Computers & Operations Research 136 (2021) 105480

Contents lists available at ScienceDirect

Computers and Operations Research

journal homepage: www.elsevier.com/locate/cor

Mature offshore oil field development: Solving a real options problem using stochastic dual dynamic integer programming

Steffen J. Bakker *, Andreas Kleiven, Stein-Erik Fleten, Asgeir Tomasgard

Department of Industrial Economics and Technology Management, Norwegian University of Science and Technology, Alfred Getz veg 3, NO-7491 Trondheim, Norway

ARTICLE INFO

Keywords:
SDDiP
Real options
Mature oil field development
Multistage stochastic integer programming

ABSTRACT

Oil and gas companies are facing low output prices and are forced to focus on the development of mature fields. Relevant investment decisions for operators include lifetime-enhancing activities, such as drilling new wells or permanent shutdown. We study the problem of optimal timing of investments in mature oil and gas fields in the presence of price uncertainty, which is an example of a complex real options problem consisting of a portfolio of interdependent options. We formulate a multistage stochastic integer programming model that incorporates a detailed representation of the uncertain oil price, and demonstrate how such a complex real options problem can be efficiently solved using the Stochastic Dual Dynamic Integer Programming algorithm. The paper presents a numerical example based on realistic data and discusses our computational results. We find that only a small number of Markov states are required to represent the uncertain price process, while obtaining convergence of the lower and upper bounds of the objective function. The value of stochastic solution of 11% is considerable in this example. It is concluded that the shutdown decision tends to be postponed as a result of high decommissioning costs, high discount rates, high price uncertainty and low operational expenditures, while it generally is accelerated if the decommissioning costs increase over time.

1. Introduction

In 2018, an approximate volume of 70% of the world's oil and gas production came from mature fields (O'Brien et al., 2016). Mature fields are defined to be those that have reached the plateau production phase, the decline phase may have started and the fields are reaching the end of their economic life. The number of such fields has been steadily increasing and can be expected to grow in the near future. Together with the relatively low recent output prices, this has forced oil and gas operators to focus on mature fields. Optimal development of mature fields means that operators have to make strategic decisions, such as starting production enhancing activities or field shutdown. Oil price fluctuations can have a considerable effect on these decisions, particularly for marginal fields.

The field of real options has attempted to address similar problems (Ekern, 1988; Clarke and Reed, 1990; Lund, 2000), but the model formulations tend to restrict either the price processes or the options available. In mathematical programming, related work has been done in the setting of general offshore oil field development (Nygreen et al., 1998; Gupta and Grossmann, 2012). However, this literature tends to focus on decisions taken at the beginning of a field's lifetime. In this work, we take elements from real options and mathematical programming and develop a multistage stochastic integer programming model (MSIP) that maximizes the net present value of an offshore mature asset. The model allows a detailed representation of the uncertain price process as well as the inclusion of a portfolio of options. In particular, the price process is represented using the Short-Term Long-Term (STLT) model from Schwartz and Smith (2000). We solve the developed MSIP using the Stochastic Dual Dynamic integer Programming (SDDiP) algorithm (Zou et al., 2019).

The main contribution of this work is to demonstrate how complex real options problems, consisting of a portfolio of interdependent options, can be formulated as an MSIP and efficiently solved using SDDiP. This approach allows for evaluating multiple activities, several stochastic factors and problem-specific constraints. The resulting model can be efficiently solved using existing SDDiP packages (Ding and Ahmed, 2019; Dowson and Kapelevich, 2021). We present a new model formulation for the problem of optimally developing a mature offshore oil field. To describe price uncertainty, we include the STLT model and estimate its parameters based on recent oil futures data. Based on input from a Norwegian operator, we construct a numerical example to perform a computational study. Here, we find that only a small number

* Corresponding author.
E-mail address: steffen.bakker@ntnu.no (S.J. Bakker).

https://doi.org/10.1016/j.cor.2021.105480
Received 27 March 2020; Received in revised form 30 June 2021; Accepted 12 July 2021
Available online 16 July 2021

of Markov states are needed to represent the uncertain price process. In addition, we analyze the optimal policy to show how decisions may adapt to realizations of the uncertain price process. Using the optimal MSIP-policy as opposed to the solution of a deterministic problem that uses expectations of uncertain factors leads to an improvement in the objective function of around 11%. Also, we investigate the effect of different factors on investment timing. Besides the traditional finding that uncertainty delays investment, we find that increasing decommissioning costs on mature fields can lead to acceleration of the shutdown decision.

The remainder of this paper is organized as follows: Section 2 provides an extensive literature review. The MSIP model is presented in Section 3. This section includes a problem description, treatment of the uncertain price process as well as a formulation of the model. In Section 4 we describe how our model can be solved using the SDDiP algorithm. Further, Section 5 presents estimation results of the STLT model as well as computational results from a numerical example, including a discussion based on investment timing. The conclusions are presented in Section 6.

2. Literature review

In this section we review the relevant literature in the fields of real options and mathematical programming with a particular focus on the development of oil fields. For a more general review of oil field development we refer to Babadagli (2007) and Tavallali et al. (2016), while Vrålstad et al. (2019) and Bakker et al. (2019) focus on shutdown decisions and plug and abandonment (P&A) operations on mature offshore oil fields.

2.1. Real options

In the context of oil fields, the field of real options has focused on the option to prematurely abandon (or to shut down) a field (Paddock et al., 1988; Ekern, 1988). Traditional real options literature is based on financial option pricing and the assumption that a replicating portfolio of financial instruments can be constructed to replicate the returns of the real option. Myers and Majd (1983) and McDonald and Siegel (1985) use this approach to value an abandonment option employing standard no-arbitrage arguments. The prevailing stochastic processes, such as the oil price or production rate, are typically assumed to evolve as a Geometric Brownian Motion (GBM). Analytical results such as the optimal abandonment time can be calculated (Olsen and Stensland, 1988; Clarke and Reed, 1990). More recently, this approach has been used to value the option of switching from oil to gas production (Støre et al., 2018).

Other solution approaches for real options problems are dynamic programming and decision analysis (Dixit and Pindyck, 1994). Early work has been conducted by Bonini (1977), who formulated a discrete time dynamic programming model for abandoning a project with uncertain future revenues. Smith and Mccardle (1998, 1999) integrated the option pricing approach and the decision analytic approach in an elegant way. They distinguished between market uncertainties and private uncertainties, where the respective typical examples are oil prices and production rates. Market risks can be perfectly hedged by traded financial instruments, while private risks cannot be perfectly hedged.

Applications of real options to petroleum field development and abandonment include the following. Lund (2000) describes a stochastic dynamic programming model for evaluating the development of a marginal offshore oil field under uncertainty, where the author treats both market and private risk. Several simplifications in the representation of the problem are required to obtain a manageable model in terms of size. Most notably, the uncertain factors are restricted to have only two or three realizations. Dias (2004) presents an overview of real options that occur in petroleum exploration, development and production. As an example, Dias (2004) discusses how the option to expand production by drilling additional wells can be included in traditional real options literature.

When a real options problem is formulated as a stochastic dynamic program, it is often solved using an Approximate Dynamic Programming (ADP) technique (Powell, 2007). The optimal policy value obtained from ADP approaches is often assessed by estimating a dual bound using information relaxations (Brown et al., 2010). The most commonly used technique to value real options in the last two decades is the Least Squares Monte Carlo Approach (LSM) by Longstaff and Schwartz (2001), which is also known as the simulation-and-regression method. Applications of this approach to the petroleum industry are given by Fleten et al. (2011) and Jafarizadeh and Bratvold (2012), where the latter implements the STLT model.

2.2. Mathematical programming

Within the field of mathematical programming, an extensive amount of research has been undertaken on the development of oil and gas fields (Sullivan, 1988; Khor et al., 2017). Here, the main focus tends to be on decision making during early phases of the life-cycle of a field. In contrast, the present work contributes to the literature by focusing on the development of mature fields.

A related topic is the planning and scheduling of (infrastructure) investment and operation in offshore oil/gas fields. This includes the design and operation of pipelines, fields, wells and installations. Decisions have to be made regarding issues such as the network structure, platform capacities, well locations and production rates. In addition, both physical and financial restrictions have to be taken into account. Most of the work on this problem uses a deterministic approach, and the formulated models tend to be multiperiod mixed-integer programming models (Nygreen et al., 1998; Iyer and Grossmann, 1998; Lin and Floudas, 2003; Gupta and Grossmann, 2012). Integrality conditions, multiple periods, detailed descriptions of operational or technical parts and other features make these problems generally difficult to solve. When introducing uncertainty, their complexity increases even further.

Several attempts have been made to analyze planning problems in the petroleum industry under uncertainty and different decomposition techniques have been used to solve multistage stochastic programming models. Jonsbråten (1998) presents a mixed-integer programming model for the optimal design and operation of an oil field under price uncertainty. The problem is solved using a variant of the progressive hedging algorithm that allows for integer variables. However, the complexity of this problem forces the author to allow for only three scenarios for the uncertain price process. Goel and Grossmann (2004) and Gupta and Grossmann (2014) make use of a Lagrangian Decomposition to solve a multistage stochastic program for the planning of an offshore oil or gas field infrastructure, while Tarhan et al. (2009) use a duality based branch-and-bound procedure for the same problem. These problems are also extremely difficult to solve. As a result, the uncertainties are described using high and low values, leading to scenario sets of size four to eight. Even though these models are formulated as multistage problems, in the case studies, two-stage stochastic programs are solved to ease the computational burden.

Another approach to solve multistage stochastic optimization problems is the Stochastic Dual Dynamic Programming (SDDP) algorithm (Pereira and Pinto, 1991). This algorithm is based on a dynamic programming formulation with a nested cost-to-go function. It can be considered an adaptation of the nested Benders decomposition (Birge, 1985). The applicability of this algorithm is restricted to linear problems with continuous variables. Recent advances have made it possible to include integer variables in this framework. Zou et al. (2019) extended the SDDP algorithm to allow for binary state variables and refer to this as Stochastic Dual Dynamic Integer Programming (SDDiP). A conceptual treatment of the SDDiP algorithm is given in Section 4. Applications of the SDDiP algorithm can be found in hydropower plant scheduling (Hjelmeland et al., 2019) and electric power infrastructure planning (Lara et al., 2019). We contribute by presenting a new application of the SDDiP algorithm to an oil field development problem.

2.3. Real options vs. stochastic programming

When options are available in portfolios, the standard approaches to real options have their scope limitations, particularly when there are interdependencies between options. One approach to resolve these challenges and to allow for interdependencies in a portfolio of real options is to extend the traditional LSM method (Maier et al., 2019, 2020). The authors formulate the model as a multistage stochastic integer program (MSIP) and describe the uncertainty using a Markov process. Another example of a work cast as a real options problem in a stochastic mixed-integer programming framework is by Wang and de Neufville (2004). The authors make use of a two-step procedure, where first options are identified within a simulation model, and subsequently these options are used in a stochastic mixed integer linear program. This highlights a key difference between real options and stochastic programming, as outlined by Wallace (2010). Real options require options to be predefined, while stochastic programming might find options that are not directly clear. Another difference is that stylized real options analysis more easily allows for gaining analytical insights, while stochastic programming can be used to solve larger and more complex problems. This work aims to combine both fields by solving a complex real options problem using the SDDiP algorithm from stochastic programming.

3. Mathematical model

In this section, we provide a problem definition, discuss the price process that we use and formulate the main sets and variables, constraints and objective function for the corresponding MSIP.

3.1. Problem description

We consider a mature offshore oil field that is approaching the end of its economic life. The field has gone through primary and secondary recovery phases, but there are still certain activities available to extend its lifetime. These activities to enhance production include the drilling of new wells, sidetracking existing production wells (slot recovery) or other techniques such as Enhanced Oil Recovery (EOR). These activities are considered to be available real options. The production facility can be a platform or a floating/subsea production system. When the installation is producing, it incurs a yearly operational expenditure (OPEX). Shutting down the field requires permanently plugging and abandoning the wells as well as decommissioning the installation, which are costly operations. Considering this problem and making clear dissemination of our results, we restrict ourselves to a field that solely produces oil. However, the analysis can quite easily be extended to other commodities such as natural gas.

The two most important variables determining the value of a petroleum asset are future prices and future production. Since oil prices are by nature stochastic and have considerable fluctuation, we allow for uncertainty in prices. In addition, production profiles might fluctuate as well. To model production profiles from existing wells as well as potential targets, operators make use of different methods, such as complex reservoir simulators or more standard decline curves. Since we are considering mature fields, operators have probably been able to collect a considerable amount of data on historical production, reservoir properties and the possibilities for activities that can extend the lifetime of a particular field. On the Norwegian Continental Shelf, operators are obliged to present an annual status report to the Norwegian Petroleum Directorate, which compiles forecasted production profiles for existing wells, as well as optional activities (Norwegian Petroleum Directorate, 2019). As a result, we assume that estimated production profiles are given for the field, as well as the optional activities.

Focusing on the planning level of the problem at hand, strategic decisions such as shutting down a field, or further developing it, are typically taken on a yearly basis. These decisions require considerable planning and preparation and are therefore taken at least a year in advance. On the Norwegian Continental Shelf, the production license is granted for a period of ten years (Petroleumsloven, 1996, § 3-9). Depending on field conditions, the time horizon for a mature field can also be up to ten years. As a result, it is natural to formulate this problem in discrete time with annual time steps and a finite time horizon.

Finally, the problem of strategic decision making on mature offshore oil fields aims to maximize the net present value of the field. This can be achieved by optimally deciding when development activities should be started and/or when the field has to be shut down. Yearly profits equal the difference between revenues and cost, where the revenues are calculated based on realized spot prices and given production profiles.

3.2. Price process

To model price behavior, both the GBM and mean reverting processes such as the Ornstein–Uhlenbeck (OU) process, have often been used. The GBM tends to be able to reflect the long-term behavior of the oil price, while a mean reverting process can capture short-term fluctuations. However, these simple models fail to capture the full dynamics of the spot price for oil. More realistic price behavior can be captured by the two-factor Short-Term Long-Term (STLT) model by Schwartz and Smith (2000).

The STLT model has a stochastic equilibrium price level, while short-term deviations, which are the difference between the spot price and the equilibrium level, are also stochastic and mean revert to zero. In the STLT model, the logarithm of the spot price at time t (p_t) can be decomposed as follows:

$$\log p_t = \chi_t + \xi_t, \tag{1}$$

where χ_t is referred to as the short-term factor (short-term deviations) in prices and ξ_t denotes the long-term factor (equilibrium price level). The short-term deviations are assumed to revert to zero following an OU process, while the equilibrium level is assumed to follow a Brownian Motion Process. These factors are not directly observable, but they can be estimated from historical future contracts. Schwartz and Smith (2000) present a risk neutral version of the STLT model that can be used for option valuation. The advantage of using the risk neutral process is that we can discount cash flows at a risk-free rate. We have to discretize the process to make use of the STLT model in a multistage optimization problem. A possible discretization of this process is given by Jafarizadeh and Bratvold (2012), who present a risk neutral STLT model in discrete-time:

$$\tilde{p}_t = \exp\left(\tilde{\chi}_t + \tilde{\xi}_t\right), \tag{2}$$

$$\tilde{\chi}_t = \tilde{\chi}_{t-1} e^{-\kappa \Delta t} - \left(1 - e^{-\kappa \Delta t}\right) \frac{\lambda_\chi}{\kappa} + \sigma_\chi \epsilon_\chi \sqrt{\frac{1 - e^{-2\kappa \Delta t}}{2\kappa}}, \tag{3}$$

$$\tilde{\xi}_t = \tilde{\xi}_{t-1} + \mu_\xi^* \Delta t + \sigma_\xi \epsilon_\xi \sqrt{\Delta t}. \tag{4}$$

Here, $\tilde{p}_t$, $\tilde{\chi}_t$ and $\tilde{\xi}_t$ are the risk-neutral equivalents to the spot price, short-term factor and long-term factor respectively, while σ_χ and σ_ξ are volatilities and ϵ_χ and ϵ_ξ are correlated standard normal random variables with correlation coefficient $\rho_{\chi\xi}$. In addition, κ is the mean-reversion coefficient, λ_χ is a risk premium that specifies a reduction in the drift for the short-term process and μ_ξ^* is the risk neutral drift of the equilibrium level, $\tilde{\xi}_t$. Finally, Δt denotes the length of time period t in years, allowing for time stages with different lengths. Note that in the long run, under the risk neutral measure, the expected spot prices grow at a rate $\exp(\mu_\xi^* + \frac{1}{2}\sigma_\xi^2)$. For a thorough discussion related to model parameters, we refer to Schwartz and Smith (2000).

3.3. Sets and variables

We give a brief definition of some of the main sets and variables, and their corresponding domains. For a complete overview over all sets and indices, parameters and variables used in the MSIP, we refer to Appendix B.

The time structure of the problem is flexible with regards to both resolution and horizon. We therefore define the set $\mathcal{T}$, indexed by t and τ, representing time periods or stages of possibly different lengths. The set $\mathcal{A}$, indexed by a, contains the activities that are available to increase production. These activities include sidetracks (slot recovery), drilling of new wells and/or enhanced oil recovery (EOR) methods. Some activities cannot be started in a certain stage, so $\mathcal{A}_t$ gives the activities that can be started in stage t. Moreover, we define $\Gamma_t = \{(a,\tau) \in \mathcal{A} \times \mathcal{T} \mid a \in \mathcal{A}_\tau, \tau \le t-1\}$.

We have two types of binary *state* variables. First, we have the shutdown variables y_t equaling one if the installation has been shut down in previous time periods, or if the shutdown decision is made in time period t. Second, the activity variables x_{at} equal one if activity a is undertaken in time period t. We make use of a time-lag of one period in the decision process, since the decision to shut down the installation, or to start an activity, affects production from the next period and onward. The binary state variables have the following domains:

$$x_{at}, y_t \in \{0,1\}, \quad \forall a \in \mathcal{A}_t, t \in \mathcal{T}. \tag{5}$$

In addition, we have *stage* or *local* variables capturing the production (q_t) and cost (c_t) in each time period, with the following domains:

$$q_t, c_t \in \mathbb{R}_0^+, \quad \forall t \in \mathcal{T}. \tag{6}$$

3.4. Constraints

The constraints of the MSIP can be related to the calculation of production profiles or the relationships between activities.

3.4.1. Production profile

The production in time period t, denoted by the continuous variable q_t, can be described by the following set of constraints:

$$q_t \le Q_t^{BASE} + \sum_{(a,\tau)\in\Gamma_t} Q_{a\tau t}^{ADD} x_{a\tau} \quad \forall t \in \mathcal{T}, \tag{7}$$

$$q_t \le M_t^{PROD}(1 - y_{t-1}), \quad \forall t \in \mathcal{T}, \tag{8}$$

where M_t^{PROD} equals the maximum achievable production in time period t, and Q_t^{BASE} and $Q_{a\tau t}^{ADD}$ are parameters that give the quantity of the base production and additional production due to activities, respectively. So, the production equals zero if the shutdown decision has been made in the previous period. Otherwise, it equals a base production level increased with additional quantities if activities have been started.

There might be restrictions on the production capacity. In particular, the lower bound (Q^L), might represent a minimum production that should be obtained to keep the installation in operation.

$$Q^L \le q_t \le Q^U, \quad \forall t \in \mathcal{T}. \tag{9}$$

3.4.2. Activities

The decision maker can make life-enhancing *activities*, defined by x_{at} for all $a \in \mathcal{A}_t, t \in \mathcal{T}$, or perform permanent *shutdown*, denoted y_t for all $t \in \mathcal{T}$. The first restriction on the activities is that they can only be executed once:

$$\sum_{t\in\mathcal{T}} x_{at} \le 1, \quad \forall a \in \mathcal{A}. \tag{10}$$

We might have relations between different activities. A simple example can be a restriction on the maximum number of activities that can be started each year, given by $L \in \mathbb{Z}_+$. This is enforced by:

$$\sum_{a\in\mathcal{A}_t} x_{at} \le L, \quad \forall t \in \mathcal{T}. \tag{11}$$

However, we can easily add other relationships or logical constraints. This makes it easy to deal with real option portfolios, where there are interdependencies between the real options. For example, when activity a_1 has been started, activity a_2 cannot be started anymore:

$$x_{a_2 t} \le 1 - \sum_{\tau\le t} x_{a_1\tau}, \quad \forall t \in \mathcal{T}. \tag{12}$$

For the shutdown decision, we impose a constraint that the field has to be shut down at some point during the time horizon, i.e.:

$$y_T = 1, \tag{13}$$

and

$$y_t \ge y_{t-1}, \quad \forall t \in \mathcal{T} \setminus \{0\}, \tag{14}$$

where the last equation defines the irreversibility of the shutdown decision.

3.5. Objective

We aim to maximize the expected net present value (NPV) of the asset, which consists of the expected discounted net profits in each stage. Here, the profits equal the revenues due to the sale of the extracted hydrocarbons minus total costs (c_t):

$$\max_{\{y_t, x_{at}\}_{a\in\mathcal{A}_t, t\in\mathcal{T}}} \mathbb{E}\left[\sum_{t\in\mathcal{T}} D_t\left(\tilde{p}_t q_t - c_t\right)\right], \tag{15}$$

where D_t is the discount factor for time period t and the expectation is with respect to the exogenous price factors. The discount factor can also be expressed in terms of the risk-free interest rate (R): $D_t = \frac{1}{(1+R)^t}$. Total costs consist of operational expenditures (C_t^{OPEX}), costs of undertaking activities (C_{at}^{ACT}) and decommissioning costs (C_t^{DECOM}):

$$c_t = C_t^{OPEX}\left(1 - y_{t-1}\right) + \sum_{a\in\mathcal{A}_t} C_{at}^{ACT} x_{at} + C_t^{DECOM}\left(y_t - y_{t-1}\right), \quad \forall t \in \mathcal{T}. \tag{16}$$

4. SDDiP Algorithm

The SDDiP algorithm (Zou et al., 2019) is used to solve MSIPs and makes use of a stage-wise decomposition to prevent a combinatorial explosion of the number of states/scenarios. SDDiP is based on a dynamic programming formulation that includes *expected cost-to-go functions* for each time stage. The expected cost-to-go functions are not known in advance and have to be approximated. Given an approximated cost-to-go function and state, solving the corresponding stage optimization problem gives a decision for that stage. So implicitly the approximated cost-to-go function maps a given state of the model to a decision. We refer to such a mapping as a *policy*. The cost-to-go functions are approximated in an iterative fashion using the solutions of the optimization problem in each stage, which can be considered to be Benders cuts. Each iteration consists of a *forward pass* and a *backward pass*. In the forward pass a set of scenarios is sampled. The current *policy* is evaluated on each of these scenarios, leading to a set of policy values. Based on these realized policy values, a statistical bound on the objective function value can be constructed. This is an upper bound for a minimization problem, as the cost-to-go function is represented by an outer approximation. The backward pass then generates cuts that outer approximate the expected cost-to-go function. The backward pass consists of solving relaxed subproblems from the last to the first stage, where the solutions of future stages are used to generate cuts and approximations to the cost-to-go function. The relaxed solution for the first time stage gives a lower bound to the problem, as only a subset of scenarios is considered. This procedure is repeated until a convergence criterion has been reached. The performance of the obtained policy can be tested by means of a simulation, which finally gives an expected policy value.

Table 1
Kalman Filter parameter estimates based on ICE Brent Crude Futures.

	κ	σ_χ	λ_χ	σ_ξ	μ_ξ^*	$\rho_{\chi\xi}$
Estimate	0.407	0.273	−0.147	0.149	−0.007	0.306
Standard error	0.004	0.014	0.015	0.006	0.001	0.054

SDDiP explicitly distinguishes between two types of variables in each stage. That is, *state* variables linking successive stages, and *local* (or *stage*) variables. The SDDiP algorithm requires the optimization problem at each stage to be linear and the state variables to be binary. The latter is not a large restriction for real options problems, as the decisions (to wait, abandon or expand) are binary. Moreover, SDDiP restricts the stochastic process to a finite number of realizations (discrete process). If this is not the case, such a process can be approximated using methods such as sample average approximation, k−means or other constructive methods (Pflug, 2001). The original (possibly continuous) stochastic process, is sometimes referred to as the *true* distribution (Ding and Ahmed, 2019).

The STLT model presented above is a stagewise dependent process, satisfying the Markov property. In particular, the spot price in stage t is dependent on the realizations of the short-term and long-term factors in the previous stage. There are two approaches to incorporate stagewise dependence (Löhndorf and Shapiro, 2019; Downward et al., 2020) in the SDDiP algorithm. The first approach, referred to as the *Time Series* (TS) approach, incorporates the autoregressive model defined in (3)–(4) as constraints and adds the short- and long-term factors as state variables. The second approach, referred to as the *Markov Chain* (MC) approach, uses an MC approximation of the random process in Eqs. (3)–(4), capturing the correlation between both factors. Both Löhndorf and Shapiro (2019) and Ding and Ahmed (2019) find that the MC approach yields better convergence of the bounds and gives a superior optimal policy compared to the TS approach. We implemented both approaches and obtained similar results. Therefore, this work uses the MC approach. Nevertheless, Appendix C presents a short discussion about how to implement the TS approach, as well as a short computational study.

Finally, the SDDiP algorithm can be almost directly applied to the MSIP presented above. Extra care has to be taken when defining the state variables that link the different stage problems. In particular, it is important to keep track of past activities affecting current production profiles, so we have to add all activity variables $\{x_{at} : a \in \mathcal{A}_t, t \in \mathcal{T}\}$ as state variables in each stage.

5. Computational study

5.1. STLT estimation

We make use of future contracts for Brent Crude to estimate the parameters of the STLT model for oil prices. Brent Crude is the main trading classification of oil produced from the North Sea and is being traded on the Intercontinental Exchange (ICE). We retrieved historical data on these contracts from Thomson Reuters Eikon (2020), where we used data from January 2006 until December 2019. The maturity of these contracts has been on average around eight years, which gives a good match with the time horizon of our decision problem. In case of missing observations, we have made use of linear interpolation and extrapolation. Finally, we were able to estimate the parameters of the STLT model using the Kalman Filtering technique (Harvey, 1990) as implemented in Matlab by Goodwin (2020) for the STLT model. Parameter estimates and standard errors are given in Table 1. We only report parameter estimates for the risk-neutral process, which is the process we use for valuation.

Focusing on the standard errors, we can conclude that all the parameters are being estimated with relatively high accuracy. When comparing the parameter estimates with previously published estimates of the STLT model for oil prices (Schwartz and Smith, 2000; Jafarizadeh and Bratvold, 2012), we observe that we obtain a much lower short-term mean reversion rate (κ). This means that short-term deviations are more persistent. The estimated value of κ implies a half-life of deviations of 1.7 years. When comparing the volatility and correlation coefficients ($\sigma_\chi, \sigma_\xi, \rho_{\chi\xi}$), we find fairly similar estimates. In addition, we notice that we have a negative risk-neutral equilibrium drift rate. Finally, we obtain initial values for the price factors given by $\chi_0 = 0.534$ and $\xi_0 = 3.633$, implying an initial spot price (p_0) of 64US$/bbl.

5.2. Numerical example

In order to demonstrate the performance of the developed model, we have constructed a numerical example representing a mature offshore field on the Norwegian Continental Shelf. The parameter estimates are based on input from a Norwegian operator as well as a publicly available database (Norwegian Petroleum Directorate, 2020).

The field under consideration consists of a stand-alone production facility with 15 well slots. There is a single reservoir from which mainly oil is being produced. The yearly operational expenditures (OPEX) of this installation are 50 million dollars (US$MM). The time horizon of the problem is ten years, during which the installation has to be shut down at some point. The risk-free interest rate for projects with this time horizon is based on Norwegian 10-year government bonds and set to 2%. We are mainly interested in the decisions during the first number of years. Therefore, we set the first three stages to have an annual resolution, the next two to have a bi-annual resolution, while the last stage has a duration of three years. This leads to a total of six stages.

The production profiles used in the numerical example are based on actual realized production profiles (Norwegian Petroleum Directorate, 2020). In the first year, the base production level equals 0.81 million barrels (MMbbl), while in year ten, this equals just 0.28 MMbbl. We can increase production by undertaking three different activities that have been quantified. That is, we have two possibilities to perform slot recovery ($sidetrack1$ and $sidetrack2$), where the activities target different parts of the reservoir. In addition, four new wells can be drilled ($wells4$). The drilling of new wells leads to a higher increase in production compared to slot recovery. However, there is also a significant higher cost related to this activity. The first two slot recovery activities cost US$15MM and US$30MM respectively, while the four wells cost US$200MM.

Finally, we assume that the fifteen wells require a rig to perform plugging operations. With a day rate of US$0.5MM and an average of 25 days per well, this leads to a P&A cost of US$150MM. Removal of the platform has been estimated to cost US$200MM, which leads to a decommissioning cost of US$350MM.

5.3. Computational results

To analyze the numerical example, we implemented the model in Python 3.8, formulated it using the MSPPy package (Ding and Ahmed, 2019) and solved it with Gurobi 9.0.1. The analyses have been carried out on a Dell PowerEdge R640 computer with an Intel Xeon Gold 5115 CPU, 2.4 GHz processor, 96 Gb RAM, running CentOS Linux 7. We did not make use of any of the parallelization schemes available in the MSPPy package.

5.3.1. Markov Chain representation

The stochastic process in (3)–(4) defines a continuous Markov process, which we refer to as the true distribution. To solve the model using SDDiP, this needs to be represented by a discrete Markov Chain. The Markov states in each stage are represented using pairs $(\tilde{\chi}_t, \tilde{\xi}_t)$. In our analyses, we used the k-means approach to train the Markov Chains.

In each backward pass in the SDDiP algorithm we add both strengthened Benders cuts and Lagrangian cuts, which was found to give satisfactory performance. Every K iterations we simulate the obtained

Table 2
Computational results for different numbers of Markov states per stage, including number of iterations (# its.), solution time, deterministic upper bound and optimality gap. The presented optimality gap is based on a simulation of the discrete Markov Chain, while the expected policy value and confidence interval are based on a simulation of the true distribution.

Markov states	First stage decision	# Its	Time (s)	UB (MM US$)	Opt. gap	EPV^* (MM US$)	$CI^*_{95\%}$
1	shutdown	9	0.36	−358.19	0.00%	−358.19	(−358.19 , −358.19)
2	sidetrack2	12	0.94	−331.55	0.93%	−326.95	(−329.87 , −324.03)
5	sidetrack2	12	6.35	−320.40	0.61%	−316.70	(−319.33 , −314.07)
10	sidetrack2	12	24.36	−318.11	1.07%	−316.54	(−319.14 , −313.94)
15	sidetrack2	12	39.84	−317.56	0.99%	−316.38	(−319.00 , −313.76)
20	sidetrack2	12	56.22	−317.19	1.96%	−316.12	(−318.75 , −313.48)
25	sidetrack2	12	70.67	−317.33	0.51%	−316.10	(−318.73 , −313.47)
30	sidetrack2	12	127.02	−317.25	0.37%	−316.06	(−318.70 , −313.43)

policy N^{SIM} times, to evaluate the quality of the policy. In this way, we can construct a $(1-\alpha)$-confidence interval on the Expected Policy Value (EPV), which gives us a statistical lower bound for the problem. The optimality gap can then be constructed based on the deterministic upper bound (UB) obtained from the forward passes and the statistical lower bound from the simulations. In the same way, we can construct a $(1-\alpha)$-confidence interval for the *improvement* in the EPV. We use this as a stabilization stopping criterion for the algorithm. When the upper bound of this EPV-improvement interval does not exceed a specified tolerance level ϵ, we assume that the policy has stabilized and the algorithm is terminated. Table 2 presents computational results for different representations of the Markov Chain in terms of the number of Markov states per stage.

The obtained policy is valid for the Markov states of the discrete Markov Chain. When extending this policy to the true distribution, we match realizations of the true distribution with Markov states that are closest in terms of distance and use the corresponding approximation of the cost-to-go function (Löhndorf and Shapiro, 2019). Using this approach, we can construct a $(1-\alpha)$-confidence interval for the expected policy value (EPV*), denoted by $CI^*_{1-\alpha}$, where the asterisk refers to the evaluation on the true distribution. Simulating the obtained policy is computationally expensive, so it is not warranted to do this each iteration or to use a very large N^{SIM}. However, choosing a large value of K in turn increases the number of iterations and hence the solution time. In the sequel, we have used $K=3$, $N^{SIM}=5000$, $\alpha = 5\%$ and $\epsilon = 0.01\%$, which provided a good balance towards computational performance.

First, we note that considering only one Markov state is equivalent to solving a deterministic problem which includes the expected spot price as single scenario. This is the only case that is solved to optimality. For the other cases, the stabilization stopping criterion is reached at a point where there still is a relatively small optimality gap remaining. Fig. 1 visualizes the convergence of the SDDiP algorithm for the case with 10 Markov states. For visualization purposes we simulated the obtained policy each iteration (i.e. $K=1$), which made the algorithm terminate after 10 iterations. We observe that the confidence interval for the EPV is relatively small, as we used a relatively large value for N^{SIM}. The statistical lower bound decreases in the fifth iteration, which might be due to the algorithm reaching an unexplored part of the feasible region. The algorithm terminates when there is no more improvement in the performance of the obtained policy, but there is still a marginal optimality gap.

Turning back to Table 2, we observe that the models solve in mere seconds. When using 30 Markov states per stage, the problem still solves in around 2 min. Also, the solution times scale approximately linearly with the number of states. Moreover, we see that there is little improvement in the EPV^* when adding more than around 10 Markov states. This implies that for the computational study relatively few Markov states are needed to represent the uncertain price process.

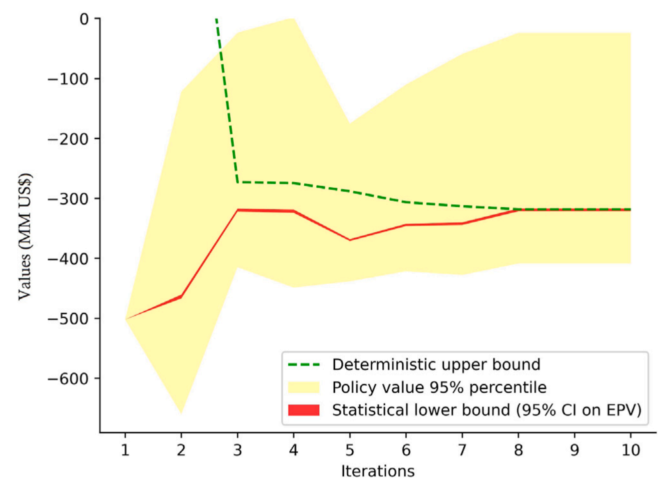

Fig. 1. Convergence of the SDDiP algorithm in terms of development of the deterministic upper bound and statistical lower bound for 10 Markov states.

5.3.2. Larger instances

The numerical example includes three activities, which is a typical size of a real options or investment timing problem. Nevertheless, the example can be considered to be relatively small. To test the performance of the methodology on a larger instance, we extended the numerical example to include a total of twelve activities and used 10 Markov states per state. The policy stabilized after 18 iterations and 55 s, with an optimality gap of 1.67%. Even though the computational results indicate that larger instances can be solved, such instances are considered to be less realistic for the problem under consideration as operators only quantify the most promising activities on a field. In addition, it is challenging to obtain data estimates for such large instances (in particular production profiles).

5.3.3. Value of the stochastic solution

As mentioned before, the expected value problem is represented by the use of a single Markov state, where the uncertainties are replaced by their mean. Common practice for oil companies is to perform net present value calculations based on expected price realizations and/or on percentiles of price distributions such as the 10% percentile (P10, low prices) and 90% percentile (P90, high prices). While this approach aims to account for uncertainty in the price process, its actual use in decision making is limited. In our numerical example, both the expected value problem and the low price scenario (P10) suggest shutting down the field immediately. However, if a high price scenario (P90) occurs, it would be optimal to invest in 4 new wells. So, the solutions from these deterministic models differ, giving inconsistent information to the decision maker. Moreover, this deterministic approach fails to suggest the optimal first stage solution, which is to perform two sidetracks.

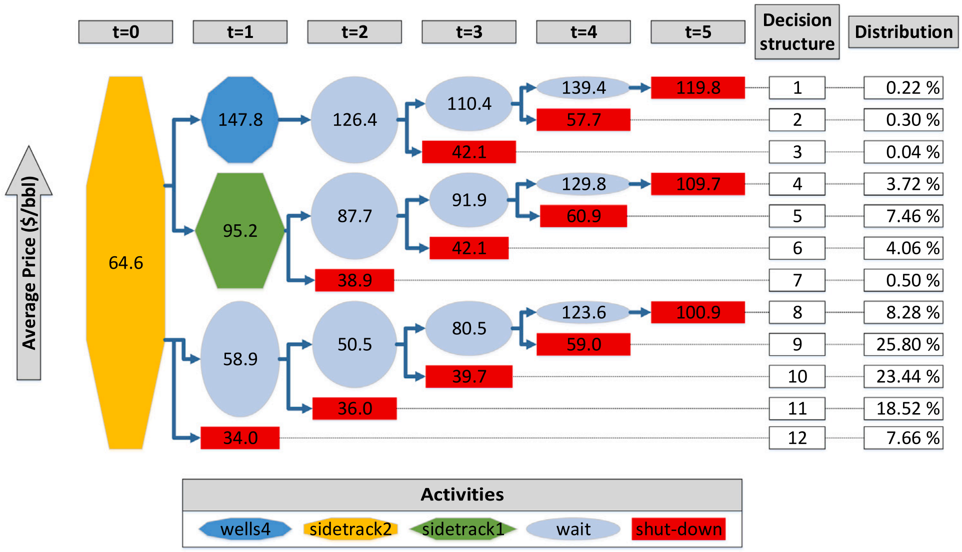

Fig. 2. Structure of the decisions (state variables) when simulating the optimal policy. We present the average spot price (in US$/bbl) in each stage for each pathway.

This decision is chosen to obtain more information about the uncertain process in subsequent stages. For a high price realization, the option to expand can still be exercised, while for low price realizations the installation can still be shut down.

The expected benefit of using a representation of the prevailing uncertainty instead of the expected value problem, is usually referred to as the Value of the Stochastic Solution (VSS). In our numerical example, the expected value problem results in a cost of US$358.2MM, while the optimal policy when considering 10 Markov states leads to an expected cost of US$318.1MM. This implies an expected VSS of US$40.1MM, which reflects a 11% decrease in costs.

5.3.4. Decision structure

From Table 2 we have seen that the optimal first stage decision is to perform two sidetracks. In addition, when solving the model we obtain a policy, represented by a family of cuts that estimate the cost-to-go function. We can use this policy in a Monte-Carlo study to find out which decisions are being made for different scenarios at the various stages. Fig. 2 visualizes the structures of the decisions that appear in a Monte-Carlo study of the optimal policy using 5000 simulations. A pathway represents a unique decision structure, while a box/node represents a decision taken in a certain stage. The last column presents the distribution of occurrences of the different decision structures. Within each node, we present the average spot price for the pathways that go through this particular node.

The diagram illustrates the richness in decision structures for the fairly simple numerical example presented. When spot prices increase drastically, we see that it can be optimal to drill four extra wells, while for more moderate price changes, the optimal decision tends to be to wait or to drill a new sidetrack. For low prices, we observe that the field tends to be shut down before the time horizon of the problem. In only around 12% of the scenarios, the field is shut down in the last stage.

5.3.5. Factors influencing the average shutdown year

This section investigates the effect of different factors on the average shutdown year. We solve the model for different parametrizations and simulate the obtained policy. The average shutdown year is obtained by translating the average shutdown stage into its duration in years. Note that the figures in this section display some non-monotonic behavior. As the model parameters are changed (the shutdown cost, discount rate, OPEX and price volatility), the course of the SDDiP algorithm is affected. As a result, the policies will evolve in different ways, even when using the same seeds in these analyses.

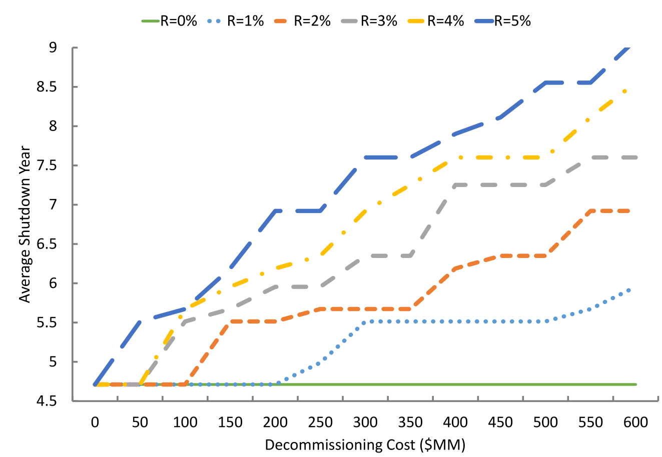

Fig. 3. Average year in which the field is shut down, as a function of decommissioning costs and discount rate.

Using data from a large set of oil and gas wells located in Canada, Muehlenbachs (2015) uses a dynamic programming real options model to show that it is common practice for oil and gas companies to temporarily shut down wells to prevent the high costs that accompany permanent P&A. By postponing the permanent shutdown decision, permanent P&A costs are shifted forward and can be discounted by a high discount factor. As a result, we are interested in the effect that the shutdown cost and discount rate have on postponing the shutdown decision in our problem setting.

Fig. 3 shows the average shutdown year, while varying the decommissioning costs and discount rate. With no monetary time value (i.e. $R = 0$), the field is expected to be shut down after approximately 4.7 years. In this case, the decommissioning costs have no effect on the shutdown decision. However, the main observation is that the average shutdown year increases both as a function of the decommissioning costs as well as the interest rate. In the base-case with C^{DECOM} = US$350MM and $R = 2\%$, the field is expected to be shut down after

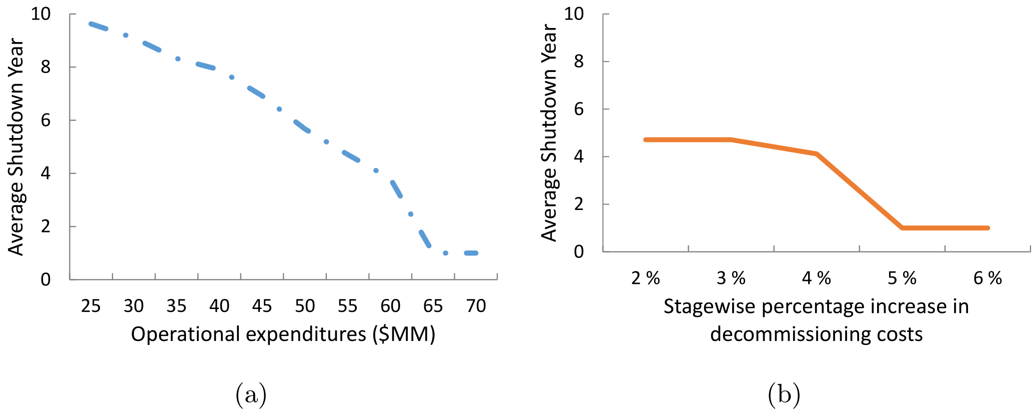

Fig. 4. Average year in which the field is shut down. Both as a function of different realizations of the operational expenditures 4(a) as well as the percentage increase in decommissioning costs per stage 4(b).

5.67 years. For moderate increases in the parameters, this can be extended by several years. Nevertheless, for these parametrizations, permanent abandonment is not extended all the way to the horizon (10 years). The main reason being that there is a significant OPEX relative to decommissioning costs. That is, when the revenue from the saleable production and the savings from postponing the shutdown decision due to discounting do not cover this OPEX, the field will be shut down.

The effect of operational expenditures on the average shutdown year is illustrated in Fig. 4. With an OPEX of US$25MM, the shutdown decision is pushed all the way back to the final year in the time horizon. If it is legally possible to temporarily plug all wells connected to an installation, the operational expenditures of the installation could be reduced significantly, giving an economic incentive to postpone the costly P&A operations.

While high decommissioning costs, high interest rates and low operational expenditures extend the shutdown decision, increasing decommissioning costs over time might have the opposite effect. As we focus on mature fields, we are facing ageing infrastructure and wells that have an increased risk of integrity problems (Vignes and Aadnøy, 2010). Examples of such problems are corrosion on the platform or deformation of the casing in the borehole. It requires more time to plug wells with integrity problems and this may lead to higher decommissioning costs. Such a cost increase cannot be easily quantified, and we are not aware of any works describing this. To evaluate the effect, we define a percentage stagewise cost increase in decommissioning costs. This is based on duration estimates of plugging operations for different well complexities given by Øia et al. (2018), which can be seen as a proxy for well integrity. The use of a 6% cost increase per stage implies that the decommissioning cost will be 1.34 times higher in stage 5 (years 8 through 10) than in stage 0, which is in accordance with the data presented by Øia et al. (2018). Fig. 4 visualizes the effect of a percentage stagewise cost increase in decommissioning costs on the average shutdown year. As expected, we observe that with increasing decommissioning costs, the shutdown decision is expedited. When using a 5% cost increase per stage or higher, this would lead to immediate shutdown in stage 0.

Finally, an important topic in the real options literature is the relationship between uncertainty and investment timing. The effect of increasing volatility on the price level at which an investment is triggered or the expected time until the investment can be analyzed. Traditionally, it is found that uncertainty delays the investment. For example, Clarke and Reed (1990) consider the option to abandon an oil well, assuming that the uncertain process follows a GBM. In this case, the authors show that the oil price level triggering shutdown decreases as the volatility of the GBM increases. However, as these calculations quickly become complex, such comparative statistics are usually only performed for single options. When considering multiple options in a portfolio, the effect of an increase in volatility is not so easily determined. We therefore perform a sensitivity analysis on the volatility of the short-term deviations (σ_χ) and the long-term equilibrium level (σ_ξ), respectively. In Fig. 5 we present the average shutdown year for varying volatility, where the base cases with $\sigma_\chi = 0.273$ and $\sigma_\xi = 0.149$ result in an average shutdown year of 5.67. Even though we observe some non-monotonic behavior, the figure confirms the traditional relationship found in real options literature that uncertainty delays the shutdown decision (Dixit and Pindyck, 1994). This holds for an increase in short-term as well as long-term volatility.

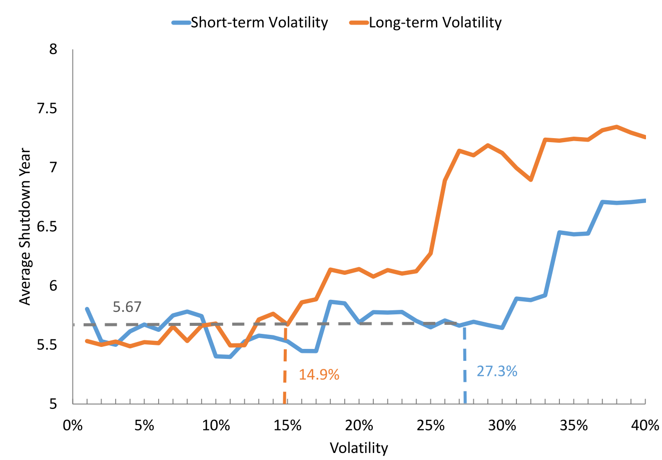

Fig. 5. Average year in which the field is shut down, as a function of the volatility of the short-term factor (σ_χ) and the long-term factor (σ_ξ). For each curve, the other volatility is fixed at its base level.

6. Conclusions

The problem of developing a mature offshore oil field can be considered to be a real options problem consisting of a portfolio of interdependent options. In this paper, we showed that this problem can be formulated as a multistage stochastic integer linear program and efficiently solved with the SDDiP algorithm. The oil price was considered to be the main uncertain factor and is represented by the STLT model of Schwartz and Smith (2000) to account for oil price dynamics.

To analyze the problem, we considered a numerical example based on realistic input from a Norwegian operator. We solved the problem using the SDDiP algorithm and Markov chain approach. We found that the algorithm tends to converge quickly to a relatively small optimality gap. We also investigated the number of Markov states that are needed to represent the true distribution satisfactorily, based on the expected policy value. We observed that the use of around 10 Markov states is sufficient to represent the uncertain price process. The model solves in mere minutes and scales linearly with the number of states. For the numerical example, we found that the value of the stochastic solution lies around 11%, which means that there is significant value in using this model rather than deterministic alternatives.

The SDDiP framework allows for simulating obtained policies, giving the opportunity to perform different kinds of analyses. We evaluated solution structures and found that the decisions in each stage are highly dependent on the realizations of the price process. Moreover, we investigated factors that influence the average shutdown year of the field by means of a sensitivity analysis. Our conjectures were confirmed in that high decommissioning costs, high discount rates and low operational expenditure tend to postpone the shutdown decision. Nevertheless, these effects might be negated if there are increasing well integrity problems, represented by decommissioning costs that increase over time, leading to acceleration of the shutdown decision. In line with traditional real options theory, we found that an increase in the volatility of the short-term deviations and long-term equilibrium level postpones the shutdown decision.

To conclude, we demonstrated that industry-sized problems involving infrastructure development under uncertainty can be solved by framing the problem as an MSIP and solving it using SDDiP. While the computational results showed that the presented approach is effective when only considering price as stochastic, it remains to be seen how the approach performs when considering multiple uncertain elements, e.g. cost and production profile. Future work might also look at applying the presented methodology on other cases that give rise to larger instances with different configurations. This can be in terms of, for example, the definition of the profits and costs, the number of time periods or a larger set of interdependencies between options. Finally, a limitation of the SDDiP method is that it requires binary state variables, which in the context of a pure real options problem tends to be satisfied. When, however, both continuous and binary state variables are required, other solution methods are needed, such as binarization of continuous state variables (Zou et al., 2019). Future research might focus on the development and performance of such methods.

CRediT authorship contribution statement

Steffen J. Bakker: Conceptualization, Methodology, Software, Formal analysis, Investigation, Writing – original draft, Writing – review & editing, Visualization. **Andreas Kleiven:** Conceptualization, Writing - review & editing. **Stein-Erik Fleten:** Conceptualization, Writing – review & editing, Funding acquisition. **Asgeir Tomasgard:** Writing – review & editing, Funding acquisition.

Acknowledgments

The authors are grateful to the anonymous referees, whose comprehensive reports helped to significantly improve the structure and presentation of this paper.

This paper was prepared as a part of the project "Economic Analysis of Coordinated Plug and Abandonment Operation" (ECOPA), funded by the Research Council of Norway through the PETROSAM2 and PETROMAKS2 programs (p-nr: 247589).

Appendix A. Abbreviations

ADP	Approximate Dynamic Programming
bbl	Barrel
EOR	Enhanced Oil Recovery
EPV	Expected Policy Value
GBM	Geometric Brownian Motion
LSM	Least Squares Monte Carlo
MAE	Mean Absolute Error
MC	Markov Chain
MM	Million
MSIP	Multistage Stochastic Integer Programming
NPV	Net Present Value
OPEX	Operational Expenditures
OU	Ornstein Uhlenbeck
P&A	Plug and Abandonment
SDDiP	Stochastic Dual Dynamic (Integer) Programming
SDDP	Stochastic Dual Dynamic Programming
STLT	Short-term Long-term
TS	Time Series
VSS	Value of Stochastic Solution

Appendix B. Nomenclature

B.1. Sets and indices

$\mathcal{T}$	=	Set of time periods or stages, indexed by t and τ.
$\mathcal{A}$	=	Set of activities (sidetracks, new wells, EOR/IOR) that can be undertaken, indexed by a.
$\mathcal{A}_t$	=	Subset of activities that can be undertaken at time period t.
Γ_t	=	Set of pairs of activities and stages (a, τ) that could have been undertaken before time period t.

B.2. Parameters

D_t	=	Discount factor in time period t.
R	=	Risk neutral interest rate.
Q_t^{BASE}	=	Base production level in time period t.
$Q_{a\tau t}^{ADD}$	=	Additional production in time period t, when activity a has been started in time period τ.
M_t^{PROD}	=	Theoretically maximum production in time period t.
Q^L, Q^U	=	Lower and upper bound on the production capacity.
L	=	Constant representing the maximum number of activities that can be performed in a certain stage.
C_{at}^{ACT}	=	Cost of undertaking activity a in time period t.
C_t^{OPEX}	=	Operational expenditure if the installation is producing in time period t.
C_t^{DECOM}	=	Decommissioning costs in time period t.
$\tilde{p}_t$	=	Spot price for oil in time period t.
K	=	Number of iterations between policy evaluations (using simulation).
N^{SIM}	=	Number of simulations to evaluate a policy.

B.3. Variables

B.3.1. Binary state variables

y_t	=	Shutdown variable, equaling one if the decision to shut down the installation is made in period t or any of the preceding time periods, and zero otherwise.
x_{at}	=	Activity variable, equaling one if the decision to start activity a is taken in time-period t, and zero otherwise.

B.3.2. Control variables

q_t = Total production when the installation is in operation, equals zero otherwise.

c_t = Costs in time period t, consisting of OPEX, abandonment- and activity costs.

Appendix C. Time series approach

To solve the MSIP presented in Section 3 using a Time Series approach, several steps have to be undertaken:

1. Eqs. (2)–(4) have to be included in the model and $\tilde{\chi}_t$ and $\tilde{\xi}_t$ have to be included as state variables. The error terms in these equations are now stagewise independent.
2. Eq. (2), defining the price as a function of the two factors, includes the non-linear exponential function. This function has to be approximated by a piecewise linear function.
3. The price is now modeled as a variable, instead of an uncertain parameter. This leads to a product of the price and production variable. But, since the production variable is based on binary variables (the activity variables), this can be rewritten in an exact manner.

A more formal description of this procedure can be obtained from the authors.

References

Babadagli, T., 2007. Development of mature oil fields - a review. J. Pet. Sci. Eng. 57 (3–4), 221–246. http://dx.doi.org/10.1016/j.petrol.2006.10.006.

Bakker, S., Vrålstad, T., Tomasgard, A., 2019. An optimization model for the planning of offshore plug and abandonment campaigns. J. Pet. Sci. Eng. 180, 369–379. http://dx.doi.org/10.1016/j.petrol.2019.05.042.

Birge, J., 1985. Decomposition and partitioning methods for multistage stochastic linear programs. Oper. Res. 33 (5), 989–1007. http://dx.doi.org/10.1287/opre.33.5.989.

Bonini, C.P., 1977. Capital investment under uncertainty with abandonment options. The J. Financial Quant. Anal. 12 (1), 39–54. http://dx.doi.org/10.2307/2330286.

Brown, D.B., Smith, J.E., Sun, P., 2010. Information relaxations and duality in stochastic dynamic programs. Oper. Res. 58 (4-part-1), 785–801. http://dx.doi.org/10.1287/opre.1090.0796.

Clarke, R., Reed, J., 1990. Oil-well valuation and abandonment with price and extraction rate uncertain. Resour. Energy 12, 361–382. http://dx.doi.org/10.1016/0165-0572(90)90029-I.

Dias, M.A.G., 2004. Valuation of exploration and production assets: An overview of real options models. J. Pet. Sci. Eng. 44 (1–2), 93–114. http://dx.doi.org/10.1016/j.petrol.2004.02.008.

Ding, L., Ahmed, S., 2019. A python package for multi-stage stochastic programming. Georgia Institute of Technology, http://www.optimization-online.org/DB_HTML/2019/05/7199.html.

Dixit, A.K., Pindyck, R.S., 1994. Investment under Uncertainty. Princeton University Press, Princeton, New Jersey.

Downward, A., Dowson, O., Baucke, R., 2020. Stochastic dual dynamic programming with stagewise dependent objective uncertainty. Oper. Res. Lett. 48 (1), 33–39. http://dx.doi.org/10.1016/j.orl.2019.11.002.

Dowson, O., Kapelevich, L., 2021. SDDP.jl : A julia package for stochastic dual dynamic programming. INFORMS J. Comput. 33 (1), 27–33. http://dx.doi.org/10.1287/ijoc.2020.0987.

Ekern, S., 1988. An option pricing approach to evaluating petroleum projects. Energy Econ. 10 (2), 91–99. http://dx.doi.org/10.1016/0140-9883(88)90023-0.

Fleten, S.-E., Gunnerud, V., Hem, O.D., Svendsen, A., 2011. Real option valuation of offshore petroleum field tie-ins. J. Real Options 1, 1–17.

Goel, V., Grossmann, I.E., 2004. A stochastic programming approach to planning of offshore gas field developments under uncertainty in reserves. Comput. Chem. Eng. 28 (8), 1409–1429. http://dx.doi.org/10.1016/j.compchemeng.2003.10.005.

Goodwin, D., 2020. Schwartz-smith 2-factor model - parameter estimation. https://www.mathworks.com/matlabcentral/fileexchange/43352-schwartz-smith-2-factor-model-parameter-estimation.

Gupta, V., Grossmann, I.E., 2012. An efficient multiperiod MINLP model for optimal planning of offshore oil and gas field infrastructure. Ind. Eng. Chem. Res. 51 (19), 6823–6840. http://dx.doi.org/10.1021/ie202959w.

Gupta, V., Grossmann, I.E., 2014. Multistage stochastic programming approach for offshore oilfield infrastructure planning under production sharing agreements and endogenous uncertainties. J. Pet. Sci. Eng. 124, 180–197. http://dx.doi.org/10.1016/j.petrol.2014.10.006.

Harvey, A.C., 1990. Forecasting, Structural Time Series Models and the Kalman Filter. Cambridge University Press., Cambridge, http://dx.doi.org/10.1017/CBO9781107049994.

Hjelmeland, M.N., Zou, J., Helseth, A., Ahmed, S., 2019. Nonconvex medium-term hydropower scheduling by stochastic dual dynamic integer programming. IEEE Trans. Sustain. Energy 10 (1), 481–490. http://dx.doi.org/10.1109/TSTE.2018.2805164.

Iyer, R., Grossmann, I.E., 1998. Optimal planning and scheduling of offshore oil field infrastructure investment and operations. Ind. Eng. Chem. Res. 5885 (97), 1380–1397. http://dx.doi.org/10.1021/ie970532x.

Jafarizadeh, B., Bratvold, R.B., 2012. Two-factor oil-price model and real option valuation: An example of oilfield abandonment. SPE Economics Manag. 4 (3), 158–170. http://dx.doi.org/10.2118/162862-PA.

Jonsbråten, T.W., 1998. Oil field optimization under price uncertainty. J. Oper. Res. Soc. 49 (8), 811–818. http://dx.doi.org/10.1057/palgrave.jors.2600562.

Khor, C.S., Elkamel, A., Shah, N., 2017. Optimization methods for petroleum fields development and production systems: a review. Opt. Eng. 18 (4), 907–941. http://dx.doi.org/10.1007/s11081-017-9365-2.

Lara, C.L., Siirola, J.D., Grossmann, I.E., 2019. Electric power infrastructure planning under uncertainty: Stochastic dual dynamic integer programming (SDDiP) and parallelization scheme. Opt. Eng. In Press, http://dx.doi.org/10.1007/s11081-019-09471-0.

Lin, X., Floudas, C., 2003. A novel continuous-time modeling and optimization framework for well platform planning problems. Opt. Eng. 4 (1/2), 65–95. http://dx.doi.org/10.1023/A:1021860330145.

Löhndorf, N., Shapiro, A., 2019. Modeling time-dependent randomness in stochastic dual dynamic programming. European J. Oper. Res. 273 (2), 650–661. http://dx.doi.org/10.1016/j.ejor.2018.08.001.

Longstaff, F.A., Schwartz, E.S., 2001. Valuing American options by simulation: A simple least-squares approach. Rev. Financ. Stud. 14 (1), 113–147. http://dx.doi.org/10.1093/rfs/14.1.113.

Lund, M.W., 2000. Valuing flexibility in offshore petroleum projects. Ann. Oper. Res. 99 (1–4), 325–349. http://dx.doi.org/10.1023/A:1019284119505.

Maier, S., Pflug, G.C., Polak, J.W., 2019. Valuing portfolios of interdependent real options under exogenous and endogenous uncertainties. European J. Oper. Res. In press, http://dx.doi.org/10.1016/j.ejor.2019.01.055.

Maier, S., Polak, J.W., Gann, D.M., 2020. Valuing portfolios of interdependent real options using influence diagrams and simulation-and-regression: A multi-stage stochastic integer programming approach. Comput. Oper. Res. 115, 104505. http://dx.doi.org/10.1016/j.cor.2018.06.017.

McDonald, R.L., Siegel, D.R., 1985. Investment and the valuation of firms when there is an option to shut down. Internat. Econom. Rev. 26 (2), 331–349. http://dx.doi.org/10.2307/2526587.

Muehlenbachs, L., 2015. A dynamic model of cleanup: estimating sunk costs in oil and gas production. Internat. Econom. Rev. 56 (1), 155–185. http://dx.doi.org/10.1111/iere.12098.

Myers, S., Majd, S., 1983. Calculating abandonment value using option pricing theory. In: Sloan School of Management. 1462–83, Sloan School of Management, pp. 35–38, http://hdl.handle.net/1721.1/2056.

Norwegian Petroleum Directorate, 2019. Annual status report 2019 for fields in production (ASR2019). https://www.npd.no/en/regulations/reporting_and_applications/production/annual-status-report-2019-for-fields-in-production-asr2019/.

Norwegian Petroleum Directorate, 2020. Factpages [yearly field production]. https://factpages.npd.no/Default.aspx?culture=en&nav1=field&nav2=TableView%7CProduction%7CSaleable%7CYearly.

Nygreen, B., Christiansen, M., Haugen, K., Bjørkvoll, T., Kristiansen, O., 1998. Modeling Norwegian petroleum production and transportation. Ann. Oper. Res. 82, 251–268. http://dx.doi.org/10.1023/A:1018962703587.

O'Brien, J., Sayavedra, L., Mogollon, J.L., Lokhandwala, T., Lakani, R., 2016. Maximizing mature field production - a novel approach to screening mature fields revitalization options. In: SPE Europec Featured At 78th EAGE Conference and Exhibition. Vienna, Austria, http://dx.doi.org/10.2118/180090-ms.

Øia, T.M., Aarlott, M.M., Vrålstad, T., 2018. Innovative approaches for full subsea P&A create new opportunities and cost benefits. In: SPE Norway One Day Seminar. Bergen, Norway, http://dx.doi.org/10.2118/191315-MS.

Olsen, T.E., Stensland, G., 1988. Optimal shutdown decisions in resource extraction. Econom. Lett. 26, 215–218. http://dx.doi.org/10.1016/0165-1765(88)90137-1.

Paddock, J.L., Siegel, D.R., Smith, J.L., 1988. Option valuation of claims on real assets: The case of offshore petroleum leases. Q. J. Econ. 103 (3), 479–508. http://dx.doi.org/10.2307/1885541.

Pereira, M.V., Pinto, L.M., 1991. Multi-stage stochastic optimization applied to energy planning. Math. Program. 52 (1–3), 359–375. http://dx.doi.org/10.1007/BF01582895.

Petroleumsloven, 1996. Lov 21. november 1996 nr. 72 om petroleumsvirksomhet. https://lovdata.no/dokument/NL/lov/1996-11-29-72.

Pflug, G.C., 2001. Scenario tree generation for multiperiod financial optimization by optimal discretization. Math. Program. 89, 251–271. http://dx.doi.org/10.1007/PL00011398.

Powell, W.B., 2007. Approximate Dynamic Programming: Solving the Curses of Dimensionality (Wiley Series in Probability and Statistics), second ed. Wiley-Interscience, USA, http://dx.doi.org/10.1002/9780470182963.

Schwartz, E., Smith, J.E., 2000. Short-term variations and long-term dynamics in commodity prices. Manage. Sci. 46 (7), 893–911. http://dx.doi.org/10.1287/mnsc.46.7.893.12034.

Smith, J.E., Mccardle, K.F., 1998. Valuing oil properties: integrating option pricing and decision analysis approaches. Oper. Res. 46 (2), 198–217. http://dx.doi.org/10.1287/opre.46.2.198.

Smith, J.E., Mccardle, K.F., 1999. Options in the real world: Lessons learned in evaluating oil and gas investments. Oper. Res. 47 (1), 1–15. http://dx.doi.org/10.1287/opre.47.1.1.

Støre, K., Fleten, S.E., Hagspiel, V., Nunes, C., 2018. Switching from oil to gas production in a depleting field. European J. Oper. Res. 271 (2), 710–719. http://dx.doi.org/10.1016/j.ejor.2018.05.043.

Sullivan, J., 1988. The application of mathematical programming methods to oil and gas field development planning. Math. Program. 42, 189–200. http://dx.doi.org/10.1007/BF01589403.

Tarhan, B., Grossmann, I.E., Goel, V., 2009. A multistage stochastic programming approach for the planning of offshore oil or gas field infrastructure under decision dependent uncertainty. Ind. Eng. Chem. Res. 48, 3078–3097. http://dx.doi.org/10.1021/ie8013549.

Tavallali, M.S., Karimi, I.A., Baxendale, D., 2016. Process systems engineering perspective on the planning and development of oil fields. AICHE J. 62 (8), 2586–2604. http://dx.doi.org/10.1002/aic.15209.

Thomson Reuters Eikon, 2020. [Ice brent crude futures]. https://eikon.thomsonreuters.com/.

Vignes, B., Aadnøy, B.S., 2010. Well-integrity issues offshore Norway. SPE Product. Operations 25 (2), 145–150. http://dx.doi.org/10.2118/112535-PA.

Vrålstad, T., Saasen, A., Fjær, E., Øia, T., Ytrehus, J.D., Khalifeh, M., 2019. Plug & abandonment of offshore wells: ensuring long-term well integrity and cost-efficiency. J. Pet. Sci. Eng. 173, 478–491. http://dx.doi.org/10.1016/j.petrol.2018.10.049.

Wallace, S.W., 2010. Stochastic programming and the option of doing it differently. Ann. Oper. Res. 3–8. http://dx.doi.org/10.1007/s10479-009-0600-x.

Wang, T., de Neufville, R., 2004. Building real options into physical systems with stochastic mixed-integer programming. In: 8th Real Options Annual International Conference. Montreal, pp. 1–35, http://hdl.handle.net/1721.1/102772.

Zou, J., Ahmed, S., Sun, X.A., 2019. Stochastic dual dynamic integer programming. Math. Program. 175 (1–2), 461–502. http://dx.doi.org/10.1007/s10107-018-1249-5.