Hyldmo, Tobias
Skudal, Markus Hoff

Development of a portfolio of marginal fields under reservoir and price uncertainty: A compound real options approach with Bayesian learning

Master's thesis in Industrial Economics and Technology Management
Supervisor: Verena Hagspiel
Co-supervisor: Semyon Fedorov, Olga Noshchenko and Sophie Haseldonckx (NPD)
June 2023

Master's thesis

NTNU
Norwegian University of Science and Technology
Dept. of Industrial Economics and Technology Management

NTNU
Norwegian University of
Science and Technology

NORWEGIAN PETROLEUM DIRECTORATE

Hyldmo, Tobias
Skudal, Markus Hoff

Development of a portfolio of marginal fields under reservoir and price uncertainty: A compound real options approach with Bayesian learning

Master's thesis in Industrial Economics and Technology Management
Supervisor: Verena Hagspiel
Co-supervisor: Semyon Fedorov, Olga Noshchenko and Sophie Haseldonckx (NPD)
June 2023

Norwegian University of Science and Technology
Dept. of Industrial Economics and Technology Management

NTNU
Norwegian University of
Science and Technology

Preface

This academic paper is delivered as a Master's Thesis in Financial Engineering as a concluding part of our MSc. degree in Industrial Economics and Technology Management at the Norwegian University of Science and Technology (NTNU). The paper is planned to be submitted to a peer-reviewed international journal on energy economics. This work is also related to BRU21 - NTNU's research and innovation program in digital and automation solutions for the offshore energy sector (`www.ntnu.edu/bru21`). Moreover, this work was presented at the BRU21 conference on the 31st of May, 2023, in Trondheim.

The paper proposes a solution for the valuation and optimal timing of tying back a portfolio of marginal fields on the Norwegian continental shelf (NCS) under reservoir- and price uncertainty. We utilize our diverse backgrounds in finance and computer science to solve this problem in a realistic manner. Furthermore, we include an appendix outlining a solution approach for incorporating active reservoir size learning through appraisal well options in ROA (see Appendix A). We propose a machine learning algorithm to determine whether and when drilling an appraisal well to optimize project value is beneficial.

We want to extend our warm appreciation to our main supervisor, Professor Verena Hagspiel. With her excellent competence and experience, she has created a thriving learning environment for us to develop. We would also like to thank our co-supervisors, Doctor Semyon Fedorov and Ph.D. candidate Olga Noshchenko for all their help, discussions, and insights related to real options and petroleum economics. Moreover, we would like to thank Sophie Haseldonckx at the Norwegian Petroleum Directorate (NPD) for contributing with expertise from the oil and gas industry and helping us to improve the industry relevance for this paper.

Norwegian University of Science and Technology
Trondheim, June 2023

Tobias Hyldmo

Markus Hoff Skudal

Abstract

The decreasing average size of oil and gas (O&G) discoveries on the Norwegian continental shelf (NCS) has increased interest in mitigating the uncertainties involved in O&G investment decisions. Marginal fields, in particular, will not be able to support the same data gathering as larger fields, as the cost of appraisal wells may outweigh the potential revenues of the field. This lack of information increases subsurface uncertainty, making the investment decision more complex. Therefore, handling this uncertainty through realistic modeling and accounting for managerial flexibility is essential when evaluating O&G investments. Using existing infrastructure through tiebacks is one of the most cost-efficient solutions for developing smaller O&G discoveries. As many maturing production facilities on the NCS obtain spare production capacity, tiebacks from marginal fields to these facilities can be an appropriate investment option.

This paper presents a compound real options analysis (CROA) that evaluates a portfolio of two marginal fields under reservoir- and oil price uncertainty. It allows us to identify the additional value from managerial flexibility when investing in tiebacks to an existing host. Based on updates of crude spot prices and sizes of the reservoirs, we apply a least-squares Monte Carlo algorithm (LSM) to evaluate whether to exercise the option to invest in either field or wait and reevaluate in consecutive years. A decision-maker has to choose which field should be tied back first while knowing there is an option to develop the other field afterward. We analyze the parameters influencing this choice and propose decision rules considering a portfolio view of O&G projects. The CROA approach is applied to a realistic case study resulting in improved decisions and, thereby, a return on investment of 25.4 % higher than the industry standard myopic valuation approach. This increase in value origins from the ability the CROA has to 1) better adjust the investment timing to what is favorable for the field given the current risk factors, 2) reduce downside risk by never investing if the field never seems profitable, and 3) incorporate information about the second tieback to better account for the total portfolio value. Overall, the paper provides a new perspective on real options analysis (ROA) of a portfolio of marginal oil fields.

Keywords - Real options, compound options, petroleum economics, marginal fields, reservoir uncertainty, price uncertainty

Sammendrag

Den synkende, gjennomsnittlige størrelsen på olje- og gassfunn på den norske kontinentalsokkelen har økt interessen for å redusere usikkerheten knyttet til investeringsbeslutninger innen olje- og gasssektoren. Marginale felt vil spesielt ikke kunne forsvare den samme innsamlingen av data som større felt, da kostnadene for brønnboring kan overstige de potensielle inntektene fra feltet. Denne mangelen på data øker usikkerheten på havbunnen, noe som gjør investeringsbeslutningen mer kompleks. Derfor er det essensielt å redusere denne usikkerheten gjennom realistisk modellering og å ta hensyn til tidsfleksibilitet når man evaluerer investeringer i marginale oljefelt. Bruk av eksisterende infrastruktur er en kostnadseffektiv løsning for utvikling av mindre olje- og gassfunn. På den norske kontinentalsokkelen er det mange produksjonsplattformer som får ledig produksjonskapasitet etter hvert som de blir eldre, og tilknytning (tiebacks) av marginale felt til disse installasjonene kan være et hensiktsmessig investeringsalternativ.

Denne artikkelen presenterer en analyse av sammensatte realopsjoner (CROA) som vurderer investeringer utsatt for usikkerhet knyttet til oljereservoaret og oljepriser for å identifisere den ekstra verdien av tidsfleksibilitet ved investeringer i en portefølje av to marginale felt. Basert på kontinuerlige oppdateringer av spotpriser og størrelser på reservoarene, bruker artikkelen en least-squares Monte Carlo-algoritme (LSM) for å vurdere om man skal innløse opsjonen til å investere eller vente og revurdere påfølgende år. En beslutningstaker må velge hvilket felt som skal knyttes tilbake først, hensyntatt at det finnes en opsjon til å utvikle det andre feltet senere. Vi analyserer parametrene som påvirker dette valget og foreslår derav beslutningsregler som tar hensyn til porteføljeperspektivet i O&G-prosjekter. CROA-tilnærmingen blir brukt på en syntetisk, men realistisk casestudie som resulterer i forbedrede beslutninger og en økt avkastning på investeringen på 25,4 % høyere enn bransjestandarden for verdivurderinger. Denne økningen i verdi stammer hovedsakelig fra CROA sin evne til å 1) bedre justere investeringstidspunktet til det som er gunstig for feltet gitt størrelsen på risikofaktorene, 2) redusere nedsiderisikoen ved å ikke investere hvis feltet ikke virker lønnsomt, og 3) inkludere informasjon om den antatte verdien på det andre feltet for å bedre ta hensyn til den totale porteføljeverdien. Totalt sett gir artikkelen et nytt perspektiv på realopsjonsanalyse av en portefølje av marginale oljefelt.

Nøkkelord - Realopsjoner, sammensatte opsjoner, petroleumsøkonomi, marginale oljefelt, reservoar-risiko, pris-risiko

Table of contents

List of Figures

List of Tables

Acronyms

ABEX	Abandonment expenditures
BBL	Barrels of drilling fluids (oil)
CAPEX	Capital expenditures
CF	Cash flow
CROA	Compound real options analysis
DA	Decision analysis
DCF	Discounted cash flow
E&P	Exploration and production (oil and gas)
ESG	Environmental, social, and governance
EUR	Estimated ultimate recovery (of oil from a field)
GBM	Geometric Brownian motion
LSM	Least-squares Monte Carlo
MAE	Mean absolute error
MC	Monte Carlo (simulation)
ML	Machine learning
MMUSD	Million United States dollar
MSE	Mean squared error
NCS	Norwegian continental shelf
NN	Neural network
NPD	Norwegian petroleum directorate
NPV	Net present value
NTNU	Norwegian University of Science and Technology
O&G	Oil and gas
OPEX	Operational expenses
PDO	Plan for development and operation
RMSE	Root mean squared error
ROA	Real options analysis
RoI	Return on investment
SD	Standard deviation
SHC	Spare host capacity
Sm^3	Standard cubic meters (SM3 in figures)
STLT	Short-term/long-term (two-factor price model)
SVM	vector-support machine
USD	United States dollar
VoI	Value of information

Selected symbols

t	a given year during the host's lifetime
t_0	the starting year in a given reference frame
T	time in years
$\mathbb{E}$	expected value
R_α	discrete discount rate
e	Euler's number
CF_t	cash flow throughout the year t
p_t	oil price at time t
q_t	quantity oil produced at time t
α	inflation rate
γ	continuous discount rate
J	production start of the first field
K	production start of the second field
ω	project lead time
ϕ	last investment year
L	host facility abandonment year
M	number of reservoir cases
N	number of simulations
β_0	fixed cost parameter (tariff/OPEX)
β_1	variable, oil quantity-dependent cost parameter (tariff/OPEX)
β_2	variable, oil price-dependent cost parameter (tariff)
χ_t	short-term oil price parameter
ξ_t	long-term oil price parameter
δ	minimum uncertainty of reservoir trajectories
$B(t)$	best guess reservoir size distribution
V	the "true" reservoir distribution
v*	"true" reservoir sample from V
$I(t)$	trajectory-specific information distribution created around $\nu*$
$i*$	reservoir information sample from $I(t)$
θ	initial uncertainty of $B(t_0)$
ϵ	reservoir learning rate
μ	mean of a given distribution
σ	standard deviation of a given distribution
λ	field depletion rate
F	value of the option to invest
X	state vector of a field
ν	decision variable for the timing of the first field
ψ	decision variable for the timing of the second field
Π	payoff function in option evaluation
k	second field value contribution bins
R^2	the statistical measure R-squared

1 Introduction

Since the first discovery of oil and gas (O&G) on the Norwegian continental shelf (NCS) in 1969, the industry has played a critical role in developing Norway's public welfare through its supply chains, social spill-over effects, and a tax rate of 78 percent (The Norwegian Tax Administration, 2022). Today the Norwegian petroleum industry is being debated and scrutinized from different perspectives in later years: the Norwegian government receives high revenues from O&G sales due to political conflict (The Economist, 2022) and a global need to shift towards sustainable and renewable energy sources (UN Secretary-General António Guterres, 2023). Despite global technological advancements, many industries and countries still rely on O&G as primary input factors and energy sources and are not yet in the position to be able to fully transition to electrified systems (The International Energy Agency, 2022*b*). The Norwegian O&G sector has a lower carbon footprint in production compared to most other producing regions in the world (Masnadi et al., 2018; The International Energy Agency, 2022*a*). Thus, one could argue that the Norwegian O&G sector plays a crucial role as a transitioning industry by offering hydrocarbons with a lower carbon intensity during the transition period in the global economy. This includes serving people's energy needs and applications that cannot be readily replaced by alternatives to fossil fuels.

The NCS has matured as a basin, and there is a clear trend of diminishing average sizes of O&G discoveries (see Figure 1). Following this trend, the interest in marginal field[1] investments and the subsequent decision process have increased. However, high investment costs related to delineation and infrastructure make traditional standalone development for marginal fields often economically unfeasible. One way to develop marginal fields is through tiebacks to existing facilities (The Norwegian Petroleum Directorate, 2019). These hosts can provide spare capacity to neighboring fields in return for economic compensation in the form of tariffs. When considering a decision to tie in a new field, time is a critical component due to the limited lifetime of the host. Thus, finding the optimal time to invest in a marginal field is highly important.

Specific to marginal fields, it is often very costly to gather additional data about the field, for instance, by drilling appraisal wells before issuing a Plan for Development and Operation (PDO). This results in greater exposure to uncertainty in terms of profitability for marginal fields compared to larger ones, which impacts the field's potential. This prompts the industry to find methods to improve decision-making for marginal field development under uncertainty.

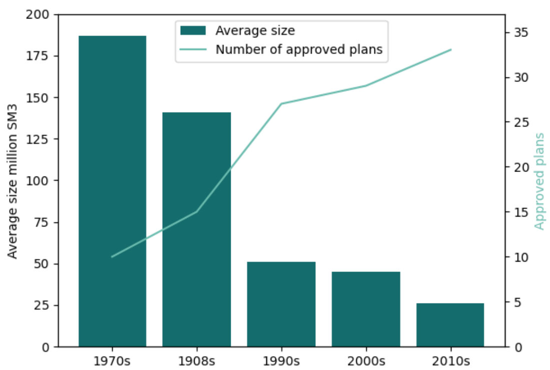

Figure 1: Decreasing size of discoveries on the Norwegian continental shelf. Visualization based on publicly available discovery data from The Norwegian Petroleum Directorate (2019).

[1]Marginal oil fields are characterized by discoveries that have remained unexploited for various reasons. This includes their limited reservoir sizes, high development costs, and technological limitations (Society of Petroleum Engineers, 2023).

In this paper, we consider optimal field development for a portfolio of marginal fields under uncertainty. Specifically, we study the optimal tieback timing of two marginal fields to one available host with capacity constraints under technical and market uncertainty. We aim to answer the following. Firstly, how can this novel problem modeling and solution approaches better can answer whether, when, and in what order to tie back the two fields accounting for more extensive uncertainty? Additionally, in that regard, how the industry can apply this methodology to make better decisions. Secondly, what key factors drive field selection for oil exploration and production (E&P) companies, enabling them to make optimal decisions? We apply a real options analysis (ROA) and adapt the widely recognized least-squares Monte Carlo simulation (LSM) solution approach proposed by Longstaff and Schwartz (2001) to solve these questions.

We develop a methodology that allows valuing tieback investments while accounting for uncertainty about future market prices and the estimated ultimate recovery (EUR) of the fields. Firstly, we model future oil prices by using simulations of the two-factor price model of Schwartz and Smith (2000). Secondly, we use continuously updated probability distributions of the EUR using Bayesian updating to construct production rates. In order to model production rates resulting from the Bayesian updating, we present an optimization model that accounts for constraints related to spare host capacity, recoverable volume, and decline rate. The resulting production rates are then, together with the market price simulations, used to calculate the project's net present value (NPV). To address the recurring evaluations of whether and when to tie back the two marginal fields, we apply our custom LSM algorithm. Finally, we apply compound options to the ROA framework. This makes a decision-maker able to concurrently and better evaluate whether, when, and in which order the marginal field should be tied back to a host facility to maximize the expected return on investment (RoI). The ROA framework allows us to quantify the value of timing flexibility in terms of the tieback of the two fields. Additionally, it allows us to construct decision rules for management to use in their investment evaluation. The proposed methodology provides practical implications for stakeholders in the O&G industry, making it highly relevant for regulatory authorities such as the Norwegian Petroleum Directorate (NPD) and E&P companies.

This paper presents a valuable contribution to the literature by addressing the portfolio challenge of tieback investments under market price and reservoir uncertainty. This research extends to two important aspects of the literature: 1) bridging the gap between market price and technical uncertainty in a ROA by accounting for both at the same time and 2) proposing a novel ROA framework applying compound options to address the problem of optimally allocating a project portfolio of marginal fields under uncertainty. The contributions take an important step in accounting for market and reservoir uncertainty utilizing our Bayesian learning modeling. We denote the framework using compound options in ROA as compound real options analysis (CROA).

In order to validate the methodology, we compare our results to those resulting from a myopic decision approach. We find that considering the optimal timing and the order of tiebacks for the two fields has considerable value and allows us to reduce downside risk. Our findings consistently demonstrate that flexibility holds greater value in the face of increased project risk. The robustness of our ROA approach, combined with our risk modeling, ensures reliable and meaningful results. The exploration of compound options for assessing the marginal field project portfolio has generated interesting findings, opening up paths for future research in the field of ROA.

1.1 Literature review

This paper seeks to contribute to two key research domains. Firstly, we investigate the optimal timing for developing tieback projects in marginal oil fields. Our approach integrates ROA to account for managerial flexibility under oil price and reservoir uncertainty, ultimately maximizing the project's value. While previous works such as Fedorov et al. (2021, 2022a); Jafarizadeh and Bratvold (2015, 2020); Sales et al. (2021); and Lin et al. (2013) have addressed closely related aspects, none have encompassed all the aforementioned elements in their research. Secondly, our work contributes to the literature by extending a classical ROA framework to a portfolio perspective using compound real options. To the best of our knowledge, few studies have analyzed the optimal allocation of a portfolio of O&G projects. Lin et al. (2013) presents a portfolio perspective that accounts for multiple risk factors but does not utilize a ROA approach. We propose a novel methodology for concurrently considering multiple fields, and we denote it as CROA, a compound real options analysis framework.

Real options originated from the application of traditional financial options valuation methods to project valuation. The term was first introduced by Myers (1977) and has since gained widespread recognition for its ability to assess projects with significant risks and emphasize the value of incorporating risk through timing and managerial flexibility by viewing them as options. Tourinho et al. (1979) was the first to apply ROA to the O&G industry, estimating the value of natural resource reserves while accounting for price uncertainty. In 1988, Paddock et al. (1988) and Ekern (1988) used ROA for O&G projects. Paddock et al. (1988) compares ROA to the traditional discounted cash flow approach (DCF) and highlight its advantages, including the need for less proprietary data, lower computational costs, reduced error susceptibility, guidance on investment timing, and insights for government policy and company behavior. Ekern (1988) demonstrate how an ROA gives different results than a DCF approach due to accounting for managerial flexibility when choosing the timing of investment. His focus was on options related to the development and operation of tiebacks for satellite fields. Dixit and Pindyck (1994) further expands the flexibility of the ROA framework through a comprehensive review explicitly formulated using differential equations. They concluded that ROA and Decision Analysis (DA) produce similar results when DA can explicitly value the project. DA is a method used for decision-making, evaluating, and comparing potential options, risks, and benefits, and is often employed for similar problems as ROA. Galli et al. (1999) compare three solution approaches for ROA: Monte Carlo simulations (MC), option pricing, and decision trees. Together with dynamic programming, these are the most commonly used solution methods. This paper adopts a simulation-based approach to address our decision problem, building upon the abovementioned studies.

ROA has gained significant traction in the O&G industry in recent decades, specifically for E&P projects, due to their long time horizons, substantial capital expenditures (CAPEX), and the ability to incorporate multiple uncertainties. Bjerksund and Ekern (1990) expand on previous approaches, exploring various investment and operational options and their impact on project value. Dias (1997) integrate market price and technical uncertainty in ROA, while Cortazar and Schwartz (1998) assess the optimal timing for oil field production. Galli et al. (2001) study drilling decisions, and Chorn and Shokhor (2006) use dynamic programming with ROA for exploration opportunities. Jafarizadeh and Bratvold (2009, 2012, 2015) provide insights on timing, managerial decision-making, and field sizes. Overall, ROA proves valuable for its ability to optimize timing and valuation in the O&G sector.

MC simulation, developed by John von Neumann and Stanislaw Ulam during World War II, has been widely utilized in various research fields for many years (Eckhardt, 1987). Tilley (1993) pioneers its use in valuing American Options, which laid the foundation for simulating real options. Before this development, ROA faced several challenges. Lander and Pinches (1998) identifies the most significant hurdles: limited understanding of ROA outside academia, violations of modeling assumptions, and the inapplicability of additional mathematical constraints in real projects. A significant breakthrough came with the introduction of the least squares Monte Carlo (LSM) framework by Longstaff and Schwartz (2001). Utilizing a two-factor price model by Schwartz and Smith (2000) to account for oil price uncertainty, their work propelled LSM as a state-of-the-art evaluation method for real options. LSM enables unbiased consideration of multiple uncertainties without excessive computational demands (Fedorov, 2021). Notable works employing LSM to explore investment under uncertainty in the O&G sector include those by Fleten et al. (2011); Jafarizadeh and Bratvold (2012, 2015), and Hong et al. (2019).

Reservoir uncertainty is a critical factor in evaluating the value of E&P projects, as the extractable volume of a field highly impacts the revenues and costs. Dias (1997) pioneered the inclusion of both price and technical uncertainty in ROA for E&P projects, emphasizing the importance of technical uncertainties in exploration decisions and economic uncertainties in production decisions. Galli et al. (2001) introduce a sequential updating approach using appraisal wells to mitigate technical uncertainty, while Armstrong et al. (2004) incorporate Bayesian updating and copulas to refine reservoir volume distributions. Santos et al. (2017, 2018) focus on robustness in well quantity and placement, while Jafarizadeh and Bratvold (2020) update scenario probabilities with new well drilling. Modeling reservoir size and production rates involved probabilistic and stochastic analysis by Chang and Lin (1999) and Subbey et al. (2004), while Lin et al. (2012) have a mathematical model to capture evolving reservoir uncertainty. Inspired by these works, we simulate different realizations of our marginal field using Bayesian updating and an analytical approach similar to Lin et al. (2012, 2013)'s method. We utilize a normal-normal conjugate pair of Bayesian updating and incorporate it into a ROA framework.

As with reservoir uncertainty, O&G prices are also a main uncertainty driver of future revenues in E&P projects. A common challenge in the industry is that O&G prices often are simplistically represented in traditional valuation methods (Bickel and Bratvold, 2008). This motivates us to use a stochastic price model to replicate the development of future oil prices. Early adoptions of price uncertainty in ROA were to model the price of the O&G commodities as a GBM (Brennan and Schwartz, 1985; Paddock et al., 1988). Later, studies like Pindyck (1999); Smith and McCardle (1999) propose that commodities have a mean-reverting price property, as they often are balanced by the mechanisms of supply and demand. Newer research combines the two abovementioned properties in factor models, with two or more factors, to even further increase the realism of the modeling. In this paper, we use the short-term/long-term (STLT) two-factor stochastic price model as of Schwartz and Smith (2000) for modeling the oil prices. This model is calibrated on historical oil prices using a Kalman filter. Including the two-factor model of Schwartz and Smith (2000) into a ROA framework has become a standard in most of the recent literature on ROA that include price uncertainty (Jafarizadeh and Bratvold, 2015; Fedorov et al., 2021; Bakker et al., 2021). We apply the STLT two-factor model calibrated using a Kalman filter in our ROA.

Despite the popularity of ROA in E&P projects, the academic literature on ROA with marginal field development is relatively sparse (Fedorov, 2021). For example, Laine (1997) analyze Norwegian marginal fields and demonstrates ROAs relevance when valuing marginal fields. Lund (1999, 2000) emphasize the value of operational flexibility in marginal oil fields, while Stoisits et al. (2010) recommend tiebacks as the most profitable approach. Fleten et al. (2011) evaluate a tieback offshore development with expansion and early decommissioning options. Lin et al. (2013) explore subsea tieback developments under uncertainty and showed how a portfolio view of this outperforms a myopic benchmark with 76% in value. Jafarizadeh and Bratvold (2015) examine smaller, steeper-decline marginal fields and compare their sensitivities to larger fields. Other studies focus on mathematical modeling approaches (Lei et al., 2021) and combining ROA with Decision Analysis (Fedorov, 2021). Additionally, Fedorov et al. (2022a) investigate tiebacks and their timing for marginal fields. Lin et al. (2013) argue many smaller discoveries only make sense to evaluate as a set. We extend this reasoning by addressing a portfolio view of marginal fields using the CROA framework.

Evaluating project investments under uncertainty has been under scrutiny for a long time. The seminal work on portfolio theory by such of Markowitz, Sharpe, Fama, and French[2] provides a foundation for understanding the benefits of diversification and risk management in project selection and allocation of resources. In the world of physical investment projects, like allocating a portfolio of investment opportunities in O&G, we find the research area of incorporating managerial flexibility into the investment opportunities as under-examined. Anand et al. (2007) contributes to the literature by exploring the concept of option portfolios, considering factors such as project interdependencies and resource constraints. Lin et al. (2013) offer insights into the dynamic nature of project portfolio selection, considering evolving uncertainties and decision-making over time while using MC to simulate NPV outcomes. Fedorov et al. (2022a) provide a recent perspective on portfolio optimization in the context of ROA, though optimizing the best total combination of individual real option evaluations. Gamba (2003) offers a valuable and broad examination of formulating ROA to be solved with LSM, including how to incorporate compound options. We apply a compound real options approach equaling Gamba's formulation of *options on options* to adapt the LSM in our CROA framework.

This paper uses insights from the papers mentioned in this section. Drawing from the real options literature, we incorporate ideas from papers such as Fedorov et al. (2021, 2022a), Jafarizadeh and Bratvold (2015), and Longstaff and Schwartz (2001), which provide frameworks for determining the optimal timing of portfolio investments through ROA and LSM. Additionally, we extend the concept of marginal field tiebacks, as discussed in Lund (2000) and Fleten et al. (2011). We use the STLT price modeling as proposed by Schwartz and Smith (2000), heavily applied throughout the literature. We incorporate the principles of Bayesian updating as applied to the estimation of recoverable volumes, as demonstrated in Lin et al. (2013), Armstrong et al. (2004), and Bhattacharya et al. (1986). Lastly, taking on a multiple-field portfolio perspective, we draw inspiration from works by Gamba (2003), Lin et al. (2013), and Fedorov et al. (2022a), exploring how uncertainty influences the value of portfolios comprising marginal fields.

[2]These influential fathers of portfolio theory enhanced the understanding of portfolio allocation, risk, diversification, and valuation. Despite their work with liquid and market-priced assets, the application of their theories still bleeds into project investment (Markowitz, 1991; Sharpe, 1963, 1964; Fama, 1970; Fama and French, 1993).

2 Background

The relatively higher uncertainty in the development of marginal fields highlights the need for innovative approaches for valuation and optimal investment timing. Subsea tiebacks are crucial for developing marginal fields, connecting them to facilities over long distances and deep waters. Despite technical challenges such as capacity constraints and flow assurance, they are often the only economically viable solution[3]. According to Scott (2017); Maslin (2019), and Bacati et al. (2020), the CAPEX for tiebacks depends on various factors, but ongoing technology developments aim to improve productivity and reduce costs even further. Compared to other methods, tiebacks may offer a 30% to 60% reduction in CAPEX. Figure 2 illustrates an example of a tieback connecting a field to its host facility in the Gulf of Mexico.

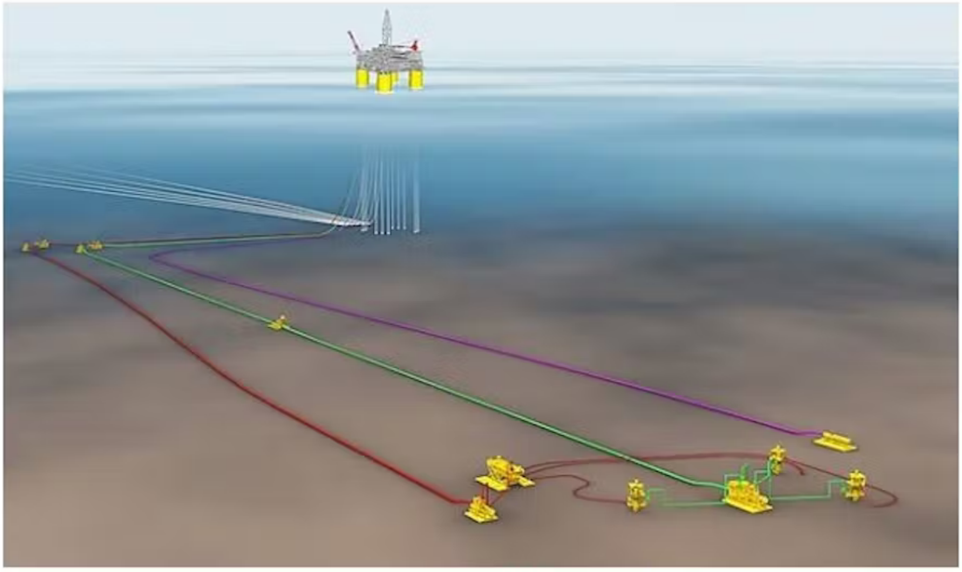

Figure 2: Visualization of a subsea tieback solution. Showing the PowerNap tieback to Olympus in the Gulf of Mexico. The Figure is retrieved from Dittrick (2019) published in the OIL & GAS JOURNAL.

Historically, the O&G industry standard for project evaluation only considers a limited number of scenarios for recoverable volumes of a field, these being the base-, high-, and low-case scenarios. Often this is not the fault of reservoir engineers but rather the link between their technical modeling and economic management. This simplification limits decision-making, as it leaves out useful information. To overcome this limitation, this paper adopts the approach of representing reservoir uncertainty as evolving probability distributions, as in Armstrong et al. (2004) and Lin et al. (2012). To model the uncertainty in recoverable volumes of reservoirs, the literature presents a few different approaches (Ian and Noeth, 2004; Sales et al., 2021; Lin et al., 2012). Some studies, like Armstrong et al. (2004), use Bayesian updating to account for volume uncertainty while using a simple triangular probability distribution, which is updated when learning about the field. We incorporate Bayesian learning, similar to Lin et al. (2012), applying Bayesian updating to create a reverse Wiener process with jumps. We will refer to the potential oil that can be extracted from a field as the estimated ultimate recovery (EUR), thus modeling evolving EUR distributions.

[3]For further reading, technical challenges with tiebacks have been studied previously by Lin et al. (2013); Lei et al. (2021), and Husy (2011)

The O&G industry faces high uncertainty in revenues due to unpredictable oil prices and recoverable reserves. Price volatility is attributed to the low responsiveness of supply and demand to short-term price changes (US. Energy Information Administration, 2022). Thus, O&G prices are relatively more volatile compared to other commodities (World Bank Group, 2014). DCF analysis, resulting in project NPV estimates, is the industry standard for O&G project valuation. In this approach, the expected cash flows of a project are discounted and summed over the lifetime of the project. However, this simple and intuitive method has limitations when the cash flows are subject to high risks. The NPV formula is given by:

$$NPV = \sum_{t=t_0}^{T} \frac{\mathbb{E}(\mathrm{CF}_t)}{(1+R_\alpha)^t} \tag{1}$$

where $\mathbb{E}(\mathrm{CF}_t)$ is the expected cash flow of period t, t_0 is the timing of the investment decision, T is the number of periods, and R_α is the discount rate that risk-adjusts the project for the time value of money and risk (Brealey et al., 2012). We will refer to the DCF approach producing an NPV estimate as the myopic approach.

The myopic approach is a straightforward method for calculating project value in stable and predictable conditions. However, it has limitations, as it assumes that all decisions are made with only the available information at the time of evaluation. Research by Paddock et al. (1988) and Cortazar and Schwartz (1998) shows that waiting for more information can add significant value to investments in O&G projects. This idea is supported by Fedorov et al. (2021) and other studies on flexible strategies. Infrastructure investments are often irreversible, making it unlikely to recover the entire capital amount if resold (Dixit and Pindyck, 1994). Irreversibility in the O&G industry, combined with large technical and market risks, increases the value of flexibility and the incentive to utilize other methods.

The ROA approach is a more precise method to evaluate project value than the myopic approach when dealing with uncertain revenues, costs, and irreversible investments (Dixit and Pindyck, 1994). This approach factors in managerial flexibility and waiting for new information, which is crucial in the O&G industry with high uncertainty. Therefore, we adopted ROA in our study, which builds on DCF calculations as part of the methodology. The NPV of the tieback solution used further in a ROA is given by:

$$NPV_{field} = \sum_{t=t_0}^{T} (q_t^{oil} \cdot p_t^{oil} - \mathrm{CAPEX}_t - \mathrm{tariff}_t - \mathrm{OPEX}_t - \mathrm{ABEX}_t) \cdot e^{(\alpha-\gamma)\cdot t} \tag{2}$$

where t is the given year, t_0 is the investment year, T is the life expectancy of the host, the quantity of oil produced is q_t^{oil}, the oil price is given by p_t^{oil}, CAPEX_t is a fixed investment cost for developing each field, tariff_t is the bilateral lease agreement between the host and field operator, OPEX is the variable cost of operating the field dependent on q_t^{oil}, ABEX_t is a fixed abandonment cost for decommissioning production infrastructure, α is the inflation, and γ is the cost of capital. CAPEX_t and ABEX_t only apply at $t = t_0$ and $t = T$, respectively. In contrast, the host's expected profits come from the annual tariff paid by all partnering field operators. Hence, the NPV for the host will be:

$$NPV_{host} = \sum_{fields} \sum_{t=t_0}^{T} \text{tariff}_t \cdot e^{(\alpha-\gamma)\cdot t} \tag{3}$$

We extend the use of ROA to evaluate both market and technical uncertainty, bridging the technical and economic perspectives of marginal field development in our paper. Therefore, this paper aims to position itself at the crossroads between the technical and economic perspectives of marginal field development in O&G. We further explore the untapped potential of analyzing marginal field investment opportunities as compound real options, referring to unlocking the opportunity to tie back the second field once you have invested in the first one.

In a problem involving multiple hosts and fields, stakeholders are anticipated to have diverse objectives and advocate for their own favorable conditions. In our problem perspective, we assume the host companies' economic objective is to maximize the tariff received from the fields, limited by what is regulated. Thus, the host is incentivized to tie in as many marginal fields as early as possible. On the other hand, NPD would, from a regulator perspective, aim to maximize the total value creation in the area. We simplify this to the sum of the value of the fields and host facility. Thus, NPD's perspective does not necessarily need to align with the host's. When we specifically examine a case where two companies hold licenses for all hosts and marginal fields, the complexity of the problem is reduced. Consequently, the companies share a mutual interest in connecting as many fields as possible, with minor deviations in timing being of relatively minor significance. However, if different combinations of controlling entities were involved, it could lead to misaligned interests and, thereby, introduce elements of competition.

Typically, joint ventures are given the right to operate licenses on the NCS. As a simplification, we will in this paper refer to the licensees of the host facility as the host or host operator(s) and the licensees of the discovery and developers of a field as the field operator(s).

The paper proceeds with a problem description in Section 3. The methodology in Section 4 covers the general premises, modeling, and solution approach. Moreover, the case study in Section 5 introduces parameters to apply for the methodology before results and sensitivities are presented and discussed in Section 6. Limitations of our work are discussed in Section 7 before concluding with potential future work and a summary in Section 8.

3 Problem description

To solve the portfolio problem, we consider the partial decision problem with the opportunity to invest in two production facilities exposed to technical and market uncertainty. We subsequently evaluate the expected return on investment (RoI) from our proposed decision sequence compared to a myopic benchmark. The modeling approach of the problem is chosen to highlight the novelty of using compound options with both market and technical uncertainty. Our modeling of the problem realistically allows decision-makers the flexibility to invest in any of the two projects at any time step, thus making a substantial contribution to solving the project portfolio allocation problem in the literature. We present our CROA solution approach in Section 4, enabling E&P companies to consider investments in multiple assets ahead of time, leading to more optimal decisions.

We consider the decision situation with one host facility and two nearby marginal fields, which we will refer to as Field A and Field B, respectively. In the context of our study, we make the assumption that the two fields are technically compatible for tiebacks to the infrastructure of the host. The problem situation is illustrated in Figure 3. We assume one operator to own and operate the host facility (Company I) and a separate operator owning and considering developing the licenses to the marginal fields available (Company II). Later in this paper we simply refer to these as the *host* and the *field operator*, respectively. This removes the complexity of competition in the problem and presents a clearer image of the problem aspects we focus on. Three terms of factors restrict the tieback decision. These are the SHC of the host, the market price uncertainty of future oil prices, and the reservoir uncertainty of the size of each field in the portfolio.

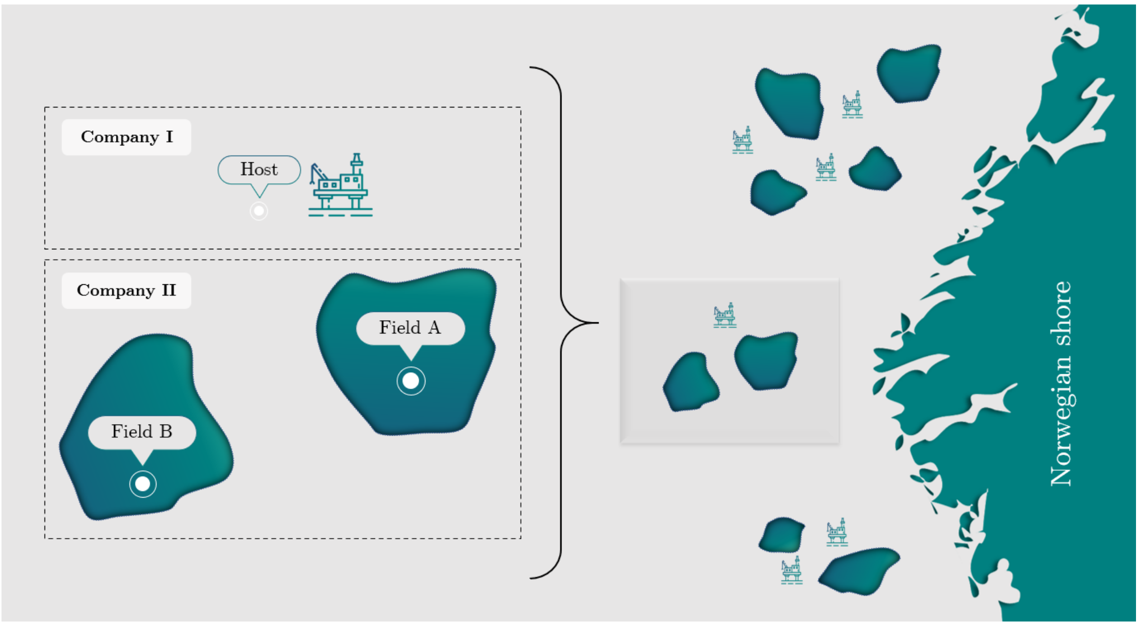

Figure 3: Problem situation with one host facility and two nearby marginal fields. The host and the fields are controlled by different entities.

The host opens up for tiebacks toward the end of its life when it gradually frees up its SHC [4]. The field operator has to decide *whether* and *when* to invest in either or both of the marginal fields, free of any ordering constraint. The investment decision is reconsidered once a year, accounting for the updated external uncertainties and the host's SHC. In Figure 4, we illustrate the timeline of the decision problem. *Phase 1* and *Phase 2* describe the decision sequence of the problem of connecting the first field and the second field, respectively. We examine the time period starting from when the host possesses spare capacity until the end of its operational lifespan (L). The final year for any investment is $L - \phi$ years before the host's abandonment. In mathematical terms, we can write this as $t \in [1, \phi], \phi < L$, where ϕ is the last year investment from the field operator can occur, and L is the host's lifetime.

[4] A host usually experiences SHC after some years of production, depending on its EUR, as the original production declines (The Norwegian Petroleum Directorate, 2022).

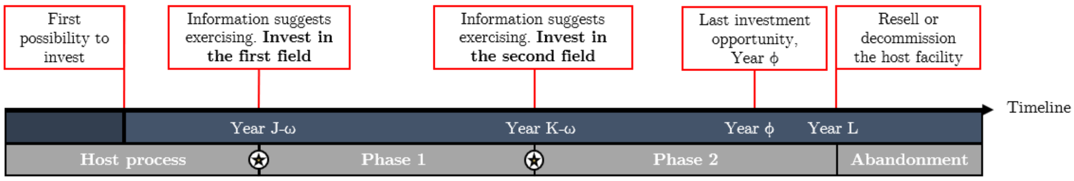

Figure 4: Decision problem timeline. It depicts two project phases, one for each field being connected.

In Figure 5, we illustrate the decision problem in more detail with visualized production rates. At each year $0 < t < J - \omega$, the field operator receives information regarding risk factors and can choose whether to exercise the option to invest in the first tieback development based on the new information. The risk factors include the current oil price in the market, an updated EUR estimate, and an expected value contribution for connecting the second alternative field later. We denote the investment decision timing of tying back the first field by $t_1 = J - \omega$. At year $t_2 = K - \omega$, the second field is chosen to be tied back to the host based on its updated EUR and oil price information. Notably, t_1 and t_2 can be inapplicable if investment conditions are unfavorable. We assume oil production starts with a lead time of ω years later for both fields; thus, the first and second filed start producing at time $t = J$ and $t = K$, respectively.

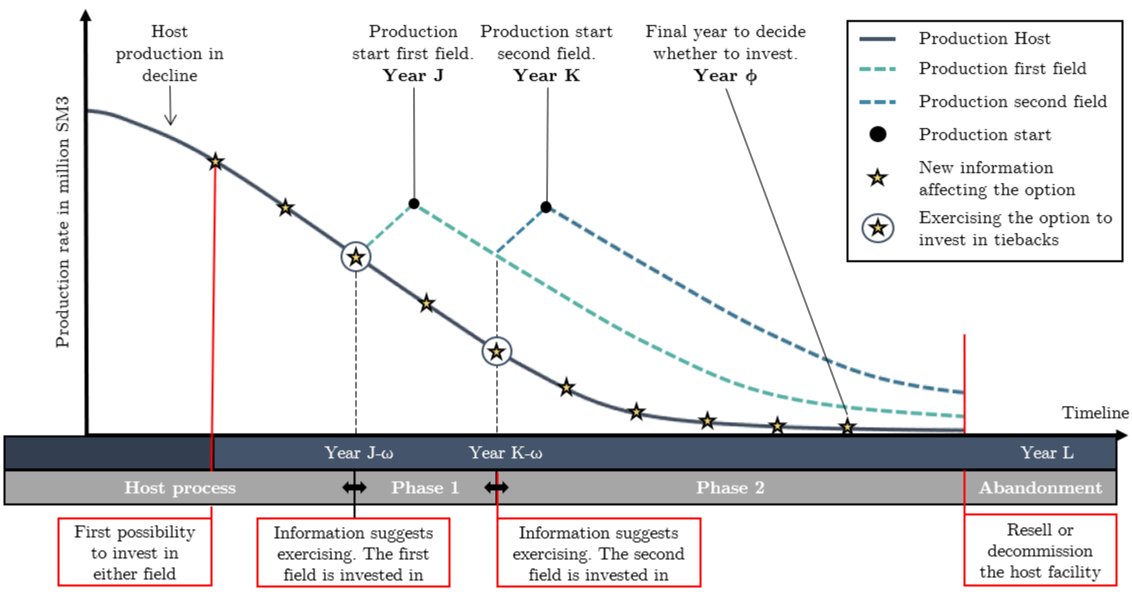

Figure 5: Detailed decision problem timeline with production overview. In the context of two production phases, we provide an overview of the sequence of decisions in our problem.

To make our model computationally feasible, we have made several assumptions. Most notably, we assume that the field operator begins learning passively about reservoir uncertainty for the second field only. This means we model reservoir uncertainty differently for the two fields depending on their order and characteristics, as described in detail in Section 4.3. We assume the SHC of the host to be known and deterministic before any tiebacks. As a simplification, we model it such that the host has shared the fixed terms for utilizing its spare capacity with the field operator(s) with licenses for the neighboring marginal fields. Thus, the field operator decides the timing of tiebacks to its fields in our model.

Figure 6 exemplifies a case where the option to tie back is exercised for Field A and Field B with a two years difference. For this example, we can assume Field A is tied back first. The production of Field B is choked[5] in the first years, whereas Field A's production still takes up most of the available capacity. In this case, the combined volume of both fields is extracted prior to the decommissioning of the host in year L.

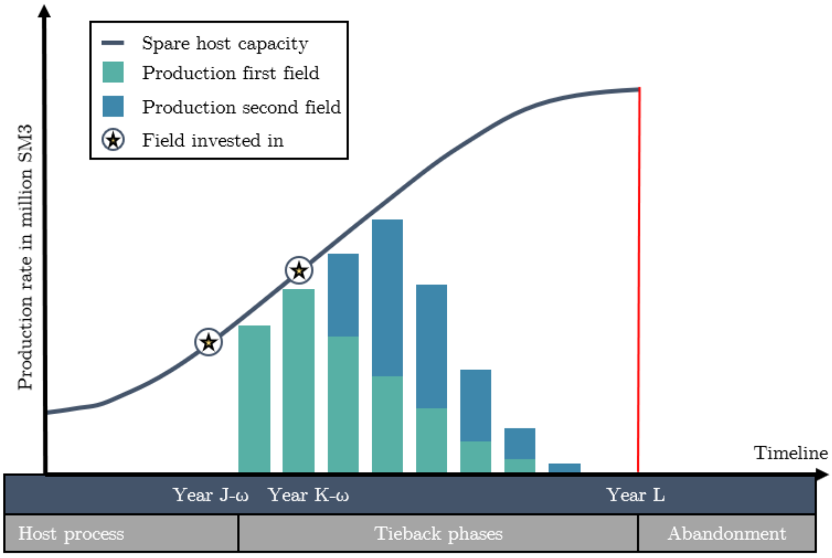

Figure 6: Example of a decision sequence. It exemplifies a total production volume when both marginal fields are developed. The first field's production is shown in teal, after receiving investment in year $J - \omega$ and starting production in year J. The second field is invested in at year $K - \omega$ and starts production in year K.

4 Methodology

This section describes how we solve the decision problem. Section 4.1 introduces the overall model and framework for estimating the value and optimal tieback timing of connecting Fields A and B. Then we present the market price model in Section 4.2 and explain how reservoir uncertainty and optimization of the production profiles are modeled in Section 4.3 and Section 4.4. The integration of all components in the CROA, including the valuation, optimal timing, and field order within our project portfolio, is detailed in Section 4.5 using the LSM method as a solution approach. Lastly, Section 4.6 describes the benchmarks we use to evaluate our proposed model.

4.1 Solution approach to the two-field decision problem

In the following, we introduce the proposed model to optimally evaluate whether and when to tie back either Field A, Field B, both, or neither of them to the host facility. The possibility of tying back a field under uncertainty regarding future framework conditions is similar to an investment option and, therefore, eligible to be evaluated with a ROA. Since ordering the two fields is flexible, the case can be viewed as two compound options where either is connected first. To the best of our knowledge, this is a novel usage of compound options in a real options framework accounting

[5]When a field holds an unrestricted production rate surpassing the host's extraction capacity, taking control of the oil extraction rate becomes viable. This can be achieved by implementing measures to restrict the flow of oil, thereby preventing excessive production. This is called to choke the field.

for project portfolio allocation. Going forward, we will denote this approach as a compound real options analysis (i.e. CROA), shown in the gray *CROA* boxes in Figure 7.

Due to its computational demand, the reservoir uncertainty of the first field is strenuous to consider as anything other than a static value. Thus, to account for reservoir uncertainty for the first field, we run M cases shown in Figure 7, one for each EUR guess of the first field. We return to a detailed framework described in the real options section in Section 4.5.

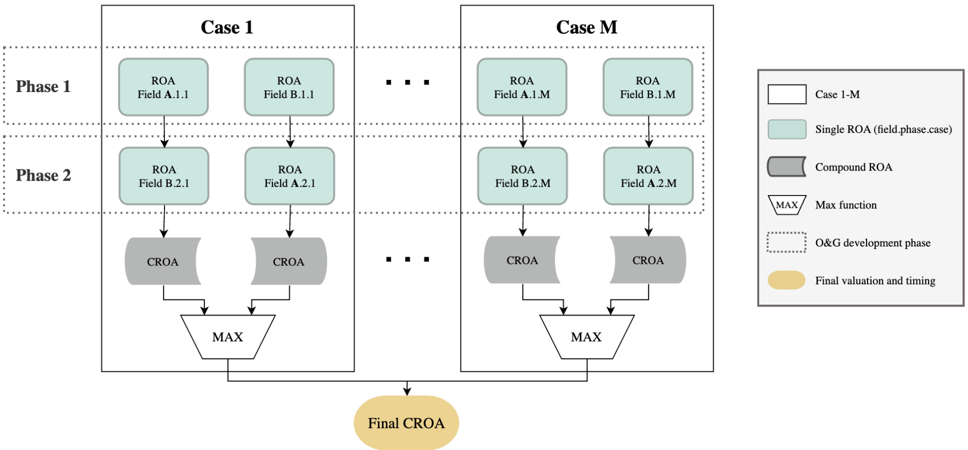

Figure 7: Total compound real options valuation and timing framework (CROA). We model two compound options to account for both possible order sequences of the fields. We can extend the framework with M cases to account for multiple first-field volumes.

The structure and components of our Single ROA approach are illustrated in Figure 8. This model is equivalent to each of the teal *ROA boxes* showing individual ROAs in Figure 7 with different input parameters. We start with the modeling of the size of the field, equivalent to the reservoir engineers' best guess. This is assumed to be constant for the first field while represented as a probability distribution for the second field, as we have started to learn about its size after tying back the first. Utilizing our present belief of the reservoir EUR, we construct the production rates for the field by applying a production optimization model. The objective of this model is to maximize the potential production rate, considering the constraints imposed by the available capacity of the host and the declining field deliverability of oil, which will be explained in Section 4.3 and Section 4.4. Thirdly, we model the project's cash flows and deploy an LSM as a solution approach to evaluate the project, as described in Section 4.5. Finally, we solve the ROA with the LSM using N simulations to see the marginal fields under different price- and reservoir realizations.

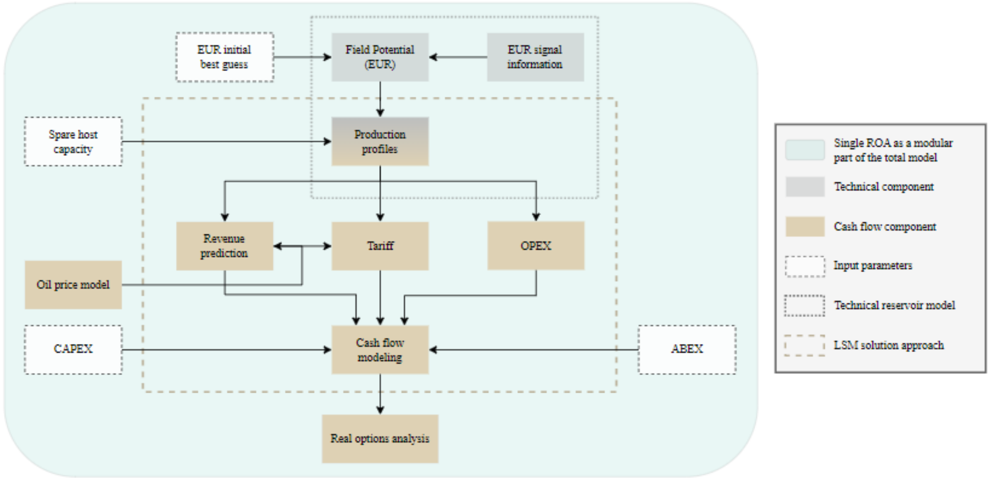

Figure 8: The Single real options analysis (ROA). The approach acts both as an independent solution approach and as a modular part of the CROA framework.

Figure 8 illustrates the different components of the Single ROA. Within the dashed, gray circumference, we observe the components of the reservoir model, as described in Section 4.3. The estimated *production profiles* influence the following *revenue prediction*, tariff, and OPEX. The revenue results from the field operator's sale of O&G and is the product of price and quantity in the *revenue prediction* component. The oil price is simulated with the Schwartz and Smith (2000) two-factor model, as described in Section 4.2.

The tariff represents a mutual agreement between a field operator and a host operator, as they are separate entities in our case. The tariff function, illustrated in Equation 4, follows a commonly employed format in practical applications (Fedorov et al., 2022a). In this function, β_0 denotes a fixed fee, while β_1 and β_2 represent coefficients for the oil volume and oil price, respectively:

$$\text{tariff}_t = \beta_0 + \beta_1 \cdot q_t^{oil} + \beta_2 \cdot p_t^{oil} \tag{4}$$

We model OPEX to encompass multiple components that combine various operational expenses into an annual cost. Primarily, these costs include staffing, fuel, facility maintenance, and leasing of storage vessels. This approach imitates a realistic representation of the overall operational expenditures. As these costs normally increase with higher production volumes, we assume OPEX consists of a fixed and an oil quantity-dependent component similar to Fedorov et al. (2021), and is represented as the components β_0 and β_1, respectively. Hence, the OPEX is given by:

$$\text{OPEX}_t = \beta_0 + \beta_1 \cdot q_t^{oil} \tag{5}$$

In our modeling, both CAPEX and ABEX are regarded as one-time costs associated with the establishment and abandonment of production infrastructure, with the possibility of ABEX being positive in certain cases. We assume that OPEX and tariff are influenced by the production profiles, as larger production volumes naturally entail higher costs for managing increased product quantities. However, for the sake of simplification, we assume that CAPEX and ABEX remain unaffected by the production rates.

4.2 Oil price modeling

In this section, we detail our approach to modeling market uncertainty in our analysis, with a focus on the risk associated with oil prices. To simulate this uncertainty, we employ a stochastic price model while conducting the ROA. The paper utilizes the Schwartz-Smith two-factor model to represent the future oil price. This model incorporates short-term deviations from a constant long-term equilibrium, as suggested by Schwartz and Smith (2000). This approach accounts for both a short-term mean-reversion effect and uncertainty in the long-term equilibrium. The mathematical expression for the Schwartz-Smith two-factor model is as follows:

$$\ln(S_t) = \chi_t + \xi_t \tag{6}$$

where χ_t is the short-term deviation from the long-term equilibrium ξ_t. This deviation allows for short-term mean-reverting behavior, which enables producers to respond to changes in market conditions. On the other hand, the long-term equilibrium prices are modeled as Brownian motion. This modeling accounts for adjustments to changes in supply and demand, macroeconomic factors, technological developments, and other relevant variables over the long term.

$$d\chi_t = -\kappa\chi_t dt + \sigma_\chi dz_\chi \tag{7}$$

The short-term deviations from the long-term equilibrium price are described in Equation 7. These deviations are expected to converge to zero over time, like an Ornstein-Uhlenbeck process. The parameters of this process include the mean-reversion coefficient κ and the volatility parameter σ_χ.

$$d\xi_t = \mu_\xi dt + \sigma_\xi dz_\xi \tag{8}$$

The long-term equilibrium price is described in Equation 8. This is modeled as a Brownian motion process with a drift rate μ_ξ and volatility σ_ξ. The correlated increments of the Brownian motion processes are represented by dz_ξ and dz_χ and are correlated with $dz_\xi dz_\chi = \rho_{\xi\chi} dt$.

$$d\chi_t = (-\kappa\chi_t - \lambda_\chi)dt + \sigma_\chi dz_\chi^* \tag{9}$$

$$d\xi_t = (\mu_\xi - \lambda_\xi)dt + \sigma_\xi dz_\xi^* \tag{10}$$

In our analysis, we adopt a risk-neutral pricing approach to account for uncertainty in the investment's market and technical aspects, as proposed by Cox et al. (1985). This approach enables us to risk-adjust the cash flows, and we make adjustments to Equation 7 and Equation 8 as shown in Equation 9 and Equation 10. In these adjusted equations, dz_ξ^* and dz_χ^* represent the correlated increments of standard Brownian motion processes, and they are correlated by $dz_\xi^* dz_\chi^* = \rho_{\xi\chi} dt$. We introduce risk premiums λ_ξ and λ_χ, subtracted from the mean μ to adjust for risk. In the short term, the risk-adjusted premium reverts to $-\lambda_\chi/\kappa$. In the long term, we obtain a drift of $\mu_\xi^* = (\mu_\xi - \lambda_\xi)$.

$$\chi_t^* = \chi_{t-1}^* e^{-\kappa\Delta t} - (1 - e^{-\kappa\Delta t})\frac{\lambda_\chi}{\kappa} + \sigma_\xi \epsilon_\xi \sqrt{\frac{1 - e^{-2\kappa\Delta t}}{2\kappa}} \tag{11}$$

$$\xi_t^* = \xi_{t-1}^* + \mu_\xi^* \Delta t + \sigma_\xi \epsilon_\xi \sqrt{\Delta t} \tag{12}$$

We must discretize the prices and cash flows when conducting the Monte Carlo simulations. The price modeling's short- and long-term components are discretized as shown in Equation 11 and Equation 12. In each time period, ϵ_ξ and ϵ_χ represent correlated standard normal random variables with a coefficient of correlation $\rho_{\xi\chi}$.

Summarizing this subsection, we have two initial values, ξ_0 and χ_0, and seven parameters ($\kappa, \omega_\xi, \omega_\chi$, $\mu_\xi, \lambda_\chi, \lambda_\xi, \rho_\xi\chi$) that need to be estimated. These are usually estimated through a Kalman filter [6]. As of Goodwin (2013), the Kalman filter is used to tune the parameters for a set of historical spot or future price data. The output is a posterior conditional distribution of the historical data. Figure 9 visualizes our price simulation by percentiles after 2015. We adopt the two-factor modeling approach because it is considered superior to one-factor models, as it can account for both short- and long-term uncertainty. Additionally, the two-factor model is relatively easy to calibrate and communicate (Fedorov et al., 2021). The parameters used in our case are further described in Section 5.

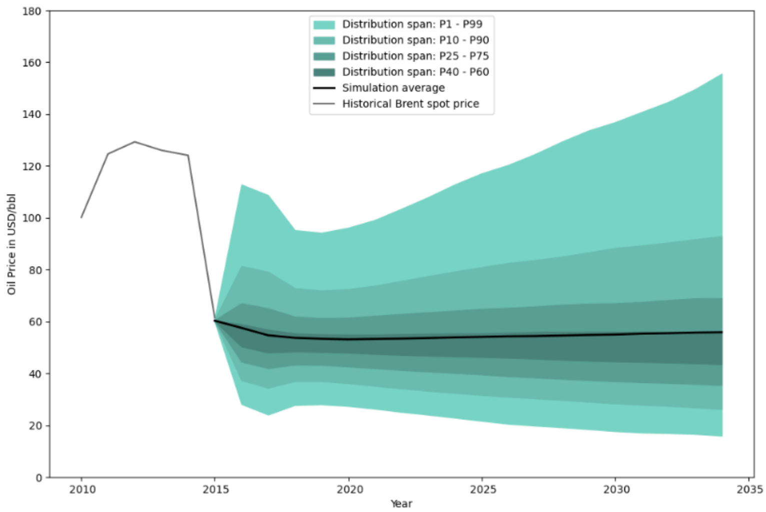

Figure 9: Price percentile plot calibrated on historical data.

4.3 Reservoir uncertainty

As we model the reservoir uncertainty for the first field in M cases of fixed sizes, as of Section 4.1, we will, in this subsection, focus on the modeling of reservoir uncertainty for the second field. This involves assigning a probability distribution to the possible values of the EUR based on the information available. This distribution is updated as more information becomes available, such as from drilling results, production data, and seismic surveys. The intention is to model the reservoir uncertainty in a way that mimics how a company's belief of the field's EUR updates over time. In

[6] The Kalman filter is an estimation technique that estimates state variables in a continuous and changing environment. It is often used in commodity market price modeling Fedorov et al. (2021); Schwartz and Smith (2000); Rankin (1986); Neumann et al. (2006)

this paper, the EUR is modeled using a normal distribution with a mean and standard deviation (SD) that are updated using Bayesian learning as more information becomes available. The mean and standard deviation of the prior distribution are determined based on the current geological and engineering knowledge, while the posterior distribution is an updated distribution that adjusts the prior distribution to the new information we just received. The result is a probability distribution of possible EUR values that reflects the current state of knowledge and uncertainty about the reservoir that develops over time.

The firm receives information through the production of the first field. This allows the field operator to better assess the EUR of the second field. We make the assumption that the field operator revises its EUR once per year. This updated information can lead to the field operator's EUR expectation being lower, higher, or unchanged in comparison to the current belief. To model this process, we consider that reservoir engineers are more likely to obtain increasingly certain information over the course of time. We incorporate incremental reservoir learning, where the uncertainty progressively decreases each year until reaching a minimum level that cannot be further reduced through passive learning (represented as $SD_{min} := \delta$). As a result, we anticipate more significant variations between the updated information during the initial years compared to the subsequent years.

We adopt a *reversed Wiener process* to model the EUR, similar to Lin et al. (2012, 2013). In our simulation, we generate realizations of the oil reservoir, each commencing from the initial estimate at time $t = 0$ and updating until it reaches convergence around the true value, guided by the acquisition of information. Each of these realizations is referred to as a trajectory. To capture a broader range of scenarios for the reservoir's EUR, the model simulates N trajectories for the same reservoir. These N trajectories align with the N simulation paths utilized in oil price simulations and the LSM, enabling their integration into the ROA and CROA frameworks.

The process of updating the belief about the EUR starts at the initial time step t_0, where we have a probability distribution of the EUR $B(t_0)$, representing the initial best guess from the reservoir engineers. We sample a true value of the EUR v^* from a distribution V containing all possible EURs that the reservoir may have. For each year t, we draw a sample from a normal distribution $I(t)$ centered around the true value v^* of the EUR for each trajectory. The standard deviation of $I(t)$ is decreasing over time, reflecting the rate of human learning. We update the belief about the EUR for each trajectory at time t by combining the sample from $I(t)$ with the belief at the previous time $B(t-1)$ using Bayesian learning. The resulting distribution $B(t)$ is the posterior distribution of the EUR at time t given the information up to that point. This procedure is repeated for each subsequent year until the end of the investment horizon. In Figure 10, we visualize this process. Mathematically, this can be written as such:

1. Define an initial probability distribution $B(t_0)$ of EUR based on the reservoir engineer's best guess at time t_0. $B(t_0)$ is normally distributed with mean $\mu_{B(t_0)}$ and SD $\sigma_{B(t_0)}$.

2. Set a true value v^* for each trajectory, drawn from a distribution $V \sim N(\mu_V, \sigma_V)$. For simplicity, we set V to be the same as $B(t_0)$.

3. Construct a distribution $I(t)$ representing the origin of the yearly signal information, with mean $\mu_{I(t)} = v^*$ and standard deviation $\sigma_{I(t)} = v^* \cdot \theta \cdot \epsilon^t$, where θ is the initial standard deviation as a percentage of the mean, and ϵ is the learning rate for each year t that passes. Finally, we sample the yearly information i^* from $I(t)$.

4. Update the distribution from $B(t-1)$ to $B(t)$ at each year t using Bayesian updating with $B(t-1)$ as the prior distribution and i^* as the new information.

5. Produce N trajectories to depict diverse realizations of the reversed Wiener process $B(t)$.

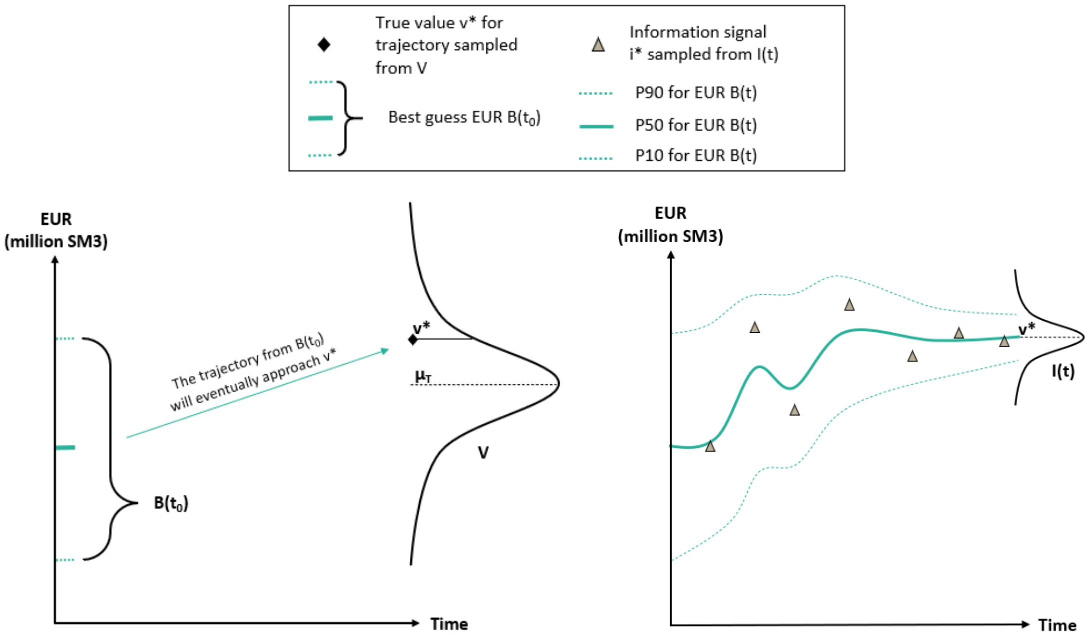

(a) (b)

Figure 10: Visualization of the assembly of a reservoir trajectory. Figure 10a refers to the first and second steps, and Figure 10b refers to the third and fourth steps described above.

Bayesian updating, as performed in Step 4 above, is a method used to update our belief of a certain parameter or hypothesis as new information becomes available. In this case, we are updating our belief about the EUR of the reservoir based on new information obtained each year, using Equation 13 and Equation 14 that is based upon Moldovan (2010). The updating of the posterior mean of $B(t)$ is based on $B(t-1) \sim N(\mu_{B(t-1)}, \sigma_{B(t-1)})$ and our new information i^*. The prior and the signal distribution are both normally distributed, forming a normal-normal conjugate pair, which also implies that the posterior is normally distributed. The equation is as follows:

$$\mu_{B(t)}|B(t-1), I(t) = \frac{\sigma^2_{I(t)}}{\sigma^2_{I(t)} + \sigma^2_{B(t-1)}} \cdot (i^*) + \frac{\sigma^2_{B(t-1)}}{\sigma^2_{I(t)} + \sigma^2_{B(t-1)}} \cdot (\mu_{B(t-1)}) \tag{13}$$

where $\sigma_{B(t-1)}$ is the standard deviation of the prior distribution $B(t-1)$, and $\sigma_{I(t)}$ is the standard deviation of the new-information distribution $I(t)$. The updating of the posterior standard deviation of $B(t)$ is given by:

$$\sigma_{B(t)}|B(t-1), I(t) = \frac{\sigma_{I(t)} + \sigma_{B(t-1)}}{2} \tag{14}$$

where $\sigma_{B(t-1)}$ and $\sigma_{I(t)}$ are the standard deviations of the prior and evidence distributions, respectively. For simplicity, we assume that $\sigma_{B(t)}$ is equal to the average of $\sigma_{I(t)}$, the signal $\sigma_{I(t)}$, which is assumed to decrease over time. Thus, $\sigma_{B(t-1)}$ declines as t grows. We adjust the updating of

$\sigma_{B(t)}$ to align with our problem while considering our intuition regarding the impact of the signal on $B(t)$. With the reduction of uncertainty in the EUR over time, the current best guess $B(t)$ provides a more accurate estimate of the field's true volume as t increases.

The term "reversed" in reversed Wiener process is used to indicate that each trajectory in the ensemble converges to its true value instead of diverging from a starting value, as in a standard Wiener process (Lin et al., 2013). This property of a reversed Wiener process allows us to model the uncertainty in the estimates of the EUR of a reservoir while also ensuring that each trajectory converges towards a single true value. The initial distribution of the reversed Wiener process is chosen to be the same for each trajectory, and the true value for each trajectory is drawn from a distribution, which retains the diffusion property of a standard Wiener process. In Figure 11, we present two realizations of a field's EUR modeled through this process.

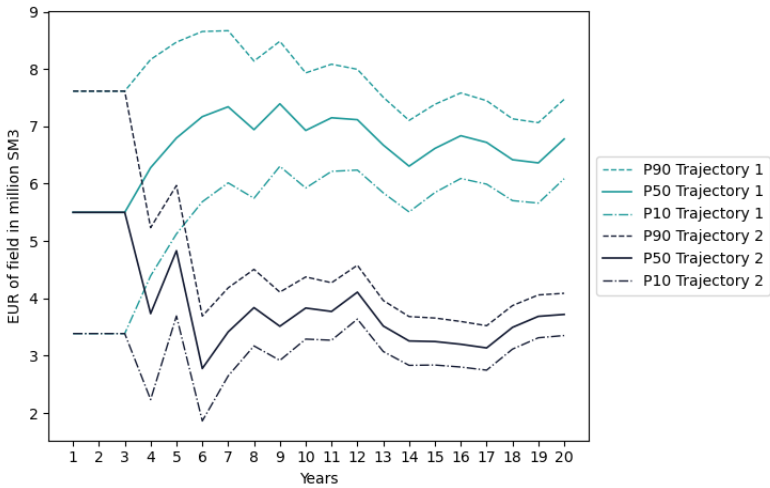

Figure 11: Example including two volume trajectories of a reservoir. Showing how a company's belief about the EUR of the second field develops over time. In this example, we invest in the first field in year 4, resulting in the start of reservoir learning shown in the figure.

According to Lin et al. (2013), our model preserves the random walk characteristic of the Wiener process but with decreasing volatility that follows an exponential pattern over time. The damping factor of this decline in volatility is interpreted as the speed of human learning. The collection of all trajectories provides decision-makers with insight into the evolution of the EUR trajectories over time. We draw samples from these trajectories to optimize production, which in turn becomes an input for the LSM solution approach.

4.4 Production profiles

To determine the optimal timing and value of a potential tieback at each time step, we sample a reservoir volume from the most recent EUR distribution at time t, denoted as s^{oil}. This distribution is taken from the trajectory at time t. Subsequently, using a production optimization model, we convert the new volume estimate, s^{oil}, into production profiles, given the constraints of the host. The aim is to maximize the total volume extracted, subject to the field potential and host constraints. The entire process is depicted in Figure 12 and elaborated further in the subsequent paragraphs.

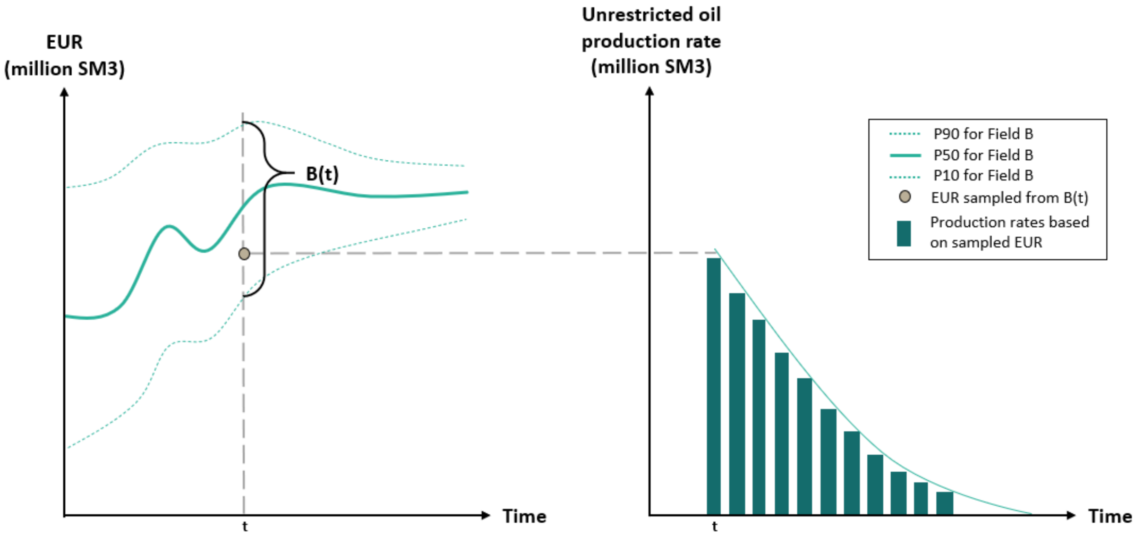

Figure 12: Transitioning from EUR uncertainty to production rates. We model the EUR uncertainty over time by sampling the $B(t)$ distribution. Our proposed optimization model takes in the sampling and produces the production rates used for decision-making at time t.

The production optimization model aims to optimize the quantity of oil produced based on the EUR and capacity constraints from the field and the host over the host's lifetime. The model is presented in Equation 15 through Equation 21 and is accompanied by Table 1, which explains the variables and parameters utilized in the model. The model is based on Hyldmo and Skudal (2022).

Table 1: Variables as parameters for the optimization model. Similar to the variables and parameters as used in Hyldmo and Skudal (2022).

Model input	Type	Description	Unit
q_t^{oil}	Variable	Oil production rate at time t	Sm^3
f_t^{oil}	Variable	Field oil potential at time t	Sm^3
s^{oil}	Parameter	Sample from an EUR distribution	Sm^3
C_t^{oil}	Parameter	SHC of oil at time t	Sm^3
λ	Parameter	Reservoir depletion rate	%
ω	Parameter	Oil production lead time	years
T	Parameter	Lifetime of host	years
T_{CAPEX}	Parameter	Year of investment	year

$$\max_{q_t^{oil}} \quad \sum_{t=0}^{T} q_t^{oil} \cdot e^{-\gamma \cdot t} \tag{15}$$

$$f_t^{oil} = \lambda \cdot (s^{oil} - \sum_{t=0}^{t-1} q_t^{oil}), \quad \forall t \in T \tag{16}$$

$$f_t^{oil} = 0, \quad \forall t \in [0, T_{CAPEX} + \omega] \tag{17}$$

$$q_t^{oil} \leq f_t^{oil}, \quad \forall t \in T \tag{18}$$

$$q_t^{oil} \leq C_t^{oil}, \quad \forall t \in T \tag{19}$$

$$\sum_{t=0}^{T} q_t^{oil} \leq s^{oil}, \quad \forall t \in T \tag{20}$$

$$q_t^{oil}, f_t^{oil} \geq 0, \quad \forall t \in T \tag{21}$$

The primary goal of the production optimization model is achieved through the objective function defined in Equation 15, where the total oil production is maximized while taking into account the time value of oil. To achieve this, the quantity produced is discounted using γ to account for that present oil is worth more than future oil. In this paper, we assume that the production is maximized without considering the economical aspect. However, optimizing the total profit would be an intuitive decision. We suggest to address this in future work, perhaps through the ability to purchase futures contracts enabling profit-based optimization.

The constraints of the production optimization model take into account the field potential, which is modeled by its decline rate (λ), and the host's capacity to receive oil (C_t^{oil}). The first two constraints, presented in Equation 16 and 17, define the nature of the field potential at time t (f_t^{oil}). Constraint Equation 16 specifies that the field potential at time t cannot exceed λ (in %) of the remaining oil volume at that time. Equation 17 ensures that the field does not start production before the year of investment (T_{CAPEX}) or before a lead time after the investment decision is made (ω). Subsequently, Equation 18 limits the oil production (q_t^{oil}) at time t by the field potential (f_t^{oil}). The constraints are simplified by assuming that the depletion rate λ already accounts for well capacity and well placement.

As stated, we assume that the SHC is deterministic and known before considering the tieback developments. In Equation 19, the oil production at time t (q_t^{oil}) is restricted to the host's capacity (C_t^{oil}). Additionally, a limit for the accumulated production volume is added to avoid exceeding the initial oil in place. This constraint is included in Equation 20. Finally, Equation 21 imposes non-negativity constraints on the variables to ensure the model's validity.

As a simplification, we will assume a deterministic, exponentially declining production rate, similar to the approach taken by Thakur (2017). This assumption implies that we obtain perfect information about the true reservoir size and subsequent production profiles in the year we begin exploiting the field. In this paper, we run the model independently for Field A and Field B when determining their production profiles. The resulting production profiles are then used further in the ROA as the quantity of oil produced at each possible investment year, as explained in Section 4.5.

4.5 The ROA implementation

The potential upside of using real options in O&G investment decisions is highly acknowledged (Paddock et al., 1988; Cortazar and Schwartz, 1998; Lund, 2000; Jafarizadeh and Bratvold, 2009; Fedorov et al., 2021). ROA allows for accounting for managerial flexibility and, thus, positioning itself as a preferred choice of analysis faced with substantial uncertainty and irreversible decisions.

Modeling our problem of allocating the project portfolio as cases with two compound options allows us great flexibility in our solution. In this section, we propose a CROA framework combining individual ROA instances to decide the timing and valuation of two marginal fields. We divide the framework into three main components: a CROA case ($\in$ M cases), a CROA, and a Single ROA. These components are outlined from a top-down perspective, where the Single ROA is highly modular throughout the framework. Our CROA case is illustrated in Figure 13. It is referred to as a case due to its assumed fixed reservoir EUR for the first field in Phase 1. Phases 1 and 2 depict the first and second field being *tied in* to the host facility. The result of each case stems from the maximum function used on the result of the two compound options representing each ordering of the fields, where we select the ordering which, on average, yields the highest value. Specifically, this means we select the ordering of the fields that is preferred in the realization of uncertain parameters.

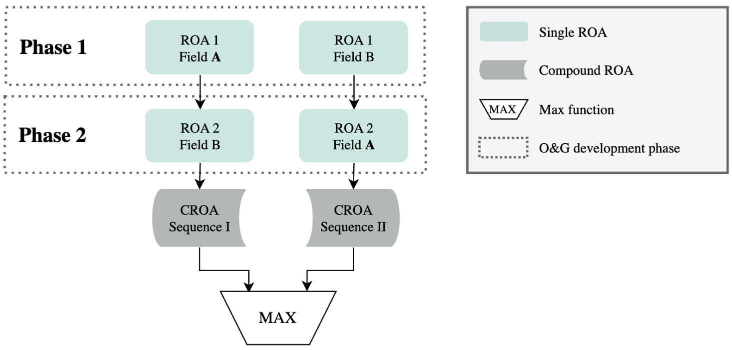

Figure 13: Compound real options analysis case. We zoom in on each case and describe its components in detail.

In our context, for each of the Single ROA instances, the decision of whether to invest or not is evaluated once a year when receiving updated uncertainty parameters. This evaluation is similar to a financial Bermudan option, which can be exercised once every year. How the terms from real options investment projects relate to financial options is summarized in Table 2.

Table 2: Terminology for real and financial options. We can view opportunities to invest in O&G field development as call options. The translations and table format are inspired by McDonald (2013).

Financial option term	**ROA in O&G terms**
Call Option	Investment in the field development project
Strike Price	CAPEX
Underlying Asset Price	Present value of the project
Expiration	The lifetime of the host or field

Longstaff and Schwartz (2001)'s least squares Monte Carlo (LSM) approach is chosen as the solution approach to evaluate our Single ROA instances. LSM is based on generating numerous simulations of the underlying asset to calculate the present value of the option value to invest at each timestep. For each simulation, the option to invest can be exercised at any timestep until maturity, which means that we estimate both the immediate payoff and the value of keeping the option alive, the continuation value, for each year. The determination of exercise boundaries involves a backward analysis, where we examine the values of risk factors to identify the conditions under which the option should be exercised. For each path, the optimal exercise time is determined by identifying the earliest time the path value crosses the exercise threshold, where the investment value exceeds the continuation value of waiting an additional year. The final option value is computed by taking the average of the present value of all the payoffs from the various paths.

We have adjusted the LSM to fit the requirements and circumstances challenging our O&G environment. For all investment options in this paper, the immediate payoff of the project is unavailable. Jafarizadeh and Bratvold (2009) suggest regressing the continuation value and the immediate payoff separately on the information from each timestep. We implement these regressions as well. However, in our implementation with Bayesian learning, we have information that is correlated with the immediate payoff. We refer to the current best guess of the EUR of the reservoir, which is an indication of the true size of the reservoir, and today's oil price, which is an indication of future oil prices. We apply these variables in the regressions.

The following paragraphs highlight the differences between the real options estimates we use in Phase 1 and Phase 2 in the CROA for our real options implementation. Gamba (2003) formulated the adjusted Bellman equation[7] for accounting for an additional option ahead in time, here shown in Equation 22 and Equation 23, where t is the current year of evaluation, υ is the investment year of the first field, ψ is the investment year of the second field, X is the current information of the fields, and T is the last investment years of either field. $F(t, X)$ represents the option value given the time step and the state of the fields.

$$F_{\text{Phase 1}}(t, X_t) = \max_{\upsilon \in \{t, T_{\text{Phase 1}}\}} \{e^{-r(\upsilon - t)} \cdot \mathbb{E}[\Pi_{\text{Phase 1}}(\upsilon, X_\upsilon) + F_{\text{Phase 2}}(\upsilon, X_\upsilon)], 0\} \tag{22}$$

$$F_{\text{Phase 2}}(\upsilon, X_\upsilon) = \max_{\psi \in \{\upsilon, T_{\text{Phase 2}}\}} \{e^{-r(\psi - \upsilon)} \cdot \mathbb{E}[\Pi_{\text{Phase 2}}(\psi, X_\psi)], 0\} \tag{23}$$

To obtain a portfolio view of the value of the first option, we adjust the ROA in *Phase 1* (ROA 1) in Equation 22 to account for the future option values from *Phase 2* (ROA 2) in Equation 23. This means the payoff and regression in *ROA 1* add the final values from *ROA 2*, shown as $F_{Phase\ 2}(\psi, X_\psi)$. This also implies that we need an estimate of *ROA 2* before *ROA 1* when making the investment decision of the first field. As the decision-maker does not hold perfect information about the second field, our approach is to simulate an estimation of the value of the option to add the second field later by binning the true value into k bins.

As the field operator has the ability to trade oil futures based on the current observable oil prices at the time of investment, it follows a correlation between today's oil price and the compound value of adding the second field later. Thus, a decision-maker may base such a guess of the value of the second field on the information about the current belief of the EUR and future oil prices. Given that the decision-maker manages to get a best-effort estimate of the second field contribution within a rough range for *ROA 2*, the CROA can incorporate this value and impact the exercise decision in *ROA 1*.

Given this adjustment of *ROA 1*, we create a compound real option for a given ordering of the fields. Thus, it will resemble what is visualized in Figure 14, where we first compute the values of different simulations of the second field, conditioned on various tieback years of the first field. This

[7] The Bellman equation is expressed by the value of a state which equals the maximum of the immediate reward and the value of future states (Dixit and Pindyck, 1994). In our context, the LSM is using the least-squares regression estimate of the continuation value versus the immediate payoff to decide on the investment timing providing the best total expected profit (Gamba, 2003).

comes together to form our CROA evaluation of the total portfolio. In the Single ROA benchmark, we use two equations identical to Equation 23 for both phases.

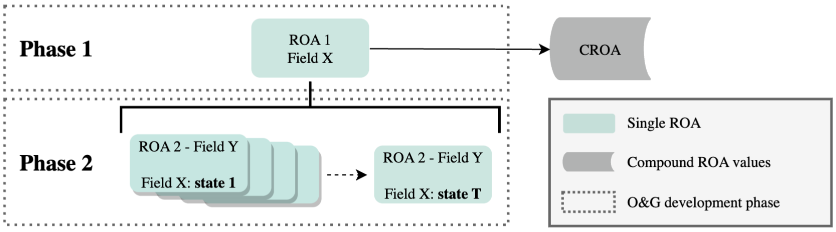

Figure 14: Compound real option analysis. We closely examine how each compound option is modeled.

Recalling the model overview in Figure 7, Section 4.1, we assemble both *ROA 1* and *ROA 2* and further into the compound real options cases, finally making up our total model.

4.6 Benchmarking the solution approach

We have included two benchmarks as reference points to showcase how the CROA works regarding total valuation, timing, and the value of flexibility under uncertainty. First, we include a benchmark valuation with a myopic view of marginal field development with tiebacks. The benchmark utilizes the NPV approach, as presented in Equation 2, and uses the expected size of the reservoirs and the mean of the future price paths, as described in Section 4.2 and Section 4.3, to calculate the expected revenues and costs for investment in each year. To catch the portfolio view of this, we have tested the combination of all investment years for Field A and Field B and selected the combination yielding the highest value as the benchmark.

Second, we include the Single ROA as a benchmark. The Single ROA approach is based on the same assumptions and methodology as the CROA. However, we first optimize Field A using ROA without considering that we can add Field B later. Then we utilize a second ROA to optimize the timing and value of Field B before adding these values together as the portfolio's value. Hence, the reader should be aware that the Single ROA is also an approach proposed by this paper, which works perfectly if one is to only optimize the tieback of a single field. Moreover, one can view the Single ROA approach as natural for what would happen if the two fields were operated by two distinct field operators, where each wants to optimize their tieback developments, regardless of their competitor.

5 Case study

We analyze a case study with one host and two marginal fields, similar to the one found in Hyldmo and Skudal (2022), using the methodology described in Section 4. Although the data is synthetic, the case resembles real development projects on the NCS and is designed together with O&G experts at NTNU. The case highlights the critical aspects of a typical O&G project, and the paper aims to demonstrate the relevance and validity of the methodology proposed.

Case description and general economic parameters

We assume that both the host and the fields are feasible for a tieback and that all licenses and permissions are in place. In our case, we have set the lead time ω to 1 year. We learn about the size of the second field from the same year as the investment in the first field. The decommissioning of the host will occur in year 20, meaning that year 20 is the last year available for production. Moreover, we have set the last possible investment year in either field to be year 18.

We use a discount rate of 7%, following Fedorov et al. (2022a) and the recommendation from The Norwegian Petroleum Directorate (2020) for the risk-adjusted discount rate. The literature suggests using a risk-free rate in ROA under the assumptions of efficient markets. As the markets are non-efficient and the literature does not offer a way to adjust appropriately for this (Jafarizadeh and Bratvold, 2015), we choose to use the discount rate noted above, although it may be argued to be slightly too high. The inflation rate in Norway and the United States is set to 2%, assuming the Norwegian Central Bank and The Federal Reserve meet their long-term goals (Norges Bank, 2022; Federal Reserve, 2020). The inflation rate is represented by α.

To calculate the myopic benchmark, we use the expected EUR of our current best information $B(t)$ and apply it in the optimization model to generate production rates. The NPV is then calculated using these production rates and the mean of the simulated oil price paths. We iterate through all possible exercise years and select the highest NPV as our benchmark. For both the myopic approach and the real options approaches, we use N = 10 000 simulated paths for the oil price and the second field's EUR[8].

Recall that we place the expected second field value contribution into bins when evaluating the investment decision of the first field. The estimating of the second field's value is binned into $k = 3$ bins. In our case, these are yearly selected to be the P20, P50, and P80 of the estimated value distribution of the second field. We do this to obtain a rough but still broad representation of the full state space of possible second-field valuations. This binning is necessary to make the model applicable to a field operator to utilize in a real use case. Our main message is not to use these percentiles exactly but rather to showcase the concept of a rough estimation of the compound option value and how it affects the CROA. These aforementioned general parameters for the case study are summarized in Table 3.

Table 3: General parameter values

Parameter	Value
Lead time (ω)	1 year
Last possible year of exercising (ϕ)	Year 18
Lifetime Host (L)	Year 20
Discount rate (γ)	7%
Inflation rate (α)	2%
Number of simulations (N)	10 000
Number of cases (M)	1
Bins of the value of the second field (k)	3 - {P20, P50, P80}

[8]We choose 10 000 simulations based on robustness tests performed with the algorithm. However, for making the sensitivity analysis in Section 6.3 computationally feasible, we have used 1 000 simulations. A more detailed summary can be found in Appendix B.

Oil price simulation

Applying the STLT two-factor price model of Schwartz and Smith (2000), we use a Kalman filter to calibrate the parameters from historical oil prices. As this paper aims to find interesting patterns in our ROA and CROA approaches, we use the same parameters previously estimated within our team and earlier presented in Fedorov et al. (2022b). These parameters were estimated with Brent crude oil spot prices from March 2006 to June 2021 and are listed in Table 4.

Table 4: Parameter values for oil price simulation.

Two-factor parameter	Value
ξ_0	4.07
χ_0	0.10
σ_ξ	0.12
σ_χ	0.56
λ_χ	0.11
μ_ξ^*	-0.0045
κ	0.45
$\rho_{\xi\chi}$	0.12

Production profiles and host constraints

We assume a learning rate of 20% for the production profiles, giving us $\epsilon = 0.8$. The minimum uncertainty for each trajectory is set to 25 % of the initial uncertainty ($B(t_0)$), representing a field operator never can become certain of a field's EUR. The depletion rate λ, used to generate production profiles, is set to 20 %. That assumes a maximum extraction rate of 20% of the remaining EUR for each field per year. All parameters related to production profiles, sizes, and uncertainty of the fields are summarized in Table 5.

Table 5: Parameter values for generating production profiles.

Parameter	Value	
Depletion rate (λ)	20	%
Minimum uncertainty (δ)	25	%
Information volatility discount factor (ϵ)	0.80	-
First field when Field A - EUR	6.0	million Sm^3
First field when Field B - EUR	5.5	million Sm^3
Second field when Field A - EUR	$N(6.0, 1.80)$	million Sm^3
Second field when Field B - EUR	$N(5.5, 1.65)$	million Sm^3

As the original production of the host decreases over time, it leaves capacity that can be utilized by Field A and Field B. We denote the SHC for the base case as *increasing*, reflected by its shape shown in Figure 15. We assume the SHC to consist of oil volumes, having already taken into consideration the water and fluid constraints of the host. The *increasing* case is modeled to reflect a maturing host with available oil capacity and not so much of a water constraint.

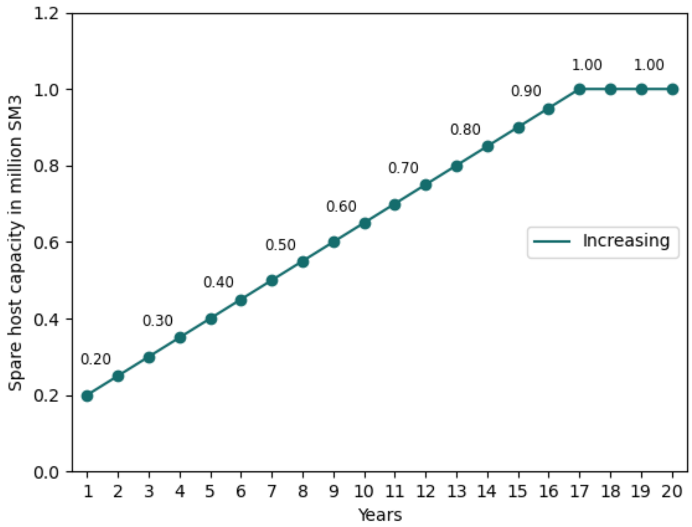

Figure 15: Visualization of the increasing spare host capacity.

Costs and abandonment

Table 6 lists fixed tariff, OPEX, CAPEX, and ABEX parameters. The *beta* parameters describing the tariff and OPEX, earlier mentioned in Equation 4 and Equation 5, are selected to make as realistic costs as possible.

Table 6: Parameter values for costs.

Parameter	**Field**	β_0	β_1	β_2	**Value**	
CAPEX	A	-	-	-	600	MMUSD
CAPEX	B	-	-	-	550	MMUSD
ABEX	A&B	-	-	-	100	MMUSD
Tariff	A&B	20.0	3.0	0.03	-	MMUSD/year
OPEX	A&B	10.0	5.0	-	-	MMUSD/year

6 Results

Key to stakeholders of this decision problem is whether and when to invest in either field. Applying the methodology to the case presented in Section 5, we first derive the total value and optimal timing for a portfolio of two marginal fields and one host in Section 6.1. Secondly, in Section 6.2, we examine decision rules and the practical application of our proposed CROA approach. Thirdly, we analyze the decision of which field to tie back first and explore how important model parameters affect this choice in Section 6.3. Further, we dedicate a section comparing the CROA to a Single ROA approach and evaluate its robustness for different SHC variations in Section 6.4. Finally, in Section 6.5, we further compare the CROA to a CROA which has perfect information about the value of the second field when making its decision to invest in the first field.

6.1 Base case results

The results for the base case with an *increasing* SHC are presented in Figure 16. The figure displays the total value estimate for the portfolio for the CROA and the myopic benchmark, including the average percentage value contribution of both fields to the total value. The valuation of the fields stems from choosing to tie back Field A before Field B for those simulation paths where investment was considered optimal, as most simulations results in a higher valuation with this ordering of fields. How the ordering results from the CROA approach can be interpreted is discussed in Section 6.3.

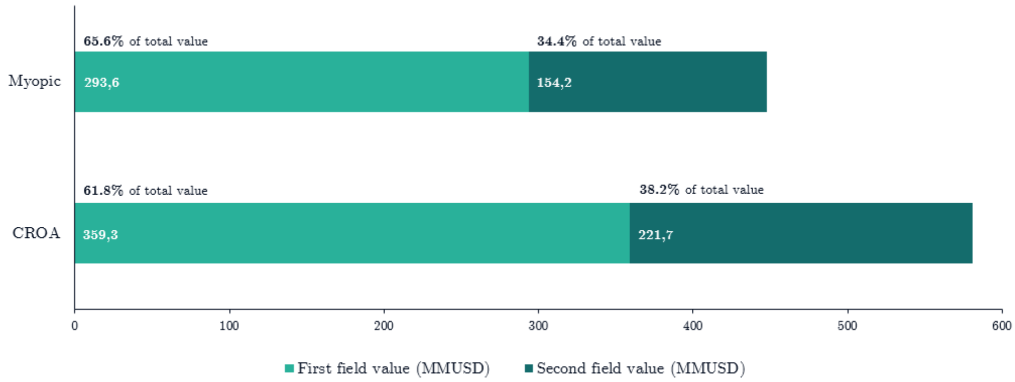

Figure 16: Valuation of the portfolio comparing the CROA to the myopic approach. This reflects the average expected valuation from choosing Field A first and then having the option of investing in Field B later.

The myopic approach yields a portfolio value of 447.8 MMUSD, with an average composition of 65.6% profits from Field A and 34.4% from Field B. In contrast, utilizing the CROA approach for investment in the two fields results in an expected value of 581.2 MMUSD, representing a 29.8% increase compared to the myopic approach. On average, this valuation comprises 61.8% profits from Field A and 38.2% from Field B.

An overview of the solution approaches' most common investment timing statistics can be found in Table 7. The statistics need proper context to be a valuable comparison due to the nature of both approaches. For the myopic approach, the displayed years are the timing yielding the highest expected portfolio value. This suggests decision-makers use the most common investment timing for all future realizations. In contrast, the CROA results suggest investments are more probable to occur in some years compared to others. This does not suggest any specific decisions for a decision-maker today but rather gives a picture of likely outcomes in the future. The investment decisions, according to the CROA framework, should be made by observing risk factors and applying decision rules resulting from the analysis. That means a decision-maker today believes it is optimal for the myopic approach to invest in Field A in year 6 and in Field B in year 13, regardless of future observable risk factors. Most likely, a decision-maker will tend to invest in both fields earlier by applying the CROA approach but will observe risk factors and follow decision rules to make the final decision.

In Figure 17, we present a visualization of the different realizations of the DCF calculations in the myopic approach. Investing in Field A in year 6 yields the highest valuation, but the expected value in years 4, 5, and 7 is close to this optimum. The dark line shows the expected Field B contribution for each investment decision for the first field. The second field timing is conditioned on the decision for the first field, leading to investment in Field B in year 13.

Table 7: Most common tieback timing results of the myopic and CROA approach. In the majority of the simulation cases, it is optimal to start with tying back Field A.

Solution approach	Most common investment year for the first field	Most common investment year for the second field
Myopic	6	13
CROA	1	12

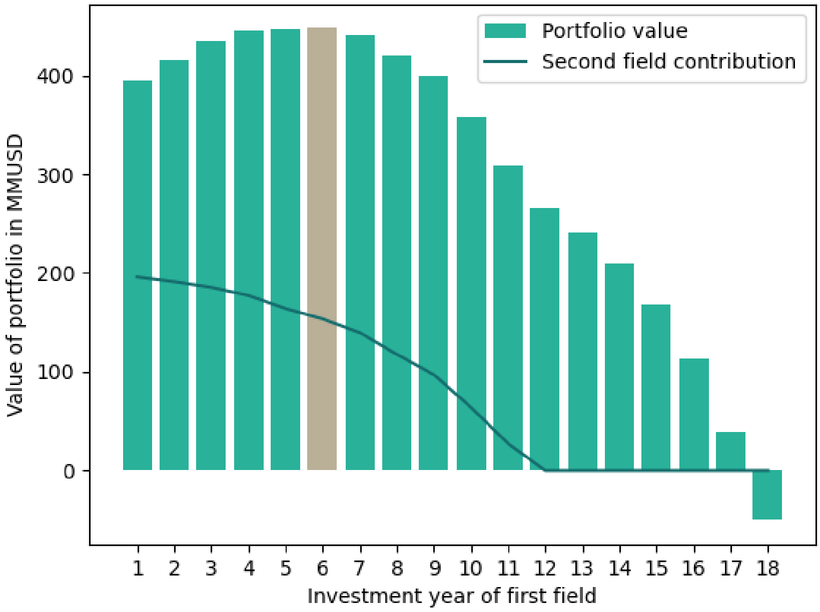

Figure 17: The optimal timing to invest in the first field as of the myopic approach. It shows all cash flow combinations of each investment year. The results suggest investing in year 6 in Field A and in year 13 for Field B as a direct consequence.

A visualization of the investment timing of the CROA is illustrated in Figure 18. The timing results for Field A as the field first field and Field B as the second field is shown in Figure 18a and Figure 18b, respectively. One notable difference from the myopic approach is that in certain simulation paths of the CROA, none of the fields are selected for investment. In the case of the CROA, it is optimal to invest in Field A in the majority of simulation cases as opposed to not doing so. Specifically, it is optimal to invest in year 1 for 24.4% of the simulation cases, in year 2 for 15.9% of the simulation cases, and so forth, while never investing in 13.2%. This is shown in Figure 18a. For the CROA, Figure 18b shows the investment timing for Field B, conditioned on the investment in Field A in year 1. Recall that we calculate all conditional results, and therefore we choose to present the most common condition. Interestingly, we can note that the most probable outcome for the CROA is that Field B is exercised in years 10 to year 13, but with year 12 as the most common investment year. Moreover, ∼30% of the simulations will never be invested in.

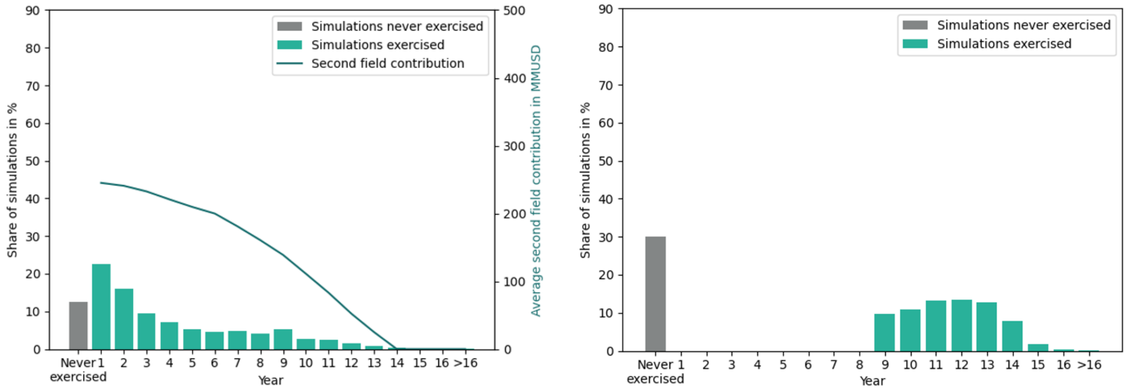

(a) Share of simulations for investing in the first field, showing the average contribution of the second field.

(b) Share of simulations for investing in the second field, given that we invest in the first field in year 1.

Figure 18: The optimal timing to invest in the first and second fields as of the CROA solution approach.. Given all simulation cases, the results show the share of simulations exercised at each time step or not at all. The results do not suggest an investment timing today but rather probable future outcomes. The CROA generates decision rules based on risk factors, and these are outlined in Section 6.2.

When comparing the simulation outcomes of the CROA and the myopic approach from applying the respective timing results, a decision-maker will observe significant differences in the distribution of values. By applying the price paths from Section 4.2 and the true values v^* for field EUR from Section 4.3, we can simulate the outcomes of making these respective decisions. Figure 19 shows how these outcomes compare in a histogram of possible values. The timing suggestions from the myopic approach often turn out as non-profitable outcomes represented by the left tail of the distribution. Meanwhile, the CROA approach limits the downside by not investing when the simulations reflect a negative expected cash flow. The yearly reevaluation of the CROA allows for a good hedge against the downside when the investment decision is deferred. Moreover, the CROA has a rightmost skew representing a higher probability of capturing upside potential from the investment. This value comes from adjusting the investment year based on finding a state profitable enough to justify the downside risk of the investment.

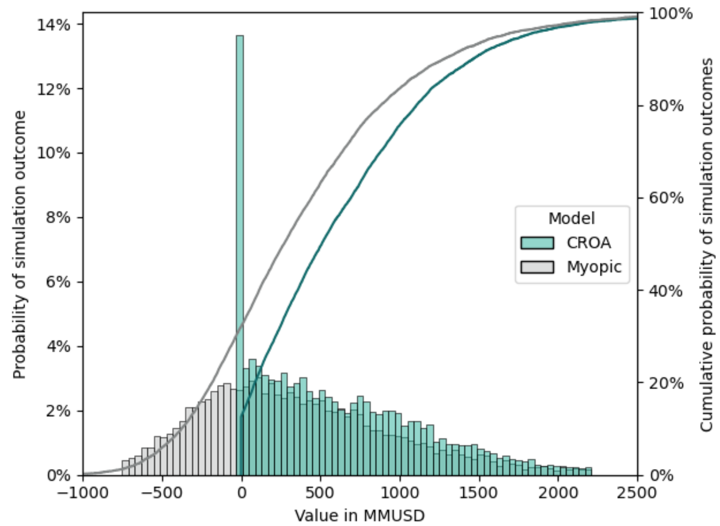

Figure 19: Value distributions of the myopic and the CROA approach. The histplot shows how the values originating from the suggested decisions of the myopic and the CROA approach result for all simulations. The respective darker lines illustrate the cumulative distributions.

The CROA approach identifies additional value in timing flexibility amid market and reservoir uncertainty, mitigating monetary risks while capturing upside potential. This value is evident in three key aspects: 1) Some paths are never exercised, avoiding unprofitable reservoirs. 2) Paths are exercised earlier than the myopic approach suggests, capitalizing on high expected revenues due to a favorable size of the risk parameters. 3) Paths can be exercised after the myopic approach's indicated year, benefiting from receiving information and deferred decision-making. Thus, the CROA enables field operators to optimize investment strategies and increase the portfolio value of marginal fields.

Results using three EUR cases

Thus far, we have considered the results for the base case. We now analyze how the expected value changes when increasing or decreasing the EUR of the first fields by 50%. We choose to adjust the second field's EUR distribution mean correspondingly. This is equivalent to increasing the number of volume cases to $M = 3$. Our methodology can be straightforwardly applied to any other choice of EUR volume cases. We choose $M = 3$ to showcase the concept and still prioritize examining other aspects of the approach, like decision thresholds and field selection drivers. The reasoning for choosing 50% up or down is to present two extreme cases compared to the nuanced parameters characterizing the base case. We introduce a *high* case and *low* case in terms of size for both fields. The *high* case represents an increase of the fields EUR by 50% of the size of the base case, while the *low* case is a decrease of 50%. The new case parameters are displayed in Table 8.

Table 8: Parameter values for generating production profiles in the low and high case.

Parameter	**Low**	**High**	
Field A when first - EUR	3.0	9.0	million Sm^3
Field B when first - EUR	2.25	8.25	million Sm^3
Field A when second - EUR	$N(3.0, 0.90)$	$N(9.0, 2.70)$	million Sm^3
Field B when second - EUR	$N(2.25, 0.68)$	$N(8.25, 2.48)$	million Sm^3
Case probability weights	20	20	%

These changes in expected EUR for both fields have significant implications for both valuation and optimal timing results, which are presented Table 9 and Table 10. Interestingly, the most common timing decision clearly changes in the base case but not in the two others. This supports our initial presumptions about the relevance of the use of real options in edge cases evaluating marginal fields. Using the probability weights from our case study section (*low* - 20%, *base* - 60%, and *high* - 20%), Section 5, we can estimate a more realistic valuation of the portfolio using domain knowledge about future scenarios of outcomes. Our weighted portfolio results in 556.1 MMUSD, a 25.2% increase compared to the myopic approach.

Table 9: Expected NPV of the portfolio in MMUSD, resulting from the myopic and CROA approach for the three different volume cases, respectively.

Solution approach	**Low**	**Base**	**High**	**Weighted average**
Myopic	0	447.8	877.6	444.2
CROA	56.1	581.2	980.7	556.1 (+25.2%)

For the *high* case, the estimated value applying the CROA is equal to 980.7 MMUSD, while an expected NPV of 877.6 MMUSD results from the myopic approach. In both cases, it is optimal to

Table 10: Most common tieback timing results for the three cases and two solutions approaches. We present the timing results for tying back Field A first, as this is optimal in the majority of the simulations.

EUR case	Solution approach	Most common timing first field	Most common timing second field
Low	Myopic	Never	Never
	CROA model	Never	Never
Base	Myopic	6	13
	CROA model	1	12
High	Myopic	1	13
	CROA	1	13

immediately invest in year 1 and in the second field in year 13 for the majority of simulation cases. The EURs for the *low* case result in the optimal decision to never invest in any of the fields for both solution approaches. The CROA will, for most of the simulations, never exercise the option to invest in the first field. Consequently, we are most likely to never exercise the second field, either. However, in case the oil price rises sufficiently high in later years, it may invest in the first field later, most commonly made around years 10 through 14. The resulting total project value is equal to 56.1 MMUSD, showing the value of incorporating the yearly reevaluation in our solution approach, compared to the static decision in the myopic approach, which is zero. More detailed timing results for the *low* and *high* cases regarding the aforementioned discussion are shown in Figure 20.

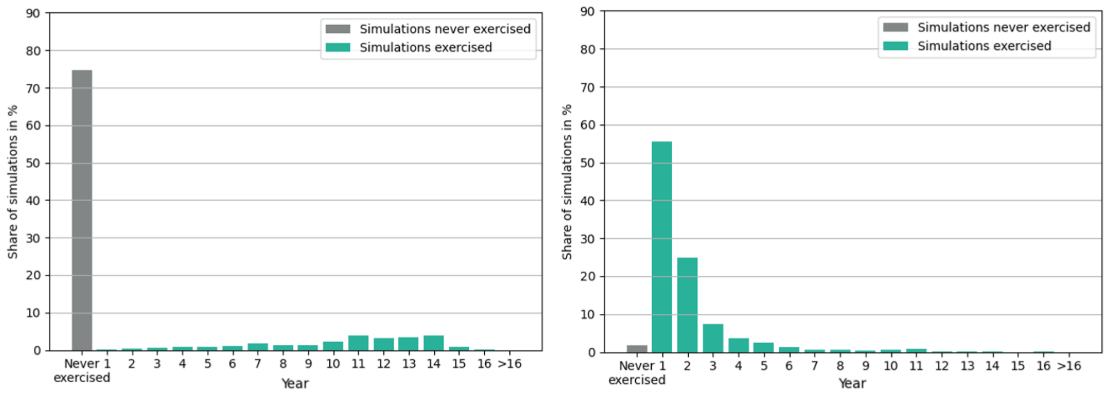

(a) Low case timing **(b)** High case timing

Figure 20: Exercise timing of tying back the first field in the low and high for the CROA. In both cases, Field A is developed first and Field B second.

Our findings demonstrate the substantial value gained by incorporating flexibility in investment timing. Specifically, our analysis reveals that the project value derived from the CROA consistently exceeds that obtained through a myopic investment approach across all EUR cases examined. We emphasize the significant disparity in investment timing between CROA and the myopic approach, with the former suggesting earlier investment for the first field. This distinction bears profound implications for the decision-making process undertaken by field operators.

6.2 Investment thresholds and decision rules

We now analyze the investment thresholds and decision rules for the CROA approach in detail. We answer for which values of risk factors have the LSM regressions estimated that it is optimal to invest as opposed to deferring. First, we present the optimal exercise thresholds for all years using the expected future revenues. The optimal timing of investment for a field can vary substantially for different simulation paths, depending on the combination of the current oil price, the current best guess of field EUR, and the binned expected value of the second field. In Figure 21, we illustrate this exercise threshold as a function of the expected revenue calculated at every time step. As we are unable to see the true profit of a path at the decision point, we use the expected revenue in our regressions in the LSM as an indicator of the true profit for a path, as the expected revenue is observable and the true profit of a path is not. We have explored different versions of regressions that can be a good indicator of the true profit of a path, and the expected revenue had the best and most stable performance regarding this estimation of the true profit for a path. The expected revenue is calculated as the sum of produced oil, based on the current best guess of the EUR, multiplied by today's observable oil price and added to the binned expected value of tying back the second field later.

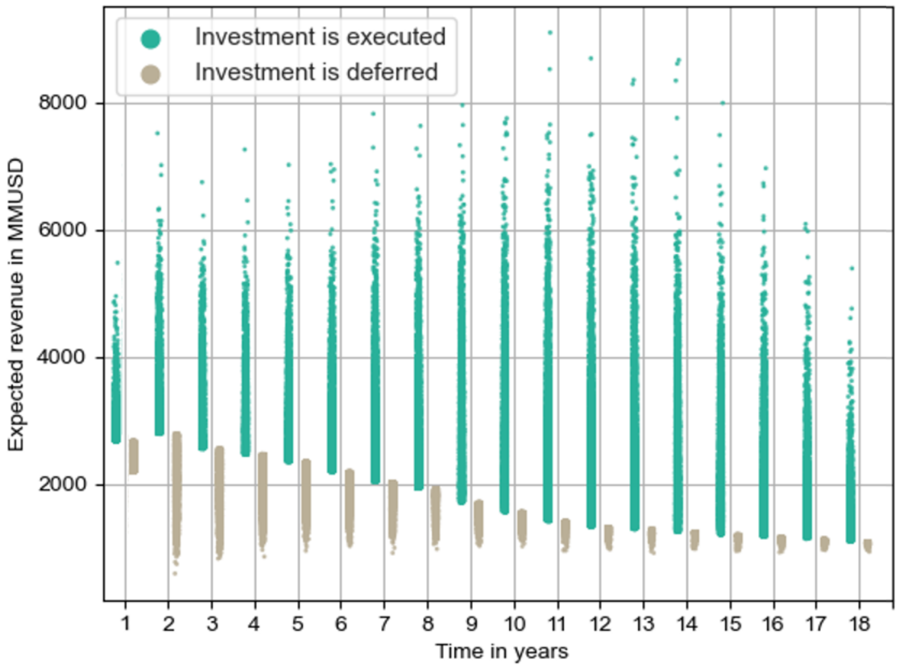

Figure 21: Investment thresholds for each year for the first field. The expected revenue is relative to each year it is considered, e.g., revenue in year 8 is discounted back to year 8.

The data points are divided into two groups: an upper investment region, where the optimal decision at this state is to invest, and a lower continuation region indicating where it is optimal to defer and reconsider investment next year. The boundary between these two regions indicates the investment threshold. In our base case, the investment threshold is downward-sloping. This is due to the upward-sloping SHC, indicating that there is value in waiting to invest when the host has more available capacity in later years.

We can examine in more detail how regression in the CROA approach derives the investment decision using the exogenous risk factors that together compose the expected revenues. The risk factor for tying back the first field will be the oil price and the value of the second field, as the EUR of the first field is static for each of the M scenarios in the CROA approach. Recall the binning of the expected value of adding the second field later into k bins, as of Section 4.5. Given that we can roughly estimate the value of adding the second field later, the model can present clear decision rules for decision-makers regarding the necessary oil price needed to invest for each bin.

In Figure 22, we present the proposed decision rules following the investment threshold for tying back the first field, given the available risk factors present at the time of evaluation. For a decision-maker to make use of these decision rules, we propose they should consider the figure for the current year and see if the combination of today's oil price and the expected value of tying back the second field suggests exercising the option to invest or to defer and reevaluate again next year. Given our three expected binned values of the second field, a decision rule in the form of a threshold oil price is found on the y-axis. For example, in year 2 in Figure 22b, the oil price is approximately 75 USD/bbl for the P20, 70 USD/bbl for the median, and 63 USD/bbl for the P80. Similar to Figure 21, we see the pattern of a decreasing investment threshold for the increment of years.

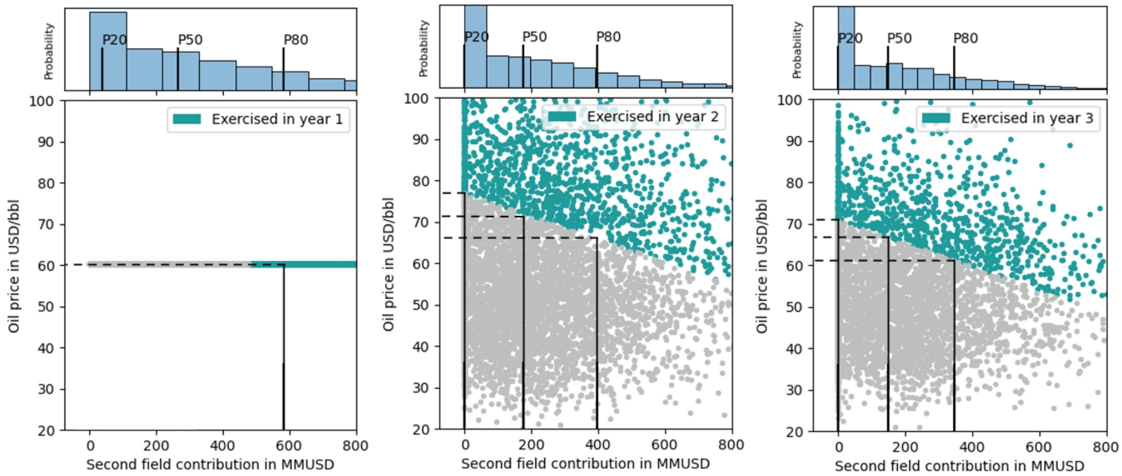

(a) Investment threshold for Field A in year 1

(b) Investment threshold for Field A in year 2

(c) Investment threshold for Field A in year 3

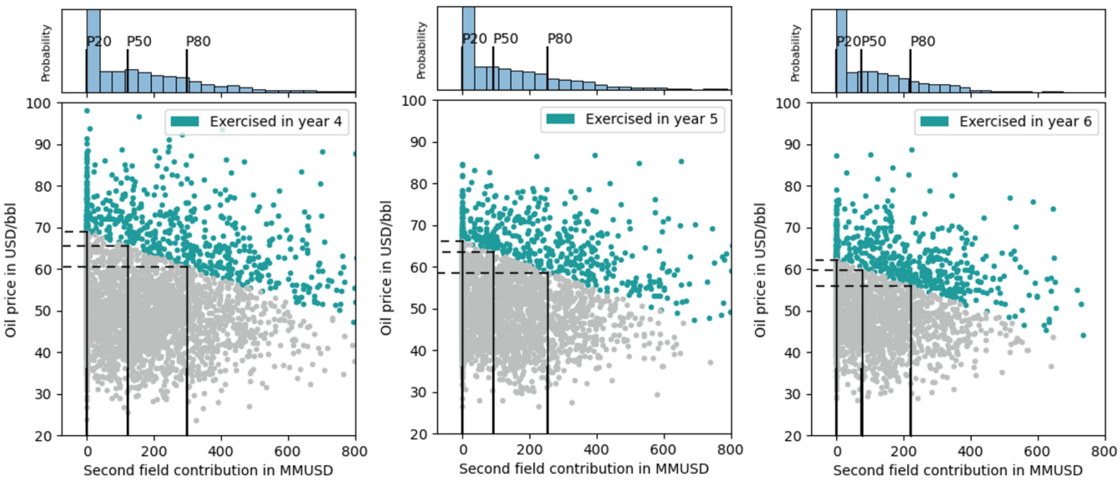

(d) Investment threshold for Field A in year 4

(e) Investment threshold for Field A in year 5

(f) Investment threshold for Field A in year 6

Figure 22: The collection of investment thresholds in years 1-6 specifying decision rules for the first field. The evaluation uses the oil price and the estimated second field value contribution. The percentiles are recalculated for each year as the second field contribution distribution is adjusted when conditioned on the first field.

The overall trend of the investment threshold for the first field is summarized in Figure 23. We generally observe a trade-off between the required oil price and the required expected value of tying back the second field. This means that a field operator can lower the required current oil price if the expected value contribution of the second field is significantly large.

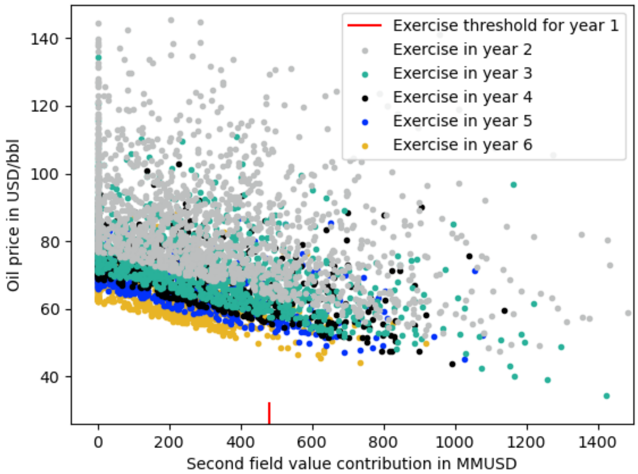

Figure 23: Investment threshold summary plot for the first field in years 1-6. It displays the annual exercise thresholds for the first field using the oil price and estimated second field value contribution.

When considering tying back the second field, the evaluation is conditioned on the decision to tie back the first field and the updated SHC. In this decision situation, the only risk factors to consider are the oil price at this time and the best guess of the second field's EUR. These both can be observed by a decision-maker in reality. Thus, we can present investment thresholds for the second field in Figure 24, having the field's best guess EUR on the x-axis and the oil price on the y-axis. For a field operator, the decision rule comes from selecting the current year and best-guess EUR leading to the threshold oil price on the y-axis. For example, in year 12, the investment threshold for the oil price is approximately 55 USD/bbl, given a best-guess EUR of 4 million Sm^3.

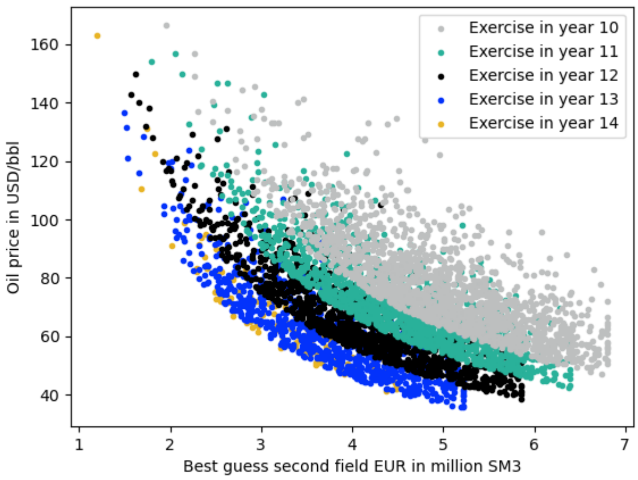

Figure 24: Investment threshold plot for the second field in years 10-14. It displays the annual exercise thresholds for the second field using the oil price and the second field's best guess reservoir volume.

In Figure 24, we also observe the trade-off between risk factors, but now considering a single option to invest, we can use today's oil price and the current best guess of the recoverable volume of oil. This investment threshold is non-linear and convex, showing the trade-off between the required current oil price and belief of the size of the field to invest in each year. Moreover, the general threshold decreases with time. As the visualizations for both the first and second fields clearly show a decision boundary, we encourage decision-makers to apply CROA decision rules in order to invest appropriately based on the risk factors present at the time.

Decision rules from ROA and CROA equip decision-makers with a comprehensive and flexible framework to evaluate investment opportunities by explicitly considering the option value. This approach enables informed decisions aligned with effective risk management and maximized value creation. By acknowledging embedded real options, a field operator can assess investments within a dynamic environment and adapt to changing circumstances.

6.3 Decision statistics and drivers of field selection

Essential to our problem is whether we want to tie back Field A or Field B first. The CROA model provides insight into this by outputting statistics of how our simulations pathwise compare when choosing Field A or Field B first. This is done by pathwise selecting the one with the highest present value. Table 11 show these statistics for the CROA in the base case. For a decision-maker who wants to order fields that result in the highest average value, this decision almost always corresponds to the statistics of what share of simulations pathwise is higher than the other.

Table 11: CROA simulation statistics. We include the *weighted* column as a tool to decide the field ordering if a decision-maker wants to account for future outcomes

	Low case	Base case	High case	Weighted	
A first	14.4	55.7	73.7	51.0	%
B first	10.2	31.4	24.6	25.8	%
Both zero	75.4	12.9	1.8	23.2	%

Vital for a decision-maker is also the risk of their investments and the relative risk from comparable investment alternatives. I.e., what risk am I, as a decision-maker, facing if I make the alternative choice rather than what is suggested as the most profitable by the model, given that I will invest? Choosing whether to go with Field A or Field B first could result in substantially different outcomes. A decision-maker does not know at the time of investment if the true outcome of choosing A first or B first is higher or lower than what is recommended as the highest average value. In the base case, choosing to tie back Field A first gives the highest average portfolio value. We simulate that we move forward with this decision and see how it pathwise compares to the alternative ordering of choosing to tie back Field B first.

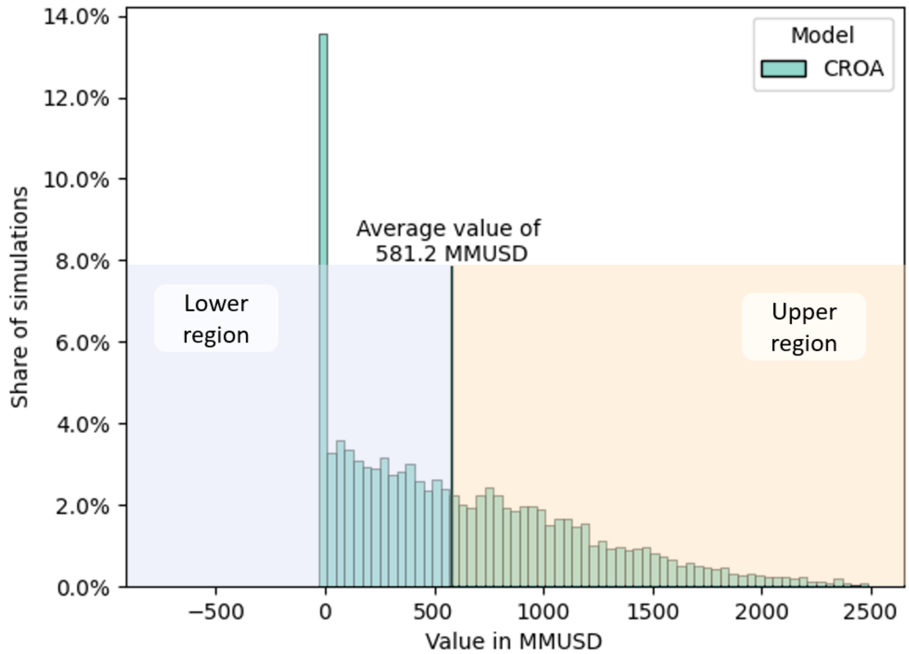

Figure 25: Value distribution from the CROA, divided into a lower and upper region. It shows a histogram of pathwise maximum values from the simulations. The average of 581.2 MMUSD is the artificial average value if we were able to pathwise always select the highest valued ordering of the fields for all simulations.

Figure 25 displays the resulting RoI for the CROA from each simulation, similar to Figure 19, but highlighting a lower region below the average payoff of 581.2 MMUSD and an upper region above this average. First, we will analyze the general trend in Figure 26. Later, we find different characteristics for choosing between Field A or Field B first in these two regions, which we will discuss later, based on Figure 27.

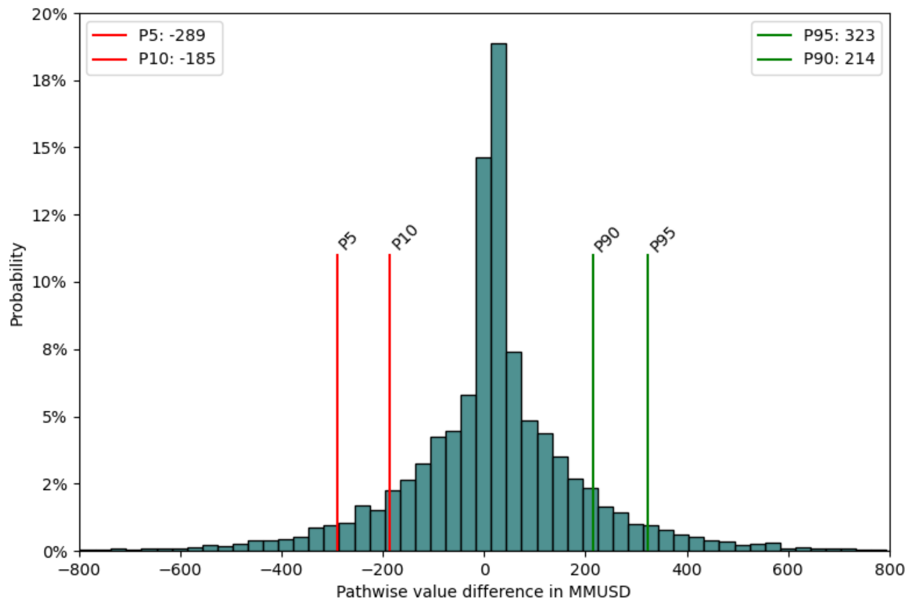

Figure 26: Pathwise difference distribution of Field A-first compared to Field B-first values. The histogram shows how much better it is to choose Field A first compared to choosing Field B first.

By visually representing the subtracted simulation outcomes of tying back Field A before Field B, Figure 26 shows the quantitative benefit of this decision. We can see a difference between the right and left sides of the distribution around zero, but it is not that substantial. The right side of the distribution contains slightly higher values than the left. That means both orderings provide both an upside potential and a downside risk with either order, but Field A first yields a higher value on average. This can be seen by carefully inspecting the figure or by comparing the absolute values of the percentiles noted on the figure. The reader should note that the equivalent figure of comparing the selection of tying back Field B before Field A would be the same as Figure 26, but mirrored.

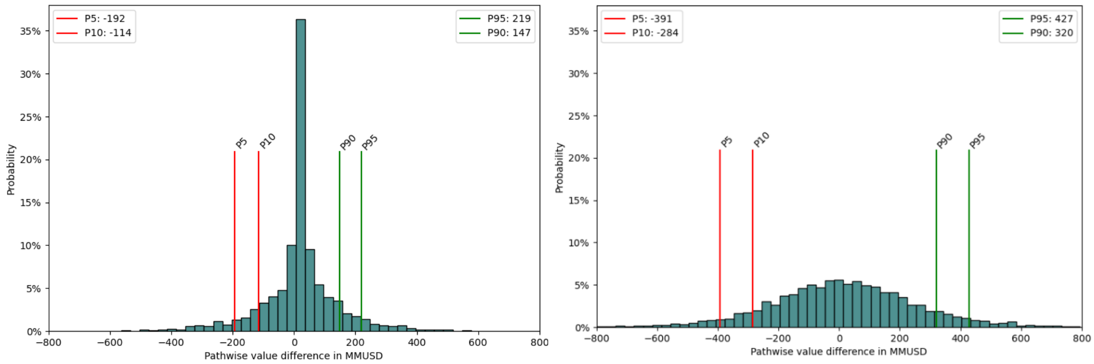

(a) Lower region

(b) Upper region

Figure 27: Pathwise difference distribution divided into the lower and upper regions. The histograms show how much better it was to choose Field A first compared to choosing Field B first in both regions.

Figure 27 presents a histogram illustrating subsets categorized into the *upper* and *lower* regions. The decision is based on whether the simulation path results in a portfolio value below or above the average. We see the same patterns as in Figure 26 regarding the difference between the left and right sides of the distribution around zero, where the right tail is slightly larger than the left tail. This pattern is also validated by the percentiles shown in each figure. However, comparing the *lower* region to the *upper* region, we see substantial differences in terms of risk. In the *lower* region, the differences between choosing Field A or Field B first are not that large, indicated by the spike in values around zero. In the *upper* region, the consequence of making the wrong choice is larger in absolute terms. The nature of larger portfolio values in the *upper* region explains why the deviation is higher. However, we encourage decision-makers to exercise caution when selecting the ordering of fields, particularly in scenarios where the total portfolio value is substantial and the disparity in downside risk associated with the choice of field sequence is comparatively greater.

As the modeling that leads to the statistics of tying back Field A or Field B first is rather complex, it is challenging to isolate specific circumstances that influence the optimal field order. Thus, we choose the order sequence with the highest average value. With this in mind, we want to run sensitivities on parameters from the case study in Section 5 to find patterns in what parameters are more important for the choice of tying back either field first.

Drivers of field selection

Moreover, we aim to answer what parameters are most influential when deciding to tie back Field A or Field B first. This is intended to answer the problem an oil producer faces when needing to know what field it wants to start tying back among a portfolio of alternatives. Perhaps the most natural parameter is the field EUR, and hence to tie back the field with the largest EUR. We have tested this in the base case in Section 6, where Field A and Field B have different sizes. This favors Field A in 55.7% and Field B in 31.4% of the simulations. While the EUR of the field may be the most influential factor, the following paragraphs will see how other parameters influence the choice of what field to exercise first.

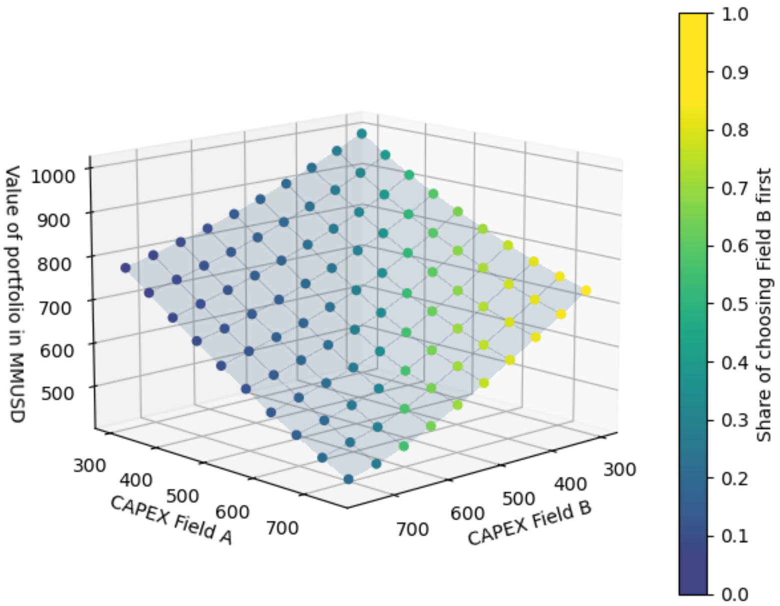

Figure 28: CAPEX sensitivity of both fields in the portfolio. While the fields' CAPEX values are shown in the xy-plane, we can observe the expected portfolio value on the z-axis. Additionally, the share of simulations suggesting Field B first is represented by the color grading.

Firstly, we analyze the sensitivity of the fields' CAPEX. CAPEX is a parameter that can be considered field-specific. Figure 28 shows how the dynamic of altering the CAPEX of the two fields does have a substantial impact on which field to tie back first. If the CAPEX of Field A is relatively high compared to Field B, we will go with Field B first in most cases. As expected from an essential cost parameter, lower CAPEX values result in a higher portfolio valuation and vice versa.

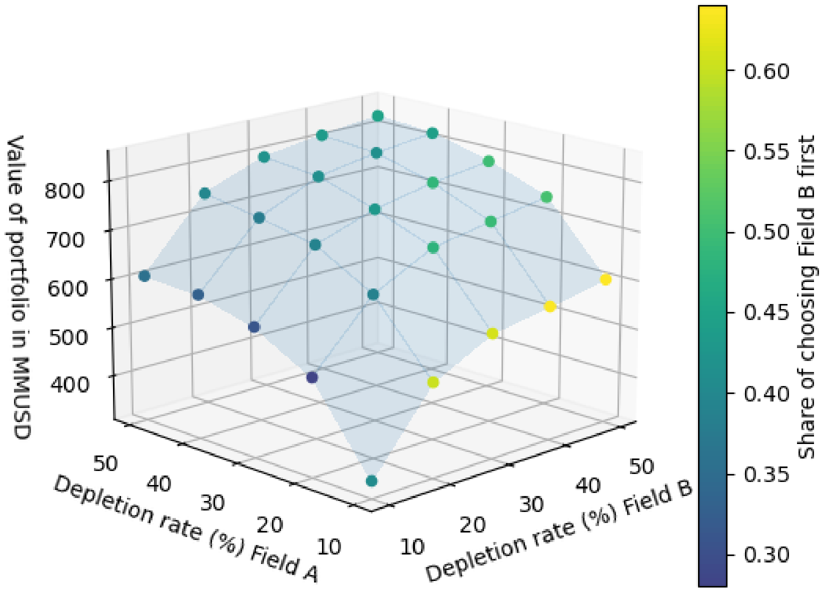

Figure 29: Depletion rate sensitivity of both fields in the portfolio. While the fields' depletion rates are shown in the xy-plane, we can observe the expected portfolio value on the z-axis. Additionally, the share of simulations suggesting Field B first is represented by the color grading.

Secondly, we look at the sensitivity of the depletion rate for each field. This parameter is also field-specific. We see the same pattern as with the CAPEX and EUR, as an increased depletion rate for a field favors selecting that field first, as we can see in Figure 29. When more oil can be recovered earlier, the cash flow turns positive at an earlier point in time, which favors the biggest field the most.

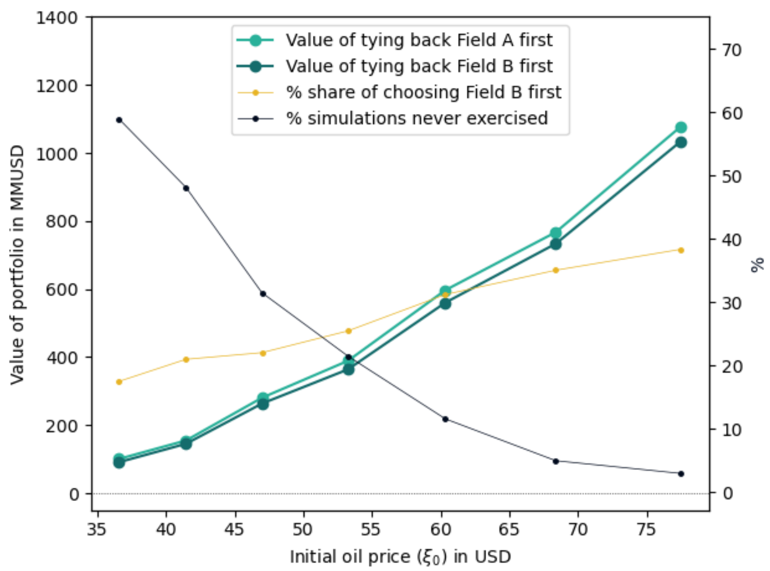

Figure 30: Sensitivity on the initial oil price. We can observe the expected portfolio value on the leftmost y-axis. The rightmost y-axis represents the share of simulations. This is used to show both the share of simulations suggesting Field B first and the share of simulations never exercised.

Thirdly, Figure 30 shows the sensitivity of the oil starting price, ξ_0. This parameter can be characterized as exogenous as it stems from the market pricing of oil. This means that it will affect both fields similarly and that it is not dependent on the field itself. The trend is that a higher ξ_0 increases the portfolio's total value and increases the relative share of choosing Field B first. The percentage of never-exercised paths decreases as the ξ_0 increases. In the far left end, few simulations prefer Field B first. We interpret this as the model favoring the bigger field when the surroundings tend towards non-profitability.

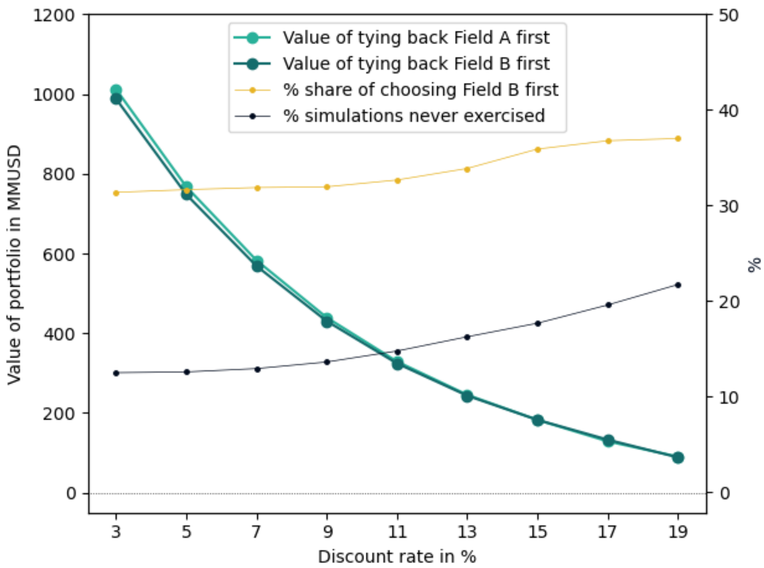

Figure 31: Discount rate sensitivity. We can observe the expected portfolio value on the leftmost y-axis. The rightmost y-axis represents the share of simulations. This is used to show both the share of simulations suggesting Field B first and the share of simulations never exercised.

Moreover, we want to show a sensitivity analysis on the discount rate in Figure 31. This is also an exogenous parameter, regardless if a risk-free or risk-adjusted rate is utilized. However, this parameter acts differently compared to the starting price sensitivity displayed above. As the discount rate increases, the portfolio's total value decreases, and the share of choosing B first and the percentage of paths never exercised increase. We believe that it is crucial to consider the increasing SHC as a key factor in explaining this trend. A higher discount rate lowers the present value of future revenues and costs. This can be troublesome if we pay a high CAPEX today while the revenues arrive far in the future, which could be the case with the increasing SHC. Thus, the model regression prioritizes tying back the field with the lowest EUR and the lowest CAPEX first and then allowing tying back the larger field later on. Then the available SHC is higher, and the costs and revenues can occur closer in time. Thus, the model's regression estimates a higher value for tying back Field B first when the discount rate increases as the SHC increases.

Based on our CROA approach, we have performed three additional sensitivities: the learning rate, the host's lifetime, and the volatility of the oil price in both the short and long term. As they mostly follow the same patterns as the field-specific and exogenous parameters shown above, we have put them in Appendix C to strengthen our results and to still keep this section concise and to the point.

Table 12: Field selection drivers summary. Based on the presented sensitivities and their impact on selecting which field to tie back first. The order of the parameters is intended to be in prioritized order of significance. The bottom three sensitivities are included in Appendix C.

Driver	Type	Impact	Decision-making insights
Field EUR	Field-specific	Yes, clear	The largest field EUR is favored.
CAPEX	Field-specific	Yes, clear	Lower CAPEX is favored.
Depletion rate (λ)	Field-specific	Yes, clear	High depletion rate is favored.
Initial oil price (ξ_0)	Exogenous	Yes, some	Overall, the largest field is favored. The trend is stronger when the initial oil price is low.
Discount rate (γ)	Exogenous	Yes, some	Incentivizes tying back a large proportion of the available EUR, translating to revenues, close to the time of paying the CAPEX.
Learning rate (ϵ)	Field-specific	Yes, some	Overall, the largest field is favored. Higher learning for the second field means the coherent first field will be slightly favored.
Host lifetime (L)	Exogenous	Yes, some	Overall, the largest field is favored. The trend is stronger when the lifetime of the host is shorter.
Oil price volatility (σ_ξ and σ_χ)	Exogenous	Yes, minor	Overall, the largest field is favored. The trend is stronger when the volatility is low, especially the long-term volatility.

Table 12 summarizes our insights from all sensitivities outlined in this section. Each parameter is categorized as what type of parameter it is and how much impact they have on the choice of field to tie back first and commented on how altering the parameter affects the field selection. Based on this insight, we see patterns of field-specific and exogenous parameters influencing choosing which field to tie back first. When adjusting field-specific parameters, the field that exhibits the most favorable parameter in terms of profitability tends to be prioritized for development. When altering the exogenous parameters, the largest field tend to be relatively more preferred when the exogenous parameters indicate lower portfolio values. A symptom of a portfolio's transitioning towards lower portfolio values is an increase in the percentage of paths never exercised. As the exogenous parameters shift the total portfolio value upwards, the relative share of selecting the smallest field first rises until reaching a threshold of about 35-40%. In summary, the field with the largest EUR is preferred unless field-specific parameters tilt the favor otherwise. The solution approach may prefer to first tie back the smallest field for the right field-specific and exogenous parameters. Also, the probability of this outcome is reduced if exogenous parameters transition the portfolio toward lower values.

6.4 The CROA versus Single ROA approach

In Section 6.1, we presented our results and discussed the implications of our CROA methodology compared to the myopic approach. In the essence of further evaluating our CROA approach, it is natural to also question how it weighs up against another ROA approach that disregards the value of the second field when making the decision to invest in the first one. In our case, this is the single ROA approach, which adopts the methodology of this paper and is based upon the concept of ROA already shown to be effective within O&G (Fleten et al., 2011; Jafarizadeh and Bratvold, 2020; Fedorov et al., 2021). We show the CROA model compares to a single ROA in terms of valuation and optimal investment timing. Furthermore, we explore this comparison under different variations of the SHC to see how these models perform in different environments.

Figure 32 shows the portfolio's total value for both the single ROA and the CROA approaches. A single ROA approach of each field individually results in a total value of 558.2 MMUSD. This can be compared to the results of the CROA that had an expected portfolio value of 581.2 MMUSD, 4.1% higher compared to the single ROA. Moreover, we see that the value of the first field in the Single ROA is larger than the the first field in the CROA. However, the contribution of the second field for each approach on average values only 31.7% of the total for the single ROA, while being 38.0% of the total value for the CROA. Thus, the value of the second field outweighs the loss of the first field from the single ROA to the CROA, bringing the total value of the CROA to be higher.

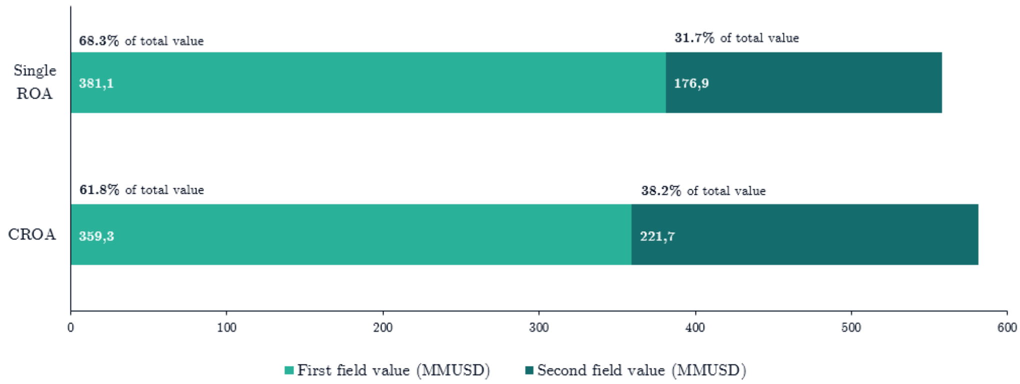

Figure 32: Valuation of the portfolio comparing the CROA to the Single ROA approach. Equal to the previous benchmark, this reflects the average expected valuation from choosing Field A first and then having the option of investing in Field B later.

Compared to the Single ROA, the CROA more often finds it optimal to invest earlier in the first field, as can be seen in Figure 33. The early investment is to provide enough time to develop both fields before the host is decommissioned. Thus, if a decision maker neglects the second field in its decision on when to tie back the first, it would tie back the first field too late for the second to produce sufficiently within the lifetime of the host. Furthermore, we continue by looking at the conditional investment timing of the second field, given that the licensees have decided to invest in the first field in year 1, the most common timing of the first field in the CROA, and in year 9, which is the most common timing for the first field in the single ROA. Investment in the first field in year 1 results in about 75.1% of the simulation cases for the second field being invested in at a later stage. The equivalent decision to invest in the first field in year 9 for the single ROA results only in a share of 63.9% of the simulations being exercised for the second field, which is much lower than the CROA. The most common tieback timing results are summarized in Table 13.

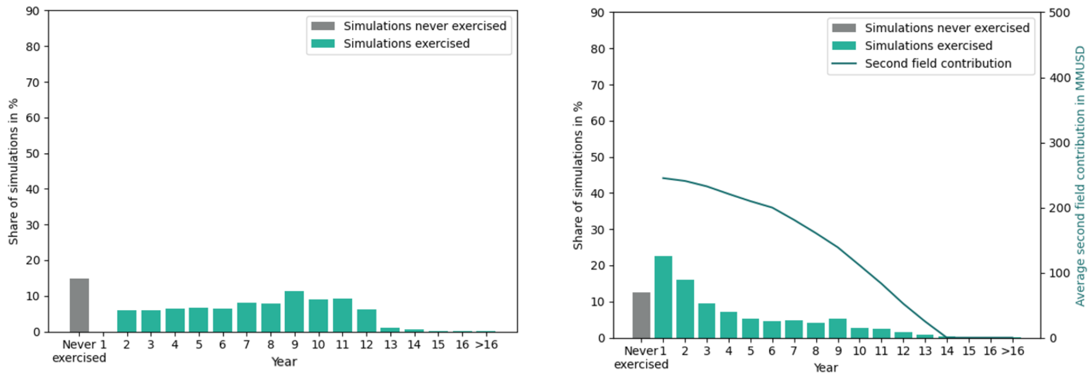

(a) Increasing with Single ROA

(b) Increasing with CROA

Figure 33: The optimal timing for investment in the first for all simulation cases as of the Single ROA and CROA approach. Equal to the CROA approach, the Single ROA approach generates decision rules based on risk factors. The Single ROA lacks the second field contribution estimation as it does not incorporate this in its evaluation.

Table 13: Most common tieback timing results of the Single ROA and CROA approach. In the majority of the simulation cases, it is optimal to start with tying back Field A.

Solution approach	Most common investment year for the first field	Most common investment year for the second field
Single ROA	9	13
CROA	1	12

While other papers have studied the production constraint of oil, water, and gas (see Fedorov et al. (2022a)), we have chosen to focus on production optimization on oil volumes. To still include the perspective on the mix of oil, water, and gas, we introduce three sensitivities of the SHC to show how water constraints affect the available production capacity. In Figure 34, we introduce two other shapes of the SHC to further investigate how the CROA and Single ROA behave under different conditions. We introduce a constant *flat high* SHC, reflecting a host with a lot of available capacity and not so much of a water constraint. Furthermore, we also introduce a decreasing *decreasing* SHC, reflecting a host with a restrictive future water constraint that limits the oil production.

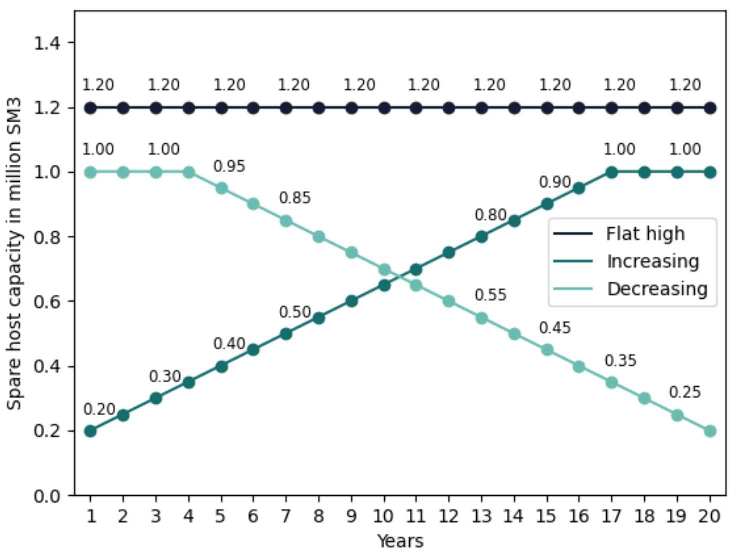

Figure 34: Spare host capacity with multiple variations. We now include a flat high and decreasing SHC to evaluate the robustness of the CROA framework.

In Hyldmo and Skudal (2022), we analyzed the effects of altering the SHC's shape and size in our ROA approach under reservoir uncertainty. The findings indicate a strong relation to the increasing value of waiting when the SHC has a positive slope. We see this also in Figure 35, where we present the valuation results by comparing the Single ROA and CROA approach for the three variations of the SHC. It shows how the CROA approach captures additional value in the *increasing* and the *flat high* case and is equal in the *decreasing* case. The performance is percentage-wise the strongest when there is value in waiting to invest, as with the *increasing* case.

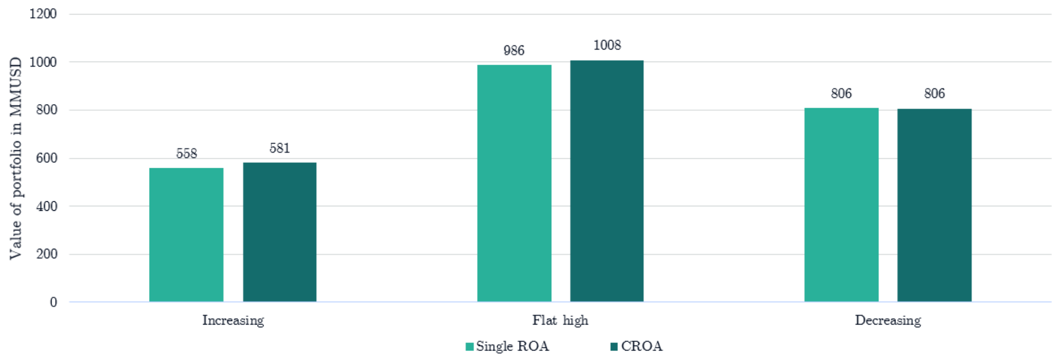

Figure 35: Visualization of the total portfolio value of the Single ROA and CROA approach using three variations of the spare host capacity. The marginal benefit of applying the CROA approach is decreasing along with the incentive to invest later.

Figure 35 shows a comparison of the Single ROA to the CROA for the different SHCs. The results show that the CROA percentage-wise outperforms the single ROA the most in the *increasing* case. In the *flat high* and the *decreasing* case, there is generally an incentive for both models to invest as early as possible. In the *flat high* case, the incentive to exercise early is partly outweighed by the ability to wait for better oil prices and more learning of the reservoir EUR for the second field, making the CROA slightly outperform the single ROA. In the *decreasing* case, the incentive of early exercise is so strong that 99.6% of the simulations are exercised in the first year for both models, meaning that the exercising and the total portfolio value for both approaches are exactly the same. Thus, this indicates that the CROA has a robust performance under different SHCs.

6.5 The CROA approach with perfect information

Until now, we have used a version of CROA with only a rough estimate of the value of the second field to mimic decision-making under uncertainty. However, we will show how this applied version of CROA compares to a CROA with perfect information about the value of the second field (Perfect CROA). We will do this for the same SHC versions as presented in Section 6.4. The corresponding valuation results can be seen in Figure 36.

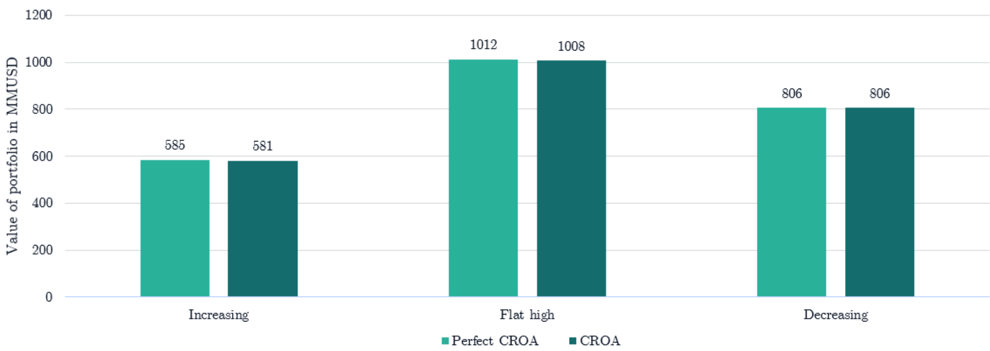

Figure 36: Visualization of the total portfolio value of the Perfect CROA and CROA approach using three variations of the spare host capacity. The Perfect CROA has perfect information on the simulated value distribution of the second field when making investment decisions.

The CROA is slightly lower or equal to the Perfect CROA in all cases. For the *increasing* and *flat high* cases, the CROA shows a slightly lower total result than the Perfect CROA. This comes from the binning effect of guessing the size of the second field in the CROA, where some of the simulations will be binned to a higher or lower second field value of what is exactly true. This may cause the simulation to be exercised slightly earlier or later than what it would optimally be with perfect information. However, this does not seem to impact the average values significantly. For the *decreasing* SHC case, both solution approaches score precisely the same when the result is reduced to three significant figures. In the same way as in Section 6.4, all approaches find no value in waiting to invest, and nearly all exercise in the first possible year. Thus, these empirical results conclude that the CROA approximates the Perfect CROA well.

The CROA approach demonstrates remarkable robustness when applied to multiple variations of the SHC. From comparing the CROA approach to the single ROA and the Perfect CROA, we observe a trend of CROA being valuable and robust in altering SHC variations. These findings highlight the strength and versatility of the CROA approach, making it a valuable tool for various applications requiring robust analysis and decision-making in dynamic market environments.

7 Final remarks

Real options analysis serves as a valuable instrument for assessing the optimal timing and valuation of marginal fields. Nevertheless, it is important to acknowledge that it also has its limitations. Our problem modeling simplifies a complex environment and does not account for all relevant risk factors. Therefore, an interesting avenue for future work would be to incorporate additional risk factors into the model to enhance its accuracy and robustness. Specifically, we suggest reservoir modeling and accounting for operational risk in more detail.

Quality and selection of parameter values used in testing the solution approach are of utmost importance for the model's validity. The simplification of parameters, such as static field potential and discount rate, may not fully capture the dynamic nature of the O&G industry. It is crucial to recognize that the accuracy of ROA is highly dependent on the input parameters, which can be challenging to estimate accurately in advance. Limiting our study is the potential bias in selecting parameters used in the base case, as they are based on subjective expert knowledge and intuition. To address this limitation, we recommend that companies adopting our methodology customize the parameter set to their specific case and perform stress-testing and sensitivity analysis accordingly, ensuring a more unbiased and tailored analysis.

A characteristic of the CROA approach is the necessary assumption about a static EUR of the first field. This is a limiting factor, as we optimally would like to model this EUR estimate as a distribution similar to the second field. This assumption can be relaxed by including multiple EUR cases for the first field. This equals increasing M, as we have done with $M = 3$. With computational capabilities, it is possible to extend this model with many more. Another possibility is to explore heuristics in how the second field contribution is incorporated into the ROA of the first field. We suggest such an endeavor as future work to be especially beneficial to the literature.

Several minor factors limiting the methodology can easily be addressed in future work. A limitation lies in the assumption that all costs related to infrastructure removal will be handled at the end of the host's lifetime. In reality, the timing of infrastructure decommissioning can vary, and this aspect should be considered on a case-by-case basis. At the current state, the tieback to both fields inherent equal tariff policy and assumed OPEX model. Realistically, these will differ, which can be reflected in future research. Other considerations could include modeling the CAPEX and the host's lifetime as stochastic processes. These limitations highlight areas for further research and improvement in real options analysis for evaluating marginal field investments.

8 Conclusion

This paper has investigated the application of ROA and CROA to optimize economic value and make better decisions related to marginal O&G development. By examining a case of one host and two marginal fields, we have contributed to the research by showing the significant impact of accounting for uncertainties in oil prices, field EUR, and the value of later adding a second field to decision-making processes. The paper focused on two main objectives; developing novel modeling and solution approaches that lead to better decision-making and providing insight into field selection drivers. We address this by 1) assessing the ability of the CROA-adjusted LSM to capture the value of managerial flexibility by determining the optimal exercise timing, 2) exploring decision rules resulting in a superior valuation of marginal fields in a project portfolio, and 3) stress-testing the LSM model to identify under what circumstances one field should be tied back before another, contributing to the understanding of field selection drivers.

Firstly, we designed a Single ROA model accounting for oil price simulations with a two-factor stochastic model and a Bayesian learning model of field EUR. Internally in the ROA, we utilized an optimization model to generate production profiles based on the updated estimates of the fields' EUR. The ROA model was solved using an LSM solution approach. We expanded this framework by creating a CROA, a compound options approach incorporating ROA evaluations of the second field while analyzing the decision for the first field. We ran one CROA for each order of choosing to tie back Fields A and B before choosing the ordering yielding the highest value. To assess the effectiveness of our approach, we compared the results against a Single ROA approach, based on the same methodology as here, but without the portfolio view, and a myopic approach, which is widely utilized in the O&G industry.

The results demonstrate that our methodology is a valuable decision-making tool in the O&G industry. By including an estimate of the value of the second field in the LSM, we observed significant improvements in the evaluation of timing and valuation for marginal fields, leading to reduced downside potential and enhanced consideration of managerial flexibility. Following the optimal decision rules with regard to timing from the weighted CROA model, the average portfolio valuation is 25.4 % higher than the myopic DCF approach. Our sensitivity analysis further supports the model validity, which examined the effects of various factors such as CAPEX, depletion rate, and discount rate and how they affect the choice of which field to tie back first. We also found that higher risk corresponds to a higher economic value of managerial flexibility.

Additionally, we explore the concept of active learning of field EUR by gathering more data about the marginal field by adding the option to drill an appraisal well. This can be found in Appendix A. Here, decision-makers are able to purchase reduced reservoir uncertainty by gathering more information about the field, which they can use in the decision-making process. The approach utilizes machine-learning techniques and a regression decision tree to establish an investment policy for appraisal wells based on the current price and reservoir information available. Our findings indicate that the quality of the data obtained through our modeling approach lacks the reliable attributes necessary for constructing a robust and dependable ML model. This work contributes to exploring what data and analysis models may prove effective for evaluating options to drill appraisal wells.

Several intriguing directions for future research can be pursued based on this paper to further enhance the realism of the decision problem. Besides the future work suggested in relation to the mentioned limitations of the methodology, we propose the following. Firstly, it would be interesting to see how different tariff policies impact the selection of which field to tie back first. Secondly, incorporating more market uncertainty factors, such as making the costs, for instance, the CAPEX, be correlated with the market prices to better account for the supply and demand effects of the costs. Thirdly, adding gas and water production to the model can make the production rates and constraints applied in the optimization model more appropriate. Furthermore, the integration of environmental, social, and governance (ESG) parameters into the model can provide another perspective on the decision problem. Finally, the exploration of competition modeling can shed light on the decision strategies employed by fields and hosts in forming partnerships.

Our CROA offers a valuable tool for decision-makers, enabling them to assess the economic viability of a portfolio of assets more effectively and holistically. By considering multiple risk factors and incorporating the value of waiting to invest, our methodology enhances decision-making and leads to more optimal outcomes. These findings have important implications for regulators, field operators, and host operators. Our study provides valuable insights into the importance of accounting for uncertainty and the benefits of employing a concurrent portfolio perspective using a CROA approach in the O&G industry.

Bibliography

Anand, J., Oriani, R. and Vassolo, R. S. (2007), Managing a portfolio of real options, *in* 'Real Options Theory', Vol. 24, Emerald Group Publishing Limited, pp. 275–303.

Armstrong, M., Galli, A., Bailey, W. and Couët, B. (2004), 'Incorporating technical uncertainty in real option valuation of oil projects', *Journal of Petroleum Science and Engineering* **44**(1-2), 67–82.

Bacati, F., Arcangeletti, G., Breuskin, B., Giolo, R., Muguerra, P., Radicioni, A., Shaiek, S. and Sibaud, E. (2020), 'How technology enables a lower cost for subsea tiebacks', *OTC Offshore Technology Conference - Day 1 Mon, May 04, 2020* . D011S011R001.doi: `https://doi.org/10.4043/30667-MS`.
URL: *https://doi.org/10.4043/30667-MS*

Bakker, S. J., Kleiven, A., Fleten, S.-E. and Tomasgard, A. (2021), 'Mature offshore oil field development: Solving a real options problem using stochastic dual dynamic integer programming', *Computers & Operations Research* **136**, 105480.

Bhattacharya, S., Chatterjee, K. and Samuelson, L. (1986), 'Sequential research and the adoption of innovations', *Oxford Economic Papers* **38**, 219–243.

Bickel, J. E. and Bratvold, R. B. (2008), 'From uncertainty quantification to decision making in the oil and gas industry', *Energy Exploration & Exploitation* **26**(5), 311–325.

Bjerksund, P. and Ekern, S. (1990), 'Managing investment opportunities under price uncertainty: From" last chance" to" wait and see" strategies', *Financial Management* pp. 65–83.

Brealey, R., Myers, S. C., Allen, F. and Mohanty, P. (2012), *Corporate finance*, Tata McGraw-Hill Education.

Brennan, M. J. and Schwartz, E. S. (1985), 'Evaluating natural resource investments', *Journal of business* pp. 135–157.

Chang, C.-P. and Lin, Z.-S. (1999), 'Stochastic analysis of production decline data for production prediction and reserves estimation', *Journal of Petroleum Science and Engineering* **23**(3-4), 149–160.

Chorn, L. and Shokhor, S. (2006), 'Real options for risk management in petroleum development investments', *Energy Economics* **28**(4), 489–505.

Cortazar, G. and Schwartz, E. S. (1998), 'Monte carlo evaluation model of an undeveloped oil field', *Journal of Energy Finance & Development 3(1), 73–84.* . `https://www.sciencedirect.com/science/article/pii/S1085744399800696`.

Cox, J. C., Ingersoll, J. E. and Ross, S. A. (1985), 'An intertemporal general equilibrium model of asset prices', *Econometrica* **53**(2), 363–384.
URL: *http://www.jstor.org/stable/1911241*

Dias, M. (1997), The timing of investment in e & p: Uncertainty, irreversibility, learning, and strategic consideration, *in* 'SPE Hydrocarbon Economics and Evaluation Symposium', OnePetro.

Dittrick, P. (2019), 'Operators commit to investment in deepwater gulf of mexico'. Retrieved 03 November 2022, from https://www.ogj.com/exploration-development/article/14071431/operators-commit-to-investment-in-deepwater-gulf-of-mexico.

Dixit, A. and Pindyck, R. (1994), *Investment Under uncertainty*, Princeton University Press.

Eckhardt, R. (1987), 'Stan ulam, john von neumann, and the monte carlo method', *Los Alamos Science* **15**(131-136), 30.

Ekern, S. (1988), 'An option pricing approach to evaluating petroleum projects', *Energy Economics* **10**(2), 91–99.

Fama, E. F. (1970), 'Efficient capital markets: A review of theory and empirical work', *The journal of Finance* **25**(2), 383–417.

Fama, E. F. and French, K. R. (1993), 'Common risk factors in the returns on stocks and bonds', *Journal of financial economics* **33**(1), 3–56.

Federal Reserve (2020), 'Why does the federal reserve aim for inflation of 2 percent over the longer run?'. Retrieved 20 May 2022, from https://www.federalreserve.gov/faqs/economy_14400.htm.

Fedorov, S. (2021), 'Cost effective development of small discoveries on the norwegian continental shelf'. https://ntnuopen.ntnu.no/ntnu-xmlui/handle/11250/2987791.

Fedorov, S., Hagspiel, V., Haseldonckx, S., Rønning, A., Haugsgjerd, J. and Rogstad, R. (2022a), Evaluation of tieback developments for marginal oil fields with timing flexibility.

Fedorov, S., Hagspiel, V. and Lerdahl, T. (2021), 'Real options approach for a staged field development with optional wells', *Journal of Petroleum Science and Engineering, 205, 108837* . https://doi.org/10.1016/j.petrol.2021.108837.

Fedorov, S., Lavrutich, M., Hagspiel, V. and Lerdahl, T. (2022b), 'Risk and benefit sharing schemes in oil exploration and production', *Energy Economics* **116**, 106401.

Fleten, S.-E., Gunnerud, V., Hem, Ø. D. and Svendsen, A. (2011), 'Real option valuation of offshore petroleum field tie-ins', *Journal of real options* **1**, 1–17.

Galli, A., Armstrong, M. and Jehl, B. (1999), 'Comparison of three methods for evaluating oil projects', *Journal of Petroleum Technology* **51**(10), 44–49.

Galli, A., Jung, T., Armstrong, M. and Lhote, O. (2001), Real option evaluation of a satellite field in the north sea, *in* 'SPE Annual Technical Conference and Exhibition', OnePetro.

Gamba, A. (2003), 'Real options valuation: A monte carlo approach', *Faculty of Management, University of Calgary WP* (2002/3).

Goodwin, D. (2013), Schwartz-smith two-factor model in the copper market: before and after the new market dynamics. Master's thesis at Lund University. https://lup.lub.lu.se/luur/download?func=downloadFile&recordOId=3809118&fileOId=3809122.

Hong, A., Bratvold, R. B. and Lake, L. W. (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* **22**(03), 1144–1160.

Husy, S. (2011), 'Marginal Fields: Technology enables profitability/Marginal Fields and their Challenges', *Paper presented at the Offshore Technology Conference, Houston, Texas, USA, May 2011. Paper Number: OTC-21382-MS* . https://doi.org/10.4043/21382-MS.

Hyldmo, T. and Skudal, M. H. (2022), 'Sequential tieback of marginal fields under reservoir uncertainty: a real options approach with bayesian learning'. Project Thesis submitted at NTNU in the Fall semester of 2022 at the Department of Industrial Economics and Technology Management.

Ian, L. and Noeth, S. (2004), *Economics of Petroleum Production: Profit and Risk, vol. 1*, Multi-Science Publishing, Brentwood, United Kingdom.

IBM (2023), 'Decision trees'. Retrieved 24 April 2023, from https://www.ibm.com/topics/decision-trees.

Jafarizadeh, B. and Bratvold, R. B. (2009), Taking real options into the real world: asset valuation through option simulation, *in* 'SPE Annual Technical Conference and Exhibition', OnePetro.

Jafarizadeh, B. and Bratvold, R. B. (2012), 'Two-factor oil-price model and real option valuation: an example of oilfield abandonment', *SPE Economics & Management* **4**(03), 158–170.

Jafarizadeh, B. and Bratvold, R. B. (2015), 'Oil and gas exploration valuation and the value of waiting', *The Engineering Economist* **60**(4), 245–262.

Jafarizadeh, B. and Bratvold, R. B. (2020), 'Sequential exploration: valuation with geological dependencies and uncertain oil prices', *SPE Journal* **25**(05), 2401–2417.

Laine, J. (1997), Option valuation of field development projects, *in* 'SPE Hydrocarbon Economics and Evaluation Symposium', OnePetro.

Lander, D. M. and Pinches, G. E. (1998), 'Challenges to the practical implementation of modeling and valuing real options', *The quarterly review of economics and finance* **38**(3), 537–567.

Lei, G., Silva, T. L. and Stanko, M. (2021), A mathematical programming approach to determine optimal tieback development concepts of two offshore oil fields in the barents sea, *in* 'The 31st International Ocean and Polar Engineering Conference', OnePetro.

Lin, J., de Weck, O., de Neufville, R. and Yue, H. K. (2013), 'Enhancing the value of offshore developments with flexible subsea tiebacks', *Journal of Petroleum Science and Engineering, 102, 73-83* . https://doi.org/10.1016/j.petrol.2013.01.003.

Lin, J., de Weck, O. and MacGowan, D. (2012), 'Modeling epistemic subsurface reservoir uncertainty using a reverse wiener jump–diffusion process', *Journal of Petroleum Science and Engineering* **84-85**, 8–19. https://doi.org/10.1016/j.petrol.2012.01.015.

Longstaff, F. A. and Schwartz, E. S. (2001), *Valuing American Options by Simulation: A Simple Least-Squares Approach*, The Review of Financial Studies Spring 2001 Vol. 14. No. 1, pp. 113-147.

Lund, M. W. (1999), Real options in offshore oil field development projects, *in* '3rd Annual Real Options Conference, Leiden', Citeseer.

Lund, M. W. (2000), 'Valuing flexibility in offshore petroleum projects', *Annals of Operations Research* **99**(1), 325–349.

Markowitz, H. M. (1991), 'Foundations of portfolio theory', *The journal of finance* **46**(2), 469–477.

Maslin, E. (2019), 'Tapping tiebacks', *Offshore Engineer* . Retrieved 30 October 2022, from `https://www.oedigital.com/news/467643-tapping-tiebacks`.

Masnadi, M. S., El-Houjeiri, H. M., Schunack, D., Li, Y., Englander, J. G., Badahdah, A., Monfort, J.-C., Anderson, J. E., Wallington, T. J., Bergerson, J. A. et al. (2018), 'Global carbon intensity of crude oil production', *Science* **361**(6405), 851–853.

McDonald, R. (2013), *Derivatives Markets*, third edn, Pearson.

Moldovan, T. (2010), 'The conjugate prior for the normal distribution'. Retrieved 18 October 2022, from `https://people.eecs.berkeley.edu/~jordan/courses/260-spring10/lectures/lecture5.pdf`.

Myers, S. C. (1977), 'Determinants of corporate borrowing', *Journal of financial economics* **5**(2), 147–175.

Neumann, A., Siliverstovs, B. and von Hirschhausen, C. (2006), 'Convergence of european spot market prices for natural gas? a real-time analysis of market integration using the kalman filter', *Applied Economics Letters* **13**(11), 727–732.
URL: *https://doi.org/10.1080/13504850500404258*

Norges Bank (2022), 'Inflasjon'. Retrieved 24 November 2022, from `https://www.norges-bank.no/tema/pengepolitikk/Inflasjon/`.

Paddock, J. L., Siegel, D. R. and Smith, J. E. (1988), 'Valuation of claims on real assets: The case of offshore petroleum leases', *The Quarterly Journal of Economics 103(3), 479–508.* . `https://www.jstor.org/stable/1885541`.

Pindyck, R. S. (1999), 'The long-run evolutions of energy prices', *The energy journal* **20**(2).

Rankin, J. M. (1986), 'Kalman filtering approach to market price forecasting'. 8291. `https://lib.dr.iastate.edu/rtd/8291`.

Sales, L., Jäschke, J. and Stanko, M. (2021), 'Early field planning using optimisation and considering uncertainties: Study case: Offshore deepwater field in brazil.'. `https://doi.org/10.1016/j.petrol.2021.109058`.

Santos, S. M., Gaspar, A. T. F. and Schiozer, D. J. (2017), 'Risk management in petroleum development projects: Technical and economic indicators to define a robust production strategy', *Journal of Petroleum Science and Engineering* **151**, 116–127.

Santos, S. M., Gaspar, A. T. and Schiozer, D. J. (2018), 'Managing reservoir uncertainty in petroleum field development: Defining a flexible production strategy from a set of rigid candidate strategies', *Journal of Petroleum Science and Engineering* **171**, 516–528.

Schwartz, E. S. and Smith, J. E. (2000), 'Short-term variations and long-term dynamics in commodity prices', *Management Science* **46**(7), 893–911.

Scott, S. L. (2017), 'Tie-backs and seafloor pumping improve economics in deepwater developments'. Retrieved 30 October 2022, from `https://www.worldoil.com/magazine/2017/february-2017/features/tie-backs-and-seafloor-pumping-improve-economics-in-deepwater-developments`.

Sharpe, W. F. (1963), 'A simplified model for portfolio analysis', *Management science* **9**(2), 277–293.

Sharpe, W. F. (1964), 'Capital asset prices: A theory of market equilibrium under conditions of risk', *The journal of finance* **19**(3), 425–442.

Smith, J. E. and McCardle, K. F. (1999), 'Options in the real world: Lessons learned in evaluating oil and gas investments', *Operations Research* **47**(1), 1–15.
URL: *https://doi.org/10.1287/opre.47.1.1*

Society of Petroleum Engineers (2023), 'Field developments and technical solutions – marginal fields'. Retrieved 28. May 2023, from `https://www.spe.org/en/training/courses/field-developments-and-technical-solutions-marginal-fields/#:~:text=Marginal%20fields%20refer%20to%20discoveries,the%20vicinity%20and%20profitable%20consumers`.

Stoisits, R. F., Bashagour, H. M. and Su, C.-G. (2010), Application of integrated production and reservoir modeling to optimize deepwater development, *in* 'International Oil and Gas Conference and Exhibition in China', OnePetro.

Subbey, S., Christie, M. and Sambridge, M. (2004), 'Prediction under uncertainty in reservoir modeling', *Journal of Petroleum Science and Engineering* **44**(1-2), 143–153.

Thakur, P. (2017), Chapter 6 - fluid flow in cbm reservoirs, *in* P. Thakur, ed., 'Advanced Reservoir and Production Engineering for Coal Bed Methane', Gulf Professional Publishing, pp. 75–90.
URL: *https://www.sciencedirect.com/science/article/pii/B9780128030950000065*

The Economist (2022), 'Norway is profiting embarrassingly from war in europe'. Retrieved 10 October 2022, from `https://www.economist.com/europe/2022/09/08/norway-is-profiting-embarrassingly-from-war-in-europe`.

The International Energy Agency (2022*a*), 'Norway 2022 - energy policy review'. Retrieved 28 October 2022, from `https://iea.blob.core.windows.net/assets/de28c6a6-8240-41d9-9082-a5dd65d9f3eb/NORWAY2022.pdf`.

The International Energy Agency (2022*b*), 'The oil and gas industry in energy transitions'. Retrieved 10 October 2022, from `https://iea.blob.core.windows.net/assets/4315f4ed-5cb2-4264-b0ee-2054fd34c118/The_Oil_and_Gas_Industry_in_Energy_Transitions.pdf`.

The Norwegian Petroleum Directorate (2019), 'Resource report 2019'. Retrieved 10 October 2022, from `https://www.npd.no/en/facts/publications/reports/resource-report/resource-report-2019/fields/`.

The Norwegian Petroleum Directorate (2020), 'Resource report 2020: Exploration'. Retrieved 10 October 2022, from `https://www.npd.no/en/facts/publications/reports/resource-report/resource-report-2020/`.

The Norwegian Petroleum Directorate (2022), 'Resource report 2022: Norway is a competitive and long-term supplier of oil and gas to europe'. Retrieved 10 October 2022, from `https://www.npd.no/en/facts/news/general-news/2022/resource-report-2022-norway-is-a-competitive-and-long-term-supplier-of-oil-and-gas-to-europe/`.

The Norwegian Tax Administration (2022), 'Petroleum tax regulations - the norwegian tax administration'. Retrieved 10 October 2022, from https://www.skatteetaten.no/en/business-and-organisation/reporting-and-industries/industries-special-regulations/oil-and-gas-exploration-and-petroleum-tax/the-petroleum-tax-office-and-the-petroleum-tax-system/petroleum-tax-regulations/.

Tilley, J. A. (1993), Valuing american options in a path simulation model, *in* 'Transactions of the Society of Actuaries', Citeseer.

Tourinho, O. A. et al. (1979), The option value of reserves of natural resources, Technical report, University of California at Berkeley.

UN Secretary-General António Guterres (2023), 'Secretary-general calls on states to tackle climate change 'time bomb' through new solidarity pact, acceleration agenda, at launch of intergovernmental panel report'. Retrieved 28. May 2023, from https://press.un.org/en/2023/sgsm21730.doc.htm.

US. Energy Information Administration (2022), 'What drives crude oil prices?'. Retrieved 14 October 2022, from eia.gov/finance/markets/crudeoil/spot_prices.php.

World Bank Group (2014), 'Commoditiy markets outlook'. Retrieved 14 October 2022 https://www.worldbank.org/content/dam/Worldbank/GEP/GEPcommodities/commodity_markets_outlook_2014_july.pdf.

Appendices

A Active learning of reservoir size with appraisal well options - A machine learning approach

In the paper, we have presented a robust framework for evaluating marginal oil fields using compound real options with LSM as a solution approach. Still, we do not consider any possibilities to reduce the EUR uncertainty further. Hence, we want to test the effect of doing active learning on fields, i.e., gathering further data through, for example, the drilling of an appraisal well. The benefit of drilling an appraisal well is to reduce the uncertainty of the size of the reservoir. In our case, this reduces the standard deviation of the current best guess distribution. Such a possibility can be modeled as a periodical, in our case annual, option to reduce EUR uncertainty in exchange for the cost of drilling such a well. Thus, the challenge becomes when to drill an appraisal well and if our solution approach is appropriate to indicate optimal decisions.

We analyze the problem from the perspective of decision tree regression from the field of supervised machine learning (ML). Decision trees are a powerful ML tool that can help with "(...) data mining and knowledge discovery tasks" (IBM, 2023). At their core, decision tree regression is a series of interconnected nodes representing splits into attributes that aim to increase the purity of the subsequent nodes.

The main advantage of decision trees is their simplicity and interpretability. They provide a clear and intuitive way to understand how a particular decision was made, which can be especially helpful for decision-making in complex scenarios. Our regression tree splits nodes on the attribute most predictive of dividing the expected value of drilling into separate baskets. We have an intuitive understanding of when appraisal wells can be worthwhile to drill, and this is reflected in our hypothesis for this approach:

Hypothesis

> We presume it only should be useful to drill in certain cases. The decision maker has lower incentives to drill the appraisal wells when new information is less likely to change the initial investment decision: when the best guess predicts either a large or a small reservoir. In contrast, information about middle-sized reservoirs close to a break-even state will yield a high value to obtain. In general, it is beneficial to drill as early as possible.

A.1 Solution approach

We assume that drilling an appraisal well will reveal each path's true volume ($v*$). For all available exercise years, we do a path-wise evaluation of whether it was beneficial to drill an appraisal well or not.

From there, the active-learning path will converge to the true size of the reservoir, while in the passive-learning path, the decision maker will remain uncertain about the true volume, still learning about it passively in the following years. This is illustrated in Figure 37.

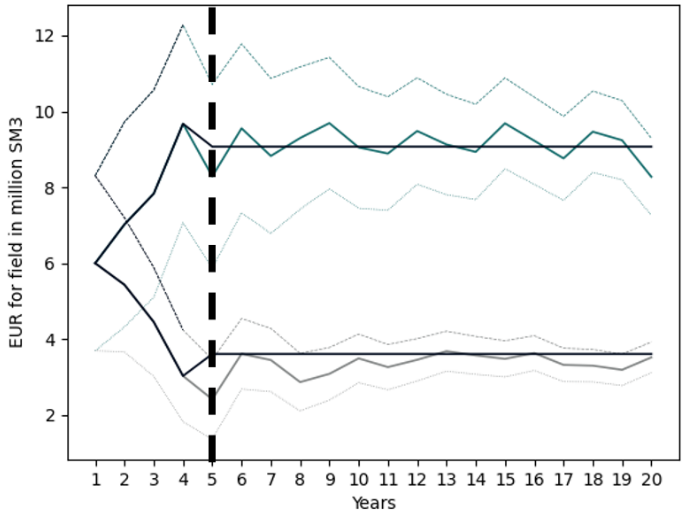

Figure 37: Appraisal drilling example from drilling in year 5 for two trajectories.

To find out whether drilling a well is valuable, we make a path-wise comparison of the active-learning paths to the passive-learning paths. The value of information (VoI) is then calculated from the following equation:

$$\text{VoI} = V_{\text{appraisal}} - V_{\text{no appraisal}} \tag{24}$$

$$\text{Reward} = \text{VoI} - C_{\text{drilling}} \tag{25}$$

where V is the project value with and without drilling, and C is the cost of drilling the appraisal well. We want to train an ML model, more specifically, a decision tree regression, to make a rule of whether and when to invest. As far as we know, this is a new contribution to the literature and has never been done before. After conducting the decision tree analysis, trying other supervised learning algorithms is also relevant.

We want to force active learning on all paths to see the behavior of the results. Moreover, we want to do this for all paths in all available years for investment. Each state is represented by the year, current oil price, and current best guess of the reservoir size. The goal is to evaluate the path-wise reward, as shown in the equation above. We input the state and the corresponding reward to the decision tree, and the tree chooses the best splitting points by Mean Squared Error (MSE) as a purity criterion for splitting the nodes. We have two alternative decision tree models.

We train on 90% of our simulations and test our model on 10% of our simulations to verify our results. Thus, the verified model visualizes decision rules appropriate for decision-makers facing the challenge of investing in an appraisal well.

A.2 Results

The hypothesis cannot be confirmed at this time due to the lack of distinct patterns in the results. Despite extensive efforts, including the implementation of Decision trees and parameter tuning, the intended result was unfortunately not achieved. Our efforts to improve the accuracy of our model were hampered by several factors, including the quality of available data, the performance limitations of the Decision tree we used, and the complexity of the problem we were trying to solve. Most profoundly, we postulate that the data quality generated through our LSM solution approach does not allow us to provide any significant pattern due to the complexity of the problem. Naturally, our attributes of the main uncertainties seem to be too weak indicators of whether it is beneficial to drill an appraisal well or not.

In Figure 38, we present our training data used and the subsequent feature importance analysis performed on our parameters. The data is generated from drilling in year 3, as an example representative of other drilling years. Figure 38b shows our parameters' relative importance and does not directly indicate how well they predict our target data, i.e., the VoI from drilling. Noteworthy still is *best guess EUR* minuscule impact on prediction. We can view the parameters in the context of two-parameter pairs with a random reference. The *current oil price* and *best guess EUR* are values we know in the model at any point in time. These are related to *expected future oil prices* and *real EUR*, respectively, which are unknown to the model but that we include highlighting their impact.

	Expected future oil prices (USD$)	Current oil price (USD$/SM3)	Real EUR (MSM3)	Best guess EUR (MSM3)	Random
0	3852.209524	242.953573	6.803610	7.392068	50
1	6257.241836	392.082688	5.106937	3.825532	4
2	11415.586877	310.441478	5.635583	6.534949	56
3	3716.414630	224.749599	8.080870	8.745472	86
4	6293.914546	347.455343	5.890813	7.414813	89
...	...	...	...	...	...
8995	4932.793688	294.079424	5.861789	7.121111	40
8996	8433.740990	459.205227	6.083321	3.282217	27
8997	10449.609905	304.258663	5.618081	5.776655	88
8998	6044.157630	382.101321	6.811712	1.841985	54
8999	8594.445907	385.761511	9.895105	11.197786	3

9000 rows × 5 columns

(a) Training data with attributes (year 3).

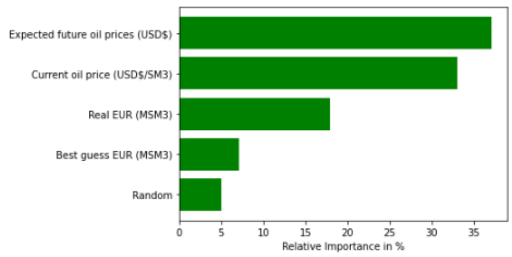

(b) Feature importance (year 3).

Figure 38: Visualization of data and feature analysis used in the data cleaning and data engineering steps of ML.

As expected, we observe that *real EUR* is more critical for performance than *best guess EUR*. Additionally, the *expected future oil prices* are higher than the *current oil price*. In Figure 39, we observe an example of a regression decision tree generated from the training data shown in Figure 38a. This is an example tree showing the outcome of drilling in year 4. We can seemingly spot a degree of conformity with our hypothesis in the tree. In particular, the most distinct, orange node represents the highest average VoI of drilling an appraisal well.

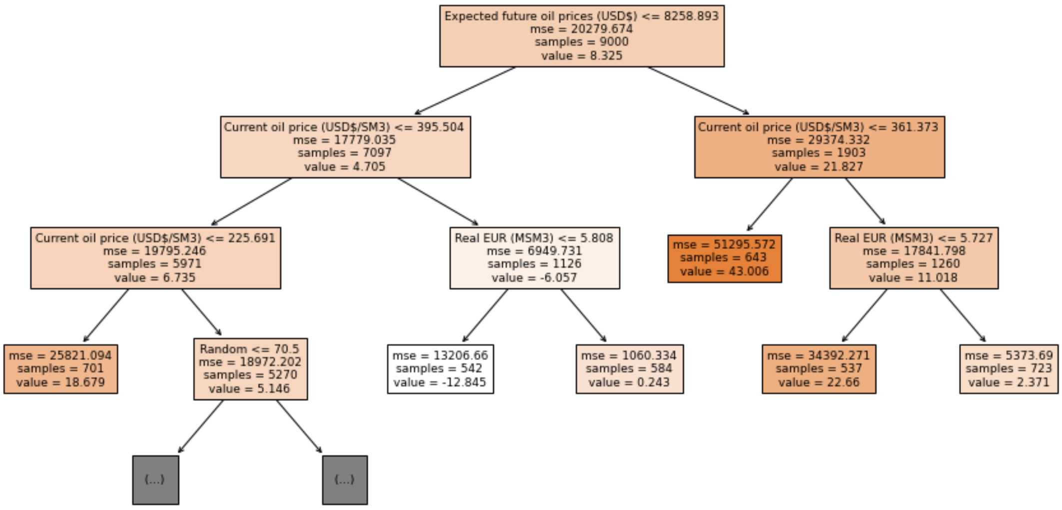

Figure 39: Regression tree for appraisal drilling example year 3.

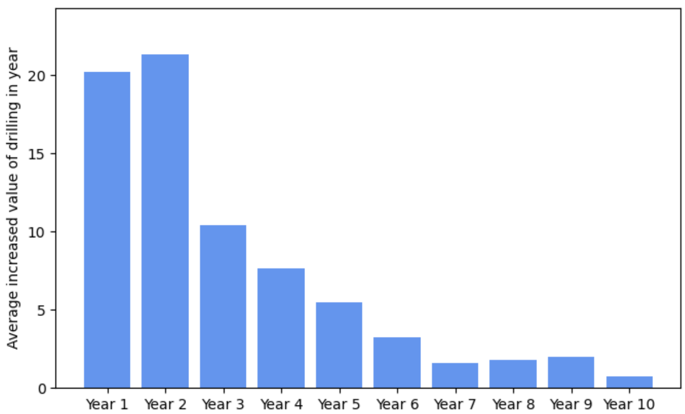

Figure 40: Drilling year sensitivity.

One result that makes sense is sensitivity to the drilling year. Figure 40 shows that, on average, early drilling seems profitable. This indicates that you should do it early if you want to drill. Furthermore, we want to try to relax the constraint of getting the true volume t^* after drilling the appraisal well. As perfect information about the reservoir size is the most valuable, as of Hyldmo and Skudal (2022), we would expect the value of drilling to decrease if the standard deviation is not zero after drilling. Thus, driving the standard deviation up to 5% of the size of the true field after drilling lowers the value of drilling from 21.3 MMUSD down to 19.3 MMUSD when drilling in year 2 and from 6.0 MMUSD down to 4.0 MMUSD when drilling in year 5.

However, looking closer into trees with different shuffles of data in the train and test sets, our study suggests that the current approach using decision trees is insufficient to provide accurate predictions. We highlight three considerations confirming that the proposed method is a poor fit for our problem: (1) inconsistent categorization of valuable states, (2) too high error and error variance, and (3) the model does not differentiate cases with a high enough value on average to break even with the costs of a realistic appraisal well.

The categorization inconsistency is observed from tweaking model parameters, and we experience high sensitivity to any change. Given a new random state in the train-test-split function, we can obtain different root splits. This means minor differences in training data result in a separate main decision point, pointing to a lack of robustness.

In Table 14, we can see the performance measure of the model. When examining regression models, R^2 is the most common statistic as it measures the goodness of fit. We find it startlingly low, as it resides between zero and one, and one reflects a perfect fit.

Table 14: Model performance measures for the test set for drilling in year 4.

Performance measure	Value	Unit
R-squared (goodness of fit)	0.01	
Mean Squared Error (MSE)	20 163	
Root Mean Squared Error (RMSE)	142	MMUSD
Mean Absolute Error (MAE)	72	MMUSD

None of the other measures speaks up for the model either, with them all being immense. Root Mean Squared Error (RMSE) and Mean Absolute Error (MAE) are most easily interpretable, as these are denoted in MMUSD. RMSE is more sensitive to outliers, but both show a monetary measure of the average error of the model. Comparing these to the highest-valued average node at 43 MMUSD, we can see the downside stretches far down. Recall that we are talking about the VoI; thus, this substantial downside risk is present before any cost is accounted for. Lastly, the MSE score in Table 14 only represents the average, but this error equates to double that for the most interesting node. It results in a node RMSE of 226.5 MMUSD.

A.3 Conclusions

In this section, we proposed an active learning approach to reduce EUR uncertainty through drilling appraisal wells using ML techniques. We used a regression decision tree to create a policy for when to invest in an appraisal well, but the modeling did not verify our hypothesis. We found that data quality obtained through our modeling approach does not hold enough reliable attributes to build a trustworthy ML model.

For future work, we recommend revisiting the topic of active learning through the option to drill an appraisal well in depth. As we propose that our data quality is our main challenge, we encourage future studies to examine realistic models of the decision problem with data attributes more potent for an ML approach, for instance, by adding more geological-dependent attributes. Additionally, it could prove efficient to expand on the complexity of decision trees or explore other ML approaches, such as neural networks (NN) and support-vector machines (SVM).

B Simulation robustness

Assessing the model's robustness is essential when making investment decisions in ROA. A robust model can produce consistent and dependable results even when underlying assumptions or input parameters vary. This helps decision-makers identify potential risks and opportunities associated with investment projects and make informed decisions. In the LSM solution approach, model robustness is closely linked to the number of simulations (N) used in the analysis. The number of simulations required in LSM can affect the accuracy and precision of the results, making it crucial to consider when evaluating investment projects. In this section, we summarize our model's robustness and discuss the impact of the number of simulations on LSM's accuracy and reliability.

In Figure 41, we see a convergence in the results as N increases sufficiently. Above 1 000 simulations, the increment in simulations does not affect the validity of the valuation results to any large degree.

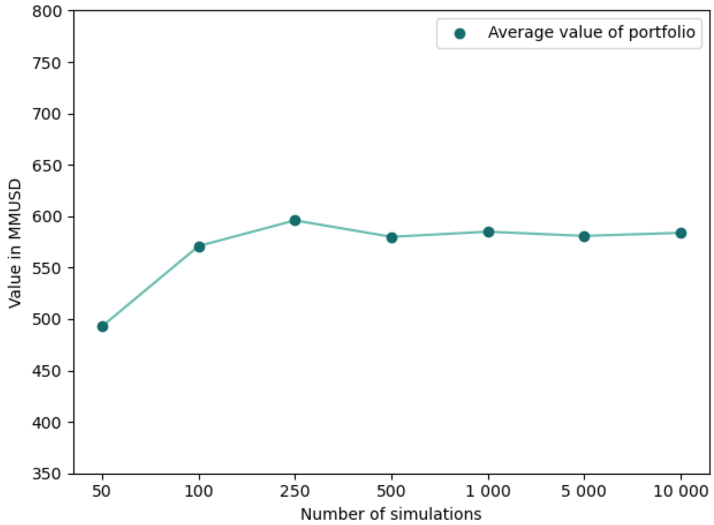

Figure 41: Convergence plot for the number of simulations in the LSM. It shows the value output for our model in MMUSD for different numbers of simulations.

This trend is also visible when examining the exercise times for the three scenarios in the CROA. However, small changes are visible in the timing results from 1 000 to 10 000 simulations. This is shown in Figure 42 and Figure 43, for 1 000 and 10 000 simulations, respectively. As expected, the curves are more smooth for a higher number of simulations.

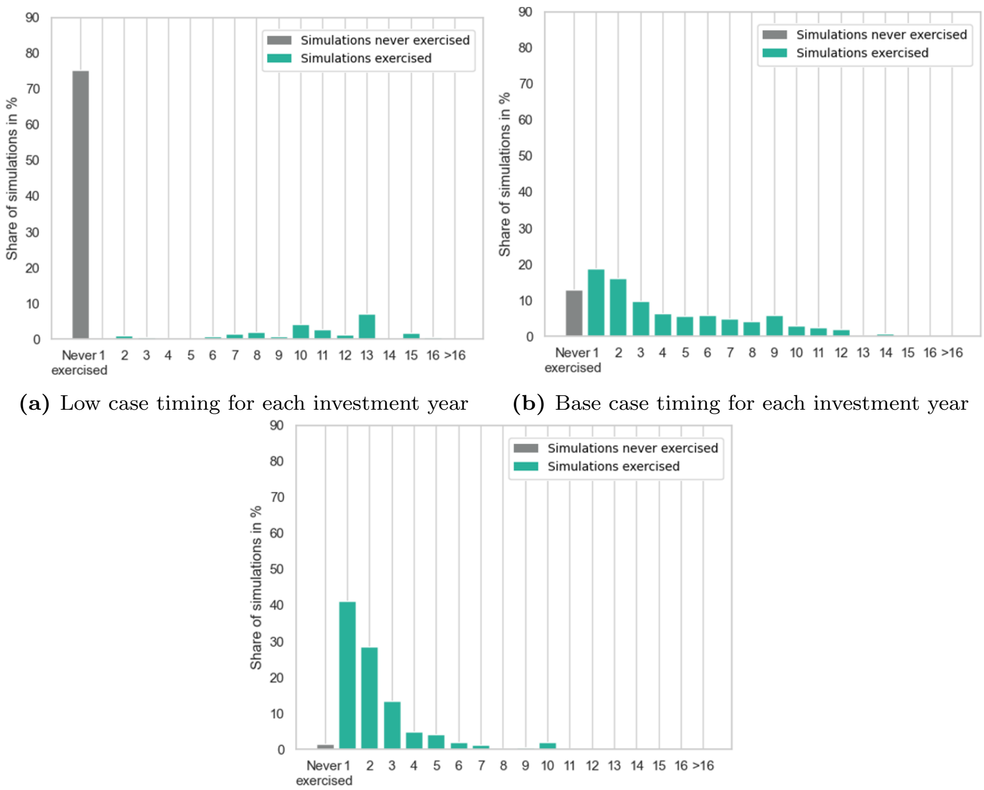

(a) Low case timing for each investment year

(b) Base case timing for each investment year

(c) High case timing for each investment year

Figure 42: Optimal tieback timing - 1 000 simulations.

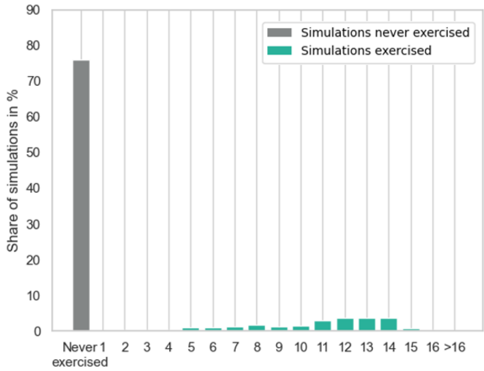

(a) Low case timing for each investment year

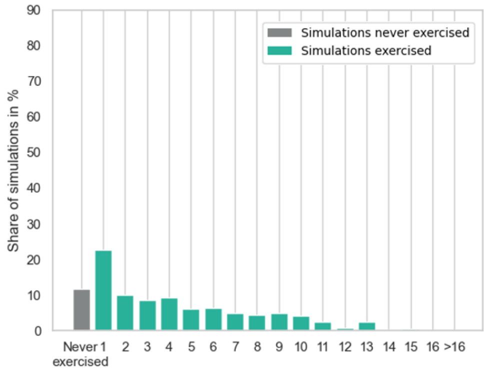

(b) Base case timing for each investment year

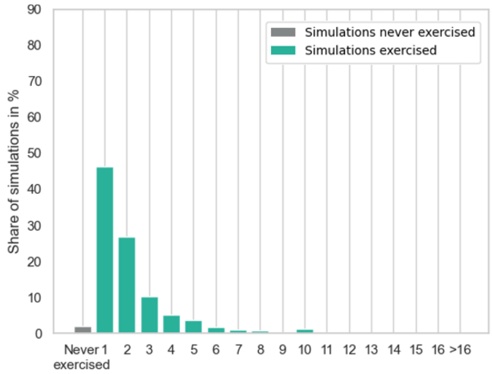

(c) High case timing for each investment year

Figure 43: Optimal tieback timing - 10 000 **simulations**.

Therefore, we employ a total of 10 000 simulations to evaluate all our results in Section 6. However, for the sensitivity analysis, we find it appropriate to proceed with 10 000 simulations. This number allows us to effectively illustrate the desired patterns while maintaining a manageable computational workload.

C Additional sensitivities on drivers of field selection

As a continuation to Section 6.3, we present an additional three sensitivities on a hunt to see what parameters affect the choice of which field to tie back first.

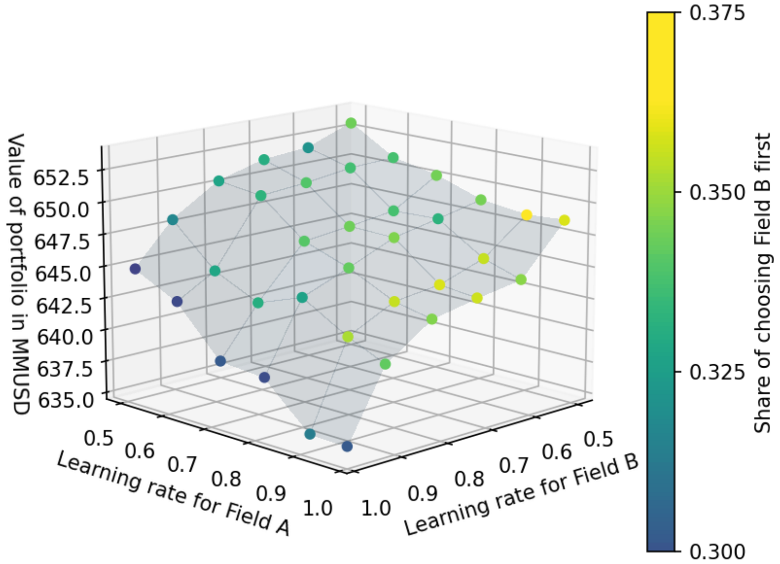

Figure 44: Learning rate sensitivity.

Figure 44 show a general pattern that learning is valuable. This confirms the intuition highlighted by the learning rate sensitivity in Hyldmo and Skudal (2022). Moreover, the trend is that relatively bigger field-specific learning of Field B compared to Field A increases the relative value of Field B compared to Field A. As our model is such that it first learns of Field B after tying back Field A, an increase in the relative value of Field B will imply that we more often tie back Field A first such that we can benefit from the increased learning in Field B.

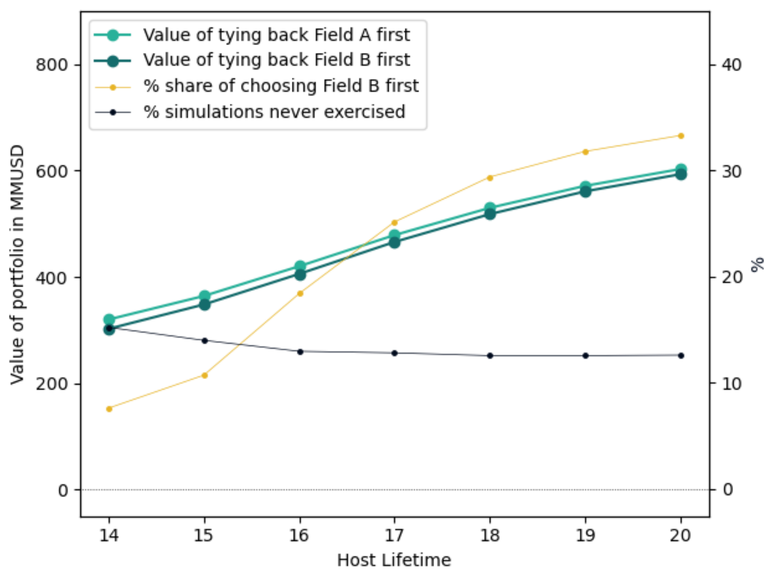

Figure 45: Sensitivity on the lifetime of the host.

Secondly, we have a sensitivity to the lifetime of the host. The host's lifetime is an exogenous parameter outside the field operator's control. Figure 45 shows that the portfolio's total value decreases as both the value of flexibility, the amount of oil the field operator can extract, and the probability of tying back the second field decreases. Consequently, we see that the share of paths never exercised goes up as the host's lifetime goes down. We also see that the relative share of choosing to tie back Field B goes down with the decrease in profitability.

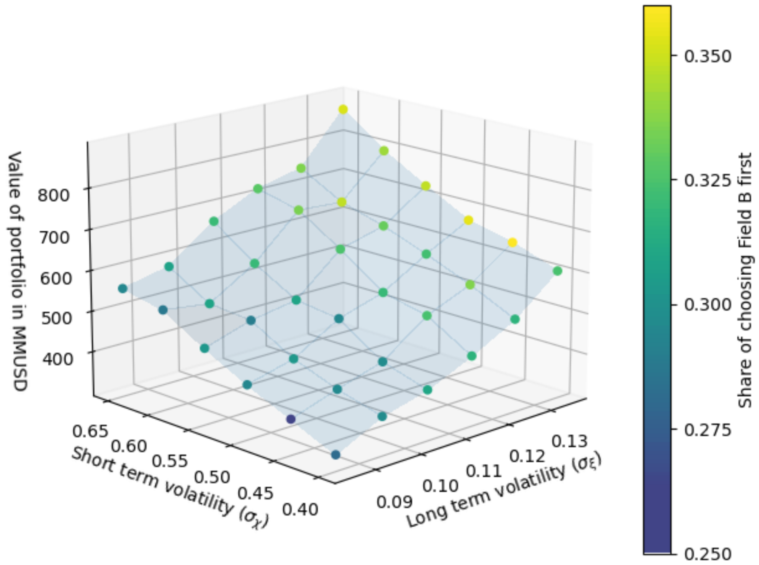

Figure 46: Sensitivity on the oil price volatility.

Finally, we also see the sensitivity of the oil price volatility. Based on Figure 46, we can see a weak pattern of the share of choosing Field B first going down when the volatility goes down. This makes sense, as the quantity produced will be relatively more important for the revenues when the oil prices are flatter. Thus, we would expect the share of choosing the bigger field to go up, which makes sense with what we see in Figure 46. Moreover, we see the same pattern as of Dixit and Pindyck (1994), where a general increase in both the short- and long-term volatility increases the portfolio's total profitability.

NTNU
Norwegian University of
Science and Technology

NORWEGIAN PETROLEUM
DIRECTORATE