A general framework for modelling decision-dependent information revelation in stochastic programming

Rishi Adiga, Professor Andy Philpott
Department of Engineering Science
University of Auckland
New Zealand
radi130@auckland.ac.nz

Abstract

The stochastic programming framework for optimising decision-making assumes a predetermined information structure, where certain information is inevitably revealed at fixed time stages regardless of the decisions made. This assumption is not valid for problems with endogenous uncertainty, in which the sequence of information revelation is not fixed beforehand and is instead governed by model decisions. In this paper we present a comprehensive framework for defining information structures for such programs, based on modelling decision-dependent information discovery in an exploration-exploitation paradigm. This enables modellers to keep track of information flows and ensure that the desired causality between model decisions and information revelation is modelled appropriately. This framework is applied on a SMIP optimising well placement for a geothermal reservoir, and results demonstrating the effect of the information structure are discussed.

Key words: stochastic programming, endogenous uncertainty, decision-dependent information revelation, exploration-exploitation.

1 Introduction

Uncertainty in stochastic programming (SP) problems is modelled by how the scenarios relate to each other (the information structure), and the assumed probability distribution over the scenarios (our belief state). Typically the uncertainty is treated as being static, in that both the information structure (the sequence of information revelation about the random variables) and the belief state are assumed to be fixed, regardless of the decisions we make. In other words, we do not anticipate what the outcomes of the random variables will be right from the start but we do anticipate when they will be realized, and we assume the belief state remains constant over all stages. In this setting the uncertainty in a problem is *exogenous* - it applies the same regardless of the decisions we take and is not affected by them. By contrast *endogenous* uncertainty is that which is affected by the decisions we are optimising.

This terminology was first introduced in the SP literature by (Jonsbråten 1998), who worked on optimising petroleum field exploitation under reservoir uncertainty. Endogenous uncertainty is generally classified into two types in the literature: that where model decisions affect the probability distribution of the uncertain parameter outcomes or scenarios (type 1), and that where they affect the realization of the random variable outcomes (type 2). With type 1 endogenous uncertainty some decisions can alter our belief state, and with type 2 they determine if and when the random variables are realized. This distinction was made by (Goel and Grossmann 2006) who first proposed an explicit framework and solution method aimed at tackling type 2 endogenous uncertainty.

These authors and their colleagues have also put out a series of papers modelling this kind of uncertainty for offshore oil and gas field drilling problems. Another related application that was modelled in this manner is open-cast pit mining (Boland, Dumitrescu, and Froyland 2008). We will provide a detailed and up to date discussion of the literature on endogenous uncertainty in stochastic optimisation and related fields in a future paper; for now we refer interested readers to (Apap and Grossmann 2017) for a review of previous work on the subject. However the examples mentioned serve to show that it is often preferrable to model endogenous uncertainty in drilling and excavation applications. This is because geological or reservoir uncertainty is only resolved once the subsurface is actually investigated. If the decisions regulating this are the ones being optimised, then the information is only revealed as a result of those model decisions, and not otherwise. This is what motivates our work as well.

Previously we have developed a reduced order method for efficiently forecasting production outcomes for different combinations of wells from a small number of geothermal reservoir simulations (Adiga, O'Sullivan, and Philpott 2019). We showed that the forecasts generated using this method can be converted into NPV contributions and used in MIP models for optimising the selection of geothermal production wells (Adiga, O'Sullivan, and Philpott 2018). When the reduced order method is applied with an ensemble of different models (calibrations) of the same reservoir, the forecasts can be used in a SMIP hedging over different scenarios representing the calibrations. The SMIP models presented in (Adiga, O'Sullivan, and Philpott 2018) assumed a fixed information structure and treated the reservoir uncertainty as exogenous. In this paper we focus on modelling type 2 endogenous uncertainty in the SP framework, and develop a general template formulation for doing this from an exploration-exploitation perspective.

We first build up the key ingredients of this general formulation using an example Newsvendor problem. We then apply this for the geothermal well placement problem presented in (Adiga, O'Sullivan, and Philpott 2018), comparing the model with the original one, and demonstrate the effect of modelling information revelation in this way as a proof of concept. We do not go into detail on the specific problem data and results for the examples we present in terms of comparing solution metrics or policies. Rather the emphasis is on the modelling aspect, showing how this impacts the information structure of a problem and why it is important. Our proposed framework allows modellers to keep track of information flows and make sure that the appropriate causality between model decisions and information revelation is enforced. This can be difficult for large and complicated problems. Previous treatments of such problems tend not to elaborate on the details of modelling the problem and ensuring the information structure is correct, instead focusing on eliminating non-anticipativity constraints (NACs) for model reduction and decomposition-based solution methods.

2 Motivating Application

We now discuss the well placement problem from (Adiga, O'Sullivan, and Philpott 2018). We do not go into detail on the formulation, focusing just on the information structure. An ensemble of reservoir models was developed for the Kerinci geothermal system in Indonesia, and NPV forecasts were calculated for a set of candidate production wells from all of them, by applying the reduced order method from (Adiga, O'Sullivan, and Philpott 2019) on each calibration. The reservoir is bounded by four intersecting faults: two running vertically down the western and eastern ends, and two going across the northern and southern ends. The calibrations differ in the locations of deep upflows, which are boundary conditions defining the upward injection of hot fluid into the reservoir from below, but have the same physical structure and grid discretization. They were constructed starting from one particular calibration of the Kerinci reservoir and varying the locations of two adjacent deep upflows, which were moved along the eastern and western faults, defining a different calibration for each location and giving 16 calibrations in total. Calibrations 1 to 8 had the deep upflows on the western fault, with them being at the top of the fault in calibration 1 and moved down the fault by one block from the preceding calibration till calibration 8 in which they were at the bottom of the fault. Likewise the upflows were moved down the eastern fault in calibrations 9 to 16.

The stochastic process is defined by the realization of random variables, which in this problem model the deep upflow locations. We considered a multistage problem with $T = 5$ stages and four random variables modelling the location of the deep upflows in each calibration, given in the vector $\boldsymbol{\xi}(\omega) = [\xi_1, \xi_2, \xi_3, \xi_4]$. They are defined as follows.

$$\xi_1 = \begin{cases} \text{W} & \text{if the deep upflows lie on the western fault} \\ \text{E} & \text{if they lie on the eastern fault} \end{cases},$$

$$\xi_2 = \begin{cases} \text{N} & \text{if the deep upflows lie on the northern half of the fault selected by } \xi_1 \\ \text{S} & \text{if they lie on the southern half} \end{cases},$$

$$\xi_3 = \begin{cases} \text{N} & \text{if the deep upflows lie on the northern half of those selected by } \xi_2 \\ \text{S} & \text{if they lie on the southern half} \end{cases},$$

$$\xi_4 = \begin{cases} \text{N} & \text{if the deep upflows are on the northern half of the those selected by } \xi_3 \\ \text{S} & \text{if they lie on the southern half} \end{cases}.$$

The scenarios represent each calibration in the ensemble, being defined by every combination of these random variable outcomes. The full scenario enumeration is given in Table 1. We assume the random variables are realized sequentially, after each of the first four stages. Since each random variable has two possible outcomes, the branching after each stage in the scenario tree is binary, shown in Figure 1.

Table 1: Kerinci ensemble scenarios

Scenario	Random outcomes	Scenario	Random outcomes
ω_1	$(\xi_1 = \text{W}, \xi_2 = \text{N}, \xi_3 = \text{N}, \xi_4 = \text{N})$	ω_9	$(\xi_1 = \text{E}, \xi_2 = \text{N}, \xi_3 = \text{N}, \xi_4 = \text{N})$
ω_2	$(\xi_1 = \text{W}, \xi_2 = \text{N}, \xi_3 = \text{N}, \xi_4 = \text{S})$	ω_{10}	$(\xi_1 = \text{E}, \xi_2 = \text{N}, \xi_3 = \text{N}, \xi_4 = \text{S})$
ω_3	$(\xi_1 = \text{W}, \xi_2 = \text{N}, \xi_3 = \text{S}, \xi_4 = \text{N})$	ω_{11}	$(\xi_1 = \text{E}, \xi_2 = \text{N}, \xi_3 = \text{S}, \xi_4 = \text{N})$
ω_4	$(\xi_1 = \text{W}, \xi_2 = \text{N}, \xi_3 = \text{S}, \xi_4 = \text{S})$	ω_{12}	$(\xi_1 = \text{E}, \xi_2 = \text{N}, \xi_3 = \text{S}, \xi_4 = \text{S})$
ω_5	$(\xi_1 = \text{W}, \xi_2 = \text{S}, \xi_3 = \text{N}, \xi_4 = \text{N})$	ω_{13}	$(\xi_1 = \text{E}, \xi_2 = \text{S}, \xi_3 = \text{N}, \xi_4 = \text{N})$
ω_6	$(\xi_1 = \text{W}, \xi_2 = \text{S}, \xi_3 = \text{N}, \xi_4 = \text{S})$	ω_{14}	$(\xi_1 = \text{E}, \xi_2 = \text{S}, \xi_3 = \text{N}, \xi_4 = \text{S})$
ω_7	$(\xi_1 = \text{W}, \xi_2 = \text{S}, \xi_3 = \text{S}, \xi_4 = \text{N})$	ω_{15}	$(\xi_1 = \text{E}, \xi_2 = \text{S}, \xi_3 = \text{S}, \xi_4 = \text{N})$
ω_8	$(\xi_1 = \text{W}, \xi_2 = \text{S}, \xi_3 = \text{S}, \xi_4 = \text{S})$	ω_{16}	$(\xi_1 = \text{E}, \xi_2 = \text{S}, \xi_3 = \text{S}, \xi_4 = \text{S})$

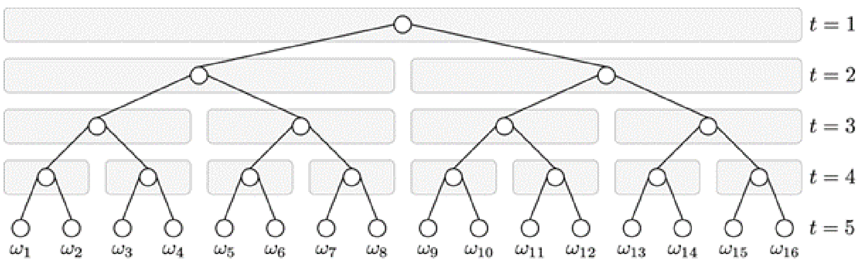

Figure 1: Scenario tree for (SP1)

The branching after the first stage corresponds to the realization of ξ_1 and tells us whether the deep upflows are on the western fault or the eastern fault. Practically this represents finding out whether the productive areas of the reservoir are near the western fault or the eastern fault. For the two nodes in the second stage, NACs link decisions between the scenarios in which the deep upflows are on the west (ω_1 through ω_8), and those in which they are on the east (ω_9 through ω_{16}). The branching after the second stage is due to the realization of ξ_2, which reveals whether the deep upflows are on the northern or southern end of each fault. This narrows down the exact location of a productive area, and leads to four nodes at the third stage, which group scenarios by whether the deep upflow is on the northern end of the western fault (ω_1 through ω_4), the southern end (ω_5 through ω_8), the northern end of the eastern fault (ω_9 through ω_{12}), or the southern end (ω_{13} through ω_{16}). This process continues till the last random variable ξ_4 is realized, causing the branching into the leaf nodes after the fourth stage and leading to full information for each scenario.

The NACs enforcing this information structure are given by:

$$\begin{aligned}
x_i^{(1)}(\omega) &= \bar{x}_i^{(1)}, && \omega \in \Omega, \\
x_i^{(2)}(\omega) &= \bar{x}_i^{(2)}(\xi_1), && \omega \in \Omega, \\
x_i^{(3)}(\omega) &= \bar{x}_i^{(3)}(\xi_1, \xi_2), && \omega \in \Omega, \\
x_i^{(4)}(\omega) &= \bar{x}_i^{(4)}(\xi_1, \xi_2, \xi_3), && \omega \in \Omega,
\end{aligned} \qquad (\mathcal{I}_{\text{SP1}})$$

where $x_i^{(t)}(\omega)$ is the decision to drill a well to feedzone (block) i in scenario $\omega \in \Omega$ and stage t. We do not detail the other constraints that define the problem dynamics. The objective is to maximise the NPV return from the selected wells, with the decision to select a well and the decision to drill it being one and the same. See (Adiga, O'Sullivan, and Philpott 2018) for more details.

3 Newsvendor Example

Consider a Newsvendor problem with $n = 3$ magazine titles, each with its own cost, selling price, and inventory. All magazines must be procured in advance of demand occurring, and within an overall budget b. We denote the set indexing the magazines by $Q = [1, \ldots, n]$. Each magazine $i \in Q$ costs c_i to procure and can be sold for a price p_i or returned for a refund r_i. The demands are given by the random vector $\boldsymbol{\xi}(\omega) = [\xi_1, \ldots, \xi_n]$, which has a multi-variate probability distribution with finite support, giving a finite set of scenarios $\omega \in \Omega$. Let the demand ξ_i for each magazine i have N possible realizations $\xi_i \in \Xi_i$, each with a probability $\mathbb{P}(\xi_i)$ of occuring. The scenarios define every combination of demand outcomes $\omega : (\xi_1, \ldots, \xi_n) \in \prod_{i=1}^{n} \Xi_i$, giving N^n possible realizations of $\boldsymbol{\xi}(\omega)$. We also assume the demands are independent, and therefore each scenario ω has probabilility $\mathbb{P}(\omega) = \mathbb{P}(\xi_1, \ldots, \xi_n) = \prod_{i=1}^{n} \mathbb{P}(\xi_i)$ of occuring.

The optimization problem faced by the newsvendor is a two-stage stochastic program that can be formulated as follows.

(NVPI) :

$$
\begin{array}{llcll}
\min & \sum_{\omega\in\Omega} \mathbb{P}(\omega) \ \sum_{i=1}^{3} & & (c_i x_i(\omega) - p_i y_i(\omega) - r_i w_i(\omega)) & \\
\text{s.t.} & y_i(\omega) & \leq & \xi_i, & i \in Q; \omega \in \Omega, \\
 & y_i(\omega) + w_i(\omega) & = & x_i(\omega), & i \in Q; \omega \in \Omega, \\
 & \sum_{i=1}^{3} c_i x_i(\omega) & \leq & b, & \omega \in \Omega, \\
 & x_i(\omega) & = & u_{1i}(\xi_2, \xi_3), & i \in Q; \omega \in \Omega, \\
 & x_i(\omega) & = & u_{2i}(\xi_1, \xi_3), & i \in Q; \omega \in \Omega, \\
 & x_i(\omega) & = & u_{3i}(\xi_1, \xi_2), & i \in Q; \omega \in \Omega, \\
 & x_i(\omega), y_i(\omega), w_i(\omega) & \geq & 0, & i \in Q; \omega \in \Omega.
\end{array}
$$

where $x_i(\omega)$, $y_i(\omega)$ and $w_i(\omega)$ are the quantities of each magazine to be purchased, sold and returned in each scenario. We define (ξ_{-j}) as the $(n-1)$-tuple of all random variables $(\xi_1, \ldots, \xi_{j-1}, \xi_{j+1}, \ldots, \xi_n)$, ignoring ξ_j. Here the variables $u_{ji}(\xi_{-j}) = u_{ji}(\xi_1, \ldots, \xi_{j-1}, \xi_{j+1}, \ldots, \xi_n)$ must take the same value for all realizations of ξ_j. Thus the constraint $x_i(\omega) = u_{ji}(\xi_{-j})$ for magazine i implies that purchases $x_i(\omega)$ cannot vary with the outcome of demand ξ_j. Including these constraints for each demand ξ_j ensures that $x_i(\omega)$ is the *same* purchase decision for every $\omega \in \Omega$.

The advantage of the formulation (NVPI) is that it enables information revelation for any combination of demands. Relaxing the constraint $x_i(\omega) = u_{1i}(\xi_2, \xi_3)$ gives a model that can adapt to ξ_1 without adapting to other ξ_j. Then relaxing constraint $x_i(\omega) = u_{2i}(\xi_1, \xi_3)$ reveals ξ_2 as well. Choosing different sets of NACs to relax gives information about different combinations of demand outcomes, and iteratively relaxing different NACs gives different sequences of information revelation. Alternatively, we can include exploration decisions which can relax the constraints, modelled as conditional NACs (CNACs), within the formulation. These actions are modelled as binary variables, defined below.

$$
z_i = \begin{cases} 1 & \text{if demand for magazine } i \text{ is investigated} \\ 0 & \text{otherwise} \end{cases}.
$$

They are costed at e_i for each magazine i, and must be made within an exploration budget b_2. With these exploration decisions and costs we formulate the following three-stage information acquisition model for the three magazine example.

(NVIA) :

$$
\begin{array}{llcll}
\min & \sum_{\omega\in\Omega}\mathbb{P}(\omega)\ \sum_{i=1}^{3} & & \big(c_i x_i(\omega) - p_i y_i(\omega) & \\
 & & & \quad - r_i w_i(\omega) + e_i z_i\big) & \\
\text{s.t.} & y_i(\omega) & \leq & \xi_i, & i\in Q;\omega\in\Omega,\\
 & y_i(\omega)+w_i(\omega) & = & x_i(\omega), & i\in Q;\omega\in\Omega,\\
 & \sum_{i=1}^{3} c_i x_i(\omega) & \leq & b_1, & \omega\in\Omega,\\
 & \sum_{i=1}^{3} e_i z_i(\omega) & \leq & b_2, & \omega\in\Omega,\\
 & |x_i(\omega) - u_{1i}(\xi_2,\xi_3)| & \leq & Mz_1, & i\in Q;\omega\in\Omega,\\
 & |x_i(\omega) - u_{2i}(\xi_1,\xi_3)| & \leq & Mz_2, & i\in Q;\omega\in\Omega,\\
 & |x_i(\omega) - u_{3i}(\xi_1,\xi_2)| & \leq & Mz_3, & i\in Q;\omega\in\Omega,\\
 & x_i(\omega), y_i(\omega), w_i(\omega) & \geq & 0, & i\in Q;\omega\in\Omega,\\
 & u_{ji}(\xi_{-j}) & \geq & 0, & \xi_k\in\Xi_k; k\in Q\setminus j; i,j\in Q,\\
 & z_i & \in & [0,1], & i\in Q.
\end{array}
$$

We solved different instances of this model with slightly different problem data. Based on the value added by discovering a demand and cost of doing so, different exploration decisions are selected, giving different *resulting scenario trees*. The trees for two instances are given in Figure 2. In Instance 1 of the problem the CNACs at the second stage enforce non-anticipativity only over scenarios with the same outcome of demand ξ_3. This is because that is the demand that is discovered in the first stage ($z_3 = 1$), thereby breaking non-anticipativity over the different outcomes of ξ_3. The constraints $x_i(\omega) \leq u_{3i}(\xi_1,\xi_2)+Mz_3$ and $x_i(\omega) \geq u_{3i}(\xi_1,\xi_2)-Mz_3$ become $x_i(\omega) \leq u_{3i}(\xi_1,\xi_2)+M$ and $x_i(\omega) \geq u_{3i}(\xi_1,\xi_2)-M$ respectively, allowing $x_i(\omega)$ and $u_{3i}(\xi_1,\xi_2)$ to be different for all $\omega\in\Omega$. Then at the third stage all the demands have been realized and the sale and refund decisions are fully anticipative in each scenario. Similarly, in Instance 2 of the problem the CNACs enforce non-anticipativity over scenarios sharing the same outcomes for both ξ_1 and ξ_3, as those are the demands that are investigated in the first stage ($z_1, z_3 = 1$).

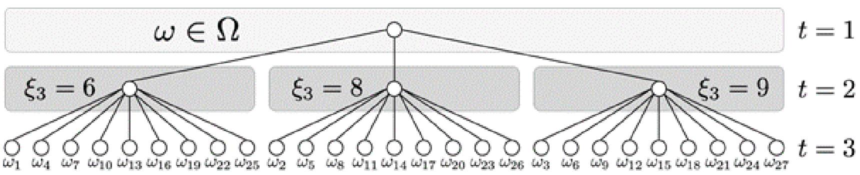

(a) Instance 1

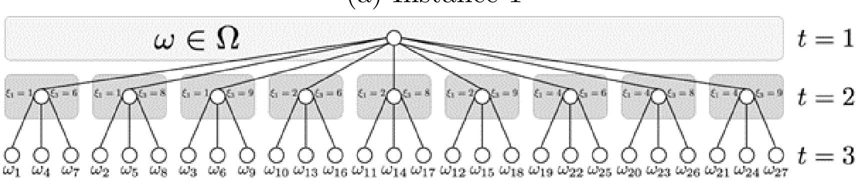

(b) Instance 2

Figure 2: (NVIA) scenario trees for each problem instance

While the scenario trees for the two problem instances look similar, the nodes at the second stage are different, covering different sets of scenarios in each tree. The scenario trees that emerge from the exploration decisions are different for the different instances of the problem depending on the data chosen for the costs e_i. Thus the *resulting scenario tree* is dependent on the solution itself, specifically on the exploration decisions. A helpful way to think about this is that the constraints define a scenario fan containing a number of possible trees, and which of these trees materializes depends on the value of the decision variables that the CNACs are contingent upon in the optimal solution. A scenario tree is just a visual depiction of the information structure of a problem. This distinction applies to the information structures as well. A constraint formulation with CNACs defines an *underlying information structure* in which the relationships between the decision variables govern the information flows that can manifest in the solution, and the particular *resulting information structure* that emerges depends on the decisions in the optimal solution. The information structure of this problem is adaptive in later stages to previously revealed information. By solving both the exploration (z) and exploitation (x, y, w) decisions together in (NVIA), we can optimise the decision variables and the resulting information structure at the same time.

Note that the exploration decisions only gave a *subset* of magazine demand outcomes to discover, not a *sequence*. This is because all the exploration decisions are made at the same stage with the same knowlege (no information in the first stage). (NVIA) was a three-stage model with the exploration decisions z being made in the first stage, the purchase decisions x in the second stage and the sale and return decisions y and v in the third stage. If we extend this so the exploration decisions can be made in multiple stages and apply the information constraints on the exploration variables as well, then we get a model in which the discovery of random variable outcomes follows an optimal policy that can give a different sequence of exploration depending on the outcomes of the investigation. For example, we could get perfect information on the demand ξ_1 and then obtain information on ξ_2 only if ξ_1 is low. We can seek a solution with recourse in the exploration decisions as well, as opposed to that given by (NVIA) which only allowed recourse in the exploitation decisions.

We now consider a model with three exploration stages where we can discover demand information followed by two exploitation stages where we make the purchase, sale and return decisions. We call this model (MAIA). It gives a policy consisting of a series of discovery decisions z followed by the purchase, sale and return decisions x, y and z in each scenario ω. If the chosen exploration decisions remove all demand uncertainty, then $x_i(\omega)$ will be the Wait-and-See purchase solution for each scenario ω. Alternatively the purchase decisions will be adapted to the history of information that has been accrued, without anticipating demand that is still unknown. For brevity we do not present the formulation for (MAIA), but we present the resulting scenario trees from solving the same two problem instances, given in Figure 3.

In instance 1 the CNACs at the second stage enforce non-anticipativity over all the scenarios as none of the demands are investigated in the first stage and therefore no information is obtained. Then the CNACs at the third stage enforce non-anticipativity over scenarios that share the same outcome for demand ξ_3, due to it being discovered in the second stage. This is also the case in the fourth stage for scenarios with the medium and high outcomes of ξ_3 ($\xi_3 \in [8, 9]$), but for the scenarios with the low outcome the decision to discover ξ_1 in the third stage breaks non-anticipativity over ξ_1. Then non-anticipativity applies over the scenarios sharing the same outcomes for ξ_1 and with $\xi_3 = 6$. The purchase decisions are therefore adapted to this history of information revelation.

In Instance 2 the second stage CNACs enforce non-anticipativity over scenarios with the same outcome of demand ξ_1, as it is discovered in the first stage. The

CNACs at the third stage make the scenarios sharing the same outcomes for both ξ_1 and ξ_3 non-anticipative, due to demand ξ_3 being discovered in the second stage. This remains the case in the fourth stage for most of the scenarios. However, for those with the low and medium outcomes of ξ_1 and the low outcome of ξ_3, the decision to discover ξ_2 in the third stage breaks non-anticipativity over ξ_2 as well in the fourth stage, giving full information for those scenarios. The purchase decisions are adapted to this history of information revelation.

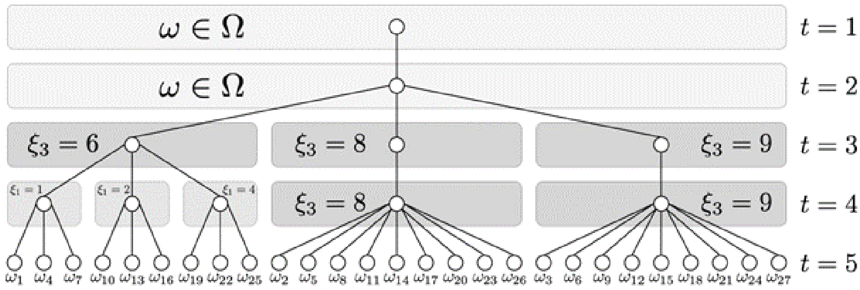

(a) Instance 1

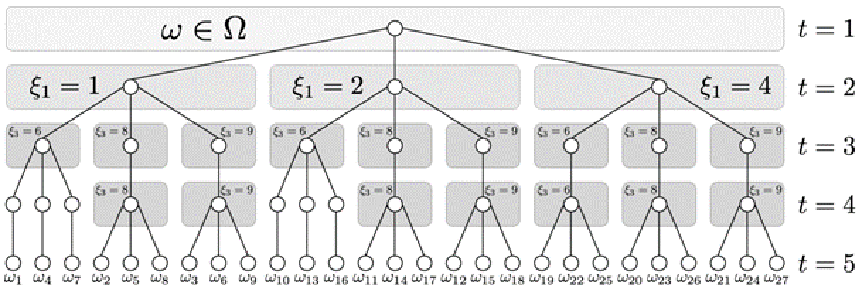

(b) Instance 2

Figure 3: (MAIA) scenario trees for each problem instance

The distinction between the (MAIA) and (NVIA) solutions is that now exploration decisions can also be adapted to the information revealed from previous exploration decisions, giving different future exploration decisions for different outcomes of previous exploration. This can be seen in the scenario trees, which are symmetric for the (NVIA) solutions (see Figure 2), while those for the (MAIA) solutions trees branch asymmetrically. Depending on the previous decisions, there are is no branching after some nodes (no information revealed) in some parts of the trees, and child sub-trees following from some nodal branching are different as they follow different sequences of later stage exploration decisions.

4 Prescriptive Information Revelation Models

We call these kind of models Prescriptive Information Revelation Models (PIRMs), and provide a general formulation for modelling their information constraints. Consider a problem with two types of decisions, exploitation decisions x and exploration decisions z, and scenarios $\omega \in \Omega$. We denote by (ξ_{-J}) the tuple of random variables (ξ_j), $j \in [1, \dots, n] \setminus J$. Let $i \in I$ be the indices defining the subset of exploitation decisions $x_i^t(\omega)$ on which the NACs are applied, with $t \in [1, \dots, T]$ denoting the stage a decision is made in. For a collection $\mathcal{J}$ of such subsets of random variables for which constraints must be enforced, we can write the first stage constraints as

$$x_i^1(\omega) = u_{J,i}^1(\xi_{-J}), \qquad i \in I; J \in \mathcal{J}; \omega \in \Omega.$$

From the second stage onwards the NACs are conditioned upon previous exploration decisions. There may be different exploration actions that can relax these constraints for different scenarios. Therefore the sets of exploration actions need to be defined in terms of both the random variables *over* whose outcomes they will relax a NAC if chosen ($\xi_j, j \in J$), and the specific random variable outcomes corresponding to the scenarios *for* which the constraint will be relaxed. This is defined by a particular outcome $\hat{\xi}_k$ for a random variable indexed by k. There may be multiple random outcomes defining the subset of scenarios for which decisions relax these constraints, and so we can define each set of exploration decisions as z_p, $p \in P_{J,\mathcal{K}}$. Here $\mathcal{K}$ is the set of tuples containing the indices of the random variables and their particular outcomes which correspond to this subset of scenarios.

A NAC over the random variables ξ_j, $j \in J$ could be relaxed by different exploration actions for different subsets of scenarios. Therefore we define by $\mathcal{K}_J$ the collection of sets $\mathcal{K}$, with the tuples in each $\mathcal{K}$ defining a particular subset of scenarios for which the NAC over the random variables defined in J will be relaxed if any exploration decision z_p, $p \in P_{J,\mathcal{K}}$ is selected. This gives the following general definition for CNACs in the second stage and onwards.

$$\left|x_i^t(\omega) - u_{J,i}^t(\xi_{-J})\right| \leq M \sum_{s=1}^{t-1} \sum_{p \in P_{J,\mathcal{K}}} z_p^s(\omega), \qquad i \in I; t \in [2, \dots, T];$$
$$\xi_k = \hat{\xi}_k^l; \xi_m \in \Xi_m; m \in [1, \dots, n] \setminus K; k \in K; (k,l) \in \mathcal{K}; \mathcal{K} \in \mathcal{K}_J; J \in \mathcal{J}.$$

These constraints do not apply over all scenarios $\omega \in \Omega$ because they are defined for the specific realizations $\xi_k = \hat{\xi}_k^l, (k,l) \in \mathcal{K}$. As such, the constraints only apply over $\xi_m \in \Xi_m, m \in [1, \dots, n] \setminus K$, where K indexes the random variables which are defined in $\mathcal{K}$ ($k \in K, (k,l) \in \mathcal{K}$). The set K is defined in relation to $\mathcal{K}$ and not in relation to J; K and J can be overlapping in $[1, \dots, n]$ or disjoint. The same information structure can be applied on the exploration variables as well.

5 Applying the PIRM for the Kerinci reservoir problem

We can use the PIRM framework to model the information structure of the problem with endogenous uncertainty, defining the information constraints on the collection $\mathcal{J} = \{[1], [2], [3], [4]\}$. Then, the linking variables u_J in the constraints for each $J \in \mathcal{J}$ are functions of all the random variables except those that the respective constraints apply over and those that depend on them, (ξ_{-J}). In the first stage the NACs are equality constraints and cannot be relaxed, and as such the constraints for $J \in \mathcal{J} \setminus [1]$ are superfluous and may be omitted, since applying $x_i^{(1)}(\omega) = u_{[1],i}^{(1)}$ is sufficient. However the constraints need to be applied conditionally for all $J \in \mathcal{J}$ from the second stage onwards, since any of them can be relaxed in any stage. Adding these as pairs of CNACs results in the set of information constraints given overleaf.

The constraints from the second stage onwards can be relaxed by various drilling decisions indexed in $P_{J,\mathcal{K}}, \mathcal{K} \in \mathcal{K}_J$ for each $J \in \mathcal{J}$ in any stage, allowing the random variables to be realized at any stage after the first. The set of candidate feedzones to which wells can be drilled, C, is partitioned into subsets of feedzones that when drilled to reveal different pieces of information. Let C_{W} contain the feedzones along the western fault and C_{E} those on the eastern fault. We denote both of them by C_{ξ_1}, with realization of ξ_1 specifying which one. Then let C_{NW} contain the feedzones in the northern half of the western fault and C_{SW} be the set of those in the southern half. Similarly we define C_{NE} and C_{SE}. We denote these sets by $C_{\xi_2\xi_1}$. Continuing this partitioning further, we define the sets $C_{\xi_3\xi_2\xi_1}$ and $C_{\xi_4\xi_3\xi_2\xi_1}$.

$$
\begin{aligned}
&x_i^{(1)}(\omega) = u_{[1],i}^{(1)}, && i \in C; \omega \in \Omega,\\
&\left| x_i^{(t)}(\omega) - u_{[1],i}^{(t)} \right| \leq \textstyle\sum_{s=1}^{t-1} \sum_{p \in P_{J,\mathcal{K}}} x_p^{(s)}(\omega), && i \in C; t \in [2,\ldots,T]; \xi_k = \hat{\xi}_k^l; \xi_m \in \Xi_m;\\
& && m \in Q \setminus K; k \in K; (k,l) \in \mathcal{K}; \mathcal{K} \in \mathcal{K}_{[1]},\\
&\left| x_i^{(t)}(\omega) - u_{[2],i}^{(t)}(\xi_1) \right| \leq \textstyle\sum_{s=1}^{t-1} \sum_{p \in P_{J,\mathcal{K}}} x_p^{(s)}(\omega), && i \in C; t \in [2,\ldots,T]; \xi_k = \hat{\xi}_k^l; \xi_m \in \Xi_m;\\
& && m \in Q \setminus K; k \in K; (k,l) \in \mathcal{K}; \mathcal{K} \in \mathcal{K}_{[2]},\\
&\left| x_i^{(t)}(\omega) - u_{[3],i}^{(t)}(\xi_1,\xi_2) \right| \leq \textstyle\sum_{s=1}^{t-1} \sum_{p \in P_{J,\mathcal{K}}} x_p^{(s)}(\omega), && i \in C; t \in [2,\ldots,T]; \xi_k = \hat{\xi}_k^l; \xi_m \in \Xi_m;\\
& && m \in Q \setminus K; k \in K; (k,l) \in \mathcal{K}; \mathcal{K} \in \mathcal{K}_{[3]},\\
&\left| x_i^{(t)}(\omega) - u_{[4],i}^{(t)}(\xi_1,\xi_2,\xi_3) \right| \leq \textstyle\sum_{s=1}^{t-1} \sum_{p \in P_{J,\mathcal{K}}} x_p^{(s)}(\omega), && i \in C; t \in [2,\ldots,T]; \xi_k = \hat{\xi}_k^l; \xi_m \in \Xi_m;\\
& && m \in Q \setminus K; k \in K; (k,l) \in \mathcal{K}; \mathcal{K} \in \mathcal{K}_{[4]}.
\end{aligned}
\tag{$\mathcal{I}_{\mathrm{SP2}}$}
$$

Now drilling to feedzones in the northernmost two columns where deep upflows can be on the western fault should reveal four distinct pieces of information about the deep upflow locations. These are: if they are on the western fault or not (revealing the outcome of ξ_1) since the feedzones are indexed in C_{W}, if they are on the northern half of the western fault or not (revealing the outcome of ξ_2 for $\xi_1 = \mathrm{W}$) since the feedzones are indexed in C_{NW}, if they are on the northern quarter on the western fault or not (revealing the outcome of ξ_3 for $\xi_1 = \mathrm{W}, \xi_2 = \mathrm{N}$), since the feedzones are indexed in C_{NNW}, and if they are on the northern eighth of the western fault or not (revealing the outcome of ξ_4 for $\xi_1 = \mathrm{W}, \xi_2, \xi_3 = \mathrm{N}$) since the feedzones are indexed in C_{NNNW}.

To model this information revelation in the CNACs from the second stage onwards, we define the collections $\mathcal{K}_{[j]} = \{[(1,l_1),\ldots,(j,l_j)], l_1 \in L_1,\ldots,l_j \in L_j\}, j \in Q$, which pick out every combination of outcomes of the respective random variable and those on whose outcomes its realization depends. For $J = [1]$, the collection $\mathcal{K}_{[1]} = \{[(1,1)],[(1,2)]\}$ picks out both the outcomes of ξ_1. For $J = [2]$, $\mathcal{K}_{[2]} = \{[(1,1),(2,1)],[(1,1),(2,2)],[(1,2),(2,1)],[(1,2),(2,2)]\}$ picks out every combination of outcomes of ξ_1 and ξ_2. Similarly we define collections $\mathcal{K}_{[3]}$ and $\mathcal{K}_{[4]}$. Then for example, choosing a decision indexed in $P_{[3],[(1,1),(2,2),(3,1)]}$ will relax the constraints over the outcomes of ξ_3 for the scenarios with $\xi_1 = \mathrm{W}$, $\xi_2 = \mathrm{S}, \xi_3 = \mathrm{N}$ (ω_5 and ω_6), splitting them from the scenarios with $\xi_1 = \mathrm{W}$, $\xi_2, \xi_3 = \mathrm{S}$ (ω_7 and ω_8).

We then assign each of the partitioned subsets as the set indexing decisions that relax a particular constraint, $P_{J,\mathcal{K}}$, for different $\mathcal{K} \in \mathcal{K}_J$. For $J = [1]$ we assign $P_{[1],[(1,1)]} = C_{\mathrm{W}}$ and $P_{[1],[(1,2)]} = C_{\mathrm{E}}$ as the sets indexing decisions that can relax the constraints over ξ_1, and for $J = [2]$ we assign $P_{[2],[(1,1),(2,1)]} = C_{\mathrm{NW}}$, $P_{[2],[(1,1),(2,2)]} = C_{\mathrm{SW}}$, $P_{[2],[(1,2),(2,1)]} = C_{\mathrm{NE}}$ and $P_{[2],[(1,2),(2,2)]} = C_{\mathrm{SE}}$ as those indexing decisions that can relax the constraints over ξ_2. In general we assign $P_{[j],[(1,l_1),\ldots,(j,l_j)]} = C_{\hat{\xi}_j\ldots\hat{\xi}_1}$ as the set indexing the decisions that can relax a constraint over ξ_j for the scenarios with $\xi_1 = \hat{\xi}_1,\ldots,\xi_j = \hat{\xi}_j$, where $\hat{\xi}_i$ denotes the realization indexed by l_i of the random variable ξ_i. As another example, $P_{[4],[(1,2),(2,1),(3,1),(4,2)]} = C_{\mathrm{SNNE}}$ is the set indexing decisions which can relax the constraint over ξ_4 for $\xi_1 = \mathrm{E}, \xi_2, \xi_3 = \mathrm{N}$ and $\xi_4 = \mathrm{S}$

(scenario ω_{10}), splitting it from the scenario with $\xi_4 = \text{N}$ for the same outcomes of the other random variables (ω_9).

Using these definitions, we can define the specific information constraints for the second stage and onwards. The two pairs of constraints given below are for $J = [1]$. The first is for $\mathcal{K} = [(1,1)]$, applying for scenarios ω_1 through ω_8. The second is for $\mathcal{K} = [(1,2)]$, applying for scenarios ω_9 through ω_{16}.

$$\left| x_i^{(t)}(\omega) - u_{[1],i}^{(t)} \right| \leq \sum_{s=1}^{t-1} \sum_{p \in C_{\text{W}}} x_p^{(s)}(\omega), \qquad i \in C; t \in [2, \dots, T]; \xi_1 = \text{W}; \xi_m \in [\text{N}, \text{S}]; m \in Q \backslash [1],$$

$$\left| x_i^{(t)}(\omega) - u_{[1],i}^{(t)} \right| \leq \sum_{s=1}^{t-1} \sum_{p \in C_{\text{E}}} x_p^{(s)}(\omega), \qquad i \in C; t \in [2, \dots, T]; \xi_1 = \text{E}; \xi_m \in [\text{N}, \text{S}]; m \in Q \backslash [1].$$

Likewise we do this for $J \in [2, 3, 4]$, using the other collections to assign the relevant partitioned subsets of decisions to the appropriate constraints. We do not enumerate them but all of these information constraints, along with the first stage ones given before, are collectively denoted by the set ($\mathcal{I}_{\text{SP2}}$). They define the information structure for the problem given that the reservoir uncertainty is endogenous and that each drilling decisions reveals specific information about each random variable. Along with the constraints defining the problem dynamics and the objective function, which remain the same as in the original problem, they define the the formulation (SP2). Solving it gives the resulting scenario tree shown in Figure 4.

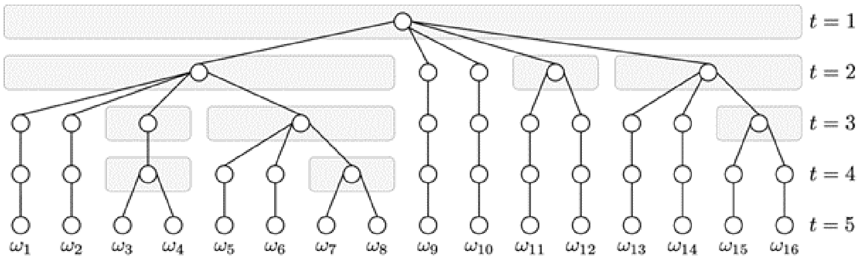

Figure 4: Resulting scenario tree for (SP2) solution

The first stage decision in all scenarios is the same as in the (SP1) solution. It is to drill in the northeast of the reservoir, near the deep upflow locations in scenarios ω_9 and ω_{10}. In (SP1) we then get information only on ξ_1, allowing the second stage decisions to adapt to its outcome (splitting scenarios ω_1 to ω_8 from scenarios ω_9 to ω_{16}). However the act of drilling in the northeast in the (SP2) solution means that not only is ξ_1 revealed, but ξ_2 is revealed for $\xi_1 = \text{E}$ (splitting scenarios ω_9 to ω_{12} from scenarios ω_{13} to ω_{16}), ξ_3 is revealed for $\xi_1 = \text{E}$ and $\xi_2 = \text{N}$ (splitting scenarios ω_9 and ω_{10} from ω_{11} and ω_{12}), and ξ_4 is revealed for $\xi_1 = \text{E}$, $\xi_2 = \text{N}$ and $\xi_3 = \text{N}$ (splitting scenarios ω_9 and ω_{10} from each other). This is because the chosen drilling decision is in the sets C_{E}, C_{NE}, C_{NNE}, C_{NNNE} and C_{SNNE}. Therefore choosing it reveals perfect information on ξ_1, on ξ_2 for $\xi_1 = \text{E}$, on ξ_3 for $\xi_1 = \text{E}$ and $\xi_2 = \text{N}$, and on ξ_4 for $\xi_1 = \text{E}$ and $\xi_2, \xi_3 = \text{N}$.

The second stage decision in the (SP1) solution is then to select and drill in the northwest in scenarios with the outcome $\xi_1 = \text{W}$ (ω_1 to ω_8), and again in the northeast in scenarios with $\xi_1 = \text{E}$ (ω_9 to ω_{16}). However in the (SP2) solution we can now make decisions independently for scenarios with different outcomes of all the random variables revealed in the first stage. Thus the selection for the first eight scenarios ($\xi_1 = \mathrm{W}$) is the same as in (SP1), but is quite varied for the rest. The wells drilled in scenarios ω_9 ($\xi_1 = \mathrm{E}$, ξ_2,ξ_3,$\xi_4 = \mathrm{N}$) and ω_{10} ($\xi_1 = \mathrm{E}$, ξ_2,$\xi_3 = \mathrm{N}$, $\xi_4 = \mathrm{S}$) are different, even thought they both target the northeast of the reservoir. This is due to the more local information revealed about this area in the first stage. For scenarios ω_{11} and ω_{12} ($\xi_1 = \mathrm{E}$, $\xi_2 = \mathrm{N}$, $\xi_3 = \mathrm{S}$) the decision is to drill in the middle of the eastern fault, and for scenarios ω_{13} to ω_{16} ($\xi_1 = \mathrm{E}$, $\xi_2 = \mathrm{S}$) the decision is to drill in the southeast.

These choices then reveal different pieces of information going forward. The decision to drill in the northwest in scenarios ω_1 to ω_8 reveals the outcome of ξ_2 for $\xi_1 = \mathrm{W}$ (splitting scenarios ω_1 to ω_4 from ω_5 and ω_8) since it is indexed in C_{NW}. It reveals the outcome of ξ_3 for $\xi_1 = \mathrm{W}$ and $\xi_2 = \mathrm{N}$ (splitting scenarios ω_1 and ω_2 from ω_3 and ω_4) and the outcome of ξ_4 for $\xi_1 = \mathrm{W}$, $\xi_2 = \mathrm{N}$ and $\xi_3 = \mathrm{N}$ (splitting scenarios ω_1 and ω_2 from each other) as well, since it is also indexed in the sets C_{NNW} and C_{NNNW} respectively. Similarly, drilling in the middle of the eastern fault in scenarios ω_{11} and ω_{12} reveals the outcome of ξ_4 for $\xi_1 = \mathrm{E}$, $\xi_2 = \mathrm{N}$ and $\xi_3 = \mathrm{S}$ (splitting scenarios ω_{11} and ω_{12}) since that decision is indexed in C_{NSNE}. Likewise, drilling in the southeast in scenarios ω_{13} to ω_{16} reveals the outcome of ξ_3 for $\xi_1 = \mathrm{E}$ and $\xi_2 = \mathrm{S}$ (splitting scenarios ω_{13} and ω_{14} from ω_{15} and ω_{16}) since it is indexed in C_{NSE}, and the outcome of ξ_4 for $\xi_1 = \mathrm{E}$, $\xi_2 = \mathrm{S}$ and $\xi_3 = \mathrm{N}$ (splitting scenarios ω_{15} and ω_{16}) since it is indexed in C_{NNSE} and C_{SNSE}.

This process continues through the stages, giving different sequences of information discovery for each scenario which are determined by previous decisions. In general, the (SP2) solution makes more use of information than that of (SP1), with the clustering of selected feedzones in the vicinity of the deep upflows for each scenario happening earlier on. This is because more information is revealed in earlier stages, allowing more adaptive decisions to be made earlier on. As a result, the objective value of the (SP2) solution is higher than that of the (SP1) solution.

References

Adiga, Rishi, John O'Sullivan, and Andy Philpott. 2018. "Stochastic MIP Models for Geothermal Well Selection and Scheduling." *ORSNZ.* Palmerston North.

———. 2019. "Reduced Models for Optimizing Well Placement and Scheduling." *44th Workshop on Geothermal Reservoir Engineering.* Stanford, California.

Apap, Robert M., and Ignacio E. Grossmann. 2017. "Models and computational strategies for multistage stochastic programming under endogenous and exogenous uncertainties." *Computers & Chemical Engineering* 103 (August): 233–274.

Boland, Natashia, Irina Dumitrescu, and Gary Froyland. 2008. "A multistage stochastic programming approach to open pit mine production scheduling with uncertain geology." November.

Goel, Vikas, and Ignacio E. Grossmann. 2006. "A Class of stochastic programs with decision dependent uncertainty." *Mathematical Programming* 108 (2): 355–394 (September).

Jonsbråten, Tore Wiig. 1998. "Optimization Models for Petroleum Field Exploitation." Ph.D. diss.