Tilburg University

Efficient Methods for Several Classes of Ambiguous Stochastic Programming Problems under Mean-MAD Information

Postek, Krzysztof; Romeijnnders, Ward; den Hertog, Dick; van der Vlerk, Maartne H.

Publication date:
2016

Document Version
Early version, also known as pre-print

[Link to publication in Tilburg University Research Portal](#)

Citation for published version (APA):

Postek, K., Romeijnnders, W., den Hertog, D., & van der Vlerk, M. H. (2016). *Efficient Methods for Several Classes of Ambiguous Stochastic Programming Problems under Mean-MAD Information*. (CentER Discussion Paper; Vol. 2016-039). Operations research.

General rights

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Take down policy

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

No. 2016-039

**EFFICIENT METHODS FOR SEVERAL CLASSES OF
AMBIGUOUS STOCHASTIC PROGRAMMING PROBLEMS
UNDER MEAN-MAD INFORMATION**

By

Krzysztof Postek, Ward Romeijnders, Dick den Hertog,
Maarten H. van der Vlerk

22 September 2016

ISSN 0924-7815
ISSN 2213-9532

Efficient methods for several classes of ambiguous stochastic programming problems under mean-MAD information

Krzysztof Postek

CentER and Department of Econometrics and Operations Research, Tilburg University, k.postek@tilburguniversity.edu

Ward Romeijnders

Department of Operations, Faculty of Economics and Business, University of Groningen, w.romeijnders@rug.nl

Dick den Hertog

CentER and Department of Econometrics and Operations Research, Tilburg University, d.denhertog@tilburguniversity.edu

Maarten H. van der Vlerk

Department of Operations, Faculty of Economics and Business, University of Groningen, m.h.van.der.vlerk@rug.nl

We consider decision making problems under uncertainty, assuming that only partial distributional information is available - as is often the case in practice. In such problems, the goal is to determine here-and-now decisions, which optimally balance deterministic immediate costs and worst-case expected future costs. These problems are challenging, since the worst-case distribution needs to be determined while the underlying problem is already a difficult multistage recourse problem. Moreover, as found in many applications, the model may contain integer variables in some or all stages.

Applying a well-known result by Ben-Tal and Hochman (1972), we are able to efficiently solve such problems without integer variables, assuming only readily available distributional information on means and mean-absolute deviations. Moreover, we extend the result to the non-convex integer setting by means of convex approximations (see Romeijnders et al. (2016a)), proving corresponding performance bounds. Our approach is straightforward to implement using of-the-shelf software as illustrated in our numerical experiments.

Key words: robust; ambiguous; integer; recourse; stochastic; multi-stage

1. Introduction

In practice, many decision makers are faced with uncertainty in some parameters of their model. Consider, for example, customer demand in production planning, supply of renewable energy in unit commitment problems, the precision of physical devices in engineering design, and the return on investment in finance. In these problems, the information on the uncertain parameters, based on e.g. historical data or expert opinions, can either be limited or extensive. In the latter case, uncertainty is typically modelled using random parameters with known probability distributions, whereas in the former case the distributions are *ambiguous*, i.e. only partly known (Knight 1921). In this paper, we address mathematical optimization problems where (some of) the parameter

distributions are ambiguous. One of our main contributions is that for such problems we provide efficient solution methods that are easy to implement using off-the-shelf software.

If the uncertain random parameters are revealed gradually over time, then we can model the decision problem as a multi-stage stochastic programming (SP) problem (see, e.g., the textbooks Birge and Louveaux (1997), Prékopa (1995), Shapiro et al. (2009)) in which the planning horizon consists of multiple time stages. Under the assumption that the probability distributions of the uncertain random parameters are known, the problem is to determine so-called here-and-now decisions implemented before (some of) these uncertain parameters are revealed and new decisions have to be made. This process repeats itself over several stages and the objective is to minimize the sum of the here-and-now costs and the *expected* future costs, taking the distributions of the uncertain random parameters and the decisions in later stages into account.

We, however, assume in line with practical experience that only limited information on the distributions of the uncertain random parameters is available, and instead of minimizing the expected costs, we take a *distributionally robust* (ambiguous) approach and minimize the *worst-case expected costs* over all possible (or admissible) probability distributions.

Early contributions to the minimization of worst-case expectations are the works of Scarf (1958) and Žáčková (1966). In the SP literature it has been referred to as *minimax* problems, and is mainly considered within the framework of generalized moment problems (Kemperman 1968), where the objective is to determine the worst-case expectation of a given function under conditions on some (generalized) moments of the uncertain random parameters. For a discussion on cases in which exact bounds or approximation procedures are available, we refer to Edirisinghe (2011).

Moreover, Shapiro and Kleywegt (2002) and Shapiro and Ahmed (2004) use duality to show that distributionally robust SP problems can be equivalently reformulated as a standard SP problem in which the probability distributions of the random parameters are known. The difference with our approach is that we can utilize the explicitly known worst-case distributions under the given distributional information, instead of using duality.

Recently, distributional uncertainty has gained the attention of the Robust Optimization (RO) community. They treat the sets of admissible probability distributions as uncertainty sets and use conic duality (see the references above and also, e.g., Isii (1963), Shapiro (2001), and Ben-Tal et al. (2015)) to derive equivalent, computationally tractable forms of constraints on the worst-case expectation. Prominent examples of this approach are the papers of Delage and Ye (2010) and Wiesemann et al. (2014), who provide also good surveys of the existing approaches.

Despite these developments, solving distributionally robust SP problems remains challenging for several reasons. First, it may be hard to determine the worst-case probability distribution (maximizing the expected costs), even for given here-and-now decisions. Second, it may be computationally

intractable to determine optimal here-and-now decisions taking into account all decisions in future time stages under all possible realizations of the uncertain parameters. Third, the problem may contain integer decision variables, and fourth, solving the problem may require special purpose algorithms that are not available in standard software packages.

In this paper we overcome these four challenges for a large class of distributionally robust SP problems. In the remainder of this introduction we discuss each of these major challenges separately.

In our setting, the first challenge (evaluating the worst-case expectation) is void because of the particular distributional information we use — the supports, means, mean absolute deviations from the means (MADs), and the probability that a given variable is at least equal to its mean, which are easy to estimate using e.g. the procedures given in Postek et al. (2015). Under such distributional information, we can use a result of Ben-Tal and Hochman (1972, referred to as BTH72 below), to prove that the worst-case and best-case marginal distributions are discrete with at most three possible realizations if the distributionally robust SP problem only contains continuous decision variables. The well-known Edmundson-Madansky upper bound (Edmundson 1956) and Jensen lower bound (Jensen 1906) are similar in spirit to the results we use. The main difference between our results and the work of e.g. Shapiro and Kleywegt (2002) and Shapiro and Ahmed (2004) is that in our case we obtain simple worst-case and best-case distributions that are the same for every here-and-now decision. This indeed simplifies matters considerably: the problem becomes much more tractable and from the practical perspective it is much more intuitive that the distributions do not depend on the (initial) decisions. Our approach is applicable only if the random parameters are stochastically independent. This is obviously a restrictive assumption but an advantage of our approach is that it can exploit this property.

The difference between the worst-case and best-case expectations of the objective gives an easy-to-calculate upper bound on the *value of distributional information* (VDI); see e.g. Delage et al. (2015). The VDI is related to the value of the stochastic solution (VSS) introduced by Birge (1982). However, the VSS measures the added value of solving the stochastic problem instead of its deterministic version, whereas the VDI measures the added value of (or the willingness to pay for) knowing the probability distributions of the random parameters. The VDI is particularly relevant in a data-driven environment where it can be used to assess the costs of gathering more data.

If all worst-case distributions are discrete with three possible realizations (as we will find), then the distributionally robust SP problem reduces to a standard SP problem with 3^n realizations (or scenarios) of the joint distribution of the n random parameters. This exponential number of scenarios underlies the second major challenge (computational tractability): it is the reason why the problem is computationally intractable from an RO point of view. For this reason, the problem is often approximated by imposing *decision rules* on future decisions as a function of the

revealed random parameters; see Garstka and Wets (1974) for the first contribution in the SP literature. In the RO literature Ben-Tal et al. (2004) first formulated the decision rules as *affine functions* of the revealed parameters, and their approach has been extended to other function classes by e.g. Chen and Zhang (2009), Ben-Tal et al. (2009) and Bertsimas et al. (2011). From an SP perspective, however, dealing with 3^n scenarios is not unusual and there exist many solution methods to (approximately) solve such problems. In Section 4 we give a brief overview of these methods. In particular, in Appendix B we present a particularly efficient implementation of such methods tailored to the problem at hand.

The third major challenge (inclusion of integer variables) is relevant since many decision problems require integer decision variables to be modelled realistically. Consider e.g. unit commitment decisions in electric power generation (see e.g. Römisch and Schultz (1996), Bertsimas et al. (2013) and many others) or lot sizing decisions in inventory control (see, e.g., Postek and den Hertog (2016)). Within the SP literature, stochastic mixed-integer programming (SMIP) problems have been studied by e.g. Laporte and Louveaux (1993), Carøe and Schultz (1999), and Ahmed et al. (2004), (see also the surveys by Schultz (2003), Klein Haneveld and van der Vlerk (1999), and Sen (2005)), while in the RO literature systematic approaches have been developed to allow for integer decision variables in future time stages; see e.g. Bertsimas and Georghiou (2015), Hanasusanto et al. (2015), and Postek and den Hertog (2016).

For SMIP problems the main difficulty is that due to the integer variables in future time stages, the optimal objective value is generally not convex in the uncertain parameters. For this reason, van der Vlerk (2004), Klein Haneveld et al. (2006), Romeijnnders et al. (2015), and Romeijnnders et al. (2016b) have proposed convex approximations for several classes of SMIP problems. For these approximations error bounds have been derived that depend on the total variations of the probability density functions of the random parameters in the model. We use the idea of convex approximations to provide a framework for solving a large class of two-stage distributionally robust SMIP problems in which the distributions of some random parameters are known and others are ambiguous. We derive error bounds for two approximations of which one is obtained by (incorrectly) assuming that the worst-case distributions are the same as in the continuous distributionally robust SP case, i.e., assuming convexity. In Section 5 we apply this framework to an operating room scheduling problem.

In that section we carry out numerical experiments on an inventory control problem as well. Dealing with the fourth major challenge (ease of implementation), we show that we can obtain good solutions using off-the-shelf software, despite the exponential number of scenarios. Moreover, we show that for problems of realistic size we may obtain exact optimal solutions, using the specific

structure of the problem to speed up existing algorithms. Furthermore, we provide additional managerial insights (i) by calculating the VDI, (ii) by graphically depicting the so-called *Pareto-stripe*, an extension of the Pareto curve, which shows the tradeoff between various types of objectives, and (iii) by comparing various approaches from the SP and RO literature.

To summarize, we provide a framework for solving distributionally robust SP problems, satisfying the following properties:

- (1) the required parameters of the independent probability distributions in the ambiguity set can be estimated from data;
- (2) there is a simple worst-case distribution that is the same for all here-and-now decisions;
- (3) future decisions depend on the observed values of the uncertain parameters;
- (4) the solution method is able to accommodate for integer decision variables in two-stage problems;
- (5) the value of distributional information can be quantified;
- (6) the solution method is easy to implement using off-the-shelf software and known SP techniques.

The structure of our paper is as follows. In Section 2.1 we introduce our approach for two-stage continuous problems and we extend it to the multi-stage setting in Section 2.2. Section 3 includes our new theoretical results on convex approximations of two-stage stochastic programs with integer decision variables. In Section 4 we discuss general techniques helpful in dealing with the number of scenarios which grows exponentially with the dimension of the vector of uncertain parameters. In Section 5 we present three numerical experiments involving operating room planning and inventory management. Each of the experiments illustrates our approach for a particular class of distributionally robust SP problems: two-stage problems with continuous and with integer variables, and a continuous multi-stage problem.

2. Distributionally robust SP problems

In this section we describe our approach for solving distributionally robust SP problems in case all decision variables are continuous. For ease of exposition, we first consider two-stage problems in Section 2.1; multi-stage problems are discussed in Section 2.2. Although the results in this section appear to be known in the SP literature (Ben-Tal and Hochman 1976), we are the first — to our knowledge — to make these results explicit in a multi-stage setting.

2.1. Two-stage problems

The distributionally robust SP problem that we consider is

$$\inf_{x \in X} \sup_{\mathbb{P}_z \in \mathcal{P}_z} \mathbb{E}_{\mathbb{P}_z} [c^\top x + v(x, z)], \quad (1)$$

where $X = \{x \in \mathbb{R}_+^{n_1} : Ax = b\}$ represents the set of feasible first-stage solutions, $\mathcal{P}_z$ is the ambiguity set for probability distributions, and $v(x, z)$ is the second-stage value function defined as a function of the first-stage variables x and the random parameters $z = (\xi, \omega)$:

$$v(x, z) = \inf_{y \in Y} \left\{ q(\xi)^\top y : Wy = h(\omega) - T(\omega)x \right\}.$$

Here, y are the second-stage (or recourse) variables and $Y \subset \mathbb{R}_+^{n_2}$ is a polyhedral set. The second-stage costs $q(\xi)$, the technology matrix $T(\omega)$, and the right-hand side $h(\omega)$ depend on the random vector $z = (\xi, \omega)$. We assume that q , T , and h are affine functions of z and that *all components of z are independent*. Thus, in particular, $q(\xi)$ is independent from $T(\omega)$ and $h(\omega)$. Moreover, since the recourse matrix W is deterministic, we say that the problem has *fixed recourse* (see, e.g., Shapiro et al. (2009)).

In problem (1), the here-and-now decisions x have to be made while the parameter z are unknown, and after the uncertain parameter z is revealed we are allowed to take recourse actions y to compensate for possible violations of the constraints $T(\omega)x = h(\omega)$. The objective is to minimize the sum of the direct costs $c^\top x$ and the worst-case expected costs $\sup_{\mathbb{P}_z \in \mathcal{P}_z} \mathbb{E}_{\mathbb{P}_z}[v(x, z)]$.

Here, the ambiguity set $\mathcal{P}_z$ is defined as

$$\mathcal{P}_z = \left\{ \mathbb{P}_z : \text{supp}(z_i) \subseteq [a_i, b_i], \quad \mathbb{E}_{\mathbb{P}_z}[z_i] = \mu_i, \quad \mathbb{E}_{\mathbb{P}_z}[|z_i - \mu_i|] = d_i, \quad \mathbb{P}_z\{z_i \geq \mu_i\} = \beta_i, \quad z_i \perp z_j, i \neq j \right\}, \quad (2)$$

where $z_i \perp z_j$ means that z_i and z_j are stochastically independent. Postek et al. (2015) explain procedures to estimate these parameters from historical data and the conditions on a, b, μ, d, β such that $\mathcal{P}_z$ is non-empty. Throughout this paper we refer to the ambiguity set $\mathcal{P}_z$ in (2) as a (μ, d, β) ambiguity set.

EXAMPLE 1. Consider an inventory manager who needs to order a specific amount x of products. Later, when the uncertain customer demand z is known, he can order an additional amount y of the products, however, at unknown but likely higher prices. The objective of the manager is to minimize the expected total cost. However, due to lack of knowledge on the true distribution, he chooses the ‘safe option’ and minimizes the worst-case expected cost over the set $\mathcal{P}_z$ of distributions that can be ‘true’ based on the data.

2.1.1. Worst-case expectation Problem (1) is difficult to solve because the worst-case probability distribution $\mathbb{P}_z \in \mathcal{P}_z$ may depend on the first-stage decision $x \in X$, and we need to optimize over x . However, for the (μ, d, β) ambiguity set $\mathcal{P}_z$ in (2), the worst-case distribution $\mathbb{P}_{\bar{z}}$ turns out to be the same for every first-stage decision so that the distributionally robust SP problem in (1) reduces to

$$\inf_{x \in X} \mathbb{E}_{\mathbb{P}_{\bar{z}}}[c^\top x + v(x, \bar{z})],$$

where each component of $\bar{z}$ follows a known discrete distribution with at most three realizations. This result is summarized in Proposition 1 below. Its proof combines the fact that the second-stage value function $v(x, z)$ is convex in ω and concave in ξ (see, e.g., Fiacco and Kyparisis 1986) with results from BTH72, who provide closed-form expressions for the worst-case expectations maximizing and minimizing the expectations of convex and concave functions.

Proposition 1 *The two-stage distributionally robust SP problem*

$$\inf_{x \in X} \sup_{\mathbb{P}_z \in \mathcal{P}_z} \mathbb{E}_{\mathbb{P}_z} \left[c^\top x + \inf_{y \in Y} \left\{ q(\xi)^\top y : Wy = h(\omega) - T(\omega)x \right\} \right]$$

with (μ, d, β) ambiguity set $\mathcal{P}_z$ for $z = (\xi, \omega) \in \mathbb{R}^{n_\xi} \times \mathbb{R}^{n_\omega}$ as defined in (2) is equivalent to

$$\inf_{x \in X} \mathbb{E}_{\bar{\mathbb{P}}_{\bar{z}}} \left[c^\top x + \inf_{y \in Y} \left\{ q(\bar{\xi})^\top y : Wy = h(\bar{\omega}) - T(\bar{\omega})x \right\} \right],$$

where the worst-case random vector $\bar{z} = (\bar{\xi}, \bar{\omega}) \in \mathbb{R}^{n_\xi} \times \mathbb{R}^{n_\omega}$ has independent components with marginal distributions

$$\mathbb{P} \left\{ \bar{\xi}_i = \mu_i - \frac{d_i}{2(1 - \beta_i)} \right\} = 1 - \beta_i, \quad \text{and} \quad \mathbb{P} \left\{ \bar{\xi}_i = \mu_i + \frac{d_i}{2\beta_i} \right\} = \beta_i, \quad i = 1, \dots, n_\xi,$$

and

$$\begin{aligned} \mathbb{P} \left\{ \bar{\omega}_i = a_{n_\xi+i} \right\} &= \frac{d_{n_\xi+i}}{2(\mu_{n_\xi+i} - a_{n_\xi+i})}, & \mathbb{P} \left\{ \bar{\omega}_i = b_{n_\xi+i} \right\} &= \frac{d_{n_\xi+i}}{2(b_{n_\xi+i} - \mu_{n_\xi+i})}, \\ \mathbb{P} \left\{ \bar{\omega}_i = \mu_{n_\xi+i} \right\} &= 1 - \frac{d_{n_\xi+i}}{2(\mu_{n_\xi+i} - a_{n_\xi+i})} - \frac{d_{n_\xi+i}}{2(b_{n_\xi+i} - \mu_{n_\xi+i})} \end{aligned}$$

for $i = 1, \dots, n_\omega$.

Proof. See Appendix A. □

Since the worst-case distribution $\mathbb{P}_{\bar{z}}$ has finite support, we can enumerate all $K := 2^{n_\xi} \times 3^{n_\omega}$ scenarios of $\bar{z}$ and rewrite the distributionally robust SP problem in (1) as

$$\begin{aligned} \inf_{x \in X} \sup_{\mathbb{P}_z \in \mathcal{P}_z} \mathbb{E}_{\mathbb{P}_z} [c^\top x + v(x, z)] &= \inf_{x \in X} \mathbb{E}_{\mathbb{P}_{\bar{z}}} [c^\top x + v(x, \bar{z})] \\ &= \inf_{x \in X} \sum_{k=1}^K p_k [c^\top x + v(x, \bar{z}^k)], \end{aligned}$$

where p_k denotes the probability of scenario $\bar{z}^k$, $k = 1, \dots, K$. The latter problem can be rewritten in its deterministic equivalent form, yielding

$$\inf_{x \in X, y_k \in Y} \left\{ c^\top x + \sum_{k=1}^K p_k q(\bar{\xi}^k) y_k : Wy_k = h(\bar{\omega}^k) - T(\bar{\omega}^k)x, k = 1, \dots, K \right\}.$$

REMARK 1. Note that the worst-case expectation of the random vector ω does not require information on parameter β , i.e. in case one deals with uncertainty in the constraints of the problems only, it suffices to estimate the parameter a , b , μ , and d of the probability distribution ω . This means that having the knowledge on β does not change the worst-case expectation value. An even more striking fact is that if β is known, the three point worst-case distribution of ω may not satisfy this probability bound, i.e., it may hold that $\mathbb{P}(\bar{\omega}_i \geq \mu_{n_\xi+i}) < \beta_{n_\xi+i}$ for some $1 \leq i \leq n_\omega$. This is because the worst-case probability bound is tight but it need not be attained, see BTH72.

2.1.2. Best-case expectation Similar as for the worst-case expectation we can obtain the best-case expectation over all probability distributions in the (μ, d, β) ambiguity set $\mathcal{P}_z$ by using results of BTH72. Again, the best-case distribution P_z is a discrete distribution with at most three realizations per component that does not depend on the first-stage decision x .

Proposition 2 *The two-stage distributionally robust SP problem*

$$\inf_{x \in X} \inf_{\mathbb{P}_z \in \mathcal{P}_z} \mathbb{E}_{\mathbb{P}_z} \left[c^\top x + \inf_{y \in Y} \left\{ q(\xi)^\top y : Wy = h(\omega) - T(\omega)x \right\} \right]$$

with (μ, d, β) ambiguity set $\mathcal{P}_z$ for $z = (\xi, \omega) \in \mathbb{R}^{n_\xi} \times \mathbb{R}^{n_\omega}$ as defined in (2) is equivalent to

$$\inf_{x \in X} \mathbb{E}_{\mathbb{P}_z} \left[c^\top x + \inf_{y \in Y} \left\{ q(\xi)^\top y : Wy = h(\omega) - T(\omega)x \right\} \right], \quad (3)$$

where the best-case random vector $z = (\xi, \omega) \in \mathbb{R}^{n_\xi} \times \mathbb{R}^{n_\omega}$ has independent components with marginal distributions

$$\mathbb{P}\{\xi_i = a_i\} = \frac{d_i}{2(\mu_i - a_i)}, \quad \mathbb{P}\{\xi_i = b_i\} = \frac{d_i}{2(b_i - \mu_i)}, \quad \mathbb{P}\{\xi_i = \mu_i\} = 1 - \frac{d_i}{2(\mu_i - a_i)} - \frac{d_i}{2(b_i - \mu_i)}$$

for $i = 1, \dots, n_\xi$ and

$$\mathbb{P}\left\{\omega_i = \mu_{n_\xi+i} - \frac{d_{n_\xi+i}}{2(1 - \beta_{n_\xi+i})}\right\} = 1 - \beta_{n_\xi+i}, \quad \mathbb{P}\left\{\omega_i = \mu_{n_\xi+i} + \frac{d_{n_\xi+i}}{2\beta_{n_\xi+i}}\right\} = \beta_{n_\xi+i}.$$

for $i = 1, \dots, n_\omega$.

Proof. See Appendix A. □

Notice that since $v(x, z)$ is concave in ξ and convex in ω the worst-case distribution of ξ has the same structure as the best-case distribution of ω and vice versa. Moreover, we can derive the deterministic equivalent form of (3) analogous to that of the worst-case expectation.

Best-case expectation is a useful complement to the worst-case expectation since they bound the actual expected costs in the stochastic program, which is unknown since the probability distribution of the random vector of parameters z is unknown. The difference between the worst-case and best-case expectation can be interpreted as an upper bound on the *value of distributional information* (VDI, Delage et al. (2015)), i.e. the amount we are willing to pay for complete knowledge of the probability distribution of z . We illustrate this concept in the numerical experiments of Section 5.

2.2. Multi-stage problems

We consider now the general multi-stage linear problem. For ease of exposition we limit ourselves to the uncertainty in the constraints driven by random vector $z = \omega$. The results, however, extend easily to the case including also uncertainty in the cost vector driven by a random vector ξ , as in the two-stage problem (1).

Here, $x_t \in \mathbb{R}_+^{n_t}$ denote the decision vectors implemented at time $t = 1, 2, \dots, T$. The uncertain parameter $z \in \mathbb{R}^{n_z}$ has a corresponding structure $z = (z_1, \dots, z_{T-1})$, $z_t \in \mathbb{R}^{n_{z,t}}$ for $t = 1, \dots, T-1$, with $n_z = \sum_{t=1}^{T-1} n_{z,t}$. The time sequence of decisions and uncertainty revealing is

$$x_1 \rightarrow z_1 \rightarrow x_2 \rightarrow z_2 \rightarrow \dots \rightarrow z_{T-1} \rightarrow x_T.$$

Since all random parameters are independently distributed, our formulation of the multi-stage case has a nested form as in Shapiro et al. (2009). The problem to solve at time $t = 1$ is:

$$\inf_{x_1 \in X_1} \left\{ c_1^\top x_1 + \sup_{\mathbb{P}_{z_1} \in \mathcal{P}_{z_1}} \mathbb{E}_{\mathbb{P}_{z_1}} v_2(x_1, z_1) \right\}, \quad (4)$$

where $X_1 = \{x_1 \in \mathbb{R}_+^{n_1} : A_{11}x_1 = b_1\}$ and the ambiguity set $\mathcal{P}_{z_t}$, $t = 1, \dots, T-1$, is defined as

$$\mathcal{P}_{z_t} = \{\mathbb{P}_{z_t} : \text{supp}(z_{ti}) \subseteq [a_{ti}, b_{ti}], \quad \mathbb{E}_{\mathbb{P}}(z_{ti}) = \mu_{ti}, \quad \mathbb{E}_{\mathbb{P}}|z_{ti} - \mu_{ti}| = d_{ti}, \quad \forall i, \quad z_{ti} \perp z_{t,j}, \quad \forall i \neq j, \}.$$

The value function $v_t(x_{t-1}, z_{t-1})$, $t = 2, \dots, T-1$ is defined as the optimal value of the optimization problem to solve at time t :

$$v_t(x_{t-1}, z_{t-1}) = \inf_{x_t \in X_t} \left\{ c_t^\top x_t + \sup_{\mathbb{P}_{z_t} \in \mathcal{P}_{z_t}} \mathbb{E}_{\mathbb{P}_{z_t}} v_{t+1}(x_t, z_t) \right\}, \quad (5)$$

where $X_t = \{x_t \in \mathbb{R}_+^{n_t} : \sum_{s=1}^{t-1} A_{ts}(z_{t-1})x_s + A_{tt}x_t = b_t(z_{t-1})\}$ and $v_T(x_{T-1}, z_{T-1})$ is the optimal value of the optimization problem at stage T :

$$v_T(x_{T-1}, z_{T-1}) = \inf_{x_T \in X_T} \{c_T^\top x_T\}, \quad (6)$$

with $X_T = \{x_T \in \mathbb{R}_+^{n_T} : \sum_{s=1}^{T-1} A_{Ts}(z_{T-1})x_s + A_{TT}x_T = b_T(z_{T-1})\}$. At each stage t the objective function consists of a linear component involving the decisions x_t and (with the exception of stage T) the worst-case expected value of the optimal value of the problem to be solved at the next stage. At each stage, a system of constraints is to hold that involves the decision vectors $x_1, \dots, x_t$, and the coefficients $A_{ts}(\cdot)$, and $b_t(\cdot)$ which depend on the outcome of the uncertain parameter z_{t-1} , observed before x_t is implemented. We assume that the functions $A_{ti}(\cdot)$, and $b_t(\cdot)$ are linear. The assumption that matrices A_{tt} , $t = 1, \dots, T$ are fixed is the multi-stage equivalent of the two-stage *fixed recourse* restriction.

In the two-stage case of Section 2.1, in order to reformulate problem (1) to the closed-form equivalent, the function $v(x, \cdot)$ has to be convex. A similar property is needed here to reformulate the multi-stage problem to a closed form and holds for the functions $v_t(x_{t-1}, \cdot)$; moreover, at each time t the decision maker is solving a tractable convex optimization problem in the decision variables, as stated by the following proposition.

Proposition 3 *Functions $v_t(x_{t-1}, z_{t-1})$ and $v_T(x_{T-1}, z_{T-1})$ are convex in z_{t-1} , $t = 2, \dots, T-1$ and z_{T-1} , respectively, and the optimization problems (5) and (6) are convex in x_t for $t = 1, \dots, T-1$ and in x_T , respectively.*

Proof. See Appendix A. □

Proposition 3 implies that we can use the results of BTH72 to give a closed form of the multi-stage problem (4). We do this by recursively inserting the worst-case distributions of z_t from Proposition 1, considering the problem at stages $t = T-1, T-2, \dots, 2$. The final result is stated in the following proposition.

Proposition 4 *Distributionally robust SP problem formulated in (4), (5), and (6) is equivalent to the following problem:*

$$\inf_{x_1 \in X_1} \left\{ c_1^\top x_1 + \mathbb{E}_{\mathbb{P}_{\bar{z}_1}} v_2(x_1, \bar{z}_1) \right\}, \quad (7)$$

where

$$v_t(x_{t-1}, z_{t-1}) = \inf_{x_t \in X_t} \left\{ c_t^\top x_t + \mathbb{E}_{\mathbb{P}_{\bar{z}_t}} v_{t+1}(x_t, \bar{z}_t) \right\}, \quad t = 1, \dots, t-1, \quad (8)$$

where the worst-case distributions $\mathbb{P}_{\bar{z}_t}$ are defined as in Proposition 1, and $v_T(x_{T-1}, z_{T-1})$ is given by (6).

Formulations (7) and (8), together with the final-stage problem (6) constitute a single big optimization problem with a tree structure. In this structure, the first-stage problem refers via $v_2(x_1, \bar{z}^k)$, $k = 1, \dots, 3^{n_{z,1}}$ to $3^{n_{z,1}}$ second-stage problems, each of which links to $3^{n_{z,2}}$ stage 3 problems, and so on. This corresponds to the tree structure of the worst-case distribution of the uncertain parameter, depicted in Figure 1.

REMARK 2. As mentioned in the beginning of this section, it is possible, similar as in the two-stage case, (i) to consider also uncertainty in the objective function coefficients $c_2, \dots, c_T$ since the solutions of the optimization problems at each stage are concave in $c_2, \dots, c_T$, respectively; (ii) to construct a closed form of the problem in which the best-case expectation is minimized with respect to the distribution of parameter $z = (\xi, \omega)$.

Figure 1 Scenario tree for the worst-case distribution of $\bar{z} \in \mathbb{R}^2$ in the multi-stage setting $\bar{z} = (\bar{z}_1, \bar{z}_2) \in \mathbb{R}^2$.

3. Two-stage mixed-integer recourse models

Mixed-integer recourse models arise when the optimization problem involves integer decision variables. The advantage of incorporating such variables in the model is that they may be used to model e.g. indivisibilities or on/off decisions, the disadvantage however is that solving the model becomes much more complicated because generally the second-stage value function is non-convex. In the distributionally robust context of this paper, this implies that the result of BTH72 cannot be applied directly. Nevertheless, their result may be of use when we consider two-stage mixed-integer recourse models where some of the distributions of the random parameters in the model are known and others are unknown. The key observation in the underlying analysis is that under specific conditions the expected value function of a mixed-integer recourse model allows for a good convex approximation.

3.1. Problem formulation

Consider a two-stage mixed-integer recourse model with second-stage value function $v(x, z)$ defined as

$$v(x, z) = \inf_y \left\{ q(\xi)^\top y : Wy = h(\omega) - T(\omega)x, \ y \in \mathbb{Z}_+^p \times \mathbb{R}_+^{n_2-p} \right\},$$

where the vector $z = (\xi, \omega)$ represents the random parameters in the model. Similar to (1) we assume that q , h , and T are affine functions of these parameters. The distributionally robust mixed-integer recourse model that we consider in this section is

$$\inf_{x \in X} \sup_{\mathbb{P}_z \in \mathcal{P}_z} \mathbb{E}_{\mathbb{P}_z} [c^\top x + v(x, z)], \quad (9)$$

where $\mathcal{P}_z$ represents the (μ, d, β) ambiguity set and $X = \{x \in \mathbb{R}_+^{n_1} : Ax = b\}$. Since $v(x, z)$ is concave in ξ for fixed x and ω , it follows from the same reasoning as in Section 2.1 that $\mathbb{P}_\xi$ defined in Proposition 2 is the worst-case distribution of ξ . However, $v(x, z)$ is in general not convex in ω , so that the result of BTH72 cannot be applied to derive the worst-case expectations with respect to the distribution of ω . Nevertheless, we are able to use the result if some of the distributions of the random parameters are known and the other distributions are contained in a (μ, d, β) uncertainty set. For ease of exposition we assume in this section that the distribution of the right-hand side

random vector $h(\omega)$ is fully known, whereas only limited information is available on the distribution of the technology matrix $T(\omega)$, i.e. $\mathbb{P}_\omega \in \mathcal{P}_\omega$. Moreover, we assume that $h(\omega)$ is independent from $T(\omega)$. Furthermore, since we already discussed the worst-case distribution of ξ we assume that the second-stage costs parameters q are deterministic. For notational convenience, we drop the dependence of h and q on ω and ξ , respectively, and write T as a function of z instead of ω .

Under these assumptions, the distributionally robust mixed-integer recourse model in (9) reduces to

$$\inf_{x \in X} \sup_{\mathbb{P}_z \in \mathcal{P}_z} \mathbb{E}_{\mathbb{P}_z} [c^\top x + Q(x, z)], \quad (10)$$

where Q is defined for every realization of z as

$$Q(x, z) = \mathbb{E}_{\mathbb{P}_h} \left[\inf_y \left\{ q^\top y : Wy = h - T(z)x, y \in \mathbb{Z}_+^p \times \mathbb{R}_+^{n_2-p} \right\} \right]. \quad (11)$$

This expected value function Q is key to solving (10), since if Q is convex in z , then we may apply the result of BTH72 to obtain the worst-case distribution $\mathbb{P}_z$ of z . For example, Klein Haneveld et al. (2006) show that under specific conditions on h the expected value function Q of a simple integer recourse model is convex in the tender variables $u := T(z)x$, and thus indeed also convex in z . In general, however, $Q(x, \cdot)$ is not convex, but it may allow for a good convex approximation $\hat{Q}(x, \cdot)$. By replacing Q by $\hat{Q}$ we obtain an approximation of (10) for which the objective is convex in z , and thus $\mathbb{P}_z$ defined in Proposition 2 is its worst-case distribution. We derive error bounds on the optimality gaps of the approximating solutions that depend on $\|Q - \hat{Q}\|_\infty$, the maximum difference between Q and $\hat{Q}$.

In Section 3.2 we discuss the case where the simple integer recourse function Q is convex in z , and in Section 3.3 we derive the error bounds for using convex approximations $\hat{Q}$ for the general two-stage mixed-integer case. In the remainder of this section we briefly survey literature on convex approximations for mixed-integer recourse models and their corresponding error bounds; see also Romeijnders et al. (2014) for an overview.

Klein Haneveld et al. (2006) derived the first error bound for so-called α -approximations of simple integer recourse models which decreases with the total variations $|\Delta|f_i$ of the marginal probability densities functions f_i of the right-hand side random variables h_i . For example for normally distributed random variables this implies that the error bound decreases if the variances of the random variables increase. A similar error bound is obtained by Romeijnders et al. (2015, 2016b) for two different types of convex approximations for totally unimodular integer recourse models. The latter convex approximation is generalized by Romeijnders et al. (2016a) to the general two-stage mixed-integer case. The error bound corresponding to this convex approximation is asymptotic in nature: it converges to zero if all total variations of the probability density functions of the random variables in the model converge to zero.

3.2. Simple integer recourse models

The one-sided simple integer recourse model, introduced in Louveaux and van der Vlerk (1993), is a special case of (10) for which a closed-form expression for the second-stage value function can be obtained. The expected value function Q is given by

$$Q(x, z) = \sum_{i=1}^m \mathbb{E}_{\mathbb{P}_{h_i}} \left[q_i [h_i - T_i(z)x]^+ \right], \quad x \in \mathbb{R}^{n_1}, \quad (12)$$

where $[s]^+ := \max\{0, [s]\}$, $s \in \mathbb{R}$ and $T_i(z)$ is the i -th row of the matrix $T(z)$. Interestingly, Klein Haneveld et al. (2006) show that this simple integer recourse function Q may be convex in the tender variables $u = T(z)x$, and thus in z , if the underlying random vector h is continuously distributed and every marginal probability density function f_i can be expressed as

$$f_i(s) = H_i(s+1) - H_i(s), \quad s \in \mathbb{R}, \quad (13)$$

for some cumulative distribution function H_i with finite mean. This implies that under these conditions the worst-case distribution $\mathbb{P}_{\bar{z}}$ of z can be derived using the results of BTH72 (this worst-case distribution is the same for every first-stage decision x).

Proposition 5 *Consider the distributionally robust simple integer recourse model*

$$\inf_{x \in X} \sup_{\mathbb{P}_z \in \mathcal{P}_z} \mathbb{E}_{\mathbb{P}_z} [c^\top x + Q(x, z)], \quad (14)$$

where Q is defined in (12), $X = \{x \in \mathbb{R}_+^{n_1} : Ax = b\}$, and the ambiguity set $\mathcal{P}_z$ for the distributions $\mathbb{P}_z$ of z is defined analogously to (2). Then, if each random variable h_i has a pdf f_i satisfying (13), then the optimization problem in (14) is equivalent to

$$\inf_{x \in X} \mathbb{E}_{\mathbb{P}_{\bar{z}}} [c^\top x + Q(x, \bar{z})],$$

where the worst-case distribution $\mathbb{P}_{\bar{z}}$ of z is defined analogously as in Proposition 1. $\square$

In case a marginal density function f_i does not satisfy (13) a natural approach is to approximate it by a density function $\hat{f}_i$ that is approximately the same as f_i , but does satisfy (13), yielding a convex approximation $\hat{Q}$ of Q . This is the main idea behind the so-called α -approximations derived in Klein Haneveld et al. (2006), and their generalization to complete integer recourse models by van der Vlerk (2004). For these convex approximations upper bounds on $\|Q - \hat{Q}\|_\infty$ have been derived. Accordingly, in the next section we assume that a convex approximation $\hat{Q}$ and corresponding upper bound on $\|Q - \hat{Q}\|_\infty$ are available.

3.3. Convex approximations

Again consider the general distributionally robust two-stage mixed-integer recourse model defined in (10):

$$\eta^* := \inf_{x \in X} \sup_{\mathbb{P}_z \in \mathcal{P}_z} \mathbb{E}_{\mathbb{P}_z} [c^\top x + Q(x, z)], \quad (15)$$

where the expected value function Q , defined in (11), is generally non-convex. We assume that Q allows for a good convex approximation $\hat{Q}$ for which $\|Q - \hat{Q}\|_\infty$ is small. Then, we may approximate (15) by replacing Q by $\hat{Q}$, obtaining

$$\hat{\eta} := \inf_{x \in X} \sup_{\mathbb{P}_z \in \mathcal{P}_z} \mathbb{E}_{\mathbb{P}_z} [c^\top x + \hat{Q}(x, z)] \quad (16)$$

$$= \inf_{x \in X} \mathbb{E}_{\mathbb{P}_{\bar{z}}} [c^\top x + \hat{Q}(x, \bar{z})], \quad (17)$$

where the equality in (17) follows from applying the result of BTH72 to the convex objective in (16). The approximating problem is a convex optimization problem for which the distributions of the random parameters are known. It can be solved efficiently using existing solution methods from SP; see Section 4. To guarantee the quality of the approximate solution $\hat{x}$ obtained from solving the optimization problem in (17), we derive an error bound on the optimality gap $G(\hat{x}) - \eta^*$, where $G(\hat{x})$ represents the objective value of the solution $\hat{x}$:

$$G(x) := \sup_{\mathbb{P}_z \in \mathcal{P}_z} \mathbb{E}_{\mathbb{P}_z} [c^\top x + Q(x, z)], \quad x \in X. \quad (18)$$

In fact, we show that $|\hat{\eta} - \eta^*| \leq \|Q - \hat{Q}\|_\infty$ and $G(\hat{x}) - \eta^* \leq 2\|Q - \hat{Q}\|_\infty$; see Theorem 1 below.

Interestingly, we may approximate the optimization model in (17) by replacing $\hat{Q}$ by the original mixed-integer recourse function Q to obtain the approximating model

$$\tilde{\eta} = \inf_{x \in X} \mathbb{E}_{\mathbb{P}_{\bar{z}}} [c^\top x + Q(x, \bar{z})]. \quad (19)$$

This model indirectly approximates the original mixed-integer recourse model (15), but it can also be derived directly from (15) by assuming that $\mathbb{P}_{\bar{z}}$ is the worst-case distribution in that model. However, using the interpretation of an indirect approximation via the convex approximating model in (17), we can derive an error bound for the approximate solution $\tilde{x}$ obtained from solving (19).

Theorem 1 *Consider the distributionally robust mixed-integer recourse model defined in (15) and let $\hat{Q}$ be any convex approximation of the mixed-integer expected value function Q defined in (11). Let $\hat{x}$ and $\tilde{x}$ denote optimal solutions of the approximating models defined in (17) and (19), respectively. Then,*

$$\begin{aligned} (i) \quad & |\hat{\eta} - \eta^*| \leq \|Q - \hat{Q}\|_\infty & \text{and} & & G(\hat{x}) - \eta^* \leq 2\|Q - \hat{Q}\|_\infty, \\ (ii) \quad & 0 \leq \eta^* - \tilde{\eta} \leq 2\|Q - \hat{Q}\|_\infty & \text{and} & & G(\tilde{x}) - \eta^* \leq 2\|Q - \hat{Q}\|_\infty. \end{aligned}$$

Furthermore, since the upper bound on $G(\tilde{x}) - \eta^*$ holds for every convex approximation $\hat{Q}$, it actually holds for the best convex approximation:

$$G(\tilde{x}) - \eta^* \leq 2 \inf_{\hat{Q}} \{\|Q - \hat{Q}\|_{\infty} : \hat{Q} \text{ is convex}\}.$$

Proof. See Appendix A. □

From a computational point of view, the approximating model in (17) is easiest to solve since it is a convex optimization model. The approximating model in (19) is a non-convex two-stage mixed-integer recourse model for which the distributions of the random parameters are known. The latter is the main advantage of this approximating model over the original distributionally robust model in which the worst-case distribution of $\mathbb{P}_{\bar{z}}$ still has to be determined and may possibly be different for every first-stage decision x . Nevertheless, solving (19) can be a very challenging task. The error bound for this approximating model, however, is the same as for the convex approximating model in (17). The fact that the optimality gap of $G(\tilde{x}) - \eta^*$ does not depend on the particular $\hat{Q}$ implies that even if no good convex approximation $\hat{Q}$ of Q is known, we might still approximate the distributionally robust mixed-integer recourse model in (15) by assuming that $\mathbb{P}_{\bar{z}}$ is the worst-case distribution of z . If a good convex approximation $\hat{Q}$ of Q is available, then we can use it in the convex approximating model (17).

4. Solution methods for continuous stochastic programming models with exponentially many scenarios

In Sections 2 and 3 we have shown how to reduce a distributionally robust optimization problem to an SP problem for which the distributions of the random variables in the model are known. In particular, in case all decision variables are continuous, we need to solve a continuous stochastic programming model

$$\inf_{x \in X} \left\{ c^{\top} x + \mathbb{E}_{\mathbb{P}_{\bar{z}}} [v(x, \bar{z})] \right\}, \quad (20)$$

where the joint distribution of $\bar{z}$ has exponentially many scenarios in the number of random parameters. From a robust optimization point of view this means that the problem in (20) is intractable. Indeed, Dyer and Stougie (2006) show that these SP problems are $\#P$ -hard. Nevertheless, there has been a vast amount of work in the SP literature that deals with this kind of problems, yielding efficient (approximate) solution methods to these problems, in particular for two-stage problems.

In this section we first discuss so-called simple recourse problems in Section 4.1 and we show that for stochastic programming models with such structure, the size of (20) does not increase exponentially in the number of random parameters. In Section 4.2 we discuss techniques from the SP literature to solve two-stage and multi-stage stochastic programming problems with exponentially many scenarios.

4.1. Simple recourse models

In this section we consider so-called simple recourse models introduced by Wets (1983), where the recourse matrix $W = [I_m, -I_m]$ with I_m denoting the m -dimensional identity matrix. For this model, the second-stage value function is given by

$$v(x, z) = \inf_{y^+, y^-} \left\{ (q^+(\xi))^\top y^+ + (q^-(\xi))^\top y^- : y^+ - y^- = h(\omega) - T(\omega)x, y^+, y^- \in \mathbb{R}_+^m \right\},$$

with the conventional indices ‘+’ and ‘-’ representing the surplus and the shortage, respectively. We can obtain an exact expression for this second-stage value function, using among others the *separability of the second-stage problem*:

$$v(x, z) = \sum_{i=1}^m \left(q_i^+(\xi) (h_i(\omega) - T_i(\omega)x)^+ + q_i^-(\xi) (h_i(\omega) - T_i(\omega)x)^- \right),$$

where $T_i(\omega)$ denotes the i -th row of $T(\omega)$, and $(h_i(\omega) - T_i(\omega)x)^+$ and $(h_i(\omega) - T_i(\omega)x)^-$ denoting the nonnegative and the nonpositive parts of $h_i(\omega) - T_i(\omega)x$, respectively. Suppose that only the right-hand side random vector h is random, then if we drop the dependence of h on z , (20) reduces to

$$\inf_{x \in X} \left\{ c^\top x + \sum_{i=1}^m \mathbb{E}_{\mathbb{P}_{\bar{h}_i}} \left[q_i^+(\bar{h}_i - T_i x)^+ + q_i^-(\bar{h}_i - T_i x)^- \right] \right\}, \quad (21)$$

where $\mathbb{P}_{\bar{h}_i}$ is the worst-case distribution of h_i as defined in Proposition 1. Since it is a three-point distribution, the size of the problem in (21) only increases linearly in m . The key observation here is that due to the separability of the second-stage problem the simple recourse model in (21) only involves the m marginal distributions $\mathbb{P}_{\bar{h}_i}$, each with three scenarios, instead of the joint distribution $\mathbb{P}_{\bar{h}}$ with 3^m scenarios.

In case there is also uncertainty in the technology matrix T , the simple recourse problem is not completely separable in the random parameters. However, we show in the operating room experiment of Section 5.1 that we can use the structure of the problem to substantially speed up the existing algorithms.

4.2. Stochastic Programming approaches

The fact that the size of the problem grows exponentially in the number of random parameters is common in SP, and many SP approaches are aimed at reducing the number of scenarios. In this section we survey some relevant SP literature.

One of the most frequently used solution methods is the sample average approximation (SAA), discussed in e.g. Shapiro et al. (2009). The idea of this method is to replace the original worst-case distribution of $\bar{z}$ in (20) by a sample z^s , $s = 1, \dots, N_s$, where N_s is much smaller than the number of scenarios of $\bar{z}$, yielding

$$\inf_{x \in X} \left\{ c^\top x + \frac{1}{N_s} \sum_{s=1}^{N_s} v(x, z^s) \right\}. \quad (22)$$

If the sample size N_s is small, then the approximation in (22) is easier to solve than the original model in (20). We may solve (22) for several different samples of $\bar{z}$ yielding (possibly) different first-stage solutions x , and use an out-of-sample test to determine the best among them. In the operating room experiment of Section 5.1 we show that the SAA method may give near-optimal solutions.

Alternatively, we may use other approaches to reduce the number of scenarios. For example, Dupačová et al. (2003) and Heitsch and Römisch (2003) do so by combining similar scenarios. Pflug (2001) uses the Wasserstein metric to construct a discrete probability distribution (with few scenarios) that minimizes the distance between the original and approximating distribution. His method can also be applied to multi-stage stochastic programming models. Approximations relying on a reduced scenario set are justified by stability results of e.g. Römisch (2003) which shows that a small change in the distributions of the random parameters only result in a small change in the optimal first-stage solutions.

For two-stage stochastic programming models with only a modest number of scenarios efficient solution methods are available. Most of them rely on decomposition of the problem and are variants of the L-shaped algorithm of van Slyke and Wets (1969); see e.g. Ruszczyński (1986) and Hingle and Sen (1991) for well-known examples. We refer to Zverovich et al. (2012) for a recent survey comparing several decomposition methods. Although multi-stage stochastic programming models are considerably more difficult to solve than two-stage models, several solution methods do exist. For the interested reader we mention progressive hedging (Rockafellar and Wets 1991), nested Benders' decomposition (Birge 1985), and stochastic dual dynamic programming (Pereira and Pinto 1991).

So far we have only discussed how to obtain a first-stage solution. However, when this solution is obtained by solving an approximation of the original stochastic programming problem, then we may use sampling to assess the quality of the solution; see e.g. the Multiple Replications Procedure (MRP) of Bayraksan and Morton (2009). Different sampling methods, such as Latin Hypercube sampling, may be used to reduce the bias and sample variance of the optimality gap of the approximating solution. We use the MRP to assess the quality of a surgery-to-OR assignment in the operating room experiment of Section 5.2.

5. Numerical experiments

In this section we present three numerical experiments to illustrate the advantages of the approach developed. The first experiment, a modified version of the operating room (OR) scheduling problem of Denton et al. (2010), illustrates (i) how to reduce the computational effort related to the exponential number of scenarios by using SP techniques and exploiting the problem's properties

and (ii) the differences in the performance of distributionally robust solutions compared to other methods used in OR management.

The second experiment, related also to OR management and involving integer recourse variables shows (i) how the novel theoretical results of Section 3 can be used to construct intuitive convex approximations of this integer recourse model, (ii) how to solve it efficiently, and (iii) how to use additional existing techniques to obtain better bounds on the performance of the optimal solutions.

In the third and last experiment, which is a continuation of the inventory management experiment from Postek et al. (2015), we show (i) how our approach is applied to multi-stage problems, (ii) how feasible decisions can be constructed for uncertainty realizations not belonging to the discrete worst-case support, and (iii) we provide managerial insights regarding the value of distributional information and the trade-off between worst-case objective value and worst-case expected objective value.

5.1. Operating room scheduling under uncertainty

We apply the method proposed in Section 2.1 to the OR scheduling problem introduced by Denton et al. (2010). In this problem, surgeries with random durations have to be assigned to ORs before the durations of these surgeries are known. Fixed costs are incurred for every OR that is opened, and for each OR overtime costs are incurred if the actual total duration of the surgeries exceeds a regular work day of T minutes. Contrary to Denton et al. (2010), we assume that the probability distributions of the surgery durations are (partially) unknown and, hence, we minimize the total worst-case expected costs using the result of this paper. We carry out numerical experiments to show that for problem instances with 10 or 15 surgeries as in Denton et al. (2010), we are able to obtain the optimal surgery-to-OR allocation with reasonable computational effort.

In Section 5.1.1 we define the OR scheduling problem and list the various solution methods we use and which are detailed in Appendix B. In Section 5.1.2 we carry out the numerical experiments.

5.1.1. Problem formulation The OR scheduling problem can be formulated as a two-stage recourse model, where in the first stage we have to determine how many ORs to open and the assignment of the surgeries to the ORs. With N denoting the number of surgeries that have to be performed, we define y_{ij} for every $i, j = 1, \dots, N$, as a binary variable equal to 1 if surgery j is assigned to the i -th OR, and 0 otherwise. Thus, we assume that there N ORs available. Accordingly, we define x_i for every $i = 1, \dots, N$, as a binary variable equal to 1 if the i -th OR is opened, and 0 otherwise. Furthermore, for every opened OR we incur fixed costs c_f and for every minute of overtime exceeding a regular workday of T minutes we incur variable costs c_v per OR.

Let z represent the random vector of surgery durations and θ_i the minutes of overtime in the i -th OR. Then, in case the surgery durations z would be deterministic the OR scheduling problem reads

$$\begin{aligned} \min_{x,y,\theta} \quad & \sum_{i=1}^N c_f x_i + \sum_{i=1}^N c_v \theta_i \\ \text{s.t.} \quad & \sum_{i=1}^N y_{ij} = 1, & j = 1, \dots, N, \end{aligned} \quad (23)$$

$$y_{ij} \leq x_i, \quad i, j = 1, \dots, N, \quad (24)$$

$$\theta_i \geq \sum_{j=1}^N z_j y_{ij} - T x_i, \quad i = 1, \dots, N, \quad (25)$$

$$x_i \in \{0, 1\}, \quad y_{ij} \in \{0, 1\}, \quad \theta_i \geq 0, \quad i, j = 1, \dots, N. \quad (26)$$

Constraint (23) means that every surgery j is assigned to exactly one OR, constraint (24) models that surgery j can only be assigned to the i -th OR if it is opened, and constraint (25) defines θ_i as the minutes of overtime for the i -th OR.

We let X denote the set of feasible first-stage decisions x and y satisfying (23), (24), and (26). In addition, we assume that X includes several symmetry breaking constraints introduced in Denton et al. (2010). For example, we assume without loss of generality that $x_1 \geq \dots \geq x_N$. Moreover, if the surgeries j_1 and j_2 with $j_1 < j_2$ are of the same type, then we assume that surgery j_1 (j_2) is assigned to OR k_1 (k_2), with $k_1 \leq k_2$, respectively:

$$\sum_{k=1}^i y_{k,j_1} \geq \sum_{k=1}^i y_{k,j_2}, \quad i = 1, \dots, N.$$

Similar as Denton et al. (2010) we assume that the surgery durations z are random and unknown when the surgery-to-OR assignment has to be made. Contrary to this reference, however, we assume that the distribution $\mathbb{P}_z$ of the random vector z is unknown and belongs to a (μ, d, β) ambiguity set $\mathcal{P}_z$ as defined in Section 2.1. The objective is to find a surgery-to-OR assignment, i.e., to determine $(x, y) \in X$, that minimizes the worst-case expected total costs. Given a first-stage decision $(x, y) \in X$ and a realization z of surgery durations, the number of minutes of overtime in the i -th OR is

$$\theta_i(x, y, z) = \left(\sum_{j=1}^N z_j y_{ij} - T x_i \right)^+. \quad (27)$$

The OR scheduling problem minimizing worst-case expected costs is thus given by

$$\min_{(x,y) \in X} \sup_{\mathbb{P}_z \in \mathcal{P}_z} \left\{ \sum_{i=1}^N c_f x_i + c_v \mathbb{E}_{\mathbb{P}_z} \left[\sum_{i=1}^N \left(\sum_{j=1}^N z_j y_{ij} - T x_i \right)^+ \right] \right\}. \quad (28)$$

Since the objective function in (28) is convex in z for every $(x, y) \in X$, we can use the result of BTH72 to obtain the worst-case distribution $\mathbb{P}_{\bar{z}}$ as defined in Proposition 1, and thus the optimization problem in (28) reduces to

$$\min_{(x, y) \in X} \left\{ \sum_{i=1}^N c_f x_i + c_v \mathbb{E}_{\mathbb{P}_{\bar{z}}} \left[\sum_{i=1}^N \left(\sum_{j=1}^N \bar{z}_j y_{ij} - T x_i \right)^+ \right] \right\}. \quad (29)$$

The optimization problem in (29) is a two-stage recourse model with binary first-stage variables and continuous second-stage variables, and where the random vector $\bar{z}$ of surgery durations has 3^N scenarios. In fact, this problem has the simple recourse structure discussed in Section 4.1. However, contrary to Section 4.1, here the randomness is in the technology matrix and not in the right-hand side, so that the number of scenarios does not necessarily reduce to $3N$ as in Section 4. Nevertheless, we will use the structure of the problem to deal with exponentially many scenarios.

We use several solution methods to solve the optimization problem in (29). The first is a Sample Average Approximation (SAA) method, see Shapiro et al. (2009), which is very easy to implement in practice. The second method (LDR-WCEC: Linear Decision Rules - Worst-Case Expected Cost) uses linear decision rules (LDR) for the overtime costs θ_i so that the optimal surgery-to-OR assignment in this approximating optimization problem can be obtained very fast. The drawback of these two methods is that they only yield an approximate solution to (29). Therefore we also use an L-shaped algorithm, see van Slyke and Wets (1969), which yields the optimal solution to (29). The challenge of this exact algorithm is to deal with an exponential number of scenarios. In Appendix B, we discuss the SAA and LDR-WCEC methods, and the L-shaped algorithm for this OR experiment. There, we also present several ideas to deal with the exponential number of scenarios, speeding up computations for the L-shaped algorithm considerably. In Section 5.1.2 we carry out numerical experiments and find, among others, that the SAA method yields near-optimal solutions within reasonable time limits.

5.1.2. Numerical experiments In this section we carry out numerical experiments on problem instances of similar size as in Denton et al. (2010), i.e. with $N = 10$ and $N = 15$. In all experiments we assume that $c_f = 1$ and $c_v = 0.333$ or $c_v = 0.0833$, similar as in Denton et al. (2010). To obtain the parameters of the (μ, d, β) ambiguity set $\mathcal{P}_z$ we use data on the surgery duration distributions given in Gul et al. (2011). In this reference, estimates of surgery duration distributions are given for several types of surgeries. We use these estimates to compute μ, d, a, b , and β , where a and b represent the 0.1% and 99.9% quantile of the distribution. In Table 1 the data of the four types of surgeries that we consider in our experiments are given.

For all four combinations of N and c_v we generate 50 problem instances by randomly sampling with equal probabilities N surgery types from Table 1. We only report results for $N = 15$, since

Table 1 OR experiment - parameters of the surgery duration distributions (Gul et al. 2011).

Surgical group	μ	a	b	d	β
Oral Maxillofacial procedure	36.00	2.2	307.2	22.3	0.34
Pain Medicine	20.93	2.3	125.3	10.6	0.37
Ophthalmology	41.63	11.9	125.5	12.6	0.42
Urology	138.16	37.7	433.1	43.3	0.42

results for $N = 10$ are similar. For every problem instance we use the SAA method with $N_s = 1000$, the LDR-WCEC method, and the L-shaped algorithm to obtain surgery-to-OR assignments (x, y) . In addition, we also obtain $(x, y) \in X$ minimizing the best-case expected costs (min-BCEC) using a similar L-shaped algorithm as for minimizing the worst-case expected costs, and we obtain the surgery-to-OR assignment $(x, y) \in X$ minimizing the worst-case costs (min-WC). For all these first-stage solutions (x, y) , we calculate the fixed costs (FC), the worst-case expected costs (WCEC), the best-case expected costs (BCEC), the expected costs (EC), the worst-case costs (WCC) over the support:

$$\mathcal{Z} = [a_1, b_1] \times \dots \times [a_N, b_N],$$

and the running time (RT) of the algorithm in seconds. Here, the fixed costs (FC) represent the number of opened ORs since $c_f = 1$. Moreover, the expected costs (EC) are estimated using a sample of 100,000 from the surgery duration distributions given in Gul et al. (2011). Furthermore, to facilitate comparison with their results, the worst-case costs (WCC- τ), depending on a parameter τ , are calculated using the same uncertainty set as in Denton et al. (2010):

$$\mathcal{Z} = \left\{ z \in \mathbb{R}_+^N : z_j \in [a_i, b_i] \ \forall j, \sum_{j=1}^N \frac{z_j - a_j}{b_j - a_j} \leq \tau \right\}.$$

Here, τ is a parameter representing how many surgeries can attain their maximum duration. The averages of these performance measures over the 50 problem instances are given in Table 2.

Table 2 OR experiment - average values of the performance measures over 50 problem instances with $N = 15$ and $c_v = 0.333$. For every solution, we report the fixed costs (FC), the worst-case expected costs (WCEC), the best-case expected costs (BCEC), the expected costs (EC), the worst-case costs (WCC), and the running time (RT) of the algorithm in seconds.

Solution method	FC	BCEC	EC	WCEC	WCC-1	WCC-2	WCC-4	RT (in seconds)
min-BCEC	2.50	2.69	3.11	4.66	3.00	15.49	30.50	8.06
SAA-WCEC	3.28	3.28	3.41	4.27	3.36	14.62	28.89	63.26
L-shaped-WCEC	3.30	3.30	3.42	4.25	3.37	14.73	29.30	167.19
LDR-WCEC	8.08	8.08	8.09	8.08	8.08	8.70	10.56	1.40
min-WC	8.98	8.98	8.99	8.98	8.98	8.98	9.12	0.29

We conclude from Table 2 that the SAA method and the L-shaped algorithm yield very similar results. This implies that, although the SAA solution does not necessarily minimize the worst-case expected costs, its solution is (near-)optimal. Moreover, the surgery-to-OR assignment obtained

using linear decision rules for minimizing worst-case expected costs (LDR-WCEC) is more stable since the worst-case costs with $\tau = 2$ and $\tau = 4$ is much smaller. However, in expectation this LDR-WCEC solution is not good for these problem instances. This is because the number of ORs that are opened in this solution, i.e. the fixed costs (FC), are much larger than for the min-BCEC, SAA, and the L-shaped approaches. We observe in Table 2 that the fixed costs are smallest for the min-BCEC solution. This is as expected since the solution minimizes the best-case expectation corresponding to surgery duration distributions for which the longest possible surgery durations are smaller than for the worst-case expectation. Because fewer ORs are opened for this solution, its worst-case expectation is larger than for the L-shaped algorithm and SAA method. On average however, i.e. sampling from the estimated surgery duration distributions of Gul et al. (2011), the min-BCEC solution performs better. Comparing the running times of the algorithms we observe that the LDR-WCEC, min-BCEC and min-WC methods run within several seconds, whereas the SAA method and L-shaped algorithm require on average one minute and almost three minutes, respectively. Given that the SAA method can be implemented more efficiently than in our experiments (using e.g. a decomposition algorithm) and that the L-shaped algorithm minimizes the exact worst-case expected costs under 3^N scenarios, these methods run within reasonable time limits.

Table 3 OR experiment - average values of the performance measures over 50 problem instances with $N = 15$ and $c_v = 0.0833$. Terminology as in Table 2.

Solution method	FC	BCEC	EC	WCEC	WCC-1	WCC-2	WCC-4	RT (in seconds)
min-BCEC	2.1	2.4	2.53	2.9	2.34	5.51	9.69	4.57
SAA-WCEC	2.08	2.42	2.53	2.89	2.32	5.31	10.2	30.93
L-shaped-WCEC	2.08	2.42	2.53	2.89	2.33	5.34	10.2	24.7
LDR-WCEC	3.34	3.4	3.45	3.69	3.35	6.05	9.85	0.65
min-WC	8.44	8.44	8.44	8.44	8.44	8.44	8.65	0.45

In Table 3 the results for $N = 15$ and $c_v = 0.0833$ are given. Comparing to the previous case with $c_v = 0.333$, we observe that on average the number of opened ORs is smaller than in Table 2 since the per minute overtime costs are smaller. Moreover, the LDR-WCEC solution does not have as large worst-case expected costs, but the approximation is still not very good in expectation. Furthermore, the difference between the optimal worst-case expected costs and optimal best-case expected costs, displayed in bold face in Tables 2 and 3, is smaller now. Since this difference yields an upper bound on the *value of distributional information* of the surgery durations, we conclude that for $c_v = 0.333$ we would be willing to spend more time and effort to better estimate the distributions of the surgery durations.

Finally, we report on the efficiency of our tailored implementation of the L-shaped algorithm, as described in Appendix B.3. To illustrate the reductions in the number of scenarios that need to

be used by our L-shaped algorithm, we computed the average number of scenarios that had to be evaluated per L-shaped iteration over all 50 runs. The results are given in Table 4 and we can see that on average we need between 5% and 15% of the scenarios if the number of surgeries is 10 and less than 1% of the scenarios in case $N = 15$.

Table 4 OR experiment - average numbers of scenarios to evaluate per iteration in the L-shaped algorithm.

Experiment	N	c_v	Evaluations	Evaluations/ $3^N \times 100\%$
1	10	0.333	2833	4.80
2	10	0.0833	7146	12.10
3	15	0.333	8509	0.059
4	15	0.0833	25685	0.179

5.2. Numerical experiments for two-stage ambiguous integer recourse models

In this section we again consider the OR scheduling problem of Section 5.1. However, here we assume that overtime wages are paid in full hours: if the overtime in a given OR is only a few minutes, then still its OR staff has to be paid a full hour of overtime work. In addition we assume that there is uncertainty in the regular work day duration T . This duration may be interpreted as the effective time spent on performing surgeries and may be smaller (or larger) than the targeted 480 minutes due to inefficiency (or efficiency) of the OR staff.

In Section 5.2.1 we show that this problem can be modelled as a distributional robust integer recourse model of Section 3, and we derive a convex approximation for this problem. In Section 5.2.2 we evaluate this convex approximation using numerical experiments.

5.2.1. Problem definition and convex approximation Since we do not have detailed information about the efficiency of the OR staff, we assume that for every OR i the duration of a regular work day equals $T + \epsilon_i$, where the probability distribution $\mathbb{P}_{\epsilon_i}$ of ϵ_i belongs to a (μ, d, β) -uncertainty set with $E[\epsilon_i] = 0$. Moreover, contrary to Section 5.1, we assume that the probability distributions of the surgery durations z are known, for example based on historical data. Under these assumptions, the problem can be cast into the framework of Section 3, where the distributions of the surgery durations z are known and the distributions of the work day durations $(T + \epsilon_i)x_i$ are contained in a (μ, d, β) ambiguity set. Again, letting X denote all feasible surgery-to-OR assignments (x, y) , the optimization problem we consider is given by

$$\inf_{(x, y) \in X} \left\{ \sum_{i=1}^N c_f x_i + \sup_{\mathbb{P}_{\epsilon} \in \mathcal{P}_{(\mu, d)}} \left\{ \mathbb{E}_{\mathbb{P}_{\epsilon}} \left[\mathbb{E}_{\mathbb{P}_z} \left[\sum_{i=1}^N 60c_v \left\lceil \left(\sum_{j=1}^N z_j y_{ij} - (T + \epsilon_i)x_i \right) / 60 \right\rceil^+ \right] \right] \right\} \right\}.$$

Here, all durations are in minutes and the round-up operator ensures that overtime wages are paid in full hours.

Because of the round-up operator, the objective function is not convex in ϵ and thus the results of Ben-Tal and Hochman (1972) cannot be applied. This means that we do not know the worst case distribution of $\mathbb{P}_\epsilon$. In fact, the worst-case distribution may be different for every surgery-to-OR assignment $(x, y) \in X$. Following Section 3, we define the expected value function Q as

$$Q(x, y, \epsilon) = \mathbb{E}_{\mathbb{P}_z} \left[\sum_{i=1}^N 60c_v \left\lceil \left(\sum_{j=1}^N z_j y_{ij} - (T + \epsilon_i)x_i \right) / 60 \right\rceil^+ \right],$$

and we consider its convex approximation

$$\hat{Q}(x, y, \epsilon) = \mathbb{E}_{\mathbb{P}_z} \left[\sum_{i=1}^N c_v \left(\sum_{j=1}^N z_j y_{ij} - (T + \epsilon_i - 30)x_i \right)^+ \right].$$

Here, we simultaneously relax the integrality of the overtime hours and subtract 30 minutes from the work day duration (if the i -th OR is opened). The rationale of doing so is that on average we have to pay approximately 30 minutes of additional overtime if overtime is paid in full hours.

For the convex approximating model with Q replaced by $\hat{Q}$ we can apply the results from Section 3 to conclude that the worst-case distribution equals $\mathbb{P}_\epsilon$ for every $(x, y) \in X$. The approximating model becomes

$$\begin{aligned} & \inf_{(x, y) \in X} \left\{ \sum_{i=1}^N c_f x_i + \mathbb{E}_{\mathbb{P}_\epsilon} [\hat{Q}(x, y, \bar{\epsilon})] \right\} \\ &= \inf_{(x, y) \in X} \left\{ \sum_{i=1}^N c_f x_i + \mathbb{E}_{\mathbb{P}_\epsilon} \left[\mathbb{E}_{\mathbb{P}_z} \left[\sum_{i=1}^N c_v \left(\sum_{j=1}^N z_j y_{ij} - (T + \epsilon_i - 30)x_i \right)^+ \right] \right] \right\}. \end{aligned} \quad (30)$$

This model can be solved e.g. using SAA yielding an approximating surgery-to-OR assignment $(\hat{x}, \hat{y})$.

5.2.2. Numerical experiments For our numerical experiments we consider the same setting as in Section 5.1. That is, the fixed costs for opening an OR are normalized, i.e. $c_f = 1$, and we consider two different cases for the overtime costs: $c_v = 0.0333$ and $c_v = 0.00833$. Moreover, we only consider the types of surgeries presented in Table 1. For every OR i , we assume that the ambiguity set of $\mathbb{P}_{\epsilon_i}$ is defined by $a_i = -60$, $b_i = 60$, $\mu_i = 0$, $d_i = 30$, and $\beta_i = 0.5$. This means that the regular work day duration $T + \epsilon_i$ will be between 420 and 540 minutes (i.e. 7 and 9 hours).

To solve the convex approximating model in (30) we use SAA with sample size $\hat{N}_s$ to obtain an approximating surgery-to-OR assignment. We repeat this procedure ten times, obtaining ten possibly different surgery-to-OR assignments, and use an out-of-sample test of size 10,000 to obtain the best among them. We let $(\hat{x}, \hat{y})$ denote this surgery-to-OR assignment. Contrary to Section 5.1 we are not able to determine the optimal surgery-to-OR assignment. That is why we have to use a different approach to determine the quality of the solution $(\hat{x}, \hat{y})$. It turns out that using the

maximum difference between the expected value function Q and its convex approximation $\hat{Q}$, as suggested in Theorem 1 yields error bounds that are too large for this particular problem. That is why we instead use a combination of the Multiple Replications Procedure (MRP) discussed in e.g. Bayraksan and Morton (2009) and total variation error bounds. In Appendix C.1 we discuss this approach in more detail. The result is an (approximate) 95% confidence interval on the optimality gap of $(\hat{x}, \hat{y})$.

Table 5 Integer OR experiment - numerical results for the integer OR problem with $N = 15$ over 10 problem instances, where the surgery types are randomly generated based on Table 1, and $c_v = 0.0333$. Here, $\hat{N}_s$ denotes the sample size used to obtain the approximating solution, and FC denotes the fixed costs of this solution. Next, ELB OBJ VAL gives an expected lower bound on the optimal objected value, and next we have an (approximate) 95% confidence interval on the absolute optimality gap and an upper bound on the relative optimality gap. Finally, RT denotes the average running time in seconds of solving the SAA of the convex approximating model with a sample size of $\hat{N}_s$.

$\hat{N}_s$	FC	ELB OBJ VAL	95% CI OPT GAP	REL OPT GAP	RT (in sec)
10	2.8	3.40	[0,0.214]	6.38%	1.16
100	3.0	3.40	[0,0.086]	2.52%	11.5
1000	3.0	3.40	[0,0.067]	1.98%	143.5

Table 6 Integer OR experiment - numerical results for the integer OR problem with $N = 15$ over 10 problem instances, where the surgery types are randomly generated based on Table 1, and $c_v = 0.00833$. Terminology the same as in Table 5.

$\hat{N}_s$	FC	ELB OBJ VAL	95% CI OPT GAP	REL OPT GAP	RT (in sec)
10	2.8	3.41	[0,0.243]	6.68%	1.43
100	2.9	3.41	[0,0.093]	2.72%	9.98
1000	3.0	3.41	[0,0.074]	2.20%	176.5

The results in Tables 5 and 6 are obtained by solving 10 problem instances with $c_v = 0.0333$ and $c_v = 0.00833$, respectively, and $N = 15$ surgeries, each randomly selected from Table 1. We obtained three approximating surgery-to-OR assignment by solving the SAA of (30) with sample size $\hat{N}_s = 10, 100$, and 1000. As can be observed in the tables, the computation time required to obtain these surgery-to-OR assignments increases in the sample size $\hat{N}_s$. As expected, the quality of the surgery-to-OR assignments also increases in the sample size since the relative optimality gap (REL OPT GAP) decreases from approximately 6% for $\hat{N}_s = 10$ to 2% for $\hat{N}_s = 1000$. This difference is not caused by the number of ORs that is opened in the two different cases but rather by the different surgery-to-OR assignments. Overall we conclude that the 95% confidence intervals on the optimality gap are surprisingly small compared to the expected lower bound on the objective value for $\hat{N}_s = 1000$, in particular since these distributionally robust integer problems are extremely hard to solve and we are not able to calculate the optimal solution. We would like to stress that the

values in Tables 5 and 6 are upper bounds on the optimality gap that hold with high probability. The actual value of the optimality gaps might be even smaller. In Appendix C.1 more details can be found on how the confidence interval on the optimality gap is obtained.

5.3. Inventory experiment

5.3.1. Introduction Our final experiment concerns a multi-stage problem - an inventory management example adapted from Ben-Tal et al. (2005), used also in Postek et al. (2015), comprising a single product with inventory managed over T stages. At the beginning of each stage t the decision maker has an inventory of size I_t and he orders a quantity x_t for unit price c_t . The customers then place their demands z_t . The retailer's status at the beginning of the planning horizon is given by the parameter I_1 (initial inventory). Apart from the ordering cost, the following costs are incurred over the planning horizon: (i) holding cost $h_t(I_{t+1})^+$, where h_t is the unit holding cost, (ii) shortage cost $s_t(-I_{t+1})^+$, where s_t is the unit shortage cost.

Inventory I_{t+1} left at the end of stage T has a unit salvage value s . Also, one must impose $h_t - s \geq -s_t$ to maintain the problem's convexity. A practical interpretation of this constraint is that in the last stage it is more profitable to satisfy the customer demand rather than to be left with an excessive amount of inventory. The constraints in the model include (i) balance equations linking the inventory in each stage to the inventory, order quantity, and demand in the preceding stage, (ii) upper and lower bounds on the order quantities in each stage $L_t \leq x_t \leq U_t$, (iii) upper and lower bounds on the cumulative order quantity up to stage $\hat{L}_t \leq \sum_{\tau=1}^t x_\tau \leq \hat{U}_t$.

The problem to solve without uncertainty in the demand is

$$\begin{aligned} \min_x \quad & \sum_{t=1}^T \{c_t x_t + h_t(I_{t+1})^+ + s_t(-I_{t+1})^+\} \\ \text{s.t.} \quad & I_{t+1} = I_t + x_t - z_t, & t = 1, \dots, T \\ & L_t \leq x_t \leq U_t, & t = 1, \dots, T \\ & \hat{L}_t \leq \sum_{\tau=1}^t x_\tau \leq \hat{U}_t, & t = 1, \dots, T. \end{aligned} \tag{31}$$

To model uncertainty about demands $z = (z_1, \dots, z_T)$, we assume that $\mathcal{Z}$ is the support defined as $\mathcal{Z} = \mathcal{Z}_1 \times \dots \times \mathcal{Z}_T$, where $\mathcal{Z}_t = [a_t, b_t]$, $t = 1, \dots, T$, which corresponds to z being a random variable with independent components. For the ambiguity set of the uncertain demand distribution, we set $\mu_t = (a_t + b_t)/2$, $\mathbb{E}_{\mathbb{P}}|z_t - \mu_t| = (b_t - a_t)/4$, and $\mathbb{P}(z_t \geq \mu_t) = \beta$ for $t = 1, \dots, T$. We use the same 50 problem instances with $T = 6$ as Postek et al. (2015); the ranges for the uniform sampling of parameters are given in Table 7.

Our goal is to obtain and compare decisions corresponding to various solution approaches for this multi-stage problem with distributional uncertainty. Among others, there are two questions related to a multi-stage problem in such a setting: (i) what should be the minimized objective criterion?; and (ii) how to make the later-stage decisions adjust to the observed demands?

Table 7 Inventory experiment - ranges for parameter sampling in the inventory experiment.

Parameter	Range	Parameter	Range	Parameter	Range
a_t	$[0, 20]$	L_t	0	$\hat{U}_t$	$0.8 \sum_{i=1}^T U_t$
b_t	$[a_t, a_t + 100]$	U_t	$[50, 70]$	h_t	$[0, 5]$
c_t, s_t	$[0, 10]$	$\hat{L}_t$	0	I_1	$[20, 50]$

With respect to the first question, we propose to minimize either the worst-case expectation as the most pessimistic value in our setting, or the best-case expectation as the most optimistic one. The closed-form formulation of the worst-case and best-case expectations of the objective function in (31) can be evaluated using the methodology of Section 2.2.

The second issue – adjustability of decisions – is important as (i) in this way the later-stage decisions can be better for each outcome of the uncertainty, (ii) the best here-and-now decisions can be different if later-stage adjustability is accounted for. Adjustability is typically achieved by formulating the later-stage decisions as functions (*decision rules*) of the realized demand and then, optimizing the parameters of these functions as decision variables

A simple and common choice for the decision rules is to define them as linear functions $x_{t+1} := x_{n+1,0} + \sum_{j=1}^t x_{n+1,j} z_j$, where $t = 1, \dots, T$, of the observed uncertainties. However, as affinity may be too restrictive, we propose also piecewise-linear decision rules obtained by interpolating the decisions in the finite worst-case (or best-case) support, described in detail in Appendix D.1. If the coefficients of the decision rules are determined before the planning horizon and they are not altered later, we denote this approach by *evaluation*.

Alternatively to fixing the decision rules and not changing them later, it is possible, having found the optimal solution and implemented the initial decision x_1 and observed z_1 , to solve a new optimization problem where new decisions (and decision rules) for stages $2, \dots, T$ are determined. In contrast to the *evaluation approach*, we denote it as *reoptimization*.

In this setting, we consider three solutions:

- minimizing the worst-case expectation with linear decision rules (results are taken from Postek et al. (2015)). We denote this approach as L-WCE (‘L’ stands for ‘linear’ and ‘WCE’ for ‘worst-case expectation’);
- minimizing the worst-case expectation using the piecewise-linear decision rules of Appendix D.1. We denote this approach as PL-WCE (‘PL’ stands for ‘piecewise-linear’);
- minimizing the best-case expectation using the piecewise-linear decision rules of Appendix D.1. We denote this approach as PL-BCE- β (‘BCE’ for ‘best-case expectation’ and β is the skewness parameter for each t as defined in (2), which we assume to be the same for all $t = 1, \dots, T$).

5.3.2. Intervals for the expected value of total cost Due to the distributional uncertainty in our problem, it is not possible to know the exact value of the expected total cost incurred over the planning horizon. However, due to the convexity of the objective function in z it is possible to evaluate both the worst-case and the best-case expectations of the total cost, which gives an interval for the expectation of the objective. Such intervals allow us to compare the three solutions with respect to (i) optimality: minimizing expected costs, lower values are preferable, (ii) range: the narrower an interval, the less ambiguity about the ‘true’ expected cost.

Intervals for solutions in the evaluation approach. Table 8 presents the results on the performance of the three solutions. The PL-WCE solution achieves a better worst-case (maximum total cost over the entire support) objective value (2358 versus 2384), and for each β the upper and lower endpoints of the interval for PL-WCE are smaller than upper and lower endpoints of the interval for L-WCE, for example, $[943, 1007]$ versus $[970, 1049]$ for $\beta = 0.5$. This provides strong evidence that restricting the decision rules to linear functions can have a negative effect on the quality of the solution as measured by the objective function.

We now compare the widths of the intervals corresponding to different solutions, which is our proxy for the value of distributional information and the ‘riskiness’ of each solution. We observe that PL-BCE solutions give expectation intervals that are overall much more dispersed than the PL-WCE solutions, compare e.g. $[908, 1133]$ (width 225) and $[943, 1007]$ (width 64) in the third row. On average, the intervals corresponding to the PL-BCE solutions are 5 times wider than the ones from PL-WCE solutions. This indicates that minimization of the worst-case expectation (pessimistic approach) may have a ‘compressing’ impact on the expectation interval, whereas the solutions obtained by minimizing the best-case expectation (optimistic approach) come with a much wider range.

With respect to the value of distributional information (VDI), we can approximate it as follows on the example of the PL-WCE solution. The width of the interval is for $\beta = 0.5$ is given by $1007 - 943 = 64$ which is the VDI. This value, divided by the upper bound on the worst-case performance yields $64/1007 \times 100\% \approx 6.35\%$ – it is the remaining relative uncertainty about the objective expectation. It is questionable whether profits can be gained by knowing or gathering exact data on the distribution since (i) computational handling of this extra information in the optimization problem would be significantly more complicated, (ii) the resulting more precise expectation value would be much more sensitive to estimation errors.

Intervals for solution values in the reoptimization approach. We also consider the intervals for the objective function value assuming that the decision maker can reoptimize the solution over time, described in more detail in Appendix D.2. The results are given in Table 9. Compared to Table 8, it is clear that for each solution and each value of β the corresponding upper and

Table 8 Inventory experiment - evaluation intervals - ranges for the expectation of the objective over $\mathcal{P}_z$ ('expectation range' is computed for a given solution using the upper and lower bound results of (BTH72) under given assumptions) and worst-case cost ('worst-case value' is the maximum total cost obtained for the single worst-case scenario out of $\mathcal{Z}$). All numbers are averages over the 50 instances.

Objective type	β	Solution				
		L-WCE	PL-WCE	PL-BCE-0.25	PL-BCE-0.5	PL-BCE-0.75
Expectation range	0.25	[1004, 1049]	[973, 1007]	[940 , 1178]	[976, 1133]	[1087, 1228]
Expectation range	0.5	[970, 1049]	[943, 1007]	[1009, 1178]	[908 , 1133]	[978, 1228]
Expectation range	0.75	[994, 1049]	[960, 1007]	[1157, 1178]	[986, 1133]	[905 , 1228]
Worst-case value	-	2384	2358	2628	2553	3005

Table 9 Inventory experiment - reoptimization intervals. All numbers are averages over the 50 instances.
Terminology as in Table 8.

Objective type	β	Solution				
		L-WCE	PL-WCE	PL-BCE-0.25	PL-BCE-0.5	PL-BCE-0.75
Expectation range	0.25	[972, 1011]	[965, 1007]	[940 , 1040]	[975, 1037]	[995, 1038]
Expectation range	0.5	[941, 1011]	[933, 1007]	[938, 1040]	[903 , 1037]	[933, 1038]
Expectation range	0.75	[955, 1011]	[952, 1007]	[982, 1040]	[942, 1037]	[905 , 1038]

lower bounds of the intervals are not larger than the ones from Table 8 (compare for example the lower endpoints for L-WCE). Partly due to this change, the intervals obtained for various solutions become more similar.

5.3.3. Simulation study Apart from simply knowing the intervals to which the expected total costs are guaranteed to belong, one may be interested in the performance of the three solutions in a 'reasonable' simulation setting. Since we do not know the exact distributions of the uncertain random parameters, we use the following two distributions to sample from:

- uniform sample: demand scenarios $\hat{z}$ are sampled from a uniform distribution on the support $\mathcal{Z}$;
- (μ, d) sample: demand scenarios $\hat{z}$ are sampled from a randomly sampled distribution $\hat{\mathbb{P}} \in \mathcal{P}_z$ – the details of the sampling methodology are given in Appendix D.3.

As the (μ, d) sample involves the distributional uncertainty, it is 'broader' than the uniform sample, i.e. it encompasses more than one possible choice for the probability distribution out of the given ambiguity set.

Evaluation. Table 10 presents the results in the evaluation approach. The PL-WCE solution again gives better values than L-WCE, both in terms of the mean values - an improvement of 2.81% on the uniform sample, and the standard deviations of the objective function value - for example, an improvement of 3.01% on the (μ, d) sample. Interesting results are also obtained for the PL-BCE solutions: on the (μ, d) sample they perform better than the L-WCE and PL-WCE solutions, despite their focus on the best-case expectation. Also, the PL-BCE solutions provide substantial

Table 10 Inventory experiment - evaluation simulation results. Numbers in brackets denote the % change compared to the L-WCE solution. All numbers are averages over all 50 problem instances.

Value	Sample	Solution				
		L-WCE	PL-WCE	PL-BCE-0.25	PL-BCE-0.5	PL-BCE-0.75
Mean	Uniform	994	966 (−2.81%)	1019 (+2.54%)	999 (+0.54%)	1026 (+3.29%)
Standard deviation	Uniform	259	251 (−3.08%)	272 (+5.09%)	255 (−1.45%)	286 (10.49%)
Mean	(μ, d)	1003	971 (−3.19%)	976 (−2.66%)	962 (−4.00%)	986 (−1.60%)
Standard deviation	(μ, d)	265	257 (−3.01%)	241 (−8.83%)	223 (−15.62%)	248 (−6.40%)

Table 11 Inventory experiment - reoptimization simulation results. Numbers in brackets denote the % change compared to the L-WCE solution. All numbers are averages over all 50 instances.

Value	Solution				
	L-WCE	PL-WCE	PL-BCE-0.25	PL-BCE-0.5	PL-BCE-0.75
Mean value	977.8	972.8 (−0.5%)	986.6 (+0.90%)	974.5 (−0.34%)	980.6 (+0.29%)
Standard deviation	251.2	257.5 (+2.51%)	272.6 (+8.51%)	237.7 (−5.34%)	236.0 (−6.03%)

decreases in the standard deviation of the estimator of the expected total costs compared to the other solutions.

Reoptimization. We also consider the reoptimization version of our experiment where we use 500 demand samples from the (μ, d) distributions (we do not report on the results for the uniform sample as they are nearly the same). Table 11 presents the results. The means of the simulated total cost are almost the same for all solutions, differing by less than 1%. This small difference is in line with the results of the previous subsection where the intervals in the reoptimization experiment turn out to be similar as well.

Choosing the right solution. The results of this and the previous subsection give rise to the question whether any of the three solutions is preferable to others. We suggest that this choice depends on three factors: (i) risk-aversion of the decision maker, (ii) amount of computational power available, and (iii) possibility (or not) to re-optimize.

For a risk-averse decision maker, the PL-WCE solution is more flexible (the implied decision rule is piecewise-linear instead of linear) than the L-WCE and gives better worst-case expected performance in evaluation settings, as shown in Table 8. On the other hand, it requires a larger computational effort (second criterion) as each worst-case demand scenario requires a separate ordering trajectory. The number of optimization variables in the problem with L-WCE decisions equals $T(T+1)/2$ (one variable for time 1 decisions, 2 for time 2 decisions, etc.) whereas for PL-WCE this number equals $(3^T - 1)/2$ (enumerating all the 3^T trajectories and elimination of some double-counted decisions through the *nonanticipativity constraints*). With respect to the third criterion, if reoptimization is possible, then we see that the differences between the three solutions are very small.

Figure 2 Inventory experiment - upper (WCE) and lower (BCE) bounds on the expectation of the objective function values for a single problem instance (left panel) and aggregated over all 50 instances (right panel). The best-case expectations (lower bounds) have been computed for $\beta \in \{0.25, 0.5, 0.75\}$.

5.3.4. Pareto stripe A common tool in decision support is the Pareto curve, illustrating a tradeoff between two criteria. It is obtained by finding, for fixed bounds on one objective (for example, the worst-case cost over the entire demand support), the minimum of another objective (for example, the mean cost). A strong feature of our approach is that it allows to evaluate both the worst-case and the best-case expectation of a convex function. That is, for a given bound on the worst-case value of the total cost, we are able to identify the entire interval for the expected cost. This gives rise to an extension of the Pareto curve, denoted as the *Pareto stripe*, which in our case depicts how a bound on the worst-case cost affects (i) the best(worst)-case expectation, (ii) the value of distributional information, as measured by the width of this interval.

Mathematically, the Pareto stripe is obtained by minimizing, for a given (fixed) upper bound $C \in \mathbb{R}$ on the worst-case value of the objective function:

$$\sup_{z \in \mathcal{Z}} \left\{ \sum_{t=1}^T (c_t x_t + h_t (I_{t+1})^+ + s_t (-I_{t+1})^+) \right\} \leq C, \quad (32)$$

the worst-case expectation (or the best-case expectation) by means of the PL-WCE solution (PL-BCE solution, respectively, with different possible values for the skewness β).

The left panel in Figure 2 presents such a stripe for a single problem instance. An interesting feature is that the best-case expectations obtained for various values of parameter β need not preserve any monotonicity relation. For example, the best-case expectation when $\beta = 0.75$, obtained for the worst-case bound C of around 2400 (horizontal axis) is the highest of all, but is smaller than best-case expectations for $\beta = 0.25$ and $\beta = 0.5$ when the worst-case bound C is 3000.

The right panel in Figure 2 presents the Pareto stripe aggregated over all 50 problem instances. For orientation, we note that the rightmost value of the continuous black curve in the right panel

of Figure 2 corresponds to the first row of Table 8, whereas the rightmost value of the horizontal axis is the worst-case objective value of PL-WCE solutions from Table 8. Figure 2 provides an assessment of the value of distributional information. We observe that as the bound on the worst-case performance (horizontal axis) grows, the width of the Pareto stripe increases slightly, corresponding to a growth in the VDI (width of the interval compared to its upper bound value) from about 7% to about 11%.

6. Conclusion

In this paper we have considered stochastic programming problems with distributional ambiguity. We have shown that under mean - MAD distributional information, the problem admits a closed form reformulation as the corresponding worst-case distributions consist of 3 points per component and is independent from the first stage decisions. This holds both for two-stage and multi-stage continuous models. We have proposed methods to deal with the exponential number of scenarios that perform well in the numerical experiments. For two-stage problems with integer recourse variables, we show how good convex approximations can be derived that have a provable performance guarantee. Our numerical experiments entailing operating room scheduling and inventory management provide also simple yet powerful managerial insights such as (i) the easy-to-calculate value of distributional information (difference between the worst- and best-case expectation under the given information) and (ii) the Pareto stripe, which shows how the interval containing the true expected objective function changes relative to a bound on a certain performance measure. Overall, we have proposed a practical framework of solving a wide class of problems that can easily be implemented in a variety of real-world applications.

References

- Ahmed, S., M. Tawarmalani, N.V. Sahinidis. 2004. A finite branch-and-bound algorithm for two-stage stochastic integer programs. *Mathematical Programming* **100** 355–377.
- Bayraksan, G., D.P. Morton. 2009. Assessing solution quality in stochastic programs via sampling. *Tutorials in Operations Research*. INFORMS, 102–122.
- Ben-Tal, A., D. den Hertog, J.-Ph. Vial. 2015. Deriving robust counterparts of nonlinear uncertain inequalities. *Mathematical Programming* **1-2** 265–299.
- Ben-Tal, A., L. El Ghaoui, A. Nemirovski. 2009. *Robust optimization*. Princeton University Press.
- Ben-Tal, A., B. Golany, A. Nemirovski, J.-Ph. Vial. 2005. Retailer-supplier flexible commitments contracts: a robust optimization approach. *Manufacturing & Service Operations Management* **7(3)** 248–271.
- Ben-Tal, A., A. Goryashko, E. Guslitzer, A. Nemirovski. 2004. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming* **99(2)** 351–376.

- Ben-Tal, A., E. Hochman. 1972. More bounds on the expectation of a convex function of a random variable. *Journal of Applied Probability* **9** 803–812.
- Ben-Tal, A., E. Hochman. 1976. Stochastic programs with incomplete information. *Operations Research* **24**(2) 336–347.
- Bertsimas, D., A. Georghiou. 2015. Design of near optimal decision rules in multistage adaptive mixed-integer optimization. *Operations Research* **63**(3) 610–627.
- Bertsimas, D., D.A. Iancu, P.A. Parrilo. 2011. A hierarchy of near-optimal policies for multistage adaptive optimization. *IEEE Transactions on Automatic Control* **56**(12) 2809–2824.
- Bertsimas, D., E. Litvinov, X.A. Sun, J. Zhao, T. Zheng. 2013. Adaptive robust optimization for the security constrained unit commitment problem. *IEEE Transactions on Power Systems* **28**(1) 52–63.
- Birge, J.R. 1982. The value of the stochastic solution in stochastic linear programs with fixed recourse. *Mathematical Programming* **24**(3) 314–325.
- Birge, J.R. 1985. Decomposition and partitioning methods for multistage stochastic linear programs. *Operations Research* **33** 989–1007.
- Birge, J.R., F. Louveaux. 1997. *Introduction to Stochastic Programming*. Springer, New York.
- Carøe, C.C., R. Schultz. 1999. Dual decomposition in stochastic integer programming. *Operations Research Letters* **24** 37–45.
- Chen, X., Y. Zhang. 2009. Uncertain linear programs: Extended affinely adjustable robust counterparts. *Operations Research* **57**(6)(6) 1469–1482.
- Delage, E., S. Arroyo, Y. Ye. 2015. The value of distribution information in distributionally robust optimization. <http://web.hec.ca/pages/erick.delage/ValueDistributionInfo.pdf>.
- Delage, E., Y. Ye. 2010. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research* **58**(3) 595–612.
- Denton, B.T., A.J. Miller, H.J. Balasubramanian, T.R. Huschka. 2010. Optimal allocation of surgery blocks to operating rooms under uncertainty. *Operations Research* **58** 802–816.
- Dupačová, J., N. Gröwe-Kuska, W. Römisch. 2003. Scenario reduction in stochastic programming: An approach using probability metrics. *Mathematical Programming* **95** 493–511.
- Dyer, M., L. Stougie. 2006. Computational complexity of stochastic programming problems. *Mathematical Programming* **106** 423–432.
- Edirisinghe, N.C.P. 2011. *Stochastic Programming: The State of the Art In Honor of George B. Dantzig*, chap. Stochastic Programming Approximations Using Limited Moment Information, with Application to Asset Allocation. Springer New York, 97–138.
- Edmundson, H.P. 1956. Bounds on the expectation of a convex function of a random variable. Technical Report Paper 982, The RAND Corporation.

- Fiacco, A.V., J. Kyparisis. 1986. Convexity and concavity properties of the optimal value function in parametric nonlinear programming. *Journal of Optimization Theory and Applications* **48(1)** 95–126.
- Garstka, S.J., R.J.-B. Wets. 1974. On decision rules in stochastic programming. *Mathematical Programming* **7(1)** 117–143.
- Gul, S., B.T. Denton, J.W. Fowler, T. Huschka. 2011. Bi-criteria scheduling of surgical services for an outpatient procedure center. *Production and Operations Management* **20** 406–417.
- Hanasusanto, G.A., D. Kuhn, W. Wiesemann. 2015. K -adaptability in two-stage robust binary programming. *Operations Research* **63(4)** 877–891.
- Heitsch, H., W. Römisch. 2003. Scenario reduction algorithms in stochastic programming. *Computational Optimization and Applications* **24** 187–206.
- Higle, J.L., S. Sen. 1991. Stochastic decomposition: An algorithm for two-stage linear programs with recourse. *Mathematics of Operations Research* **16** 650–669.
- Isii, K. 1963. On sharpness of Tchebycheff-type inequalities. *Annals of the Institute of Statistical Mathematics* **14(1)** 185–197.
- Jensen, J.L.W.V. 1906. Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Mathematica* **30** 175–193.
- Kemperman, J.H.B. 1968. The general moment problem, a geometric approach. *Annals of Mathematical Statistics* **39** 93–122.
- Klein Haneveld, W.K., L. Stougie, M.H. van der Vlerk. 2006. Simple integer recourse models: convexity and convex approximations. *Mathematical Programming* **108** 435–473.
- Klein Haneveld, W.K., M.H. van der Vlerk. 1999. Stochastic integer programming: General models and algorithms. *Annals of Operations Research* **85** 39–57.
- Klein Haneveld, W.K., M.H. van der Vlerk. 2006. Integrated chance constraints: reduced forms and an algorithm. *Computational Management Science* **3(4)** 245–269.
- Knight, F. 1921. *Risk, uncertainty and profit*. New York.
- Laporte, G., F.V. Louveaux. 1993. The integer L-shaped method for stochastic integer programs with complete recourse. *Operations Research Letters* **13** 133–142.
- Louveaux, F.V., M.H. van der Vlerk. 1993. Stochastic programming with simple integer recourse. *Mathematical Programming* **61** 301–325.
- Pereira, M.V.F., L.M.V.G. Pinto. 1991. Multi-stage stochastic optimization applied to energy planning. *Mathematical Programming* **52** 359–375.
- Pflug, G.Ch. 2001. Scenario tree generation for multiperiod financial optimization by optimal discretization. *Mathematical Programming* **89** 251–271.

- Postek, K., A. Ben-Tal, D. den Hertog, B. Melenberg. 2015. Exact robust counterparts of ambiguous stochastic constraints under mean and dispersion information. CentER Discussion Paper No. 2015-030, available at: http://www.optimization-online.org/DB_FILE/2015/06/4946.pdf.
- Postek, K., D. den Hertog. 2016. Multi-stage adjustable robust mixed-integer optimization via iterative splitting of the uncertainty set. *INFORMS Journal on Computing* **28(3)** 553–574.
- Prékopa, A. 1995. *Stochastic Programming*. Kluwer Academic Publishers, Dordrecht.
- Rockafellar, R.T., R.J.-B. Wets. 1991. Scenarios and policy aggregation in optimization under uncertainty. *Mathematics of Operations Research* **16** 119–147.
- Romeijnders, W., R. Schultz, M.H. van der Vlerk, W.K. Klein Haneveld. 2016a. A convex approximation for two-stage mixed-integer recourse models with a uniform error bound. *SIAM Journal on Optimization* **26** 426–447.
- Romeijnders, W., L. Stougie, M.H. van der Vlerk. 2014. Approximation in two-stage stochastic integer programming. *Surveys in Operations Research and Management Science* **19** 17–33.
- Romeijnders, W., M.H. van der Vlerk, W.K. Klein Haneveld. 2015. Convex approximations of totally unimodular integer recourse models: A uniform error bound. *SIAM Journal on Optimization* **25** 130–158.
- Romeijnders, W., M.H. van der Vlerk, W.K. Klein Haneveld. 2016b. Total variation bounds on the expectation of periodic functions with applications to recourse approximations. *Mathematical Programming* **157** 3–46.
- Römisch, W. 2003. Stability of stochastic programming problems. A. Ruszczyński, A. Shapiro, eds., *Handbooks in Operations Research and Management Science*, vol. 10: Stochastic Programming. Elsevier, 483–554.
- Römisch, W., R. Schultz. 1996. Decomposition of a multi-stage stochastic program for power dispatch. *ZAMM-Zeitschrift für Angewandte Mathematik und Mechanik* **76(3)** 2932.
- Ruszczyński, A. 1986. A regularized decomposition method for minimizing a sum of polyhedral functions. *Mathematical Programming* **35** 309–333.
- Scarf, H.E. 1958. A min-max solution of an inventory problem. K. J. Arrow, S. Karlin, H. E. Scarf, eds., *Studies in the Mathematical Theory of Inventory and Production*. Stanford University Press, 201–209.
- Schultz, R. 2003. Stochastic programming with integer variables. *Mathematical Programming* **97** 285–309.
- Sen, S. 2005. Algorithms for stochastic mixed-integer programming models. K. Aardal, G.L. Nemhauser, R. Weismantel, eds., *Handbooks in Operations Research and Management Science*, vol. 12: Discrete Optimization. Elsevier, 515–558.
- Shapiro, A. 2001. *On duality theory of conic linear problems*. Springer US, 135–165.
- Shapiro, A., S. Ahmed. 2004. On a class of minimax stochastic programs. *SIAM Journal on Optimization* **14(4)** 1237–1249.

- Shapiro, A., D. Dentcheva, A. Ruszczyński. 2009. *Lectures on Stochastic Programming: Modeling and Theory*. SIAM, Philadelphia.
- Shapiro, A., A. Kleywegt. 2002. Minimax analysis of stochastic problems. *Optimization Methods and Software* **17(3)** 523–542.
- Smith, R. L. 1984. Efficient monte-carlo procedures for generating points uniformly distributed over bounded regions. *Operations Research* **32** 1296–1308.
- van der Vlerk, M.H. 2004. Convex approximations for complete integer recourse models. *Mathematical Programming* **99** 297–310.
- van Slyke, R.M., R. Wets. 1969. L-shaped linear programs with applications to optimal control and stochastic programming. *SIAM Journal on Applied Mathematics* **17** 638–663.
- Žáčková, J. 1966. On minimax solutions of stochastic linear programming problems. *Časopis Pro Pěstování Matematiky* **91** 423–430.
- Wets, R.J.-B. 1983. Solving stochastic programs with simple recourse. *Stochastics* **10** 219–242.
- Wiesemann, W., D. Kuhn, M. Sim. 2014. Distributionally robust convex optimization. *Operations Research* **62(6)** 1358–1376.
- Zverovich, V., C.I. Fábián, E.F.D. Ellison, G. Mitra. 2012. A computational study of a solver system for processing two-stage stochastic LPs with enhanced Benders decomposition. *Mathematical Programming Computation* **4** 211–238.

Appendix

A. Proofs

Proof of Proposition 1. For simplicity, assume first that $n_\omega = 1$ and $n_\xi = 0$. Since $v(x, z)$ is a convex function of z then by result of BTH72 we have that:

$$\begin{aligned} \sup_{\mathbb{P}_{z_1} \in \mathcal{P}_{z_1}} \mathbb{E}_{\mathbb{P}_{z_1}} v(x, z_1) &= \\ &= \frac{d_1}{2(\mu_1 - a_1)} v(x, a_1) + \left(1 - \frac{d_1}{2(\mu_1 - a_1)} - \frac{d_1}{2(b_1 - \mu_1)}\right) v(x, \mu_1) + \frac{d_1}{2(b_1 - \mu_1)} v(x, b_1), \end{aligned} \quad (33)$$

that is, the worst-case expectation of $v(x, z)$ is achieved by a three-point distribution with support $\{a_1, \mu_1, b_1\}$ and probabilities $d_i/2(\mu_i - a_i)$, $1 - d_i/2(\mu_i - a_i) - d_i/2(b_i - \mu_i)$ and $d_i/2(b_i - \mu_i)$, respectively. For $n_z \geq 2$ we observe that due to the independence of z_i 's we have:

$$\mathcal{P}_z = \mathcal{P}_{z_1} \times \dots \times \mathcal{P}_{z_{n_\omega}},$$

where

$$\mathcal{P}_{z_i} = \{\mathbb{P}_{z_i} : \text{supp}(z_i) \subseteq [a_i, b_i], \quad \mathbb{E}_{\mathbb{P}}(z_i) = \mu, \quad \mathbb{E}_{\mathbb{P}}|z_i - \mu_i| = d_i\}, \quad i = 1, \dots, n_\omega.$$

Therefore, we can apply formula (34) component-wise w.r.t. z :

$$\sup_{\mathbb{P}_z \in \mathcal{P}_z} \mathbb{E}_{\mathbb{P}_z} v(x, z) = \sup_{\mathbb{P}_{z_1} \in \mathcal{P}_{z_1}} \mathbb{E}_{\mathbb{P}_{z_1}} \left\{ \sup_{\mathbb{P}_{z_2} \in \mathcal{P}_{z_2}} \mathbb{E}_{\mathbb{P}_{z_2}} \left\{ \dots \sup_{\mathbb{P}_{z_{n_\omega}} \in \mathcal{P}_{z_{n_\omega}}} \mathbb{E}_{\mathbb{P}_{z_{n_\omega}}} v(x, (z_1, \dots, z_{n_\omega})) \dots \right\} \right\}.$$

Therefore, the support of the worst-case distribution of z is a product of the worst-case distributions of z_i , equal to $\bar{\mathcal{Z}}$, and the probability of a single z^k is equal to the product of the worst-case probabilities of the respective components of z^k , as defined in Proposition 1. A similar argument holds for the worst-case expectation w.r.t. ξ and since it is assumed that the components of ω and ξ are mutually independent, the claim follows.

Proof of Proposition 2. The proof is analogous to Proposition 1, therefore, we only consider the case $n_\omega = 1$ and $n_\xi = 0$. Since $v(x, z)$ is a convex function of z then by result of BTH72 we have that:

$$\begin{aligned} \inf_{\mathbb{P}_{z_1} \in \mathcal{P}_{z_1}} \mathbb{E}_{\mathbb{P}_{z_1}} v(x, z_1) &= \\ &= (1 - \beta_1)v(x, \mu_1 - d_1/2(1 - \beta_1)) + \beta_1 v(x, \mu_1 + d_1/2\beta_1), \end{aligned} \quad (34)$$

that is, the best-case expectation of $v(x, z)$ is achieved by a two-point distribution with support $\{\mu_1 - d_1/2(1 - \beta_1), \mu_1 + d_1/2\beta_1\}$ and probabilities $(1 - \beta_1)$ and β_1 , respectively.

Proof of Proposition 3. Consider first problem (6) solved at time T . The problem is linear, hence convex in x_T and by Fiacco and Kyparisis (1986) it holds that the optimal value of (6) is convex in x_{T-1} and z_{T-1} . Next, consider the problem to solve at time $T - 1$:

$$\inf_{x_{T-1}} \left\{ c_{T-1}^T x_{T-1} + \sup_{\mathbb{P}_{z_{T-1}} \in \mathcal{P}_{z_{T-1}}} \mathbb{E}_{\mathbb{P}_{z_{T-1}}} v_T(x_{T-1}) : \sum_{s=1}^{T-2} A_{T-1s}(z_{T-2})x_s + A_{T-1T-1}x_{T-1} = b_{T-1}(z_{T-2}) \right\}. \quad (35)$$

Since $v_T(x_{T-1}, z_{T-1})$ is convex in x_{T-1} , the objective function in (35) is also convex in x_{T-1} and, since the remaining constraints are linear in x_{T-1} , the problem is convex in x_{T-1} . Again, by Fiacco and Kyparisis (1986) it holds that $v_{T-1}(x_{T-2}, z_{T-2})$ is convex in z_{T-2} and x_{T-2} . The same argument is applied recursively to time stages $T - 1, T - 2, \dots, 2$ which proves the claim.

Proof of Theorem 1. Let x^* denote an optimal solution to the optimization problem in (10). Then,

$$\begin{aligned} \eta^* \leq G(\hat{x}) &= \sup_{\mathbb{P}_z \in \mathcal{P}_z} \mathbb{E}_{\mathbb{P}_z} \left[c^T \hat{x} + Q(\hat{x}, z) \right] \\ &\leq \sup_{\mathbb{P}_z \in \mathcal{P}_z} \mathbb{E}_{\mathbb{P}_z} \left[c^T \hat{x} + \hat{Q}(\hat{x}, z) \right] + \|Q - \hat{Q}\|_\infty \\ &= \hat{\eta} + \|Q - \hat{Q}\|_\infty. \end{aligned}$$

Here, the first inequality holds since $\hat{x}$ is not necessarily optimal in the original model. Similarly, we have

$$\begin{aligned} \hat{\eta} &\leq \sup_{\mathbb{P}_z \in \mathcal{P}_z} \mathbb{E}_{\mathbb{P}_z} \left[c^T x^* + \hat{Q}(x^*, z) \right] \\ &\leq \sup_{\mathbb{P}_z \in \mathcal{P}_z} \mathbb{E}_{\mathbb{P}_z} \left[c^T x^* + Q(x^*, z) \right] + \|Q - \hat{Q}\|_\infty \\ &= \eta^* + \|Q - \hat{Q}\|_\infty. \end{aligned}$$

Combining $\eta^* \leq \hat{\eta} + \|Q - \hat{Q}\|_\infty$ and $\hat{\eta} \leq \eta^* + \|Q - \hat{Q}\|_\infty$ yields the first inequality in (i). Furthermore, using $G(\hat{x}) \leq \hat{\eta} + \|Q - \hat{Q}\|_\infty$ and $\hat{\eta} \leq \eta^* + \|Q - \hat{Q}\|_\infty$, it follows that

$$G(\hat{x}) \leq \eta^* + 2\|Q - \hat{Q}\|_\infty,$$

and from this the second inequality in (i) follows immediately.

Further, observe that $\tilde{\eta}$ is a lower bound for η^* since $\mathbb{P}_{\bar{z}}$ is not necessarily the worst-case distribution in model (10), and thus $0 \leq \eta^* - \tilde{\eta}$. Next, let $\hat{Q}$ be a convex approximation of Q . Then, the remaining inequalities in (ii) follow directly from

$$\begin{aligned}
 \eta^* \leq G(\tilde{x}) &= \sup_{\mathbb{P}_z \in \mathcal{P}_z} \mathbb{E}_{\mathbb{P}_z} \left[c^T \tilde{x} + Q(\tilde{x}, z) \right] \\
 &\leq \sup_{\mathbb{P}_z \in \mathcal{P}_z} \mathbb{E}_{\mathbb{P}_z} \left[c^T \tilde{x} + \hat{Q}(\tilde{x}, z) \right] + \|Q - \hat{Q}\|_{\infty} \\
 &= \mathbb{E}_{\mathbb{P}_{\bar{z}}} \left[c^T \tilde{x} + \hat{Q}(\tilde{x}, \bar{z}) \right] + \|Q - \hat{Q}\|_{\infty} \\
 &\leq \mathbb{E}_{\mathbb{P}_{\bar{z}}} \left[c^T \tilde{x} + Q(\tilde{x}, \bar{z}) \right] + 2\|Q - \hat{Q}\|_{\infty} \\
 &= \tilde{\eta} + 2\|Q - \hat{Q}\|_{\infty}.
 \end{aligned} \tag{36}$$

where the first inequality holds since $\tilde{x}$ is not necessarily optimal in model (10), and where we apply the result of BTH72 in (36).

B. Solution methods for the OR experiment of Section 5.1

In this appendix we discuss several solution methods for solving the OR scheduling problem.

B.1. SAA method

The main difficulty in solving the problem in (29) is to deal with the 3^N scenarios of the surgery durations $\bar{z}$. A well known approach in the SP literature to circumvent this difficulty is to use sampling to approximate $\bar{z}$ by a random vector having fewer scenarios. Thus, we sample N_s scenarios from the worst-case distribution $\mathbb{P}_{\bar{z}}$ to obtain the sample z^s , $s = 1, \dots, N_s$. Then, letting $\theta_{i,s}$ denote the minutes of overtime in the i -th OR under scenario s , we can derive a large-scale deterministic equivalent MILP formulation:

$$\begin{aligned}
 (DEF) \quad \min_{x,y,\theta} \quad & \sum_{i=1}^N c_f x_i + \frac{1}{N_s} \sum_{s=1}^{N_s} \sum_{i=1}^N c_v \theta_{i,s} \\
 \text{s.t.} \quad & \theta_{i,s} \geq \sum_{j=1}^n z_j^s y_{ij} - T x_i, & i = 1, \dots, N, s = 1, \dots, N_s, \\
 & (x, y) \in X, \theta_{i,s} \geq 0, & i = 1, \dots, N, s = 1, \dots, N_s.
 \end{aligned}$$

This deterministic equivalent formulation contains $N + N^2$ binary variables, corresponding to x and y , and $N \times N_s$ continuous variables, corresponding to θ . In the numerical experiments in Section 5.1.2 we solve this MILP for $N_s = 1000$ and $N = 10$ or $N = 15$ using Gurobi. For these parameters the number of binary variables in the deterministic equivalent formulation is small so that the MILP can be solved within reasonable time limits. Of course, using decomposition algorithms, such as for example an L-shaped algorithm, we may solve this model faster. However, we prefer to use the current method to show that good solutions may be obtained (for problems of reasonable size) using this straightforward, easy-to-implement algorithm.

B.2. Linear decision rules

Another way to deal with the 3^N scenarios is to use *linear decision rules* for the overtime costs θ_i . Instead of using the exact expression for θ_i , given in (27), we approximate θ_i by an affine function $\hat{\theta}_i$ of $\bar{z}$:

$$\hat{\theta}_i(z) = u_i + \sum_{j=1}^N v_{ij} \bar{z}_j. \quad (37)$$

Here, u_i and v_{ij} denote the coefficients of the linear decision rule $\hat{\theta}_i$. These coefficients are determined here-and-now, i.e. at the same time as the surgery-to-OR assignment (x, y) . Hence, the first-stage decision variables in the resulting (approximating) optimization problem are $(x, y) \in X$ and the coefficients u and V , where $u = (u_1, \dots, u_N)^\top$ and V is a matrix containing the elements v_{ij} for $i, j = 1, \dots, N$. The approximating optimization problem using $\hat{\theta}_i$ instead of θ_i is given by

$$\begin{aligned} (LDR) \quad & \min_{x, y, u, V} \sum_{i=1}^N c_f x_i + c_v \mathbb{E}_{\bar{z}} \left[\sum_{i=1}^N \hat{\theta}_i(\bar{z}) \right] \\ \text{s.t.} \quad & \hat{\theta}_i(\bar{z}) \geq \sum_{j=1}^N \bar{z}_j y_{ij} - T x_i, & \bar{z} \in \mathcal{Z}, i = 1, \dots, N, \end{aligned} \quad (38)$$

$$\begin{aligned} & \hat{\theta}_i(\bar{z}) \geq 0, & \bar{z} \in \mathcal{Z}, i = 1, \dots, N, \\ & (x, y) \in X. \end{aligned} \quad (39)$$

Here, constraints (38) and (39) make sure that for every $i = 1, \dots, N$ the approximation $\hat{\theta}_i(\bar{z})$ for the overtime costs is non-negative and at least as large as the actual overtime costs $\theta_i(\bar{z})$ for all 3^N possible realizations of $\bar{z}$. Moreover, using the linear decision rule in (37), the expected overtime costs in the objective of (LDR) become

$$\begin{aligned} c_v \mathbb{E}_{\bar{z}} \left[\sum_{i=1}^N \hat{\theta}_i(\bar{z}) \right] &= c_v \mathbb{E}_{\bar{z}} \left[\sum_{i=1}^N \left(u_i + \sum_{j=1}^N v_{ij} \bar{z}_j \right) \right] \\ &= c_v \sum_{i=1}^N \left(u_i + \sum_{j=1}^N v_{ij} \mu_j \right) \\ &= c_v \sum_{i=1}^N u_i + c_v \sum_{i=1}^N \sum_{j=1}^N v_{ij} \mu_j. \end{aligned}$$

The optimization problem (LDR) can thus be rewritten as

$$\begin{aligned} (LDR) \quad & \min_{x, y, u, V} c_f \sum_{i=1}^N x_i + c_v \sum_{i=1}^N u_i + c_v \sum_{i=1}^N \sum_{j=1}^N v_{ij} \mu_j \\ \text{s.t.} \quad & u_i + \sum_{j=1}^N v_{ij} \bar{z}_j \geq \sum_{j=1}^N \bar{z}_j y_{ij} - T x_i, & \bar{z} \in \mathcal{Z}, i = 1, \dots, N, \end{aligned} \quad (40)$$

$$\begin{aligned} & u_i + \sum_{j=1}^N v_{ij} \bar{z}_j \geq 0, & \bar{z} \in \mathcal{Z}, i = 1, \dots, N, \\ & (x, y) \in X. \end{aligned} \quad (41)$$

Observe that this is a MILP with only a small number of decision variables, but with exponentially many constraints, since (40) and (41) are defined for every $\bar{z} \in \mathcal{Z}$, and $\bar{z}$ is a discrete random vector with 3^N realizations. However, since the convex hull of $\mathcal{Z}$ is a box uncertainty set, we can use standard techniques from Robust Optimization (Ben-Tal et al. 2009) to obtain the robust counterpart of this problem.

B.3. Adjusted version of the L-shaped algorithm

The L-shaped algorithm solves the optimization problem in (29) exactly. In this section we discuss our tailored implementation of the L-shaped algorithm, see e.g. van Slyke and Wets (1969). In this algorithm we approximate the expected overtime costs

$$Q(x, y) = c_v \mathbb{E}_{\bar{z}} \left[\sum_{i=1}^N \left(\sum_{j=1}^N \bar{z}_j y_{ij} - T x_i \right)^+ \right], \quad (x, y) \in X,$$

by an artificial decision variable $\eta \geq 0$. Using optimality cuts of the form

$$\eta \geq \pi^k x + \lambda^k y + \beta^k, \quad k = 1, \dots, K,$$

where π^k and λ^k are row vectors and β^k is a constant, the variable η will represent a lower bound of $Q(x, y)$.

At every iteration k , we solve the master problem

$$\begin{aligned} (\text{MP}_k) \quad & \min_{x, y, \eta} \quad c_f \sum_{i=1}^N x_i + \eta \\ \text{s.t.} \quad & \eta \geq \pi^l x + \lambda^l y + \beta^l, \quad l = 1, \dots, k-1, \\ & (x, y) \in X, \quad \eta \geq 0, \end{aligned}$$

which includes all optimality cuts of previous iterations, to obtain the current solution (x^k, y^k, η^k) . For this current solution, we evaluate $Q(x^k, y^k)$. Note that η^k is a lower bound of $Q(x^k, y^k)$ and that if $Q(x^k, y^k) = \eta^k$, then (x^k, y^k) is the optimal solution to (29). In general, we stop the algorithm if $Q(x^k, y^k) - \eta^k < \epsilon$ for some small number ϵ . If this optimality criterion does not hold, then we derive an optimality cut $\eta \geq \pi^k x + \lambda^k y + \beta^k$ and solve the master problem MP_{k+1} .

The challenge in this algorithm is to evaluate $Q(x^k, y^k)$ and to derive an optimality cut $\eta \geq \pi^k x + \lambda^k y + \beta^k$, dealing with the 3^N scenarios of the random vector $\bar{z}$. We only discuss why $Q(x^k, y^k)$ may be evaluated fast, despite this exponential number of scenarios. For similar reasons we may derive optimality cuts in an efficient way.

Evaluation of $Q(x^k, y^k)$ In this section we discuss the evaluation of $Q(x^k, y^k)$. For convenience we drop the index k and use (x, y) to refer to the current solution. We speed up the evaluation of $Q(x, y)$ in two different ways. First, by using the simple recourse structure of the problem, and second by using an efficient data structure for the scenarios.

Since the operating room scheduling problem can be modeled as a simple recourse model with uncertainty in the technology matrix, we can use the results of Klein Haneveld and van der Vlerk (2006) to evaluate $Q(x, y)$ faster. Similar as in Section 2 this is possible since we are not dealing with the joint distribution of $\bar{z}$ but with several marginal distributions of total surgery durations in operating rooms. To show this more formally, we first introduce additional notation.

Let $J_i = \{j : y_{ij} = 1\}$ denote the set of surgeries carried out in the i -th OR with $N_i := |J_i|$ denoting the number of surgeries. Moreover, let ζ^i denote the subvector of $\bar{z}$ containing these N_i surgeries. That is, ζ^i contains the j -th component of $\bar{z}$ if and only if $j \in J_i$. Then, by separability of the expected overtime costs,

$$\begin{aligned} Q(x, y) &= \mathbb{E}_{\mathbb{P}_{\omega^*}} \left[\sum_{i=1}^N Q_i(x, y) \right] \\ &= \sum_{i=1}^N \mathbb{E}_{\mathbb{P}_{\zeta^i}} \left[\left(\sum_{j=1}^N \zeta_j^i y_{ij} - T x_i \right)^+ \right]. \end{aligned}$$

Since the marginal distribution of ζ^i has 3^{N_i} realizations, it follows immediately that it suffices to compute overtime costs for $\sum_{i=1}^N 3^{N_i}$ ‘scenarios’. For example, if $N_i = 1$ for all $i = 1, \dots, N$, i.e. if in every OR only a single surgery is carried out, then this number reduces to $3N$, whereas if $N_i = N$, then it equals 3^N .

The second way to speed up computations is based on the following two special cases, which we consider for the i -th OR only. If overtime costs are zero for every scenario, then $Q_i(x, y) = 0$, and if overtime costs are positive for all scenarios, then

$$Q_i(x, y) = \mathbb{E}_{\mathbb{P}_{\zeta^i}} \left[\left(\sum_{j=1}^N \zeta_j^i y_{ij} - T x_i \right)^+ \right] = \sum_{j \in J_i} \left(\mu_j y_{ij} - T x_i \right).$$

The first case hold if $\sum_{j \in J_i} b_j \leq T$ and the second case if $\sum_{j \in J_i} a_j > T$.

The main idea for a general approach, exploiting these special cases, is to iteratively condition on surgery durations ζ_l^i until one of the two special cases applies, i.e. until overtime costs corresponding to all scenarios under consideration are either all zero or all positive. In this way we do not necessarily have to compute all overtime costs for each scenario.

An alternative interpretation of the same idea is to assume that the random vector ζ^i is ordered chronologically, meaning that the first component of ζ^i corresponds to the surgery that is carried out first and the last component of ζ^i corresponds to the surgery that is carried out last. Thus, surgery durations are revealed gradually over time, and we may represent this process by a scenario tree, see e.g. Figure 1. This scenario tree represents all possible realizations of surgery durations at every stage $l = 0, \dots, N_i$, where stage l corresponds to the situation where the first l surgery durations have been observed (i.e. the first l surgeries have been carried out). For example, at the root node in stage 0, no surgery durations are observed yet, whereas in the leave nodes at stage N_i all surgery durations are completely specified.

We iteratively construct the scenario tree, keeping track of the probability $p(n)$ of reaching each node n , and the total surgery duration $D(n)$ of all surgeries carried out before reaching node n . For every such node n at stage l , we compute whether

$$(i) \quad D(n) + \sum_{j=l+1}^{N_i} b_j \leq T,$$

and

$$(ii) \quad D(n) + \sum_{j=l+1}^{N_i} a_j > T.$$

If (i) holds, then the average overtime costs $q(n)$ over all scenarios corresponding to subleaves of node n equal zero. If (ii) holds, then

$$q(n) = D(n) + \sum_{j=l+1}^{N_i} \mu_j,$$

and if (i) and (ii) both do not hold, then we expand the scenario tree, creating three subnodes of n at stage $l + 1$, by conditioning on the three possible realizations of the $(l + 1)$ -th surgery duration. We repeat this process until all nodes are evaluated. In practice much fewer evaluations than 3^{N_i} will be required, as shown in Table 4 in the main part of the paper.

C. Error bounds for the integer experiment of Section 5.2

C.1. General description

In this appendix we discuss how we obtain the error bounds in Tables 5 and 6 of Section 5.2. This error bound is derived by combining the Multiple Replications Procedure (MRP) of Bayraksan and Morton (2009) and the total variation error bounds discussed in Section 3. To our knowledge, this is the first attempt to combine these two approaches. Moreover, it may be interesting to apply this error bound to other applications involving integer decision variables and uncertain random parameters.

A straightforward error bound

Before we derive this error bound, we first discuss why direct application of the error bound of Section 3 is not sufficient to obtain a tight bound.

Let $(\hat{x}, \hat{y})$ denote the optimal surgery-to-OR assignment in (30) and let η^* denote the optimal objective value of the original problem. Then, by Theorem 1 in Section 3 and defining

$$G(\hat{x}, \hat{y}) := \sum_{i=1}^N c_f \hat{x}_i + \sup_{\mathbb{P}_\epsilon \in \mathcal{P}_\epsilon} \mathbb{E}_{\mathbb{P}_\epsilon} Q(\hat{x}, \hat{y}, \epsilon), \quad (42)$$

we have

$$G(\hat{x}, \hat{y}) - \eta^* \leq 2\|Q - \hat{Q}\|_\infty := 2 \sup_{x, y, \epsilon} |Q(x, y, \epsilon) - \hat{Q}(x, y, \epsilon)|.$$

We can obtain an upper bound on $|Q(x, y, \epsilon) - \hat{Q}(x, y, \epsilon)|$ by straightforward application of the total variation error bounds derived in Romeijnders et al. (2016b). However, this bound depends significantly on the surgery-to-OR assignment (x, y) . For example, if every surgery is carried out in a separate OR then the bound reduces to

$$|Q(x, y, \epsilon) - \hat{Q}(x, y, \epsilon)| \leq \frac{1}{2} \sum_{j=1}^N 60c_v h(60|\Delta|f_j), \quad (43)$$

where f_j is the marginal density of the random surgery duration z_j in minutes, and $h(t) = t/8$ if $t \leq 4$ and $h(t) = 1 - 2/t$, otherwise. The value of 60 is present in the error bound since overtime wages are paid in full hours (of 60 minutes). In contrast, if all surgeries are carried out in a single OR, then the error bound reduces to

$$|Q(x, y, \epsilon) - \hat{Q}(x, y, \epsilon)| \leq 30c_v h(60|\Delta|\bar{g}), \quad (44)$$

where $\bar{g}$ is the marginal density of the sum of all surgery durations.

For the numerical experiments in Section 5.2, the bound in (43) is much larger than the bound in (44). Both, however, are by definition larger than the upper bound on $G(\hat{x}, \hat{y}) - \eta^*$. This implies that the error bound $2\|Q - \hat{Q}\|_\infty$ may too large for practical purposes. However, at the same time we do not expect such an extreme surgery-to-OR assignment (x, y) as in (43) to be optimal. This is relevant, since for computing a valid error bound we only require the difference between Q and $\hat{Q}$ in the approximating solution $(\hat{x}, \hat{y})$ and the optimal solution (x^*, y^*) . The problem, however, is that we do not know the optimal solution (x^*, y^*) . Thus, although we do know the approximating solution $(\hat{x}, \hat{y})$, to obtain a valid upper bound we need to take into account the worst-case surgery-to-OR assignment.

C.2. Multiple Replications Procedure

To avoid this problem we will apply the Multiple Replications Procedure (MRP) described in e.g. Bayraksan and Morton (2009). This method cannot be readily applied since it requires us to determine the worst-case probability distribution in the integer model, but combined with the total variation error bound it will yield a much tighter (probabilistic) bound. Below we describe the main outline of the approach.

To assess the quality of the approximating solution we will use (an adjusted version) of the Multiple Replications Procedure (MRP) described in e.g. Bayraksan and Morton (2009). The goal is to evaluate $G(\hat{x}, \hat{y}) - \eta^*$, where G is defined in (42). Since Q is not convex in ϵ we cannot use the results of this paper to determine the worst-case distribution of ϵ . That is why we approximate Q by $\hat{Q}$ and obtain

$$\begin{aligned} G(\hat{x}, \hat{y}) &= \sum_{i=1}^N c_f \hat{x}_i + \sup_{\mathbb{P}_\epsilon \in \mathcal{P}_\epsilon} \left\{ \mathbb{E}_{\mathbb{P}_\epsilon} \left[\hat{Q}(\hat{x}, \hat{y}, \epsilon) + \left(Q(\hat{x}, \hat{y}, \epsilon) - \hat{Q}(\hat{x}, \hat{y}, \epsilon) \right) \right] \right\} \\ &\leq \sum_{i=1}^N c_f \hat{x}_i + \sup_{\mathbb{P}_\epsilon \in \mathcal{P}_\epsilon} \mathbb{E}_{\mathbb{P}_\epsilon} \hat{Q}(\hat{x}, \hat{y}, \epsilon) + \sup_{\mathbb{P}_\epsilon \in \mathcal{P}_\epsilon} \mathbb{E}_{\mathbb{P}_\epsilon} \left[Q(\hat{x}, \hat{y}, \epsilon) - \hat{Q}(\hat{x}, \hat{y}, \epsilon) \right] \\ &= \hat{G}(\hat{x}, \hat{y}) + \sup_{\mathbb{P}_\epsilon \in \mathcal{P}_\epsilon} \mathbb{E}_{\mathbb{P}_\epsilon} \left[Q(\hat{x}, \hat{y}, \epsilon) - \hat{Q}(\hat{x}, \hat{y}, \epsilon) \right], \end{aligned}$$

where $\hat{G}$ equals G with Q replaced by $\hat{Q}$. Since $\hat{Q}$ is convex in ϵ it follows that $\mathbb{P}_\epsilon$ is the worst-case distribution of ϵ in $\hat{G}$. This implies that $\hat{G}(\hat{x}, \hat{y})$ does not contain any optimization problem and that it can easily be estimated using (Monte Carlo) sampling.

To eliminate the supremization of $\mathbb{P}_\epsilon$ over $\mathcal{P}_{(\mu, d)}$ in η^* we assume that $\mathbb{P}_\epsilon$ is the worst-case distribution, so that we obtain a lower bound $\tilde{\eta}$ for η^* , see Theorem 1 in Section 3. To obtain $\tilde{\eta}$ we still have to minimize over all feasible surgery-to-OR assignments $(x, y) \in X$. However, the MRP is able to deal with such problems.

Combining both results we obtain an upper bound on the optimality gap:

$$G(\hat{x}, \hat{y}) - \eta^* \leq \hat{G}(\hat{x}, \hat{y}) - \tilde{\eta} + \sup_{\mathbb{P}_\epsilon \in \mathcal{P}_\epsilon} \mathbb{E}_{\mathbb{P}_\epsilon} \left[Q(\hat{x}, \hat{y}, \epsilon) - \hat{Q}(\hat{x}, \hat{y}, \epsilon) \right].$$

We will use the MRP to bound $\hat{G}(\hat{x}, \hat{y}) - \tilde{\eta}$ and we use a total variation error bound for the last term. This bound may be much tighter than the previous total variation bound, since we only have to compute the difference between Q and $\hat{Q}$ for a *fixed* surgery-to-OR assignment $(\hat{x}, \hat{y})$.

C.3. Total variation bounds

In this section we consider the total variation error bounds mentioned in the previous sections. Suppose that a feasible surgery-to-OR assignment (x, y) is given and assume for the moment that ϵ is fixed. We consider

$$Q(x, y, \epsilon) - \hat{Q}(x, y, \epsilon) = c_v \sum_{i=1}^N \psi_i(x, y, \epsilon),$$

where

$$\psi_i(x, y, \epsilon) = \mathbb{E}_{\mathbb{P}_\epsilon} \left[60 \left[\left(\sum_{j=1}^N z_j y_{ij} - (T + \epsilon_i) x_i \right) / 60 \right]^+ - \left(\sum_{j=1}^N z_j y_{ij} - (T + \epsilon_i - 30) x_i \right)^+ \right].$$

Obviously, if $x_i = 0$, then $y_{ij} = 0$ for all $j = 1, \dots, N$, so that $\psi_i(x, y, \epsilon) = 0$. If not, then let J_i denote the set of surgeries that are carried out in OR i , i.e. $j \in J_i$ if and only if $y_{ij} = 1$. Then, define ξ_i as the total surgery duration in the i -th OR in minutes:

$$\xi_i = \sum_{j=1}^N z_j y_{ij} = \sum_{j \in J_i} \omega_j.$$

Using this definition, and defining $s_i = (T + \epsilon_i)x_i = T + \epsilon_i$, the expression for ψ_i reduces to

$$\psi_i(x, y, \epsilon) = \mathbb{E}_{\mathbb{P}_{\xi_i}} \left[60 \left\lceil \frac{\xi_i - s_i}{60} \right\rceil^+ - (\xi_i - s_i + 30)^+ \right].$$

Applying the total variation error bound yields

$$\psi_i(x, y, \epsilon) \leq 30h(60|\Delta|g_i), \quad (45)$$

where $|\Delta|g_i$ is the total variation of the probability density function g_i of the total surgery duration ξ_i in the i -th OR. Thus,

$$Q(x, y, \epsilon) - \hat{Q}(x, y, \epsilon) \leq \frac{1}{2}c_v \sum_{i=1}^N 60h(|\Delta|g_i)x_i. \quad (46)$$

Here, we add x_i in the expression to ensure that $\psi_i(x, y, \epsilon) = 0$ if $x_i = 0$. Observe that if each surgery is carried out in a separate OR, that in this case $g_i = f_i$ for every $i = 1, \dots, N$ and we obtain the bound given in (43). On the other hand, if all surgeries are carried out in a single OR, then we obtain the bound in (44). Since the bound in (46) holds independent of the value of ϵ , we conclude that

$$\begin{aligned} \sup_{\mathbb{P}_\epsilon \in \mathcal{P}_\epsilon} \mathbb{E}_{\mathbb{P}_\epsilon} [Q(\hat{x}, \hat{y}, \epsilon) - \hat{Q}(\hat{x}, \hat{y}, \epsilon)] &\leq \sup_{\mathbb{P}_\epsilon \in \mathcal{P}_\epsilon} \mathbb{E}_{\mathbb{P}_\epsilon} \left[\frac{1}{2}c_v \sum_{i=1}^N 60h(|\Delta|g_i)\hat{x}_i \right] \\ &= \frac{1}{2}c_v \sum_{i=1}^N 60h(|\Delta|g_i)\hat{x}_i. \end{aligned} \quad (47)$$

This error bound may be much smaller than the one described in (43). For one, since $\hat{x}_i$ may be zero for many ORs. In addition, g_i is the pdf of the sum of several independent random variables, and increasing the number of surgeries in the i -th OR will decrease the total variation of g_i .

Tighter total variation error bounds The error bounds in Tables 5 and 6 are still even tighter than the one in (47). Surprisingly, we derive these bounds by applying the result of Ben-Tal and Hochman (1972) once more.

First, consider again

$$\psi_i(x, y, \epsilon) = \mathbb{E}_{\xi_i} \left[60 \left\lceil \frac{\xi_i - s_i}{60} \right\rceil^+ - (\xi_i - s_i + 30)^+ \right],$$

where $s_i = T + \epsilon_i$ and observe that its underlying value function equals zero if $\xi_i \leq s_i - 30$. For this reason, omitting the technical details, we show that for fixed $s_i = T + \epsilon_i \in \mathbb{R}$,

$$\begin{aligned} \psi_i(x, y, \epsilon) &\leq 30h \left(60|\Delta|g_i \left([s_i - 30, +\infty) \right) \right) \\ &= 30h \left(60|\Delta|g_i \left([T + \epsilon_i - 30, +\infty) \right) \right), \end{aligned}$$

where $g_i([T + \epsilon_i - 30, +\infty))$ denotes the total variation of g_i on the interval $[T + \epsilon_i - 30, +\infty)$. This bound is tighter than the one in (45) and is attained if $\epsilon_i \rightarrow -\infty$. Since the bound is non-increasing in ϵ_i , we may conclude that for every $\epsilon_i \in [a_i, b_i]$,

$$\psi(x, y, \epsilon) \leq 30h \left(60|\Delta|g_i \left([T + a_i - 30, +\infty) \right) \right),$$

and thus

$$\sup_{\mathbb{P}_\epsilon \in \mathcal{P}_\epsilon} \mathbb{E}_{\mathbb{P}_\epsilon} [Q(\hat{x}, \hat{y}, \epsilon) - \hat{Q}(\hat{x}, \hat{y}, \epsilon)] \leq \frac{1}{2} \sum_{i=1}^N 60c_v h \left(60|\Delta|g_i \left([T + a_i - 30, +\infty) \right) \right) \hat{x}_i.$$

Figure 3 Piecewise linear decision rules. Having the defined the decisions for a_t , μ_t , and b_t , the decision for points outside $\{a_t, \mu_t, b_t\}$ are convex combinations of the decisions for a_t , μ_t or μ_t , b_t .

However, we may obtain a tighter bound if the bound on $\psi(x, y, \epsilon)$ is convex in ϵ_i for $\epsilon_i \in [a_i, b_i]$ since it allows us to apply the result of Ben-Tal and Hochman (1972) in a surprising way. For example, if for all opened ORs i , the bounds are convex for $\epsilon_i \in [a_i, b_i]$, then

$$\sup_{\mathbb{P}_\epsilon \in \mathcal{P}_\epsilon} \mathbb{E}_{\mathbb{P}_\epsilon} [Q(\hat{x}, \hat{y}, \epsilon) - \hat{Q}(\hat{x}, \hat{y}, \epsilon)] \leq \frac{1}{2} \mathbb{E}_{\mathbb{P}_\epsilon} \left[\sum_{i=1}^N 60c_v h \left(60|\Delta|g_i \left([T + \bar{\epsilon}_i - 30, +\infty) \right) \right) \hat{x}_i \right]. \quad (48)$$

Of course, the bound $h(60|\Delta|g_i([T + \epsilon_i - 30, +\infty))$ is in general not convex, but it may be in special cases. Notice, for example, that h is linear on $[0, 4]$ so that the bound is convex if $|\Delta|g_i([t + \epsilon_i - 30, +\infty))$ is convex in $\epsilon_i \in [a_i, b_i]$ and this total variation is small enough. In our numerical experiments, g_i is the pdf of the sum of several independent lognormal random variables, so that by the Central Limit Theorem, it is approximately normally distributed. Since a normal density function has a convex decreasing right tail it may satisfy the requirements. In our numerical experiments we check numerically for every opened OR i whether convexity holds; if not then we replace ϵ_i^* by a_i in the error bound of (48).

D. Appendix - inventory experiment

D.1. Decisions for uncertainty realizations outside the finite worst-case support

In this Appendix, we provide a detailed procedure to obtain a sequence of feasible decisions for arbitrary outcome of uncertainty $\hat{z}$, based on the solution to worst-case expectation version of problem (7) with 3^T points in the support. The idea is to use the solutions $\bar{x}_1, \bar{x}_2, \dots$ to (7) and the convexity of the feasible set of (7) to construct a feasible sequence of decisions for all stages. In this setting, the later-stage decisions become piecewise-affine functions of the observed uncertainties.

The way we accomplish this is the following. Each realization of $\hat{z}$ is a convex combination of some of the elements z_t of the discrete worst-case distribution support. For a given component $\hat{z}_t \in \mathbb{R}$ we define that (i) if $\hat{z}_t \in [a_t, \mu_t]$ then $\hat{z}_t$ is formulated as a convex combination of a_t and μ_t , (ii) if $\hat{z}_t \in [\mu_t, b_t]$ then $\hat{z}_t$ is formulated as a convex combination of μ_t and b_t . Then, for each realization $\hat{z}_t$ we obtain a unique set of coefficients of the convex combination. A way to implement feasible decisions corresponding to an arbitrary realization $\hat{z}_t$ is to use the same convex combination coefficients with respect to decision vectors corresponding to the uncertainty realizations in the support of the worst-case distribution of z , which is illustrated in Figure 3.

The following example illustrates this idea mathematically.

EXAMPLE 2. Imagine a problem instance where $T = 2$ and $z \in \mathbb{R}$ belongs to a (μ, d) ambiguity set with $a = 0$, $\mu = 1$, and $b = 2$ and the realized uncertainty is $\hat{z} = 0.5$. Then we can see that $\hat{z} \in [a_1, \mu_1]$. So, we have that $\hat{z}$ is a convex combination of a_1 , and μ_1 . Indeed, take

$$\hat{z} = \lambda a_1 + (1 - \lambda) \mu_1.$$

Since each of the points a, μ has its own stage 2 decision, let us denote them as $x_2(a)$ and $x_2(\mu)$:

$$x_2(\hat{z}) = \lambda x_2(a) + (1 - \lambda) x_2(\mu).$$

Note that such a policy is feasible for $\hat{z}$ because:

$$\begin{aligned} A_{2,1}(a)x_1 + A_{2,2}x_2(a) &= b_2(a), \quad A_{2,1}(\mu)x_1 + A_{2,2}x_2(\mu) = b_2(\mu) \quad \Rightarrow \\ \Rightarrow \quad A_{2,1}(\lambda a_1 + (1 - \lambda)\mu_1)x_1 + A_{2,2}x_2(\lambda a_1 + (1 - \lambda)\mu_1) &= b_2(\lambda a_1 + (1 - \lambda)\mu_1) \end{aligned}$$

by linearity of $A_{2,1}(\cdot)$ and $b_2(\cdot)$. □

More formally, the decisions are implemented as:

$$x(\hat{z}) = \sum_{\kappa \in \{1,2\}^{nz}} \lambda(\kappa) x(z(\kappa)),$$

where $x(z(\kappa))$ is the sequence of decisions related to the demand trajectory $z(\kappa)$ in the worst-case support, defined as:

$$z(\kappa) = \begin{bmatrix} z_1(\kappa_1) \\ \vdots \\ z_{T-1}(\kappa_{T-1}) \end{bmatrix}, \quad \kappa = \begin{bmatrix} \kappa_1 \\ \vdots \\ \kappa_{T-1} \end{bmatrix} \quad z(\kappa_t) = \begin{bmatrix} z_{t1}(\kappa_{t1}) \\ \vdots \\ z_{tn_{z,t}}(\kappa_{t,n_{z,t}}) \end{bmatrix}, \quad \kappa_t = \begin{bmatrix} \kappa_{t1} \\ \vdots \\ \kappa_{tn_{z,t}} \end{bmatrix}$$

and

$$\begin{aligned} z_{tj}(1) &= \begin{cases} a_{tj} & \text{if } \hat{z}_{tj} \leq \mu_{tj} \\ \mu_{tj} & \text{otherwise,} \end{cases} \\ z_{tj}(2) &= \begin{cases} \mu_{tj} & \text{if } \hat{z}_{tj} \leq \mu_{tj} \\ b_{tj} & \text{otherwise.} \end{cases} \end{aligned}$$

and $\lambda(\kappa)$ are unique coefficients such that:

$$\hat{z} = \sum_{\kappa \in \{1,2\}^{nz}} \lambda(\kappa) z(\kappa),$$

where

$$\lambda(\kappa) = \prod_{t=1}^{T-1} \lambda_t(\kappa_t), \quad \lambda_t(\kappa_t) = \prod_{j=1}^{n_{z,t}} \lambda_{tj}(\kappa_{tj}), \quad \lambda_{tj}(1) = \begin{cases} \frac{\hat{z}_{tj} - a_{tj}}{\mu_{tj} - a_{tj}} & \text{if } \hat{z}_{tj} \leq \mu_{tj} \\ \frac{\hat{z}_{tj} - \mu_{tj}}{b_{tj} - \mu_{tj}} & \text{otherwise,} \end{cases}$$

and $\lambda_{tj}(2) = 1 - \lambda_{tj}(1)$. In this way the resulting decision always satisfies the problem's constraints for a given realization of the uncertain parameter.

D.2. Reoptimization to compute the intervals of Section 5.3.2

We explain the computation of the ends of the intervals using reoptimization on the example of computing the lower bound for the PL-WCE solutions. The question corresponding to computation of the lower end of the interval is: what is the expected total cost if the true demand distribution is the best-case distribution, but the decision maker assumes all the time that it is the worst-case distribution?

To answer it, for each of the 2^T best case demand trajectories we compute the ordering decisions in a reoptimizing fashion. That is, at time 1 decisions are determined such that the worst-case expectation is minimized. The corresponding time 1 decision is implemented. At time 2 it turns out that the demand in time 1 was one of the demands belonging to the best-case support $\{\mu_1 - d_1/2(1 - \beta_1), \mu_1 + d_1/2\beta_1\}$ of z_1 . In this situation, decisions for stages 2-6 are constructed (thus, a new optimization problem is solved) that minimize again the worst-case expectation and the corresponding time 2 decision is implemented. At time 3 it turns out that the demand at time 2 belonged to the support $\{\mu_2 - d_2/2(1 - \beta_2), \mu_2 + d_2/2\beta_2\}$ of the best-case distribution and so on. In the end, for each of the 2^T demand trajectories in the support of the best-case distribution, a corresponding decision trajectory is obtained and the objective function values for each of the best-case trajectories are weighted with the corresponding best-case probability.

D.3. Simulating the (μ, d) sample

The (μ, d) sample in the inventory experiment is constructed as follows. First, a discretized distribution $\hat{\mathbb{P}} \in \mathcal{P}_z$ is sampled using the hit-and-run method (Smith 1984). The hit-and-run method is implemented as follows. For the $[0, 1]$ interval (from which the demands on the relevant support intervals are sampled using the inverse transform) we construct a grid of 51 equidistant points. For a fixed (μ, d) the set of probability masses for the points of the grid satisfying the μ and d values constitutes a polytope. We sample 10 probability distributions uniformly from this polyhedron by iteratively choosing a random direction and sampling uniformly a point on the segment of the line along this direction belonging to the polytope. Then, we sample the demand in each period in two steps, by sampling first one of the distributions, and then by sampling a point in the $[0, 1]$ interval with the given distribution.