scholes, 1973, The Pricing of Options and Corporate Liabilities: *The Journal of political economy*, v. 81, no. 3, p. 637-654, doi: 10.1086/260062.
- Boyle, P. P., 1977, Options: A Monte Carlo approach: *Journal of financial economics*, v. 4, no. 3, p. 323-338, doi: 10.1016/0304-405X(77)90005-8.
- Brandao, L. E., J. S. Dyer, and W. J. Hahn, 2005, Using Binomial Decision Trees to Solve Real-Option Valuation Problems: *Decision analysis*, v. 2, no. 2, p. 69-88, doi: 10.1287/deca.1050.0040.
- Bratvold, R., and S. H. Begg, 2010, *Making Good Decisions*, Society of Petroleum Engineers.
- Breiman, L., 2001, Random Forests: *Machine Learning*, v. 45, p. 5-32, doi: <https://doi.org/10.1023/A:1010933404324>.
- Breiman, L., J. H. Friedman, R. A. Olshen, and C. J. Stone, 1984, *Classification And Regression Trees*, 1 ed., v., CRC Press.
- Chen, L., M. Zhou, J. She, Z. Liu, F. Dong, and K. Hirota, 2018, Three-Layer Weighted Fuzzy Support Vector Regression for Emotional Intention Understanding in Human–Robot Interaction: *IEEE Transactions on fuzzy systems*, v. 26, no. 5, p. 2524-2538, doi: <https://doi.org/10.1109/TFUZZ.2018.2809691>.
- Cook, R. D., and S. Weisberg, 1982, Criticism and Influence Analysis in Regression: *Sociological methodology*, v. 13, p. 313-361, doi: 10.2307/270724.
- Cox, J. C., S. A. Ross, and M. Rubinstein, 1979, Option pricing: A simplified approach: *Journal of financial economics*, v. 7, no. 3, p. 229-263, doi: 10.1016/0304-405X(79)90015-1.
- Cunningham, P., and S. J. Delany, 2021, k-Nearest Neighbour Classifiers - A Tutorial: *ACM Comput. Surv.*, v. 54, no. 6, Article 128, 25 p., doi: <https://doi.org/10.1145/3459665>.
- Diederich, A., 2001, Sequential Decision Making, Smelser, N.J., P.B. (Eds), *International Encyclopedia of the Social and Behavioral Sciences.*, p. 13917-13922.
- Hong, A., and R. B. Bartvold, 2023, Impact of Risk Attitude on Reservoir Management Decisions, Manuscript in preparation.

- Hong, A., R. B. Bratvold, and L. W. Lake, 2019, Fast analysis of optimal improved-oil-recovery switch time using a two-factor production model and least-squares Monte Carlo algorithm.: SPE Reservoir Evaluation & Engineering, v. 22(03), p. 1144-1160.
- Hong, A. J., R. B. Bratvold, P. Thomas, and R. G. Hanea, 2018, Value-of-information for model parameter updating through history matching: Journal of petroleum science & engineering, v. 165, p. 253-268, doi: 10.1016/j.petrol.2018.02.004.
- Hongrui, C., R. Lun, and Z. Ran, 2016, Evaluating CCS Investment of China by a Novel Real Option-Based Model: Mathematical problems in engineering, v. 2016, doi: 10.1155/2016/8180674.
- Howard, R. A., 1966, Decision Analysis: Applied Decision Theory: Fourth International Conference in Operational Research, New York City: Wiley-Interscience, p. 55-71.
- Jafarizadeh, B., and R. Bratvold, 2012, Two-Factor Oil-Price Model and Real Option Valuation: An Example of Oilfield Abandonment: SPE economics & management, v. 4, no. 3, p. 158-170, doi: 10.2118/162862-PA.
- Jafarizadeh, B., and R. B. Bratvold, 2013, Sell Spot or Sell Forward? Analysis of Oil-Trading Decisions With the Two-Factor Price Model and Simulation: SPE Economics & Management, no. SPE 165581, p. 80-88, doi: <https://doi.org/10.2118/165581-PA>.
- Jordaan, I. J., 2004, Decisions under uncertainty : probabilistic analysis for engineering decisions, Cambridge, Cambridge University Press.
- Kullawan, K., R. B. Bratvold, and J. E. Bickel, 2018, Sequential geosteering decisions for optimization of real-time well placement: Journal of petroleum science & engineering, v. 165, p. 90-104, doi: 10.1016/j.petrol.2018.01.068.
- Longstaff, F., and E. Schwartz, 2001, Valuing American Options by Simulation: A Simple Least-Squares Approach.: The Review of Financial Studies, v. 14, no. 1, p. 113-147.
- Loy, A., L. Follett, and H. Hofmann, 2016, Variations of Q-Q Plots: The Power of Our Eyes: The American statistician, v. 70, no. 2, p. 202-214, doi: 10.1080/00031305.2015.1077728.
- Mokhtar, A., M. Jalali, H. He, N. AL-Ansari, A. Elbeltagi, K. Alsafadi, and H. G. Abdo, 2021, Estimation of SPEI Meteorological Drought Using Machine Learning Algorithms: IEEE Access, v. 9, p. 65503-65523, doi: <https://doi.org/10.1109/ACCESS.2021.3074305>.
- Neely III, J. E., and R. de Neufville, 2001, Hybrid real options valuation of risky product development projects: Int. J. of Technology. 1. 10.1504/IJTPM.2001.001743, p. 29-46.
- Parra-Sanchez, C., 2010, A Life Cycle Optimization Approach to Hydrocarbon Recovery: Master's thesis, The University of Texas at Austin, 122 p.
- Powell, W. B., 2011, Approximate dynamic programming: solving the curses of dimensionality, second edition: Wiley series in probability and statistics, 2nd ed., v.: Hoboken, N.J, Hoboken, N.J: John Wiley & Sons.
- Schützenmeister, A., U. Jensen, and H. P. Piepho, 2012, Checking Normality and Homoscedasticity in the General Linear Model Using Diagnostic Plots: Communications in statistics. Simulation and computation, v. 41, no. 2, p. 141-154, doi: 10.1080/03610918.2011.582560.

- scikit-learn.org, Model selection. scikit-learn manual. Web. Accessed: February 16th, 2023, https://scikit-learn.org/stable/modules/model_evaluation.html#r2-score.
- scikit-learn.org, R² score, the coefficient of determination. scikit-learn manual. Web. Accessed: February 17th, 2023, https://scikit-learn.org/stable/modules/model_evaluation.html#r2-score.
- Shahali, R., 2022, Machine learning methods for assessing value-of-information: Master's thesis, University of Stavanger, 71 p.
- Shamshirband, S., D. Petkovi', H. Javidnia, and A. Gani, 2015, Sensor Data Fusion by Support Vector Regression Methodology—A Comparative Study: IEEE SENSORS JOURNAL, v. 15, no. 2, p. 850-854, doi: <https://doi.org:10.1109/JSEN.2014.2356501>.
- statsmodels.org, Graphics, gofplots, qqplot. Statsmodels manual. Web. Accessed: January 31st, 2023, <https://www.statsmodels.org/0.9.0/generated/statsmodels.graphics.gofplots.qqplot.html>.
- Von Neumann, J., and O. Morgenstern, 1947, Theory of games and economic behavior: Princeton, NJ, Princeton University Press.
- Von Winterfeldt, D., and W. Edwards, 2007, Defining a Decision Analytic Structure, Cambridge University Press, p. 81-103.
- Wang, C., G. Li, and A. C. Reynolds, 2009, Production Optimization in Closed-Loop Reservoir Management: SPE journal (Society of Petroleum Engineers (U.S.) : 1996), v. 14, no. 3, p. 506-523, doi: 10.2118/109805-PA.
- Wang, Y., J. Li, and Y. Li, 2017, Choosing Between Two Classification Learning Algorithms Based on Calibrated Balanced 5X2 Cross-Validated F-Test: Neural processing letters, v. 46, no. 1, p. 1-13, doi: 10.1007/s11063-016-9569-z.
- Wiggins, M. L., and R. A. Startzman, 1990, An Approach to Reservoir Management: SPE Annual Technical Conference and Exhibition. New Orleans, Louisiana, 23–26 September.
- Wilk, M. B., and R. Gnanadesikan, 1968, Probability Plotting Methods for the Analysis of Data: Biometrika, v. 55, no. 1, p. 1-17, doi: 10.2307/2334448.
- Witten, I. H., E. Frank, M. A. Hall, and C. J. Pal, 2011, Data Mining: Practical Machine Learning Tools and Techniques: The Morgan Kaufmann Series in Data Management Systems: San Francisco, San Francisco: Elsevier Science, xxxiii-xxxiii p.
- Wong, T.-T., and P.-Y. Yeh, 2020, Reliable Accuracy Estimates from k-Fold Cross Validation: IEEE transactions on knowledge and data engineering, v. 32, no. 8, p. 1586-1594, doi: 10.1109/TKDE.2019.2912815.
- Wui, N. C. S., 2019, Using the Least-Squares Monte Carlo Algorithm to Optimize IOR Initiation Time: Master's thesis, Norwegian University of Science and Technology, 202 p.
- Yang, L., S. Liu, S. Tsoka, and L. G. Papageorgiou, 2016, Mathematical programming for piecewise linear regression analysis: Expert Systems With Applications, v. 44, p. 156-167, doi: <https://doi.org/10.1016/j.eswa.2015.08.034>.

Appendix A: Sensitivity analysis of the non-linear regression hyperparameters on the fitting scores.

The following supplementary figures show additional sensitivity analysis on the fitting scores obtained with the different non-linear regression algorithms and their respective material hyperparameters, as part of the testing phase documented in Chapter 7: Selection of regression method, Section 7.1: K-fold cross-validation.

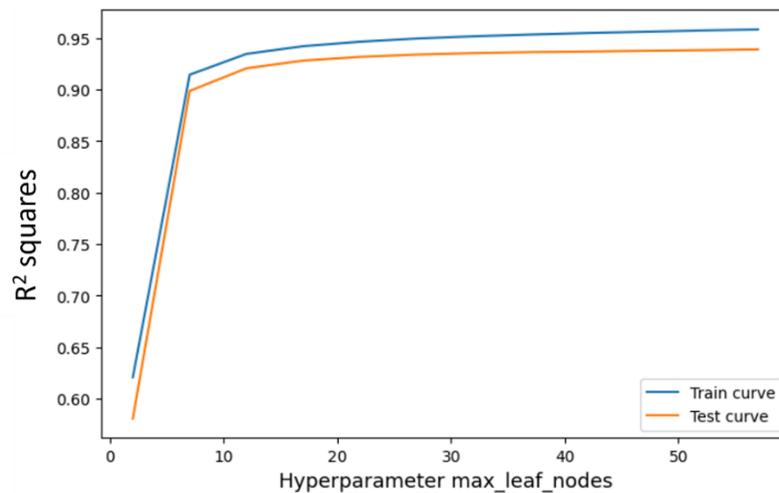


Figure 69. Sensitivity analysis of hyperparameter ‘maximum leaf nodes’ from Random Forest algorithm on the fitting scores.

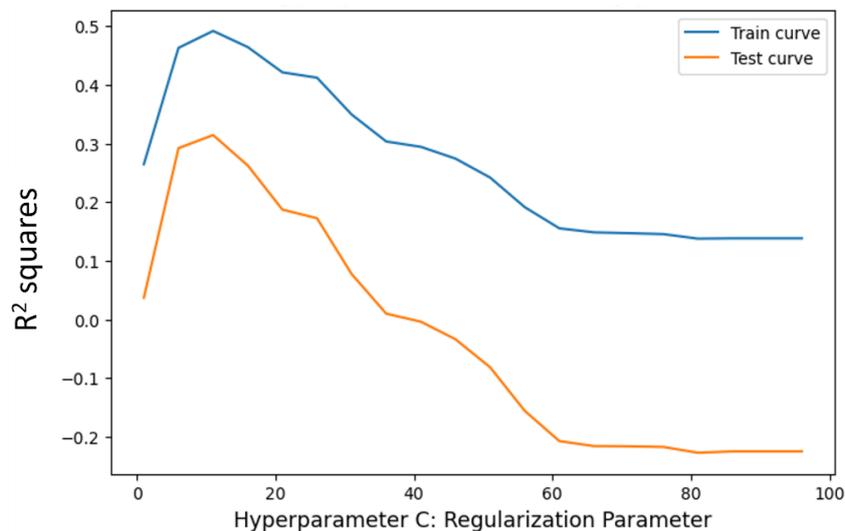


Figure 70. Sensitivity analysis of hyperparameter ‘C: Regularization from SVR algorithm on the fitting scores.

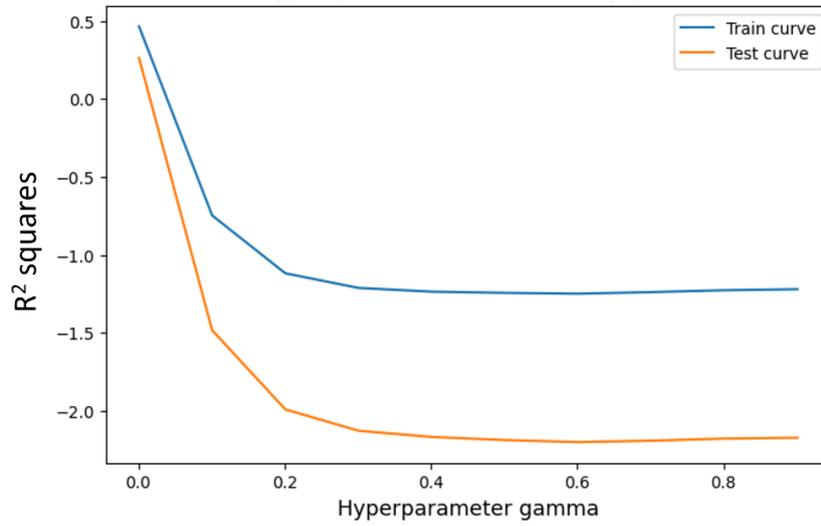


Figure 71. Sensitivity analysis of hyperparameter 'gamma' from SVR algorithm on the fitting scores.

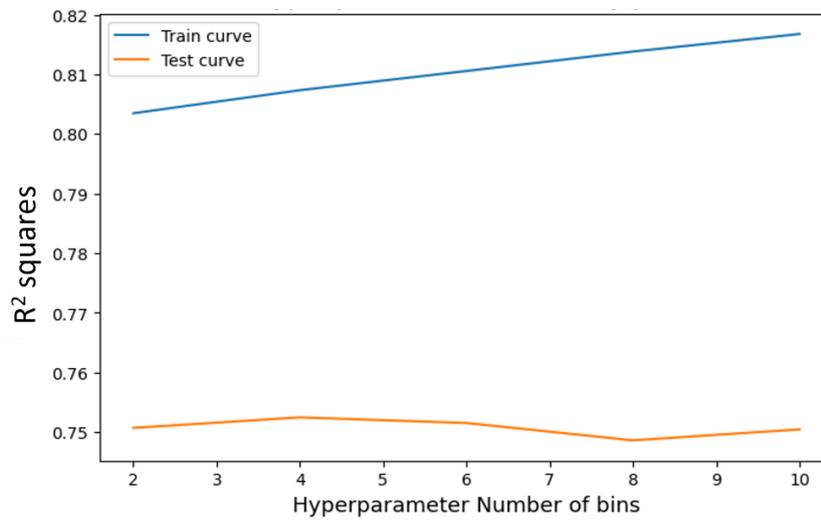


Figure 72. Sensitivity analysis of hyperparameter 'number of bins' from Piecewise algorithm (RF estimator) on the fitting scores.

Appendix B: Bi-dimensional regression 3D fit plots

These supplementary images show additional 3D regression fits using different non-linear algorithms, as part of the testing phase documented in Chapter 7: Selection of regression method, Section 7.3: Comparison of different regression methods.

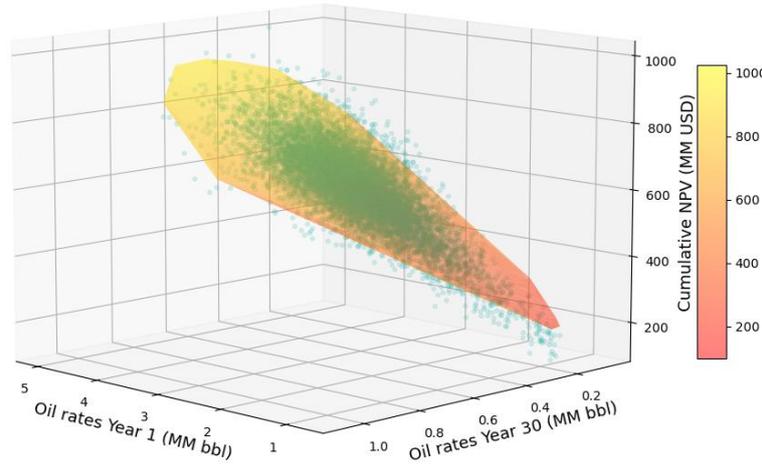


Figure 73. Bi-dimensional regression 3D fit plot, using OLS for NPV versus oil rates (risk-neutral case). Blue dots: input data. Warmer colors: regression surface.

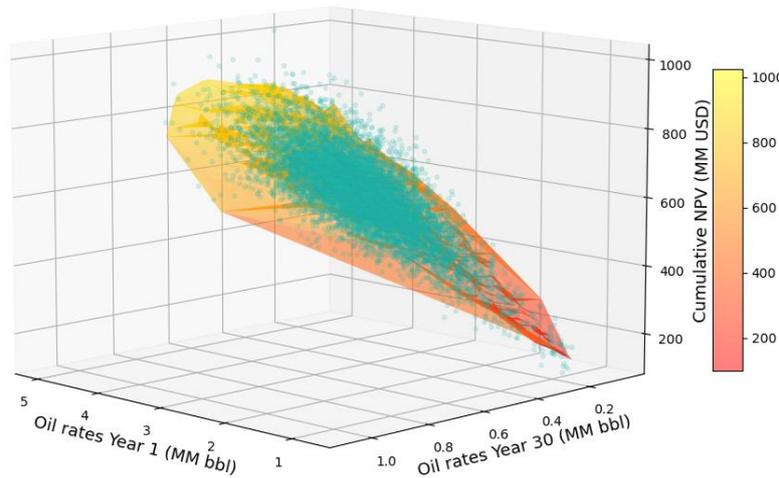


Figure 74. Bi-dimensional regression 3D fit plot, using XGB for NPV versus oil rates (risk-neutral case). Blue dots: input data. Warmer colors: regression surface.

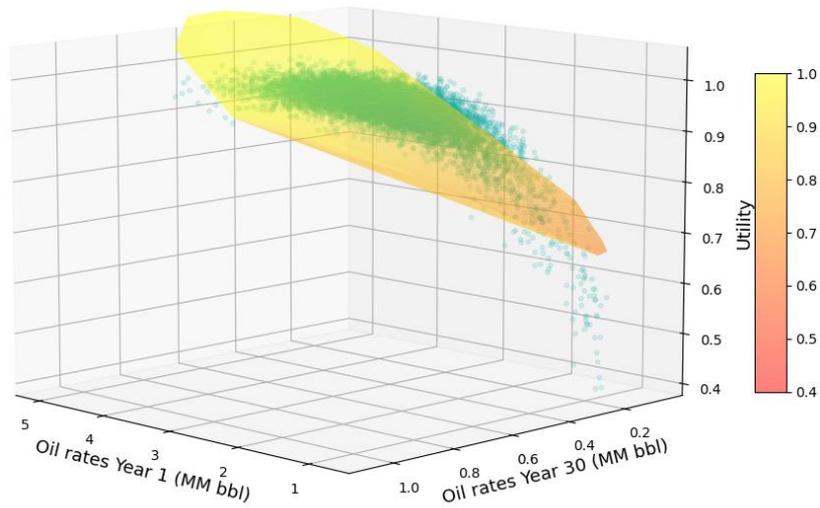


Figure 75. Bi-dimensional regression 3D fit plot, using SVR for utilities versus oil rates (risk-averse case). Blue dots: input data. Warmer colors: regression surface.

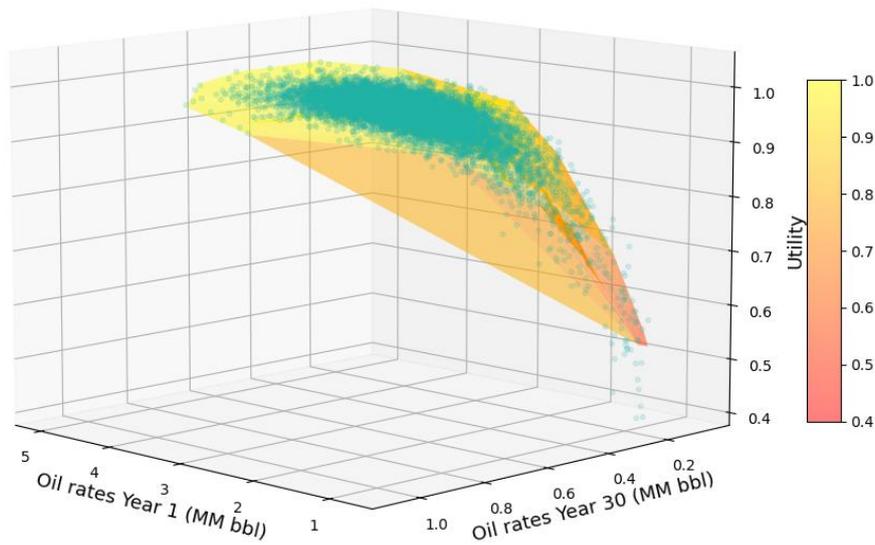


Figure 76. Bi-dimensional regression 3D fit plot, using Piecewise 10 slopes (OLS regressor) for utilities versus oil rates (risk-averse case). Blue dots: input data. Warmer colors: regression surface.

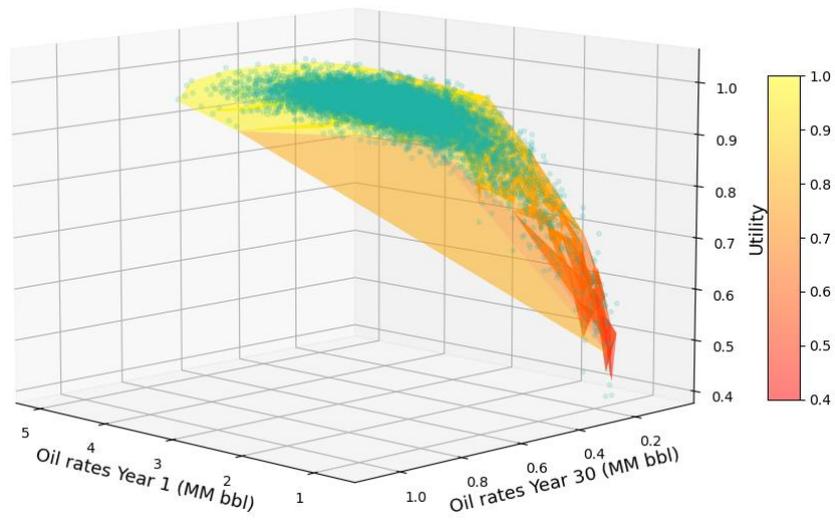


Figure 77. Bi-dimensional regression 3D fit plot, using Piecewise 4 slopes (XGB regressor) for utilities versus oil rates (risk-averse case). Blue dots: input data. Warmer colors: regression surface.

Appendix C: LSM Utilities validation

These supplementary images show the utilities validation for additional cases of risk tolerances, as documented in Chapter 8: Case study, Section 8.3: Case study results, Sub-section 8.3.2.1: LSM utilities validation.

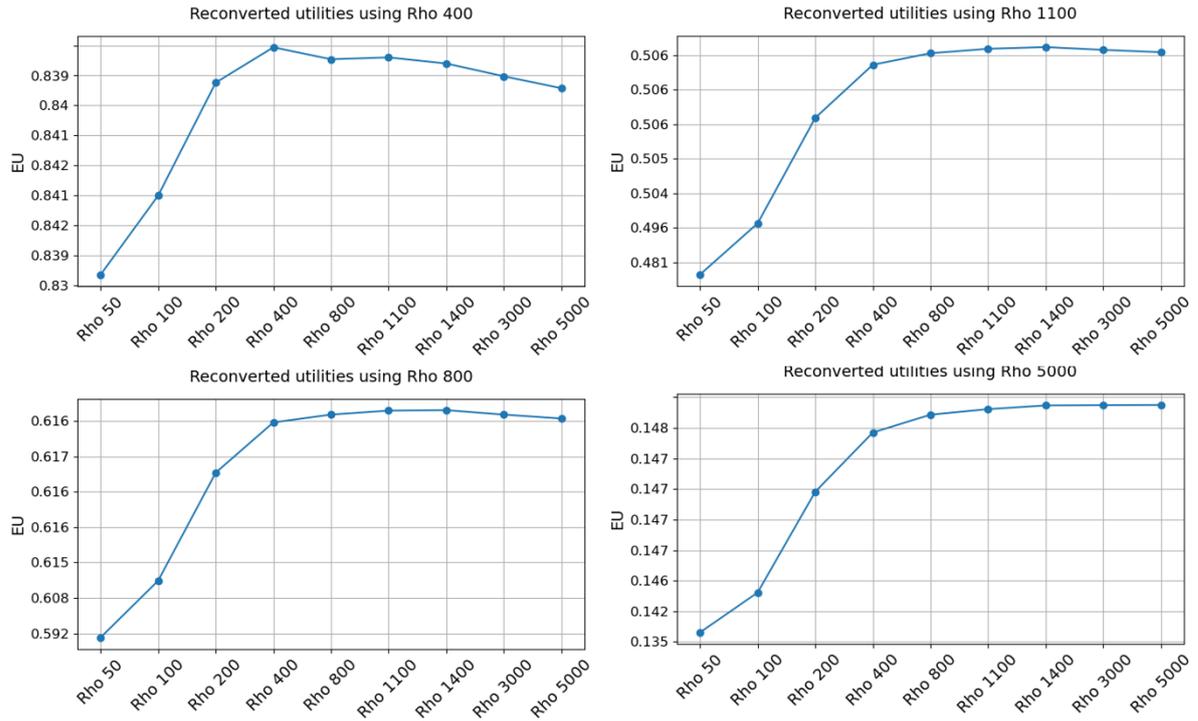


Figure 78: LSM utilities validation: Risk-averse EU re-converted using different risk tolerances. Maximum EU is found by re-converting with its respective Rho.

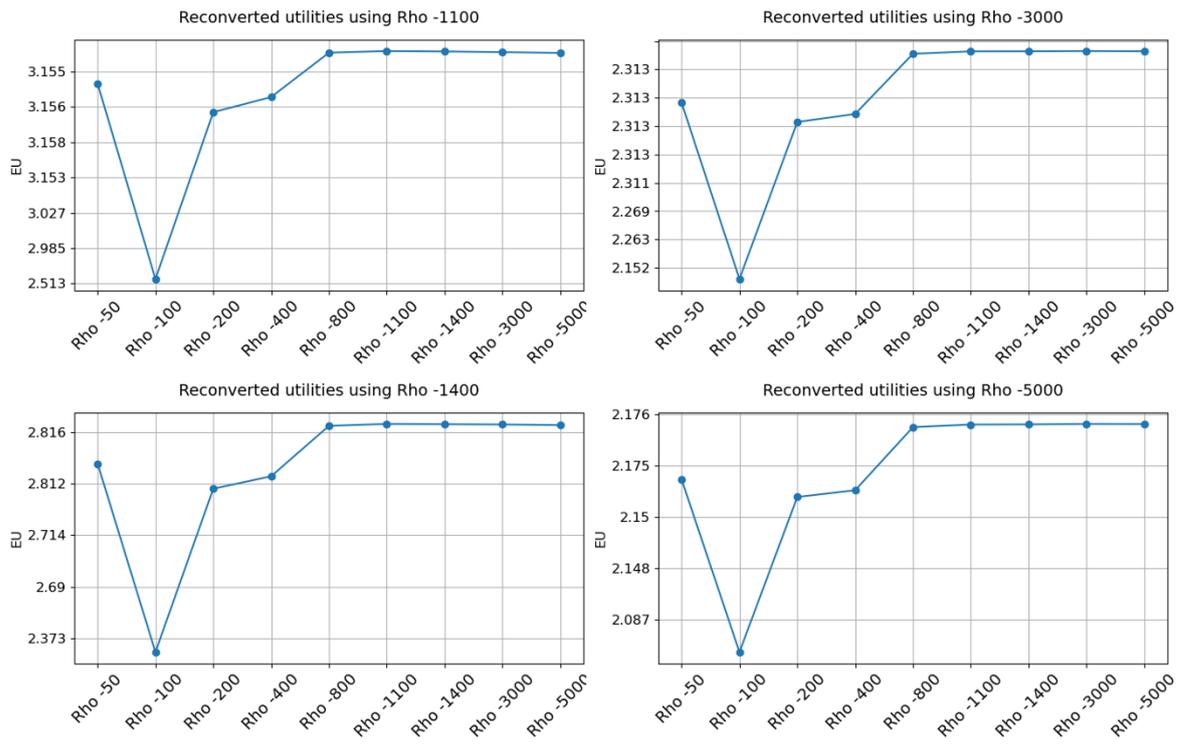


Figure 79: LSM utilities validation: Risk-seeking EU re-converted using different risk tolerances. Maximum EU is found by re-converting with its respective Rho.

Appendix D: LSM Python code for the IOR initiation and termination time case study with risk-neutral attitude, using OLS regression.

This code solves the IOR initiation and termination time case study for a risk-neutral decision-maker, using OLS regression in the LSM algorithm, as per results documented in Chapter 8. This code also exports 10,000 Monte Carlo realizations of the model parameters, to be used by the subsequent cases to make direct comparisons.

https://github.com/marianneyanez/LSM-two-factor-model-optimization/blob/main/LSM_risk_neutral_OLS_and_export_10000_MC_samples.ipynb

Appendix E: LSM Python code for the IOR initiation and termination time case study with risk-neutral attitude, using XGB regression.

This code solves the IOR initiation and termination time case study for a risk-neutral decision-maker, using XGB regression in the LSM algorithm, as per results documented in Chapter 8. This code reads 10,000 Monte Carlo samples for the model parameters generated in code linked in Appendix D.

https://github.com/marianneyanez/LSM-two-factor-model-optimization/blob/main/LSM_risk_neutral_XGB_reading_fixed_10000_MC_samples.ipynb

Appendix F: LSM Python code for the IOR initiation and termination time case study with risk-averse attitude, using XGB regression.

This code solves the IOR initiation and termination time case study for a risk-averse decision-maker, using XGB regression in the LSM algorithm, as per results documented in Chapter 8. This code reads 10,000 Monte Carlo samples for the model parameters generated in code linked in Appendix D.

https://github.com/marianneyanez/LSM-two-factor-model-optimization/blob/main/LSM_risk_averse_XGB_reading_fixed_10000_MC_samples.ipynb

Appendix G: LSM Python code for the IOR initiation and termination time case study with risk-seeking attitude, using XGB regression.

This code solves the IOR initiation and termination time case study for a risk-seeking decision-maker, using XGB regression in the LSM algorithm, as per results documented in Chapter 8. This code reads 10,000 Monte Carlo samples for the model parameters generated in code linked in Appendix D.

https://github.com/marianneyanez/LSM-two-factor-model-optimization/blob/main/LSM_risk_seeking_XGB_reading_fixed_10000_MC_sample.s.ipynb