Algorithmic Operations Research Vol.4 (2009) 1–18

Solving some Multistage Robust Decision Problems with Huge Implicitly Defined Scenario Trees

Michel Minoux

University Paris 6, France

Abstract

This paper describes models and solution algorithms for solving robust multistage decision problems under a special type of uncertainty model referred to here as parsimonious. The main interest of such a model is to provide compact representations of potentially huge scenario trees, leading to efficient dynamic programming-based computation of optimal strategies. Also, contrary to the case of most previously published work on similar problems, which essentially require an independence assumption (on the occurrences of uncertain events in different time periods) our model handles - and properly exploits - some form of dependence over time via a concept of uncertainty budget constraints. Examples of application are discussed including optimal inventory management and the search for robust shortest paths in directed acyclic graphs. Computational results illustrating and validating the proposed approach are also presented.

Key words: robust optimization, robust dynamic programming, uncertainty models, multistage decision models, inventory management.

1. Introduction

How to take the best possible decisions on how to manage a system when information available on this system is partial, unreliable and subject to all kinds of uncertainty, has long been a major concern in Decision Sciences, Automatic Control and Operations Research. To handle such problems, a huge variety of models and solution methods has been proposed in the past including probabilistic models such as:

- two-stage and multistage stochastic programming (see [5], [20]);
- chance-constrained programming (see [8]);
- stochastic dynamic programming (see e.g. [16]).

A well-known limitation of all the approaches based on probabilistic models is that, in many contexts of application, the probability distributions which are assumed to be known to run the solutions algorithms, *are not available*.

One way to bypass the lack of information about probability distributions is the so-called *scenario-based* approach. A complete set of values assigned to each of the uncertain parameters involved in the problem is called a *scenario*. Scenarios can be obtained either by analyzing past data on the behavior of the system, or by simulating how the system works, or by resorting to some expert's experience and skills. It should be noted that most of the contributions concerning scenario-based two-stage or multistage stochastic programming (see [5], [20]) have been focused on optimizing expected values : extensions to handle some measures of risk (variance, CVar), in these probabilistic models have been proposed, essentially for the 2-stage case (see for instance [17]) but, to the best of our knowledge, similar proposals for the multistage case do not seem to have been explored, up to now. In this context, the approach described in the present paper may be viewed as a step towards filling this gap (i.e. developing tools capable of handling risk in the context of multistage optimization problems).

Once a set of scenarios is available, various definitions of robustness for a solution (i.e. a sequence of decisions to be taken over a finite discretized period of time) can be considered. In the model proposed by Kouvelis and Yu [12] which applies to combinatorial optimization problems with uncertainty on the objective function coefficients only, two criteria for selecting an optimal robust solution with respect to a given set of scenarios are proposed: the Min-Max criterion (choose the solution leading to the best objective function value in the worst scenario possible); the Min-Max regret criterion (for a given solution the regret w.r.t. a given scenario is the difference between

Email: Michel Minoux [Michel.Minoux@lip6.fr].

the solution cost for this scenario and the optimal solution cost for this scenario; the Max-regret is the maximum of the regret over all scenarios; we look for the solution minimizing this value). Kouvelis and Yu also investigate extensions of the above robustness criteria to the case of infinite sets of scenarios *defined by interval data*. However, even when applied to well-solved combinatorial problems (such as shortest paths, minimum spanning tree, assignment etc ...), Kouvelis-Yu's approach most often leads to difficult (NP-hard) optimization problems, and this eventually stimulated many subsequent research works, in particular along the idea of approximation.

A different type of approach investigated by Bertsimas and Sim [3], [4], features a twofold interest as compared with Kouvelis-Yu's namely: it preserves polynomial solvability of well-solved problems (such a shortest paths, networks flows etc ...); and it applies more generally to linear programming problems with uncertainty both in the objective function *and* in the constraints. The Bertsimas-Sim approach is concerned with *rowwise uncertainty* i.e. there is an uncertainty set associated with each row of the problem (the objective function itself being viewed as a row). For a given row i, this is defined by requiring that the number of uncertain coefficients in row i which are allowed to deviate from their nominal (="average") value should be less than a predetermined prescribed value Γ_i (in applications, the values Γ_i associated with the various rows have to be specified by the decision-maker). A nice feature of the above uncertainty model is that the robust version of an uncertain linear programming problem can be reduced to standard linear programming with just moderate increase in size (a few additional variables and constraints). On the other hand, a limitation of the approach is that it does not handle uncertain linear programs with column-wise uncertainty, and, in particular linear programs with uncertainty on the right handsides. Such problems were investigated by Soyster [18], [19], but as observed by many authors, the solutions produced by Soyster's models tend to be very "conservative". As an alternative, we have proposed in a previous paper (see [13]) the concept of 2-stage robust decision model in which the set of decision variables is partitioned into two subsets: the set of variables corresponding to *immediate or primary decision* (those to be taken prior to any realization of uncertainty), sometimes also referred to as "here and now" variables; and the set of variables corresponding to *subsidiary decisions* (or adjustments) which can be performed after some realization of uncertainty is observed, sometimes called "wait and see" variables. As a typical example of application of our 2-stage model we mention robust PERT scheduling under uncertainty on the task durations: a robust earliest termination date for the whole project has to be determined, so that it is achievable for any possible combination of task durations out of a given uncertainty set. (This problem turns out to be solvable in polynomial time when the uncertainty set for the task durations is of the Bertsimas-Sim type but *not* by using Bertsimas and Sim's approach which is shown not to be applicable in this case).

The purpose of the present paper is to investigate an extension of the 2-stage model to *multistage* robust decision problems. Since such an extension is concerned with how to take best sequences of decisions on dynamic systems subject to uncertainty, it will be presented in the context of dynamic programming, assuming finite state-space and discrete-time finite horizon. Our model will handle uncertainties, both on the state transition function *and* on the reward function.

Observe that some formally similar multistage robust decision problems have already been considered in the literature, in particular in [11] and [15], where extensions to the standard MDP model to handle uncertainty (in the form of ambiguity in the transition functions) are discussed. However, these works rely on the standard assumptions used in the context of Markov Decision Processes, in particular they require perfect knowledge of the probability measures defining the state transition functions. As will be shown below, the model proposed here turns out to be far less exacting in terms of necessary input data. Moreover the analysis carried out in [11], [15] heavily relies on an assumption, which basically amounts to assuming independence of the outcomes of uncertainty in different time periods. This independence assumption is explicitly stated and referred to as the "rectangularity assumption" in [11] but turns out to be also implicit in [15]. In this respect, a distinctive feature of the present work is that it is based on an uncertainty model which naturally and appropriately handles some specific type of dependency via a concept of uncertainty budget constraints. Our work also appears to be closely related to the so-called *Minimax (or Maximin) control approach* in Dynamic Programming, as described e.g. in [2]. However, the Minimax counterpart to the standard Dynamic programming recursion stated in the above reference (see Chapt. 1, §1.6) again relies on the independence of the uncertainty sets from one time instant to the next.

From this perspective, our approach may be viewed as an extension of the basic Minimax (or Maximin) DP model to handle more complicated uncertainty sets (with non independent occurrences of uncertainty over time), implicitly defined, based on a kind of state-space representation of uncertainty.

The paper is organized as follows. Section 2. provides a general overview of the problem addressed and introduces our notation. In Section 3. we describe the general model of uncertainty against which optimal robust strategies will have to be determined, and which will be referred to as the (multidimensional) *parsimonious uncertainty model*. This model is based on a representation of the uncertainty set as the solution set of a system of linear inequality constraints (referred to as 'uncertainty budget constraints') and leads to compact (implicit) state-space representations of potentially huge scenario trees. A solution algorithm to determine an optimal robust strategy (or: "closed-loop solution") is then described in Section 4.. It is based on a backwards dynamic programming recursion, and its running time is not proportional to the cardinality of the (implicitly defined) set of scenarios, but to the cardinality of a *usually much smaller set* arising from the definition of the parsimonious uncertainty model, and referred to as the *uncertainty status space*. Two typical applications of our model and algorithm are described in Section 5.: one concerning optimal inventory management under uncertainty, the other, some robust shortest path problems in directed circuitless graphs. Computational results on series of randomly generated instances of the inventory management problem under uncertainty are also presented and discussed.

2. Problem Statement: a robust dynamic programming approach

We consider a dynamic system evolving over a discretized finite time period $t = 0, 1, \dots, T$. At each time instant $t = 1, \dots, T$, the system can be in any possible state in a discrete finite set of states $\mathcal{S} = \{1, \dots, N\}$. The state S_0 of the system at time $t = 0$ is supposed to be known. In a classical (deterministic) dynamic programming model, at each time instant t, when the system is in state S at time t, its evolution over the $(t + 1)^{th}$ time period (i.e. between time t and time $t + 1$) is described by a *state transition equation* providing the new state $S1$ of the system at time $t + 1$ as a function of:

- the state at time t;
- the decision x_t taken at the beginning of the $(t+1)^{th}$ time period.

Moreover there is a reward associated with such a transition which also depends on the state S at time t and on the decision x_t.

The set $X(t, S)$ of possible decisions starting from state S at time t is supposed to be a known finite discrete set.

To define a robust version of such a problem we now consider that both the state transition function and the reward function between t and $t + 1$ depend on one or several uncertain parameters ω_t taken in some given finite uncertainty set. (A more precise definition of the uncertainty set will be given in Section 3. below).

In view of preserving the genericity of the model, as much as possible, we will consider that each ω_t is a q-component vector, some of the components of ω_t influencing the state transition function, and some of the components influencing the reward function. Observe that the set of components of ω_t influencing the state transition function and the set of components of ω_t influencing the reward function are not assumed to be disjoint, but they are not assumed to be identical either. The model is thus capable of representing all kinds of situations which may be viewed as intermediates between full dependence and full independence of the uncertainty factors acting on the state transition function and on the reward function.

We will denote:

$$S1 = F(t, S, x_t, \omega_t) \in \mathcal{S} \tag{1}$$

the state reached at time $t + 1$ when starting from $S \in \mathcal{S}$ at time t after taking the decision $x_t \in X(t, S)$ and for the value ω_t of the uncertainty vector. The corresponding reward will be denoted

$$R(t, S, x_t, \omega_t) \in \mathbb{R}. \tag{2}$$

We assume that the decision x_t for period $[t, t+1]$ has to be taken prior to the occurrence (or: realization) of an uncertainty vector ω_t. Therefore it will be assumed that the set of possible decisions $X(t, S)$ is defined in such a way that for any $x_t \in X(t, S)$ and any ω_t in the uncertainty set, the value $F(t, S, x_t, \omega_t)$ is well defined and leads to a feasible state $S1 \in \mathcal{S}$ at time $t + 1$.

3. The parsimonious uncertainty model: compact representation of huge implicitly defined scenario trees

We now discuss the way the uncertainty set Ω is defined in our model. A first simple way of defining the uncertainty set Ω containing all possible uncertain vectors $\omega = (\omega_0, \omega_1, \ldots, \omega_{T-1})$ would be to take Ω as a cartesian product $\Omega = \Omega_0 \times \Omega_1 \times \cdots \times \Omega_{T-1}$.

In such a model, for any $t \in \{0, \ldots, T-1\}$, any $\omega_t \in \Omega_t$ can occur, irrespective of which occurrences of the other uncertainty parameters ω_t, actually arise. This is an independence assumption which, as already mentioned in the introduction, has been somewhat systematically adopted so far in the literature (in [9] and [11] it is referred to as the "rectangularity" property). Of course, under such an assumption, the analysis of the problem is simplified, but the resulting models may not be very realistic in all situations. For instance, if for each t, some $\bar{\omega}_t \in \Omega_t$ represents a worst-case situation, the model will tend to find the best possible solutions against the occurrence of $(\bar{\omega}_1, \bar{\omega}_2, \ldots, \bar{\omega}_T)$ i.e. assuming that at every step t the uncertainty corresponding to the worst-case situation occurs. As a result, the robust solutions proposed by the model will tend to be very " conservative". By contrast, we believe that an uncertainty model capable of representing at least some kinds of dependence among the possible outcomes of uncertain events in various successive time periods would be of potential interest to a number of applications.

Consider as an example the case of optimal daily management of a power distribution system under uncertainty, induced by weather conditions. A worst-case situation with respect to the uncertainty on weather conditions, would be to have every day in winter an extreme-low temperature. Past records show that this never occurs. Similarly, alternating between extreme-low and extreme high temperature from one day to the next during the 3 winter months is not a realistic scenario. More generally it is clear that *any reasonable a priori knowledge on the structure of the uncertainty set should be taken into account* to make the model (and the solutions produced) more realistic.

For instance, referring to our example above, over 90 winter days, assuming that at most k (e.g. $k = 15$) feature extreme-low temperatures and at most k' (e.g. $k' = 15$) feature extreme-high temperatures would be certainly more realistic.

This means that *some kind of dependence* among the occurrences of the various ω_t vectors has to be included in the definition of the set Ω. Referring, once more, to our example above, if we are in the process of constructing a scenario of uncertainty, and if, at some stage t ($t = 30$ say) , we know that the scenario already includes k extreme-low temperatures, we know that (under the uncertainty model suggested above) only average and extreme-high temperatures will be allowed as the subsequent values of temperature in the scenario.

The definition of the uncertainty set which we are going to propose below in connection with our robust dynamic programming model certainly does not pretend to be appropriate for all kinds of applications; however it does capture a wide variety of the time-dependence phenomena among uncertainties which one may wish to take into account in robust dynamic optimization problems. The various examples discussed at the end of the paper will illustrate this capability.

Considering $\Omega_0, \ldots, \Omega_{T-1}$ T finite subsets of $\mathbb{N}^q$, our proposal in to take the uncertainty set Ω as an *implicitly defined subset of the cartesian product* $\Omega_0 \times \Omega_1 \times \cdots \times \Omega_{T-1}$ and, more specifically as the set of all $\omega = (\omega_0, \omega_1, \ldots, \omega_{T-1})$ satisfying:

$$\begin{cases} \omega_t \in \Omega_t \qquad t = 0, \ldots, T-1 \\ \displaystyle\sum_{t=0}^{T-1} \omega_t \leq B \end{cases} \tag{3}$$

for some given $B = (B_1, B_2, \ldots, B_q)^T \in \mathbb{N}^q$. Thus the uncertainty set corresponds to the solution set of an associated system of linear inequalities which will be referred to as the *uncertainty budget constraints*.

An intuitive explanation of the above definition is as follows.

Each $\omega_t \in \Omega_t$ corresponds to a possible outcome, at time t, of an uncertain process influencing the values of some parameters in the evolution of our dynamic system, which correspond to some components of the ω_t vectors (those components $i \in [1, q]$ such that $\omega_t(i) \neq 0$).

Considering first the case where all ω_t $(t = 0, \ldots, T-1)$ are $0-1$ vectors, the i^{th} inequality (3) which reads:

$$\sum_{t=0}^{T-1} \omega_t(i) \leq B_i \tag{4}$$

essentially imposes a limitation on the number of occurrences of the uncertainty process corresponding to the i^{th} component, over the whole time period $[0, T]$ under consideration. In other words, we have a global

”uncertainty budget” B divided among q ”uncertainty features” $(B_1, B_2, \dots, B_q)$, each possible outcome of the underlying uncertain processes influencing the problem being characterized by an ”uncertainty profile” $\omega_t \in \{0,1\}^q$ specifying which uncertainty features are involved (those for which $\omega_t(i) = 1$) or not (those for which $\omega_t(i) = 0$).

Indeed, our model is still slightly more general in that the ω_t vectors (”uncertainty profiles”) may have general (nonnegative) integral components, thus providing additional flexibility in the definition of the uncertainty set Ω (replacing the cardinality constraints (4) by weighted sums with integral weights).

As will be shown in §4. below, for practical applicability of the model and associated solution algorithm, the main limitation will be that the components of the ω_t vectors and the B vector be sufficiently small integers, in order that the quantity

$$\prod_{i=1}^{q}(B_i + 1)$$

remains sufficiently small (typically less than 10^3 to 10^4). In section 4 below, the above quantity will turn out to be the cardinality of the *uncertainty status space*.

However, even with this restriction, it is worth pointing out that an attractive feature of the proposed model is its capability of handling (implicitly) huge scenario trees, much larger than those which can be used in scenario-based stochastic programming problems (see e.g. [5], [6]). Consider, for instance, a 12 period $(T = 12)$ optimal inventory management problem (such as the one discussed in Section 5.1. below) with uncertainty on procurement costs and on requirements: at each time period t, we have: a nominal value α_t^0 and two extreme value α_t^- and α_t^+ for the procurement cost; a nominal value d_t^0 and two extreme values d_t^- and d_t^+ for the requirements. We have here two uncertainty features $(q = 2)$, one corresponding to uncertainty on procurement costs, the other corresponding to requirements. Assuming independence of the two uncertain processes, we have to consider at each time period 9 combinations of uncertainty corresponding to the 9 possible outcomes:

$$\begin{pmatrix}\alpha_t^0\\ d_t^0\end{pmatrix} \begin{pmatrix}\alpha_t^0\\ d_t^-\end{pmatrix} \begin{pmatrix}\alpha_t^0\\ d_t^+\end{pmatrix} \begin{pmatrix}\alpha_t^-\\ d_t^0\end{pmatrix} \begin{pmatrix}\alpha_t^-\\ d_t^-\end{pmatrix} \begin{pmatrix}\alpha_t^-\\ d_t^+\end{pmatrix} \begin{pmatrix}\alpha_t^+\\ d_t^0\end{pmatrix} \begin{pmatrix}\alpha_t^+\\ d_t^-\end{pmatrix} \begin{pmatrix}\alpha_t^+\\ d_t^+\end{pmatrix}.$$

Suppose now we want to define Ω by allowing at most 6 deviations from nominal value for each uncertainty feature, i.e. $B_1 = B_2 = 6$.

In that case, for each t, we would represent the 9 combinations of uncertainty by the following vectors:

$$\Omega_t = \left\{ \overset{w_t^1}{\begin{pmatrix}0\\0\end{pmatrix}} \overset{w_t^2}{\begin{pmatrix}0\\1\end{pmatrix}} \overset{w_t^3}{\begin{pmatrix}0\\1\end{pmatrix}} \overset{w_t^4}{\begin{pmatrix}1\\0\end{pmatrix}} \overset{w_t^5}{\begin{pmatrix}1\\1\end{pmatrix}} \overset{w_t^6}{\begin{pmatrix}1\\1\end{pmatrix}} \overset{w_t^7}{\begin{pmatrix}1\\0\end{pmatrix}} \overset{w_t^8}{\begin{pmatrix}1\\1\end{pmatrix}} \overset{w_t^9}{\begin{pmatrix}1\\1\end{pmatrix}} \right\}.$$

Now, the number of distinct solutions to (3), i.e. the total number of scenarios corresponding to the above definition turns out to be greater than

$$C_{12}^6 \times 9^6 \simeq 491 \times 10^6.$$

On the other hand, since the cardinality of the status space is only $(6 + 1) \times (6 + 1) = 49$ in this case, using the dynamic programming recursion presented in the next section, an exact optimal robust strategy taking into account the uncertainty corresponding to all the scenarios in the implicitly defined uncertainty set Ω defined above, can be computed *in a matter of seconds* on a standard PC workstation (assuming that the cardinality of the state space representing the possible inventory levels is not too large, typically less than 10^3 to 10^4).

4. A dynamic programming-type recursion to find an optimal robust strategy

We now address the question of actually computing an optimal robust strategy for a dynamic optimization problem given by specifying:

(a) the uncertainty set Ω implicitly defined by providing, $\forall t$, the list of all possible $\omega_t \in \Omega_t$ together with the ”uncertainty budget” constraint (3):

$$\sum_{t=0}^{T-1} \omega_t \leq B;$$

(b) the state transition function F defined in (1);

(c) the reward function R defined in (2).

In the same way as for the case of stochastic multiperiod optimization problems, a solution to the above stated problem does not correspond to some well-defined sequence of decisions $(x_1, x_2, \dots, x_T)$ leading to maximum reward (in the worst situation created by uncertainty).

Indeed a solution corresponds to a *strategy* (also sometimes referred to as "closed loop" solution) which is specified by associating with each state S, at each time instant t, the value x_t of an optimal decision to be taken during period $[t, t+1]$ given the current *uncertainty status* of state S at time t.

The *uncertainty status* of a state S at time t provides the necessary and sufficient information concerning the past occurrences of uncertainty, in view of properly restricting the future possible scenarios to be considered (from instant t to T) to only those in Ω.

To do so, it is easily seen that we only have to record the q-component vector $\sigma = (\sigma_1, \sigma_2, \dots, \sigma_q) \in \mathbb{N}^q$ corresponding to the fraction of the uncertainty budget which is still available to represent the uncertainties to be taken into account between instants t and T.

Therefore, for each state S at each time t the total number of possible σ vectors (each corresponding to an *uncertainty status*) will be:

$$K = (B_1 + 1) \times (B_2 + 1) \times \cdots \times (B_q + 1).$$

In the following, the set of all possible σ vectors will be denoted Σ (the uncertainty status space).

Now a *strategy*, i.e. a solution to our robust dynamic optimization problem, will be defined by associating a decision $x_t \in X(t, S)$ with each state S, at each time t and for each uncertainty status σ. Thus, if φ is a strategy, we will denote $\varphi(t, S, \sigma) \in X(t, S)$ the decision to be taken under strategy φ at time t when the system is in state S and for the uncertainty status $\sigma \in \Sigma$.

An optimal robust strategy φ^* is a strategy such that, for each state S at time t and each uncertainty status σ, $\varphi^*(t, S, \sigma)$ is the decision leading to the maximum reward over the period $[t, T]$ in the worst possible scenario in the set $\Omega^{\sigma}_{[t,T]}$ of all $(\omega_t, \omega_{t+1}, \dots, \omega_{T-1}) \in \Omega_t \times \Omega_{t+1} \times \cdots \times \Omega_{T-1}$ such that:

$$\sum_{\theta=t}^{T-1} \omega_\theta \leq \sigma.$$

We denote $z^*(t, S, \sigma)$ the maximum reward which can be obtained using the optimal strategy φ^* when starting from state S at time t with the uncertainty status σ.

We now show that the z^* and φ^* values can be determined via the following backward dynamic programming recursion:

$$z^*(t, S, \sigma) = \operatorname*{Max}_{x_t \in X(t,S)} \{\psi(t, S, \sigma, x_t)\} \quad (5)$$

where:

$$\psi(t, S, \sigma, x_t) = \operatorname*{Min}_{\substack{\omega_t \in \Omega_t \\ \omega_t \leq \sigma}} \{R(t, S, x_t, \omega_t) + \quad (6)$$

$$z^*(t+1, S1, \sigma - \omega_t)] \quad (7)$$

$$\text{and} \qquad S1 = F(t, S, x_t, \omega_t)$$

$$\varphi^*(t, S, \sigma) = \operatorname*{argmax}_{x_t \in X(t,S)} \{\psi(t, S, \sigma, x_t)\}. \quad (8)$$

We assume of course that for $t = T$ the $z^*(T, S, \sigma)$ values are known and given (they can be interpreted as "end of game" rewards. Of course, since the uncertain phenomena occurring during the period $[0, T]$ do not influence what possibly takes place after time T, it is legitimate to assume that, for all $\sigma \in \Sigma$, the $z^*(T, S, \sigma)$ values are equal to a single reference value, $z^*(T, S, 0)$, which is the unique "end of game" value for state S at time T.

Proposition 1 *The $z^*(t, S, \sigma)$ and $\varphi^*(t, S, \sigma)$ values computed from the backward recursion* (5)-(8)*, starting with given $z^*(T, S, 0)$ for all $S \in \mathcal{S}$, define the reward and decision functions associated with an optimal robust strategy over $[0, T]$. In particular, the best (robust) decision to be taken in the first stage, starting with the system in state S_0 at time 0, is $\varphi^*(0, S_0, B) \in X(0, S_0)$, and the corresponding optimal worst-case reward (against the proposed uncertainty model $\omega \in \Omega$) is $z^*(0, S_0, B)$.*

Proof: The result is obtained by induction, along the same lines as for the case of Minimax (or Maximin) Dynamic Programing (see [2]). The main difference lies in the fact that the uncertainty sets involved at each step of the recursion now depend (via the uncertainty status σ) on the various possible past occurrences of uncertainty. □

The following result shows that the recursion (5)-(8) is a pseudopolynomial algorithm for solving the robust multistage decision problem.

Proposition 2 *Let us denote $\Omega_{\max} = \operatorname*{Max}_{t=0,\dots,T-1} \{|\Omega_t|\}$ the maximum number of possible realizations of uncertain events which can occur at time t. The computational complexity of the recursion* (5)-(8) *is $\mathcal{O}(T \times |\mathcal{S}| \times |X| \times |\Sigma| \times \Omega_{\max})$ where $|\mathcal{S}|$ is the state space cardinality, $|X|$ the cardinality of the set of possible decisions and $|\Sigma| = (B_1 + 1) \times (B_2 + 1) \times \cdots \times (B_q + 1)$ is the cardinality of the uncertainty status space.*

Proof: At each stage $t = 1, \ldots, T$, for each state $S \in \mathcal{S}$, for each uncertainty status $\sigma \in \Sigma$, and for each decision x_t, we have to compute the value $\psi(t, S, \sigma, x_t)$ given by (7). Each such value is obtained as the minimum of at most $\Omega_{\max}$ terms. Assuming (which is realistic for many applications) that each term can be computed in time $\mathcal{O}(1)$, the result follows. □

It is of interest to compare the above result with the complexity of the standard dynamic programing algorithm when applied to a problem of a comparable size but without uncertainty, namely $\mathcal{O}(T \times |\mathcal{S}| \times |X|)$. It is seen that the approach proposed here for handling uncertainty increases complexity by a factor $|\Sigma| \times \Omega_{\max}$ which is typically much smaller than the number of scenarios implicitly represented by the parsimonious uncertainty model introduced here (for the example given at the end of Section 3., this factor is only 49×9, whereas the total number of scenarios is greater than 400 millions !).

It is also interesting to compare the above result with the complexity of the dynamic programming-based approach to the robust 0-1 knapsack problem proposed in [21], which features a computational effort growing exponentially with the cardinality of the set of scenarios. By contrast, the complexity of our procedure *only grows linearly with the cardinality of the uncertainty status space*, which, as already mentioned, can be considerably smaller than the number of scenarios (again refer to the example given at the end of Section 3. for typical figures).

5. Some applications

We describe in this section two typical applications of our robust dynamic programming model, one concerning optimal inventory management under uncertainties (§5.1.) the other concerning some new variants of the robust shortest path problem (§5.2.). Computational results on the optimal robust inventory management problem will be presented and discussed in §5.1.3..

5.1. *Robust optimal inventory management*

We consider a multiperiod inventory problem for a single product, assuming discretized time over a finite horizon $[0, T]$. At each time period [t-1,t] $(t = 1, \ldots, T)$ the product under consideration:

- can be bought at unit price α_t ;
- can be sold at unit price β_t.

We also have to satisfy the requirements of the customers over time. The product quantity required in the t^{th} time period (i.e. between the time instants $t-1$ and t) is denoted d_t.

In the standard deterministic version of the problem, all the quantities α_t, β_t and d_t are supposed to be exactly known. By contrast, we will consider an extended version of the problem where some of (or all) the quantities α_t, β_t, d_t are subject to uncertainty. Indeed, since the prices β_t at which the product is sold to customers are in control of the decision-makers, it is legitimate to assume that they are not subject to uncertainty. Therefore we will only consider uncertainty on the d_t and α_t values.

To illustrate the flexibility of our model, two different ways of describing uncertainty on the d_t and α_t values will be successively described: the first one will assume independence of the sources of uncertainty influencing the requirements and the prices ; the second one will get rid of this independence assumption.

5.1.1. *A first uncertainty model: the independent case*

Consistent with our general model, the uncertainty domain $\mathcal{D}$ for the d_t values will typically be defined as follows.

For each t a set of possible values for d_t

$$D_t = \left\{d_t^1, d_t^2, \ldots, d_t^{\nu}\right\}$$

is considered. (This set includes, but is not necessarily limited to a nominal value and one or two extreme values for d_t). Note that, for the sake of notational simplicity, we assume the cardinality of the set D_t equal to the same integer ν for all t, but the model would readily accommodate time-varying cardinalities. To each $d_t^k \in D_t$ we attach a p-component vector $v_t^k \in \mathbb{N}^p$ corresponding to the *uncertainty profile* associated with the occurrence of the value d_t^k of d_t. Given a p-component vector $B_d \in \mathbb{N}^p$ (uncertainty budget for requirements) the uncertainty set $\mathcal{D}$ is then defined as the set of T-vectors of the form: $d = \left\{d_1^{k_1}, d_2^{k_2}, \ldots, d_T^{k_T}\right\}$ such that:

$$\sum_{j=1}^{T} v_t^{k_j} \le B_d \tag{9}$$

(in the above, $k_j \in [1, \nu] \quad \forall j$).

In a similar way, the uncertainty domain $\mathcal{A}$ for the prices α_t would be defined by considering for each t a set of possible values of α_t, namely:

$$A_t = \left\{\alpha_t^1, \alpha_t^2, \ldots, \alpha_t^{\nu'}\right\}$$

and by associating a q-component vector $w_t^k \in \mathbb{N}^q$ (uncertainty profile) with each possible occurrence α_t^k of α_t (again, the assumption $|A_t| = \nu'$, $\forall t$ is just for notational simplicity). Then the uncertainty set $\mathcal{A}$ is the set of T-vectors of the form:

$$\alpha = \left(\alpha_1^{k_1}, \alpha_2^{k_2}, \ldots, \alpha_T^{k_T}\right)$$

satisfying

$$\sum_{j=1}^{T} w_t^{k_j} \le B_\alpha \tag{10}$$

where $B_\alpha \in \mathbb{N}^q$ is the given "uncertainty budget" for prices.

To subsume, in this way of representing uncertainty, we have at each step t, $\nu \times \nu'$ combinations of values for α_t and d_t and two independent sets of uncertainty budget constraints: (9) for the requirements, involving p-component vectors v_t^k and B_d ; and (10) for the prices, involving q-component vectors w_t^k and B_α.

In order to illustrate the above model and the application of the solution procedure described in §4, we will consider the following small numerical example. In this example there are $T = 4$ time periods, the product can be bought or sold only by integer amounts and the maximum capacity of the inventory is $c = 10$ units. (So, at each time instant, the set of possible states for the inventory is $\mathcal{S} = \{0, 1, \ldots, 10\}$). x_t $(t = 1, \ldots, T)$ denoting the number of units purchased at the beginning of period t, the objective is to find an optimal robust strategy in terms of the x_t variables to maximize profit (selling returns -procurement costs-stockout penalties + end-of-game value).

At each time period t, we assume that d_t can take only 3 values, one "nominal value" d_t^2, and two "extreme values" an extreme low-value $d_t^1 < d_t^2$, and an extreme-high value $d_t^3 > d_t^2$ (thus, $\nu = 3$). Similarly, α_t can take either a "nominal value" α_t^1 or an "extreme-high value" $\alpha_t^2 > \alpha_t^1$ (thus $\nu' = 2$).

The v_t^k and w_t^k vectors ("uncertainty profiles") are just 1-dimensional (scalars) with 0-1 values: $v_t^1 = 1$, $v_t^2 = 0$, $v_t^3 = 1$, and $w_t^1 = 0$, $w_t^2 = 1$ (thus $p = 1$ and $q = 1$). The uncertainty budget for requirements is: $B_d = 2$, and for prices: $B_\alpha = 2$. In other words, in this example, the uncertainty budget constraint on d_t (resp.: on α_t) simply allows at most 2 occurrences of an extreme value for d_t (resp.: α_t) over the period $[0, 4]$ under consideration. Finally we note that, due to uncertainty, we have to take into account the possibility of *stockout*. We will therefore consider that, at each time period t, each missing unit of the product will incur a penalty cost π_t. The following table provides the numerical values for α_t, β_t, d_t and π_t, for $t = 1$ to $T = 4$. Note that we have taken a constant unit penalty cost $\pi_t = 20$ (big enough so that the model will tend to produce an optimal strategy avoiding stockout as much as possible).

	$t=1$	$t=2$	$t=3$	$t=4$
α	$\alpha_1^1 = 2$ $\alpha_1^2 = 4$	$\alpha_2^1 = 5$ $\alpha_2^2 = 6$	$\alpha_3^1 = 3$ $\alpha_3^2 = 8$	$\alpha_4^1 = 4$ $\alpha_4^2 = 6$
β	$\beta_1 = 4$	$\beta_2 = 6$	$\beta_3 = 6$	$\beta_4 = 7$
d	$d_1^1 = 2$ $d_1^2 = 3$ $d_1^3 = 5$	$d_2^1 = 5$ $d_2^2 = 6$ $d_2^3 = 8$	$d_3^1 = 3$ $d_3^2 = 4$ $d_3^3 = 6$	$d_4^1 = 5$ $d_4^2 = 6$ $d_4^3 = 9$
π	$\pi_1 = 20$	$\pi_2 = 20$	$\pi_3 = 20$	$\pi_4 = 20$

The initial inventory level is assumed to be $S_0 = 0$. Moreover each unit of product remaining in the inventory at the end of the last (4^{th}) period is supposed to have an "end of game" value equal to 3. So we can start the application of the recursion (5)-(8) of §4 by taking:

$$z^*(4, S, \sigma) = 3.S, \text{ for all } \sigma = \begin{pmatrix}\sigma_1 \\ \sigma_2\end{pmatrix} \in [0, 2]^2.$$

A detailed account of all the calculations resulting from the recursion (5)-(8) would take too much room, so we only provide part of the intermediate results obtained.

Table 1 below provides the values $z^*(0, 0, \sigma)$ and $\varphi^*(0, 0, \sigma)$ obtained at the end of the recursion for the first stage, and for an initial inventory level 0. $\varphi^*(0, 0, \sigma)$ is the optimal decision to be taken (the amount of the product to be bought) in the first stage for each uncertainty status, to achieve the corresponding z^* value.

It is seen from this table that, following the optimal strategy φ^*, the worst-case optimal return which can be expected is 15 corresponding to $\sigma = \begin{pmatrix}B_\alpha \\ B_d\end{pmatrix} = \begin{pmatrix}2 \\ 2\end{pmatrix}$. Also the value $z^*\left(0, 0, (0, 0)\right) = 67$ indicates the optimal solution value for the deterministic problem with all prices and requirements equal to their nominal values (the corresponding solution is $x_1 = 10$, $x_2 = 0$, $x_3 = 9$, $x_4 = 0$, with an inventory level 0 at the end of last period).

Table 1

$\sigma = (\sigma_1, \sigma_2)$	(0,0)	(1,0)	(2,0)	(0,1)	(1,1)	(2,1)	(0,2)	(1,2)	(2,2)
$z^*(0, 0, \sigma)$	67	47	28	58	37	17	55	35	15
$\varphi^*(0, 0, \sigma)$	10	10	10	10	10	10	10	10	10

The $z^*(0, 0, \sigma)$ and $\varphi^*(0, 0, \sigma)$ values as obtained from the recursion (5)-(8) *on the inventory problem. σ_1 (resp. σ_2) denotes the uncertainty status w.r.t. the prices (resp. the requirements).*

Indeed an uncertainty status $\begin{pmatrix} 0 \\ 0 \end{pmatrix}$ at time 0 corresponds to a problem for which no deviation from nominal value is allowed for the prices and for the requirements. The difference between the two values 67 and 15 can thus be interpreted as the *"price of robustness"*.

Now, suppose that after taking the decision $x_1 = 10$ corresponding to the optimal strategy (see Table 1) we are informed that an extreme-high requirement value $d_1 = 5$ and a nominal price $\alpha_1 = 2$ occurred in stage 1. So, at the end of stage 1, the state of the inventory is $S = 5$, and, moreover, 1 unit of the uncertainty budget for requirements has been consumed, therefore the optimal decision to be taken for stage 2 is the one corresponding to the uncertainty status $\begin{pmatrix} 2 \\ 1 \end{pmatrix}$.

From Table 2, it is thus seen that the optimal decision to be taken is

$$x_2 = \varphi^* \left(1, 5, \begin{pmatrix} 2 \\ 1 \end{pmatrix} \right) = 5,$$

the corresponding z^* value being

$$z^* \left(1, 5, \begin{pmatrix} 2 \\ 1 \end{pmatrix} \right) = 24.$$

These values $z^* \left(1, 5, \begin{pmatrix} 2 \\ 1 \end{pmatrix} \right)$ and $\varphi^* \left(1, 5, \begin{pmatrix} 2 \\ 1 \end{pmatrix} \right)$ can easily be deduced, using the recursion (5)-(8), from the values $z^*(2, S, \sigma)$ displayed in Table 3 below.

5.1.2. *A possibly more realistic model: the dependent case*

In order to make the uncertainty model even closer to reality, it may be desirable to take into account some dependencies among the various sources of uncertainty which may be observed in practice. In the context of inventory management, it is frequently the case that demands and prices are correlated: prices tend to rise as a result of demand growth, and tend to fall down as demand is reduced. We now show that our parsimonious uncertainty model is flexible enough to take into account such phenomena.

The idea is the following. Instead of considering independently, at each stage t, possible realizations of the prices and possible realizations of the requirements, we will consider a list of possible realizations of the pair $\begin{pmatrix} \alpha_t \\ d_t \end{pmatrix}$ which we denote $\begin{pmatrix} \alpha_t^1 \\ d_t^1 \end{pmatrix} \begin{pmatrix} \alpha_t^2 \\ d_t^2 \end{pmatrix} \cdots \begin{pmatrix} \alpha_t^\nu \\ d_t^\nu \end{pmatrix}$. (Again, for notational simplicity, we assume that the lists for $t = 1, 2, \ldots, T$ have the same cardinality ν).

In practice, these pairs should be chosen so as to reflect the type of dependencies observed in reality (e.g. α_t^k will be high when d_t^k is high, and low when d_t^k is low).

Associated with each pair $\begin{pmatrix} \alpha_t^k \\ d_t^k \end{pmatrix}$ in the above list, we will also consider a p-component integer vector $v_t^k \geq 0$ (uncertainty profile), the uncertainty set defining all the possible realizations for the sequences of pairs $\begin{pmatrix} \alpha_1 \\ d_1 \end{pmatrix} \begin{pmatrix} \alpha_2 \\ d_2 \end{pmatrix} \cdots \begin{pmatrix} \alpha_T \\ d_T \end{pmatrix}$ being specified as the set of all $\begin{pmatrix} \alpha_1^{k_1} \\ d_1^{k_1} \end{pmatrix} \begin{pmatrix} \alpha_2^{k_2} \\ d_2^{k_2} \end{pmatrix} \cdots \begin{pmatrix} \alpha_T^{k_T} \\ d_T^{k_T} \end{pmatrix}$ such that $\sum_{i=1}^{T} v^{k_i} \leq B$ where $B > 0$ is a given p-component integer vector ("uncertainty budget").

Illustrating this on our 4-stage inventory problem, we might consider for instance the numerical values given in Table 5 (here $\nu = 4$ and $p = 2$) and $B = \begin{pmatrix} 2 \\ 2 \end{pmatrix}$. As compared with the model considered, in §5.1.1. for the independent case, it is seen that such a representation essentially excludes pairs $\begin{pmatrix} \alpha_t \\ d_t \end{pmatrix}$ where α_t would take on an extreme-high (resp. an extreme-low) value when d_t takes an extreme-low (resp. an extreme-high) value. This is consistent with the above-mentioned observed correlation between α_t and d_t.

We observe that taking dependencies into account amounts to excluding some outcomes of uncertainty from the list of possible outcomes considered under the

Table 2

(σ_1, σ_2)	(0,0)	(1,0)	(2,0)	(0,1)	(1,1)	(2,1)	(0,2)	(1,2)	(2,2)
$z^*(1,5,\sigma)$	67	46	40	56	35	24	53	34	20
$\varphi^*(1,5,\sigma)$	1	4	5	3	4	5	3	5	5

The $z^*(1,5,\sigma)$ *and* $\varphi^*(1,5,\sigma)$ *values as obtained from the recursion* (5)-(8) *on the inventory problem.* σ_1 *(resp.* σ_2*) denotes the uncertainty status w.r.t. the prices (resp. the requirements).*

Table 3

S \ (σ_1, σ_2)	(0,0)	(1,0)	(2,0)	(0,1)	(1,1)	(2,1)	(0,2)	(1,2)	(2,2)
0	36 (10)	10 (4)	−2 (4)	29 (10)	−5 (6)	−19 (6)	27 (10)	−7 (6)	−19 (6)
1	39 (9)	18 (3)	6 (3)	32 (9)	3 (5)	−11 (5)	30 (9)	1 (5)	−11 (5)
2	42 (8)	24 (2)	14 (2)	35 (8)	11 (4)	−3 (4)	33 (8)	9 (4)	−3 (4)
3	45 (7)	30 (2)	22 (1)	38 (7)	19 (3)	5 (3)	36 (7)	17 (3)	5 (3)
4	48 (6)	34 (2)	30 (0)	41 (6)	23 (2)	13 (2)	39 (6)	23 (2)	13 (2)
5	51 (5)	39 (1)	36 (0)	44 (5)	29 (2)	21 (1)	42 (5)	29 (2)	21 (1)
6	54 (4)	45 (1)	42 (0)	47 (4)	35 (2)	29 (0)	45 (4)	33 (2)	29 (0)
7	57 (3)	50 (1)	48 (0)	50 (3)	39 (2)	35 (0)	48 (3)	38 (1)	35 (0)
8	60 (2)	54 (0)	54 (0)	53 (2)	44 (1)	41 (0)	51 (2)	44 (1)	41 (0)
9	63 (1)	60 (0)	60 (0)	56 (1)	50 (1)	47 (0)	54 (1)	49 (1)	47 (0)
10	66 (0)	66 (0)	66 (0)	59 (0)	53 (0)	53 (0)	57 (0)	53 (0)	53 (0)

The $z^*(2,S,\sigma)$ *values (and the* $\varphi^*(2,S,\sigma)$ *values in parenthesis) as obtained from the recursion* (5)-(8) *on the inventory problem.* σ_1 *(resp.* σ_2*) denotes the uncertainty status w.r.t. the prices (resp. the requirements).*

Table 4

	$t=1$				$t=2$				$t=3$				$t=4$			
$\begin{pmatrix}\alpha_t^k\\ d_t^k\end{pmatrix} \rightarrow$	$\begin{pmatrix}2\\2\end{pmatrix}$	$\begin{pmatrix}2\\3\end{pmatrix}$	$\begin{pmatrix}4\\3\end{pmatrix}$	$\begin{pmatrix}4\\5\end{pmatrix}$	$\begin{pmatrix}5\\5\end{pmatrix}$	$\begin{pmatrix}5\\6\end{pmatrix}$	$\begin{pmatrix}6\\6\end{pmatrix}$	$\begin{pmatrix}6\\8\end{pmatrix}$	$\begin{pmatrix}3\\3\end{pmatrix}$	$\begin{pmatrix}3\\4\end{pmatrix}$	$\begin{pmatrix}8\\4\end{pmatrix}$	$\begin{pmatrix}8\\6\end{pmatrix}$	$\begin{pmatrix}4\\5\end{pmatrix}$	$\begin{pmatrix}4\\6\end{pmatrix}$	$\begin{pmatrix}6\\6\end{pmatrix}$	$\begin{pmatrix}6\\9\end{pmatrix}$
$v_t^k \rightarrow$	$\begin{pmatrix}0\\1\end{pmatrix}$	$\begin{pmatrix}0\\0\end{pmatrix}$	$\begin{pmatrix}1\\0\end{pmatrix}$	$\begin{pmatrix}1\\1\end{pmatrix}$	$\begin{pmatrix}0\\1\end{pmatrix}$	$\begin{pmatrix}0\\0\end{pmatrix}$	$\begin{pmatrix}1\\0\end{pmatrix}$	$\begin{pmatrix}1\\1\end{pmatrix}$	$\begin{pmatrix}0\\1\end{pmatrix}$	$\begin{pmatrix}0\\0\end{pmatrix}$	$\begin{pmatrix}1\\0\end{pmatrix}$	$\begin{pmatrix}1\\1\end{pmatrix}$	$\begin{pmatrix}0\\1\end{pmatrix}$	$\begin{pmatrix}0\\0\end{pmatrix}$	$\begin{pmatrix}1\\0\end{pmatrix}$	$\begin{pmatrix}1\\1\end{pmatrix}$

previous model (§5.1.1.). The resulting robust strategies can thus only be less conservative, leading to higher worst-case returns. This will be confirmed by the computational experiments reported in Section 5.1.3. below.

Also we note that there is much modeling flexibility offered in the specification of the uncertainty sets via the choice of the v_t^k vectors (uncertainty profiles) ; indeed the number p of components and the values of the components themselves may be given all kinds of interpretations and consequently ”modulated” to fit the needs of the application under consideration in the best way possible.

5.1.3. *Computational experiments*

We present here various computational experiments illustrating and validating the flexibility of the proposed approach on series of randomly generated instances of the inventory management problem under uncertainty.

a) The independent case

The first series of experiments concerns the independent case (§5.1.1.), the corresponding instance involving 10 period problems ($T = 10$) with an inventory of maximum capacity $c = 30$ units and initial inventory level $S_0 = 0$.

Each instance has been generated as follows. For each time period $t = 1, \ldots, T$, the extreme values $d_{\min}(t)$ and $d_{\max}(t)$ for demands, $\alpha_{\min}(t)$ and $\alpha_{\max}(t)$ for procurement costs, are obtained by drawing at random uniformly distributed independent integer valued random variables $\gamma \in [5, 15]$, $\delta \in [3, 9]$, $\theta \in [10, 20]$, $\eta \in [2, 8]$ and setting

$$d_{\min}(t) = \gamma \; ; \; d_{\max}(t) = d_{\min}(t) + \delta \; ;$$
$$\alpha_{\min}(t) = \theta \; ; \; \alpha_{\max}(t) = \alpha_{\min}(t) + \eta \; ;$$

Finally, the selling prices β_t are obtained by drawing at random integer valued independent uniformly distributed variables in the interval $[20, 30]$.

Two uncertainty budget constraints are considered ($q = 2$), the first one corresponding to demands and the second one to procurement costs. For any given time period t, the various possible combinations $\begin{pmatrix} d \\ \alpha \end{pmatrix}$ of a demand value and a procurement cost value are all pairs of integers belonging to $[d_{\min}(t), d_{\max}(t)] \times [\alpha_{\min}(t), \alpha_{\max}(t)]$. These pairs are indexed by $k = 1, \ldots, K_t$ (where $K_t = (d_{\max}(t) - d_{\min}(t) + 1)\,(\alpha_{\max}(t) - \alpha_{\min}(t) + 1)$). With each such pair $\begin{pmatrix} d_t^k \\ \alpha_t^k \end{pmatrix}$ we associate the coefficients $v_t^k \in \mathbb{R}$ and $w_t^k \in \mathbb{R}$ in the two uncertainty budget constraints. The value of each of these coefficients depends on how much d_t^k (resp. α_t^k) deviates from the nominal value $d_{\text{nom}}(t) = \left\lfloor \dfrac{d_{\max}(t) + d_{\min}(t)}{2} \right\rfloor$ for demands and : $\alpha_{\text{nom}}(t) = \left\lfloor \dfrac{\alpha_{\max}(t) + \alpha_{\min}(t)}{2} \right\rfloor$ for procurement costs. Their precise values are defined as:

$$v_t^k = \left| d_t^k - d_{\text{nom}}(t) \right|$$
$$w_t^k = \left| \alpha_t^k - \alpha_{\text{nom}}(t) \right| .$$

It is seen that, in the above, occurrences of demands (resp. of procurement costs) and their weights in the associated budget constraint do not depend on the values taken by procurement costs (resp. by demands). Also, observe that, for each t, there is a single pair $\begin{pmatrix} d_t^k \\ \alpha_t^k \end{pmatrix}$ for which both coefficients v_t^k and w_t^k are zero, namely the pair corresponding to nominal values, both for demands and for procurement costs.

Table 5 displays the results obtained on a series of 12 instances (numbered P1 to P12). For each instance, 9 different values of $B = \begin{pmatrix} B_d \\ B_\alpha \end{pmatrix}$, the right hand side of the budget constraints are considered. Each entry in the table provides:

- the value of the optimal robust strategy resulting from the application of the dynamic programming recursion (5)-(8);
- in parenthesis, the corresponding optimal decision in the first time period (the amount of product to be purchased in the first time period according to the optimal strategy).

We first observe that the figures displayed in the first column, for which $B = \begin{pmatrix} 0 \\ 0 \end{pmatrix}$, correspond to the *deterministic case*, where there is no uncertainty (demands and procurement cost taking their nominal value at each time instant).

We also note that the figures in the last column for which $B = \begin{pmatrix} \infty \\ \infty \end{pmatrix}$ correspond to optimal strategies under the classical Max-Min dynamic programming approach. Indeed, the uncertainty budget constraints are inactive in this case, therefore at each step of the DP recursion, the worst case situation to be considered is independent of the past occurrences of uncertainty (this contrasts with what occurs in our parsimonious uncertainty model for smaller values of $B = \begin{pmatrix} B_d \\ B_\alpha \end{pmatrix}$, when the uncertainty budget constraints become active).

The figures shown in columns 2-8 of Table 5 correspond to cases where the uncertainty budget constraints tend to become more and more effective as the B_d and B_α values are decreased. They properly illustrate the significance of the impact of the size of the budget set on the quality of the resulting optimal strategies. For each column in Table 5 (for each value of $B = \begin{pmatrix} B_d \\ B_\alpha \end{pmatrix}$) the difference in maximum rewards as compared with the figures in the first column can be interpreted as the "*price of robustness*".

Big differences are observed, in terms of maximum profit, between the optimal deterministic solutions (first column) and the optimal Max-Min strategies, confirming the commonly accepted idea that the latter tends to produce fairly conservative solutions.

In this respect, an interesting feature of the new uncertainty model proposed in the present paper is to provide a systematic way of exploring "intermediate" and less conservative robust strategies between these two extremes.

b) The dependent case

The second series of test problems concerns the case where, at each time period t, procurement costs and demands appear to be correlated. The intensity of the correlation will be controlled by means of a real parameter $u \in [0, 1]$, the previously described representation of uncertainty being modified accordingly. In the new resulting uncertainty model for demands and procurement costs, the value $u = 0$ corresponds (for fixed t) to totally correlated values of d_t^k and α_t^k; the value $u = 1$ corresponds to uncorrelated values, in other words, for $u = 1$, we find again the uncertainty model considered for the independent case.

For any $u \in [0, 1]$, and $t = 1, \ldots, T$, the definition of the set of allowed occurrences of pairs $\begin{pmatrix} d_t \\ \alpha_t \end{pmatrix}$ is changed as follows. Instead of allowing all the pairs $\begin{pmatrix} d_t \\ \alpha_t \end{pmatrix}$ in $[d_{\min}(t), d_{\max}(t)] \times [\alpha_{\min}(t), \alpha_{\max}(t)]$ we only allow those pairs which satisfy the conditions:

$$\left| \frac{d_t - \bar{d}_t}{d_{\max}(t) - d_{\min}(t)} - \frac{\alpha_t - \bar{\alpha}_t}{\alpha_{\max}(t) - \alpha_{\min}(t)} \right| \leq u \tag{11}$$

where:

$$\bar{d}_t = \frac{d_{\min}(t) + d_{\max}(t)}{2}$$

$$\bar{\alpha}_t = \frac{\alpha_{\min}(t) + \alpha_{\max}(t)}{2}$$

Observe that $u = 0$ in the above amounts to requiring that the only allowed $\begin{pmatrix} d_t \\ \alpha_t \end{pmatrix}$ pairs are those which correspond through a linear function (more precisely, the linear function φ for which $\alpha_{\max}(t) = \varphi(d_{\max}(t))$ and $\alpha_{\min}(t) = \varphi(d_{\min}(t))$, and this is the case of perfect correlation. On the other hand, for $u = 1$, any pair $\begin{pmatrix} d_t \\ \alpha_t \end{pmatrix}$ in $[d_{\min}(t), d_{\max}(t)] \times [\alpha_{min}(t), \alpha_{\max}(t)]$ satisfies (11), therefore this corresponds to the case where there is no dependence between demands and procurement costs.

To better understand the role of the parameter u in controlling the amount of dependency, let us consider the example (illustrated in Figure 1) where: $[d_{\min}(t), d_{\max}(t)] = [2, 8]$, $[\alpha_{min}(t), \alpha_{\max}(t)] = [5, 9]$, and $u = 0.2$.

The points corresponding to the 11 pairs $\begin{pmatrix} d \\ \alpha \end{pmatrix}$ satisfying (11) are shown as bold dots. If we interpret this set of pairs in terms of a standard probabilistic model, assuming each pair can occur with equal probability (1/11), then the corresponding correlation coefficient has value 0.943. More generally, there is a direct relationship between the value of u and the correlation coefficient ρ as shown in Table 6 below (columns 1 and 2). (The third column in Table 6 also provides the values of ρ as a function of u in the case of infinitely many points uniformly distributed in a box of $\mathbb{R}^2$ of the form $\left[\underline{d}, \bar{d}\right] \times [\underline{\alpha}, \bar{\alpha}]$ with $\underline{d} < \bar{d}$, $\underline{\alpha} < \bar{\alpha}$; interestingly this function does not depend on the ratio $\left(\bar{d} - \underline{d}\right) / \left(\bar{\alpha} - \underline{\alpha}\right)$). This confirms that u is indeed a relevant parameter to control the intensity of the dependence between demands and procurement costs in our robust dynamic programming model for inventory management.

Table 7 displays the results obtained on some of the instances taken from the previous series of experiments for 3 distinct choices for B (the right handside of the budget constraints) and 6 distinct values of u ranging from 0.1 to 1.

For each instance considered, the intervals of variation of demands and procurement costs are not influenced by the parameter u, they are therefore the same as for the independent case (which corresponds to $u = 1$).

However it is seen that, if an a priori knowledge about possible dependence between demands and procurement costs is available, it can be appropriately taken into account by our model, possibly leading to significantly improved optimal strategies. More precisely, Table 9 summarizes the average relative improvements (in terms of optimal worst-case benefits) over the standard Maximin DP model (column $u = 1$) obtained for

Table 5

$B \rightarrow$	$\binom{0}{0}$	$\binom{1}{1}$	$\binom{3}{3}$	$\binom{5}{5}$	$\binom{8}{8}$	$\binom{10}{10}$	$\binom{12}{12}$	$\binom{20}{20}$	$\binom{\infty}{\infty}$
P1	980 (28)	908 (29)	803 (28)	713 (27)	608 (16)	558 (16)	508 (16)	379 (16)	320 (16)
P2	1630 (30)	1560 (30)	1435 (22)	1358 (21)	1259 (16)	1206 (16)	1159 (16)	1013 (21)	956 (26)
P3	1357 (12)	1297 (12)	1177 (12)	1081 (13)	989 (14)	940 (14)	896 (14)	765 (15)	668 (15)
P4	1289 (30)	1219 (30)	1116 (29)	1029 (28)	911 (22)	850 (21)	796 (21)	660 (29)	626 (30)
P5	737 (9)	674 (9)	581 (10)	491 (10)	415 (12)	376 (12)	337 (12)	209 (12)	100 (11)
P6	1583 (12)	1527 (12)	1433 (16)	1352 (16)	1252 (20)	1188 (20)	1124 (13)	977 (13)	943 (13)
P7	1266 (29)	1217 (29)	1128 (20)	1058 (18)	965 (16)	909 (15)	858 (16)	704 (16)	630 (16)
P8	1403 (30)	1343 (30)	1229 (29)	1133 (29)	1004 (27)	934 (19)	877 (19)	701 (21)	610 (27)
P9	1142 (9)	1072 (10)	974 (18)	908 (17)	814 (18)	754 (17)	702 17)	547 (12)	478 (12)
P10	1185 (10)	1115 (10)	1007 (11)	923 (12)	817 (16)	762 (13)	716 (13)	598 (14)	523 (14)
P11	1027 (14)	957 (14)	847 (15)	744 (16)	644 (17)	592 (17)	545 (17)	396 (17)	366 (17)
P12	1244 (29)	1174 (29)	1063 (29)	968 (26)	849 (21)	787 (19)	733 (18)	579 (30)	546 (30)

Impact of the size of the uncertainty set on the values of optimal robust strategies.

Table 6

Values of u	Correlation coefficient ρ integer points in $[2,8] \times [5,9]$	Correlation coefficient ρ continuous uniform distribution in a box of $\mathbb{R}^2$
0	1	1
0.05	1	0.994
0.1	0.982	0.978
0.2	0.943	0.911
0.3	0.872	0.801
0.4	0.783	0.658
0.5	0.565	0.5
0.6	0.466	0.344
0.7	0.368	0.207
0.8	0.254	0.098
0.9	0.138	0.027
1	0	0

Correspondence between the values of the parameter u and the correlation coefficient ρ under the standard probabilistic interpretation

the various values of B and of u. It is observed that these improvements, while relatively modest for smaller values of B, can be as large as 15 % for $u = 0.2$, and close to 20 % for $u = 0.1$ when $B = \binom{12}{12}$.

Table 8

$u \rightarrow$	0.1	0.2	0.4	0.6	0.8	1
$B = \binom{5}{5}$	9.7 %	5.7 %	1.4 %	0.15	0	0
$B = \binom{8}{8}$	14 %	9.8 %	3.5	0.36	0.1	0
$B = \binom{12}{12}$	19.8 %	15.5	7.5	1.5	0.5	0

Average improvements over the Maximin DP model as deduced from the results of Table 7.

The above results illustrate well the possible impact of dependence among the various sources of uncertainty in a multistage robust decision model.

5.2. *Some robust shortest path models on circuitless graphs*

We consider a circuitless directed graph $G = [\mathcal{N}, \mathcal{U}]$ with $n = |\mathcal{N}|$ nodes and $m = |\mathcal{U}|$ arcs, in which

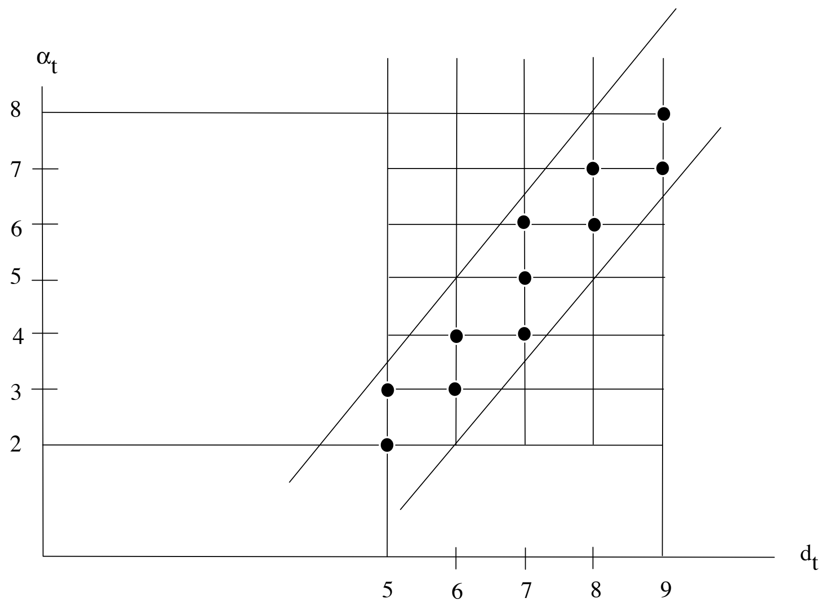

Fig. 1. The 11 integer points in $[5, 9] \times [2, 8]$ satisfying condition (11) for $u = 0.2$. This is a case featuring a significant amount of correlation between demands and procurement costs (high demands tend to correspond to rather high procurement costs, low demands to rather low procurement costs).

we distinguish a root node (a node having zero in-degree) and a target node. We assume that the nodes are numbered according to a topological ordering (i.e. $(i, j) \in \mathcal{U} \Rightarrow i < j$), and without loss of generality, that the root node is indexed 1, and the target note is indexed n.

Each arc $u = (i, j) \in \mathcal{U}$ has an associated length l_u which is not exactly known but which can take any value from a given finite set of values: $L_u = \{l_u^1, l_u^2, \ldots, l_u^\nu\}$ (for the sake of notational simplicity, we assume that all L_u have equal cardinality ν, but of course the proposed model is more general and readily extends to the case of nonuniform cardinalities).

Associated with each l_u^k value in L_u, we assume that we are given a p-component integer vector $w_u^k \geq 0$ representing the "uncertainty profile" of the corresponding realization $l_u = l_u^k$ for the length of arc u. The various components of the w_u^k vectors may be interpreted for instance as corresponding to various possible sources of uncertainty (weather conditions, measure of congestion of arc u on a transportation network, etc) and, for a given realization l_u^k of l_u in L_u, the i^{th} component of w_u^k is 1 (or more generally a positive integer) if the i^{th} source of uncertainty is a factor contributing to the outcome $l_u = l_u^k$, 0 otherwise. In addition to the above, a global nonnegative p-component vector B is given ("uncertainty budget"), the uncertainty set for the arc lengths being defined as the set of $l = (l_u)_{u \in U}$ of the form $(l_1^{k_1}, l_2^{k_2}, \ldots, l_m^{k_m})$ where $\sum_{i=1}^{m} w_i^{k_i} \leq B$.

We note that such a model is sufficiently general to handle situations for which a given realization $l_u = l_u^k$ is the result of joint influences of several distinct sources of uncertainty: the corresponding w_u^k vector will have several components equal to 1 (or, more generally, non zero). Also we note that in most applications, the largest among the l_u^k values in L_u (those related to the most unfavorable situations with respect to finding the shortest path solution) will tend to correspond to the w_u^k vectors of largest "weight" (as measured e.g. in terms of number of non zero components, or in terms of L_1 norm).

Also worth mentioning is another special case of the above general model, potentially useful in applications, where there are several (p) independent sources of uncertainty acting on *disjoint subsets of arcs* $U_1, U_2, \ldots, U_p$. Then, for each $u \in U_i$, the corresponding w_u^k vectors will have all components 0 except the i^{th} component which can be non zero. In such a case the global uncertainty budget constraint just

Table 7

	$u \rightarrow$	0.1	0.2	0.4	0.6	0.8	1
P2	$B = \binom{5}{5}$	1584 (30)	1423 (30)	1376 (17)	1358 (21)	1358 (21)	1358 (21)
	$B = \binom{8}{8}$	1482 (23)	1362 (30)	1300 (19)	1259 (16)	1259 (16)	1259 (16)
	$B = \binom{12}{12}$	1400 (20)	1306 (18)	1225 (21)	1161 (16)	1159 (16)	1159 (16)
P6	$B = \binom{5}{5}$	1473 (13)	1406 (22)	1368 (19)	1352 (16)	1352 (16)	1352 (16)
	$B = \binom{8}{8}$	1412 (13)	1345 (22)	1279 (22)	1252 (20)	1252 (20)	1252 (20)
	$B = \binom{12}{12}$	1358 (18)	1301 (20)	1203 (17)	1132 (13)	1125 (13)	1124 (13)
P7	$B = \binom{5}{5}$	1162 (17)	1128 (22)	1071 (16)	1059 (17)	1058 (18)	1058 (18)
	$B = \binom{8}{8}$	1122 (18)	1077 (17)	1000 (22)	969 (16)	966 (16)	965 (16)
	$B = \binom{12}{12}$	1082 (17)	1029 (17)	935 (16)	874 (16)	864 (16)	858 (16)
P9	$B = \binom{5}{5}$	986 (11)	958 (16)	910 (18)	909 (18)	908 (17)	908 (17)
	$B = \binom{8}{8}$	945 (12)	902 (15)	832 (16)	816 (18)	815 (18)	814 (18)
	$B = \binom{12}{12}$	889 (15)	844 (19)	761 (16)	716 (16)	708 (16)	702 (17)
P10	$B = \binom{5}{5}$	1037 (13)	1017 (13)	955 (12)	923 (12)	923 (12)	923 (12)
	$B = \binom{8}{8}$	985 (14)	948 (13)	883 (13)	822 (16)	820 (17)	817 (16)
	$B = \binom{12}{12}$	938 (14)	886 (14)	806 (14)	739 (14)	727 (14)	716 (13)
P12	$B = \binom{5}{5}$	1052 (30)	1025 (30)	980 (26)	975 (26)	968 (26)	968 (26)
	$B = \binom{8}{8}$	978 (30)	948 (30)	873 (20)	857 (21)	849 (21)	849 (21)
	$B = \binom{12}{12}$	904 (19)	868 (25)	784 (27)	745 (18)	735 (18)	733 (18)

Results obtained for various values of the parameter u controlling the intensity of the dependence between demands and procurement costs.

amounts to imposing one uncertainty budget constraint for each subset of arcs separately.

Once defined the uncertainty set on the arc lengths as explained above, the problem consists in determining a best possible (robust) strategy for choosing paths in the graph from each node i to the terminal node n in such a way that the worst possible lengths of these paths (over the set of all possible eventual realizations of uncertainty) is minimized. This problem is easily recognized as a special case of the general model presented in Sections 2. to 4. above ; indeed, identifying the nodes of the circuitless graph G with the states of a dynamic system, this corresponds to the case where there is no uncertainty on the state transition function (uncertainty only influences the ”return” function). From this, we deduce that the robust shortest path problem can be solved by the following recursion in which the nodes of the graph are examined according to a reverse topological ordering, starting by assigning shortest path values equal to 0, to node n

(whatever its uncertainty status) and then assigning shortest path values to nodes $n-1, n-2, \ldots$ etc., for each possible uncertainty status $\sigma \leq B$.

$$z^*(n,\sigma) = 0, \text{ for all } \sigma \leq B \tag{12}$$

and then, for $i = n-1, n-2, \ldots, 1$:

$$z^*(i,\sigma) = \min_{(i,j)\in\omega^+(i)} \left\{ \max_{\substack{k=1\cdots\nu \\ \text{s.t.} \\ w_{ij}^k \leq \sigma}} \left\{ l_{ij}^k + z^*(j, \sigma - w_{ij}^k) \right\} \right\}. \tag{13}$$

$$\varphi^*(i,\sigma) = \operatorname*{argmin}_{(i,j)\in\omega^+(i)} \left\{ \max_{\substack{k=1\cdots\nu \\ \text{s.t.} \\ w_j^k \leq \sigma}} \left\{ l_{ij}^k + z^*(j, \sigma - w_{ij}^k) \right\} \right\}. \tag{14}$$

At each step of the recursion, $\varphi^*(i,\sigma)$ denotes the best arc to take to leave node i, given that the corresponding uncertainty status of node i is σ (due to the observed realizations of uncertainty on the path from node 1 to i).

Let us illustrate the recursion (12)-(14) on the simple example graph shown on Figure 2 below (note that the nodes are numbered according to a topological order). We take $B = \begin{pmatrix}1\\1\end{pmatrix}$. Therefore, there are 4 possible distinct values for σ in this example: $\begin{pmatrix}0\\0\end{pmatrix} \begin{pmatrix}1\\0\end{pmatrix} \begin{pmatrix}0\\1\end{pmatrix} \begin{pmatrix}1\\1\end{pmatrix}$.

The application of recursion (12)-(14) then provides the $z^*(i,\sigma)$ values and corresponding $\varphi^*(i,\sigma)$ values displayed in Table 9.

It is seen that, under the given uncertainty model, the optimal strategy leads to the decision of leaving node 1 via arc (1,3). Now, in node 3, the decision to take depends on which realization of uncertainty is actually observed while traversing arc (1,3). If the nominal length 8 was observed, we are in node 3 with uncertainty status $\begin{pmatrix}1\\1\end{pmatrix}$ and the optimal decision which should be taken then is to use arc (3,4). If the extreme-high value 11 was observed, we are at node 3 with status $\begin{pmatrix}1\\0\end{pmatrix}$ and then we should leave node 3 using arc (3,7). In the former case the path followed $(1 \rightarrow 3 \rightarrow 4 \rightarrow 5 \rightarrow 7)$ has worst-case length 28 ; in the latter case the path followed $(1 \rightarrow 3 \rightarrow 7)$ has worst case length 26. The figure $z^*\left(1, \begin{pmatrix}1\\1\end{pmatrix}\right) = 28$ shown in Table 5 corresponds to the worst case between the two alternatives (since, when starting at node 1 we do not know in advance, which length will actually be realized on arc (1,3)).

This example thus illustrates the fact that the solution produced by recursion (12)-(14) is not a well-defined path between origin and destination in the graph, but a set of optimal paths defining a *strategy* (closed-loop solution).

It also illustrates the fact that the robust shortest path model investigated here appears to be significantly different from the one proposed in [3], since in the latter case, an “open-loop” solution is looked for instead of a ’closed-loop’ solution.

According to such a strategy the optimal path to be followed depends on the information about uncertainty collected during the graph traversal process itself.

6. Conclusions

A class of multistage robust decision problems has been investigated in connection with a special type of uncertainty model referred to here as the *parsimonious uncertainty model*. To the best of our knowledge, this is the first time such a way of modeling uncertainty is proposed in the context of robust dynamic programming problems. In particular, it has been shown that a key interest of such a model is to provide compact representations of potentially huge scenario trees, leading to an efficient (pseudopolynomial) dynamic-programming-based algorithm for computing optimal strategies (“closed-loop” solutions). From the point-of-view of applications, the uncertainty model proposed here has been shown to offer modelling flexibility in various ways: (a) it is capable of representing, in any given time period, dependence among several parameters influenced by uncertainty (refer to the example of prices and requirements as discussed in §5.1.2.); (b) it is designed to take into account, via the uncertainty budget constraints, dependence among uncertain events occurring in *different time periods*; (c) by varying the components of the right hand side of the uncertainty budget constraints, and by exploiting the intermediate results of the dynamic programming recursion, it can be used to generate a variety of more or less conservative solutions featuring various robustness levels.

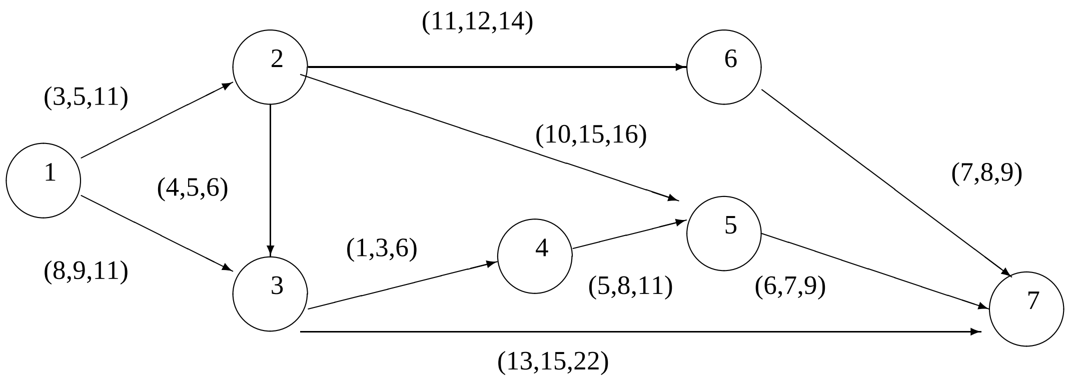

Fig. 2. A 7-node graph to illustrate the robust shortest path computation. On each arc there are 3 possible values of the length (l_u^1, l_u^2, l_u^3) corresponding to the uncertainty profiles $w_u^1 = \begin{pmatrix}0\\0\end{pmatrix}$ $w_u^2 = \begin{pmatrix}1\\0\end{pmatrix}$ $w_u^3 = \begin{pmatrix}0\\1\end{pmatrix}$. The uncertainty budget is $B = \begin{pmatrix}1\\1\end{pmatrix}$.

Table 9

	$z^*\left(i, \begin{pmatrix}0\\0\end{pmatrix}\right)$	$z^*\left(i, \begin{pmatrix}1\\0\end{pmatrix}\right)$	$z^*\left(i, \begin{pmatrix}0\\1\end{pmatrix}\right)$	$z^*\left(i, \begin{pmatrix}1\\1\end{pmatrix}\right)$
$i = 7$	0 (-)	0 (-)	0 (-)	0 (-)
$i = 6$	7 (6,7)	8 (6,7)	9 (6,7)	9 (6,7)
$i = 5$	6 (5,7)	7 (5,7)	9 (5,7)	9 (5,7)
$i = 4$	11 (4,5)	14 (4,5)	17 (4,5)	18 (4,5)
$i = 3$	12 (3,4)	15 (3,7)	18 (3,4)	20 (3,4)
$i = 2$	16 (2,5)	19 (2,6)	21 (2,6)	22 (2,6)
$i = 1$	19 (1,2)	22 (1,2)	26 (1,3)	28 (1,3)

The optimal decision rules obtained for the example in Figure 1. For each pair (i, σ) the value $z^*(i, \sigma)$ is displayed, followed by the arc corresponding to $\varphi^*(i, \sigma)$.

As an additional interesting aspect of our model, we mention the fact that *it is by no means exacting in terms of input data*. For instance it does not require from the decision maker a precise knowledge of a huge number of (possibly multidimensional) probability distributions (such information is rarely at hand when dealing with applications involving uncertainty). On the contrary, it only requires much coarser and sparser information on the uncertain parameters, typically: maximum and minimum observable value, maximum number of occurrences of extreme (worst-case) values over the period of study. Clearly, assuming availability of such information appears to be much more realistic in many situations. For all the above reasons, the variety of applications which might be addressed via the model and solution approach proposed here appears to be potentially huge and this will be the subject of future research work.

Acknowledgements

An anonymous referee is gratefully acknowledged for helpful remarks and constructive comments.

References

[1] Ben-Tal A., Nemirovski A. (1999). "Robust Solutions to Uncertain Programs". Oper. Res. Letters, 25, pp. 1-13.
[2] Bertsekas D.P. (2005). "Dynamic Programming and Optimal Control". Athena Scientific, Vol. 1.
[3] Bertsimas D., Sim M. (2003). "Robust Discrete Optimization and Network Flows". Math.. Prog. B, 98, pp. 49-71.
[4] Bertsimas D., Sim M. (2004). "The price of Robustness". Oper. Res. 52, 1, pp. 35-53.
[5] Birge J.R. (1985). "Decomposition and Partitioning Methods for Multistage Stochastic Linear Programs". Oper. Res. 33, 5, pp. 989-1007.
[6] Birge J.R. (1997). " Stochastic Programming Computation and Applications". INFORMS J. on Computing, 9, 2, pp. 111-133.
[7] Birge J.R., Louveaux F.V. (1997). " Introduction to Stochastic Programming ". Springer Verlag, New York.
[8] Charnes A., Cooper W.W. (1959). "Chance Constrained Programming". Management Sciences 5, pp. 73-79.
[9] Epstein L.G., Schneider M. (2001). "Recursive Multiple Priors". Journal Econom. Theory, 113, 1, pp. 1-31.
[10] Iniguchi M., Sakawa M. (1995). "Minimax Regret Solution to Linear Programming Problems with Interval Objective Function". EJOR, 86, 3, pp. 526-536.
[11] Iyengar G.N. (2005). "Robust Dynamic Programming" Maths. of Operations Research, 30, 2, pp. 257-280.
[12] Kouvelis P., Yu G. (1997). "Robust Discrete Optimization and its Applications". Kluwer Academic Publisher, Boston.
[13] Minoux M. (2007). "Duality, Robustness and 2-Stage Robust LP Decision Models". Annales du Lamsade, Université Paris-Dauphine, France.
[14] Mulvey J.M., Vanderbei R.J., Zenios S.A. (1995). "Robust Optimization of Large Scale Systems". Oper. Res. 43, 2, pp. 264-281.
[15] Nilim A., El Ghaoui L. (2005). "Robust Control of Markov Decision Processes with Uncertain Transition Matrices". Operations Research, 53, 5, pp. 780-798.
[16] Ross S.M. (1983). "Introduction to Stochastic Dynamic Programming". Academic Press, New York.
[17] Schultz R., Tiedemann S. (2004). "Conditional Value-at-Risk in Stochastic programs with Mixed-Integer Recourse". Research report, University of Duisburg-Essen, Germany.
[18] Soyster A.L. (1973). "Convex Programming with Set-Inclusive Constraints and Applications to Inexact Linear Programming". Oper. Res. 21, pp. 1154-1157.
[19] Soyster A.L. (1979). " Inexact Linear Programming with Generalized Resource Sets". EJOR, 3, pp. 316-321.
[20] Van Slyke R., Wets R.J. (1969). "L-Shaped Linear Programs with Applications to Optimal Control and Stochastic Programming". SIAM J. on Appl. Math. 17, pp. 638-663.
[21] Yu G. (1996). "On the Max-Min 0-1 Knapsack Problem with Robust Optimization Applications". Operations Research, 44, 2, pp. 407-415.
[22] Yu G., Yang J. (1998). "On the Robust Shortest Path Problem". Computers and Oper. Res. 25, 6, pp. 457-468.

Received 7 Feb 07; revised 11 Feb 08; accepted 24 Feb 08