



Designing tractable piecewise affine policies for multi-stage adjustable robust optimization

Simon Thomä¹ · Grit Walther¹ · Maximilian Schiffer²

Received: 11 January 2023 / Accepted: 19 December 2023
© The Author(s) 2024

Abstract

We study piecewise affine policies for multi-stage adjustable robust optimization (ARO) problems with non-negative right-hand side uncertainty. First, we construct new dominating uncertainty sets and show how a multi-stage ARO problem can be solved efficiently with a linear program when uncertainty is replaced by these new sets. We then demonstrate how solutions for this alternative problem can be transformed into solutions for the original problem. By carefully choosing the dominating sets, we prove strong approximation bounds for our policies and extend many previously best-known bounds for the two-staged problem variant to its multi-stage counterpart. Moreover, the new bounds are—to the best of our knowledge—the first bounds shown for the general multi-stage ARO problem considered. We extensively compare our policies to other policies from the literature and prove relative performance guarantees. In two numerical experiments, we identify beneficial and disadvantageous properties for different policies and present effective adjustments to tackle the most critical disadvantages of our policies. Overall, the experiments show that our piecewise affine policies can be computed by orders of magnitude faster than affine policies, while often yielding comparable or even better results.

Keywords Robust optimization · Multi-stage problem · Adjustable optimization · Approximation bounds

Mathematics Subject Classification 90C47 · 90C57 · 90C59 · 49K35

✉ Simon Thomä
simon.thomae@om.rwth-aachen.de

Grit Walther
walther@om.rwth-aachen.de

Maximilian Schiffer
schiffer@tum.de

¹ Chair of Operations Management, RWTH Aachen University, Aachen, Germany

² School of Management and Munich Data Science Institute, Technical University of Munich, Munich, Germany

1 Introduction

In practice, most decision-making problems have to be solved in view of uncertain parameters. In the operations research domain, two fundamental frameworks exist to inform such decisions. The *stochastic optimization* framework captures uncertainty by using probability distributions and aims to optimize an expected objective. Initially introduced in the seminal work of Dantzig [26], the framework has been intensively studied and scholars used it to solve a vast variety of problems including production planning [6, 55], relief networks [49], expansion planning [62, 64], and newsvendor problems [52]. While stochastic optimization performs well on many problem classes, finding tractable formulations is oftentimes challenging. Additionally, the data needed to approximate probability distributions might not always be available, and gathering data is often a costly and time-consuming process. The *robust optimization* (RO) framework overcomes many of these shortcomings by capturing uncertainty through distribution-free uncertainty sets instead of probability distributions. By choosing well-representable uncertainty sets, RO often offers computationally tractable formulations that scale well on a variety of optimization problems. In recent years, these favorable properties have led to a steep increase in research interest, see, e.g. [8, 11, 20, 36, 63]. The flexibility of uncertainty sets and scalability of solution methods also make RO very attractive for applicational purposes and it has been widely applied to many operations management problems [45].

In traditional RO, all decisions must be made before the uncertainty realization is revealed. However, oftentimes some decisions can be delayed until after (part of) the uncertainty realization is known in real-world situations. As a consequence, RO may lead to excessively conservative solutions. To remedy this drawback, Ben-Tal et al. [9] introduced the concept of *adjustable robust optimization* (ARO) where some decisions can be delayed until the uncertainty realization is (partly) known. In general ARO, uncertainty realizations are revealed over multiple stages and decisions can be made after each reveal. A decision made in stage t can thus be modeled as a function of all uncertainties associated with previous stages $t' \leq t$.

While ARO improves decision-making in theory, ARO is equivalent to RO on some special problem instances, where static decision policies yield optimal adjustable solutions [9, 22]. Using similar arguments to the ones used for the optimality of static solutions, Marandi and den Hertog [46] identified conditions where optimal adjustable decisions are independent of some uncertain parameters. In general, static policies do not yield optimal solutions in the ARO setting. Elucidating this, Haddad-Sisakht and Ryan [37] identified a collection of sufficient conditions that imply the suboptimality of static policies and a strict improvement of ARO over RO.

In general, even the task of finding optimal adjustable solutions in the special case of two-stage ARO proves to be computationally intractable [9]. Accordingly, recent works developed many approximation schemes for ARO that often yield good and sometimes even optimal results in practice, see, e.g. [9, 19, 42, 58]. In the special case with only two decision stages, the first stage decisions are fixed before any uncertain parameters are known and the second stage decisions use full knowledge of the uncertainty realization. Two-stage ARO already has many applications in practice and has widely been studied in the literature, see, e.g. [14, 16, 18, 22, 23, 39, 40, 59].

Still, many real-world problems show inherent multi-stage characteristics and cannot be modeled by two-stage ARO. Examples include variants of inventory management [12], humanitarian relief [13], and facility location [5]. The transition from two-stage to multi-stage ARO introduces two main challenges. First, multi-stage ARO problems entail *nonanticipativity restrictions* that disallow decisions to utilize future information. Second, many approaches that solve the problem by iteratively splitting the uncertainty space, like scenario trees [41], and adaptive partitioning [16, 53], grow exponentially in the number of stages. As a consequence, many results found for two-stage ARO do not readily generalize to multi-stage scenarios. Against this background, we design piecewise affine policies for multi-stage ARO that overcome the previously mentioned challenges and extend, although it is not straightforward, many of the best know approximation bounds for two-stage problems to a multi-stage setting.

In the remainder of this section, we formally introduce our problem (Sect. 1.1), discuss closely related work (Sect. 1.2), and summarize our main contributions (Sect. 1.3).

1.1 Problem description

In this work we study multi-stage adjustable robust optimization with covering constraints and a positive affine uncertain right hand side. Specifically, we consider the following problem:

$$\begin{aligned}
 Z_{AR}(\mathcal{U}) = \min_{\mathbf{x}(\xi)} \max_{\xi \in \mathcal{U}} & \quad \mathbf{c}^\top \mathbf{x}(\xi) \\
 \text{s.t.} & \quad \mathbf{A}\mathbf{x}(\xi) \geq \mathbf{D}\xi + \mathbf{d} \quad \forall \xi \in \mathcal{U}
 \end{aligned} \tag{1}$$

with $\mathbf{A} \in \mathbb{R}^{l \times n}$, $\mathbf{c} \in \mathbb{R}^n$, $\mathbf{D} \in \mathbb{R}_+^{l \times m}$, $\mathbf{d} \in \mathbb{R}_+^l$, and compact $\mathcal{U} \subset \mathbb{R}_+^m$. Here, m is the number of uncertain parameters, n is the number of decisions, and l is the number of constraints. To model the problem's T stages, we split the uncertainty vector ξ into T sub-vectors $\xi = (\xi^1, \dots, \xi^T)$ with ξ^t being the uncertainty vector realized in stage t . In the following, we denote by $\underline{\xi}^t := (\xi^1, \dots, \xi^t)$ the vector of all uncertainties with known realization in stage t . Similarly, the adjustable decision vector $\mathbf{x}(\xi)$ divides into $\mathbf{x}(\xi) := (\mathbf{x}^1(\underline{\xi}^1), \dots, \mathbf{x}^T(\underline{\xi}^T))$, where the decision \mathbf{x}^t made in stage t has to preserve *nonanticipativity* and may only depend on those uncertainties $\underline{\xi}^t$ whose realization is known in stage t . We explicitly allow ξ^1 to be zero-dimensional making the initial decision \mathbf{x}^1 non-adjustable. Finally, we denote by $\underline{\mathbf{x}}^t(\xi) := (\mathbf{x}^1(\underline{\xi}^1), \dots, \mathbf{x}^t(\underline{\xi}^t))$ the vector of all decisions in the first t stages. Figure 1 visualizes the multi-stage decision process with nonanticipativity restrictions.

Unless explicitly stated otherwise, we assume w.l.o.g. that the following assumption holds throughout the paper.

Assumption 1 $\mathcal{U} \subseteq [0, 1]^m$ is convex, full-dimensional with $e_i \in \mathcal{U}$ for all $i \in \{1, \dots, m\}$, and *down-monotone*, i.e., $\forall \xi \in \mathcal{U}, \mathbf{0} \leq \xi' \leq \xi : \xi' \in \mathcal{U}$.

Down-monotonicity holds because $\mathbf{D}, \mathbf{d}, \xi$ are all non-negative, and thus constraints become less restrictive for smaller values of ξ . Convexity holds due to the linearity




Fig. 1 Illustration of multi-stage decision making over T stages. In each stage t a fraction ξ^t of the uncertainty is realized and decisions x^t are made. Here, decisions x^t may only depend on those uncertainties ξ^t whose realization is known in stage t

of the problem, and $e_i \in \mathcal{U} \subseteq [0, 1]^m$ holds as \mathcal{U} is compact and \mathbf{D} can be re-scaled appropriately. We note that the non-negativity assumption of the right-hand side does restrict the problem space. As Bertsimas and Goyal [18] point out, this assumption prevents the introduction of uncertain or constant upper bounds. However, upper bounds in other decision variables are still possible as \mathbf{A} is not restricted, and \mathbf{D} , \mathbf{d} can be zero. Overall, Problem (1) covers a wide range of different problem classes including network design [50, 61], capacity planning [44, 48, 51], as well as versions of inventory management [12, 60] where capacities are unbounded or subject to the decision makers choice.

In the context of multi-stage decision making, some works require *stagewise* uncertainty, i.e., that the uncertainty set \mathcal{U} consists of uncertainty sets $\mathcal{U}_1, \dots, \mathcal{U}_T$ for each stage, see, e.g., [19, 20, 32]. Like other approaches based on decision rules [17, 43], we do not need these restricting assumptions. However, we show how to utilize the existence of such a structure in Sect. 3.

1.2 Related work

Feige et al. [30] show that already the two-stage version of Problem (1) with $\mathbf{D} = \mathbf{1}$, $\mathbf{d} = \mathbf{0}$ and \mathbf{A} being a 0-1-matrix is hard to approximate with a factor better than $\Omega(\log m)$, even for budgeted uncertainty sets. As it is thereby impossible to find general solutions for \mathbf{x} , a common technique to get tractable formulations is to restrict the function space.

In this context, Ben-Tal et al. [9] consider \mathbf{x} to be affine in ξ . Specifically, they propose x^t to be of the form $x^t(\xi^t) = \mathbf{P}^t \xi^t + \mathbf{q}^t$. Affine policies have been found to deliver good results in practice [1, 10] and are even optimal for some special problems [19, 42, 58]. Further popular decision rules include segregate affine [24, 25], piecewise constant [20], piecewise affine [14, 31], and polynomial [4, 21] policies, as well as combinations of these [54]. For surveys on adjustable policies we refer to Delage and Iancu [28] and Yanıkoğlu et al. [63].

A key question that arises when using policies to solve ARO problems is how good the solutions are compared to an optimal unrestricted solution. To answer this, many approximation schemes for a priori and a posteriori bounds have been proposed. In the context of a posteriori bounds the focus lies on finding tight upper and lower approximation problems. Hadjiyiannis et al. [38] estimate the suboptimality of affine decision

rules using sample scenarios from the uncertainty set. Similar sample lower bounds are used by Bertsimas and Georghiou [17] to bound the performance of piecewise affine policies. Kuhn et al. [43] investigate the optimality of affine policies by using the gap between affine solutions on the primal and the dual of the problem. Georghiou et al. [31] generalize this primal-dual approach to affine policies on lifted uncertainty sets. Building on both of the previous approaches, Georghiou et al. [33] propose a convergent hierarchy of policies that combine affine policies with extreme point scenario samples. Daryalal et al. [27] construct lower bounds by relaxing nonanticipativity and stage-connecting constraints in multi-stage ARO. They then use these lower bounds to construct primal solutions in a rolling horizon manner.

In the context of a priori bounds, most approximation schemes have been proposed for the two-staged version of Problem (1). For general uncertainty sets on the two-stage version of (1), Bertsimas and Goyal [18] show that affine policies yield an $O(\sqrt{m})$ approximation if c and x are non-negative. They further construct a set of instances where this bound is tight, showing that no better general bounds for affine policies exist. Using geometric properties of the uncertainty sets, Bertsimas and Bidkhori [15] improve on these bounds for some commonly used sets including budgeted uncertainty, norm balls, and intersections of norm balls. Ben-Tal et al. [14] propose new piecewise affine decision rules for the two-stage problem that on some sets improve these bounds even further. In addition to strong theoretical bounds, this new approach also yields promising numerical results that can be found by orders of magnitude faster than solutions for affine adjustable policies. For budgeted uncertainty sets and some generalizations thereof, Housni and Goyal [40] show that affine policies even yield optimal approximations with an *asymptotic bound*, i.e., asymptotic behavior of the approximation bound, see, e.g. [56], of $O\left(\frac{\log m}{\log \log m}\right)$. This bound was shown to be tight by Feige et al. [30] for reasonable complexity assumptions, namely 3SAT cannot be solved in $2^{O(\sqrt{m})}$ time on instances of size m . We present an overview of known a priori approximation bounds for some commonly used uncertainty sets on two-stage ARO in Table 4 of “Appendix A”.

To the best of our knowledge, Bertsimas et al. [20] are the only ones that provide a priori bounds for multi-stage ARO so far. They show these bounds for piecewise constant policies using geometric properties of the uncertainty sets. More specifically, they consider multi-stage uncertainty networks, where the uncertainty realization is taken from one of multiple independent uncertainty sets in each stage. While the choice of the set selected in each stage may depend on the sets selected before, the uncertainty sets are otherwise independent. Although this assumption is fairly general, it still leaves many commonly used sets where uncertainty is dependent over multiple stages uncovered. Among others, uncovered sets include widely used hypersphere and budgeted uncertainty.

As can be seen, previous work has predominantly focused on providing tighter approximation bounds for two-stage ARO, inevitably raising the question of whether similar bounds hold for multi-stage ARO as well. This work contributes towards answering this question by extending many of the currently best-known a priori bounds on two-stage ARO to its multi-stage setting. By so doing, we are—to the best of our knowledge—the first ones to provide a priori approximation bounds for the multi-

stage ARO Problem (1), where uncertainty sets can range over multiple stages. Unless explicitly stated otherwise, we will always refer to a priori approximation bounds when we discuss approximation bounds in the remainder of this paper.

1.3 Our contributions

With this work, we extend the existing literature in multiple ways, where our main contributions are as follows.

Tractable piecewise affine policies for multi-stage ARO: Motivated by piecewise affine policies for two-stage ARO [14], we present a framework to construct policies that can be used to efficiently find good solutions for the multi-stage ARO Problem (1). Instead of solving the problem directly for uncertainty \mathcal{U} , we first approximate \mathcal{U} by a dominating set $\hat{\mathcal{U}}$. To do so, we define the concept of nonanticipative multi-stage domination and show that this new definition of domination fulfills similar properties to two-stage domination. Based on this new definition, we then construct dominating sets $\hat{\mathcal{U}}$ such that solutions on $\hat{\mathcal{U}}$ can be found efficiently. More specifically, we choose $\hat{\mathcal{U}}$ to be a polytope for which worst-case solutions can be computed by a linear program (LP) over its vertices. In order to ensure nonanticipativity, which is the main challenge of this construction, we introduce a new set of constraints on the vertices that guarantee the existence of nonanticipative extensions from the vertex solutions to the full set $\hat{\mathcal{U}}$. Finally, we show how to use the solution on the dominating set $\hat{\mathcal{U}}$ to construct a valid solution for the original uncertainty set \mathcal{U} .

Approximation bounds: To the best of our knowledge, we provide the first approximation bounds for the multi-stage Problem (1) with general uncertainty sets. More specifically, we show that our policies yield $O(\sqrt{m})$ approximations of fully adjustable policies. While this bound is tight for our type of policies in general, we show that better bounds hold for many commonly used uncertainty sets.

While our main contribution is to extend approximation bounds to multi-stage ARO, Problem (1) is further less restrictive than problems previously discussed in the literature on approximation bounds. In addition to being restricted to two-stage ARO, previous work often assumed c and x to be non-negative [14, 15, 18]. Ben-Tal et al. [14] additionally restricted parts of A , i.e., they require the parts of A associated with the first stage decision to be non-negative. Our policies do not need this assumption. However, we show that Problem (1) is unbounded whenever there is a feasible x with $c^\top x < 0$, due to the non-negativity of the right-hand side. As a consequence, our policies do not readily extend to general maximization problems.

From a theoretical perspective, mainly asymptotic bounds are of interest. In practice, however, also the exact factors of the approximation are important. Throughout the paper, we thus always give the asymptotic, as well as the exact bounds. We compare all our bounds to the previously best-known bounds for the two-stage setting given in

Ben-Tal et al. [14] and show that our constructions yield both constant factor, as well as asymptotic improvements.

Comparison with affine policies: Using the newly found bounds, we show that no approximation bound for affine policies on hypersphere uncertainty exists that is better than the bound we show for our policies. For budgeted uncertainty, on the other hand, we show that affine policies strictly dominate our piecewise affine policies. These findings confirm results that have been reported for the two-stage variant, where affine policies do not perform well for hypersphere uncertainty [18], but very well for budgeted uncertainty [40].

Improvement heuristic: Due to inherent properties of our policy construction, resulting solutions are overly pessimistic on instances where the impact on the objective varies significantly between different uncertainty dimensions. To diminish this effect, we introduce an improvement heuristic that performs at least as well as our policies and that can be integrated into the LP used to construct our policies. While these modifications come at the cost of higher solution times, they allow for significant objective improvements on some instance classes.

Tightening piecewise affine policies via lifting: We show that in the context of Problem (1) the piecewise affine policies via lifting presented by Georghiou et al. [31] yield equivalent solutions to affine policies. To prevent this from happening, we construct tightened piecewise affine policies via lifting using insights from our piecewise affine policies. These new policies integrate the approximative power of affine policies, and our piecewise affine policies and are guaranteed to perform at least as well as the individual policies they combine.

Numerical evidence: Finally, we present two sets of numerical experiments showing that our policies solve by orders of magnitude faster than the affine adjustable policies presented by Ben-Tal et al. [9], the piecewise affine policies via lifting presented by Georghiou et al. [31], and the near-optimal piecewise affine policies by Bertsimas and Georghiou [17], while often yielding comparable or improving results. First, we study a slightly modified version of the tests presented in Ben-Tal et al. [14], allowing us to demonstrate our policies' scalability and the impact of our improvement heuristic. Second, we focus on demand covering instances to demonstrate good performances of our policies for a problem that resembles a practical application. We refer to our git repository (<https://github.com/tumBAIS/piecewise-affine-ARO>) for all material necessary to reproduce the numerical results outlined in this paper.

Comparison against closely related work: Compared to the closely related work by Ben-Tal et al. [14], who first introduced the concept of domination in the context of ARO, our contributions are multifold. First, we extend domination-based piecewise affine policies to a wider class of problems by switching from a two-stage to a multi-stage setting and relaxing assumptions. We discuss the structural reasons that make this extension non-trivial at the beginning of Sect. 2. In addition to showing stronger approximation guarantees for our policies, we conduct comprehensive theoretical and numerical comparisons with other adaptable policies. Based on these comparisons, we construct two new policies that mitigate weaknesses of domination and integrate




Fig. 2 Relations between multi-stage ARO policies compared in this paper. An arc from a policy P to another policy P' states $Z_P \leq Z_{P'}$, where Z_P and $Z_{P'}$ are optimal objective values for the ARO Problem (1) solved with policy P, P' respectively. Dashed arcs only hold for hypersphere (H) or budgeted (B) uncertainty. Relations proved for the first time in this paper are highlighted (blue, bold). The compared policies are: static policies (static); piecewise affine policies via domination by Ben-Tal et al. [14] (PAPBT); our piecewise affine policies via domination (PAP), c.f., Sects. 2 and 3; affine policies [9] (AFF); our piecewise affine policies with rescaling (SPAP), c.f., Sect. 4; near-optimal piecewise affine policies [17] (BG); piecewise affine policies via lifting [31] (LIFT); our tightened piecewise affine policies via lifting (TLIFT), c.f., Sect. 5


its strength with the strength of other policies. More specifically, the first policy integrates finding a good outer approximation of the uncertainty set in the optimization process. The second policy integrates structural results from domination into lifting policies, c.f., [31]. As a result, we get a hierarchy of piecewise affine adjustable policies with provable relative performance guarantees. We give an overview of all policies constructed in our work and their relative performance guarantees compared to other policies in Fig. 2.

The rest of this paper is structured as follows. In Sect. 2, we introduce our policies and elaborate on their construction. In Sect. 3, we present our approximation bounds for the multi-stage ARO Problem (1). We present an improvement heuristic for our policies in Sect. 4. By using the results of Sects. 2 and 3, we construct tightened piecewise affine policies via lifting in Sect. 5. Finally, we provide numerical evidence for the performance of our policy compared to other state of the art policies in Sect. 6. Section 7 concludes this paper with a brief reflection of our work and avenues for future research. To keep the paper concise, we defer proofs that could possibly interrupt the reading flow to “Appendices B–O”.

2 Framework for piecewise affine multi-stage policies

In this section, we present our piecewise affine framework for the multi-stage ARO Problem (1). The main rationale of our framework is to construct new uncertainty sets $\hat{\mathcal{U}}$ that dominate the original uncertainty sets \mathcal{U} . With this, our framework follows a similar rationale as the two-stage framework from Ben-Tal et al. [14]. For a problem $Z_{AR}(\mathcal{U})$ we construct $\hat{\mathcal{U}}$ in such a way that $Z_{AR}(\hat{\mathcal{U}})$ can be efficiently solved, and a solution of $Z_{AR}(\hat{\mathcal{U}})$ can be used to generate solutions for $Z_{AR}(\mathcal{U})$.

Fig. 3 Two dimensional hypersphere uncertainty set \mathcal{U} with (dashed) dominating set $\hat{\mathcal{U}}$ (2) induced by the convex combination of vertices v_0, v_1, v_2 and dominating function h (4) that maps a point $\xi \in \mathcal{U}$ to a point $\hat{\xi} \in \hat{\mathcal{U}}$



In this context, we note that one cannot straightforwardly apply the construction scheme used by Ben-Tal et al. [14] due to nonanticipativity requirements. More specifically, Ben-Tal et al. [14] construct $\hat{\mathcal{U}}$ as polytopes, where it is well known that worst-case solutions always occur on extreme points, as any solution can be represented by convex combinations of extreme point solutions. The construction of these convex combinations, however, is not guaranteed to be nonanticipative in the multi-stage setting. To overcome this challenge, we incorporate nonanticipativity in the concept of uncertainty set domination and extend it to a multi-stage setting.

Definition 1 (Domination) Given an uncertainty set $\mathcal{U} \subseteq \mathbb{R}_+^m$, we say that \mathcal{U} is dominated by $\hat{\mathcal{U}} \subseteq \mathbb{R}_+^m$ if there is a domination function $h: \mathcal{U} \rightarrow \hat{\mathcal{U}}$ with $h(\xi) \geq \xi$, and h can be expressed as $h(\xi) = (h^1(\xi^1), \dots, h^T(\xi^T))$ where h^t maps to the uncertainties in stage t and depends on uncertainties up to that stage.

Intuitively, an uncertainty set $\hat{\mathcal{U}}$ dominates another set \mathcal{U} if for every point $\xi \in \mathcal{U}$ there is a point $\hat{\xi} \in \hat{\mathcal{U}}$ that is at least as large in each component, i.e., $\hat{\xi} \geq \xi$. Later we show that the dominating set $\hat{\mathcal{U}}$ can be constructed as the convex combination of $m + 1$ vertices v_0, \dots, v_m . We also show how to construct dominating functions h for these vertex induced dominating sets. Figure 3 illustrates the hypersphere uncertainty set $\mathcal{U} = \{\xi \in \mathbb{R}_+^m \mid \|\xi\|_2 \leq 1\}$ together with our dominating set $\hat{\mathcal{U}}$ (2) and the dominating function h (4) for $m = 2$.

Due to the non-negativity of the problem’s right-hand side, domination at most restricts the set of feasible solutions. As a consequence, each feasible solution for a realization $\hat{\xi} \in \hat{\mathcal{U}}$ is also a feasible solution for all realizations $\xi \in \mathcal{U}$ that are dominated by $\hat{\xi}$. Using this property, we can derive piecewise affine policies for $Z_{AR}(\mathcal{U})$ from solutions of $Z_{AR}(\hat{\mathcal{U}})$. Since \mathcal{U} is full-dimensional and down-monotone by Assumption 1, there always exists a factor $\beta \geq 0$ such that scaling \mathcal{U} by β contains $\hat{\mathcal{U}}$. Theorem 1 shows that with this factor β , solutions of problem $Z_{AR}(\hat{\mathcal{U}})$ are β -approximations for problem $Z_{AR}(\mathcal{U})$. It also shows that $Z_{AR}(\hat{\mathcal{U}})$ is unbounded exactly when $Z_{AR}(\mathcal{U})$ is unbounded. Thus, we assume for the remainder of this paper w.l.o.g. that both $Z_{AR}(\mathcal{U})$ and $Z_{AR}(\hat{\mathcal{U}})$ are bounded.

Theorem 1 Consider an uncertainty set \mathcal{U} from Problem (1) and a dominating set $\hat{\mathcal{U}}$. Let $\beta \geq 1$ be such that $\forall \hat{\xi} \in \hat{\mathcal{U}}: \frac{1}{\beta} \hat{\xi} \in \mathcal{U}$. Moreover, let $Z_{AR}(\mathcal{U})$ and $Z_{AR}(\hat{\mathcal{U}})$ be optimal values of Problem (1). Then, either $Z_{AR}(\mathcal{U})$ and $Z_{AR}(\hat{\mathcal{U}})$ are unbounded or

$$0 \leq Z_{AR}(\mathcal{U}) \leq Z_{AR}(\hat{\mathcal{U}}) \leq \beta \cdot Z_{AR}(\mathcal{U}).$$

We present the proof for Theorem 1 in ‘‘Appendix B’’.

In the remainder of this section, we demonstrate how the results of Theorem 1 can be used to efficiently construct β -approximations for $Z_{AR}(\mathcal{U})$. Therefore, we show in Sect. 2.1 how to construct dominating polytopes $\hat{\mathcal{U}}$ and efficiently find solutions $Z_{AR}(\hat{\mathcal{U}})$ that comply with nonanticipativity requirements. Then, we construct the dominating function $h: \mathcal{U} \rightarrow \hat{\mathcal{U}}$, which allows us to extend these solutions to solutions for $Z_{AR}(\mathcal{U})$ in Sect. 2.2.

2.1 Construction of the dominating set

In the following, we construct a dominating set in the form of a polytope for which the worst-case solution can be efficiently found by solving a linear program on its vertices. Specifically, for an uncertainty set \mathcal{U} , we consider dominating sets $\hat{\mathcal{U}}$ of the form

$$\hat{\mathcal{U}} := \text{conv}(\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_m) \tag{2}$$

where for all $i \in \{0, \dots, m\}: \frac{1}{\beta} \mathbf{v}_i \in \mathcal{U}$ and for all $i \in \{1, \dots, m\}: \mathbf{v}_i = \mathbf{v}_0 + \rho_i \mathbf{e}_i$ for some $\rho_i \in \mathbb{R}_+$. We postpone the construction of the domination function h , the base vertex \mathbf{v}_0 , and parameters ρ_1, \dots, ρ_m to Sect. 2.2 and first focus on the construction of solutions for $Z_{AR}(\hat{\mathcal{U}})$. Here, we extend the notation on \mathbf{x} and ξ introduced in Sect. 1.1 to \mathbf{x}_i and \mathbf{v}_i . Consequently, \mathbf{x}_i^t is the sub-vector of \mathbf{x}_i corresponding to decisions made in stage t , and \mathbf{v}_i^t is the sub-vector of \mathbf{v}_i corresponding to uncertainties up to stage t . Then, the key component for our construction is LP (3)

$$Z_{LP}(\hat{\mathcal{U}}) = \min_{\mathbf{x}_0, \dots, \mathbf{x}_m} z \tag{3a}$$

$$\text{s.t. } z \geq \mathbf{c}^\top \mathbf{x}_i \quad \forall i \in \{0, \dots, m\} \tag{3b}$$

$$\mathbf{A} \mathbf{x}_i \geq \mathbf{D} \mathbf{v}_i + \mathbf{d} \quad \forall i \in \{0, \dots, m\} \tag{3c}$$

$$\mathbf{x}_i^t = \mathbf{x}_j^t \quad \forall i, j \in \{0, \dots, m\}, t \in \{1, \dots, T\}, \mathbf{v}_i^t = \mathbf{v}_j^t. \tag{3d}$$

Intuitively, the Objective (3a) together with Constraints (3b) minimize the maximal cost over all vertex solutions \mathbf{x}_i . Constraints (3c) ensure that each \mathbf{x}_i is a feasible solution for the respective uncertainty vertex \mathbf{v}_i of $\hat{\mathcal{U}}$. Finally, Constraints (3d) ensure nonanticipativity by forcing vertex solutions to be equal unless different uncertainties

were observed. With these constraints, we construct LP (3) such that it is sufficient to find an optimal solution for $Z_{LP}(\hat{\mathcal{U}})$ in order to find an optimal solution for $Z_{AR}(\hat{\mathcal{U}})$.

Lemma 2 *Let $\hat{\mathcal{U}}$ be a dominating set as described in (2), $Z_{LP}(\hat{\mathcal{U}})$ be the solution of LP (3), and $Z_{AR}(\hat{\mathcal{U}})$ be the solution of Problem (1). Then the LP solution (\mathbf{x}_i) on the vertices of $\hat{\mathcal{U}}$ can be extended to a solution on the full set $\hat{\mathcal{U}}$ and we find:*

$$Z_{LP}(\hat{\mathcal{U}}) = Z_{AR}(\hat{\mathcal{U}}).$$

We present the proof for Lemma 2 in ‘‘Appendix C’’.

2.2 Construction of the domination function

In the previous section, we showed how to construct dominating sets $\hat{\mathcal{U}}$ such that $Z_{AR}(\hat{\mathcal{U}})$ can be solved efficiently. In order for $\hat{\mathcal{U}}$ to be a valid dominating set for some uncertainty set \mathcal{U} , we additionally have to construct a nonanticipative dominating function $\mathbf{h}: \mathcal{U} \rightarrow \hat{\mathcal{U}}$ according to Definition 1. Specifically, we use

$$\mathbf{h}(\boldsymbol{\xi}) := (\boldsymbol{\xi} - \mathbf{v}_0)_+ + \mathbf{v}_0 \tag{4}$$

where $(\cdot)_+$ is the element-wise maximum with 0 and \mathbf{v}_0 is the base vertex from Definition (2). By construction, \mathbf{h} maps each uncertainty realization $\boldsymbol{\xi}$ to its element-wise maximum with \mathbf{v}_0 . It directly follows that \mathbf{h} is nonanticipative, as each element in $\mathbf{h}(\boldsymbol{\xi})$ solely depends on the corresponding element in $\boldsymbol{\xi}$.

Finally, we have to ensure that $\mathbf{h}(\boldsymbol{\xi}) \in \hat{\mathcal{U}}$ for all $\boldsymbol{\xi} \in \mathcal{U}$. We do so by choosing the base vertex \mathbf{v}_0 and parameters ρ_1, \dots, ρ_m during the construction of $\hat{\mathcal{U}}$ appropriately. Using $\lambda_i(\boldsymbol{\xi}) := \frac{((\boldsymbol{\xi} - \mathbf{v}_0)_+)_i}{\rho_i}$ with the convention $\frac{0}{0} = 0$, we find

$$\mathbf{h}(\boldsymbol{\xi}) = \sum_{i=1}^m \lambda_i(\boldsymbol{\xi}) \mathbf{v}_i + \left(1 - \sum_{i=1}^m \lambda_i(\boldsymbol{\xi}) \right) \mathbf{v}_0.$$

By definition, any convex combination of \mathbf{v}_i is contained in $\hat{\mathcal{U}}$. Thus, \mathbf{h} is a valid domination function if and only if

$$\max_{\boldsymbol{\xi} \in \mathcal{U}} \sum_{i=1}^m \lambda_i(\boldsymbol{\xi}) \leq 1. \tag{5}$$

Condition (5) gives a compact criterion to check the validity of dominating sets. By doing so, it lays the basis for our optimal selection of the base vertex \mathbf{v}_0 and parameters ρ_0, \dots, ρ_m . Checking Condition (5) generally requires solving a convex optimization problem. However, in Sect. 3 we show that for many commonly used special uncertainty sets, this problem can be significantly simplified, leading to low dimensional unconstrained minimization problems or even analytical solutions.

Recall that we showed how to extend a solution (x_0, \dots, x_m) of LP (3) to the full set $\hat{\mathcal{U}}$ in the proof of Lemma 2. Combining this with \mathbf{h} and using Theorem 1 we get a piecewise affine solution for $Z_{AR}(\mathcal{U})$ by

$$\mathbf{x}(\xi) = \sum_{i=1}^m \lambda_i(\xi) \mathbf{x}_i + \left(1 - \sum_{i=1}^m \lambda_i(\xi)\right) \mathbf{x}_0 \quad (6)$$

that has an optimality bound of β .

2.3 Limitations

While our policies overcome the nontrivial challenge of nonanticipativity on extreme point solutions, they still rely on the ability to form convex combinations. As integrality is not preserved by convex combinations, there is no natural way to extend our approach to integer or binary recourse decisions \mathbf{x} . However, including non-adjustable integer or binary first-stage decisions in our framework is straightforward. Also, it is not straightforwardly possible to incorporate uncertain recourse decisions, i.e., dependence of A on ξ , into our approach, as worst-case realizations for problems with uncertain recourse are not necessarily extreme points of \mathcal{U} , see, e.g., [3, 33].

3 Optimality bounds for different uncertainty sets

In the previous section, we demonstrated how to construct nonanticipative piecewise affine policies for the multi-stage Problem (1). On this basis, proving approximation bounds mainly depends on geometric properties of the uncertainty sets \mathcal{U} . We first show approximation bounds for some commonly used permutation invariant uncertainty sets. On these sets, the dominating sets are permutation invariant and we give closed-form constructions. We then give approximation bounds for our piecewise affine policies on general uncertainty sets. Finally, we demonstrate how the bounds of an uncertainty set \mathcal{U} generalize to transformations of that set. While in theory, mostly asymptotic bounds are of interest, in practice constant factors are important as well. Thus we always state exact, as well as asymptotic bounds. Table 1 gives an overview of all bounds that are explicitly proven in Propositions 5, 7, 9, 10 and 11 of this section. We compare all our results against the results for the two-stage setting in Ben-Tal et al. [14] and show constant factor, as well as asymptotic improvements. For budgeted and hypersphere uncertainty sets, we further compare the theoretical performance of our piecewise affine policies with affine adjustable policies.

For permutation invariant uncertainty sets there exists an optimal choice of $\hat{\mathcal{U}}$ that is also permutation invariant. More specifically, \mathbf{v}_i simplifies to $\mathbf{v}_0 = \mu \mathbf{e}$ and $\mathbf{v}_i = \mathbf{v}_0 + \rho \mathbf{e}_i$, for some μ, ρ .

Lemma 3 *Let \mathcal{U} be a permutation invariant uncertainty set. Then there exist μ, ρ such that for the dominating uncertainty set $\hat{\mathcal{U}}$ spanned by $\mathbf{v}_0 = \mu \mathbf{e}$ and $\mathbf{v}_i = \mathbf{v}_0 + \rho \mathbf{e}_i$ for $i \in \{1, \dots, m\}$ there is no other dominating set $\hat{\mathcal{U}}'$ constructed as in (2) with a smaller approximation factor β .*

Table 1 Performance bounds of the piecewise affine policy for different uncertainty sets

No.	Uncertainty set \mathcal{U}	Bound	Asymptotic bound
I	$\{\xi \in \mathbb{R}_+^m \mid \ \xi\ _2^2 \leq 1\}$	$\frac{\sqrt[4]{m}}{\sqrt{2(m-\sqrt{m})}}$	$O(\sqrt[4]{m})$
II	$\{\xi \in [0, 1]^m \mid \ \xi\ _1 \leq k\}$	$\frac{k(m-1)}{m+k(k-2)}$	$O(\min\{k, \frac{m}{k}\})$
III	$\{\xi \in \mathbb{R}_+^m \mid \ \xi\ _p \leq 1\}$	$(2(m-1) + 2^p) \frac{p-1}{p^2} \frac{1-p}{p} (p-1) \left(\frac{p-1}{p}\right)^2$	$O\left(m \frac{p-1}{p^2}\right)$
IV	$\{\xi \in \mathbb{R}_+^m \mid \xi^T \Sigma \xi \leq 1\}$	$\begin{cases} \sqrt{\frac{1}{2} + \frac{am + \sqrt{(1-a)m + am^2}}{2(1-a)}} & a \leq m^{-\frac{2}{3}} \\ \frac{1}{\sqrt{a}} & a > m^{-\frac{2}{3}} \end{cases}$	$O\left(m^{\frac{1}{3}}\right)$
V	$\mathcal{U} \subset \mathbb{R}_+^m$	$2\sqrt{m} + 1$	$O(\sqrt{m})$

We prove specific bounds for uncertainty sets of the forms I) hypersphere uncertainty; II) budgeted uncertainty; III) p -norm ball uncertainty, with $p \geq 1$; IV) ellipsoid uncertainty, with $\Sigma := (1 - a)\mathbf{1} + a\mathbf{J}$ where $\mathbf{1}$ is the unity matrix and \mathbf{J} the matrix of all ones; V) general uncertainty sets

We present the proof for Lemma 3 in ‘‘Appendix D’’.

With these simplifications, Condition (5) becomes

$$\frac{1}{\rho} \max_{\xi \in \mathcal{U}} \sum_{i=1}^m (\xi_i - \mu)_+ \leq 1. \tag{7}$$

With the permutation invariance of the problem, Ben-Tal et al. [14] show that for any μ there exists a $j \leq m$, such that the maximization problem in (7) has a solution that is constant on the first j components and zero on all others components.

Lemma 4 (Lemma 4 in Ben-Tal et al. [14]) *Let $\gamma(j)$ be the maximal average value of the first j components of any $\xi \in \mathcal{U}$*

$$\gamma(j) := \frac{1}{j} \max_{\xi \in \mathcal{U}} \sum_{i=1}^j \xi_i.$$

Then for each μ there exists an optimal solution ξ^* for the maximization problem in Eq. (7) that has the form

$$\xi^* = \sum_{i=1}^j \gamma(j) e_i,$$

for some $j \leq m$.

Hypersphere uncertainty: We first use Lemma 4 to find a new dominating set for hypersphere uncertainty. By doing so, we find a new approximation bound that improves the bound of $\sqrt[4]{m}$ provided in Ben-Tal et al. [14] by a factor of $\sqrt{\frac{\sqrt{m}+1}{2\sqrt{m}}}$,

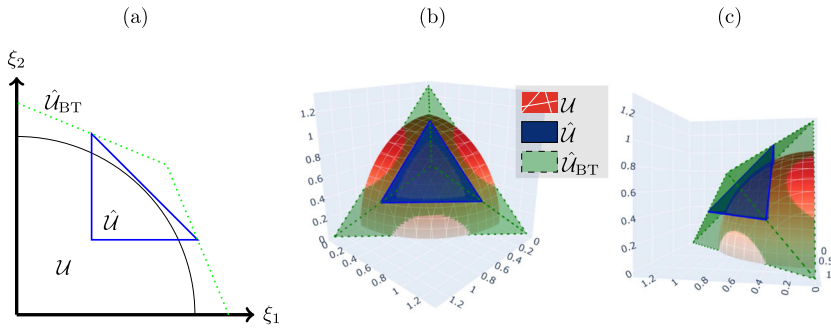


Fig. 4 Comparison of our dominating set \hat{U} (blue, solid frame) and the dominating set \hat{U}_{BT} proposed in Ben-Tal et al. [14] (green, dashed frame) for the hypersphere uncertainty set U in $m = 2$ (a) and $m = 3$ (b), (c) uncertainty dimensions (color figure online)

which for large m converges towards $\frac{1}{\sqrt{2}}$. While this improvement is irrelevant for the asymptotic complexity of the problem, the new formulation of \hat{U} does make a difference in practice. In Fig. 4 we illustrate the improvement of our dominating set \hat{U} over the dominating set \hat{U}_{BT} proposed in Ben-Tal et al. [14] for hypersphere uncertainty sets in two and three uncertainty dimensions. Note, that our sets \hat{U} are fully contained in the sets \hat{U}_{BT} , and all extreme points of \hat{U}_{BT} are located outside of \hat{U} . This implies $Z_{AR}(\hat{U}) \leq Z_{AR}(\hat{U}_{BT})$ for hypersphere uncertainty. The formal proof follows from straightforward convex containment and is left for brevity.

Proposition 5 (Hypersphere) *Consider the hypersphere uncertainty set $U = \{\xi \in \mathbb{R}_+^m \mid \|\xi\|_2^2 \leq 1\}$. Then a solution for $Z_{AR}(\hat{U})$ where \hat{U} is constructed using Criterion (7) with*

$$\mu = \frac{1}{2\sqrt[4]{m}}, \quad \rho = \frac{\sqrt[4]{m}}{2},$$

gives a $\beta = \sqrt{\frac{\sqrt{m}+1}{2}}$ approximation for problem $Z_{AR}(U)$.

We present the proof for Proposition 5 in ‘‘Appendix E’’.

We can also use this improved performance bound to show that affine adjustable policies cannot yield better bounds than piecewise affine adjustable policies for $m \geq 153$. This is because there are instances of Problem (1) with hypersphere uncertainty where affine adjustable policies perform at least $\frac{4}{5} \left(\sqrt[4]{m} - \frac{1}{\sqrt[4]{m}} \right)$ worse than an optimal policy. We formalize these results in Proposition 6. Note, that these better bounds do not imply that piecewise affine adjustable policies always yield better results than affine adjustable policies for hypersphere uncertainty.

Proposition 6 *Affine adjustable policies cannot achieve better performance bounds than $\frac{4}{5} \left(\sqrt[4]{m} - \frac{1}{\sqrt[4]{m}} \right)$ for Problem (1) with hypersphere uncertainty, even for c, x, A being non negative and A being a 0,1-matrix.*

We present the proof for Proposition 6 in “Appendix F”.

Budgeted uncertainty: Next, we tighten the bounds for budgeted uncertainty sets. Proposition 7 shows that our new bound is given by $\beta = \frac{k(m-1)}{m+k(k-2)}$. Using

$$\frac{\beta}{k} = \frac{k(m-1)}{k(m+k(k-2))} = \frac{m-1}{m-1+k^2-2k+1} = \frac{m-1}{m-1+(k-1)^2} \leq 1$$

and

$$\frac{\beta}{\frac{m}{k}} = \frac{k^2(m-1)}{m(m+k(k-2))} = \frac{k^2(m-1)}{k^2(m-1)+k^2+m^2-2km}$$

$$= \frac{k^2(m-1)}{k^2(m-1)+(m-k)^2} \leq 1,$$

we show $\beta \leq \min(k, \frac{m}{k})$, which matches the bound for the two-stage problem variant in Ben-Tal et al. [14]. As $\frac{\beta}{k}$ is decreasing in k and $\frac{\beta}{\frac{m}{k}}$ is increasing in k , we obtain a maximum improvement for $k = \frac{m}{k} \Leftrightarrow k = \sqrt{m}$. At this point the improvement of the bound reaches a factor of $\frac{1}{2}$.

Proposition 7 (Budget) Consider the budgeted uncertainty set $\mathcal{U} = \{\xi \in [0, 1]^m \mid \|\xi\|_1 \leq k\}$ for some $k \in \{1, \dots, m\}$. Then a solution for $Z_{AR}(\hat{\mathcal{U}})$ where $\hat{\mathcal{U}}$ is constructed using Criterion (7) with

$$\mu = \frac{k(k-1)}{m+k(k-2)}, \quad \rho = \frac{k(m-k)}{m+k(k-2)},$$

gives a $\beta = \frac{k(m-1)}{m+k(k-2)}$ approximation for problem $Z_{AR}(\mathcal{U})$.

We present the proof for Proposition 7 in “Appendix G”.

Note, that there is no result analogous to Proposition 6 for budgeted uncertainty as Housni and Goyal [40] showed that affine policies are in $O\left(\frac{\log(m)}{\log \log(m)}\right)$ for two-stage problems with non-negative c, x, A . Furthermore, our piecewise affine policies are strictly dominated by affine policies for integer budgeted uncertainty.

Proposition 8 Consider Problem (1) with budgeted uncertainty and an integer budget. Let Z_{PAP} be the optimal value found by our piecewise affine policy and Z_{AFF} be the optimal value found by an affine policy. Then

$$Z_{AFF} \leq Z_{PAP}.$$

We present the proof for Proposition 8 in “Appendix H”.

Norm ball uncertainty: In a similar manner as before, we construct new dominating sets for p -norm ball uncertainty and tighten the bound in Ben-Tal et al. [14] by a factor of $2^{-1+\frac{1}{p}-\frac{1}{p^2}} p^{\frac{1}{p}} (p-1)^{\frac{1}{p^2}-\frac{1}{p}}$ for sufficiently large m . This factor is always smaller than one and converges to $\frac{1}{2}$ for large p .

Proposition 9 (*p*-norm ball) Consider the *p*-norm ball uncertainty set $\mathcal{U} = \{\xi \in \mathbb{R}_+^m \mid \|\xi\|_p \leq 1\}$ with $p > 1$. Then a solution for $Z_{AR}(\hat{\mathcal{U}})$ where $\hat{\mathcal{U}}$ is constructed using Criterion (7) with

$$\begin{aligned} \mu &= 2^{\frac{1}{p}} (2(m-1) + 2^p)^{-\frac{1}{p^2}} p^{-1} (p-1)^{\frac{1}{p} + (1-\frac{1}{p})^2}, \\ \rho &= 2^{\frac{1}{p}-1} (2(m-1) + 2^p)^{\frac{1}{p} - \frac{1}{p^2}} p^{-1} (p-1)^{(1-\frac{1}{p})^2} \end{aligned}$$

gives a

$$\begin{aligned} \beta &\leq (2(m-1) + 2^p)^{\frac{1}{p} - \frac{1}{p^2}} p^{\frac{1}{p}-1} (p-1)^{(1-\frac{1}{p})^2} \\ &\leq \left((2m)^{\frac{1}{p} - \frac{1}{p^2}} + 2^{1-\frac{1}{p}} \right) p^{\frac{1}{p}-1} (p-1)^{(1-\frac{1}{p})^2} = O\left(m^{\frac{1}{p} - \frac{1}{p^2}}\right) \end{aligned}$$

approximation for problem $Z_{AR}(\mathcal{U})$.

We present the proof for Proposition 9 in ‘‘Appendix I’’.

Ellipsoid uncertainty: For the permutation invariant ellipsoid uncertainty set $\{\xi \in \mathbb{R}_+^m \mid \xi^T \Sigma \xi \leq 1\}$ with $m > 1$, $\Sigma := \mathbf{1} + a(\mathbf{J} - \mathbf{1})$, $a \in [0, 1]$, $\mathbf{1}$ being the unity matrix, and \mathbf{J} being the matrix of all ones, we construct dominating sets via a case distinction on the size of a . While for large a already a scaled simplex gives a good approximation, we construct the dominating set for small a more carefully. By doing so, we improve the previously best known asymptotic bound for the two-stage problem variant of $O(m^{\frac{2}{3}})$ [14] to $O(m^{\frac{1}{3}})$. Note, that for $a = 0$ our bounds converge to the bounds of hypersphere uncertainty in Proposition 5 and for $a = 1$ towards an exact representation.

Proposition 10 (Ellipsoid) Consider the ellipsoid uncertainty set $\mathcal{U} = \{\xi \in \mathbb{R}_+^m \mid \xi^T \Sigma \xi \leq 1\}$ with $m > 1$ and $\Sigma := \mathbf{1} + a(\mathbf{J} - \mathbf{1})$ for $a \in [0, 1]$. Here $\mathbf{1}$ is the unity matrix and \mathbf{J} is the matrix of all ones. Then a solution for $Z_{AR}(\hat{\mathcal{U}})$ where $\hat{\mathcal{U}}$ is constructed using Criterion (7) with

$$\begin{aligned} \mu &= \frac{1}{2\sqrt[4]{(1-a)^3 m + (1-a)^2 a m^2}}, & \rho &= \frac{1}{4(1-a)\mu} & \text{if } a \leq m^{-\frac{2}{3}}, \\ \mu &= 0, & \rho &= \frac{1}{\sqrt{a}} & \text{if } a > m^{-\frac{2}{3}}, \end{aligned}$$

gives a

$$\beta = \begin{cases} \sqrt{\frac{1}{2} \left(1 + \frac{1}{1-a} \left(a m + \sqrt{(1-a)m + a m^2} \right) \right)} & \text{if } a \leq m^{-\frac{2}{3}} \\ \frac{1}{\sqrt{a}} & \text{if } a > m^{-\frac{2}{3}} \end{cases} = O(m^{\frac{1}{3}})$$

approximation for problem $Z_{AR}(\mathcal{U})$.

We present the proof for Proposition 10 in ‘‘Appendix J’’.

General uncertainty sets: After having shown specific bounds for some commonly used permutation invariant uncertainty sets, we now give a general bound that holds for all uncertainty sets that fulfill the assumptions of Problem (1). We show that any uncertainty set can be dominated within an approximation factor of $\beta = 2\sqrt{m} + 1$, which improves the bound in Ben-Tal et al. [14] by a factor of $\frac{1}{2}$. As shown in Ben-Tal et al. [14] this approximation bounds is asymptotically tight, when using pure domination techniques. More precisely, for any polynomial number of vertices the budgeted uncertainty set with $k = \sqrt{m}$ cannot be dominated with some β better than $O(\sqrt{m})$.

Proposition 11 (General uncertainty) *Consider any uncertainty set $\mathcal{U} \subseteq [0, 1]^m$ that is convex, full-dimensional with $\mathbf{e}_i \in \mathcal{U}$ for all $i \in \{1, \dots, m\}$ and down-monotone. Then, there always exists a dominating uncertainty set $\hat{\mathcal{U}}$ of the form in (2) that dominates \mathcal{U} by at most a factor of $\beta = 2\sqrt{m} + 1$.*

We present the proof for Proposition 11 in ‘‘Appendix K’’.

Stagewise uncertainty sets: In general, our policies do not require stagewise uncertainty. However, the existence of such a structure can be utilized in the construction of dominating uncertainty sets, leading to approximation bounds that depend linearly on the stagewise approximation bounds.

Proposition 12 (Stagewise) *Let $\mathcal{U} = \mathcal{U}_1 \times \dots \times \mathcal{U}_T$ be a stagewise independent uncertainty set and for each \mathcal{U}_t , let $\hat{\mathcal{U}}_t$ be a dominating set constructed as in (2). Let β_t be the approximation factor for $\hat{\mathcal{U}}_t$, and let $\beta'_t = \min\{\beta' : \frac{1}{\beta'} \mathbf{e} \in \mathcal{U}_t\}$ be the constant approximation factor for set \mathcal{U}_t . Then for any partition $\mathcal{T}_1 \cup \mathcal{T}_2 = \{1, \dots, T\}$, $\mathcal{T}_1 \cap \mathcal{T}_2 = \emptyset$ of the stages, there exists a dominating set $\hat{\mathcal{U}}$ for \mathcal{U} with approximation factor*

$$\beta \leq \max \left(\sum_{t \in \mathcal{T}_1} \beta_t, \max_{t \in \mathcal{T}_2} \beta'_t \right).$$

We present the proof for Proposition 12 in ‘‘Appendix L’’.

Transformed uncertainty sets: We note that by the right-hand side $\mathbf{D}\boldsymbol{\xi} + \mathbf{d}$ of Problem (1) any positive affine transformation of uncertainty sets \mathcal{U} can be dominated by the same affine transformation of the dominating set $\hat{\mathcal{U}}$. As the approximation bounds do not depend on \mathbf{D} and \mathbf{d} , the bounds for the transformed set are the same as for the original set. One well-known uncertainty type covered by these transformations is scaled ellipsoidal uncertainty $\{\boldsymbol{\xi} \mid \sum_{i=1}^m w_i \xi_i^2 \leq 1\}$, which was first proposed by Ben-Tal and Nemirovski [7]. These sets can be constructed via transformations from hypersphere uncertainty sets with a diagonal matrix \mathbf{D} with $D_{ii} = \frac{1}{\sqrt{w_i}}$. Scaled ellipsoidal uncertainty has been applied to many robust optimization problems, including portfolio optimization [7], supply chain contracting [10], network design [47], and facility location [5].

Another widely used class that is partially covered by these positive affine transformations are factor-based uncertainties given by sets $\mathcal{U} = \{\mathbf{D}\mathbf{z} + \mathbf{d} \mid \mathbf{z} \in \mathcal{U}^z\}$. In these sets, uncertainties affinely depend on a set of factors \mathbf{z} that are drawn from a factor uncertainty set \mathcal{U}^z . Problems that were solved using such uncertainty sets include, among others, portfolio optimization [34, 35] and multi-period inventory management [2, 57]. In contrast to the general factor sets, that have no limitations on \mathbf{D} and \mathbf{d} , our approach is restricted to positive factor matrices which allows only for positive correlations between uncertainties. Nevertheless, even this subset of factor-based uncertainties has wide applicational use. As an intuitive example, one could consider component demands where the factors are demands for finished products.

4 Re-scaling expensive vertices

On some instances, uncertainties do not have a uniform impact on the objective. While accounting for high values in some of the uncertainty dimensions might drastically influence the objective, accounting for high values in others might barely have an impact. This effect can be crucial in our problem setting, because the creation of dominating sets may overemphasize single uncertainty dimensions by up to a factor of β by design. Accordingly, it might thus be beneficial to dominate these uncertainty dimensions more carefully on instances where a few critical uncertainties cause almost all the cost. In this context, we show that it is possible to shrink critical vertices \mathbf{v}_i at the cost of slightly shifting all other vertices towards their direction.

We illustrate the re-scaling process following from Lemma 13 in Fig. 5. In the depicted example we shrink the critical vertex \mathbf{v}_1 and shift the two remaining vertices $\mathbf{v}_0, \mathbf{v}_2$ towards uncertainty dimension ξ_1 . As a consequence, the cost of the vertex solution \mathbf{x}_1 decreases, while the costs of the solutions $\mathbf{x}_0, \mathbf{x}_2$ increase. The cost reduction on the critical vertex \mathbf{v}_1 leads to a reduction of the worst case vertex cost z which by the construction of LP (3) corresponds to an overall improvement of the objective function. Note, that the dominating set $\hat{\mathcal{U}}$ used in the example is not an optimal choice for the hypersphere uncertainty set depicted. However, the effect would barely be visible in two dimensions without this sub-optimal choice.

Lemma 13 *Let $\hat{\mathcal{U}} := \text{conv}(\mathbf{v}_0, \dots, \mathbf{v}_m)$ be a dominating set for \mathcal{U} . Let $\mathbf{s} \in [0, 1]^m$ be a vector of scales. Then $\hat{\mathcal{U}}' := \text{conv}(\mathbf{v}'_0, \dots, \mathbf{v}'_m)$ with $v'_{ij} := s_j(1 - v_{ij}) + v_{ij}$ is also a dominating set for \mathcal{U} .*

We present the proof for Lemma 13 in “Appendix M”.

The two extreme cases for the modified dominating sets from Lemma 13 are given by $\mathbf{s} = \mathbf{0}$ and $\mathbf{s} = \mathbf{e}$. While the dominating set does not change for $\mathbf{s} = \mathbf{0}$, all vertices become the unit vector consisting of ones in every component for $\mathbf{s} = \mathbf{e}$. Intuitively, increasing s_i leads to shifting the i th component towards one. In the same way as we constructed our dominating sets in (2) this shift of the i th component towards one increases the value for all \mathbf{v}_j with $j \neq i$ and decrease it for $j = i$. Note that this Lemma is not limited to uncertainty sets that are constructed as described in (2), but holds for any dominating set created as a convex combination of vertices.




Fig. 5 Re-scaling of the expensive vertex v_1 in a dominating set $\hat{\mathcal{U}}$ for uncertainty set \mathcal{U} . **a** Shows the change of dominating set $\hat{\mathcal{U}}$ (blue, dashed) and its vertices v_0, v_1, v_2 to the new re-scaled dominating set $\hat{\mathcal{U}}'$ (green, solid) with vertices v'_0, v'_1, v'_2 for $s_1 = 0.5, s_2 = 0$. **b** Shows the costs for the vertex solutions x_i with maximal cost z (blue, dashed) compared to the costs for the re-scaled vertex solutions x'_i with maximal cost z' (green, solid) (color figure online)

The transformation used in Lemma 13 is linear, which allows us to add s as a further decision variable to the second constraint of LP (3) for a given $\hat{\mathcal{U}}$. As $s = \mathbf{0}$ gives the original dominating set, any optimal solution found with these additional decision variables is at least as good as a non-modified solution. Thus, all performance bounds shown in Sect. 3 also hold for these re-scaled piecewise affine policies. Note that Proposition 8 also extends to the re-scaled uncertainty set; thus, all re-scaled piecewise affine policies are strictly dominated by affine policies for integer budgeted uncertainty.

Adding s to the LP increases its size, which in practice will often lead to an increase of solution times. To limit the increase of model size, it is possible to add only those s_i where one expects $s_i > 0$, as not adding a variable s_i is equivalent to fixing $s_i = 0$. Those s_i with $s_i > 0$ correspond to the critical uncertainty dimensions, and an experienced decision maker with sufficient knowledge of the problem might be able to identify them a priori.

5 Piecewise affine policies via liftings

Georghiou et al. [31] propose piecewise affine policies via liftings. In this section, we strengthen these policies by using the insights from the policies constructed in Sects. 2 and 3.

To construct piecewise affine policies via liftings, in the context of Assumption 1, we first choose $r_i - 1$ breakpoints

$$0 < z_1^i < z_2^i < \dots < z_{r_i-1}^i < 1,$$

for each uncertainty dimension $i \in \{1, 2, \dots, m\}$. For ease of notation let $z_0^i := 0$, $z_{r_i}^i := 1$. With these breakpoints, we define the lifting operator $L: \mathbb{R}^m \rightarrow \mathbb{R}^{m^L}$, where $m^L := \sum_{i=1}^m r_i$, componentwise by

$$L_{i,j}(\xi) := \left(\min(\xi_i - z_{j-1}^i, z_j^i - z_{j-1}^i) \right)_+.$$

Further, we define the linear retraction operator $R: \mathbb{R}^{m^L} \rightarrow \mathbb{R}^m$ componentwise by

$$R_i(\xi^L) := \sum_{j=1}^{r_i} \xi_{i,j}^L.$$

Here $\xi_{i,j}^L$ are the components of the m^L dimensional vector

$$\xi^L = (\xi_{1,1}^L, \xi_{1,2}^L, \dots, \xi_{1,r_1}^L, \xi_{2,1}^L, \dots, \xi_{m,1}^L, \dots, \xi_{m,r_m}^L)^\top \in \mathbb{R}^{m^L}.$$

Note that $R \circ L: \mathcal{U} \rightarrow \mathcal{U}$ is the identity. Finally, Georghiou et al. [31] construct a lifted uncertainty set \mathcal{U}^L via

$$\mathcal{U}^L := \left\{ \begin{array}{l} \xi^L \in \mathbb{R}_+^{m^L} : R(\xi^L) \in \mathcal{U}, \\ \frac{\xi_{i,j+1}^L}{z_{j+1}^i - z_j^i} \leq \frac{\xi_{i,j}^L}{z_j^i - z_{j-1}^i} \quad \forall i \in \{1, \dots, m\}, j \in \{1, \dots, r_i - 1\} \\ \xi_{i,1}^L \leq z_1^i \quad \forall i \in \{1, \dots, m\} \end{array} \right\}. \tag{8}$$

Uncertainty set (8) is an outer approximation of the lifting $L(\mathcal{U}) \subseteq \mathcal{U}^L$ and omits $R(\mathcal{U}^L) = \mathcal{U}$. Replacing the uncertainty in Problem (1) with this lifted uncertainty set yields the lifted adjustable problem

$$\begin{aligned} Z_{AR}^L(\mathcal{U}^L) &= \min_{x(\xi^L)} \max_{\xi^L \in \mathcal{U}^L} c^\top x(\xi^L) \\ \text{s.t.} \quad & Ax(\xi^L) \geq DR(\xi^L) + d \quad \forall \xi^L \in \mathcal{U}^L. \end{aligned} \tag{9}$$

Limiting x to affine policies in the lifted space, yields piecewise affine policies in the original space, which give tighter approximations than affine policies in the original space, i.e., $Z_{AFF}^L(\mathcal{U}^L) \leq Z_{AFF}(\mathcal{U})$ [31]. However, \mathcal{U}^L is not a tight outer approximation of $L(\mathcal{U})$, leading to little or no improvements over affine policies on some instances [17, 31]. In fact, we show that in the framework of ARO, the piecewise affine policies induced by lifted uncertainty (8) are equivalent to classical affine policies, in the sense that for any optimal feasible lifted affine policy, there is an affine policy with the same objective value and vice versa.

Proposition 14 *Let Z_{AFF} be the optimal objective for affine policies on $Z_{AR}(\mathcal{U})$ and let Z_{LIFT} be the optimal objective value for lifted affine policies on $Z_{AR}^L(\mathcal{U}^L)$. Then*

$$Z_{AFF} = Z_{LIFT}.$$

We present the proof for Proposition 14 in ‘‘Appendix N’’.

We overcome this shortcoming in the construction of lifted affine policies using our results on dominating uncertainty sets from Sect. 2. Consider the lifting with one break-point v_{0i} per uncertainty dimension. Here v_{0i} is the i th component of the base vector \mathbf{v}_0 from Sect. 2. Then the lifting operator L becomes

$$L_{ij}(\boldsymbol{\xi}) = \begin{cases} \min(\xi_i, v_{0i}) & \text{if } j = 1 \\ (\mathbf{h}(\boldsymbol{\xi}) - \mathbf{v}_0)_i & \text{if } j = 2. \end{cases}$$

With this construction of L it is easy to verify that by Condition (5) each $\boldsymbol{\xi}^L \in L(\mathcal{U})$ satisfies $\sum_{i=1}^m \frac{\xi_{i,2}^L}{\rho_i} \leq 1$. Accordingly, we can tighten the lifted uncertainty set \mathcal{U}^L and get the new lifted uncertainty set

$$\hat{\mathcal{U}}^L := \left\{ \boldsymbol{\xi}^L \in \mathcal{U}^L, \sum_{i=1}^m \frac{\xi_{i,2}^L}{\rho_i} \leq 1 \right\}. \tag{10}$$

By construction $\hat{\mathcal{U}}^L$ is an outer approximation of $L(\mathcal{U})$ and we have $L(\mathcal{U}) \subseteq \hat{\mathcal{U}}^L \subseteq \mathcal{U}^L$ and $R(\hat{\mathcal{U}}^L) = \mathcal{U}$. Thus, affine policies on the lifted problem with uncertainty set $\hat{\mathcal{U}}^L$ yield valid piecewise affine policies for the original problem. Furthermore, the construction of $\hat{\mathcal{U}}^L$ guarantees that the lifted policies yield tighter approximations than our piecewise affine policies via domination and classical lifted policies with the same breakpoints. Consequently, all approximation bounds for piecewise affine policies via domination also hold for the strengthened piecewise affine policies via lifting.

Proposition 15 *Let Z_{LIFT} be the optimal objective value found by the lifting policies with breakpoints \mathbf{v}_0 and lifted uncertainty set \mathcal{U}^L defined in (8), Z_{TLIFT} be the optimal objective value found by the lifting policies with breakpoints \mathbf{v}_0 and tightened lifted uncertainty set $\hat{\mathcal{U}}^L$ defined in (10), and Z_{SPAP} be the optimal objective value found by the piecewise affine policies with re-scaling described in Sect. 4. Then*

$$Z_{TLIFT} \leq Z_{LIFT}, \qquad Z_{TLIFT} \leq Z_{SPAP}.$$

We present the proof for Proposition 15 in ‘‘Appendix O’’.

6 Numerical experiments

In this section, we present two numerical experiments to compare the performance of our piecewise affine policies for the different constructions of $\hat{\mathcal{U}}$ and our tightened

Table 2 Overview of policies compared in the experiments

Policy	Dominating set or policy function	
BOX	$\hat{U}_{\text{BOX}} = \{e\}$	Constant policy
AFF	$x^t(\underline{\xi}^t) = P^t \underline{\xi}^t + q^t$	Affine adjustable policy [9]
BG	$x^t(\underline{\xi}^t) = \max_i \bar{P}_i^t \underline{\xi}^t + \bar{q}_i^t - \max_i \underline{P}_i^t \underline{\xi}^t + \underline{q}_i^t$	Near-optimal pap [17]
PAPBT	$\hat{U}_{\text{BT}} = \beta \text{conv}(e_1, \dots, e_m, v)$	Literature dominating set [14]
PAP	$\hat{U} = \text{conv}(v_0, v_1, \dots, v_m)$	Our dominating set (c.f., Sect. 2.1)
SPAP	$\hat{U} = \text{conv}(v'_0, v'_1, \dots, v'_m)$	PAP with re-scaling (c.f., Sect. 4)
TLIFT	AFF on \hat{U}^L	Tightened pap via lifting (c.f., Sect. 5)

piecewise affine policies via lifting with the performance of other policies from the literature. We compare the performance in terms of both objective value and computational time.

We run both of the following tests with hypersphere uncertainty sets and budgeted uncertainty sets. In the experiments of Ben-Tal et al. [14] piecewise affine policies performed particularly well compared to affine adjustable policies for hypersphere uncertainty and relatively bad for budgeted uncertainty. Accordingly, considering these two uncertainty types gives a good impression of the benefits and limitations of piecewise affine policies. Additionally, this experimental design allows us to analyze whether or not the new formulations with the tighter bounds presented in Propositions 5 and 7 have a significant impact in practice.

In our studies, we compare the following policies: the affine policies described in Ben-Tal et al. [9] (AFF), the constant policies resulting from a dominating set $\hat{U} = \{e\}$ with only a single point which by down-monotonicity corresponds to a box (BOX), the near-optimal piecewise affine policies with two pieces proposed in Bertsimas and Georghiou [17] (BG), our piecewise affine policies constructed as described in Propositions 5 and 7 (PAP), the piecewise affine policies constructed as described in Propositions 1 and 5 in Ben-Tal et al. [14] (PAPBT), our piecewise affine policies with the vertex re-scaling heuristic described in Sect. 4 (SPAP), and our tightened piecewise affine policies via lifting described in Sect. 5 (TLIFT). Note, that piecewise affine policies via lifting from Georghiou et al. [31] are implicitly included in the comparison by Proposition 14. In Table 2 we give an overview of all policies compared in our experiments.

For all studies, we used Gurobi Version 9.5 on a 6 core 3.70 GHz i7 8700K processor using a single core per instance.

6.1 Gaussian instances

We base our first set of benchmark instances on the experiments of Ben-Tal et al. [14] and Housni and Goyal [40]. Accordingly, we generate instances of Problem (1) by choosing $m = l = n$, $d = \mathbf{0}$, $D = \mathbf{1}_m$ and generate c , A randomly as

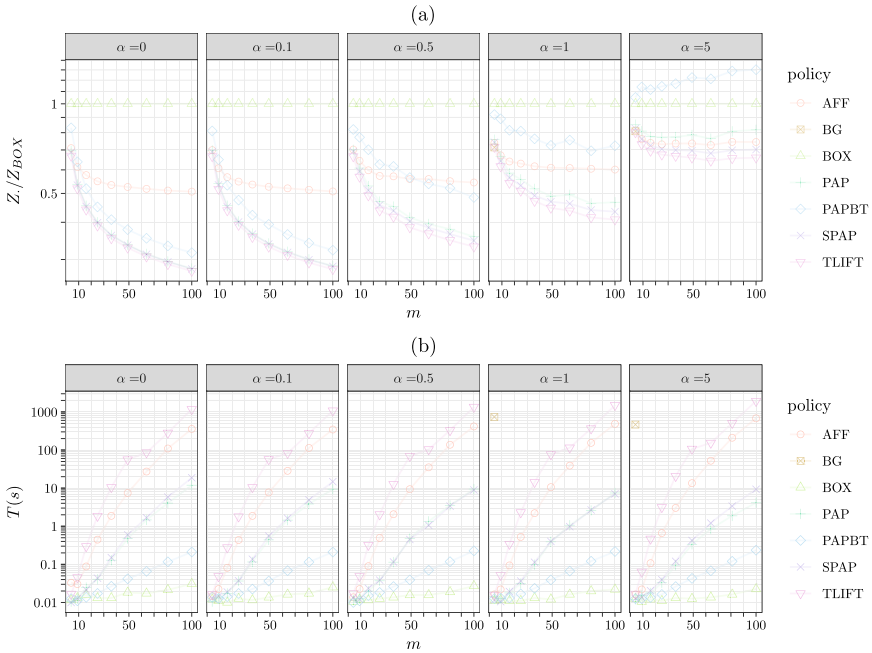


Fig. 6 Relative objective values **(a)** and solution times **(b)** for different policies on gaussian instances with hypersphere uncertainty

$$\begin{aligned}
 \mathbf{c} &= \mathbf{e} + \alpha \mathbf{g}, \\
 \mathbf{A} &= \mathbf{1} + \mathbf{G}.
 \end{aligned}$$

Here, \mathbf{e} is the vector of all ones, $\mathbf{1}$ is the identity matrix, \mathbf{g} and \mathbf{G} are randomly generated by independent and identically distributed (i.i.d.) standard Gaussians, and α is a parameter that increases the asymmetry of the problem. More specifically, \mathbf{G} is given by $G_{ij} = |Y_{ij}|/\sqrt{m}$ and \mathbf{g} is given by $g_i = |y_i|$, where Y_{ij} and y_i are i.i.d. standard Gaussians. Uncertainties ξ and decision variables \mathbf{x} are split into $\lfloor \sqrt{m} \rfloor$ stages where the i th decision always belongs to the same stage as the i th uncertainty. For the budgeted uncertainty sets we use a budget of $k = \sqrt{m}$. We consider values of $m = i^2$ for $i \in \{2, \dots, 10\}$ and values of α in $\{0, 0.1, 0.5, 1, 5\}$. For each pair of m, α , we consider 30 instances. To make the results more comparable, we scale all objective values presented by the constant policies results (i.e., Z_i/Z_{BOX}) and report averages over all solved instances. For each parameter pair m, α , we only consider those policies that found solutions on at least 75% of instances within a hard solution time limit of 4 hours. Additionally, we present all results on a logarithmic scale and artificially lower bound the scale for solution times by 0.01s to make the effects on higher solution times more visible.

Figure 6 shows the performances and solution time results on hypersphere uncertainty sets for the different policies. First, note that BG only finds solutions within the time limit for the smallest instances yielding objective values comparable or marginally

better to TLIFT. For the other policies we observe that piecewise adjustable policies perform significantly better than affine adjustable policies for small values of α . The improvement increases for larger values of m reaching almost a factor of 2 for $m = 100$ on our policies PAP, SPAP, and TLIFT. As expected, the performance of PAP and SPAP for small values of α is almost indistinguishable due to the construction. Additionally, we find that TLIFT only yields marginal improvements over PAP and SPAP for small values of α . For larger values of α the improvements of the piecewise affine adjustable policies vanish and TLIFT starts to improve over SPAP. The two policies without re-scaling (PAP and PAPBT) perform even worse than AFF for $\alpha = 5$. More severely, PAPBT even performs worse than BOX, which already is a worst-case policy. Only SPAP and TLIFT achieve better results than affine adjustable policies for all values of α .

While solution times for all policies except BOX grow exponentially in the instance size, domination-based piecewise affine adjustable policies are by orders of magnitude faster than classical affine adjustable policies (AFF) and piecewise affine policies via lifting (TPAP). These solution time improvements exceed a factor of 100 for piecewise affine policies PAP and SPAP and a factor of 1000 for PAPBT. Also, solution times of domination-based piecewise affine adjustable policies are barely influenced by values of α . This is not the case for AFF and TLIFT, which take longer to solve for increasing α . While this effect is not easily visible in Fig. 6 due to the logarithmic scale, the solution time difference for AFF and TLIFT between $\alpha = 0$ and $\alpha = 5$ reaches up to a factor of two on large instances.

Figure 7 shows the performances and solution time results on budgeted uncertainty sets. We observe that for budgeted uncertainty sets domination-based piecewise affine policies perform slightly worse than affine policies throughout all instances. This observation nicely demonstrates that our theoretical and experimental results are aligned, as the worse performance is perfectly explained by Proposition 8, which shows that affine policies strictly dominate our piecewise affine policies. Again, we observe that for higher values of α , PAP and PAPBT perform even worse than BOX. However, the solution values for SPAP stay within 5% of the affine solution values throughout all instances. We further observe that TLIFT yields the same objective values as AFF throughout all instances. Only BG yields slightly better solutions than AFF on some instances. Solution times behave similarly to solution times on hypersphere uncertainty, confirming that piecewise affine policies are found by orders of magnitude faster on different instances and uncertainty types. Only for BG solution times improve significantly, suggesting that BG is highly dependent on the shape of the uncertainty sets.

Solution time and performance results for $\alpha = 0$ align with the results found by Ben-Tal et al. [14] for the two-stage problem variant. This demonstrates that our generalized piecewise affine policies do not only extend all theoretical performance bounds, but also achieve comparable numerical results in a multi-stage setting. However, by breaking the symmetry by increasing α , we show that pure domination-based piecewise adjustable policies perform poorly on highly asymmetric instances and re-scaling (SPAP) or tightened lifting (TLIFT) constitute good techniques to overcome this shortcoming.

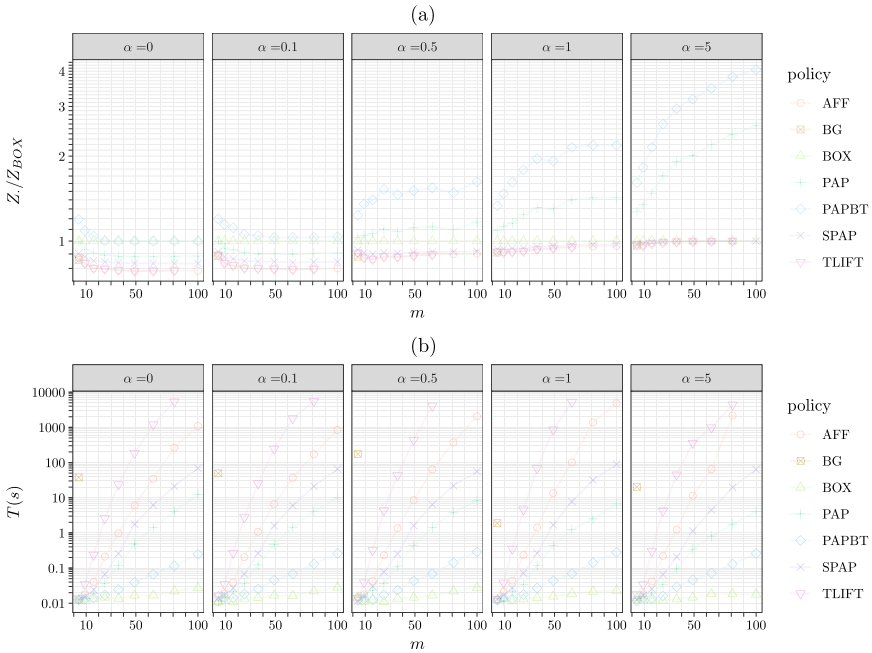


Fig. 7 Relative objective values (a) and solution times (b) for different policies on gaussian instances with budgeted uncertainty

6.2 Demand covering instances

For the second set of test instances, we consider the robust demand covering problem with non-consumed resources and uncertain demands. The problem has various applications, among others in the domains of appointment scheduling, production planning, and dispatching and is especially relevant for the optimization of service levels. Our instances consist of m^l locations, m^p planning periods, and m^e execution periods per planning period. In each execution period t , an uncertain demand ξ_{tl} has to be covered at each location l . To do so, the decision maker can buy a fixed number of resources R at a unit cost of c^R in the first stage and then distributes these R resources among the locations at the beginning of each planning period. If a demand cannot be met with the resources assigned to a location, the decision maker will either delay the demand to the next period or redirect it to another location. In either case a fraction $q_{tl}^d \in [0, 1]$ or $q_{tll'}^r \in [0, 1]$ of the demand is lost. Each unit of lost demand causes costs of c^D . Mathematically, the robust demand covering problem with non-consumed resources and uncertain demands is given by the robust LP (11), where parameters and variables are summarized in Table 3.

$$\min c^R R + c^D \left(\sum_{t \in T, l \in \mathcal{L}} q_{tl}^d s_{tl}^d + \sum_{t \in T, l \neq l' \in \mathcal{L}} q_{tll'}^r s_{tll'}^r \right) \tag{11a}$$

Table 3 Notation for the robust demand covering problem with non-consumed resources and uncertain demands

<i>Sets & elements</i>	
$l \in \mathcal{L}$	Locations
$p \in \mathcal{P}$	Planning periods
$t \in \mathcal{T}$	All $(m^p m^e)$ execution periods
<i>Parameters</i>	
c^R	Cost of buying one unit of resources
c^D	Cost of loosing one unit of demand
$q_t^d \in [0, 1]$	Fraction of demand lost when delayed in period t
$q_{ll'}^r \in [0, 1]$	Fraction of demand lost when redirected form location l to l'
<i>Uncertainties</i>	
$p(t)$	Planning period containing execution period t
ξ_{tl}	Demand at location $l \in \mathcal{L}$ in execution period $t \in \mathcal{T}$
<i>Decision variables</i>	
R	Number of resources bought (decided in the very first stage)
r_{pl}	Number of resources assigned to location $l \in \mathcal{L}$ in planning period $p \in \mathcal{P}$ (decided at the beginning of planning period p)
$s_{ill'}^r$	Demand redirected from location l to location l' (decided in decision stage t)
s_{tl}^d	Demand delayed to the next period (decided in decision stage t)

$$\text{s.t. } r_{p(t)l} + s_{tl}^d - (1 - q_{t-1}^d)s_{(t-1)l}^d + \sum_{l' \in \mathcal{L}, l' \neq l} (s_{ill'}^r - (1 - q_{ll'}^r)s_{ll'}^r) \geq \xi_{tl} \quad \forall \xi \in \mathcal{U}, \forall t \in \mathcal{T}, l \in \mathcal{L} \quad (11b)$$

$$\sum_{l \in \mathcal{L}} r_{pl} \leq R \quad \forall p \in \mathcal{P} \quad (11c)$$

$$R, r, s^r, s^d \geq 0 \quad (11d)$$

Here Objective (11a) minimizes the sum of resource costs and lost demand costs due to delay and relocation. Constraints (11b) ensure that all demands are fulfilled, delayed, or relocated, and Constraints (11c) upper bound the allocated resources in each planning period by the total number of available resources R .

In most real-world applications of the demand covering problem, some of the demand will be revealed before the actual demand occurs, e.g. due to already existing contracts, sign-ups, orders, or due to forecasting. To incorporate the increase in knowledge over time, we assume an uncertainty vector of the form $\xi = \xi^c + \xi^p + \xi^e$. Here, ξ^c is constant and known before the first stage decision, ξ^p is revealed before each planning period and ξ^e accounts for the short-term uncertainties revealed before

each execution period. Specifically, we assume that demands are given by

$$\xi_{lt} = d_{lt} \left(1 + \xi_{lt}^p + \frac{1}{2} \xi_{lt}^e \right), \tag{12}$$

where d_{lt} is the base demand for location l in execution period t . Here the uncertainty vector (ξ^p, ξ^e) is taken from one base uncertainty set \mathcal{U}^B of dimension $m = 2m^l m^p m^e$. Note, that short-term uncertainties have a less severe effect on demands than uncertainties known in advance. The resulting demand uncertainty set

$$\mathcal{U} = \left\{ \xi \mid \xi_{lt} = d_{lt} \left(1 + \xi_{lt}^p + \frac{1}{2} \xi_{lt}^e \right), (\xi^p, \xi^e) \in \mathcal{U}^B \right\}$$

is $m/2$ dimensional, where in our experiments \mathcal{U}^B is either a hypersphere uncertainty set, or a budgeted uncertainty set with budget \sqrt{m} .

To construct our instances we draw $m^l \in \{2, 4, 6, 8, 10\}$ locations at uniform random integer positions in the square $[0, 2\lfloor\sqrt{m^l}\rfloor + 1]^2$. In each of the $m^p \in \{1, 3, 5, 7\}$ planning periods, we consider $m^e = 8$ execution periods corresponding to the hours in a working day. We assume that a fraction $q_{lt}^d = 0.1$ of the demand is lost when deferred to a later execution period and consider a doubled loss rate ($q_{lt}^d = 0.2$) when demand is deferred to another planning period. Similarly, a fraction of the demand is lost when assigned to another location. We assume this fraction to be correlated to the distance and given by $q_{ll'}^r := \min(1, 0.02 \cdot \text{dist}(l, l'))$. We draw the base demands d_{lt} uniformly from the normal distribution $\mathcal{N}(10, 4)$. Finally, we set $c^R = 1$ and choose $c^D \in \{0.1, 0.25, 0.5\}$. For each combination of m^l, m^p we consider 45 instances, where we use each possible value for $c^D \in \{0.1, 0.25, 0.5\}$ in a third of these instances.

To analyze practical expected objectives, we also report a simulated average objective that the respective policies achieved on 500 randomly drawn uncertainty realizations ξ , in addition to the robust objective value. We give a detailed description of how uniform uncertainty realizations can efficiently be sampled from the budgeted and hypersphere uncertainty sets in ‘‘Appendix P’’. We again scale the results by the results achieved by constant policies and use logarithmic scales. For each instance size, we only consider those policies that found solutions on at least 75% of instances within a hard solution time limit of 2 hours.

First, we observe that BG did not solve any instance within the time limit which can be attributed to the fact that our demand covering instances are significantly larger than our gaussian instances.

For the remaining policies, Fig. 8 shows the performance and solution time results on demand covering instances with hypersphere uncertainty. Compared to our previous experiment (see Sect. 6.1) we no longer observe the strong objective improvements of piecewise affine policies over affine adjustable policies. Still, our piecewise affine formulations give similar results as affine policies, and we observe small improvements on instances with a larger number of planning periods, with TLIFT yielding strict improvements for $m^p \geq 3$. On the simulated realizations, improvements of PAP and

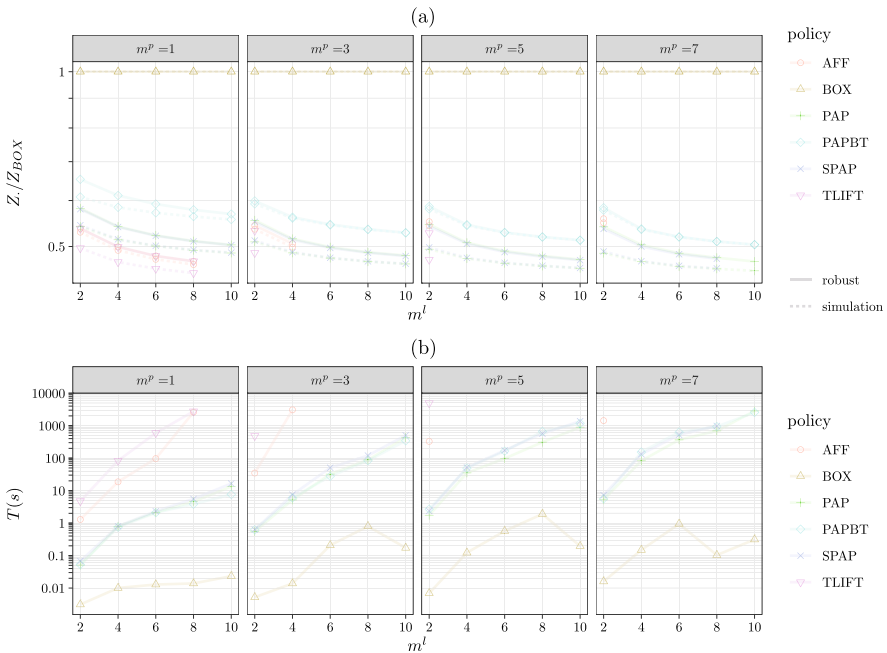


Fig. 8 Relative objective values (a) and solution times (b) for different policies on demand covering instances with hypersphere uncertainty

SPAP over affine adjustable policies can already be seen for $m^p \geq 3$, which might be of interest for a decision maker with practical interest beyond worst-case solutions.

For the solution times, we observe similar improvements to the ones observed on the gaussian instances. Still, all domination-based piecewise affine policies can be found by orders of magnitude faster than affine policies and TLIFT. Interestingly, PAP is solved similarly fast as PAPBT on these instances, while still achieving up to 15% better objective values on all instances. The largest instance that could be solved by affine adjustable policies within two hours consisted of 320 uncertainty variables and 700 decision variables, while the largest instance solved by PAPs was more than three times larger with 1, 120 uncertainty variables and 6, 230 decision variables.

Figure 9 shows the results on demand covering instances with budgeted uncertainty sets. For budgeted uncertainty sets, domination-based piecewise affine policies perform worse than affine adjustable policies throughout all instances, which again can be explained by Proposition 8. Notably, PAPBT even performs worse than constant policies on most instances. In this setting, domination-based piecewise affine policies remain better only from a solution time perspective, as they still solve by orders of magnitude faster than affine policies. Also, TLIFT does not yield any improvements over AFF.

While in this set of experiments piecewise affine policies do not show the same improvements in the objective over affine adjustable policies, they still perform slightly better with hypersphere uncertainty on larger instances. Also, they still solve by orders

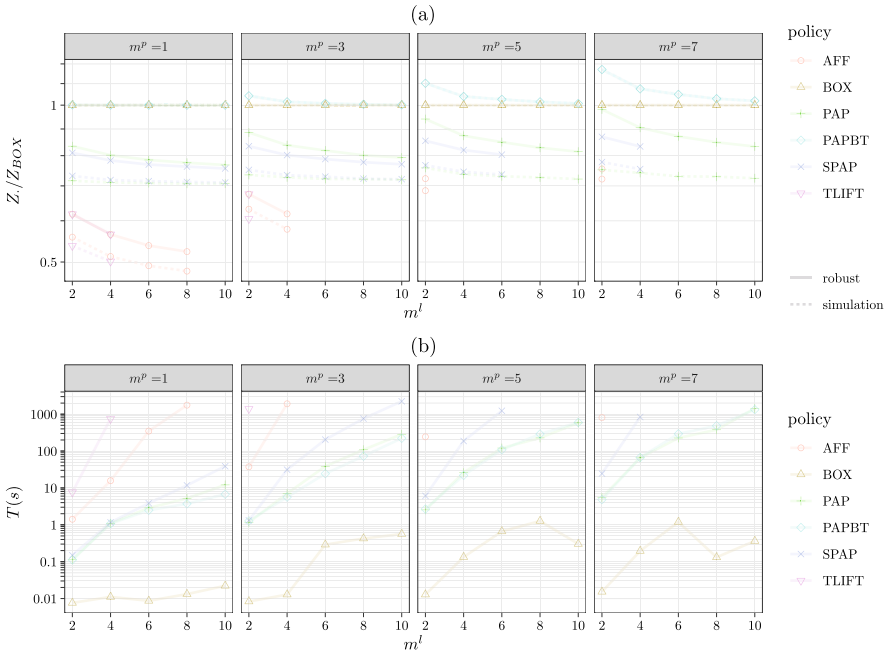


Fig. 9 Relative objective values (a) and solution times (b) for different policies on demand covering instances with budgeted uncertainty

of magnitude faster, which makes them an attractive alternative for large-scale optimization in practice.

6.3 Discussion

In the experiments presented in Sects. 6.1 and 6.2, some results, e.g. the solution time improvements of piecewise affine policies over affine policies, are consistent throughout all instances. However, other results strongly depend on the set of benchmark instances used. In the following, we give an explanation for the strong solution time improvements, discuss two of the main deviations that we observe between our results on gaussian instances and demand covering instances, and give intuitions of why these differences occur.

Size of robust counterparts: Throughout all experiments, we see strong solution time improvements of piecewise affine policies over affine policies—including affine policies on lifted uncertainty TLIFT. These can be explained by their respective robust counterparts. The robust counterpart for piecewise affine policies is given by LP (3). For the robust counterparts of affine policies, we refer to Ben-Tal et al. [9]. Counterparts of piecewise affine policies have $O(nm)$ variables compared to $O(nm + lm)$ variables for affine policies and both counterparts have $O(lm)$ constraints. More critically, constraints for piecewise affine policies contain at most $O(n)$ variables each

and feature a block structure, which solvers use to significantly speed up the solution process. This block structure is only connected by the nonanticipativity constraints. In the robust counterpart of affine policies on the other hand, $O(l)$ constraints have up to $O(nm)$ variables resulting in a denser constraint matrix and the lack of a block structure. Moreover, the robust counterpart for affine policies on hypersphere uncertainty sets is no longer linear. Instead, a quadratic program has to be solved, which tends to be computationally more challenging.

Solution times of PAPBT: In the experiments we see that PAPBT solves by a factor of 10 to 50 faster than PAP on gaussian instances, but both find solutions similarly fast on-demand covering instances. This can be explained by the construction of PAPBT and the structure of the instances' constraints. On the gaussian instances, \mathbf{A} and \mathbf{x} are non-negative and \mathbf{D} is the unit matrix. In the construction of dominating sets $\hat{\mathcal{U}}$ for PAPBT most of the vertices are chosen to be scaled unit vectors. As a consequence, most Constraints (3c) in LP (3) have a zero right-hand side, such that they trivially hold. Consequently, these constraints can be eliminated, which reduces the total number of constraints by a factor of $O(m)$. On demand covering instances, however, \mathbf{A} contains negative entries. Thus, no constraints trivially hold and no constraints can be removed. For a decision maker who is primarily interested in fast policies, this gives a good criterion on when PAPBT can improve solution times and when no such improvements can be expected.

Performance differences between gaussian and demand covering instances: We observe that the strong performance improvements of piecewise affine policies over affine policies on gaussian instances with hypersphere uncertainty do not transfer to our demand covering instances. This suggests that the relative performance between piecewise affine policies and affine policies significantly depends on the structure of the problem at hand. An intuitive explanation for this lies in the policies' construction.

Recall that piecewise affine policies derive solutions by finding vertex solutions \mathbf{x}_i that can be extended to a full solution. Thereby, \mathbf{x}_0 focuses on finding a good solution for uncertainty realizations where all uncertainties take equal values, and each \mathbf{x}_i focuses on finding a good recourse to uncertainty ξ_i . Thus, good results can be expected when there are (a) synergy effects that can be utilized by \mathbf{x}_0 , and (b) good universal recourse decisions for each uncertainty ξ_i that do not depend on the realization of other uncertainty dimensions and can be exploited by \mathbf{x}_i . On the other hand, affine policies directly find solutions on the original uncertainty set. In doing so, they do not depend as strongly on good universal recourse decisions as piecewise affine policies do. However, they also lack the ability to use synergy effects in the way vertex solutions \mathbf{x}_0 do.

The gaussian instances used fulfill both of the properties that are favorable for piecewise affine policies. By being based on a unity matrix, \mathbf{A} has relatively large values along the diagonal, leading to the existence of good universal recourse decisions. Additionally, the relatively small non-negative entries on the non-diagonals lead to synergy effects for uncertainty realizations with many small values. Demand covering instances, however, do not fulfill these properties. The question of how to redirect demand optimally heavily depends on the demand observed at other locations. Also,

the only synergistic effects that can be used solely emerge when multiple demands occur at the same location in a single planning period.

On general instances in practice, we would thus not expect to see the same performance improvements that could be observed on our gaussian benchmark instance. Still, piecewise affine policies find solutions by orders of magnitude faster than affine policies and achieve good results throughout all benchmark instances with hypersphere uncertainty. Additionally, Properties (a) and (b) give intuitive criteria on when strong objective improvements over affine policies can be expected.

7 Conclusion

In this work, we presented piecewise affine policies for multi-stage adjustable robust optimization. We construct these policies by carefully approximating uncertainty sets with a dominating polytope, which yields a new problem that we efficiently solve with a linear program. By making use of the problem's structure, we then extend solutions for the new problem with approximated uncertainty to solutions for the original problem. We show strong approximation bounds for our policies that extend many previously best-known bounds for two-stage ARO to its multi-stage counterpart. By doing so, we contribute towards closing the gap between the state of the art for two-stage and multi-stage ARO. To the best of our knowledge, the bounds we give are the first bounds shown for the general multi-stage ARO Problem. Furthermore, our bounds yield constant factor as well as asymptotic improvements over the state-of-the-art bounds for the two-stage problem variant.

In two numerical experiments, we find that our policies find solutions by a factor of 10 to 1000 faster than affine adjustable policies, while mostly yielding similar or even better results. Especially for hypersphere uncertainty sets our new policies perform well and sometimes even outperform affine adjustable policies up to a factor of two. We observe particularly high improvements on instances that exhibit certain synergistic effects and allow for universal recourse decisions. However, on some instances where few uncertainty dimensions have a high impact on the objective, pure piecewise affine policies perform particularly badly by design, sometimes even worse than constant policies. To mitigate this shortcoming, we present an improvement heuristic that significantly improves the solution quality by re-scaling the critical uncertainty dimension. Furthermore, we construct new tightened piecewise affine policies via lifting that integrate the two frameworks of piecewise affine policies via domination and piecewise affine policies via lifting and combine their approximative power.

While this work extends most of the best-known approximation results for a relatively general class of ARO problems from the two-stage to the multi-stage setting, it remains an open question whether other strong two-stage ARO results can be generalized to multi-stage ARO in a similar manner. Answering this question remains an interesting area for further research. In this context, binary and uncertain recourse decisions remain particularly relevant challenges. Our analysis in Sect. 2.3 has shown that the extension of our policies to encompass these recourse decision types is not straightforward. Nevertheless, exploring the integration of our methodology into the established approaches of piecewise constant policies and k -adaptability, which have

proven to be effective in these cases, appears as a promising starting point for future work. Another interesting area for future research is the extension of piecewise affine policies and the concept of domination to adjustable data-driven and distributionally robust optimization. More specifically, we believe that one can obtain tractable data-driven policies by directly fitting the polyhedral uncertainty sets used to construct our policies from data.

Funding Open Access funding enabled and organized by Projekt DEAL.

Declarations

Conflict of interest The authors have no conflicts of interest to declare that are relevant to the content of this article.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix A Comparison of literature approximation bounds

Table 4 summarizes existing approximation bounds for some commonly used uncertainty sets in two-stage ARO. We want to point out that in addition to being less restrictive, our bounds for multi-stage ARO presented in Table 1 outperform all these bounds except the ones by Housni and Goyal [40], which require the significant restrictive assumption of A , x , c being non-negative.

Appendix B Proof of Theorem 1

Proof We split the proof in two parts. First, we handle the cases where at least $Z_{AR}(\mathcal{U})$ or $Z_{AR}(\hat{\mathcal{U}})$ is negative and show that this already implies that both $Z_{AR}(\mathcal{U})$ and $Z_{AR}(\hat{\mathcal{U}})$ are unbounded. Second, we assume $Z_{AR}(\mathcal{U})$, $Z_{AR}(\hat{\mathcal{U}}) \geq 0$ are bounded and prove that in this case, the desired performance bounds hold.

Part 1: Let $\tilde{\mathcal{U}} \in \{\mathcal{U}, \hat{\mathcal{U}}\}$ such that $Z_{AR}(\tilde{\mathcal{U}}) < 0$. Then, there exists a solution \tilde{x} such that

$$\max_{\xi \in \tilde{\mathcal{U}}} c^\top \tilde{x}(\xi) < 0.$$

Now, assume that $Z_{AR}(\mathcal{U})$ or $Z_{AR}(\hat{\mathcal{U}})$ is bounded and let $\bar{\mathcal{U}} \in \{\mathcal{U}, \hat{\mathcal{U}}\}$ such that $Z_{AR}(\bar{\mathcal{U}})$ is bounded with an optimal solution \bar{x} . Arbitrarily fix one $\tilde{\xi} \in \tilde{\mathcal{U}}$ and consider the constant vector $\tilde{x}(\tilde{\xi})$. Then $\bar{x}(\tilde{\xi}) + \tilde{x}(\tilde{\xi})$ is a feasible solution for $Z_{AR}(\bar{\mathcal{U}})$ as for

Table 4 Approximation bounds for two-stage ARO

No	[20]	[18]	[15]	[14]	[40]
I	$1 + \sqrt{m}$	–	$\frac{1}{2}(1 + \sqrt{m})$	$\sqrt[4]{m}$	–
II	$1 + \frac{m}{k}$	–	$\frac{k^2 + km}{k^2 + m}$	$\min(k, \frac{m}{k})$	$\frac{8 \log m}{\log \log m} + 1$
III	$1 + \sqrt[p]{m}$	–	$\frac{m^{\frac{p-1}{p}} + m}{m^{\frac{p-1}{p}} + m^{\frac{1}{p}}}$	$\frac{2}{p}(p-1) \frac{p-1}{p} m^{\frac{p-1}{p^2}}$	–
IV	–	–	–	$\left(\frac{a}{2} + \frac{\sqrt{1-a}}{4\sqrt{am^2+(1-a)m}}\right)^{-1}$	–
V	m	$4\sqrt{m}$	–	$4\sqrt{m}$	–
Policy	Static	Affine	Affine	Piecewise affine	Affine

We compare bounds for uncertainty sets of the forms I) hypersphere uncertainty $\{\xi \in \mathbb{R}_+^m \mid \|\xi\|_2 \leq 1\}$; II) budgeted uncertainty $\{\xi \in [0, 1]^m \mid \|\xi\|_1 \leq k\}$; III) p -norm ball uncertainty $\{\xi \in \mathbb{R}_+^m \mid \|\xi\|_p \leq 1\}$, with $p \geq 1$; IV) ellipsoid uncertainty $\{\xi \in \mathbb{R}_+^m \mid \xi^T \Sigma \xi \leq 1\}$, with $\Sigma := (1-a)\mathbf{1} + a\mathbf{J}$ where $\mathbf{1}$ is the unity matrix and \mathbf{J} the matrix of all ones; V) general uncertainty sets $\mathcal{U} \subset \mathbb{R}_+^m$. Note, that some of the results require more restrictive assumptions on A, x, c than our Problem (1). Most prominently [40] requires A, x, c to be non-negative

all $\bar{\xi} \in \bar{\mathcal{U}}$:

$$A(\bar{x}(\bar{\xi}) + \tilde{x}(\bar{\xi})) \stackrel{(a)}{\geq} D(\bar{\xi} + \tilde{\xi}) + 2d \stackrel{(b)}{\geq} D\bar{\xi} + d.$$

Here, (a) holds because \tilde{x}, \bar{x} are feasible solutions and (b) holds as $\tilde{\xi}, D, d$ are non-negative. For the objective we then find

$$\max_{\bar{\xi} \in \bar{\mathcal{U}}} c^T(\bar{x}(\bar{\xi}) + \tilde{x}(\bar{\xi})) \leq \max_{\bar{\xi} \in \bar{\mathcal{U}}} c^T \bar{x}(\bar{\xi}) + \max_{\bar{\xi} \in \bar{\mathcal{U}}} c^T \tilde{x}(\bar{\xi}) < Z_{AR}(\bar{\mathcal{U}}).$$

This is a contradiction to \bar{x} being a minimal solution. Thus, $Z_{AR}(\bar{\mathcal{U}})$ cannot be bounded. We have thus shown that if any of $Z_{AR}(\mathcal{U}), Z_{AR}(\hat{\mathcal{U}})$ are negative, both have to be unbounded.

Part 2: Assume $Z_{AR}(\mathcal{U}), Z_{AR}(\hat{\mathcal{U}}) \geq 0$ are bounded. Let \hat{x} be an optimal solution for $Z_{AR}(\hat{\mathcal{U}})$. Furthermore, let $h: \mathcal{U} \rightarrow \hat{\mathcal{U}}$ be the domination function from Definition 1. We claim that $\tilde{x} := \hat{x} \circ h$ is a feasible solution for $Z_{AR}(\mathcal{U})$. First, we see that by the definition of h the solution \tilde{x} fulfills the nonanticipativity requirements. Specifically, we have

$$\tilde{x}(\xi) = \hat{x}(h(\xi)) = \left(\hat{x}^1\left(h^1(\xi^1)\right), \dots, \hat{x}^T\left(h^T(\xi^T)\right)\right)$$

where the decisions in stage t depend only on the uncertainty revealed up to that stage. For the constraints, we find

$$A\tilde{\mathbf{x}}(\boldsymbol{\xi}) = A\hat{\mathbf{x}}(\mathbf{h}(\boldsymbol{\xi})) \stackrel{(a)}{\geq} D\mathbf{h}(\boldsymbol{\xi}) + \mathbf{d} \stackrel{(b)}{\geq} D\boldsymbol{\xi} + \mathbf{d}.$$

Here, (a) follows from the feasibility of $\hat{\mathbf{x}}$ and (b) follows from the definition of \mathbf{h} and D being non-negative. Thus, $\tilde{\mathbf{x}}$ is a well-defined feasible solution for $Z_{AR}(\mathcal{U})$ and we have

$$Z_{AR}(\mathcal{U}) \leq \max_{\boldsymbol{\xi} \in \mathcal{U}} \mathbf{c}^\top \tilde{\mathbf{x}}(\boldsymbol{\xi}) \leq \max_{\hat{\boldsymbol{\xi}} \in \hat{\mathcal{U}}} \mathbf{c}^\top \hat{\mathbf{x}}(\hat{\boldsymbol{\xi}}) = Z_{AR}(\hat{\mathcal{U}}).$$

For the other direction let \mathbf{x}^* be an optimal solution for $Z_{AR}(\mathcal{U})$. Then, for all $\hat{\boldsymbol{\xi}} \in \hat{\mathcal{U}}$ we have $\frac{1}{\beta}\hat{\boldsymbol{\xi}} \in \mathcal{U}$ by definition of β . We define $\tilde{\mathbf{x}}(\hat{\boldsymbol{\xi}}) := \beta\mathbf{x}^*\left(\frac{1}{\beta}\hat{\boldsymbol{\xi}}\right)$ and find

$$A\tilde{\mathbf{x}}(\hat{\boldsymbol{\xi}}) = \beta A\mathbf{x}^*\left(\frac{1}{\beta}\hat{\boldsymbol{\xi}}\right) \stackrel{(a)}{\geq} \beta\left(\frac{1}{\beta}D\hat{\boldsymbol{\xi}} + \mathbf{d}\right) \stackrel{(b)}{\geq} D\hat{\boldsymbol{\xi}} + \mathbf{d}$$

where (a) follows from the feasibility of \mathbf{x}^* and (b) from the non-negativity of \mathbf{d} . Thus, $\tilde{\mathbf{x}}$ is a well-defined feasible solution for $Z_{AR}(\hat{\mathcal{U}})$ and we have

$$Z_{AR}(\hat{\mathcal{U}}) \leq \max_{\hat{\boldsymbol{\xi}} \in \hat{\mathcal{U}}} \mathbf{c}^\top \tilde{\mathbf{x}}(\hat{\boldsymbol{\xi}}) \leq \beta \max_{\boldsymbol{\xi} \in \mathcal{U}} \mathbf{c}^\top \mathbf{x}^*(\boldsymbol{\xi}) = \beta Z_{AR}(\mathcal{U}).$$

Having shown both inequalities this concludes the proof. \square

Appendix C Proof of Lemma 2

Proof We begin by showing $Z_{AR}(\hat{\mathcal{U}}) \geq Z_{LP}(\hat{\mathcal{U}})$. Let \mathbf{x} be an optimal solution of $Z_{AR}(\hat{\mathcal{U}})$. Then, (\mathbf{x}_i, z) is a valid solution for $Z_{LP}(\hat{\mathcal{U}})$, where \mathbf{x}_i, z are defined by

$$\begin{aligned} \mathbf{x}_i &:= \mathbf{x}(v_i) & \forall i \in \{0, \dots, m\}, \\ z &:= \max_{i \in \{0, \dots, m\}} \mathbf{c}^\top \mathbf{x}_i. \end{aligned}$$

The first set of LP constraints (3b) holds by definition of z . The second set of constraints (3c) holds as \mathbf{x} is a solution of $Z_{AR}(\hat{\mathcal{U}})$ and $v_i \in \hat{\mathcal{U}}$, and the last set of constraints (3d) holds by nonanticipativity of \mathbf{x} and the definition of v_i . Furthermore, we find

$$Z_{LP}(\hat{\mathcal{U}}) \leq z = \max_{i \in \{0, \dots, m\}} \mathbf{c}^\top \mathbf{x}_i \leq \max_{\boldsymbol{\xi} \in \hat{\mathcal{U}}} \mathbf{c}^\top \mathbf{x}(\boldsymbol{\xi}) = Z_{AR}(\hat{\mathcal{U}}).$$

For the other direction let (x_i, z) be an optimal solution for $Z_{LP}(\hat{U})$. Define $\lambda_i(\xi)$ for each $\xi \in \hat{U}$ by:

$$\lambda_i(\xi) := \frac{(\xi - v_0)_i}{\rho_i}.$$

For ease of notation we will in the following drop the explicit dependence on ξ and use λ_i . We directly find

$$\sum_{i=1}^m \lambda_i \leq 1$$

for all $\xi \in \hat{U}$, as by definition of \hat{U} each ξ is a convex combination of v_0, \dots, v_m . Also, note that λ_i only depends on the i th component of the uncertainty vector ξ . Using this we can define a nonanticipative decision vector $x(\xi) = (x^1(\underline{\xi}^1), \dots, x^T(\underline{\xi}^T))$ by

$$x^t(\underline{\xi}^t) := \sum_{i \in \underline{I}^t} \lambda_i x_i^t + \left(1 - \sum_{i \in \underline{I}^t} \lambda_i\right) x_0^t$$

where $\underline{I}^t := \{i \in \{1, \dots, m\} | e_i^t \neq 0\}$ is the index set corresponding to the uncertainties up to stage t .

By the nonanticipativity constraints (3d) we have $x_i^t = x_0^t$ for all $1 \leq t \leq T$ and $i \in \{1, \dots, m\} \setminus \underline{I}^t$. Thus,

$$x^t(\underline{\xi}^t) = \sum_{i \in \underline{I}^t} \lambda_i x_i^t + \left(1 - \sum_{i \in \underline{I}^t} \lambda_i\right) x_0^t = \sum_{i=1}^m \lambda_i x_i^t + \left(1 - \sum_{i=1}^m \lambda_i\right) x_0^t,$$

which shows that x is a convex combination of x_0, \dots, x_m by

$$x(\xi) = \sum_{i=1}^m \lambda_i x_i + \left(1 - \sum_{i=1}^m \lambda_i\right) x_0.$$

Using this, we find

$$\begin{aligned} Ax(\xi) &= \sum_{i=1}^m \lambda_i Ax_i + \left(1 - \sum_{i=1}^m \lambda_i\right) Ax_0 \\ &\stackrel{(a)}{\geq} \sum_{i=1}^m \lambda_i Dv_i + \left(1 - \sum_{i=1}^m \lambda_i\right) Dv_0 + d \stackrel{(b)}{=} D\xi + d. \end{aligned}$$

Here, (a) holds because all x_i are valid solutions for (3), and (b) holds as by definition of λ_i we have $\xi = \sum_{i=1}^m \lambda_i v_i + (1 - \sum_{i=1}^m \lambda_i) v_0$.

Now we find the missing inequality

$$\begin{aligned}
 Z_{AR}(\hat{\mathcal{U}}) &\stackrel{(a)}{\leq} \max_{\xi \in \hat{\mathcal{U}}} \mathbf{c}^\top \mathbf{x}(\xi) = \max_{\xi \in \hat{\mathcal{U}}} \sum_{i=1}^m \lambda_i \mathbf{c}^\top \mathbf{x}_i + \left(1 - \sum_{i=1}^m \lambda_i\right) \mathbf{c}^\top \mathbf{x}_0 \\
 &\stackrel{(b)}{\leq} \max_{\xi \in \hat{\mathcal{U}}} \max_{i \in \{0, \dots, m\}} \mathbf{c}^\top \mathbf{x}_i = \max_{i \in \{0, \dots, m\}} \mathbf{c}^\top \mathbf{x}_i = Z_{LP}(\hat{\mathcal{U}}),
 \end{aligned}$$

where we use for (a) that \mathbf{x} is a valid solution for $Z_{AR}(\hat{\mathcal{U}})$. For (b), we use that \mathbf{x} is a convex combination of $\mathbf{x}_0, \dots, \mathbf{x}_m$ with convex coefficients $(1 - \sum_{i=1}^m \lambda_i), \lambda_1, \dots, \lambda_m$. The convex sum is upper bounded by its largest summand $\mathbf{c}^\top \mathbf{x}_i$. We can thus take the maximum over all these $\mathbf{c}^\top \mathbf{x}_i$, which equals the LP solution value $Z_{LP}(\hat{\mathcal{U}})$. \square

Appendix D Proof of Lemma 3

Proof Let \mathfrak{S} be the set of all permutations on $\{1, \dots, m\}$ and for any vector ξ let $\sigma(\xi)$ be the vector with components permuted according to σ . Let $\mathfrak{S}_{ji} \subseteq \mathfrak{S}$ be the subset of all permutations mapping component j to component i . Let $\hat{\mathcal{U}}$ be a dominating uncertainty set for \mathcal{U} of the form in (2), such that the approximation factor β is minimal. Let $\mathbf{v}_0, \rho_1, \dots, \rho_m$ be the parameters defining the vertices of $\hat{\mathcal{U}}$. Define $\mu := \frac{1}{m} \mathbf{e}^\top \mathbf{v}_0, \rho := \frac{1}{m} \sum_{i=1}^m \rho_i$. Then for each $i \in \{1, \dots, m\}$

$$\begin{aligned}
 \mu \mathbf{e} + \rho \mathbf{e}_i &\stackrel{(a)}{=} \frac{1}{|\mathfrak{S}|} \sum_{\sigma \in \mathfrak{S}} \sigma(\mathbf{v}_0) + \frac{1}{m} \sum_{j=1}^m \rho_j \mathbf{e}_i \stackrel{(b)}{=} \sum_{j=1}^m \left(\frac{1}{|\mathfrak{S}|} \sum_{\sigma \in \mathfrak{S}_{ji}} \sigma(\mathbf{v}_0) + \frac{1}{m} \rho_j \mathbf{e}_i \right) \\
 &\stackrel{(c)}{=} \sum_{j=1}^m \left(\frac{1}{|\mathfrak{S}|} \sum_{\sigma \in \mathfrak{S}_{ji}} \sigma(\mathbf{v}_0) + \frac{1}{m|\mathfrak{S}_{ji}|} \sum_{\sigma \in \mathfrak{S}_{ji}} \rho_j \mathbf{e}_i \right) \\
 &\stackrel{(d)}{=} \frac{1}{|\mathfrak{S}|} \sum_{j=1}^m \sum_{\sigma \in \mathfrak{S}_{ji}} (\sigma(\mathbf{v}_0) + \rho_j \mathbf{e}_i) \stackrel{(e)}{=} \frac{1}{|\mathfrak{S}|} \sum_{j=1}^m \sum_{\sigma \in \mathfrak{S}_{ji}} \sigma(\mathbf{v}_0 + \rho_j \mathbf{e}_j),
 \end{aligned}$$

where (a) follows from symmetry and the invariance of the average component under permutations, (b) follows from $\mathfrak{S} = \bigcup_{j=1}^m \mathfrak{S}_{ji}$ for all i , (c) follows from $\sum_{\sigma \in \mathfrak{S}_{ji}} \frac{1}{|\mathfrak{S}_{ji}|} = 1$, (d) follows from $|\mathfrak{S}_{ji}| = \frac{1}{m} |\mathfrak{S}|$, and (e) follows from the linearity of permutations. For each $i \in \{1, \dots, m\}$ we now conclude

$$\begin{aligned}
 \forall j \in \{1, \dots, m\}: \frac{1}{\beta} (\mathbf{v}_0 + \rho_j \mathbf{e}_j) &\stackrel{(a)}{\in} \mathcal{U} \\
 \Rightarrow \forall j \in \{1, \dots, m\}: \sigma \left(\frac{1}{\beta} (\mathbf{v}_0 + \rho_j \mathbf{e}_j) \right) &\stackrel{(b)}{\in} \mathcal{U} \\
 \Rightarrow \frac{1}{\beta} (\mu \mathbf{e} + \rho \mathbf{e}_i) = \frac{1}{|\mathfrak{S}|} \sum_{j=1}^m \sum_{\sigma \in \mathfrak{S}_{ji}} \sigma \left(\frac{1}{\beta} (\mathbf{v}_0 + \rho_j \mathbf{e}_j) \right) &\stackrel{(c)}{\in} \mathcal{U}.
 \end{aligned}$$

Here (a) follows from $\hat{\mathcal{U}}$ being a valid dominating set with approximation factor β , (b) follows from permutation invariance of \mathcal{U} , and (c) from convexity of \mathcal{U} . Furthermore, $\frac{1}{\beta}\mu\mathbf{e} \in \mathcal{U}$ as \mathcal{U} is down-monotone. Thus the convex set induced by vertices $\mu\mathbf{e}, \mu\mathbf{e} + \rho\mathbf{e}_1, \dots, \mu\mathbf{e} + \rho\mathbf{e}_m$ is also a valid dominating set for \mathcal{U} with approximation factor β . \square

Appendix E Proof of Proposition 5

Proof Our main idea is to find μ, ρ such that (7) is always fulfilled and β is minimized. Recall that β is given by the minimal value such that $\frac{1}{\beta}\mathbf{v}_i \in \mathcal{U}$ for all \mathbf{v}_i . In the case of the hypersphere uncertainty set, this is given when $\|\mathbf{v}_i\|_2^2 \leq \beta^2$. As for each $i \in \{1, \dots, m\}$, \mathbf{v}_i is given by $\mathbf{v}_i = \mathbf{v}_0 + \rho\mathbf{e}_i$ and $\mathbf{v}_0 = \mu\mathbf{e}$ we find:

$$\|\mathbf{v}_0\|_2^2 \leq \|\mathbf{v}_i\|_2^2 = (m - 1)\mu^2 + (\mu + \rho)^2 = m\mu^2 + 2\mu\rho + \rho^2. \tag{E1}$$

This is the term for β^2 that we want to minimize in the following.

Using Lemma 4 the left-hand side in (7) becomes:

$$\begin{aligned} \frac{1}{\rho} \max_{j \in \{1, \dots, m\}} \sum_{i=1}^j (\gamma(j) - \mu)_+ &\stackrel{(a)}{=} \frac{1}{\rho} \max_{j \in \{1, \dots, m\}} j(\gamma(j) - \mu)_+ \\ &\stackrel{(b)}{=} \frac{1}{\rho} \max_{j \in \{1, \dots, m\}} j(\gamma(j) - \mu). \end{aligned}$$

For (a) we use that all summands are the same and for (b) we assume w.l.o.g. that $\max_{j \in \{1, \dots, m\}} (\gamma(j) - \mu) \geq 0$, as otherwise (7) trivially holds.

Using the property $\|\xi\|_2^2 \leq 1$ of the hypersphere uncertainty set, we find $\gamma(j) = \frac{1}{\sqrt{j}}$ for all j . Substituting $\gamma(j)$ in the above equation we find Property (7) to become:

$$\max_{j \in \{1, \dots, m\}} \sqrt{j} - j\mu \leq \rho.$$

The maximum of the left hand side is taken for $j = \frac{1}{4\mu^2}$ and for any optimal choice of μ, ρ the inequality will be tight. Thus, we find $\rho = \frac{1}{4\mu}$.

Substituting ρ in Eq. (E1), we get:

$$\beta = \sqrt{m\mu^2 + \frac{1}{2} + \frac{1}{16\mu^2}}.$$

This is minimized by $\mu = \frac{1}{2\sqrt[4]{m}}$, which gives $\rho = \frac{\sqrt[4]{m}}{2}$ and concludes the proof. \square

Appendix F Proof of Proposition 6

Proof Consider the following two-stage instance of Problem (1):

$$\begin{aligned} Z(\mathcal{U}) = \min_{\alpha, \mathbf{x}(\xi)} \max_{\xi \in \mathcal{U}} \quad & \sqrt{m}\alpha + \mathbf{e}^\top \mathbf{x}(\xi) \\ \text{s.t.} \quad & \alpha \mathbf{e} + \mathbf{x} \geq \xi \quad \forall \xi \in \mathcal{U} \\ & \alpha, \mathbf{x} \geq 0, \end{aligned}$$

where $\mathcal{U} = \{\xi \mid \|\xi\|_2^2 \leq 1\} \subset [0, 1]^m$ is the hypersphere uncertainty set, $\alpha \in \mathbb{R}$ is the first stage decision variable and $\mathbf{x} \in \mathbb{R}^m$ are the second stage decisions that may depend on the uncertainty realization ξ .

We begin by finding a lower bound for affine policies on this problem class. By symmetry of the problem there always exists an optimal affine adjustable solution $\mathbf{x}(\xi)$ of the form

$$x_i = a\xi_i + b \sum_{j \neq i} \xi_j + c,$$

for some $a, b, c \in \mathbb{R}$. To see this let \mathbf{x}^*, α be an optimal affine solution to the above problem and define

$$\mathbf{x}(\xi) := \frac{1}{|\mathfrak{S}|} \sum_{\sigma \in \mathfrak{S}} \sigma \left(\mathbf{x}^* \left(\sigma^{-1}(\xi) \right) \right),$$

where \mathfrak{S} is the set of all permutations on m elements. Then

$$\begin{aligned} \alpha \mathbf{e} + \mathbf{x}(\xi) &= \alpha \mathbf{e} + \frac{1}{|\mathfrak{S}|} \sum_{\sigma \in \mathfrak{S}} \sigma \left(\mathbf{x}^* \left(\sigma^{-1}(\xi) \right) \right) \stackrel{(a)}{=} \frac{1}{|\mathfrak{S}|} \sum_{\sigma \in \mathfrak{S}} \sigma \left(\alpha \mathbf{e} + \mathbf{x}^* \left(\sigma^{-1}(\xi) \right) \right) \\ &\stackrel{(b)}{\geq} \frac{1}{|\mathfrak{S}|} \sum_{\sigma \in \mathfrak{S}} \sigma \left(\sigma^{-1}(\xi) \right) = \xi, \end{aligned}$$

where (a) follows from the permutation invariance of $\alpha \mathbf{e}$ and (b) follows from permutation invariance of \mathcal{U} and feasibility of \mathbf{x}^* . Furthermore,

$$\begin{aligned} \max_{\xi \in \mathcal{U}} \mathbf{e}^\top \mathbf{x}(\xi) &= \frac{1}{|\mathfrak{S}|} \max_{\xi \in \mathcal{U}} \mathbf{e}^\top \sum_{\sigma \in \mathfrak{S}} \sigma \left(\mathbf{x}^* \left(\sigma^{-1}(\xi) \right) \right) \stackrel{(a)}{=} \frac{1}{|\mathfrak{S}|} \max_{\xi \in \mathcal{U}} \sum_{\sigma \in \mathfrak{S}} \mathbf{e}^\top \mathbf{x}^* \left(\sigma^{-1}(\xi) \right) \\ &\stackrel{(b)}{\leq} \frac{1}{|\mathfrak{S}|} \sum_{\sigma \in \mathfrak{S}} \max_{\xi \in \mathcal{U}} \mathbf{e}^\top \mathbf{x}^* \left(\sigma^{-1}(\xi) \right) \stackrel{(c)}{=} \frac{1}{|\mathfrak{S}|} \sum_{\sigma \in \mathfrak{S}} \max_{\xi \in \mathcal{U}} \mathbf{e}^\top \mathbf{x}^* (\xi) = \max_{\xi \in \mathcal{U}} \mathbf{e}^\top \mathbf{x}^* (\xi). \end{aligned}$$

Here, (a) follows from permutation invariance of \mathbf{e} , (b) follows from the subadditivity of a maximization function, and (c) follows from permutation invariance of \mathcal{U} .

Non-negativity of \mathbf{x} follows from $\mathbf{0} \in \mathcal{U}$. Thus \mathbf{x}, α is an optimal feasible solution. Additionally for any $\bar{\sigma} \in \mathfrak{S}$,

$$\begin{aligned} \mathbf{x}(\bar{\sigma}(\xi)) &= \frac{1}{|\mathfrak{S}|} \sum_{\sigma \in \mathfrak{S}} \sigma(\mathbf{x}^*(\sigma^{-1} \circ \bar{\sigma}(\xi))) = \bar{\sigma} \left(\frac{1}{|\mathfrak{S}|} \sum_{\sigma \in \mathfrak{S}} \bar{\sigma}^{-1} \circ \sigma(\mathbf{x}^*(\sigma^{-1} \circ \bar{\sigma}(\xi))) \right) \\ &\stackrel{(a)}{=} \bar{\sigma} \left(\frac{1}{|\mathfrak{S}|} \sum_{\sigma \in \mathfrak{S}} \sigma(\mathbf{x}^*(\sigma^{-1}(\xi))) \right) = \bar{\sigma}(\mathbf{x}(\xi)), \end{aligned}$$

where (a) follows from $(\bar{\sigma}^{-1} \circ \cdot)$ being an automorphism on \mathfrak{S} . Let σ_{ij} be the permutation switching components i and j . Then $\mathbf{x}(\sigma_{ij}(\mathbf{0})) = \sigma_{ij}(\mathbf{x}(\mathbf{0}))$ implies $x_i(\mathbf{0}) = x_j(\mathbf{0})$, $\mathbf{x}(\sigma_{ij}(\mathbf{e}_j)) = \sigma_{ij}(\mathbf{x}(\mathbf{e}_j))$ implies $x_i(\mathbf{e}_i) = x_j(\mathbf{e}_j)$ and $x_k(\mathbf{e}_i) = x_k(\mathbf{e}_j)$ for $i \neq k \neq j$, and $\mathbf{x}(\sigma_{ij}(\mathbf{e}_k)) = \sigma_{ij}(\mathbf{x}(\mathbf{e}_k))$ for $i \neq k \neq j$ implies $x_i(\mathbf{e}_k) = x_j(\mathbf{e}_k)$. This guarantees the claimed structure.

Let $\xi \in \mathcal{U}$ and for some $i \leq m$ define the vector ξ' by $\xi' := \xi - \xi_i \mathbf{e}_i$. Then ξ' is also in \mathcal{U} and by $x_i(\xi') \geq 0$ we find that every feasible solution has to fulfill

$$b \sum_{j \neq i} \xi_j + c \geq 0 \tag{F1}$$

for all $\xi \in \mathcal{U}$. Additionally, by $\alpha \mathbf{e} + \mathbf{x}(\mathbf{e}_i) \geq \mathbf{e}_i$ every feasible solution also fulfills

$$\alpha + a + c \geq 1. \tag{F2}$$

Using this we obtain a lower bound for the maximum over \mathcal{U} in the objective function as follows

$$\begin{aligned} &\max_{\xi \in \mathcal{U}} \sqrt{m}\alpha + \sum_i \left(a\xi_i + b \sum_{j \neq i} \xi_j + c \right) \\ &\stackrel{(a)}{\geq} \max_{\xi \in \mathcal{U}} \sqrt{m}\alpha + a \sum_i \xi_i = \max_{\xi \in \mathcal{U}} \sqrt{m}\alpha + a \mathbf{e}^\top \xi \\ &\stackrel{(b)}{=} \sqrt{m}\alpha + \sqrt{m}a \stackrel{(c)}{\geq} \sqrt{m}\alpha + \sqrt{m}(1 - \alpha - c) = \sqrt{m}(1 - c). \end{aligned}$$

Here (a) follows from Eq. (F1), (b) follows as the maximum is taken for $\xi = \frac{1}{\sqrt{m}}\mathbf{e}$, and (c) follows from Eq. (F2). On the other hand, using $\mathbf{0} \in \mathcal{U}$ the maximum can also be lower bounded by

$$\sqrt{m}\alpha + \sum_i c = \sqrt{m}\alpha + mc \geq mc.$$

Thus we find

$$\begin{aligned} Z_{AFF}(\mathcal{U}) &= \min_{\alpha, a, b, c} \max_{\xi \in \mathcal{U}} \sqrt{m}\alpha + \sum_i \left(a\xi_i + b \sum_{j \neq i} \xi_j + c \right) \\ &\geq \min_c \max (\sqrt{m}(1-c), mc) \stackrel{(a)}{=} \sqrt{m} - \frac{m}{m + \sqrt{m}} \geq \sqrt{m} - 1, \end{aligned}$$

where (a) holds as the maximum is taken at equality of the two terms, which is given by $c = \frac{\sqrt{m}}{m + \sqrt{m}}$.

Having found a lower bound for affine policies, we will now give an upper bound for optimal policies. Consider the policy given by $\alpha = \frac{1}{\sqrt[4]{m}}$ and $\mathbf{x} = (\xi - \alpha \mathbf{e})_+$. By construction this is feasible and we find

$$\begin{aligned} Z_{OPT}(\mathcal{U}) &\leq \max_{\xi \in \mathcal{U}} \sqrt{m}\alpha + \mathbf{e}^\top (\xi - \alpha \mathbf{e})_+ \\ &\stackrel{(a)}{=} \max_{0 \leq k \leq m} \sqrt[4]{m} + k \left(\frac{1}{\sqrt{k}} - \frac{1}{\sqrt[4]{m}} \right) \\ &= \max_{0 \leq k \leq m} \sqrt[4]{m} + \sqrt{k} - \frac{k}{\sqrt[4]{m}} \stackrel{(b)}{=} \frac{5}{4} \sqrt[4]{m}. \end{aligned}$$

Here (a) follows from Lemma 4 and (b) follows as the maximum over k is taken for $k = \frac{\sqrt{m}}{4}$. To finish the proof, we find the optimality ratio, which is given by

$$\frac{Z_{AFF}(\mathcal{U})}{Z_{OPT}(\mathcal{U})} \geq \frac{\sqrt{m} - 1}{5\sqrt[4]{m}/4} = \frac{4}{5} \left(\sqrt[4]{m} - \frac{1}{\sqrt[4]{m}} \right).$$

□

Appendix G Proof of Proposition 7

Proof We proceed analogously as in the proof for Proposition 5. Again, the idea is to find μ, ρ such that (7) is always fulfilled and β is minimized. Using the definition of \mathcal{U} and \mathbf{v}_i we find $\frac{1}{\beta} \mathbf{v}_i \in \mathcal{U}$ to be fulfilled when

$$\beta = \min \left(\frac{m}{k} \mu + \frac{1}{k} \rho, \mu + \rho \right). \quad (\text{G1})$$

Here the first term of the maximization follows from $\frac{1}{\beta} \mathbf{e}^\top \mathbf{v}_i \leq k$ and the second term from $\frac{1}{\beta} \mathbf{e}_i^\top \mathbf{v}_i \leq 1$. By Lemma 4 the left hand side in (7) becomes:

$$\frac{1}{\rho} \max_{j \in \{1, \dots, m\}} j(\gamma(j) - \mu).$$

Using the properties $\|\underline{\xi}\|_1 \leq k$ and $\xi_i \leq 1$ of the budgeted uncertainty set we find $\gamma(j) = \frac{k}{j}$ for all $j \geq k$ and $\gamma(j) = 1$ for all $j \leq k$. Inserting $\gamma(j)$ into the above, we find Property (7) to become:

$$\max_{j \in \{1, \dots, m\}} j \left(\min \left(1, \frac{k}{j} \right) - \mu \right) \leq \rho.$$

As for $\mu \geq 1$ the vector \mathbf{v}_0 would already dominate every uncertainty realization in \mathcal{U} , we can assume $\mu \leq 1$. With this, the maximum on the left-hand side is taken for $j = k$ and we find the minimal possible ρ to be $\rho = k - k\mu$.

Using this ρ together with β from Eq. (G1), we obtain:

$$\beta = \min \left(\left(\frac{m}{k} - 1 \right) \mu + 1, k - (k - 1)\mu \right).$$

As the left term in the minimization is linearly increasing in μ and the right term is linearly decreasing in μ , the minimum is taken at equality. This is the case when $\mu = \frac{k(k-1)}{m+k(k-2)}$, which gives $\rho = \frac{k(m-k)}{m+k(k-2)}$. □

Appendix H Proof of Proposition 8

To prove Proposition 8 we first introduce a more general form of piecewise affine policies. Recall that affine policies are given by $\mathbf{x}^t(\underline{\xi}^t) = \mathbf{P}^t \underline{\xi}^t + \mathbf{q}^t$, with decision variables \mathbf{P}^t and \mathbf{q}^t . By introducing additional decision variables $\underline{\zeta}$ of the same dimension as the uncertainty set, we extend affine policies to *fully piecewise affine policies* (FPAP)

$$\mathbf{x}^t(\underline{\xi}^t) = \mathbf{P}^t \left(\underline{\xi}^t - \underline{\zeta}^t \right)_+ + \mathbf{q}^t,$$

where we use the same notation for $\underline{\zeta}$, as for $\underline{\xi}$. By extending the decision matrices \mathbf{P}^t with zeroes and concatenating them to a single decision matrix

$$\mathbf{P} = \begin{pmatrix} \mathbf{P}^1 & \mathbf{0} & \mathbf{0} & \dots & \mathbf{0} \\ & \mathbf{P}^2 & \mathbf{0} & \dots & \mathbf{0} \\ & & \vdots & & \\ & & & & \mathbf{P}^T \end{pmatrix},$$

we get the compact expressions $\mathbf{x}(\underline{\xi}) = \mathbf{P}\underline{\xi} + \mathbf{q}$ and $\mathbf{x}(\underline{\xi}) = \mathbf{P}(\underline{\xi} - \underline{\zeta})_+ + \mathbf{q}$ for affine and fully piecewise affine policies, respectively.

By our initial assumption $\mathcal{U} \subseteq [0, 1]^m$, we can w.l.o.g. assume $\mathbf{0} \leq \underline{\zeta} \leq \mathbf{e}$. First, consider a policy $\mathbf{P}, \underline{\zeta}, \mathbf{q}$ with $\zeta_i < 0$. Then replacing ζ_i by 0 and \mathbf{q} by $\mathbf{q}' := \mathbf{q} - \zeta_i \mathbf{P}_{:i}$, where $\mathbf{P}_{:i}$ is the i th column of \mathbf{P} , yields an identical policy on the uncertainty set, as $(\xi_i - \zeta_i)_+ = \xi_i - \zeta_i$ for all $\underline{\xi} \in \mathcal{U}$. Similarly, if $\zeta_i > 1$, we can replace it by 1 without changing the policy, as $(\xi_i - \zeta_i)_+ = (\xi_i - 1)_+ = 0$ for all $\underline{\xi} \in \mathcal{U}$.

We can show that the fully piecewise affine adjustable policies defined above are in fact a generalization of the piecewise affine policies we introduce in this work.

Lemma 16 Consider Problem (1). Let Z_{PAP} be the optimal value found by our piecewise affine policy and let Z_{FPAP} be the optimal value found by a fully piecewise affine policy where the right hand side uncertainty is replaced by $D((\xi - \zeta)_+ + \zeta) + d$. Then

$$Z_{FPAP} \leq Z_{PAP}.$$

Proof We prove the inequality by showing that every feasible piecewise affine policy is a feasible fully piecewise affine policy. Let v_0, \dots, v_m be the chosen vertices of the piecewise affine policy and let x_0, \dots, x_m be the vertex solutions received from solving LP (3). Then the piecewise affine policy is given by Eq. (6) and we find

$$\begin{aligned} \mathbf{x}(\xi) &= \sum_{i=1}^m \lambda_i(\xi) \mathbf{x}_i + \left(1 - \sum_{i=1}^m \lambda_i(\xi)\right) \mathbf{x}_0 \\ &= \sum_{i=1}^m \lambda_i(\xi) (\mathbf{x}_i - \mathbf{x}_0) + \mathbf{x}_0 \\ &\stackrel{(a)}{=} \sum_{i=1}^m \frac{((\xi - \mathbf{v}_0)_+)_i}{\rho_i} (\mathbf{x}_i - \mathbf{x}_0) + \mathbf{x}_0 \\ &\stackrel{(b)}{=} \left(\frac{1}{\rho_1} (\mathbf{x}_1 - \mathbf{x}_0), \dots, \frac{1}{\rho_m} (\mathbf{x}_m - \mathbf{x}_0) \right) (\xi - \mathbf{v}_0)_+ + \mathbf{x}_0. \end{aligned}$$

Here (a) follows from the definition of λ_i in Sect. 2.2 and (b) follows by replacing the sum with a matrix–vector multiplication. By the nonanticipativity Constraint (3d) the vector $\mathbf{x}_i - \mathbf{x}_0$ is zero for all uncertainty dimensions associated to later stages. Thus setting

$$\begin{aligned} \mathbf{P} &:= \left(\frac{1}{\rho_1} (\mathbf{x}_1 - \mathbf{x}_0), \dots, \frac{1}{\rho_m} (\mathbf{x}_m - \mathbf{x}_0) \right) \\ \boldsymbol{\zeta} &:= \mathbf{v}_0 \\ \mathbf{q} &:= \mathbf{x}_0 \end{aligned}$$

yields a valid fully piecewise affine policy that is identical to the initial piecewise affine policy. Even with the modified right hand side, as demanded in the Lemma, the constructed fully piecewise affine policy fulfills all constraints, as by construction of the policy, and definition of the dominating function \mathbf{h} :

$$\mathbf{A}\mathbf{x}(\xi) \geq \mathbf{D}\mathbf{h}(\xi) + \mathbf{d} = \mathbf{D}((\xi - \zeta)_+ + \zeta) + \mathbf{d}.$$

□

Even though fully piecewise affine policies are slightly more flexible than affine policies in general, both are equivalent on Problem (1) under some additional assumptions.

Lemma 17 Consider Problem (1) with budgeted uncertainty and integer budget. Let Z_{FPAP} be the optimal value found by a fully piecewise affine policy with modified right hand side uncertainty as in Lemma 16 and let Z_{AFF} be the optimal value found by an affine policy. Then

$$Z_{AFF} = Z_{FPAP}.$$

For the proof of Lemma 17 we first need the following helpful result that allows us to construct a robust counterpart of Problem (1) with fully piecewise affine adjustable policies in a similar manner as robust counterparts for affine policies are constructed.

Lemma 18 Let $c \in \mathbb{R}^m$, $\zeta \in [0, 1]^m$ and $k \in \mathbb{N}_+$. Then the following two problems yield the same objective.

$$\begin{aligned} \max_{\xi} \quad & c^T (\xi - \zeta)_+ \\ \text{s.t.} \quad & e^T \xi \leq k \\ & \mathbf{0} \leq \xi \leq e \end{aligned} \tag{H1}$$

and

$$\begin{aligned} \min_{\alpha, \beta} \quad & e^T \beta + k\alpha \\ \text{s.t.} \quad & \beta + \alpha e \geq \text{diag}(e - \zeta)c \\ & \alpha \geq 0, \quad \beta \geq \mathbf{0}. \end{aligned} \tag{H2}$$

Proof Let I^k be the indices of the largest k values $c_i(1 - \zeta_i)$ and let e_{I^k} be the indicator vector of I^k . Then for Problem (H1) we find the following set of inequalities.

$$\max \{c^T (\xi - \zeta)_+ \mid e^T \xi \leq k, \mathbf{0} \leq \xi \leq e\} \tag{H3}$$

$$\stackrel{(a)}{\leq} \max \{c^T \text{diag}(e - \zeta) \xi \mid e^T \xi \leq k, \mathbf{0} \leq \xi \leq e\} \tag{H4}$$

$$\stackrel{(b)}{\leq} c^T \text{diag}(e - \zeta) e_{I^k}$$

$$\stackrel{(c)}{\leq} c^T (e_{I^k} - \zeta)_+$$

$$\stackrel{(d)}{\leq} \max \{c^T (\xi - \zeta)_+ \mid e^T \xi \leq k, \mathbf{0} \leq \xi \leq e\} \tag{H5}$$

Here (a) follows from $(\xi - \zeta)_+ \leq \text{diag}(e - \zeta) \xi$. Inequality (b) follows as e_{I^k} is an optimal choice for Problem (H4). For (c), observe that $\text{diag}(e - \zeta) e_{I^k} = (e_{I^k} - \zeta)_+$. Finally, (d) follows, as e_{I^k} is a feasible choice for Problem (H5). As (H3) and (H5) are the same problem all inequalities are in fact equalities.

To conclude the proof note that Problem (H4) is a linear optimization problem, where taking the dual yields Problem (H2). \square

Having this result, we now show Lemma 17.

Proof of Lemma 17 To prove this result we show that the robust counterparts of the two policies on Problem (1) with integer budget uncertainty are equivalent. Let $k \in \mathbb{N}_+$ be the budget.

Then Problem (1) with affine policies becomes

$$\begin{aligned}
 Z_{AFF}(\mathcal{U}) &= \min_{P,q} \max_{\xi} (c^\top P \xi | \mathbf{0} \leq \xi \leq e, e^\top \xi \leq k) + c^\top q \\
 \text{s.t.} \quad &\max_{\xi} ((D_i - A_i P) \xi | \mathbf{0} \leq \xi \leq e, e^\top \xi \leq k) \leq A_i q - d_i \quad \forall i \in \{1, \dots, l\},
 \end{aligned}$$

where A_i and D_i are the i th rows of A and D , respectively, and d_i is the i th entry of d . Dualizing the subproblems we find this to be equivalent to

$$\begin{aligned}
 Z_{AFF}(\mathcal{U}) &= \min_{P,q,\alpha^0,\dots,\alpha^l,\beta^0,\dots,\beta^l} e^\top \beta^0 + k\alpha^0 + c^\top q \\
 \text{s.t.} \quad &\beta^0 + \alpha^0 e \geq (c^\top P)^\top \\
 &e^\top \beta^i + k\alpha^i \leq A_i q - d_i \quad \forall i \in \{1, \dots, l\} \quad \text{(H6)} \\
 &\beta^i + \alpha^i e \geq (D_i - A_i P)^\top \quad \forall i \in \{1, \dots, l\} \\
 &\alpha^i \geq 0, \quad \beta^i \geq \mathbf{0} \quad \forall i \in \{1, \dots, l\}.
 \end{aligned}$$

Similarly Problem (1) with fully piecewise affine policies and modified right hand side becomes

$$\begin{aligned}
 Z_{FPAP}(\mathcal{U}) &= \min_{P,q,\zeta} \max_{\xi} (c^\top P(\xi - \zeta)_+ | \mathbf{0} \leq \xi \leq e, e^\top \xi \leq k) + c^\top q \\
 \text{s.t.} \quad &\max_{\xi} ((D_i - A_i P) (\xi - \zeta)_+ | \mathbf{0} \leq \xi \leq e, e^\top \xi \leq k) \\
 &\leq A_i q - D_i \zeta - d_i \quad \forall i \in \{1, \dots, l\} \\
 &\mathbf{0} \leq \zeta \leq e.
 \end{aligned}$$

Using Lemma 18 on the subproblems we find this to be equivalent to

$$Z_{FPAP}(\mathcal{U}) = \min_{P,q,\zeta,\alpha^0,\dots,\alpha^l,\beta^0,\dots,\beta^l} e^\top \beta^0 + k\alpha^0 + c^\top q \tag{H7a}$$

$$\text{s.t.} \quad \beta^0 + \alpha^0 e \geq (c^\top P \text{diag}(e - \zeta))^\top \tag{H7b}$$

$$e^\top \beta^i + k\alpha^i \leq A_i q - D_i \zeta - d_i \quad \forall i \in \{1, \dots, l\} \tag{H7c}$$

$$\beta^i + \alpha^i e \geq ((D_i - A_i P) \text{diag}(e - \zeta))^\top \quad \forall i \in \{1, \dots, l\} \tag{H7d}$$

$$\alpha^i \geq 0, \quad \beta^i \geq \mathbf{0} \quad \forall i \in \{1, \dots, l\} \tag{H7e}$$

$$\mathbf{0} \leq \zeta \leq e. \tag{H7f}$$

We now show that for each solution of Problem (H7) there is an equivalent solution with $\zeta = \mathbf{0}$. Let $\mathbf{P}, \mathbf{q}, \zeta, \alpha^0, \dots, \alpha^l, \beta^0, \dots, \beta^l$ be a feasible solution to Problem (H7). We replace \mathbf{P} by $\mathbf{P}' := \mathbf{P} \text{diag}(\mathbf{e} - \zeta)$, β^i by $\beta'^i := \beta^i + \text{diag}(\zeta) \mathbf{D}_i^\top$ for $i > 0$ and ζ by $\mathbf{0}$ and leave all other decision variables untouched. As none of the variables in the objective function was changed, it is sufficient to show that the new solution is again feasible. Feasibility of Constraint (H7b) directly follows from $\mathbf{P}' \text{diag}(\mathbf{e}) = \mathbf{P} \text{diag}(\mathbf{e} - \zeta)$. For Constraint (H7c), we find

$$\begin{aligned} \mathbf{e}^\top \beta'^i + k\alpha^i &\stackrel{(a)}{=} \mathbf{e}^\top \text{diag}(\zeta) \mathbf{D}_i^\top + \mathbf{e}^\top \beta^i + k\alpha^i \\ &\stackrel{(b)}{\geq} \mathbf{e}^\top \text{diag}(\zeta) \mathbf{D}_i^\top + \mathbf{A}_i \mathbf{q} - \mathbf{D}_i \zeta - d_i \stackrel{(c)}{=} \mathbf{A}_i \mathbf{q} - d_i, \end{aligned}$$

where (a) follows from the definition of β'^i , (b) from the feasibility of the original solution and (c) from $\mathbf{e}^\top \text{diag}(\zeta) \mathbf{D}_i^\top = \mathbf{D}_i \zeta$. Similarly, Constraint (H7d) is fulfilled by

$$\begin{aligned} \beta'^i + k\alpha^i \mathbf{e} &\stackrel{(a)}{=} \text{diag}(\zeta) \mathbf{D}_i^\top + \beta^i + k\alpha^i \mathbf{e} \\ &\stackrel{(b)}{\geq} \text{diag}(\zeta) \mathbf{D}_i^\top + ((\mathbf{D}_i - \mathbf{A}_i \mathbf{P}) \text{diag}(\mathbf{e} - \zeta))^\top \stackrel{(c)}{=} (\mathbf{D}_i - \mathbf{A}_i \mathbf{P}')^\top, \end{aligned}$$

where (a) and (b) again follow by the definition of β'^i and feasibility of the initial problem and (c) follows by definition of \mathbf{P}' . Finally, $\beta'^i = \beta^i + \text{diag}(\zeta) \mathbf{D}_i^\top \geq \beta^i \geq \mathbf{0}$ holds as ζ and \mathbf{D} are both non-negative. Thus we can w.l.o.g. set $\zeta = \mathbf{0}$ in Problem (H7). In doing so, we transform Problem (H7) into Problem (H6). Thus the robust counterparts of affine policies and fully piecewise affine policies with modified right hand side are equivalent. \square

Having shown these results, Proposition 8 directly follows by combining Lemmas 16 and 17. Note, that Lemma 18, which is crucial for the construction of tractable reformulations of fully piecewise affine policies, heavily depends on the linear structure of budgeted uncertainty and the existence of integer optimal solutions. Thus the results cannot easily be extended to other settings and in general one cannot hope for tractable reformulations of FPAPs.

Appendix I Proof of Proposition 9

Proof Analogously to the proofs of Propositions 5 and 7, we use Condition (7) and Lemma 4 to find a valid choice of μ and ρ . Using the property $\|\mathbf{e}\|_p \leq 1$ of the p -norm ball uncertainty set, we find $j\gamma(j)^p \leq 1$ for all $j \in \{1, \dots, m\}$ and thus

$$\gamma(j) = j^{-\frac{1}{p}}.$$

Inserting $\gamma(j)$ in Condition (7), a valid choice of ρ is given by

$$\max_{j \in \{1, \dots, m\}} j \left(j^{-\frac{1}{p}} - \mu \right) \stackrel{(a)}{\leq} \mu^{1-p} p^{-p} (p-1)^{p-1} =: \rho, \quad (11)$$

where (a) follows, as the maximum is taken at $j = (p\mu)^{-p}(p-1)^p$. Using the definition of \mathcal{U} and \mathbf{v}_i we find $\frac{1}{\beta} \mathbf{v}_i \in \mathcal{U}$ to be fulfilled when

$$\begin{aligned} \beta^p &= (m-1)\mu^p + (\mu + \rho)^p \stackrel{(a)}{\leq} (m-1 + 2^{p-1})\mu^p + 2^{p-1}\rho^p \\ &\stackrel{(b)}{=} (m-1 + 2^{p-1})\mu^p + 2^{p-1}\mu^{p-p^2} p^{-p^2} (p-1)^{p^2-p}. \end{aligned}$$

Here (a) follows from Jensen's inequality, as x^p is a convex function for $x \geq 0$, $p > 1$, and (b) holds by our choice of ρ in Eq. (11). We now choose μ to minimize the right-hand side upper bound for β and find the minimum to be realized at $\mu = 2^{\frac{1}{p}} (2(m-1) + 2^p)^{-\frac{1}{p^2}} p^{-1} (p-1)^{\frac{1}{p} + (1-\frac{1}{p})^2}$. Using this μ in Eq. (11), we find the desired ρ and we get the desired upper bound for β by taking the p th root of

$$\begin{aligned} \beta^p &\leq (2(m-1) + 2^p)^{1-\frac{1}{p}} p^{-p} (p-1)^{1+p(1-\frac{1}{p})^2} \\ &\quad + (2(m-1) + 2^p)^{1-\frac{1}{p}} p^{p^2-p} (p-1)^{1-p+(p-p^2)(1-\frac{1}{p})^2} p^{-p^2} (p-1)^{p^2-p} \\ &= (2(m-1) + 2^p)^{1-\frac{1}{p}} p^{-p} \left((p-1)^{1+p(1-\frac{1}{p})^2} + (p-1)^{1-2p+(p-p^2)(1-\frac{1}{p})^2+p^2} \right) \\ &= (2(m-1) + 2^p)^{1-\frac{1}{p}} p^{-p} \left((p-1)^{p-1+\frac{1}{p}} + (p-1)^{p-2+\frac{1}{p}} \right) \\ &= (2(m-1) + 2^p)^{1-\frac{1}{p}} p^{-p} (p-1)^{p-2+\frac{1}{p}} (1 + (p-1)) \\ &= (2(m-1) + 2^p)^{1-\frac{1}{p}} p^{1-p} (p-1)^{p-2+\frac{1}{p}}. \end{aligned}$$

□

Appendix J Proof of Proposition 10

Proof We first show that our choice of μ , ρ yields a valid dominating set using Condition (7) and Lemma 4. Using the properties $\boldsymbol{\xi}^\top \boldsymbol{\Sigma} \boldsymbol{\xi} \leq 1$ and $\boldsymbol{\Sigma} = \mathbf{1} + a(\mathbf{J} - \mathbf{1})$ of the ellipsoid uncertainty set, we find $(1-a)j\gamma(j)^2 + aj^2\gamma(j)^2 \leq 1$ and thus

$$\gamma(j) = \frac{1}{\sqrt{aj^2 + (1-a)j}}.$$

Inserting $\gamma(j)$ in Condition (7) we have to choose ρ such that

$$\rho \geq \max_{j \in \{1, \dots, m\}} j \left(\frac{1}{\sqrt{aj^2 + (1-a)j}} - \mu \right).$$

For the right-hand side term we find

$$j \left(\frac{1}{\sqrt{aj^2 + (1-a)j}} - \mu \right) \stackrel{(a)}{\leq} j \left(\frac{1}{\sqrt{aj^2}} \right) = \frac{1}{\sqrt{a}}$$

and

$$j \left(\frac{1}{\sqrt{aj^2 + (1-a)j}} - \mu \right) \leq j \left(\frac{1}{\sqrt{(1-a)j}} - \mu \right) \stackrel{(b)}{\leq} \frac{1}{4(1-a)\mu},$$

where for (a) we use $\mu \geq 0, a \leq 1$ and (b) follows as the term is maximized for $j = \frac{1}{4(1-a)\mu^2}$. Thus choosing $\rho = \frac{1}{\sqrt{a}}$ or $\rho = \frac{1}{4(1-a)\mu}$ always yields a valid dominating set. Using the definition of \mathcal{U} and v_i we find $\frac{1}{\beta} v_i \in \mathcal{U}$ to be fulfilled when

$$\beta^2 = \rho^2 + 2(1-a+am)\rho\mu + (am^2 + (1-a)m)\mu^2.$$

For $\mu = 0, \rho = \frac{1}{\sqrt{a}}$ this results in $\beta = \rho = \frac{1}{\sqrt{a}}$. Fixing $\rho = \frac{1}{4(1-a)\mu}$ yields

$$\beta^2 = \frac{1}{16(1-a)^2\mu^2} + \frac{1}{2} \left(1 + \frac{a}{1-a} m \right) + (am^2 + (1-a)m)\mu^2,$$

which is minimized for $\mu = \frac{1}{2\sqrt{4(1-a)^3m+(1-a)^2am^2}}$ resulting in

$$\beta^2 = \frac{1}{2} \left(1 + \frac{1}{1-a} \left(am + \sqrt{(1-a)m + am^2} \right) \right).$$

Finally, we close the proof by showing the desired asymptotic approximation bounds. The case $a > m^{-\frac{2}{3}}$ directly follows from $\beta = \frac{1}{\sqrt{a}} \leq \frac{1}{\sqrt{m^{-\frac{2}{3}}}} = O(m^{\frac{1}{3}})$ and for

$a \leq m^{-\frac{2}{3}}$ we find

$$\begin{aligned} \beta &= \sqrt{\frac{1}{2} \left(1 + \frac{1}{1-a} \left(am + \sqrt{(1-a)m + am^2} \right) \right)} \\ &\stackrel{(a)}{\leq} \sqrt{\frac{1}{2} \left(1 + \frac{1}{1-m^{-\frac{2}{3}}} \left(m^{\frac{1}{3}} + \sqrt{m + m^{\frac{4}{3}}} \right) \right)} \end{aligned}$$

$$\stackrel{(b)}{=} O(m^{\frac{1}{3}}),$$

where for (a) we use $a \leq m^{-\frac{2}{3}}$ and for (b) we use $m \geq 2$. \square

Appendix K Proof of Proposition 11

We prove Proposition 11 by explicitly constructing a dominating set fulfilling the desired properties. To find a good choice of $\rho_1, \dots, \rho_m, \mathbf{v}_0$ for the construction of $\hat{\mathcal{U}}$ in (2), we use the iterative approach described in Algorithm 1. Intuitively, we increase the base vertex \mathbf{v}_0 , as well as an upper bound for ρ_i , in each iteration. We bound the approximation factor β , by bounding the maximal number of iterations in the algorithm. Algorithm 1 is a refinement of Algorithm 1 in Ben-Tal et al. [14] and uses a different updating step. This modified updating step leads to a less aggressive increase of the base vertex \mathbf{v}_0 ultimately improving the approximation bound by a factor of $\frac{1}{2}$.

Algorithm 1 Generating \mathbf{v}_0 for general uncertainty sets \mathcal{U}

```

1:  $j = 0$ 
2:  $\mathbf{v}^0 = \mathbf{0}$ 
3: while  $\max_{\xi \in \mathcal{U}} \mathbf{e}^\top (\xi - \mathbf{v}^j)_+ > j + 1$  do
4:    $\xi^j \in \operatorname{argmax}_{\xi \in \mathcal{U}} \mathbf{e}^\top (\xi - \mathbf{v}^j)_+$ 
5:    $\mathbf{v}^{j+1} = \mathbf{v}^j + (\xi^j - \mathbf{v}^j)_+$ 
6:    $j = j + 1$ 
7: end while
8: return  $\beta = 2j + 1, \mathbf{v}_0 = \mathbf{v}^j$ 

```

Lemma 19 *Let β, \mathbf{v}_0 be the output of Algorithm 1. Then a solution for $Z_{AR}(\hat{\mathcal{U}})$, where $\hat{\mathcal{U}}$ is constructed according to (2) with \mathbf{v}_0 and $\rho_i = \frac{\beta+1}{2}$ is a β approximation for $Z_{AR}(\mathcal{U})$ and $\beta \leq 2\sqrt{m} + 1$.*

Proof In this proof let $J := \frac{\beta-1}{2}$ be the index j at termination of Algorithm 1. First, we show that $\hat{\mathcal{U}}$ is a valid dominating set for \mathcal{U} using Criterion (5). Let $\xi \in \mathcal{U}$. Then

$$\sum_{i=1}^m \frac{((\xi - \mathbf{v}_0)_+)_i}{\rho_i} \stackrel{(a)}{=} \frac{1}{J+1} \mathbf{e}^\top (\xi - \mathbf{v}_0)_+ \stackrel{(b)}{\leq} 1,$$

where (a) follows from our choice of ρ_i and (b) from the termination criterion of Algorithm 1. Next, we show $\frac{1}{\beta} \mathbf{v}_i \in \mathcal{U}$ for all $i \in \{1, \dots, m\}$ by

$$\frac{1}{\beta} \mathbf{v}_i = \frac{J+1}{\beta} \mathbf{e}_i + \frac{1}{\beta} \sum_{j=0}^{J-1} (\xi^j - \mathbf{v}^j)_+ \leq \frac{1}{\beta} \left((J+1) \mathbf{e}_i + \sum_{j=0}^{J-1} \xi^j \right) \stackrel{(a)}{\in} \mathcal{U}.$$

Here, the containment (a) follows from down-monotonicity, convexity and $\mathbf{e}_i \in \mathcal{U}$. Finally, we show that Algorithm 1 terminates after at most $\frac{\beta-1}{2} = J \leq \sqrt{m}$ iterations,

which implies $\beta \leq 2\sqrt{m} + 1$. Let $J' := J - 1$ be the index of the last iteration that fulfills the criterion of the loop. Then

$$\begin{aligned} J^2 &= \sum_{j=0}^{J'} J \stackrel{(a)}{\leq} \sum_{j=0}^{J'} e^\top (\xi^{J'} - v^{J'})_+ \stackrel{(b)}{\leq} \sum_{j=0}^{J'} e^\top (\xi^{J'} - v^j)_+ \\ &\stackrel{(c)}{\leq} \sum_{j=0}^{J'} \max_{\xi \in \mathcal{U}} e^\top (\xi - v^j)_+ \stackrel{(d)}{=} \sum_{j=0}^{J'} e^\top (\xi^j - v^j)_+ \\ &= \sum_{i=1}^m \sum_{j=0}^{J'} (\xi_i^j - v_i^j)_+ \stackrel{(e)}{=} \sum_{i=1}^m v_i^{J'} \stackrel{(f)}{=} \sum_{i=1}^m \max_{j \in \{0, \dots, J'\}} \xi_i^j \stackrel{(g)}{\leq} \sum_{i=1}^m 1 = m. \end{aligned}$$

Here, inequality (a) follows from the termination criterion of Algorithm 1, (b) follows as v^j only increases through the algorithm and thus $\forall j \leq J': v^{J'} \geq v^j$, (c) follows as by $\xi^{J'} \in \mathcal{U}$ the maximum over \mathcal{U} is an upper bound, (d) follows from the choice of ξ^j in Algorithm 1, (e) follows from the construction of v^j , (f) follows as by induction $v_i^{j+1} = \max(v_i^j, \xi_i^j) = \max_{j' \leq j} \xi_i^{j'}$, and (g) follows by $e^j \in \mathcal{U} \subseteq [0, 1]^m$. \square

Appendix L Proof of Proposition 12

Proof For each t , let m_t be the dimension \mathcal{U}_t . Then $m = \sum_{t=1}^T m_t$, and an uncertainty vector $\xi \in \mathcal{U}$ can be expressed as $\xi = (\xi^1, \dots, \xi^T) = (\xi_1^1, \dots, \xi_{m_1}^1, \dots, \xi_{m_T}^T)$, where $\xi^t \in \mathcal{U}_t$. For each t , let v_0^t be the base vertex and let $\rho_1^t, \dots, \rho_{m_t}^t$ be the parameters inducing $\hat{\mathcal{U}}_t$ via Construction (2). Let $\beta = \sum_{t \in \mathcal{T}_1} \beta_t$, and construct the dominating set $\hat{\mathcal{U}}$ for \mathcal{U} with base vertex $v_0' := (v_0^1, \dots, v_0^T)$, where

$$v_0'' := \begin{cases} v_0^t & \text{if } t \in \mathcal{T}_1 \\ e & \text{if } t \in \mathcal{T}_2 \end{cases}$$

and parameters

$$\rho_i'' := \begin{cases} \frac{\beta}{\beta_t} \rho_i^t & \text{if } t \in \mathcal{T}_1 \\ 0 & \text{if } t \in \mathcal{T}_2 \end{cases}$$

using Construction (2). This construction corresponds to dominating each \mathcal{U}_t with $t \in \mathcal{T}_2$ by the unit vector e . Here $\{e\}$ is dominating \mathcal{U}_t by Assumption 1. We dominate the remaining \mathcal{U}_t for $t \in \mathcal{T}_1$ with a combined polytope. Then the maximal sum of convex factors for the vertices is given by

$$\max_{\xi \in \mathcal{U}} \sum_{t \in \mathcal{T}_1} \sum_{i=1}^{m_t} \frac{((\xi - v_0')_+)_i}{\rho_i''} \stackrel{(a)}{=} \frac{1}{\beta} \max_{\xi \in \mathcal{U}} \sum_{t \in \mathcal{T}_1} \beta_t \sum_{i=1}^{m_t} \frac{((\xi^t - v_0^t)_+)_i}{\rho_i^t}$$

$$\stackrel{(b)}{=} \frac{1}{\beta} \sum_{t \in \mathcal{T}_1} \beta_t \max_{\xi^t \in \mathcal{U}_t} \sum_{i=1}^{m_t} \frac{((\xi^t - \mathbf{v}_0^t)_+)_i}{\rho_i^t} \stackrel{(c)}{\leq} \frac{1}{\beta} \sum_{t \in \mathcal{T}_1} \beta_t \stackrel{(d)}{=} 1,$$

where (a) follows from the definition of \mathbf{v}'_0 and ρ_i^t , (b) follows from \mathcal{U} being stagewise independent, (c) follows from Condition (5) and $\hat{\mathcal{U}}_t$ being a dominating set for \mathcal{U}_t , and (d) follows from the definition of β . Thus $\hat{\mathcal{U}}$ fulfills Condition (5) and is a valid dominating set for \mathcal{U} .

Let $\beta' := \max_{t \in \mathcal{T}_2} \beta'_t$. To see that $\max(\beta, \beta')$ is indeed an upper bound for the approximation factor, we find

$$\begin{aligned} \frac{1}{\max(\beta, \beta')} \mathbf{v}_0^t &\stackrel{(a)}{\leq} \frac{1}{\beta} (\mathbf{v}_0^t + \rho_i^t \mathbf{e}_i) \stackrel{(b)}{=} \frac{1}{\beta} (\mathbf{v}_0^t + \frac{\beta}{\beta_t} \rho_i^t \mathbf{e}_i) \stackrel{(c)}{\leq} \frac{1}{\beta_t} (\mathbf{v}_0^t + \rho_i^t \mathbf{e}_i) \stackrel{(d)}{\in} \mathcal{U}_t \quad \forall t \in \mathcal{T}_1 \\ \frac{1}{\max(\beta, \beta')} \mathbf{v}_0^t &\stackrel{(e)}{\leq} \frac{1}{\beta'} \mathbf{e} \stackrel{(f)}{\leq} \frac{1}{\beta'_t} \mathbf{e} \stackrel{(g)}{\in} \mathcal{U}_t \quad \forall t \in \mathcal{T}_2, \end{aligned}$$

where (a) follows from $\beta \leq \max(\beta, \beta')$ and $\rho_i^t \geq 0$, (b) follows from the definitions of \mathbf{v}_0^t and ρ_i^t , (c) follows from $\beta_t \leq \beta$ for all $t \in \mathcal{T}_1$, (d) follows as β_t is a valid approximation factor for $\hat{\mathcal{U}}_t$, (e) follows from $\beta' \leq \max(\beta, \beta')$ and the definition of \mathbf{v}_0^t , (f) follows from $\beta'_t \leq \beta'$ for all $t \in \mathcal{T}_2$, and (g) follows from the definition of β'_t . Using the stagewise structure and down monotonicity of \mathcal{U} this implies

$$\begin{aligned} \frac{1}{\max(\beta, \beta')} \mathbf{v}'_0 &= \frac{1}{\max(\beta, \beta')} (\mathbf{v}_0^1, \dots, \mathbf{v}_0^T) \in \mathcal{U}, \\ \frac{1}{\max(\beta, \beta')} (\mathbf{v}_0^t + \rho_i^t \mathbf{e}_i) &= \frac{1}{\max(\beta, \beta')} (\mathbf{v}_0^1, \dots, \mathbf{v}_0^t + \rho_i^t \mathbf{e}_i, \dots, \mathbf{v}_0^T) \in \mathcal{U}. \end{aligned}$$

□

Appendix M Proof of Lemma 13

Proof Let $h : \mathcal{U} \rightarrow \hat{\mathcal{U}}$ be the domination function of $\hat{\mathcal{U}}$. Then define $h' : \mathcal{U} \rightarrow \hat{\mathcal{U}}$ by

$$h'_j(\xi) := h_j(\xi) + s_j(1 - h_j(\xi)). \tag{M1}$$

This new dominating function h' is nonanticipative, as h is nonanticipative and the j th component of h' only depends on the j th component of h . To see that (M1) is also a domination function, consider the two cases $h_j(\xi) \geq 1$ and $h_j(\xi) \leq 1$.

1. For $h_j(\xi) \geq 1$ the second summand on the right hand side of (M1) is negative for all $s_j \in [0, 1]$. Thus $h'_j(\xi)$ is minimal for $s_j = 1$ and we find $h'_j(\xi) \geq h_j(\xi) + 1(1 - h_j(\xi)) = 1$.
2. On the other hand, for $h'_j(\xi) \leq 1$ the second summand is positive and $h'_j(\xi)$ is minimal for $s_j = 0$. Thus, we have $h'_j(\xi) \geq h_j(\xi)$.

In any case, we find $\mathbf{h}'(\xi) \geq \xi$ by $h'_j \geq \min(1, h_j(\xi)) \geq \xi_j$ where we used $\mathcal{U} \subseteq [0, 1]^m$ from Assumption 1.

Next, we show $\mathbf{h}'(\xi) \in \hat{\mathcal{U}}'$. For this let $\mathbf{h}(\xi) = \sum_{i=0}^m \lambda_i \mathbf{v}_i$ be the convex combination of $\mathbf{h}(\xi) \in \hat{\mathcal{U}}$. Then,

$$h'_j(\xi) = h_j(\xi) + s_j(1 - h_j(\xi)) \stackrel{(a)}{=} \sum_{i=0}^m (\lambda_i v_{ij} + s_j(1 - v_{ij})) = \sum_{i=0}^m \lambda_i v'_{ij}.$$

Note that in (a) we used $\sum_{i=0}^m \lambda_i = 1$. This shows $\mathbf{h}'(\xi) = \sum_{i=0}^m \lambda_i \mathbf{v}'_i$ and thus $\mathbf{h}'(\xi) \in \hat{\mathcal{U}}'$. □

Appendix N Proof of Proposition 14

Proof As $Z_{LIFT} \leq Z_{AFF}$ was already shown by Georghiou et al. [31], we are left to show $Z_{LIFT} \geq Z_{AFF}$. We do so by showing that for any optimal affine solution for $Z_{AR}^L(\mathcal{U}^L)$, there is an affine solution for $Z_{AR}(\mathcal{U})$ with the same objective value. First, we define the average lifting operator $\bar{L}: \mathcal{U} \rightarrow \mathcal{U}^L$ via

$$\bar{L}_{ij}(\xi) := (z_j^i - z_{j-1}^i)\xi_i.$$

Note, that $R_i(\bar{L}(\xi)) = \sum_{j=1}^r (z_j^i - z_{j-1}^i)\xi_i = (z_r^i - z_0^i)\xi_i = \xi_i$, which implies $R(\bar{L}(\xi)) = \xi$. Thus, $\bar{L}(\xi)$ fulfills the first condition of (8). Next, we verify the second condition by

$$\frac{\bar{L}_{i,j+1}(\xi)}{z_{j+1}^i - z_j^i} = \frac{(z_{j+1}^i - z_j^i)\xi_i}{z_{j+1}^i - z_j^i} = \xi_i = \frac{(z_j^i - z_{j-1}^i)\xi_i}{z_j^i - z_{j-1}^i} = \frac{\bar{L}_{ij}(\xi)}{z_j^i - z_{j-1}^i}.$$

Finally, the third condition holds by $\bar{L}_{i1} = z_1^i \xi_i \leq z_1^i$ which follows from $\xi_i \leq 1$ in Assumption 1. Thus $\bar{L}(\xi) \in \mathcal{U}^L$. Let \mathbf{x}^L be an optimal affine solution for $Z_{AR}^L(\mathcal{U}^L)$. As \bar{L} is a nonanticipative affine map, the concatenation $\mathbf{x}^L \circ \bar{L}$ is also nonanticipative and affine. By $\bar{L}(\mathcal{U}) \subseteq \mathcal{U}^L$ the decision $\mathbf{x} := \mathbf{x}^L \circ \bar{L}$ is a feasible solution for $Z_{AR}(\mathcal{U})$. Finally, we find

$$Z_{AFF} \leq \max_{\xi \in \mathcal{U}} \mathbf{c}^\top \mathbf{x}(\xi) = \max_{\xi^L \in \bar{L}(\mathcal{U})} \mathbf{c}^\top \mathbf{x}^L(\xi^L) \leq \max_{\xi^L \in \mathcal{U}^L} \mathbf{c}^\top \mathbf{x}^L(\xi^L) = Z_{LIFT}.$$

□

Appendix O Proof of Proposition 15

Proof To prove $Z_{TLIFT} \leq Z_{LIFT}$ and $Z_{TLIFT} \leq Z_{SPAP}$ we show that any affine solution for $Z_{AR}^L(\mathcal{U}^L)$ and any piecewise affine solution with re-scaling for $Z(\hat{\mathcal{U}})$

can be transformed to a feasible affine solution for $Z_{AR}^L(\hat{\mathcal{U}}^L)$ with at most the same objective value.

For $Z_{TLIFT} \leq Z_{LIFT}$, let $\mathbf{x}^L: \mathcal{U}^L \rightarrow \mathbb{R}^n$ be an optimal affine solution to $Z_{AR}^L(\mathcal{U}^L)$. Then \mathbf{x}^L is also a feasible affine solution for $Z_{AR}^L(\hat{\mathcal{U}}^L)$ by $\hat{\mathcal{U}}^L \subseteq \mathcal{U}^L$. For the objective value, we find

$$Z_{TLIFT} \leq \max_{\xi^L \in \hat{\mathcal{U}}^L} \mathbf{c}^\top \mathbf{x}^L(\xi^L) \stackrel{(a)}{\leq} \max_{\xi^L \in \mathcal{U}^L} \mathbf{c}^\top \mathbf{x}^L(\xi^L) = Z_{LIFT},$$

where (a) also follows from $\hat{\mathcal{U}}^L \subseteq \mathcal{U}^L$.

For $Z_{TLIFT} \leq Z_{SPAP}$ let $\hat{\mathbf{x}}(\hat{\xi})$ given by vertex solutions $\mathbf{x}_0, \dots, \mathbf{x}_m$ and scaling factor s be an optimal solution to $Z_{LP}(\hat{\xi})$ with re-scaling as described in Sect. 4. We define the map $\mathbf{h}_s^L: \hat{\mathcal{U}}^L \rightarrow \hat{\mathcal{U}}'$ from the tightened lifted uncertainty set $\hat{\mathcal{U}}^L$ defined in Eq. (10) to the re-scaled dominating uncertainty set $\hat{\mathcal{U}}'$ defined in Lemma 13 by

$$h_{si}^L(\hat{\xi}^L) := (1 - s_i)(v_{0i} + \hat{\xi}_{i2}^L) + s_i.$$

This mapping is nonanticipative. Furthermore, $\mathbf{h}_s^L(\hat{\xi}^L)$ dominates $R(\hat{\xi}^L)$ by

$$R(\hat{\xi}^L)_i = \hat{\xi}_{i1}^L + \hat{\xi}_{i2}^L \stackrel{(a)}{\leq} (1 - s_i)(\hat{\xi}_{i1}^L + \hat{\xi}_{i2}^L) + s_i \stackrel{(b)}{\leq} (1 - s_i)(v_{0i} + \hat{\xi}_{i2}^L) + s_i = h_{si}^L(\hat{\xi}^L).$$

Here, (a) follows as $\hat{\xi}_{i1}^L + \hat{\xi}_{i2}^L \leq 1$ by Assumption 1, and (b) follows from $\hat{\xi}_{i1}^L \leq v_{0i}$ as v_{0i} is the break-point. Using the definition of the re-scaled vertices \mathbf{v}'_i from Lemma 13, we find

$$\mathbf{h}_s^L(\hat{\xi}^L) = \sum_{i=1}^m \frac{\hat{\xi}_{i2}^L}{\rho_i} \mathbf{v}'_i + \left(1 - \sum_{i=1}^m \frac{\hat{\xi}_{i2}^L}{\rho_i} \right) \mathbf{v}'_0. \tag{O1}$$

By the tightening constraint in Definition (10) this is a valid convex combination for all $\hat{\xi}^L \in \hat{\mathcal{U}}^L$ and thus $\mathbf{h}_s^L(\hat{\mathcal{U}}^L) \subseteq \hat{\mathcal{U}}'$ as claimed. Thus $\mathbf{x}^L := \hat{\mathbf{x}} \circ \mathbf{h}_s^L$ is a valid solution for $Z_{AR}^L(\hat{\mathcal{U}}^L)$ and

$$\max_{\xi^L \in \hat{\mathcal{U}}^L} \mathbf{c}^\top \mathbf{x}^L(\xi^L) = \max_{\hat{\xi} \in \mathbf{h}_s^L(\hat{\mathcal{U}}^L)} \mathbf{c}^\top \hat{\mathbf{x}}(\hat{\xi}) \leq \max_{\hat{\xi} \in \hat{\mathcal{U}}'} \mathbf{c}^\top \hat{\mathbf{x}}(\hat{\xi}) = Z_{SPAP}.$$

Using that $\hat{\mathbf{x}}$ is given by the vertex solutions $\mathbf{x}_0, \dots, \mathbf{x}_m$ together with (O1) we find

$$\mathbf{x}^L(\hat{\xi}^L) = \hat{\mathbf{x}} \circ \mathbf{h}_s^L(\hat{\xi}^L) = \sum_{i=1}^m \frac{\hat{\xi}_{i2}^L}{\rho_i} \mathbf{x}_i + \left(1 - \sum_{i=1}^m \frac{\hat{\xi}_{i2}^L}{\rho_i} \right) \mathbf{x}_0 = \sum_{i=1}^m \left(\frac{\hat{\xi}_{i2}^L}{\rho_i} (\mathbf{x}_i - \mathbf{x}_0) \right) + \mathbf{x}_0,$$




Fig. 10 Acceptance rate of our proposed sampling method for budgeted uncertainty with budget $k = \sqrt{m}$

which is affine in $\hat{\mathcal{U}}^L$ and thus

$$Z_{TLIFT} \leq \max_{\hat{\xi}^L \in \hat{\mathcal{U}}^L} c^\top x^L(\hat{\xi}^L).$$

□

Appendix P Sampling Uniform Uncertainty Realizations

In this section, we describe how to generate uniform samples from the uncertainty sets used in our numerical experiments efficiently.

Hypersphere uncertainty: To sample uniformly from the hypersphere uncertainty set $\mathcal{U} = \{\xi \in \mathbb{R}_+^m \mid \|\xi\|_2 \leq 1\}$, we first draw i.i.d. samples ξ^B from the m -dimensional unit ball. It is commonly known, that this can efficiently be done by drawing m i.i.d. normal variables that are normalized and rescaled by the radius appropriately, see, e.g. [29]. We now map the samples from the unit ball to the positive orthant by taking the absolute value in each component, i.e., we get $\xi \in \mathcal{U}$ by $\xi_i := |\xi_i^B|$. By rotational invariance of the unit ball, this yields a uniform sample on the hypersphere uncertainty set.

Budgeted uncertainty: For the budgeted uncertainty set $\mathcal{U} = \{\xi \in [0, 1]^m \mid \|\xi\|_1 \leq k\}$ with budget $k = \sqrt{m}$, we first draw i.i.d. samples ξ^S from the m -dimensional unit simplex $\text{conv}(\mathbf{0}, e_1, \dots, e_m)$. This can be done efficiently by taking the first $m + 1$ i.i.d. exponentially distributed variables that were normalized with respect to the one norm, see, e.g. [29]. By scaling the samples from the unit simplex with the budget k , and rejecting all samples that do not fulfill $\xi := k\xi^S \in [0, 1]^m$, we get a uniform sample on the budgeted uncertainty set. In Fig. 10, we show the acceptance rate of the samples, which shows that the proposed rejection method is efficient for our application.

References

1. Adida, E., Perakis, G.: A robust optimization approach to dynamic pricing and inventory control with no backorders. *Math. Program.* **107**(1–2), 97–129 (2006). <https://doi.org/10.1007/s10107-005-0681-5>
2. Ang, M., Lim, Y.F., Sim, M.: Robust storage assignment in unit-load warehouses. *Manag. Sci.* **58**(11), 2114–2130 (2012). <https://doi.org/10.1287/mnsc.1120.1543>
3. Ayoub, J., Poss, M.: Decomposition for adjustable robust linear optimization subject to uncertainty polytope. *Comput. Manag. Sci.* **13**(2), 219–239 (2016). <https://doi.org/10.1007/s10287-016-0249-2>
4. Bampou, D., Kuhn, D.: Scenario-free stochastic programming with polynomial decision rules. In: *IEEE Conference on Decision and Control and European Control Conference*. IEEE, Orlando (2011). <https://doi.org/10.1109/cdc.2011.6161150>
5. Baron, O., Milner, J., Naseraldin, H.: Facility location: a robust optimization approach. *Prod. Oper. Manag.* **20**(5), 772–785 (2011). <https://doi.org/10.1111/j.1937-5956.2010.01194.x>
6. Beebe, J.H., Beightler, C.S., Stark, J.P.: Stochastic optimization of production planning. *Oper. Res.* **16**(4), 799–818 (1968). <https://doi.org/10.1287/opre.16.4.799>
7. Ben-Tal, A., Nemirovski, A.: Robust solutions of uncertain linear programs. *Oper. Res. Lett.* **25**(1), 1–13 (1999). [https://doi.org/10.1016/s0167-6377\(99\)00016-4](https://doi.org/10.1016/s0167-6377(99)00016-4)
8. Ben-Tal, A., Nemirovski, A.: Robust optimization—methodology and applications. *Math. Program.* **92**(3), 453–480 (2002). <https://doi.org/10.1007/s101070100286>
9. Ben-Tal, A., Goryashko, A., Guslitzer, E., et al.: Adjustable robust solutions of uncertain linear programs. *Math. Program.* **99**(2), 351–376 (2004). <https://doi.org/10.1007/s10107-003-0454-y>
10. Ben-Tal, A., Golany, B., Nemirovski, A., et al.: Retailer-supplier flexible commitments contracts: a robust optimization approach. *Manuf. Serv. Oper. Manag.* **7**(3), 248–271 (2005). <https://doi.org/10.1287/msom.1050.0081>
11. Ben-Tal, A., Ghaoui, L.E., Nemirovski, A.: *Robust Optimization*. Princeton University Press, Princeton (2009)
12. Ben-Tal, A., Golany, B., Shtern, S.: Robust multi-echelon multi-period inventory control. *Eur. J. Oper. Res.* **199**(3), 922–935 (2009). <https://doi.org/10.1016/j.ejor.2009.01.058>
13. Ben-Tal, A., Chung, B.D., Mandala, S.R., et al.: Robust optimization for emergency logistics planning: risk mitigation in humanitarian relief supply chains. *Transp. Res. Part B Methodol.* **45**(8), 1177–1189 (2011). <https://doi.org/10.1016/j.trb.2010.09.002>
14. Ben-Tal, A., Housni, O.E., Goyal, V.: A tractable approach for designing piecewise affine policies in two-stage adjustable robust optimization. *Math. Program.* **182**(1–2), 57–102 (2020). <https://doi.org/10.1007/s10107-019-01385-0>
15. Bertsimas, D., Bidkhorji, H.: On the performance of affine policies for two-stage adaptive optimization: a geometric perspective. *Math. Program.* **153**(2), 577–594 (2015). <https://doi.org/10.1007/s10107-014-0818-5>
16. Bertsimas, D., Dunning, I.: Multistage robust mixed-integer optimization with adaptive partitions. *Oper. Res.* **64**(4), 980–998 (2016). <https://doi.org/10.1287/opre.2016.1515>
17. Bertsimas, D., Georghiou, A.: Design of near optimal decision rules in multistage adaptive mixed-integer optimization. *Oper. Res.* **63**(3), 610–627 (2015). <https://doi.org/10.1287/opre.2015.1365>
18. Bertsimas, D., Goyal, V.: On the power and limitations of affine policies in two-stage adaptive optimization. *Math. Program.* **134**(2), 491–531 (2012). <https://doi.org/10.1007/s10107-011-0444-4>
19. Bertsimas, D., Iancu, D.A., Parrilo, P.A.: Optimality of affine policies in multistage robust optimization. *Math. Oper. Res.* **35**(2), 363–394 (2010). <https://doi.org/10.1287/moor.1100.0444>
20. Bertsimas, D., Goyal, V., Sun, X.A.: A geometric characterization of the power of finite adaptability in multistage stochastic and adaptive optimization. *Math. Oper. Res.* **36**(1), 24–54 (2011). <https://doi.org/10.1287/moor.1110.0482>
21. Bertsimas, D., Iancu, D.A., Parrilo, P.A.: A hierarchy of near-optimal policies for multistage adaptive optimization. *IEEE Trans. Autom. Control* **56**(12), 2809–2824 (2011). <https://doi.org/10.1109/tac.2011.2162878>
22. Bertsimas, D., Goyal, V., Lu, B.Y.: A tight characterization of the performance of static solutions in two-stage adjustable robust linear optimization. *Math. Program.* **150**(2), 281–319 (2015). <https://doi.org/10.1007/s10107-014-0768-y>
23. Bertsimas, D., Shtern, S., Sturt, B.: Technical note—two-stage sample robust optimization. *Oper. Res.* (2021). <https://doi.org/10.1287/opre.2020.2096>

24. Chen, X., Zhang, Y.: Uncertain linear programs: extended affinely adjustable robust counterparts. *Oper. Res.* **57**(6), 1469–1482 (2009). <https://doi.org/10.1287/opre.1080.0605>
25. Chen, X., Sim, M., Sun, P., et al.: A linear decision-based approximation approach to stochastic programming. *Oper. Res.* **56**(2), 344–357 (2008). <https://doi.org/10.1287/opre.1070.0457>
26. Dantzig, G.B.: Linear programming under uncertainty. *Manag. Sci.* **1**(3–4), 197–206 (1955). <https://doi.org/10.1287/mnsc.1.3.4.197>
27. Daryalal, M., Bodur, M., Luedtke, J.R.: Lagrangian dual decision rules for multistage stochastic mixed-integer programming. *Oper. Res.* (2022). <https://doi.org/10.1287/opre.2022.2366>
28. Delage, E., Iancu, D.A.: Robust multistage decision making. In: *The Operations Research Revolution*. INFORMS, Catonsville, pp. 20–46 (2015). <https://doi.org/10.1287/educ.2015.0139>
29. Devroye, L.: *Non-Uniform Random Variate Generation*. Springer, New York (1986). <https://doi.org/10.1007/978-1-4613-8643-8>
30. Feige, U., Jain, K., Mahdian, M., et al.: Robust combinatorial optimization with exponential scenarios. In: *Integer Programming and Combinatorial Optimization*, pp. 439–453. Springer, Berlin (2007). <https://doi.org/10.1007/978-3-540-72792-7-33>
31. Georghiou, A., Wiesemann, W., Kuhn, D.: Generalized decision rule approximations for stochastic programming via liftings. *Math. Program.* **152**(1–2), 301–338 (2015). <https://doi.org/10.1007/s10107-014-0789-6>
32. Georghiou, A., Tsoukalas, A., Wiesemann, W.: Robust dual dynamic programming. *Oper. Res.* **67**(3), 813–830 (2019). <https://doi.org/10.1287/opre.2018.1835>
33. Georghiou, A., Tsoukalas, A., Wiesemann, W.: A primal–dual lifting scheme for two-stage robust optimization. *Oper. Res.* (2020). <https://doi.org/10.1287/opre.2019.1873>
34. Ghaoui, L.E., Oks, M., Oustry, F.: Worst-case value-at-risk and robust portfolio optimization: a conic programming approach. *Oper. Res.* **51**(4), 543–556 (2003). <https://doi.org/10.1287/opre.51.4.543.16101>
35. Goldfarb, D., Iyengar, G.N.: Robust portfolio selection problems. *Math. Oper. Res.* **28**(1), 1–38 (2003). <https://doi.org/10.1287/moor.28.1.1.14260>
36. Gorissen, B.L., Yanikoglu, İ., den Hertog, D.: A practical guide to robust optimization. *Omega* **53**, 124–137 (2015). <https://doi.org/10.1016/j.omega.2014.12.006>
37. Haddad-Sisakht, A., Ryan, S.M.: Conditions under which adjustability lowers the cost of a robust linear program. *Ann. Oper. Res.* **269**(1–2), 185–204 (2018). <https://doi.org/10.1007/s10479-018-2954-4>
38. Hadjiyiannis, M.J., Goulart, P.J., Kuhn, D.: A scenario approach for estimating the suboptimality of linear decision rules in two-stage robust optimization. In: *IEEE Conference on Decision and Control and European Control Conference*. IEEE (2011) <https://doi.org/10.1109/cdc.2011.6161342>
39. Hanasusanto, G.A., Kuhn, D., Wiesemann, W.: K-adaptability in two-stage robust binary programming. *Oper. Res.* **63**(4), 877–891 (2015). <https://doi.org/10.1287/opre.2015.1392>
40. Housni, O.E., Goyal, V.: On the optimality of affine policies for budgeted uncertainty sets. *Math. Oper. Res.* **46**(2), 674–711 (2021). <https://doi.org/10.1287/moor.2020.1082>
41. Høyland, K., Wallace, S.W.: Generating scenario trees for multistage decision problems. *Manag. Sci.* **47**(2), 295–307 (2001). <https://doi.org/10.1287/mnsc.47.2.295.9834>
42. Iancu, D.A., Sharma, M., Sviridenko, M.: Supermodularity and affine policies in dynamic robust optimization. *Oper. Res.* **61**(4), 941–956 (2013). <https://doi.org/10.1287/opre.2013.1172>
43. Kuhn, D., Wiesemann, W., Georghiou, A.: Primal and dual linear decision rules in stochastic and robust optimization. *Math. Program.* **130**(1), 177–209 (2011). <https://doi.org/10.1007/s10107-009-0331-4>
44. Laguna, M.: Applying robust optimization to capacity expansion of one location in telecommunications with demand uncertainty. *Manag. Sci.* **44**(11–part–2), S101–S110 (1998). <https://doi.org/10.1287/mnsc.44.11.s101>
45. Lu, M., Shen, Z.J.M.: A review of robust operations management under model uncertainty. *Prod. Oper. Manag.* **30**(6), 1927–1943 (2021). <https://doi.org/10.1111/poms.13239>
46. Marandi, A., den Hertog, D.: When are static and adjustable robust optimization problems with constraint-wise uncertainty equivalent? *Math. Program.* **170**(2), 555–568 (2018). <https://doi.org/10.1007/s10107-017-1166-z>
47. Mudchanatongsuk, S., Ordóñez, F., Liu, J.: Robust solutions for network design under transportation cost and demand uncertainty. *J. Oper. Res. Soc.* **59**(5), 652–662 (2008). <https://doi.org/10.1057/palgrave.jors.2602362>
48. Mulvey, J.M., Vanderbei, R.J., Zenios, S.A.: Robust optimization of large-scale systems. *Oper. Res.* **43**(2), 264–281 (1995). <https://doi.org/10.1287/opre.43.2.264>

49. Noyan, N., Balci, B., Atakan, S.: A stochastic optimization model for designing last mile relief networks. *Transport. Sci.* **50**(3), 1092–1113 (2016). <https://doi.org/10.1287/trsc.2015.0621>
50. Ordóñez, F., Zhao, J.: Robust capacity expansion of network flows. *Networks* **50**(2), 136–145 (2007). <https://doi.org/10.1002/net.20183>
51. Paraskevopoulos, D., Karakitsos, E., Rustem, B.: Robust capacity planning under uncertainty. *Manag. Sci.* **37**(7), 787–800 (1991). <https://doi.org/10.1287/mnsc.37.7.787>
52. Porteus, E.L.: Stochastic inventory theory. In: *Handbooks in Operations Research and Management Science*, pp. 605–652. Elsevier, Amsterdam (1990). [https://doi.org/10.1016/s0927-0507\(05\)80176-8](https://doi.org/10.1016/s0927-0507(05)80176-8)
53. Postek, K., den Hertog, D.: Multistage adjustable robust mixed-integer optimization via iterative splitting of the uncertainty set. *INFORMS J. Comput.* **28**(3), 553–574 (2016). <https://doi.org/10.1287/ijoc.2016.0696>
54. Rahal, S., Papageorgiou, D.J., Li, Z.: Hybrid strategies using linear and piecewise-linear decision rules for multistage adaptive linear optimization. *Eur. J. Oper. Res.* **290**(3), 1014–1030 (2021). <https://doi.org/10.1016/j.ejor.2020.08.054>
55. Schmidt, R.L.: A stochastic optimization model to improve production planning and R&D resource allocation in biopharmaceutical production processes. *Manag. Sci.* **42**(4), 603–617 (1996). <https://doi.org/10.1287/mnsc.42.4.603>
56. Sedgewick, R., Flajolet, P.: *An Introduction to the Analysis of Algorithms*. Addison-Wesley, Boston (2013)
57. See, C.T., Sim, M.: Robust approximation to multiperiod inventory management. *Oper. Res.* **58**(3), 583–594 (2010). <https://doi.org/10.1287/opre.1090.0746>
58. Simchi-Levi, D., Trichakis, N., Zhang, P.Y.: Designing response supply chain against bioattacks. *Oper. Res.* **67**(5), 1246–1268 (2019). <https://doi.org/10.1287/opre.2019.1862>
59. Simchi-Levi, D., Wang, H., Wei, Y.: Constraint generation for two-stage robust network flow problems. *INFORMS J. Optim.* **1**(1), 49–70 (2019). <https://doi.org/10.1287/ijoo.2018.0003>
60. Solyali, O., Cordeau, J.F., Laporte, G.: The impact of modeling on robust inventory management under demand uncertainty. *Manag. Sci.* **62**(4), 1188–1201 (2016). <https://doi.org/10.1287/mnsc.2015.2183>
61. Wang, Z., Qi, M.: Robust service network design under demand uncertainty. *Transport. Sci.* **54**(3), 676–689 (2020). <https://doi.org/10.1287/trsc.2019.0935>
62. Xie, F., Huang, Y.: A multistage stochastic programming model for a multi-period strategic expansion of biofuel supply chain under evolving uncertainties. *Transport. Res. Part E Logist. Transport. Rev.* **111**, 130–148 (2018). <https://doi.org/10.1016/j.tre.2018.01.015>
63. Yanikoğlu, İ., Gorissen, B.L., den Hertog, D.: A survey of adjustable robust optimization. *Eur. J. Oper. Res.* **277**(3), 799–813 (2019). <https://doi.org/10.1016/j.ejor.2018.08.031>
64. Zou, J., Ahmed, S., Sun, X.A.: Partially adaptive stochastic optimization for electric power generation expansion planning. *INFORMS J. Comput.* **30**(2), 388–401 (2018). <https://doi.org/10.1287/ijoc.2017.0782>

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.