2011 50th IEEE Conference on Decision and Control and
European Control Conference (CDC-ECC)
Orlando, FL, USA, December 12-15, 2011

Decision rules for information discovery in multi-stage stochastic programming

Phebe Vayanos, Daniel Kuhn and Berç Rustem

***Abstract*— Stochastic programming and robust optimization are disciplines concerned with optimal decision-making under uncertainty over time. Traditional models and solution algorithms have been tailored to problems where the order in which the uncertainties unfold is *in*dependent of the controller actions. Nevertheless, in numerous real-world decision problems, the time of information discovery can be influenced by the decision maker, and uncertainties only become observable following an (often costly) investment. Such problems can be formulated as mixed-binary multi-stage stochastic programs with decision-dependent non-anticipativity constraints. Unfortunately, these problems are severely computationally intractable. We propose an approximation scheme for multi-stage problems with decision-dependent information discovery which is based on techniques commonly used in modern robust optimization. In particular, we obtain a conservative approximation in the form of a mixed-binary linear program by restricting the spaces of measurable binary and real-valued decision rules to those that are representable as piecewise constant and linear functions of the uncertain parameters, respectively. We assess our approach on a problem of infrastructure and production planning in offshore oil fields from the literature.**

***Index Terms*— endogenous uncertainty, binary decision rules.**

I. INTRODUCTION

Stochastic programming is a discipline that develops models and algorithms for solving decision problems affected by uncertain data (see e.g. Birge and Louveaux [1]). In most of these problems, the uncertain parameters are revealed sequentially as time progresses. The decision-making process is therefore dynamic in the sense that the decisions are allowed to depend on the observable data. Mathematically, these adaptive decisions must be modeled as functions or *decision rules* of those uncertain parameters that are known at the time of decision making.

In stochastic programming it is usually assumed that the order in which the uncertainties unfold is independent of the decision maker's actions. However, this assumption fails to hold in numerous real-world decision problems, where the decisions influence the time of information discovery. In order to establish a succinct terminology, Jonsbråten [2] coined the terms of *exogenous* and *endogenous* uncertainties, which refer to parameters whose 'time of revelation' is *in*dependent and dependent of the decisions, respectively. We will use this terminology throughout this paper. Moreover, we will refer to those decisions that trigger an information discovery as *measurement* or *observation* variables.

This work was supported by EPSRC Grant EP/H0204554/1.

The authors are with the Department of Computing, Imperial College London, UK. {fpv02,dkuhn,b.rustem}@imperial.ac.uk

We highlight the practical significance of models with endogenous uncertainties by presenting several real-world decision problems in which the time of information discovery is inherently decision-dependent.

A. Motivating examples

Oil companies spend substantial efforts on infrastructure and production planning in offshore oilfields [3], which typically consist of several reservoirs with uncertain volumes. For each reservoir, one needs to determine whether and when to extract oil. The uncertain volume of a reservoir becomes observable only when an expensive well platform is built and the drilling for oil is initiated. The drilling decisions thus control the sequence of information discovery.

Pharmaceutical companies typically maintain R&D pipelines that comprise multiple candidate drugs. Before a drug can enter the marketplace, it needs to pass several costly clinical trials that may last for many years. The outcome (success/failure) of each trial is uncertain and will only be revealed once the trial is completed. Thus, pharmaceutical companies need to orchestrate the clinical trials with the goal to maximize the rate of discovering effective drugs [4]. The decisions to proceed with different trials can thus be viewed as measurement variables which determine how the uncertainty unfolds.

A related problem is that of R&D project portfolio optimization [5]. Here, the goal is to decide how to distribute scarce resources among a number of projects with different performance characteristics. The return of any project is uncertain and will only be revealed upon the project's termination. The start times of the various projects and the resource allocations thus impact the time of information discovery.

B. Literature review

Research on stochastic programs with endogenous uncertainties began with the works of Jonsbråten et al. [6] and Jonsbråten [2] in 1998 . They studied decision problems in which the control actions can impact both the distribution of the uncertainties as well as the timing of their revelation. Problems with decision-dependent information discovery are perceived as particularly hard even if the distribution of the uncertainties is unaffected by the decisions, and therefore the literature on numerical solution procedures remains scarce. To the best of our knowledge, all existing algorithms rely on the assumption that the uncertain parameters follow a discrete distribution. In this case the decision process can be modeled through a finite scenario tree whose branching structure depends on the binary measurement decisions that determine the time of information discovery, see Jonsbråten [2]. Goel and Grossmann have shown that stochastic programs with discretely distributed endogenous uncertainties can be reformulated as deterministic mixed-binary programs [3], but unfortunately these reformulations involve an exponential number of binary variables and non-anticipativity constraints. Research efforts have consequently focused on approximation techniques that provide suboptimal but feasible solutions to the original problem. An effective approach to complexity reduction is to require that the measurement decisions be pre-committed, that is, to approximate them by here-and-now decisions. The resulting approximate problems can be solved with an enumeration-based branch-and-bound algorithm due to Jonsbråten et al. [6] or via decomposition techniques by Goel and Grossmann [3]. More recent branch-and-bound and branch-and-cut algorithms truthfully account for the adaptive nature of the measurement variables, see Goel and Grossmann [7], [8] as well as Colvin and Maravelias [9], respectively. Moreover, several iterative schemes based on relaxations of the non-anticipativity constraints for the measurement variables have been proposed by Gupta and Grossmann [10] and by Colvin and Maravelias [9].

Problems involving continuously distributed random parameters need to be discretized before any of the above solution procedures can be applied. Solak et al. propose to use a sample average approximation for this purpose [5]. While discretization appears as a promising approach for smaller problems, it may result in a combinatorial state explosion when applied to large and medium sized problems. Conversely, using only very few discretization points can result in solutions that are suboptimal or may even fail to be implementable in practice.

In this paper we develop a methodology for solving dynamic problems with endogenous uncertainties, which is inspired by techniques that recently emerged in robust optimization [11]. We suggest to approximate the adaptive measurement decisions by piecewise constant functions and the adaptive real-valued decisions by piecewise linear functions of the uncertainties. The resulting approximate problems are equivalent to mixed-binary linear programs (MBLP), which can be solved using standard optimization software. This decision rule approximation remains applicable when the uncertain parameters are continuously distributed, and it results in near-optimal solutions that are implementable in reality. The trade-off between the solution quality and the computational speed is controlled by the granularity of the partition of the uncertainty domain. Decision rule techniques have successfully been used in the context of stochastic and robust optimization with *exogenous* uncertainty, see Ben-Tal et al. [12], Shapiro and Nemirovski [13], Goh and Sim [14] and Kuhn et al. [15].

This paper is organized as follows. The remainder of this section introduces the notation, while § II and § III develop a new decision rule approximation for two- and multi-stage stochastic programs affected by endogenous uncertainty, respectively. The benefits of our approach are illustrated in § IV through an example in the area of infrastructure and production planning.

Notation Throughout this paper, uncertainty is modeled by the probability space $(\mathbb{R}^k, \mathcal{B}(\mathbb{R}^k), \mathbb{P})$, which consists of the sample space $\mathbb{R}^k$, the Borel σ-algebra $\mathcal{B}(\mathbb{R}^k)$ and the probability measure $\mathbb{P}$, whose support we denote by Ξ. We assume that Ξ is a compact polyhedral subset of $\{\xi \in \mathbb{R}^k : \xi_1 = 1\}$. This non-restrictive assumption allows us to represent affine functions of the non-degenerate uncertain parameters $(\xi_2, \ldots, \xi_k)$ in a compact way as linear functions of $\xi = (\xi_1, \ldots, \xi_k)$. We let $\mathbb{E}(\cdot)$ denote the expectation operator with respect to $\mathbb{P}$. We further denote by $\mu := \mathbb{E}(\xi)$ the first order moment vector and by $\Sigma := \mathbb{E}(\xi\xi^\top)$ the second order moment matrix of ξ under $\mathbb{P}$. For any $m, n \in \mathbb{N}$, we let $\mathcal{L}_{m,n}$ be the space of all measurable functions from $\mathbb{R}^m$ to $\mathbb{R}^n$ that are bounded on compact sets. For two vectors $x, y \in \mathbb{R}^n$, we let $x \circ y \in \mathbb{R}^n$ denote their Hadamard product and for $j \in \mathbb{N}$, we define $x_{-j} \in \mathbb{R}^{n-1}$ as $x_{-j} := (x_1, \ldots, x_{j-1}, x_{j+1}, \ldots, x_n)$. For a square matrix $A \in \mathbb{R}^{n\times n}$, we let $\mathrm{tr}(A)$ denote the trace of A. Finally, we denote by e_1 the first canonical basis vector in $\mathbb{R}^k$.

II. The two-stage case

A. Problem formulation

A two-stage stochastic program with *exogenous uncertainty* is representable as

$$\begin{array}{ll} \min & c^\top x + \mathbb{E}(q(\xi)^\top y(\xi)) \\ \text{s.t.} & x \in \mathbb{R}^{n_1},\ y \in \mathcal{L}_{k,n_2} \\ & Tx + Wy(\xi) \le h(\xi)\ \forall \xi \in \Xi, \end{array} \tag{1}$$

where $x \in \mathbb{R}^{n_1}$ stands for the vector of first-stage decisions and $y(\xi) \in \mathbb{R}^{n_2}$ denotes the vector of second-stage (or recourse) decisions, which may depend on the observed realization of the random vector $\xi \in \mathbb{R}^k$. Here, $c \in \mathbb{R}^{n_1}$ and $q(\xi) \in \mathbb{R}^{n_2}$ are interpreted as cost vectors, while $T \in \mathbb{R}^{m\times n_1}$ and $W \in \mathbb{R}^{m\times n_2}$ are referred to as the technology and recourse matrices, respectively. Moreover, $h(\xi) \in \mathbb{R}^m$ is termed the right hand side vector. We assume that $q(\xi) = Q\xi$ for some $Q \in \mathbb{R}^{n_2\times k}$ and $h(\xi) = H\xi$ for some $H \in \mathbb{R}^{m\times k}$.

The focus of this paper is a variant of problem (1) in which the random vector is not necessarily observable in the second stage. Instead, any component of ξ is observed only if the decision maker decides to observe (or measure) this particular component. A new binary decision vector $z \in \mathcal{Z} \subseteq \{0,1\}^k$ collects these measurement decisions, that is, ξ_i is observed iff $z_i = 1$. We will henceforth assume that observing random parameters incurs a cost $f^\top z$ for some $f \in \mathbb{R}^k$ and impacts the constraints through an additional term Bz for some $B \in \mathbb{R}^{m\times k}$. In this generalized model, the second stage decisions may only depend on those random parameters that have been observed, that is, they must be representable as functions of $z \circ \xi$. Note that the binary vector z "switches off" those components of ξ that remain unobserved.

A two-stage stochastic program with *endogenous uncertainty* can therefore be formalized as

$$\begin{array}{ll} \min & c^\top x + f^\top z + \mathbb{E}(q(\xi)^\top y(\xi)) \\ \text{s.t.} & x \in \mathbb{R}^{n_1},\ z \in \mathcal{Z},\ y \in \mathcal{L}_{k,n_2} \\ & \left.\begin{array}{l} Tx + Bz + Wy(\xi) \le h(\xi) \\ y(\xi) = y(z \circ \xi) \end{array}\right\} \ \forall \xi \in \Xi. \end{array} \qquad (\mathcal{P})$$

Note that $\mathcal{Z}$ can be a strict subset of $\{0,1\}^k$, that is, it may incorporate constraints requiring that a particular component of ξ is always observed or that two components of ξ must be observed simultaneously etc.

Problem $\mathcal{P}$ encapsulates the two-stage stochastic program (1) and the (deterministic) mixed-binary linear program (MBLP) as special cases, and it involves complex compositions of functional and binary decisions. Therefore, it is severely computationally intractable. In the next section we propose a conservative approximation that reduces $\mathcal{P}$ to a single-stage robust MBLP problem.

B. Decision rule approximation

We can substantially improve the tractability of problem $\mathcal{P}$ by reducing the space of admissible second-stage decisions to those presenting an affine data dependence, thus being representable as $y(\xi) = Y\xi$ for some $Y \in \mathbb{R}^{n_2 \times k}$. This radical but effective approach to complexity reduction was proposed in [12], [14], [16], [13] as a means of approximating multi-stage robust and stochastic programs affected by exogenous uncertainty. Using this approach to simplify problem $\mathcal{P}$, which is affected by endogenous uncertainty, results in the following conservative (upper bound) approximation.

$$\begin{array}{ll} \min & c^\top x + f^\top z + \text{tr}\left(\Sigma Q^\top Y\right) \\ \text{s.t.} & x \in \mathbb{R}^{n_1},\ z \in \mathcal{Z},\ Y \in \mathbb{R}^{n_2 \times k} \\ & Tx + Bz + WY\xi \le h(\xi)\ \forall \xi \in \Xi \\ & |Y_{ij}| \le M z_j \quad i = 1,\dots,n_2,\ j = 1,\dots,k. \end{array} \qquad (\mathcal{P}_u)$$

The last constraint in $\mathcal{P}_u$ enforces non-anticipativity. It stipulates that if ξ_j was not observed in the first decision stage, then the affine decision rule $y(\xi) = Y\xi$ must be independent of ξ_j. Here, $M \in \mathbb{R}_+$ denotes a suitably chosen “big-M constant” which is large enough to guarantee that Y_{ij} is unaffected by the non-anticipativity constraint if $z_j = 1$. Problem $\mathcal{P}_u$ can be viewed as a robust MBLP, which involves semi-infinite constraints parameterized by $\xi \in \Xi$. In the following section, we reformulate $\mathcal{P}_u$ as a standard MBLP.

C. MBLP reformulation

The key ingredient for reformulating $\mathcal{P}_u$ as an MBLP is the following proposition.

Proposition 2.1: For any $\phi \in \mathbb{R}^k$ the following statements are equivalent:

(i) $\phi^\top \xi \ge 0$ for all $\xi \in \Xi$;

(ii) ϕ is an element of the cone dual to the cone generated by Ξ, i.e. $\phi \in \mathcal{K} := (\text{cone}(\Xi))^*$.

Proof: As linear functions are positive homogeneous of degree 1, we have

$$\begin{aligned} \phi^\top \xi \ge 0 \quad \forall \xi \in \Xi \quad &\Leftrightarrow \quad \phi^\top \xi \ge 0 \quad \forall \xi \in \text{cone}(\Xi) \\ &\Leftrightarrow \quad \phi \in (\text{cone}(\Xi))^* \end{aligned}$$

Thus, the claim follows. ■

By Proposition 2.1, $\mathcal{P}_u$ can be reformulated as

$$\begin{array}{ll} \min & c^\top x + f^\top z + \text{tr}\left(\Sigma Q^\top Y\right) \\ \text{s.t.} & x \in \mathbb{R}^{n_1},\ z \in \mathcal{Z},\ Y \in \mathbb{R}^{n_2 \times k} \\ & H - (Tx + Bz)\text{e}_1^\top - WY \in \mathcal{K}^m \\ & |Y_{ij}| \le M z_j \quad i = 1,\dots,n_2,\ j = 1,\dots,k, \end{array} \qquad (\mathcal{P}_u')$$

where $\mathcal{K}^m$ denotes the cone of all $m \times k$-matrices whose rows are all contained in $\mathcal{K}$. Since Ξ is a polyhedral set, $\mathcal{K}^m$ is a polyhedral cone. The conic constraint in $\mathcal{P}_u'$ therefore corresponds to a finite set of linear inequality constraints. Problem $\mathcal{P}_u'$ is thus equivalent to an MBLP involving only a finite number of decision variables and constraints. Its size grows polynomially with k, m, n_1, n_2 and the number of constraints defining the uncertainty set Ξ. The decision rule approximation thus results in a conservative approximation to $\mathcal{P}$ in the form of an MBLP whose size is polynomially bounded in the size of the original problem's inputs.

III. The multi-stage case

A. Problem formulation

A multi-stage stochastic program with exogenous uncertainty over the finite planning horizon $\mathbb{T} := \{1,\dots,T\}$ is representable as

$$\begin{array}{ll} \min & \mathbb{E}\left(\sum_{t \in \mathbb{T}} c_t(\xi)^\top y_t(\xi)\right) \\ \text{s.t.} & y_t \in \mathcal{L}_{k,n_t}\ \forall t \in \mathbb{T} \\ & \left.\begin{array}{l} \sum_{\tau=1}^t A_{t\tau} y_\tau(\xi) \le h_t(\xi) \\ y_t(\xi) = y_t(z_{t-1} \circ \xi) \end{array}\right\} \ \forall \xi \in \Xi,\ t \in \mathbb{T}, \end{array} \qquad (2)$$

where $y_t(\xi) \in \mathbb{R}^{n_t}$ denotes the vector of time t decisions. The binary vector $z_t \in \{0,1\}^k$ represents the *information base* at time $t+1$, that is, it encodes the information revealed up to time t. Thus, we have $z_{t,i} = 1$ iff ξ_i has been observed at some time $\tau \in \{0,\dots,t\}$. As information is never forgotten, we require that $z_t \ge z_{t-1}$ for all $t \in \mathbb{T}$. The last constraint in (2) enforces non-anticipativity by stipulating that y_t can only depend on uncertainties that have been observed up to time $t-1$.

Without much loss of generality, we assume that the problem data satisfies $c_t(\xi) = C_t\xi$ for some $C_t \in \mathbb{R}^{n_t \times k}$, $h_t(\xi) = H_t\xi$ for some $H_t \in \mathbb{R}^{m_t \times k}$ and $A_{t\tau} \in \mathbb{R}^{m_t \times n_\tau}$.

In the remainder we investigate a variant of problem (2) that enjoys much greater modeling power since the time of information discovery is kept flexible. We now interpret the information base $z_t(\xi) \in \mathcal{Z}_t \subseteq \{0,1\}^k$ as an adaptive decision variable, which itself depends on ξ. The set $\mathcal{Z}_t$ may incorporate constraints stipulating, for example, that a specific uncertainty can only be observed after a certain stage etc. We assume that including ξ_i in the information base at time t, which happens iff $z_{t,i}(\xi) = 1$, incurs a cost $f_{t,i}(\xi) \in \mathbb{R}$. Moreover, $z_1(\xi),\dots,z_t(\xi)$ also impact the time t constraints through the additional term $\sum_{\tau=1}^t B_{t\tau} z_\tau(\xi)$ for some $B_{t\tau} \in \mathbb{R}^{m_t \times k}$. Without much loss of generality, we assume that $f_t(\xi) = F_t\xi$ for some $F_t \in \mathbb{R}^{k \times k}$.

A multi-stage stochastic program with endogenous uncertainty can therefore be formalized as

$$\begin{array}{ll} \min & \mathbb{E}\left(\sum_{t\in\mathbb{T}} c_t(\xi)^\top y_t(\xi) + f_t(\xi)^\top z_t(\xi)\right) \\ \text{s.t.} & y_t \in \mathcal{L}_{k,n_t},\ z_t \in \mathcal{L}_{k,k}\ \forall t\in\mathbb{T} \\ & \left.\begin{array}{l} \sum_{\tau=1}^{t} A_{t\tau} y_\tau(\xi) + B_{t\tau} z_\tau(\xi) \le h_t(\xi) \\ z_t(\xi) \in \mathcal{Z}_t \\ z_t(\xi) \ge z_{t-1}(\xi) \\ z_t(\xi) = z_t(z_{t-1}(\xi)\circ\xi) \\ y_t(\xi) = y_t(z_{t-1}(\xi)\circ\xi) \end{array}\right\} \begin{array}{l}\forall \xi\in\Xi, \\ t\in\mathbb{T}.\end{array} \end{array} \qquad (\mathcal{MP})$$

The fourth constraint in $\mathcal{MP}$ corresponds to an information monotonicity constraint and ensures that information is never forgotten, and the last two constraints enforce non-anticipativity of the binary and real-valued decision variables, respectively. Without loss of generality we assume that $z_0(\xi) = \mathrm{e}_1\ \forall \xi \in \Xi$, that is, only the degenerate random parameter ξ_1 is known at the beginning. Problem $\mathcal{MP}$ subsumes the multi-stage stochastic program (2) and it involves decision-dependent non-anticipativity constraints and binary recourse variables. It is therefore severely computationally intractable. In the next section, we propose a conservative approximation that reduces $\mathcal{MP}$ to a static robust MBLP.

B. Decision rule approximation

The emergence of binary recourse variables in multistage models of the type $\mathcal{MP}$ adds another level of complexity to the two-stage models considered in § II. Indeed, while continuous recourse variables can be approximated by linear decision rules [12], [16], [14], [13], there seems to be no flexible decision rule approximation for binary recourse variables which enjoys good tractability properties. Real-valued decision rules that are piecewise constant on the subsets of an *adjustable* partition of Ξ have been studied in [17]. However, this adjustability entails considerable complications in the presence of endogenous uncertainties. We therefore approximate the measurement decisions in problem $\mathcal{MP}$ by binary-valued decision rules that are piecewise constant with respect to a *preselected* partition of Ξ. Similarly, we approximate all real-valued decisions in $\mathcal{MP}$ by decision rules that are piecewise linear with respect to the same partition. Without much loss of generality, we assume that all subsets of this partition are hyper-rectangles of the form

$$\Xi_s := \{\xi \in \Xi : w^i_{s_i-1} \le \xi_i < w^i_{s_i},\ i=1,\dots,k\},$$

where $s \in \mathbb{S} := \times_{i=1}^{k}\{1,\dots,r_i\} \subseteq \mathbb{N}^k$ and

$$w^i_1 < w^i_2 < \cdots < w^i_{r_i-1} \text{ for } i=1,\dots,k$$

represent $r_i - 1$ breakpoints along the ξ_i axis. We now approximate the binary-valued decisions in $\mathcal{MP}$ by piecewise constant decision rules of the form

$$z_t(\xi) = \sum_{s\in\mathbb{S}} \mathbb{I}_{\Xi_s}(\xi) z^s_t \qquad (3)$$

for some $z^s_t \in \{0,1\}^k$, $s\in\mathbb{S}$, $t\in\mathbb{T}$, where $\mathbb{I}_{\Xi_s}$ denotes the indicator function of Ξ_s. Similarly, we approximate the real-valued decisions in $\mathcal{MP}$ by piecewise linear decision rules of the form

$$y_t(\xi) = \sum_{s\in\mathbb{S}} \mathbb{I}_{\Xi_s}(\xi) Y^s_t \xi \qquad (4)$$

for some $Y^s_t \in \mathbb{R}^{n_t\times k}$, $s\in\mathbb{S}$, $t\in\mathbb{T}$.

In order to reduce the notational overhead, we henceforth suppress the domains of the variables $\xi,\xi'\in\Xi$, $t\in\mathbb{T}$, $s,s'\in\mathbb{S}$, $j,j'\in\{1,\dots,k\}$ and $i\in\{1,\dots,n_t\}$.

Proposition 3.1: Under the approximations (3) and (4), the non-anticipativity constraints in $\mathcal{MP}$ are equivalent to

$$\left.\begin{array}{l} |z^s_{t,j'} - z^{s'}_{t,j'}| \le z^s_{t-1,j} \\ |Y^s_{t,ij'} - Y^{s'}_{t,ij'}| \le M z^s_{t-1,j}\ \forall i \end{array}\right\} \begin{array}{c}\forall j,j',s,s',t: \\ s_{-j} = s'_{-j}\end{array} \qquad (5a)$$

$$|Y^s_{t,ij}| \le M z^s_{t-1,j}\ \forall i,j,s,t, \qquad (5b)$$

where M is a sufficiently large big-M constant.

Proof: The non-anticipativity constraints in $\mathcal{MP}$ can be re-expressed as

$$\left.\begin{array}{l} z_t(\xi) = z_t(\xi') \\ y_t(\xi) = y_t(\xi') \end{array}\right\} \forall t,\xi,\xi' : z_{t-1}(\xi)\circ\xi = z_{t-1}(\xi')\circ\xi'.$$

Substituting (3) and (4) into the above expression yields

$$\left.\begin{array}{l} z^s_t = z^{s'}_t \\ Y^s_t = Y^{s'}_t \end{array}\right\} \forall t,s,s' : z^s_{t-1}\circ s = z^{s'}_{t-1}\circ s' \qquad (6a)$$

and

$$|Y^s_{t,ij}| \le M z^s_{t-1,j}\ \forall i,j,s,t. \qquad (6b)$$

Note that (6a) enforces non-anticipativity across distinct subsets of the partition, while (6b) enforces non-anticipativity for the linear decision rules within each subset and is reminiscent of the non-anticipativity constraints in $\mathcal{P}'_u$. We now demonstrate that (5a) and (6a) are equivalent.

($\Leftarrow$) Assume that (6a) holds, and choose some j,s,s' and t with $s_{-j} = s'_{-j}$. The information monotonicity constraint stipulated in $\mathcal{MP}$ implies that

$$z^s_{\tau-1,j} = z^{s'}_{\tau-1,j} = 1 \ \Rightarrow\ z^s_{\tau,j} = z^{s'}_{\tau,j} = 1, \qquad (7)$$

while (6a) and the assumption $s_{-j} = s'_{-j}$ imply that

$$z^s_{\tau-1,j} = z^{s'}_{\tau-1,j} = 0 \ \Rightarrow z^s_\tau = z^{s'}_\tau \text{ and } Y^s_\tau = Y^{s'}_\tau \qquad (8)$$

for all $\tau\in\{0,\dots,t-1\}$. Since $z^s_0 = z^{s'}_0 = \mathrm{e}_1$, we can iteratively apply (7) and the first implication in (8) to conclude that $z^s_{t-1,j} = z^{s'}_{t-1,j}$. Thus, (8) implies

$$\left.\begin{array}{l} |z^s_{t,j'} - z^{s'}_{t,j'}| \le z^s_{t-1,j} \\ |Y^s_{t,ij'} - Y^{s'}_{t,ij'}| \le M z^s_{t-1,j}\ \forall i \end{array}\right\} \forall j'.$$

As j, s, s' and t were chosen arbitrarily, (5a) follows.

($\Rightarrow$) Assume now that (5a) holds and choose some s,s' and t with $z^s_{t-1}\circ s = z^{s'}_{t-1}\circ s'$. As $s_j, s'_j \ge 1\ \forall j$, we conclude that $z^s_{t-1} = z^{s'}_{t-1}$. Thus, z^s_{t-1} and $s-s'$ satisfy the complementarity condition $z^s_{t-1}\circ(s-s') = 0$. If $s=s'$, then (6a) holds trivially true. Next, assume that $s\neq s'$ and that there exists j' with $s_{-j'} = s'_{-j'}$, that is s and s' differ only in their j'th component. The complementarity of z^s_{t-1} and $s-s'$ then ensures that $z^s_{t-1,j'} = 0$. Together with the known identity $s_{-j'} = s'_{-j'}$, this implies via (5a) that $z_t^s = z_t^{s'}$ and $Y_t^s = Y_t^{s'}$. Thus, (6a) follows. Finally, if s and s' differ in two or more components, then (6a) can be established by applying the above argument iteratively. ■

C. MBLP reformulation

Substituting the decision rules (3) and (4) into $\mathcal{MP}$ and applying Proposition 3.1 yields a conservative approximation $\mathcal{MP}_u$ for $\mathcal{MP}$. We then proceed as in § II-C to obtain the following MBLP reformulation of $\mathcal{MP}_u$,

$$\begin{array}{ll} \min & \sum_{s\in\mathbb{S}} p_s \sum_{t\in\mathbb{T}} \mu_s^\top F_t{}^\top z_t^s + \mathrm{tr}(\Sigma_s C_t{}^\top Y_t^s) \\ \text{s.t.} & z_t^s \in \mathcal{Z}_t,\ Y_t^s \in \mathbb{R}^{n_t\times k}\ \forall s,t \\ & H_t - \sum_{\tau=1}^t A_{t\tau} Y_\tau^s + B_{t\tau} z_\tau^s \mathrm{e}_1^\top \in \mathcal{K}_s^{m_t}\ \forall s,t \\ & z_t^s \geq z_{t-1}^s \forall s,t \\ & \left.\begin{array}{l} |z_{t,j'}^s - z_{t,j'}^{s'}| \leq z_{t-1,j}^s \\ |Y_{t,ij'}^s - Y_{t,ij'}^{s'}| \leq M z_{t-1,j}^s\ \forall i \end{array}\right\} \begin{array}{l} \forall s,s' : s_{-j} = s'_{-j} \\ \forall t,j,j' \end{array} \\ & |Y_{t,ij}^s| \leq M z_{t-1,j}^s\ \forall s,t,i,j, \end{array} \qquad (\mathcal{MP}'_u)$$

where $p_s := \mathbb{P}(\xi \in \Xi_s)$, $\mu_s := \mathbb{E}(\xi|\xi \in \Xi_s)$, $\Sigma_s := \mathbb{E}(\xi\xi^\top|\xi \in \Xi_s)$ and $\mathcal{K}_s := (\mathrm{cone}(\Xi_s))^*$. Problem $\mathcal{MP}'_u$ involves only a finite number of decision variables and constraints. For a fixed number of uncertain parameters k and fixed number of breakpoints along each coordinate axis in $\mathbb{R}^k$, the size of $\mathcal{MP}'_u$ remains polynomially bounded in $m := \sum_{t\in\mathbb{T}} m_t$, $n := \sum_{t\in\mathbb{T}} n_t$ and in the number of constraints defining the uncertainty set.

IV. CASE STUDY

A. Problem description

We evaluate the proposed decision rule approach on (a variant of) an infrastructure and production planning problem in offshore oil fields from the literature [3]. An oil company has identified an offshore oil extraction site for possible exploitation. This site comprises several oil fields (or oil reservoirs) with unknown reserves. The company is assumed to be aware of the exact locations of the individual oil fields and needs to plan the oil extraction and gas production process over a period ranging from 10 to 30 years. The objective is to maximize the expected net present value (NPV) of the oil exploitation project.

In order to extract oil from the fields, dedicated well platforms need to be installed and expanded. We denote the set of candidate well platforms (that are under consideration to be built) by $\mathcal{W}$. The oil extracted at the well platforms is sent through a network of directed pipelines to a (unique) production platform $p \in \mathcal{W}$ for gas production. The set of candidate links between well platforms is denoted by $\mathcal{L}$. For any platform $w \in \mathcal{W}$ we denote by $\mathcal{L}^+(w) \subseteq \mathcal{L}$ and $\mathcal{L}^-(w) \subseteq \mathcal{L}$ the sets of all ingoing pipelines *to* w and all outgoing pipelines *from* w, respectively. We assume that all expansion and construction decisions take immediate effect and that once a platform $w \in \mathcal{W}$ has been built, the size ξ^w of the associated oil field is revealed.

We assume that the planning horizon is subdivided into yearly intervals indexed by $t \in \mathbb{T}$. At the beginning of each year, the oil company decides which new platforms and pipelines to construct. We set $z_t^w(\xi) = 1$ if platform w exists at time t; $= 0$ otherwise. Similarly, we set $x_t^l(\xi) = 1$ if pipeline l exists at time t; $= 0$ otherwise. We assume that platforms and pipelines cannot be decommissioned, that is $z_t^w(\xi) \geq z_{t-1}^w(\xi)$ and $x_t^l(\xi) \geq x_{t-1}^l(\xi)$.

In year t the company selects the yearly oil extraction $y_{\mathrm{e},t}^w(\xi)$ for platform w, the yearly flow $y_{\mathrm{e},t}^w(\xi)$ through pipeline l, and the amount $y_{\mathrm{c},t}^w(\xi)$ by which the capacity of platform w is increased at the start of the year. The cumulative oil extraction at a particular field can never exceed the field size,

$$\sum_{\tau=1}^t y_{\mathrm{e},\tau}^w(\xi) \leq \xi^w\ \forall w \in \mathcal{W},$$

while the instantaneous oil extraction is limited by the field's maximum production rate p_w, that is,

$$0 \leq y_{\mathrm{e},t}^w(\xi) \leq p^w\ \forall w \in \mathcal{W}.$$

The flow conservation constraints

$$y_{\mathrm{e},t}^w(\xi) + \sum_{l\in\mathcal{L}^+(w)} y_{\mathrm{f},t}^l(\xi) \geq \sum_{l\in\mathcal{L}^-(w)} y_{\mathrm{f},t}^l(\xi)\ \forall w \in \mathcal{W}$$

ensure that no oil is created within the network, and the box-constraints

$$0 \leq y_{\mathrm{f},t}^l(\xi) \leq M x_t^l(\xi)\ \forall l \in \mathcal{L},$$

which involve a big-M constant, force the flows through yet inexistent pipelines to vanish. Similar box constraints guarantee that yet inexistent platforms cannot be expanded.

$$0 \leq y_{\mathrm{c},t}^w(\xi) \leq M z_t^w(\xi)\ \forall w \in \mathcal{W}.$$

The yearly amount of oil pumped into the network from a particular platform must not exceed that platform's capacity, that is,

$$\sum_{l\in\mathcal{L}^-(w)} y_{\mathrm{f},t}^l(\xi) \leq \sum_{\tau=1}^t y_{\mathrm{c},\tau}^w(\xi)\ \forall w \in \mathcal{W}.$$

The company chooses the design and operating decisions with the aim to maximize the project's expected net present value (NPV)

$$\sum_{t\in\mathbb{T}} d_t \mathbb{E}\left\{ c_{\mathrm{f}}^p \sum_{l\in\mathcal{L}^-(p)} y_{\mathrm{f},t}^l(\xi) - \sum_{l\in\mathcal{L}} c_{\mathrm{b}}^l (x_t^l(\xi) - x_{t-1}^l(\xi)) \right.$$
$$\left. - \sum_{w\in\mathcal{W}} f^w(z_t^w(\xi) - z_{t-1}^w(\xi)) + c_{\mathrm{c}}^w y_{\mathrm{c},t}^w(\xi) + c_{\mathrm{e}}^w y_{\mathrm{e},t}^w(\xi) \right\},$$

where c_{f}^p denotes the unit price for gas, while f^w and c_{b}^l denote the costs for building platform w and pipeline l, respectively. Moreover, c_{c}^w and c_{e}^w represent the unit expansion and extraction costs for platform w, and d_t denotes the discount factor for year t. Note also that $\sum_{l\in\mathcal{L}^-(p)} y_{\mathrm{f},t}^l(\xi)$ represents the total outflow from the production platform, which coincides with the yearly gas production. All decisions selected at the start of year t may depend only on $z_{t-1}(\xi)\circ\xi$. Thus, if $(x_t)_{t\in\mathbb{T}}$ are interpreted as measurement decisions for fictitious degenerate random parameters, the oil extraction problem can be brought to the form $\mathcal{MP}$.

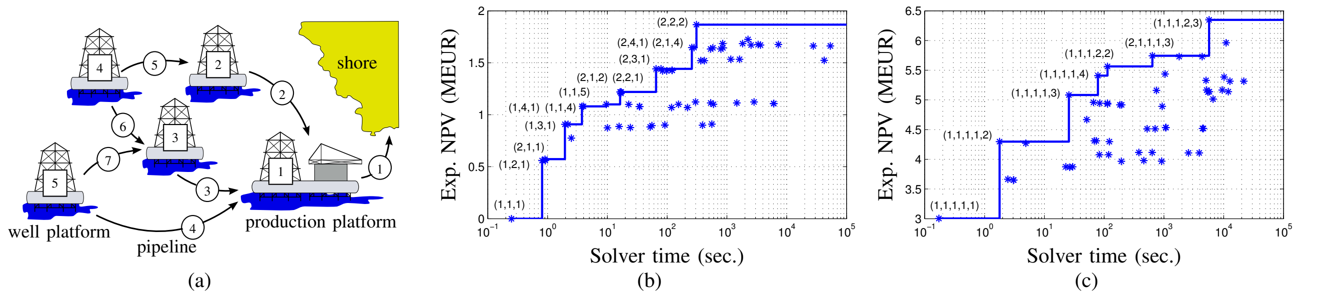

Fig. 1. Fig. (a) shows the offshore oil extraction site. The numbers in squares and circles indicate the platforms $w \in \mathcal{W}$ and the pipelines $l \in \mathcal{L}$. Figs. (b) and (c) illustrate the expected NPV in dependence of solver time for projects A and B. The labels next to the markers represent the breakpoint configuration $(r_1, \ldots, r_k)$ for the problems that achieved the tightest approximation for a given time budget.

B. Numerical results

We consider an instance of the oil extraction problem with a 15 year horizon at the offshore site shown in Fig. 1(a). The field sizes are mutually independent and uniformly distributed as $\xi^w \sim \mathcal{U}(0, u^w)\ \forall w \in \mathcal{W}$. The input parameters of the problem are summarized in Table I.

TABLE I
PROBLEM INPUT PARAMETERS

Parameter	Value	Units
$(u^w)_{w\in\mathcal{W}}$	$(10, 10, 10, 20, 20)$	10^9m^3
$(p^w)_{w\in\mathcal{W}}$	$(0.56, 0.56, 0.56, 1.1, 1.1)$	10^9m^3/year
c_f^p	1.2	EUR/10^3m^3
$(c_\text{b}^l)_{l\in\mathcal{L}}$	$(0, 2, 2, 5, 3, 3, 2)$	MEUR
$(f^w)_{w\in\mathcal{W}}$	$(5, 2, 2, 3, 3)$	MEUR
$(c_\text{c}^w)_{w\in\mathcal{W}}$	$(0.1, 0.1, 0.1, 0.1, 0.1, 0.1)$	EUR/10^3m^3
$(c_\text{e}^w)_{w\in\mathcal{W}}$	$(0, 0, 0, 0, 0, 0)$	EUR/10^3m^3
d_t	$1/(1 + 0.01)^{t-1}$	–

We consider two projects: project A aims at extracting oil from the fields 1 through 3, while project B considers all 5 fields. We proceed as described in § III-B and § III-C to obtain conservative solutions to the expected NPV maximization problems. The partitions of Ξ are constructed such that their subsets have equal probability. The results are shown on Figs. 1(b) and 1(c). For project A (B), we consider all partitions with $|\mathbb{S}| \leq 12$ ($|\mathbb{S}| \leq 6$). The figures illustrate the increase in expected NPV achieved as the solver time[1] increases. For a time budget of less than 70 secs., an increase in expected NPV of more than 1.4 MEUR is achieved relative to the non-adaptive strategy which precommits the measurement decisions at time $t = 1$ and only allows for linear extraction and capacity expansion decisions. We note that for the case $|\mathbb{S}| = 1$, project A appeared not to be profitable. Finally, we remark that exploiting the full site results in a substantially higher profit.

Acknowledgments

We are indebted to Wolfram Wiesemann for valuable discussions.

[1]All problems were solved on a 2.66GHz Intel Core i7-920 machine running CPLEX 12.2.

References

[1] J. Birge and F. Louveaux, *Introduction to Stochastic Programming*. Springer, 2000.
[2] T. Jonsbråten, “Optimization models for petroleum field exploitation,” Ph.D. dissertation, Norwegian Shool of Economics and Business Administration, 1998.
[3] V. Goel and I. Grossman, “A stochastic programming approach to planning of offshore gas field developments under uncertainty in reserves,” *Computers & Chemical Engineering*, vol. 28, no. 8, pp. 1409–1429, 2004.
[4] M. Colvin and C. Maravelias, “A stochastic programming approach for clinical trial planning in new drug development,” *Computers & Chemical Engineering*, vol. 32, no. 11, pp. 2626–2642, 2008.
[5] S. Solak, J.-P. Clarke, E. Johnson, and E. Barnes, “Optimization of R&D project portfolios under endogenous uncertainty,” *European Journal of Operational Research*, vol. 207, no. 1, pp. 420–433, 2010.
[6] T. Jonsbråten, R.-B. Wets, and D. Woodruff, “A class of stochastic programs with decision dependent random elements,” *Annals of Operations Research*, vol. 82, pp. 83–106, 1998.
[7] V. Goel and I. Grossman, “A class of stochastic programs with decision dependent uncertainty,” *Mathematical Programming*, vol. 108, no. 2, pp. 355–394, 2006.
[8] V. Goel, I. Grossman, A. El-Bakry, and E. Mulkay, “A novel branch and bound algorithm for optimal development of gas fields under uncertainty in reserves,” *Computers & Chemical Engineering*, vol. 30, no. 6–7, pp. 1076–1092, 2006.
[9] M. Colvin and C. Maravelias, “Modeling methods and a branch and cut algorithm for pharmaceutical clinical trial planning using stochastic programming,” *European Journal of Operational Research*, vol. 203, no. 1, pp. 205–215, 2010.
[10] V. Gupta and I. Grossman, “Solution strategies for multistage stochastic programming with endogenous uncertainties,” 2010, submitted for publication.
[11] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski, *Robust Optimization*, ser. Princeton Series in Applied Mathematics. Princeton University Press, 2009.
[12] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski, “Adjustable robust solutions of uncertain linear programs,” *Mathematical Programming*, vol. 99, no. 2, pp. 351–376, 2004.
[13] A. Shapiro and A. Nemirovski, *Continuous Optimization: Current Trends and Applications*. Springer, 2005, ch. On complexity of stochastic programming problems, pp. 111–114.
[14] J. Goh and M. Sim, “Distributionally robust optimization and its tractable approximations,” *Operations Research*, vol. 58, no. 4, pp. 902–917, 2010.
[15] D. Kuhn, W. Wiesemann, and A. Georghiou, “Primal and dual linear decision rules in stochastic and robust optimization,” *Mathematical Programming*, 2009.
[16] A. Georghiou, W. Wiesemann, and D. Kuhn, “Generalized decision rule approximations for stochastic programming via liftings,” 2010, available on Optimization Online, submitted for publication.
[17] D. Bertsimas and C. Caramanis, “Finite adaptability for linear optimization,” *IEEE Trans. Automat. Contr.*, vol. 55, no. 12, pp. 2751–2766, 2010.