

K -Adaptability in Two-Stage Mixed-Integer Robust Optimization

Anirudh Subramanyam¹, Chrysanthos E. Gounaris¹, and Wolfram Wiesemann²

`asubramanyam@cmu.edu`, `gounaris@cmu.edu`, `ww@imperial.ac.uk`

¹Department of Chemical Engineering, Carnegie Mellon University, United States

²Imperial College Business School, Imperial College London, United Kingdom

July 27, 2018

Abstract

We study two-stage robust optimization problems with mixed discrete-continuous decisions in both stages. Despite their broad range of applications, these problems pose two fundamental challenges: *(i)* they constitute infinite-dimensional problems that require a finite-dimensional approximation, and *(ii)* the presence of discrete recourse decisions typically prohibits duality-based solution schemes. We address the first challenge by studying a K -adaptability formulation that selects K candidate recourse policies *before* observing the realization of the uncertain parameters and that implements the best of these policies *after* the realization is known. We address the second challenge through a branch-and-bound scheme that enjoys asymptotic convergence in general and finite convergence under specific conditions. We illustrate the performance of our algorithm in numerical experiments involving benchmark data from several application domains.

Keywords: Robust Optimization, Two-Stage Problems, K -Adaptability, Branch-and-Bound.

1 Introduction

Dynamic decision-making under uncertainty, where actions need to be taken both in anticipation of and in response to the realization of a priori uncertain problem parameters, arguably forms one of the most challenging domains of operations research and optimization theory. Despite intensive research efforts over the past six decades, many uncertainty-affected optimization problems resist solution, and even our understanding of the complexity of these problems remains incomplete.

In the last two decades, robust optimization has emerged as a promising methodology to counter some of the intricacies associated with decision-making under uncertainty. The rich theory on static robust optimization problems, in which all decisions have to be taken before the uncertainty is resolved, is summarized in [2, 4, 16]. However, dynamic robust optimization problems, in which some of the decisions can adapt to the observed uncertainties, are still poorly understood.

This paper is concerned with *two-stage robust optimization problems* of the form

$$\inf_{\mathbf{x} \in \mathcal{X}} \sup_{\boldsymbol{\xi} \in \Xi} \inf_{\mathbf{y} \in \mathcal{Y}} \left\{ \mathbf{c}^\top \mathbf{x} + \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y} : \mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}(\boldsymbol{\xi})\mathbf{y} \leq \mathbf{h}(\boldsymbol{\xi}) \right\}, \quad (1)$$

where $\mathcal{X} \subseteq \mathbb{R}^{N_1}$, $\mathcal{Y} \subseteq \mathbb{R}^{N_2}$ and $\Xi \subseteq \mathbb{R}^{N_p}$ constitute nonempty and bounded mixed-integer linear programming (MILP) representable sets, $\mathbf{c} \in \mathbb{R}^{N_1}$, and the functions $\mathbf{d} : \Xi \mapsto \mathbb{R}^{N_2}$, $\mathbf{T} : \Xi \mapsto \mathbb{R}^{L \times N_1}$, $\mathbf{W} : \Xi \mapsto \mathbb{R}^{L \times N_2}$ and $\mathbf{h} : \Xi \mapsto \mathbb{R}^L$ are affine. In problem (1), the vector \mathbf{x} represents the first-stage (or ‘here-and-now’) decisions which are taken before the value of the uncertain parameter vector $\boldsymbol{\xi}$ from within the uncertainty set Ξ is observed. The vector \mathbf{y} , on the other hand, denotes the second-stage (or ‘wait-and-see’) decisions that can adapt to the realized value of $\boldsymbol{\xi}$. We emphasize that problem (1) can have a random recourse, i.e., the recourse matrix \mathbf{W} may depend on the uncertain parameters $\boldsymbol{\xi}$. Moreover, we do not assume a relatively complete recourse; that is, for some first-stage decisions $\mathbf{x} \in \mathcal{X}$, there can be parameter realizations $\boldsymbol{\xi} \in \Xi$ such that there is no feasible second-stage decision \mathbf{y} . Also, we do not assume that the sets \mathcal{X} , \mathcal{Y} or Ξ are convex.

Remark 1 (Uncertain First-Stage Objective Coefficients). *The assumption that \mathbf{c} is deterministic does not restrict generality. Indeed, problem (1) accounts for uncertain first-stage objective coefficients $\mathbf{c}' : \Xi \mapsto \mathbb{R}^{N_1}$ if we augment the second-stage decisions \mathbf{y} to $(\mathbf{y}, \mathbf{y}')$, replace the second-stage objective coefficients \mathbf{d} with $(\mathbf{d}, \mathbf{c}')$ and impose the constraint that $\mathbf{y}' = \mathbf{x}$.*

Even in the special case where \mathcal{X} , \mathcal{Y} and Ξ are linear programming (LP) representable, problem (1) involves infinitely many decision variables and constraints, and it has been shown to be NP-hard [24]. Nevertheless, problem (1) simplifies considerably if the sets \mathcal{Y} and Ξ are LP representable. For this setting, several approximate solution schemes have been proposed that replace the second-stage decisions with *decision rules*, i.e., parametric classes of linear or nonlinear functions of $\boldsymbol{\xi}$ [3, 15, 18, 19, 29]. If we further assume that \mathbf{d} , \mathbf{T} and \mathbf{W} are deterministic and Ξ is of simple form (e.g., a budget uncertainty set), a number of exact solution schemes based on Benders’ decomposition [10, 27, 34, 40] and semi-infinite programming [1, 39] have been developed.

Problem (1) becomes significantly more challenging if the set \mathcal{Y} is not LP representable. For this setting, conservative MILP approximations have been developed in [21, 36] by partitioning the uncertainty set Ξ into hyperrectangles and restricting the continuous and integer recourse decisions to affine and constant functions of $\boldsymbol{\xi}$ over each hyperrectangle, respectively. These a priori partitioning schemes have been extended to iterative partitioning approaches in [6, 32]. Iterative solution approaches based on decision rules have been proposed in [7, 8]. However, to the best of our knowledge, none of these approaches has been shown to converge to an optimal solution of problem (1). For the special case where problem (1) has a relatively complete recourse, \mathbf{d} , \mathbf{T} and \mathbf{W} are deterministic and the optimal value of the second-stage problem is quasi-convex over Ξ , the solution scheme of [39] has been extended in [41] to a nested semi-infinite approach that can solve instances of problem (1) with MILP representable sets \mathcal{Y} and Ξ to optimality in finite time.

Instead of solving problem (1) directly, we study its *K-adaptability problem*

$$\inf_{\substack{\mathbf{x} \in \mathcal{X}, \\ \mathbf{y} \in \mathcal{Y}^K}} \sup_{\boldsymbol{\xi} \in \Xi} \inf_{k \in \mathcal{K}} \left\{ \mathbf{c}^\top \mathbf{x} + \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_k : \mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}(\boldsymbol{\xi})\mathbf{y}_k \leq \mathbf{h}(\boldsymbol{\xi}) \right\}, \quad (2)$$

where $\mathcal{Y}^K = \times_{k=1}^K \mathcal{Y}$ and $\mathcal{K} = \{1, \dots, K\}$. Problem (2) determines K non-adjustable second-stage policies $\mathbf{y}_1, \dots, \mathbf{y}_K$ here-and-now and subsequently selects the best of these policies in response to the observed value of $\boldsymbol{\xi}$. If all policies are infeasible for some realization $\boldsymbol{\xi} \in \Xi$, then the solution (\mathbf{x}, \mathbf{y}) attains the objective value $+\infty$. By construction, the K -adaptability problem (2) bounds the two-stage robust optimization problem (1) from above.

Our interest in problem (2) is motivated by two observations. Firstly, problem (2) has been shown to be a remarkably good approximation of problem (1), both in theory and in numerical experiments [5, 25]. Secondly, and perhaps more importantly, the K -adaptability problem conforms well with human decision-making, which tends to address uncertainty by developing a small number of contingency plans, rather than devising the optimal response for every possible future state of the world. For instance, practitioners may prefer a limited number of contingency plans to full flexibility in the second stage for operational (e.g., in production planning or logistics) or organizational (e.g., in emergency response planning) reasons.

The K -adaptability problem was first studied in [5], where the authors reformulate the 2-adaptability problem as a finite-dimensional bilinear program and solve it heuristically. The authors also show that the 2-adaptability problem is NP-hard even if \mathbf{d} , \mathbf{T} and \mathbf{W} are deterministic, and

they develop necessary conditions for the K -adaptability problem (2) to outperform the static robust problem (where all decisions are taken here-and-now). The relationship between the K -adaptability problem (2) and static robust optimization is further explored in [9] for the special case where \mathbf{T} and \mathbf{W} are deterministic. The authors show that the gaps between both problems and the two-stage robust optimization problem (1) are intimately related to geometric properties of the uncertainty set Ξ . Finite-dimensional MILP reformulations for problem (2) are developed in [25] under the additional assumption that both the here-and-now decisions \mathbf{x} and the wait-and-see decisions \mathbf{y} are binary. The authors show that both the size of the reformulations as well as their gaps to the two-stage robust optimization problem (1) depend on whether the uncertainty only affects the objective coefficients \mathbf{d} , or whether the constraint coefficients \mathbf{T} , \mathbf{W} and \mathbf{h} are uncertain as well. Finally, it is shown in [12, 13] that for polynomial time solvable deterministic combinatorial optimization problems, the associated instances of problem (2) without first-stage decisions \mathbf{x} can also be solved in polynomial time if all of the following conditions hold: (i) Ξ is convex, (ii) only the objective coefficients \mathbf{d} are uncertain, and (iii) $K > N_2$ policies are sought. This result has been extended to discrete uncertainty sets in [14], in which case pseudo-polynomial solution algorithms can be developed.

In this paper, we expand the literature on the K -adaptability problem in two ways. From an *analytical viewpoint*, we compare the two-stage robust optimization problem (1) with the K -adaptability problem (2) in terms of their continuity, convexity and tractability. We also investigate when the approximation offered by the K -adaptability problem is tight, and under which conditions the two-stage robust optimization and K -adaptability problems reduce to single-stage problems. From an *algorithmic viewpoint*, we develop a branch-and-bound scheme for the K -adaptability problem that combines ideas from semi-infinite and disjunctive programming. We establish conditions for its asymptotic and finite time convergence; we show how it can be refined and integrated into state-of-the-art MILP solvers; and, we present a heuristic variant that can address large-scale instances. In contrast to existing approaches, our algorithm can handle mixed continuous and discrete decisions in both stages as well as discrete uncertainty, and allows for modeling continuous second-stage decisions via a novel class of highly flexible piecewise affine decision rules. Extensive numerical experiments on benchmark data from various application domains indicate that our algorithm is highly competitive with state-of-the-art solution schemes for problems (1) and (2).

The remainder of this paper develops as follows. Section 2 analyzes the geometry and tractability of problems (1) and (2), and it proposes a novel class of piecewise affine decision rules that can be modeled as an instance of problem (2) with continuous second-stage decisions. Section 3 develops a branch-and-bound algorithm for the K -adaptability problem and analyzes its convergence. We present numerical results in Section 4, and we offer concluding remarks in Section 5.

Notation. Vectors and matrices are printed in bold lowercase and uppercase letters, respectively, while scalars are printed in regular font. We use \mathbf{e}_k to denote the k^{th} unit basis vector and \mathbf{e} to denote the vector whose components are all ones, respectively; their dimensions will be clear from the context. The i^{th} row vector of a matrix \mathbf{A} is denoted by \mathbf{a}_i^\top . For a logical expression \mathcal{E} , we define $\mathbb{I}[\mathcal{E}]$ as the indicator function which takes a value of 1 if \mathcal{E} is true and 0 otherwise.

2 Problem Analysis

In this section, we analyze the geometry and tractability of the two-stage robust optimization problem (1) and its associated K -adaptability problem (2). Specifically, we characterize the continuity, convexity and tractability of both problems, as well as their relationship to the static robust optimization problem where all decisions are taken here-and-now. We also show how the K -adaptability problem with continuous second-stage decisions enables us to approximate the two-stage robust optimization problem (1) through highly flexible piecewise affine decision rules.

Table 1 summarizes our theoretical results. For the sake of brevity, we present the formal statements of these results and their proofs as supplementary material at the end of the paper. In the table, the *first-stage problem* refers to the overall problems (1) and (2), the *evaluation problem* refers to the maximization over $\boldsymbol{\xi} \in \Xi$ for a fixed first-stage decision, and the *second-stage problem* refers to the inner minimization over $\mathbf{y} \in \mathcal{Y}$ or $k \in \mathcal{K}$ for a fixed first-stage decision and a fixed realization of the uncertain problem parameters. The table reveals that despite significant differences in their formulations, the problems (1) and (2) behave very similarly. The most significant difference is caused by the replacement of the optimization over the second-stage decisions $\mathbf{y} \in \mathcal{Y}$ in problem (1) with the selection of a candidate policy $k \in \mathcal{K}$ in problem (2). This ensures that the second-stage problem in (2) is always continuous, convex and tractable, whereas the first-stage problem in (2) fails to be convex even if \mathcal{X} and \mathcal{Y} are convex. Moreover, in contrast to problem (1), the evaluation problem in (2) remains continuous as long as either the objective function or the constraints are

		Continuity	Convexity	Tractability
First-stage problem	problem (1)	if feasible	if \mathcal{X}, \mathcal{Y} convex	if $\mathcal{X}, \mathcal{Y}, \Xi$ convex and OBJ
	problem (2)	if feasible	typically not	if $\mathcal{X}, \mathcal{Y}, \Xi$ convex and OBJ
Evaluation problem	problem (1)	if OBJ	if Ξ convex and OBJ	if Ξ, \mathcal{Y} convex and OBJ
	problem (2)	if OBJ or CON	if Ξ convex and OBJ	if Ξ convex and OBJ
Second-stage problem	problem (1)	if feasible	if \mathcal{Y} convex	if \mathcal{Y} convex
	problem (2)	if feasible	always	always

Propositions A.1, B.1 Propositions A.2, B.2 Propositions A.3, B.3

Reduction to static problem		Optimality of problem (2)
problem (1)	if Ξ convex and OBJ	if \mathcal{Y}, Ξ convex and OBJ
problem (2)	if Ξ convex, OBJ and $K > \min\{N_2, N_p\}$	if Ξ convex, OBJ and $K > \min\{N_2, N_p\}$
Propositions A.4, B.4		if $ \mathcal{Y} $ finite and $K \geq \mathcal{Y} $
		Proposition B.5

Table 1. Summary of theoretical results. Here, “OBJ” refers to instances where only the objective coefficients \mathbf{d} are uncertain, while \mathbf{T}, \mathbf{W} and \mathbf{h} are constant. Similarly, “CON” refers to instances where only the constraint coefficients \mathbf{T}, \mathbf{W} and right-hand sides \mathbf{h} are uncertain, while \mathbf{d} is constant.

unaffected by uncertainty. For general problem instances, however, neither of the two evaluation problems is continuous. As we will see in Section 3.2, this directly impacts the convergence of our branch-and-bound algorithm, which only takes place asymptotically in general. Note that the convexity of the problems (1) and (2) does not depend on the shape of the uncertainty set Ξ .

2.1 Incorporating Decision Rules in the K -Adaptability Problem

Although the K -adaptability problem (2) selects the best candidate policy \mathbf{y}_k in response to the observed parameter realization $\boldsymbol{\xi} \in \Xi$, the policies $\mathbf{y}_1, \dots, \mathbf{y}_K$ —once selected in the first stage—no longer depend on $\boldsymbol{\xi}$. This lack of dependence on the uncertain problem parameters can lead to overly conservative approximations of the two-stage robust optimization problem (1) when the second-stage decisions are continuous. In this section, we show how the K -adaptability problem (2) can be used to generalize affine decision rules, which are commonly used to approximate continuous instances of the two-stage robust optimization problem (1). We note that existing schemes, such as [13, 25], cannot be used for this purpose as they require the wait-and-see decisions \mathbf{y} to be binary.

Throughout this section, we assume that problem (1) has purely continuous second-stage decisions (that is, \mathcal{Y} is LP representable), a deterministic objective function (that is, $\mathbf{d}(\boldsymbol{\xi}) = \mathbf{d}$ for all $\boldsymbol{\xi} \in \Xi$) and fixed recourse (that is, $\mathbf{W}(\boldsymbol{\xi}) = \mathbf{W}$ for all $\boldsymbol{\xi} \in \Xi$). The assumption of continuous second-stage decisions allows us to assume, without loss of generality, that $\mathcal{Y} = \mathbb{R}^{N_2}$ as any potential restrictions can be absorbed in the second-stage constraints.

The affine decision rule approximation to the two-stage robust optimization problem (1) is

$$\inf_{\substack{\mathbf{x} \in \mathcal{X}, \\ \mathbf{y} : \Xi \mapsto \mathbb{R}^{N_2}}} \left\{ \sup_{\boldsymbol{\xi} \in \Xi} \left\{ \mathbf{c}^\top \mathbf{x} + \mathbf{d}^\top \mathbf{y}(\boldsymbol{\xi}) \right\} : \mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}\mathbf{y}(\boldsymbol{\xi}) \leq \mathbf{h}(\boldsymbol{\xi}) \quad \forall \boldsymbol{\xi} \in \Xi \right\},$$

where $\mathbf{y} : \Xi \mapsto \mathbb{R}^{N_2}$ indicates that $\mathbf{y}(\boldsymbol{\xi}) = \mathbf{y}^0 + \mathbf{Y}\boldsymbol{\xi}$ for some $\mathbf{y}^0 \in \mathbb{R}^{N_2}$ and $\mathbf{Y} \in \mathbb{R}^{N_2 \times N_p}$, see [3]. This problem provides a conservative approximation to the two-stage robust optimization problem (1) since we replace the space of all (possibly non-convex and discontinuous) second-stage policies $\mathbf{y} : \Xi \mapsto \mathbb{R}^{N_2}$ with the subspace of all affine second-stage policies $\mathbf{y} : \Xi \mapsto \mathbb{R}^{N_2}$. In a similar spirit, we define the subspace of all piecewise affine decision rules $\mathbf{y} : \Xi \overset{K}{\mapsto} \mathbb{R}^{N_2}$ with K pieces as

$$\mathbf{y} : \Xi \overset{K}{\mapsto} \mathbb{R}^{N_2} \iff \left[\begin{array}{l} \exists (\mathbf{y}_k^0, \mathbf{Y}_k) \in \mathbb{R}^{N_2} \times \mathbb{R}^{N_2 \times N_p}, k = 1, \dots, K, \text{ such that} \\ \forall \boldsymbol{\xi} \in \Xi, \exists k \in \{1, \dots, K\} : \mathbf{y}(\boldsymbol{\xi}) = \mathbf{y}_k^0 + \mathbf{Y}_k \boldsymbol{\xi} \end{array} \right].$$

Note that our earlier definition of $\mathbf{y} : \Xi \overset{1}{\mapsto} \mathbb{R}^{N_2}$ is identical to our definition of $\mathbf{y} : \Xi \overset{K}{\mapsto} \mathbb{R}^{N_2}$ if $K = 1$. For $K > 1$, the decision rules $\mathbf{y} : \Xi \overset{K}{\mapsto} \mathbb{R}^{N_2}$ may be non-convex and discontinuous, and the regions where \mathbf{y} is affine may be non-closed and non-convex. We highlight that the points of nonlinearity are determined by the optimization problem. This is in contrast to many existing solution schemes for piecewise affine decision rules, such as [8, 15, 18, 19], where these points are specified ad hoc by the decision-maker.

Observation 1. *The piecewise affine decision rule problem with fixed recourse*

$$\inf_{\substack{\mathbf{x} \in \mathcal{X}, \\ \mathbf{y} : \Xi \overset{K}{\mapsto} \mathbb{R}^{N_2}}} \left\{ \sup_{\boldsymbol{\xi} \in \Xi} \left\{ \mathbf{c}^\top \mathbf{x} + \mathbf{d}^\top \mathbf{y}(\boldsymbol{\xi}) \right\} : \mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}\mathbf{y}(\boldsymbol{\xi}) \leq \mathbf{h}(\boldsymbol{\xi}) \quad \forall \boldsymbol{\xi} \in \Xi \right\} \quad (3)$$

is equivalent to the K -adaptability problem with random recourse

$$\inf_{\substack{\mathbf{x} \in \mathcal{X}, \\ (\mathbf{y}^0, \mathbf{Y}, \mathbf{z}) \in \hat{\mathcal{Y}}^K}} \sup_{\boldsymbol{\xi} \in \Xi} \inf_{k \in \mathcal{K}} \left\{ \mathbf{c}^\top \mathbf{x} + \mathbf{d}^\top \mathbf{y}_k^0 + \hat{\mathbf{d}}(\boldsymbol{\xi})^\top \mathbf{z}_{k1} : \mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}\mathbf{y}_k^0 + \hat{\mathbf{W}}(\boldsymbol{\xi})\mathbf{z}_{k2} \leq \mathbf{h}(\boldsymbol{\xi}) \right\}, \quad (4)$$

where $\hat{\mathcal{Y}} = \left\{ (\mathbf{y}^0, \mathbf{Y}, \mathbf{z}) \in \mathbb{R}^{N_2} \times \mathbb{R}^{N_2 \times N_p} \times (\mathbb{R}^{N_p} \times \mathbb{R}^{N_p L}) : \mathbf{z} = (\mathbf{z}_1, \mathbf{z}_2) \text{ with } \mathbf{z}_1 = \mathbf{Y}^\top \mathbf{d} \text{ and } \mathbf{z}_2 = [\mathbf{w}_1^\top \mathbf{Y} \dots \mathbf{w}_L^\top \mathbf{Y}]^\top \right\}$, $\hat{\mathbf{d}}(\boldsymbol{\xi}) = \boldsymbol{\xi}$ and $\hat{\mathbf{W}}(\boldsymbol{\xi}) = \text{diag}(\boldsymbol{\xi}^\top, \dots, \boldsymbol{\xi}^\top) \in \mathbb{R}^{L \times N_p L}$.

Proof. Problem (4) is infeasible if and only if for every $\mathbf{x} \in \mathcal{X}$ and $(\mathbf{y}^0, \mathbf{Y}, \mathbf{z}) \in \hat{\mathcal{Y}}^K$ there is a $\boldsymbol{\xi} \in \Xi$ such that $\mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}\mathbf{y}_k^0 + \hat{\mathbf{W}}(\boldsymbol{\xi})\mathbf{z}_{k2} \not\leq \mathbf{h}(\boldsymbol{\xi})$ for all $k = 1, \dots, K$, which in turn is the case if and only if for every $\mathbf{x} \in \mathcal{X}$ and $(\mathbf{y}_k^0, \mathbf{Y}_k) \in \mathbb{R}^{N_2} \times \mathbb{R}^{N_2 \times N_p}$, $k = 1, \dots, K$ there is a $\boldsymbol{\xi} \in \Xi$ such that $\mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}\mathbf{y}_k^0 + \mathbf{W}\mathbf{Y}_k\boldsymbol{\xi} \not\leq \mathbf{h}(\boldsymbol{\xi})$ for all $k = 1, \dots, K$; that is, if and only if problem (3) is infeasible. We thus assume that both (3) and (4) are feasible. In this case, we verify that every feasible solution $(\mathbf{x}, \mathbf{y}^0, \mathbf{Y}, \mathbf{z})$ to problem (4) gives rise to a feasible solution (\mathbf{x}, \mathbf{y}) , where $\mathbf{y}(\boldsymbol{\xi}) = \mathbf{y}_{k(\boldsymbol{\xi})}^0 + \mathbf{Y}_{k(\boldsymbol{\xi})}\boldsymbol{\xi}$ and $k(\boldsymbol{\xi})$ is any element of $\arg \min_{k \in \mathcal{K}} \left\{ \mathbf{c}^\top \mathbf{x} + \mathbf{d}^\top \mathbf{y}_k^0 + \hat{\mathbf{d}}(\boldsymbol{\xi})^\top \mathbf{z}_{k1} : \mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}\mathbf{y}_k^0 + \hat{\mathbf{W}}(\boldsymbol{\xi})\mathbf{z}_{k2} \leq \mathbf{h}(\boldsymbol{\xi}) \right\}$, in problem (3) that attains the same worst-case objective value. Similarly, every *optimal solution* (\mathbf{x}, \mathbf{y}) to problem (3) gives rise to an optimal solution $(\mathbf{x}, \mathbf{y}^0, \mathbf{Y}, \mathbf{z})$, where $\mathbf{z}_k = (\mathbf{z}_{k1}, \mathbf{z}_{k2})$ with $\mathbf{z}_{k1} = \mathbf{Y}_k^\top \mathbf{d}$ and $\mathbf{z}_{k2} = [\mathbf{w}_1^\top \mathbf{Y}_k \dots \mathbf{w}_L^\top \mathbf{Y}_k]^\top$, $k = 1, \dots, K$, in problem (4). Hence, (3) and (4) share the same optimal value and the same sets of optimal solutions. \square

We close with an example that illustrates the benefits of piecewise affine decision rules.

Example 1. Consider the following instance of the two-stage robust optimization problem (1), which has been proposed in [20, Section 5.2]:

$$\sup_{\boldsymbol{\xi} \in [-1, 1]^2} \inf_{\mathbf{y} \in \mathbb{R}_+^4} \left\{ \mathbf{e}^\top \mathbf{y} : y_1 \geq \xi_1 + \xi_2, y_2 \geq \xi_1 - \xi_2, y_3 \geq -\xi_1 + \xi_2, y_4 \geq -\xi_1 - \xi_2 \right\}.$$

The optimal second-stage policy is $\mathbf{y}^*(\boldsymbol{\xi}) = ([\xi_1 + \xi_2]_+, [\xi_1 - \xi_2]_+, [-\xi_1 + \xi_2]_+, [-\xi_1 - \xi_2]_+)$, where $[\cdot]_+ = \max\{\cdot, 0\}$, and it results in the optimal second-stage value function $Q^*(\boldsymbol{\xi}) = [\xi_1 + \xi_2]_+ + [\xi_1 - \xi_2]_+ + [-\xi_1 + \xi_2]_+ + [-\xi_1 - \xi_2]_+$ with a worst-case objective value of 2, see Figure 1. The best affine decision rule $\mathbf{y}^1 : \Xi \xrightarrow{1} \mathbb{R}_+^4$ is $\mathbf{y}^1(\boldsymbol{\xi}) = (1 + \xi_2, 1 + \xi_1, 1 - \xi_1, 1 - \xi_2)$, and it results in the constant second-stage value function $Q^1(\boldsymbol{\xi}) = 4$. The best 2-adaptable affine decision rule $\mathbf{y}^2 : \Xi \xrightarrow{2} \mathbb{R}_+^4$, on the other hand, is given by

$$\mathbf{y}^2(\boldsymbol{\xi}) = \begin{cases} (0, 1 + \xi_1, 1 + \xi_2, -\xi_1 - \xi_2) & \text{if } \xi_1 + \xi_2 \leq 0 \\ (\xi_1 + \xi_2, 1 - \xi_2, 1 - \xi_1, 0) & \text{otherwise,} \end{cases}$$

and it results in the constant second-stage value function $Q^2(\boldsymbol{\xi}) = 2$. Thus, 2-adaptable affine decision rules are optimal in this example. Figure 2 illustrates the optimal value, the affine approximation and the 2-adaptable affine approximation of the decision variable y_3 .

The piecewise affine decision rules presented here can be readily combined with discrete second-stage decisions. For the sake of brevity, we omit the details of this straightforward extension.




Figure 1. The optimal second-stage value function in Example 1 is given by the first-order cone $Q^*(\boldsymbol{\xi}) = [\xi_1 + \xi_2]_+ + [\xi_1 - \xi_2]_+ + [-\xi_1 + \xi_2]_+ + [-\xi_1 - \xi_2]_+$.




Figure 2. Plot of the optimal second-stage policy $y_3(\boldsymbol{\xi})$ in the two-stage robust optimization problem (left), the affine decision rule problem (middle) and the 2-adaptable affine decision rule problem (right).

3 Solution Scheme

Our solution scheme for the K -adaptability problem (2) is based on a reformulation as a semi-infinite disjunctive program which we present next.

Observation 2. *The K -adaptability problem (2) is equivalent to*

$$\begin{aligned}
 & \text{minimize} && \theta \\
 & \text{subject to} && \theta \in \mathbb{R}, \mathbf{x} \in \mathcal{X}, \mathbf{y} \in \mathcal{Y}^K \\
 & && \bigvee_{k \in \mathcal{K}} \left[\begin{array}{l} \mathbf{c}^\top \mathbf{x} + \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_k \leq \theta \\ \mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}(\boldsymbol{\xi})\mathbf{y}_k \leq \mathbf{h}(\boldsymbol{\xi}) \end{array} \right] \quad \forall \boldsymbol{\xi} \in \Xi.
 \end{aligned} \tag{5}$$

Moreover, if some of the constraints in problem (5) are deterministic, i.e., they do not depend on $\boldsymbol{\xi}$, then they can be moved outside the disjunction and instead be enforced for all $k \in \mathcal{K}$.

In the following, we stipulate that the optimal value of (5) is $+\infty$ whenever it is infeasible.

Proof of Observation 2. Problem (2) is infeasible if and only if (iff) for every $\mathbf{x} \in \mathcal{X}$ and $\mathbf{y} \in \mathcal{Y}^K$ there is a $\boldsymbol{\xi} \in \Xi$ such that $\mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}(\boldsymbol{\xi})\mathbf{y}_k \not\leq \mathbf{h}(\boldsymbol{\xi})$ for all $k \in \mathcal{K}$, which in turn is the case iff

for every $\mathbf{x} \in \mathcal{X}$ and $\mathbf{y} \in \mathcal{Y}^K$, the disjunction in (5) is violated for at least one $\boldsymbol{\xi} \in \Xi$; that is, iff problem (5) is infeasible. We thus assume that both (2) and (5) are feasible. In this case, one readily verifies that every feasible solution (\mathbf{x}, \mathbf{y}) to problem (2) gives rise to a feasible solution $(\theta, \mathbf{x}, \mathbf{y})$, where $\theta = \sup_{\boldsymbol{\xi} \in \Xi} \inf_{\mathbf{y} \in \mathcal{Y}^K} \{ \mathbf{c}^\top \mathbf{x} + \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_k : \mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}(\boldsymbol{\xi})\mathbf{y}_k \leq \mathbf{h}(\boldsymbol{\xi}) \}$, in problem (5) with the same objective value. Likewise, any *optimal* solution $(\theta, \mathbf{x}, \mathbf{y})$ to problem (5) corresponds to an optimal solution (\mathbf{x}, \mathbf{y}) in problem (2). Hence, (2) and (5) share the same optimal value and the same sets of optimal solutions.

We now claim that if $\mathbf{t}_l(\boldsymbol{\xi})^\top = \mathbf{t}_l^\top$, $\mathbf{w}_l(\boldsymbol{\xi})^\top = \mathbf{w}_l^\top$ and $h_l(\boldsymbol{\xi}) = h_l$ for all $l \in \{1, \dots, L\} \setminus \mathcal{L}$, where $\mathcal{L} \subseteq \{1, \dots, L\}$, then problem (5) is equivalent to

$$\begin{aligned} & \text{minimize} && \theta \\ & \text{subject to} && \theta \in \mathbb{R}, \quad \mathbf{x} \in \mathcal{X}, \quad \mathbf{y} \in \mathcal{Y}^K \\ & && \mathbf{t}_l^\top \mathbf{x} + \mathbf{w}_l^\top \mathbf{y}_k \leq h_l \quad \forall l \in \{1, \dots, L\} \setminus \mathcal{L}, \quad \forall k \in \mathcal{K} \quad (5') \\ & && \bigvee_{k \in \mathcal{K}} \left[\begin{array}{l} \mathbf{c}^\top \mathbf{x} + \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_k \leq \theta \\ \mathbf{t}_l(\boldsymbol{\xi})^\top \mathbf{x} + \mathbf{w}_l(\boldsymbol{\xi})^\top \mathbf{y}_k \leq h_l(\boldsymbol{\xi}) \quad \forall l \in \mathcal{L} \end{array} \right] \quad \forall \boldsymbol{\xi} \in \Xi. \end{aligned}$$

By construction, every feasible solution $(\theta, \mathbf{x}, \mathbf{y})$ to problem (5') is feasible in problem (5). Conversely, fix any feasible solution $(\theta, \mathbf{x}, \mathbf{y})$ to problem (5) and assume that $\mathbf{t}_l^\top \mathbf{x} + \mathbf{w}_l^\top \mathbf{y}_k > h_l$ for some $k \in \mathcal{K}$ and $l \in \{1, \dots, L\} \setminus \mathcal{L}$. In that case, the k^{th} disjunct in (5) is violated for *every* realization $\boldsymbol{\xi} \in \Xi$. We can therefore replace \mathbf{y}_k with a different candidate policy $\mathbf{y}_{k'}$ that satisfies $\mathbf{t}_l^\top \mathbf{x} + \mathbf{w}_l^\top \mathbf{y}_{k'} \leq h_l$ for all $l \in \{1, \dots, L\} \setminus \mathcal{L}$ without sacrificing feasibility. (Note that such a candidate policy $\mathbf{y}_{k'}$ exists since $(\theta, \mathbf{x}, \mathbf{y})$ is assumed to be feasible in (5).) Replacing any infeasible policy in this way results in a solution that is feasible in problem (5'). \square

Problem (5) cannot be solved directly as it contains infinitely many disjunctive constraints. Instead, our solution scheme iteratively solves a sequence of (increasingly tighter) relaxations of this problem that are obtained by enforcing the disjunctive constraints over finite subsets of Ξ . Whenever the solution of such a relaxation violates the disjunction for some realization $\boldsymbol{\xi} \in \Xi$, we create K subproblems that enforce the disjunction associated with $\boldsymbol{\xi}$ to be satisfied by the k^{th} disjunct, $k = 1, \dots, K$. Our solution scheme is reminiscent of discretization methods employed in *semi-infinite programming*, which iteratively replace an infinite set of constraints with finite subsets and solve the resulting discretized problems. Indeed, our scheme can be interpreted as a generalization of Kelley's cutting-plane method [11, 28] applied to semi-infinite disjunctive programs. In

the special case where $K = 1$, our method reduces to the cutting-plane method for (static) robust optimization problems proposed in [31].

In the remainder of this section, we describe our basic branch-and-bound scheme (Section 3.1), we study its convergence (Section 3.2), we discuss algorithmic variants to the basic scheme that can enhance its numerical performance (Section 3.3), and we present a heuristic variant that can address problems of larger scale (Section 3.4).

3.1 Branch-and-Bound Algorithm

Our solution scheme iteratively solves a sequence of scenario-based K -adaptability problems and separation problems. We define both problems first, and then we describe the overall algorithm.

The Scenario-Based K -Adaptability Problem. For a collection Ξ_1, \dots, Ξ_K of finite subsets of the uncertainty set Ξ , we define the *scenario-based K -adaptability problem* as

$$\begin{aligned} \mathcal{M}(\Xi_1, \dots, \Xi_K) = \text{minimize} \quad & \theta \\ \text{subject to} \quad & \theta \in \mathbb{R}, \quad \mathbf{x} \in \mathcal{X}, \quad \mathbf{y} \in \mathcal{Y}^K \\ & \left. \begin{aligned} \mathbf{c}^\top \mathbf{x} + \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_k &\leq \theta \\ \mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}(\boldsymbol{\xi})\mathbf{y}_k &\leq \mathbf{h}(\boldsymbol{\xi}) \end{aligned} \right\} \quad \forall \boldsymbol{\xi} \in \Xi_k, \quad \forall k \in \mathcal{K}. \end{aligned} \quad (6)$$

If \mathcal{X} and \mathcal{Y} are convex, problem (6) is an LP; otherwise, it is an MILP. The problem is closely related to a relaxation of the semi-infinite disjunctive program (5) that enforces the disjunction only over the realizations $\boldsymbol{\xi} \in \bigcup_{k \in \mathcal{K}} \Xi_k$. More precisely, problem (6) can be interpreted as a restriction of that relaxation which requires the k^{th} candidate policy \mathbf{y}_k to be worst-case optimal for all realizations $\boldsymbol{\xi} \in \Xi_k$, $k \in \mathcal{K}$. We obtain an optimal solution $(\theta, \mathbf{x}, \mathbf{y}_k)$ to the relaxed semi-infinite disjunctive program by solving $\mathcal{M}(\Xi_1, \dots, \Xi_K)$ for all partitions (Ξ_1, \dots, Ξ_K) of $\bigcup_{k \in \mathcal{K}} \Xi_k$ and reporting the optimal solution $(\theta, \mathbf{x}, \mathbf{y}_k)$ of the problem $\mathcal{M}(\Xi_1, \dots, \Xi_K)$ with the smallest objective value.

If $\Xi_k = \emptyset$ for all $k \in \mathcal{K}$, then problem (6) is unbounded, and we stipulate that its optimal value is $-\infty$ and that its optimal value is attained by any solution $(-\infty, \mathbf{x}, \mathbf{y})$ with $(\mathbf{x}, \mathbf{y}) \in \mathcal{X} \times \mathcal{Y}^K$. Otherwise, if problem (6) is infeasible for Ξ_1, \dots, Ξ_K , then we define its optimal value to be $+\infty$. In all other cases, the optimal value of problem (6) is finite and it is attained by an optimal solution $(\theta, \mathbf{x}, \mathbf{y})$ since \mathcal{X} and \mathcal{Y} are compact.

Remark 2 (Decomposability). *For K -adaptability problems without first-stage decisions \mathbf{x} , problem (6) decomposes into K scenario-based static robust optimization problems that are only coupled through the constraints referencing the epigraph variable θ . In this case, we can recover an optimal solution to problem (6) by solving each of the K static problems individually and identifying the optimal θ as the maximum of their optimal values.*

The Separation Problem. For a feasible solution $(\theta, \mathbf{x}, \mathbf{y})$ to the scenario-based K -adaptability problem (6), we define the *separation problem* as

$$\begin{aligned} S(\theta, \mathbf{x}, \mathbf{y}) &= \max_{\boldsymbol{\xi} \in \Xi} S(\theta, \mathbf{x}, \mathbf{y}, \boldsymbol{\xi}), \text{ where} \\ S(\theta, \mathbf{x}, \mathbf{y}, \boldsymbol{\xi}) &= \min_{k \in \mathcal{K}} \max \left\{ \mathbf{c}^\top \mathbf{x} + \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_k - \theta, \max_{l \in \{1, \dots, L\}} \left\{ \mathbf{t}_l(\boldsymbol{\xi})^\top \mathbf{x} + \mathbf{w}_l(\boldsymbol{\xi})^\top \mathbf{y}_k - h_l(\boldsymbol{\xi}) \right\} \right\}, \end{aligned} \quad (7)$$

for $\mathcal{S} : \mathbb{R} \cup \{-\infty\} \times \mathcal{X} \times \mathcal{Y}^K \mapsto \mathbb{R} \cup \{+\infty\}$ and $S : \mathbb{R} \cup \{-\infty\} \times \mathcal{X} \times \mathcal{Y}^K \times \Xi \mapsto \mathbb{R} \cup \{+\infty\}$. Whenever it is positive, the innermost maximum in the definition of $S(\theta, \mathbf{x}, \mathbf{y}, \boldsymbol{\xi})$ records the maximum constraint violation of the candidate policy \mathbf{y}_k under the parameter realization $\boldsymbol{\xi} \in \Xi$. Likewise, the quantity $\mathbf{c}^\top \mathbf{x} + \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_k - \theta$ denotes the excess of the objective value of \mathbf{y}_k under the realization $\boldsymbol{\xi}$ over the current candidate value of the worst-case objective, θ . Thus, $S(\theta, \mathbf{x}, \mathbf{y}, \boldsymbol{\xi})$ is strictly positive if and only if every candidate policy \mathbf{y}_k either is infeasible or results in an objective value greater than θ under the realization $\boldsymbol{\xi} \in \Xi$. Whenever θ is finite, the separation problem is feasible and bounded, and it has an optimal solution since Ξ is nonempty and compact. Otherwise, we have $S(\theta, \mathbf{x}, \mathbf{y}) = +\infty$, and the optimal value is attained by any $\boldsymbol{\xi} \in \Xi$.

Observation 3. *The separation problem (7) is equivalent to the MILP*

$$\begin{aligned} &\text{maximize} && \zeta \\ &\text{subject to} && \zeta \in \mathbb{R}, \quad \boldsymbol{\xi} \in \Xi, \quad z_{kl} \in \{0, 1\}, \quad (k, l) \in \mathcal{K} \times \{0, 1, \dots, L\} \\ & && \left. \begin{aligned} &\sum_{l=0}^L z_{kl} = 1 \\ &z_{k0} = 1 \Rightarrow \zeta \leq \mathbf{c}^\top \mathbf{x} + \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_k - \theta \\ &z_{kl} = 1 \Rightarrow \zeta \leq \mathbf{t}_l(\boldsymbol{\xi})^\top \mathbf{x} + \mathbf{w}_l(\boldsymbol{\xi})^\top \mathbf{y}_k - h_l(\boldsymbol{\xi}) \quad \forall l \in \{1, \dots, L\} \end{aligned} \right\} \forall k \in \mathcal{K}. \end{aligned} \quad (8)$$

This problem can be solved in polynomial time if Ξ is convex and the number of policies K is fixed.

Proof. Fix any feasible solution $(\theta, \mathbf{x}, \mathbf{y})$ to the scenario-based K -adaptability problem (6). For every $\boldsymbol{\xi} \in \Xi$, we can construct a feasible solution $(\zeta, \boldsymbol{\xi}, \mathbf{z})$ to problem (8) with $\zeta = S(\theta, \mathbf{x}, \mathbf{y}, \boldsymbol{\xi})$

by setting $z_{k0} = 1$ if $\mathbf{c}^\top \mathbf{x} + \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_k - \theta \geq \mathbf{t}_l(\boldsymbol{\xi})^\top \mathbf{x} + \mathbf{w}_l(\boldsymbol{\xi})^\top \mathbf{y}_k - h_l(\boldsymbol{\xi})$ for all $l = 1, \dots, L$ and $z_{kl} = 1$ for $l \in \arg \max_{l \in \{1, \dots, L\}} \{\mathbf{t}_l(\boldsymbol{\xi})^\top \mathbf{x} + \mathbf{w}_l(\boldsymbol{\xi})^\top \mathbf{y}_k - h_l(\boldsymbol{\xi})\}$ otherwise (where ties can be broken arbitrarily). We thus conclude that $\mathcal{S}(\theta, \mathbf{x}, \mathbf{y})$ is less than or equal to the optimal value of problem (8). Likewise, every feasible solution $(\zeta, \boldsymbol{\xi}, \mathbf{z})$ to problem (8) satisfies $\zeta \leq \max \{\mathbf{c}^\top \mathbf{x} + \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_k - \theta, \max_{l \in \{1, \dots, L\}} \{\mathbf{t}_l(\boldsymbol{\xi})^\top \mathbf{x} + \mathbf{w}_l(\boldsymbol{\xi})^\top \mathbf{y}_k - h_l(\boldsymbol{\xi})\}\}$ for all $k \in \mathcal{K}$; that is, $\zeta \leq S(\theta, \mathbf{x}, \mathbf{y}, \boldsymbol{\xi})$. Thus, the optimal value of problem (8) is less than or equal to $\mathcal{S}(\theta, \mathbf{x}, \mathbf{y})$ as well.

If the number of policies K is fixed and the uncertainty set Ξ is convex, then problem (8) can be solved by enumerating all $(L + 1)^K$ possible choices for \mathbf{z} , solving the resulting linear programs in ζ and $\boldsymbol{\xi}$ and reporting the solution with the maximum value of ζ . \square

The Algorithm. Our solution scheme solves a sequence of scenario-based K -adaptability problems (6) over monotonically increasing scenario sets Ξ_k , $k \in \mathcal{K}$. At each iteration, the separation problem (8) identifies a new scenario $\boldsymbol{\xi} \in \Xi$ to be added to these sets.

1. *Initialize.* Set $\mathcal{N} \leftarrow \{\tau^0\}$ (node set), where $\tau^0 = (\Xi_1^0, \dots, \Xi_K^0)$ with $\Xi_k^0 = \emptyset$ for all $k \in \mathcal{K}$ (root node). Set $(\theta^i, \mathbf{x}^i, \mathbf{y}^i) \leftarrow (+\infty, \emptyset, \emptyset)$ (incumbent solution).
2. *Check convergence.* If $\mathcal{N} = \emptyset$, then stop and declare infeasibility (if $\theta^i = +\infty$) or report $(\mathbf{x}^i, \mathbf{y}^i)$ as an optimal solution to problem (2).
3. *Select node.* Select a node $\tau = (\Xi_1, \dots, \Xi_K)$ from \mathcal{N} . Set $\mathcal{N} \leftarrow \mathcal{N} \setminus \{\tau\}$.
4. *Process node.* Let $(\theta, \mathbf{x}, \mathbf{y})$ be an optimal solution to the scenario-based K -adaptability problem (6). If $\theta \geq \theta^i$, then go to Step 2.
5. *Check feasibility.* Let $(\zeta, \boldsymbol{\xi}, \mathbf{z})$ be an optimal solution to the separation problem (8). If $\zeta \leq 0$, then set $(\theta^i, \mathbf{x}^i, \mathbf{y}^i) \leftarrow (\theta, \mathbf{x}, \mathbf{y})$ and go to Step 2.
6. *Branch.* Instantiate K new nodes τ_1, \dots, τ_K as follows: $\tau_k = (\Xi_1, \dots, \Xi_k \cup \{\boldsymbol{\xi}\}, \dots, \Xi_K)$ for each $k \in \mathcal{K}$. Set $\mathcal{N} \leftarrow \mathcal{N} \cup \{\tau_1, \dots, \tau_K\}$ and go to Step 3.

Our branch-and-bound algorithm can be interpreted as an uncertainty set partitioning scheme. For a solution $(\theta, \mathbf{x}, \mathbf{y})$ in Step 4, the sets

$$\Xi(\theta, \mathbf{x}, \mathbf{y}_k) = \left\{ \boldsymbol{\xi} \in \Xi : \mathbf{c}^\top \mathbf{x} + \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_k \leq \theta, \mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}(\boldsymbol{\xi})\mathbf{y}_k \leq \mathbf{h}(\boldsymbol{\xi}) \right\}, \quad k \in \mathcal{K},$$

describe the regions of the uncertainty set Ξ for which at least one of the candidate policies is feasible and results in an objective value smaller than or equal to θ . Step 5 of the algorithm attempts to identify a realization $\xi \in \Xi \setminus \bigcup_{k \in \mathcal{K}} \Xi(\theta, \mathbf{y}_k)$ for which every candidate policy either is infeasible or results in an objective value that exceeds θ . If there is no such realization, then the solution (\mathbf{x}, \mathbf{y}) is feasible in the K -adaptability problem (2). Otherwise, Step 6 assigns the realization ξ to each scenario subset Ξ_k , $k \in \mathcal{K}$, in turn. Figure 3 illustrates our solution scheme.




Figure 3. An illustrative example with $K = 2$ policies. Each cube represents the uncertainty set Ξ while the shaded regions represent $\Xi(\theta, \mathbf{y}_1)$ and $\Xi(\theta, \mathbf{y}_2)$. The green and blue dots represent elements of the sets Ξ_1 and Ξ_2 , respectively, while the red squares represent the candidate realizations ξ identified in Step 5 of the algorithm.

3.2 Convergence Analysis

We now establish the correctness of our branch-and-bound scheme, as well as conditions for its asymptotic and finite convergence.

Theorem 1 (Correctness). *If the branch-and-bound scheme terminates, then it either returns an optimal solution to problem (2) or correctly identifies the latter as infeasible.*

Proof. We first show that if the problem instance is infeasible, then the algorithm terminates with the incumbent solution $(\theta^i, \mathbf{x}^i, \mathbf{y}^i) = (+\infty, \emptyset, \emptyset)$. Indeed, the algorithm can only update the

incumbent solution in Step 5 if the objective value of the separation problem is non-positive. By construction, this is only possible if the algorithm has determined a feasible solution.

We now show that for feasible problem instances, the algorithm terminates with an optimal solution $(\mathbf{x}^i, \mathbf{y}^i)$ of problem (2). To this end, assume that $(\mathbf{x}^*, \mathbf{y}^*)$ is an optimal solution of problem (2) with objective value θ^* . Let \mathcal{T} be the set of all nodes of the branch-and-bound tree for which $(\theta^*, \mathbf{x}^*, \mathbf{y}^*)$ is feasible in the corresponding scenario-based K -adaptability problem (6). Note that $\mathcal{T} \neq \emptyset$ since $(\theta^*, \mathbf{x}^*, \mathbf{y}^*)$ is feasible in the root node. Let $\mathcal{T}' \subseteq \mathcal{T}$ be the set of those nodes which have children in \mathcal{T} and consider the set $\mathcal{T}'' = \mathcal{T} \setminus \mathcal{T}'$; by construction, we have $\mathcal{T}'' \neq \emptyset$. Consider an arbitrary node $\tau \in \mathcal{T}''$. By definition of \mathcal{T}'' , our algorithm has not branched τ . Since τ has been selected in Step 3, this is only possible if either (i) τ has been fathomed in Step 4 or if (ii) τ has been fathomed in Step 5. In the former case, the solution $(\theta^*, \mathbf{x}^*, \mathbf{y}^*)$ must have been weakly dominated by the incumbent solution $(\theta^i, \mathbf{x}^i, \mathbf{y}^i)$, which therefore must be optimal as well. In the latter case, the incumbent solution must have been updated to $(\theta^*, \mathbf{x}^*, \mathbf{y}^*)$. \square

We now show that our branch-and-bound scheme converges asymptotically to an optimal solution of the K -adaptability problem (2). Our result has two implications: (i) for infeasible problem instances, the algorithm always terminates after finitely many iterations, i.e., infeasibility is detected in finite time; (ii) for feasible problem instances, the algorithm eventually only inspects solutions in the neighborhood of optimal solutions.

Theorem 2 (Asymptotic Convergence). *Every accumulation point $(\hat{\theta}, \hat{\mathbf{x}}, \hat{\mathbf{y}})$ of the solutions to the scenario-based K -adaptability problem (6) in an infinite branch of the branch-and-bound tree gives rise to an optimal solution $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ of the K -adaptability problem (2) with objective value $\hat{\theta}$.*

Proof. We denote by $(\theta^\ell, \mathbf{x}^\ell, \mathbf{y}^\ell)$ and $(\zeta^\ell, \boldsymbol{\xi}^\ell, \mathbf{z}^\ell)$ the sequences of optimal solutions to the scenario-based K -adaptability problem in Step 4 and the separation problem in Step 5 of the algorithm, respectively, that correspond to the node sequence τ^ℓ , $\ell = 0, 1, \dots$, of some infinite branch of the branch-and-bound tree. Since \mathcal{X} , \mathcal{Y} and Ξ are compact, the Bolzano-Weierstrass theorem implies that $(\theta^\ell, \mathbf{x}^\ell, \mathbf{y}^\ell)$ and $(\zeta^\ell, \boldsymbol{\xi}^\ell, \mathbf{z}^\ell)$ each have at least one accumulation point.

We first show that every accumulation point $(\hat{\theta}, \hat{\mathbf{x}}, \hat{\mathbf{y}})$ of the sequence $(\theta^\ell, \mathbf{x}^\ell, \mathbf{y}^\ell)$ corresponds to a feasible solution $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ of the K -adaptability problem (2) with objective value $\hat{\theta}$. By possibly going over to subsequences, we can without loss of generality assume that the two sequences $(\theta^\ell, \mathbf{x}^\ell, \mathbf{y}^\ell)$

and $(\zeta^\ell, \boldsymbol{\xi}^\ell, \mathbf{z}^\ell)$ converge themselves to $(\hat{\theta}, \hat{\mathbf{x}}, \hat{\mathbf{y}})$ and $(\hat{\zeta}, \hat{\boldsymbol{\xi}}, \hat{\mathbf{z}})$, respectively. Assume now that $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ does *not* correspond to a feasible solution of the K -adaptability problem (2) with objective value $\hat{\theta}$. Then there is $\boldsymbol{\xi}^* \in \Xi$ such that $S(\hat{\theta}, \hat{\mathbf{x}}, \hat{\mathbf{y}}, \boldsymbol{\xi}^*) \geq \delta$ for some $\delta > 0$. By construction of the separation problem (8), this implies that

$$S(\theta^\ell, \mathbf{x}^\ell, \mathbf{y}^\ell, \boldsymbol{\xi}^\ell) = \max_{\boldsymbol{\xi} \in \Xi} S(\theta^\ell, \mathbf{x}^\ell, \mathbf{y}^\ell, \boldsymbol{\xi}) \geq S(\theta^\ell, \mathbf{x}^\ell, \mathbf{y}^\ell, \boldsymbol{\xi}^*) \geq \delta/2$$

for all ℓ sufficiently large. By taking limits and exploiting the continuity of S , we conclude that

$$S(\hat{\theta}, \hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\boldsymbol{\xi}}) \geq S(\hat{\theta}, \hat{\mathbf{x}}, \hat{\mathbf{y}}, \boldsymbol{\xi}^*) \geq \delta/2.$$

Note, however, that $S(\theta^{\ell+1}, \mathbf{x}^{\ell+1}, \mathbf{y}^{\ell+1}, \boldsymbol{\xi}^\ell) \leq 0$ since $\boldsymbol{\xi}^\ell \in \Xi_k^{\ell+1}$ for some $k \in \mathcal{K}$. Since the sequence $(\theta^{\ell+1}, \mathbf{x}^{\ell+1}, \mathbf{y}^{\ell+1})$ also converges to $(\hat{\theta}, \hat{\mathbf{x}}, \hat{\mathbf{y}})$ and $\boldsymbol{\xi}^\ell$ converges to $\hat{\boldsymbol{\xi}}$, we thus conclude that $S(\hat{\theta}, \hat{\mathbf{x}}, \hat{\mathbf{y}}, \hat{\boldsymbol{\xi}}) \leq 0$, which yields the desired contradiction.

We now show that every accumulation point $(\hat{\theta}, \hat{\mathbf{x}}, \hat{\mathbf{y}})$ of the sequence $(\theta^\ell, \mathbf{x}^\ell, \mathbf{y}^\ell)$ corresponds to an *optimal* solution $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ of the K -adaptability problem (2) with objective value $\hat{\theta}$. Assume to the contrary that $(\hat{\theta}, \hat{\mathbf{x}}, \hat{\mathbf{y}})$ is feasible but suboptimal. Then there is a feasible solution $(\theta', \mathbf{x}', \mathbf{y}')$ with $\theta' < \hat{\theta}$ that either (i) is used to update the incumbent solution after finitely many iterations, or (ii) constitutes the accumulation point of another infinite sequence $(\theta'^\ell, \mathbf{x}'^\ell, \mathbf{y}'^\ell)$. In the first case, the objective values θ^ℓ of the scenario-based K -adaptability problems will be arbitrarily close to $\hat{\theta}$ for ℓ sufficiently large, which implies that the corresponding nodes τ^ℓ will be fathomed in Step 4. Similarly, in the second case the objective values θ^ℓ and θ'^ℓ of the scenario-based K -adaptability problems will be arbitrarily close to $\hat{\theta}$ and θ' , respectively, for ℓ sufficiently large. Since $\theta' < \hat{\theta}$, the algorithm will fathom the tree nodes corresponding to the sequence $(\theta^\ell, \mathbf{x}^\ell, \mathbf{y}^\ell)$ in Step 4. The result now follows since both cases contradict the assumption that $(\hat{\theta}, \hat{\mathbf{x}}, \hat{\mathbf{y}})$ is an accumulation point. \square

Theorem 2 guarantees that after sufficiently many iterations of the algorithm, our scheme generates feasible solutions that are close to an optimal solution of the K -adaptability problem (2). In general, our algorithm may not converge after finitely many iterations. In the following, we discuss a class of problem instances for which finite convergence is guaranteed.

Theorem 3 (Finite Convergence). *The branch-and-bound scheme terminates after finitely many iterations, if \mathcal{Y} has finite cardinality and only the objective function in problem (2) is uncertain.*

Proof. If only the objective function in the K -adaptability problem (2) is uncertain, then the corresponding semi-infinite disjunctive program (5) can be written as

$$\begin{aligned}
& \text{minimize} && \theta \\
& \text{subject to} && \theta \in \mathbb{R}, \mathbf{x} \in \mathcal{X}, \mathbf{y} \in \mathcal{Y}^K \\
& && \mathbf{T}\mathbf{x} + \mathbf{W}\mathbf{y}_k \leq \mathbf{h} \quad \forall k \in \mathcal{K} \\
& && \bigvee_{k \in \mathcal{K}} \left[\mathbf{c}^\top \mathbf{x} + \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_k \leq \theta \right] \quad \forall \boldsymbol{\xi} \in \Xi,
\end{aligned}$$

see Observation 2. Thus, the scenario-based K -adaptability problem (6) becomes

$$\begin{aligned}
\mathcal{M}(\Xi_1, \dots, \Xi_K) = & \text{minimize} && \theta \\
& \text{subject to} && \theta \in \mathbb{R}, \mathbf{x} \in \mathcal{X}, \mathbf{y} \in \mathcal{Y}^K \\
& && \mathbf{T}\mathbf{x} + \mathbf{W}\mathbf{y}_k \leq \mathbf{h} \quad \forall k \in \mathcal{K} \\
& && \mathbf{c}^\top \mathbf{x} + \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_k \leq \theta \quad \forall \boldsymbol{\xi} \in \Xi_k, \forall k \in \mathcal{K},
\end{aligned}$$

and the separation problem (7) can be written as

$$\begin{aligned}
\mathcal{S}(\theta, \mathbf{x}, \mathbf{y}) &= \max_{\boldsymbol{\xi} \in \Xi} \min_{k \in \mathcal{K}} \left\{ \mathbf{c}^\top \mathbf{x} + \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_k - \theta \right\} \\
&= \mathbf{c}^\top \mathbf{x} - \theta + \max_{\boldsymbol{\xi} \in \Xi} \min_{k \in \mathcal{K}} \left\{ \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_k \right\}.
\end{aligned}$$

We now show that if \mathcal{Y} has finite cardinality, then our branch-and-bound algorithm terminates after finitely many iterations. To this end, assume that this is not the case, and let τ^ℓ , $\ell = 0, 1, \dots$ be some rooted branch of the tree with infinite length. We denote by $(\theta^\ell, \mathbf{x}^\ell, \mathbf{y}^\ell)$ and $(\zeta^\ell, \boldsymbol{\xi}^\ell, \mathbf{z}^\ell)$ the corresponding sequences of optimal solutions to the master and the separation problem, respectively. Since \mathcal{Y} has finite cardinality, we must have $\mathbf{y}^{\ell_1} = \mathbf{y}^{\ell_2}$ for some $\ell_1 < \ell_2$.

The solution $(\theta^{\ell_2}, \mathbf{x}^{\ell_2}, \mathbf{y}^{\ell_2})$ satisfies $\mathcal{S}(\theta^{\ell_2}, \mathbf{x}^{\ell_2}, \mathbf{y}^{\ell_2}) > 0$ since τ^ℓ , $\ell = 0, 1, \dots$, is a branch of infinite length. Since $\mathbf{y}^{\ell_2} = \mathbf{y}^{\ell_1}$, we thus conclude that

$$\mathbf{c}^\top \mathbf{x}^{\ell_2} - \theta^{\ell_2} + \max_{\boldsymbol{\xi} \in \Xi} \min_{k \in \mathcal{K}} \left\{ \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_k^{\ell_1} \right\} > 0.$$

Since $\boldsymbol{\xi}^{\ell_1}$ is optimal in the separation problem $\mathcal{S}(\theta^{\ell_1}, \mathbf{x}^{\ell_1}, \mathbf{y}^{\ell_1})$ and $\mathcal{S}(\theta^{\ell_1}, \mathbf{x}^{\ell_1}, \mathbf{y}^{\ell_1}, \boldsymbol{\xi}^{\ell_1}) > 0$, we have

$$\mathbf{c}^\top \mathbf{x}^{\ell_2} - \theta^{\ell_2} + \min_{k \in \mathcal{K}} \left\{ \mathbf{d}(\boldsymbol{\xi}^{\ell_1})^\top \mathbf{y}_k^{\ell_1} \right\} = \mathbf{c}^\top \mathbf{x}^{\ell_2} - \theta^{\ell_2} + \max_{\boldsymbol{\xi} \in \Xi} \min_{k \in \mathcal{K}} \left\{ \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_k^{\ell_1} \right\} > 0.$$

However, since the node $\tau^{\ell_2} = (\Xi_1^{\ell_2}, \dots, \Xi_K^{\ell_2})$ is a descendant of the node $\tau^{\ell_1} = (\Xi_1^{\ell_1}, \dots, \Xi_K^{\ell_1})$, we must have $\boldsymbol{\xi}^{\ell_1} \in \Xi_k^{\ell_2}$ for some $k \in \mathcal{K}$. This, along with the fact that $(\theta^{\ell_2}, \mathbf{x}^{\ell_2}, \mathbf{y}^{\ell_2})$ is a feasible

solution to the master problem $\mathcal{M}(\Xi_1^{\ell_2}, \dots, \Xi_K^{\ell_2})$ and that $\mathbf{y}^{\ell_2} = \mathbf{y}^{\ell_1}$, implies that

$$\mathbf{c}^\top \mathbf{x}^{\ell_2} - \theta^{\ell_2} + \min_{k \in \mathcal{K}} \left\{ \mathbf{d}(\boldsymbol{\xi}^{\ell_1})^\top \mathbf{y}_k^{\ell_1} \right\} \leq 0.$$

This yields the desired contradiction and proves the theorem. \square

We note that the assumption of deterministic constraints is critical in the previous statement.

Example 2. Consider the following instance of the K -adaptability problem (2):

$$\inf_{y_1, y_2 \in \{0,1\}} \sup_{\xi \in [0,1]} \inf_{k \in \{1,2\}} \{(\xi - 1)(1 - 2y_k) : y_k \geq \xi\}$$

On this instance, our branch-and-bound algorithm generates a tree in which all branches have finite length, except (up to permutations) the sequence of nodes $\tau^\ell = (\Xi_1^\ell, \Xi_2^\ell)$, where $(\Xi_1^0, \Xi_2^0) = (\emptyset, \emptyset)$ and $(\Xi_1^\ell, \Xi_2^\ell) = (\{\xi^0 2^{-i} : i = 0, 1, \dots, \ell - 1\}, \{0\})$, $\ell > 0$, for some $\xi^0 \in (0, 1]$. For the node τ^ℓ , $\ell > 1$, the optimal solution of the scenario-based K -adaptability problem (6) is $(\theta^\ell, y_1^\ell, y_2^\ell) = (1 - \xi^0 2^{-\ell+1}, 1, 0)$, while the optimal solution of the separation problem is $(\zeta^\ell, \xi^\ell) = (\xi^0 2^{-\ell}, \xi^0 2^{-\ell})$. Thus, our branch-and-bound algorithm does not terminate after finitely many iterations.

We note that every practical implementation of our branch-and-bound scheme will fathom nodes in Step 5 whenever the objective value of the separation problem (7) is sufficiently close to zero (within some ϵ -tolerance). This ensures that the algorithm terminates in finite time in practice. Indeed, in Example 2 the objective value of the separation problem is less than ϵ for all nodes τ^ℓ with $\ell \geq \log_2(\xi^0 \epsilon^{-1})$, and our branch-and-bound algorithm will fathom the corresponding path of the tree after $\mathcal{O}(\log \epsilon^{-1})$ iterations if we seek ϵ -precision solutions.

3.3 Improvements to the Basic Algorithm

The algorithm of Section 3.1 serves as a blueprint that can be extended in multiple ways. In the following, we discuss three enhancements that improve the numerical performance of our algorithm.

Breaking Symmetry. For any feasible solution (\mathbf{x}, \mathbf{y}) of the K -adaptability problem (2), every solution $(\mathbf{x}, \mathbf{y}')$, where \mathbf{y}' is one of the $K!$ permutations of the second-stage policies $(\mathbf{y}_1, \dots, \mathbf{y}_K)$, is also feasible in (2) and attains the same objective value. This implies that our branch-and-bound tree is highly isomorphic since the scenario-based problems (6) and (8) are identical (up to

a permutation of the policies) across many nodes. We can reduce this undesirable symmetry by modifying Step 6 of our branch-and-bound scheme as follows:

6'. *Branch.* Let $K' = 1$ if $\Xi_1 = \dots = \Xi_K = \emptyset$ and let $K' = \min \left\{ K, 1 + \max_{k \in \mathcal{K}} \{k : \Xi_k \neq \emptyset\} \right\}$ otherwise. Instantiate K' new nodes $\tau_k = (\Xi_1, \dots, \Xi_k \cup \{\xi\}, \dots, \Xi_K)$, $k = 1, \dots, K'$. Set $\mathcal{N} \leftarrow \mathcal{N} \cup \{\tau_1, \dots, \tau_{K'}\}$ and go to Step 3.

Despite generating only a subset of the nodes that our original algorithm constructs, the modification above always maintains at least one of the $K!$ solutions symmetric to every feasible solution.

Integration into MILP Solvers. Step 4 of our algorithm solves the scenario-based problem (6) from scratch in every node, despite the fact that two successive problems along any branch of the branch-and-bound tree differ only by the addition of a few constraints. We can leverage this commonality if we integrate our branch-and-bound algorithm into the solution scheme of the MILP solver used for problem (6). In doing so, we can also exploit the advanced facilities commonly present in the state-of-the-art solvers such as warm-starts and cutting planes, among others.

In order to integrate our branch-and-bound algorithm into the solution scheme of the MILP solver, we initialize the solver with the scenario-based problem (6) corresponding to the root node τ^0 of our algorithm, see Step 1. The solver then proceeds to solve this problem using its own branch-and-bound procedure. Whenever the solver encounters an *integral solution* $(\theta, \mathbf{x}, \mathbf{y}) \in \mathbb{R} \times \mathcal{X} \times \mathcal{Y}^K$, we solve the associated separation problem (8). If $\mathcal{S}(\theta, \mathbf{x}, \mathbf{y}) > 0$, then we execute Step 6 of our algorithm through a *branch callback*: we report the K new branches to the solver, which will discard the current solution. Otherwise, if $\mathcal{S}(\theta, \mathbf{x}, \mathbf{y}) \leq 0$, then we do not create any new branches, and the solver will accept $(\theta, \mathbf{x}, \mathbf{y})$ as the new incumbent solution. This ensures that only those solutions which are feasible in problem (5) are accepted as incumbent solutions.

Whenever the solver encounters a *fractional solution*, it will by default branch on an integer variable that is fractional in the current solution. However, if $\mathcal{S}(\theta, \mathbf{x}, \mathbf{y}) > 0$, it is possible to override this strategy and instead execute Step 6 of our algorithm. In such cases, a heuristic rule can be used to decide whether to branch on integer variables or to branch as in Step 6. In our computational experience, a simple rule that alternates between the default branching rule of the solver and the one defined by Step 6 appears to perform well in practice.

3.4 Modification as a Heuristic Algorithm

Whenever the number of policies K is large, the solution of the scenario-based K -adaptability problem (6) can be time consuming. In such cases, only a limited number of nodes will be explored by the algorithm in a given amount of computation time, and the quality of the final incumbent solution may be poor. As a remedy, we can reduce the size and complexity of the scenario-based K -adaptability problem (6) by fixing some of its second-stage policies. In doing so, we obtain a heuristic variant of our algorithm that can scale to large values of K .

In our computational experience, a simple heuristic that sequentially solves the 1-, 2-, \dots , K -adaptability problems by fixing in each K -adaptability problem all but one of the second-stage policies, $\mathbf{y}_1, \dots, \mathbf{y}_{K-1}$, to their corresponding values in the $(K-1)$ -adaptability problem, performs well in practice. This heuristic is motivated by two observations. First, the resulting scenario-based K -adaptability problems (6) have the same size and complexity as the corresponding scenario-based 1-adaptability problems. Second, in our experiments on instances with uncertain objective coefficients \mathbf{d} , we often found that some optimal second-stage policies of the $(K-1)$ -adaptability problem also appear in the optimal solution of the K -adaptability problem. In fact, it can be shown that this heuristic can obtain K -adaptable solutions that improve upon 1-adaptable solutions only if the objective coefficients \mathbf{d} are affected by uncertainty.

4 Numerical Results

We now analyze the computational performance of our branch-and-bound scheme in a variety of problem instances from the literature. We consider a shortest path problem with uncertain arc weights (Section 4.1), a capital budgeting problem with uncertain cash flows (Section 4.2), a variant of the capital budgeting problem with the additional option to take loans (Section 4.3), a project management problem with uncertain task durations (Section 4.4), and a vehicle routing problem with uncertain travel times (Section 4.5). Of these, the first two problems involve only binary decisions, and they can therefore also be solved with the approach described in [25]. In these cases, we show that our solution scheme is highly competitive, and it frequently outperforms the approach of [25]. In contrast, the third and fourth problems also involve continuous decisions, and there is no existing solution approach for their associated K -adaptability problems. However, the

project management problem from Section 4.4 involves *only* continuous second-stage decisions, and therefore the corresponding two-stage robust optimization problem (1) can also be approximated using affine decision rules [3], which represent the most popular approach for such problems. In this case, we elucidate the benefits of K -adaptable constant and affine decisions over standard affine decision rules. Finally, the first and last problems involve only binary second-stage decisions and deterministic constraints, and they can therefore also be addressed with the heuristic approach described in [13]. In these cases, we show that the heuristic variant of our algorithm often outperforms the latter approach in terms of solution quality.

For each problem category, we investigate the tradeoffs between computational effort and improvement in objective value of the K -adaptability problem for increasing values of K . We demonstrate that (i) the K -adaptability problem can provide significant improvements over static robust optimization (which corresponds to the case $K = 1$), and that (ii) our solution scheme can quickly determine feasible solutions of high quality.

We implemented our branch-and-bound algorithm in C++ using the C callable library of CPLEX 12.7 [26]. We used a constraint feasibility tolerance of $\epsilon = 10^{-4}$ to accept any incumbent solutions, whereas all other solver options were kept at their default values. The experiments were conducted on a single core of an Intel Xeon 2.8GHz computer with 16GB RAM.

4.1 Shortest Paths

We consider the shortest path problem from [25]. Let $G = (V, A)$ be a directed graph with nodes $V = \{1, \dots, N\}$, arcs $A \subseteq V \times V$ and arc weights $d_{ij}(\boldsymbol{\xi}) = (1 + \xi_{ij}/2)d_{ij}^0$, $(i, j) \in A$. Here, $d_{ij}^0 \in \mathbb{R}_+$ represents the nominal weight of the arc $(i, j) \in A$ and ξ_{ij} denotes the uncertain deviation from the nominal weight. The realizations of the uncertain vector $\boldsymbol{\xi}$ are known to belong to the set

$$\Xi = \left\{ \boldsymbol{\xi} \in [0, 1]^{|A|} : \sum_{(i,j) \in A} \xi_{ij} \leq \Gamma \right\},$$

which stipulates that at most Γ arc weights may maximally deviate from their nominal values.

Let $s \in V$ and $t \in V$, $s \neq t$, denote the source and terminal nodes of G , respectively. The decision-maker aims to choose K paths from s to t here-and-now, i.e., before observing the actual arc weights, such that the worst-case weight of the shortest among the chosen paths is minimized.

This problem can be formulated as an instance of the K -adaptability problem (2):

$$\inf_{\mathbf{y} \in \mathcal{Y}^K} \sup_{\xi \in \Xi} \inf_{k \in \mathcal{K}} \mathbf{d}(\xi)^\top \mathbf{y}_k$$

Here, \mathcal{Y} denotes the set of all $s - t$ paths in G ; that is,

$$\mathcal{Y} = \left\{ \mathbf{y} \in \{0, 1\}^{|A|} : \sum_{(j,l) \in A} y_{jl} - \sum_{(i,j) \in A} y_{ij} \geq \mathbb{I}[j = s] - \mathbb{I}[j = t] \quad \forall j \in V \right\}.$$

Note that this problem only contains second-stage decisions and as such, the corresponding two-stage robust optimization problem (1) may be of limited interest in practice. Nevertheless, the K -adaptability problem (2) has important applications in logistics and disaster relief [25].

For each graph size $N \in \{20, 25, \dots, 50\}$, we randomly generate 100 problem instances as follows. We assign the coordinates $(u_i, v_i) \in \mathbb{R}^2$ to each node $i \in V$ uniformly at random from the square $[0, 10]^2$. The nominal weight of the arc $(i, j) \in A$ is defined to be the Euclidean distance between the nodes i and j ; that is, $d_{ij}^0 = \sqrt{(u_i - u_j)^2 + (v_i - v_j)^2}$. The source node s and the terminal node t are defined to be the nodes with the maximum Euclidean distance between them. The arc set A is obtained by removing from the set of all pairwise links the $\lfloor 0.7(N^2 - N) \rfloor$ connections with the largest nominal weights. We set the uncertainty budget to $\Gamma = 3$. Further details on the parameter settings can be found in [25].

Table 2 summarizes the numerical performance of our branch-and-bound scheme for $K \in \{2, 3, 4\}$. Table 2 indicates that our scheme is able to reliably compute optimal solutions for small values of N and K , while the average optimality gap for large values of N and K is less than 9%. The numerical performance is strongly affected by the value of K ; very few of the 4-adaptable instances are solved to optimality within the time limit. This decrease in tractability is partly explained in Figure 4, which shows the improvement in objective value of the K -adaptability problem over the static problem (where $K = 1$). Figure 4a shows that the computed 4-adaptable solutions are typically of high quality since they improve upon the static solutions by as much as 13% for large values of N . Moreover, Figure 4b shows that these solutions are obtained within 1 minute (on average), even for the largest instances. This indicates that the gaps in Table 2 are likely to be very conservative since the majority of computation time is spent on obtaining a certificate of optimality for these solutions.

Figure 5 illustrates the quality of the solutions obtained using the heuristic variant of our algorithm, described in Section 3.3, and contrasts it with the quality of the solutions obtained

Table 2. Results for the shortest path problem. For each value of K , the “Opt (#)” column reports the number of instances (out of 100) which were solved to optimality, while the “Time (s)” column reports the average time to solve these instances to optimality. For those instances which could not be solved to optimality within the time limit of 7,200s, the average gap $|(\text{ub} - \text{lb})/\text{ub}| \times 100\%$ between the global lower bound (lb) and global upper bound (ub) of the branch-and-bound tree is reported in the “Gap (%)” column.

N	$K = 2$			$K = 3$			$K = 4$		
	Opt (#)	Time (s)	Gap (%)	Opt (#)	Time (s)	Gap (%)	Opt (#)	Time (s)	Gap (%)
20	99	6	1.23	97	408	2.51	70	539	1.74
25	91	222	4.14	64	847	2.91	33	885	2.89
30	64	744	4.40	31	1,237	4.10	16	827	4.27
35	37	1,083	5.36	14	1,020	5.01	10	896	5.23
40	10	808	6.28	6	1,670	6.43	2	39	6.10
45	9	1,152	7.70	1	16	7.06	1	15	6.61
50	2	3,307	8.55	1	2,308	7.90	0	–	7.10




Figure 4. Results for the shortest path problem using the exact algorithm. The graphs show the average improvement $|(\theta_1 - \theta_K)/\theta_1| \times 100\%$ of the objective value of the K -adaptability problem (θ_K) over the static problem (θ_1). The left graph shows the improvement after 2 hours (for increasing N), while the right graph shows the time profile of the improvement of the incumbent solution in the first 60 seconds (for $N = 50$).

using the heuristic algorithm described in [13]. Figure 5 shows that, after just one minute of computation time, the 2-, 3- and 4-adaptable solutions obtained using our heuristic algorithm

are within 0.3% of known optimal solutions and about 2% better than those obtained using the heuristic algorithm described in [13], on average. The differences in the qualities of the 6-, 8- and 10-adaptable solutions are smaller. The figure also shows that the marginal gain in objective value decreases rapidly as we increase the number of policies K . Indeed, while the 2-adaptable solutions are about 8.3% better than the 1-adaptable (i.e., static) solutions, the 10-adaptable solutions are only about 0.1% better than the 8-adaptable solutions. This may be explained by the possibility that the objective values of the corresponding K -adaptable solutions are very close to the optimal value of the two-stage robust optimization problem (1).

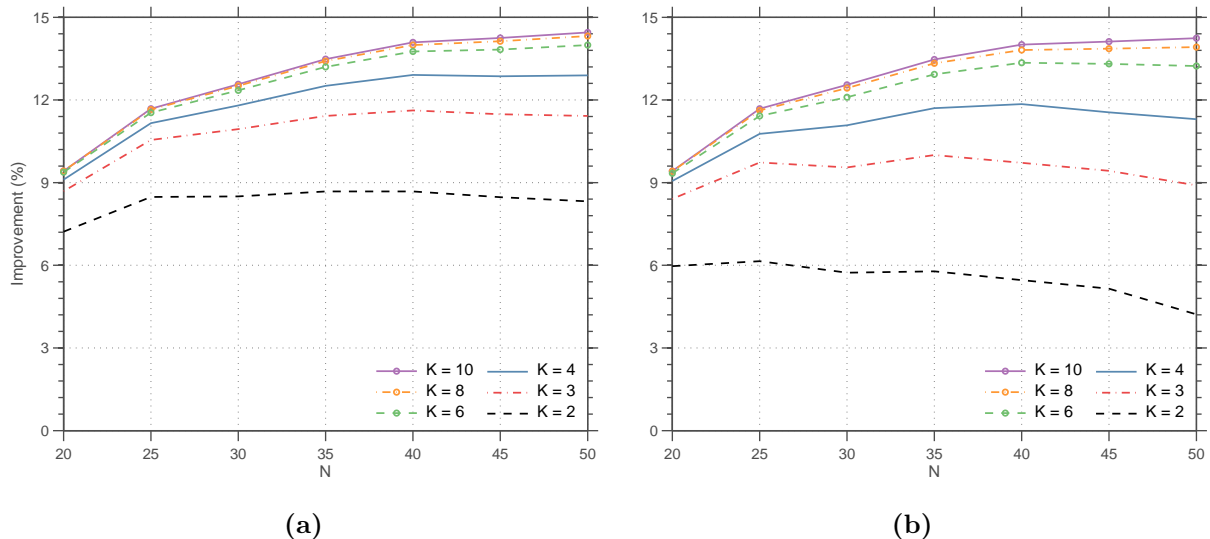


Figure 5. Results for the shortest path problem using the heuristic algorithm. The graphs show the average improvement after 1 minute obtained using the heuristic variant of our algorithm (left) and the heuristic algorithm described in [13] (right).

4.2 Capital Budgeting

We consider the capital budgeting problem from [25], where a company wishes to invest in a subset of N projects. Each project i has an uncertain cost $c_i(\boldsymbol{\xi})$ and an uncertain profit $r_i(\boldsymbol{\xi})$ that are governed by factor models of the form

$$c_i(\boldsymbol{\xi}) = \left(1 + \boldsymbol{\Phi}_i^\top \boldsymbol{\xi}/2\right) c_i^0 \quad \text{and} \quad r_i(\boldsymbol{\xi}) = \left(1 + \boldsymbol{\Psi}_i^\top \boldsymbol{\xi}/2\right) r_i^0 \quad \text{for } i = 1, \dots, N.$$

In these models, c_i^0 and r_i^0 represent the nominal cost and the nominal profit of project i , respectively, while $\boldsymbol{\Phi}_i^\top$ and $\boldsymbol{\Psi}_i^\top$ represent the i^{th} row vectors of the factor loading matrices $\boldsymbol{\Phi}, \boldsymbol{\Psi} \in \mathbb{R}^{N \times 4}$.

The realizations of the uncertain vector of risk factors $\boldsymbol{\xi}$ belong to the uncertainty set $\Xi = [-1, 1]^4$.

The company can invest in a project either before or after observing the risk factors $\boldsymbol{\xi}$. In the latter case, the company generates only a fraction κ of the profit (reflecting a penalty for postponement) but incurs the same cost as in the case of an early investment. Given an investment budget B , the capital budgeting problem can then be formulated as the following instance of the two-stage robust optimization problem (1):

$$\sup_{\mathbf{x} \in \mathcal{X}} \inf_{\boldsymbol{\xi} \in \Xi} \sup_{\mathbf{y} \in \mathcal{Y}} \left\{ \mathbf{r}(\boldsymbol{\xi})^\top (\mathbf{x} + \kappa \mathbf{y}) : \mathbf{c}(\boldsymbol{\xi})^\top (\mathbf{x} + \mathbf{y}) \leq B, \mathbf{x} + \mathbf{y} \leq \mathbf{e} \right\},$$

where $\mathcal{X} = \mathcal{Y} = \{0, 1\}^N$.

For our numerical experiments, we randomly generate 100 instances for each problem size $N \in \{5, 10, \dots, 30\}$ as follows. The nominal costs \mathbf{c}^0 are chosen uniformly at random from the hyperrectangle $[0, 10]^N$. We then set $\mathbf{r}^0 = \mathbf{c}^0/5$, $B = \mathbf{e}^\top \mathbf{c}^0/2$ and $\kappa = 0.8$. The rows of the factor loading matrices $\boldsymbol{\Phi}$ and $\boldsymbol{\Psi}$ are sampled uniformly from the unit simplex in \mathbb{R}^4 ; that is, the i^{th} row vector is sampled from $[0, 1]^4$ such that $\boldsymbol{\Phi}_i^\top \mathbf{e} = \boldsymbol{\Psi}_i^\top \mathbf{e} = 1$ is satisfied for all $i = 1, \dots, N$.

Table 3 summarizes the numerical performance of our branch-and-bound scheme for $K \in \{2, 3, 4\}$. Table 3 demonstrates that our branch-and-bound scheme performs very well for this problem class since the majority of instances is solved to optimality for $K \in \{2, 3\}$. Moreover, the optimality gaps for the unsolved instances are less than 4% for $K \in \{2, 3\}$ and less than 9% for $K = 4$ on average. Additionally, Figure 6 shows that even for the largest instances, high-quality incumbent solutions which significantly improve ($\approx 100\%$) upon the static robust solutions are obtained within 1 minute of computation time. Our results compare favorably with those of [25] as well as those of the partition-and-bound approach for the corresponding two-stage robust optimization problem presented in [6].

4.3 Capital Budgeting with Loans

We consider a generalization of the capital budgeting problem from Section 4.2 where the company can increase its investment budget by purchasing a loan from the bank at a unit cost of $\lambda > 0$ before the risk factors $\boldsymbol{\xi}$ are observed as well as purchasing a loan at a unit cost of $\mu\lambda$ (with $\mu > 1$) after the observation occurs. If the company does not purchase any loan, then the problem reduces to the one described in Section 4.2. Therefore, we expect the worst-case profits to be at least as large

Table 3. Results for the capital budgeting problem. The columns have the same interpretation as in Table 2.

N	K = 2			K = 3			K = 4		
	Opt (#)	Time (s)	Gap (%)	Opt (#)	Time (s)	Gap (%)	Opt (#)	Time (s)	Gap (%)
5	100	1	–	100	1	–	100	3	–
10	100	1	–	100	16	–	100	149	–
15	100	10	–	99	566	0.33	69	2,245	1.42
20	100	419	–	34	2,787	1.65	5	3,710	4.02
25	29	2,238	1.12	4	2,281	2.63	0	–	6.22
30	1	188	3.01	1	6,687	3.35	0	–	8.27

as in that setting. The generalized capital budgeting problem can be formulated as the following instance of problem (1):

$$\sup_{(x_0, \mathbf{x}) \in \mathcal{X}} \inf_{\xi \in \Xi} \sup_{(y_0, \mathbf{y}) \in \mathcal{Y}} \left\{ \mathbf{r}(\xi)^\top (\mathbf{x} + \kappa \mathbf{y}) - \lambda(x_0 + \mu y_0) : \begin{bmatrix} \mathbf{x} + \mathbf{y} \leq \mathbf{e} \\ \mathbf{c}(\xi)^\top \mathbf{x} \leq B + x_0 \\ \mathbf{c}(\xi)^\top (\mathbf{x} + \mathbf{y}) \leq B + x_0 + y_0 \end{bmatrix} \right\}$$

Here, $\mathcal{X} = \mathcal{Y} = \mathbb{R}_+ \times \{0, 1\}^N$. The constraint $\mathbf{c}(\xi)^\top \mathbf{x} \leq B + x_0$ ensures that the first-stage expenditures $\mathbf{c}(\xi)^\top \mathbf{x}$ are fully covered by the budget B as well as the loan x_0 taken here-and-now.

We consider problems with $N \in \{5, 10, \dots, 30\}$ projects. For each value of N , we solve the same 100 instances from Section 4.2 with $\lambda = 0.12$ and $\mu = 1.2$. Table 4 shows the computational performance of our branch-and-bound scheme for $K \in \{2, 3, 4\}$. As in the case of the problems discussed so far, the numerical tractability of our algorithm decreases as the value of K increases. However, a comparison of Tables 3 and 4 suggests that the numerical tractability is not significantly affected by the presence of the additional continuous variables x_0 and y_0 . Indeed, the majority of instances for $K = 2$ are solved to optimality and the average gap across all unsolved instances is less than 5% for $K = 3$ and less than 9% for $K = 4$. Figure 7 shows that the 4-adaptable solutions improve upon the static solutions by as much as 115% in the largest instances. Although not shown in the figure, a comparison of the objective values of the final incumbent solutions with those of the capital budgeting problem without loans (Section 4.2) reveals that for $N \geq 15$, the option to purchase loans has no effect on the worst-case profit of the static solution and results in less than 1% improvement in the worst-case profit of the 2-adaptable solution. Indeed, the option to

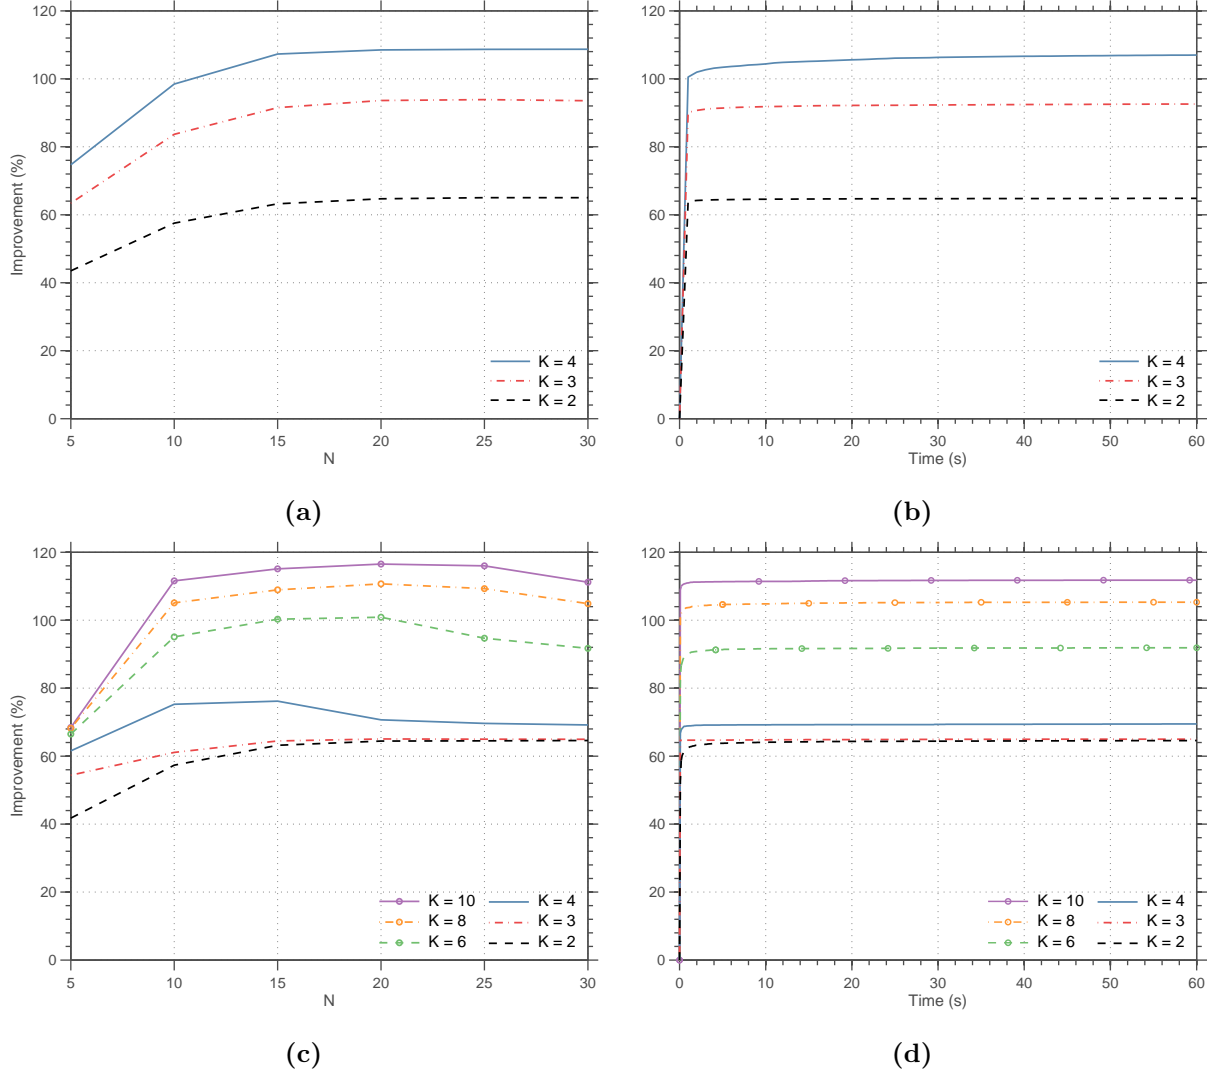


Figure 6. Results for the capital budgeting problem. The top (bottom) graphs are for the exact (heuristic) algorithm under a time limit of 2 hours (1 minute). The left graphs show the average improvement at the time limit (for increasing N), while the right graphs show the time profile of the improvement of the incumbent solution in the first 60 seconds (for $N = 30$).

purchase loans results in significantly better worst-case profits only if $K \geq 3$. The average relative gain in objective value is 4.3% for $K = 3$ and 5.9% for $K = 4$.

Table 4. Results for the capital budgeting problem with loans. The columns have the same interpretation as in Table 2.

N	$K = 2$			$K = 3$			$K = 4$		
	Opt (#)	Time (s)	Gap (%)	Opt (#)	Time (s)	Gap (%)	Opt (#)	Time (s)	Gap (%)
5	100	1	–	100	9	–	98	80	3.14
10	100	3	–	100	78	–	98	938	1.92
15	100	62	–	96	1,265	0.91	23	3,989	2.23
20	85	1,680	0.80	20	3,941	1.71	0	–	4.94
25	12	3,363	2.29	1	2,693	3.34	0	–	6.88
30	1	424	3.78	0	–	4.73	0	–	8.17

4.4 Project Management

We define a project as a directed acyclic graph $G = (V, A)$ whose nodes $V = \{1, \dots, N\}$ represent the tasks (e.g., ‘build foundations’ or ‘develop prototype’) and whose arcs $A \subseteq V \times V$ denote the temporal precedences, i.e., $(i, j) \in A$ implies that task j cannot be started before task i has been completed. We assume that each task $i \in V$ has an uncertain duration $d_i(\xi)$ that depends on the realization of an uncertain parameter vector $\xi \in \Xi$. Without loss of generality, we stipulate that the project graph G has the unique sink $N \in V$, and that the last task N has a duration of zero. This can always be achieved by introducing dummy nodes and/or arcs.

In the following, we want to calculate the worst-case makespan of the project, i.e, the smallest amount of time that is required to complete the project under the worst realization of the parameter vector $\xi \in \Xi$. This problem can be cast as the following instance of problem (1):

$$\sup_{\xi \in \Xi} \inf_{\mathbf{y} \in \mathcal{Y}} \{y_N : y_j - y_i \geq d_i(\xi) \quad \forall (i, j) \in A\}$$

Here $\mathcal{Y} = \mathbb{R}_+^N$, and y_i denotes the start time of task i , $i = 1, \dots, N$. This problem is known to be NP-hard [37, Theorem 2.1], and we will employ affine decision rules as well as K -adaptable constant and affine decisions to approximate the optimal value of this problem. Note that the problem does not contain any first-stage decisions, but such decisions could be readily included, for example, to allow for resource allocations that affect the task durations.

For our numerical experiments, we consider the instance class presented in [37, Example 2.2]. To this end, we set $N = 3m + 1$ and $A = \{(3l + 1, 3l + p), (3l + p, 3l + 4) : l = 0, \dots, m \text{ and } p = 2, 3\}$,




Figure 7. Results for the capital budgeting problem with loans. The graphs have the same interpretation as in Figure 6.

$d_{3l+2}(\xi) = \xi_{l+1}$ and $d_{3l+3}(\xi) = 1 - \xi_{l+1}$, $l = 0, \dots, m - 1$, as well as $d_{3l+1}(\xi) = 0$, $l = 0, \dots, m$. Figure 8 illustrates the project network corresponding to $m = 4$. Similar to [37], we consider the uncertainty set $\Xi = \{\xi \in \mathbb{R}_+^m : \|\xi - \mathbf{e}/2\|_1 \leq 1/2\}$.

We consider project networks of size $N \in \{3m + 1 : m = 3, 4, \dots, 8\}$. One can show that for each network size N , the optimal value of the corresponding static robust optimization problem as well as the affine decision rule problem is m , see [37, Example 2.2]. For $K \in \{2, 3, 4\}$, Figure 9 summarizes the computational performance of the branch-and-bound scheme and the improvement




Figure 8. Project network with $N = 3m + 1$ nodes for $m = 4$.

in objective value of the resulting piecewise constant and piecewise affine decision rules with K pieces over the corresponding 1-adaptable solutions. Figures 9a and 9b show that using only two pieces, piecewise constant decision rules can improve upon the affine approximation by more than 12%, while a piecewise affine decision rule can improve by more than 15%. Figures 9c and 9d show that piecewise constant decision rules require smaller computation times than piecewise affine decision rules. This is not surprising since piecewise constant decision rules are parameterized by $\mathcal{O}(KN)$ variables, whereas piecewise affine decision rules are parameterized by $\mathcal{O}(KN^2)$ variables.

4.5 Vehicle Routing

We consider the classical capacitated vehicle routing problem [22, 23, 35] defined on a complete, undirected graph $G = (V, E)$ with nodes $V = \{0, 1, \dots, N\}$ and edges $E = \{(i, j) \in V \times V : i < j\}$. Node 0 represents the unique depot, while each node $i \in V_C = \{1, \dots, N\}$ corresponds to a customer with demand $d_i \in \mathbb{R}_+$. The depot is equipped with M homogeneous vehicles; each vehicle has capacity C and it incurs an uncertain travel time $t_{ij}(\boldsymbol{\xi}) = (1 + \xi_{ij}/2)t_{ij}^0$ when it traverses the edge $(i, j) \in E$. Here, $t_{ij}^0 \in \mathbb{R}_+$ represents the nominal travel time along the edge $(i, j) \in E$, while ξ_{ij} denotes the uncertain deviation from the nominal value. Similar to the shortest path problem from Section 4.1, the realizations of the uncertain vector $\boldsymbol{\xi}$ are known to belong to the set

$$\Xi = \left\{ \boldsymbol{\xi} \in [0, 1]^{|E|} : \sum_{(i,j) \in E} \xi_{ij} \leq \Gamma \right\},$$

which stipulates that at most Γ travel times may maximally deviate from their nominal values.

A *route plan* (R_1, \dots, R_M) corresponds to a partition of the customer set V_C into M vehicle routes, $R_m = (R_{m,1}, \dots, R_{m,N_m})$, where $R_{m,l}$ represents the l^{th} customer and N_m the number of customers served by the m^{th} vehicle. This route plan is feasible if the total demand served

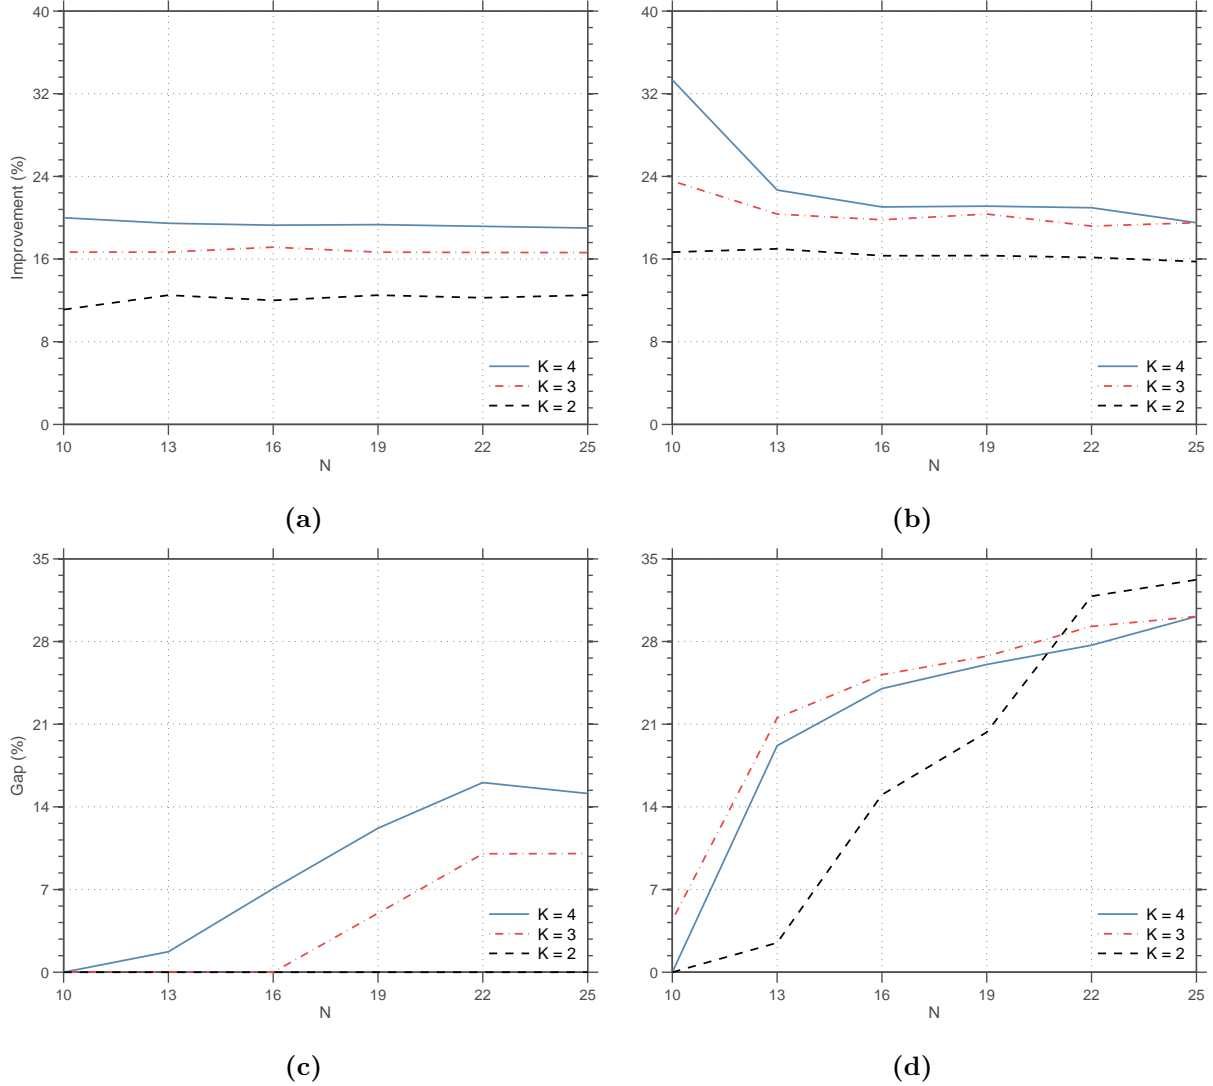


Figure 9. Results for the project management problem. The left (right) graphs show results for the piecewise constant (affine) K -adaptability problems for increasing values of N . Graphs (a) and (b) depict improvements in objective value, while graphs (c) and (d) show optimality gaps after 2 hours. The y -axes in graphs (a) and (b) have the same interpretation as those in Figure 4 while the y -axes in graphs (c) and (d) have the same interpretation as the column “Gap (%)” in Table 2.

on each route is less than the vehicle capacity; that is, if $\sum_{l=1}^{N_m} d_{R_{m,l}} \leq C$ is satisfied for all $m \in \{1, \dots, M\}$. The total travel time of a feasible route plan under the uncertainty realization ξ is given by $\sum_{m=1}^M \sum_{l=0}^{N_m} t_{R_{m,l}R_{m,l+1}}(\xi)$, where we define $R_{m,0} = R_{m,N_m+1} = 0$; that is, each vehicle starts and ends at the depot. The decision-maker aims to choose K route plans here-and-now, i.e.,

before observing the actual travel times, such that the worst-case total travel time of the shortest among the chosen route plans is minimized. This problem can be formulated as an instance of the K -adaptability problem (2):

$$\inf_{\mathbf{y} \in \mathcal{Y}^K} \sup_{\xi \in \Xi} \inf_{k \in \mathcal{K}} \mathbf{t}(\xi)^\top \mathbf{y}_k$$

Here, \mathcal{Y} denotes the set of all feasible route plans in G ; that is,

$$\mathcal{Y} = \left\{ \mathbf{y} \in \mathbb{Z}_+^{|E|} : \begin{array}{l} 0 \leq y_{ij} \leq 1 \quad \forall (i, j) \in E : i, j \in V_C, \\ \sum_{j \in V_C} y_{0j} = 2M, \\ \sum_{j \in V : (i, j) \in E} y_{ij} = 2 \quad \forall i \in V_C, \\ \sum_{(i, j) \in E : i, j \in U} y_{ij} \leq |U| - \left\lceil \frac{1}{D} \sum_{i \in U} d_i \right\rceil \quad \forall U \subseteq V_C \end{array} \right\}.$$

Similar to the shortest path problem, the K -adaptability formulation of the vehicle routing problem only contains second-stage decisions, and as such, the corresponding two-stage robust optimization problem (1) is of limited interest in practice. However, the K -adaptability problem (2) has important applications in logistics enterprises, where the time available between observing the travel times in a road network and determining the route plan is limited, or because the drivers must be trained to a small set of route plans that are to be executed daily over the course of a year.

We note that the set \mathcal{Y} represents the so-called *two-index vehicle flow* formulation of the vehicle routing problem, in which the first equation ensures that M vehicles are used; the second set of equations ensures that each customer is visited by exactly one vehicle; while the third set of inequalities ensure that there are no *subtours* disconnected from the depot and that all vehicle capacities are respected. This formulation is known to be extremely challenging to solve because it consists of an exponential number of inequalities. For $K > 1$, the corresponding K -adaptability problem is naturally even more challenging and it is practically intractable to solve it using the approach described in [25]. In contrast, the heuristic variant of our algorithm described in Section 3.3 as well as the heuristic approach of [13] only require the solution of vehicle routing subproblems that are of similar complexity as the associated 1-adaptability problems. Therefore, in the following, we only present results using these algorithms. In both cases, we solved all vehicle routing subproblems using the branch-and-cut algorithm described in [30].

For our numerical experiments, we consider all 49 instances from [30] with $N \leq 50$, which are

commonly used to benchmark vehicle routing algorithms. We set an overall time limit of 2 hours. For the heuristic variant of our algorithm, we further set a time limit of 10 minutes per vehicle routing subproblem. We note that the heuristic of [13] requires the successful termination of an expensive preprocessing step to determine good K -adaptable solutions. Therefore, to prevent bias in favor of our algorithm, Figure 10 compares the two algorithms only across the 39 instances for which this step terminated successfully. The figure shows that when the number of policies is small, the K -adaptable solutions obtained using our algorithm are about 1% better than those obtained using the heuristic algorithm of [13]. Moreover, the differences in their objective values are relatively higher for larger instances.

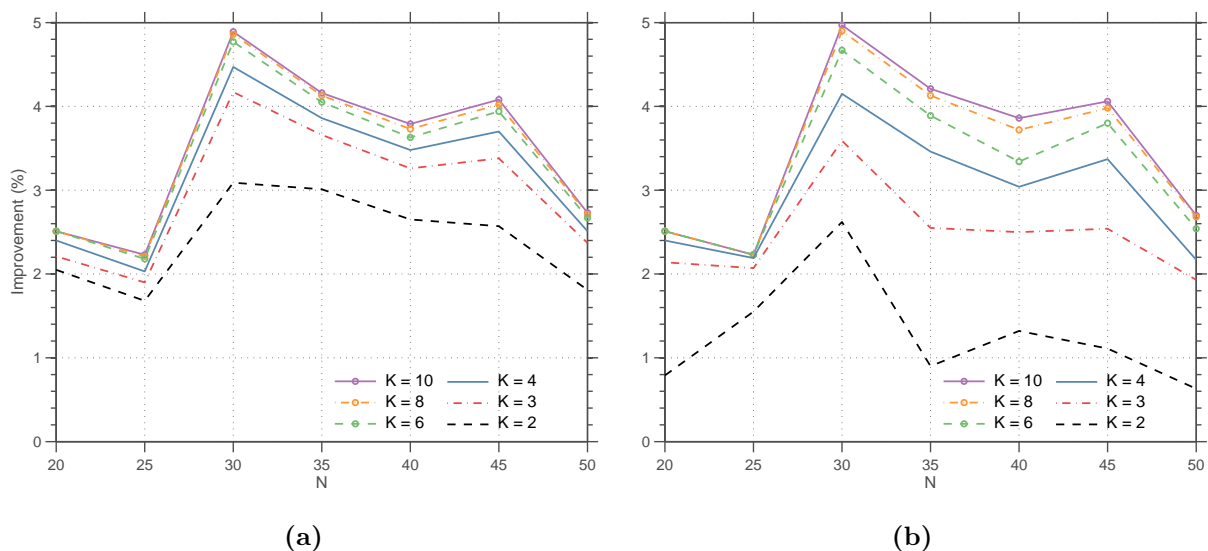


Figure 10. Results for the vehicle routing problem. The graphs show the average improvement after 2 hours obtained using the heuristic variant of our algorithm (left) and the heuristic algorithm in [13] (right).

5 Conclusions

In contrast to single-stage robust optimization problems, which are typically solved via monolithic reformulations, there is growing evidence that two-stage and multi-stage robust optimization problems are best solved algorithmically [6, 7, 8, 32, 39]. Our findings in this paper appear to confirm this observation, as our proposed branch-and-bound algorithm compares favorably with the reformulations proposed in [25]. In terms of modeling flexibility, our algorithm can accommodate mixed

continuous and discrete decisions in both stages, can incorporate discrete uncertainty, and allows us to model flexible piecewise affine decision rules. At the same time, our numerical results indicate that the algorithmic approach is highly competitive in terms of computational performance as well. From a practical viewpoint, a notable feature of our algorithm is that it admits a lightweight implementation by integrating it into the branch-and-bound schemes of commercial solvers via branch callbacks, while allowing easy modification as a heuristic for large-scale instances.

Our results open up multiple fruitful avenues for future research. The scope of the presented algorithm can be broadened further by generalizing it to two-stage *distributionally* robust optimization problems, where the uncertain problem parameters are modeled as random variables following a probability distribution that is only partially known. More ambitiously, it would be instructive to explore how the concept of K -adaptability, as well as our proposed branch-and-bound algorithm, extend to dynamic robust optimization problems with more than two stages.

Acknowledgments

Anirudh Subramanyam and Chrysanthos E. Gounaris gratefully acknowledge support from the National Science Foundation, award number CMMI-1434682. Wolfram Wiesemann gratefully acknowledges funding from the EPSRC grants EP/M028240/1, EP/M027856/1 and EP/N020030/1. Anirudh Subramanyam also acknowledges support from the John and Claire Bertucci Graduate Fellowship program at Carnegie Mellon University.

References

- [1] J. Ayoub and M. Poss. Decomposition for adjustable robust linear optimization subject to uncertainty polytope. *Computational Management Science*, 13(2):219–239, 2016.
- [2] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust Optimization*. Princeton University Press, 2009.
- [3] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming*, 99(2):351–376, 2004.

- [4] D. Bertsimas, D. B. Brown, and C. Caramanis. Theory and applications of robust optimization. *SIAM Review*, 53(3):464–501, 2011.
- [5] D. Bertsimas and C. Caramanis. Finite adaptability in multistage linear optimization. *IEEE Transactions on Automatic Control*, 55(12):2751–2766, 2010.
- [6] D. Bertsimas and I. Dunning. Multistage robust mixed integer optimization with adaptive partitions. *Operations Research*, 64(4):980–998, 2016.
- [7] D. Bertsimas and A. Georghiou. Design of near optimal decision rules in multistage adaptive mixed-integer optimization. *Operations Research*, 63(3):610–627, 2015.
- [8] D. Bertsimas and A. Georghiou. Binary decision rules for multistage adaptive mixed-integer optimization. *Mathematical Programming*, pages 1–39, 2017.
- [9] D. Bertsimas, V. Goyal, and X. A. Sun. A geometric characterization of the power of finite adaptability in multistage stochastic and adaptive optimization. *Mathematics of Operations Research*, 36(1):24–54, 2011.
- [10] D. Bertsimas, E. Litvinov, X. A. Sun, J. Zhao, and T. Zheng. Adaptive robust optimization for the security constrained unit commitment problem. *IEEE Transactions on Power Systems*, 28(1):52–63, 2013.
- [11] J. W. Blankenship and J. E. Falk. Infinitely constrained optimization problems. *Journal of Optimization Theory and Applications*, 19(2):261–281, 1976.
- [12] C. Buchheim and J. Kurtz. Min-max-min robustness: A new approach to combinatorial optimization under uncertainty based on multiple solutions. *Electronic Notes in Discrete Mathematics*, 52:45–52, 2016.
- [13] C. Buchheim and J. Kurtz. Min–max–min robust combinatorial optimization. *Mathematical Programming*, 163(1):1–23, 2017.
- [14] C. Buchheim and J. Kurtz. Complexity of minmaxmin robustness for combinatorial optimization under discrete uncertainty. *Discrete Optimization*, 28(1):1–15, 2018.

- [15] X. Chen and Y. Zhang. Uncertain linear programs: Extended affinely adjustable robust counterparts. *Operations Research*, 57(6):1469–1482, 2009.
- [16] V. Gabrel, C. Murat, and A. Thiele. Recent advances in robust optimization: An overview. *European Journal of Operational Research*, 235(3):471–483, 2014.
- [17] M. R. Garey and D. S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- [18] A. Georghiou, W. Wiesemann, and D. Kuhn. Generalized decision rule approximations for stochastic programming via liftings. *Mathematical Programming*, 152(1):301–338, 2015.
- [19] J. Goh and M. Sim. Distributionally robust optimization and its tractable approximations. *Operations Research*, 58(4):902–917, 2010.
- [20] B. L. Gorissen and D. den Hertog. Robust counterparts of inequalities containing sums of maxima of linear functions. *European Journal of Operational Research*, 227(1):30–43, 2013.
- [21] B. L. Gorissen, I. Yanıkoğlu, and D. den Hertog. A practical guide to robust optimization. *Omega*, 53:124–137, 2015.
- [22] C.E. Gounaris., P.P. Repoussis, C.D. Tarantilis, W. Wiesemann, and C.A. Floudas. An Adaptive Memory Programming Framework for the Robust Capacitated Vehicle Routing Problem. *Transportation Science*, 50(4):1239–1260, 2016.
- [23] C.E. Gounaris, W. Wiesemann, and C.A. Floudas. The Robust Capacitated Vehicle Routing Problem Under Demand Uncertainty. *Operations Research*, 61(3):677–693, 2013.
- [24] E. Guslitser. Uncertainty-immunized solutions in linear programming. Master’s thesis, Technion, Israeli Institute of Technology, 2002.
- [25] G. A. Hanasusanto, D. Kuhn, and W. Wiesemann. K -adaptability in two-stage robust binary programming. *Operations Research*, 63(4):877–891, 2015.
- [26] IBM. ILOG CPLEX Optimizer, 2018. [Online; accessed 27-July-2018].
- [27] R. Jiang, M. Zhang, G. Li, and Y. Guan. Two-stage robust power grid optimization problem. Available on Optimization Online, 2010.

- [28] J. E. Kelley. The cutting-plane method for solving convex programs. *Journal of the Society for Industrial and Applied Mathematics*, 8(4):703–712, 1960.
- [29] D. Kuhn, W. Wiesemann, and A. Georghiou. Primal and dual linear decision rules in stochastic and robust optimization. *Mathematical Programming*, 130(1):177–209, 2011.
- [30] J. Lygaard, A.N. Letchford, and R.W. Eglese. A new branch-and-cut algorithm for the capacitated vehicle routing problem. *Mathematical Programming*, 100(2):423–445, 2004.
- [31] A. Mutapcic and S. Boyd. Cutting-set methods for robust convex optimization with pessimizing oracles. *Optimization Methods and Software*, 24(3):381–406, 2009.
- [32] K. Postek and D. den Hertog. Multistage Adjustable Robust Mixed-Integer Optimization via Iterative Splitting of the Uncertainty Set. *INFORMS Journal on Computing*, 28(3):553–574, 2016.
- [33] R. T. Rockafellar and R. J.-B. Wets. *Variational Analysis*. Springer, 3rd edition, 2009.
- [34] A. Thiele, T. Terry, and M. Epelman. Robust linear optimization with recourse. Technical report, Lehigh University and University of Michigan, 2010.
- [35] P. Toth and D. Vigo. *Vehicle Routing: Problems, Methods, and Applications*. Society for Industrial and Applied Mathematics, 2nd edition, 2014.
- [36] P. Vayanos, D. Kuhn, and B. Rustem. Decision rules for information discovery in multi-stage stochastic programming. In *Proceedings of the 50th IEEE Conference on Decision and Control and European Control Conference*, pages 7368–7373, 2011.
- [37] W. Wiesemann, D. Kuhn, and B. Rustem. Robust resource allocations in temporal networks. *Mathematical Programming*, 135(1):437–471, 2012.
- [38] W. Wiesemann, D. Kuhn, and M. Sim. Distributionally robust convex optimization. *Operations Research*, 62(6):1358–1376, 2014.
- [39] B. Zeng and L. Zhao. Solving two-stage robust optimization problems using a column-and-constraint generation method. *Operations Research Letters*, 41(5):457–561, 2013.

- [40] C. Zhao, J. Wang, J.-P. Watson, and Y. Guan. Multi-stage robust unit commitment considering wind and demand response uncertainties. *IEEE Transactions on Power Systems*, 28(3):2708–2717, 2013.
- [41] L. Zhao and B. Zeng. An exact algorithm for two-stage robust optimization with mixed integer recourse problems. Technical report, University of South Florida, 2012.

Supplementary Material

A Analysis of the Two-Stage Robust Optimization Problem

To analyze problem (1), we equivalently rewrite it as

$$\inf_{\mathbf{x} \in \mathcal{X}} \mathbf{c}^\top \mathbf{x} + \mathcal{Q}(\mathbf{x}) \quad \text{with} \quad \mathcal{Q}(\mathbf{x}) = \sup_{\boldsymbol{\xi} \in \Xi} Q(\mathbf{x}, \boldsymbol{\xi}), \quad (1')$$

$$\text{where} \quad Q(\mathbf{x}, \boldsymbol{\xi}) = \inf_{\mathbf{y} \in \mathcal{Y}} \left\{ \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y} : \mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}(\boldsymbol{\xi})\mathbf{y} \leq \mathbf{h}(\boldsymbol{\xi}) \right\}$$

for $\mathcal{Q} : \mathcal{X} \mapsto \mathbb{R} \cup \{+\infty\}$ and $Q : \mathcal{X} \times \Xi \mapsto \mathbb{R} \cup \{+\infty\}$.

We first investigate whether the infima and the supremum in problem (1') are attained.

Proposition A.1 (Continuity). *Problem (1') satisfies the following properties.*

- (i) *The problem $Q(\mathbf{x}, \boldsymbol{\xi})$ attains its infimum, if it is feasible.*
- (ii) *The problem $\mathcal{Q}(\mathbf{x})$ attains its supremum, if only the objective function is uncertain. Otherwise, $\mathcal{Q}(\mathbf{x})$ does not necessarily attain its supremum, even if only the constraint right-hand sides are uncertain.*
- (iii) *The problem (1') attains its infimum, if it is feasible.*

Proof. The first statement holds since the problem $Q(\mathbf{x}, \boldsymbol{\xi})$ minimizes an affine function in \mathbf{y} over the intersection of the compact set \mathcal{Y} and the polyhedron $\{\mathbf{y} \in \mathbb{R}^{N_2} : \mathbf{W}(\boldsymbol{\xi})\mathbf{y} \leq \mathbf{h}(\boldsymbol{\xi}) - \mathbf{T}(\boldsymbol{\xi})\mathbf{x}\}$.

In view of the second statement, assume first that only the objective function in problem (1') is uncertain; that is, $\mathbf{T}(\boldsymbol{\xi}) = \mathbf{T}$, $\mathbf{W}(\boldsymbol{\xi}) = \mathbf{W}$ and $\mathbf{h}(\boldsymbol{\xi}) = \mathbf{h}$ for all $\boldsymbol{\xi} \in \Xi$. Then the inner problem $Q(\mathbf{x}, \boldsymbol{\xi})$ simplifies to $Q(\mathbf{x}, \boldsymbol{\xi}) = \inf\{\mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y} : \mathbf{y} \in \mathcal{Y}_{\mathbf{W}}\}$ with $\mathcal{Y}_{\mathbf{W}} = \{\mathbf{y} \in \mathcal{Y} : \mathbf{W}\mathbf{y} \leq \mathbf{h} - \mathbf{T}\mathbf{x}\}$. If $\mathcal{Y}_{\mathbf{W}} = \emptyset$, then $Q(\mathbf{x}, \boldsymbol{\xi}) = +\infty$ for all $\boldsymbol{\xi} \in \Xi$, and any $\boldsymbol{\xi} \in \Xi$ attains the supremum. Otherwise, we have $Q(\mathbf{x}, \boldsymbol{\xi}) = \inf\{\mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y} : \mathbf{y} \in \text{ext conv } \mathcal{Y}_{\mathbf{W}}\}$, where $\text{ext conv } \mathcal{Y}_{\mathbf{W}}$ denotes the set of (finitely many) extreme points of the convex hull of $\mathcal{Y}_{\mathbf{W}}$. We thus conclude from [33, Proposition 1.26] that $Q(\mathbf{x}, \boldsymbol{\xi})$ is upper semicontinuous in $\boldsymbol{\xi}$ for every fixed \mathbf{x} , and [33, Theorem 1.9] then implies that $\mathcal{Q}(\mathbf{x})$ attains its supremum since Ξ is a compact set.

Assume now that only the constraint right-hand sides in problem (1') are allowed to be uncertain, and consider the following instance of problem $\mathcal{Q}(\mathbf{x})$:

$$\sup_{\xi \in [0,1]} \inf_{\substack{\tau \in \mathbb{R}, \\ y \in \{0,1\}}} \{\tau - y : y \leq \xi \leq \tau\}$$

The inner minimization problem is optimized by $\tau^* = \xi$ and $y^* = \lfloor \xi \rfloor$. The supremum is 1, and it is approached by the sequence of feasible solutions $\xi \uparrow 1$. Note, however, that the limit point $\xi^* = 1$ of this sequence results in an objective function value of 0.

As for the third statement, we first note that the extended real-valued function

$$Q(\mathbf{x}, \boldsymbol{\xi}, \mathbf{y}) = \begin{cases} \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y} & \text{if } \mathbf{y} \in \mathcal{Y} \text{ and } \mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}(\boldsymbol{\xi})\mathbf{y} \leq \mathbf{h}(\boldsymbol{\xi}), \\ +\infty & \text{otherwise} \end{cases}$$

is lower semicontinuous in $(\mathbf{x}, \boldsymbol{\xi}, \mathbf{y})$ since the set $\{(\mathbf{x}, \boldsymbol{\xi}, \mathbf{y}) \in \mathbb{R}^{N_1} \times \mathbb{R}^{N_p} \times \mathcal{Y} : \mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}(\boldsymbol{\xi})\mathbf{y} \leq \mathbf{h}(\boldsymbol{\xi})\}$ is closed. From [33, Theorem 1.17] and [33, Proposition 1.26] we conclude that the lower semicontinuity is preserved by the partial minimization over \mathbf{y} and the partial maximization over $\boldsymbol{\xi}$. By [33, Theorem 1.9] and the fact that \mathcal{X} is compact we can then conclude that the problem (1') attains its infimum whenever it is feasible. \square

It is shown in [25, Example 1] that $Q(\mathbf{x})$ may not attain its supremum if both the objective function and the constraint right-hand sides in problem (1') are uncertain. Proposition A.1 (ii) strengthens this observation to instances of problem (1') where only the constraint right-hand sides are uncertain. We now consider the convexity properties of problem (1').

Proposition A.2 (Convexity). *Problem (1') satisfies the following properties.*

- (i) *The problem $Q(\mathbf{x}, \boldsymbol{\xi})$ is convex, if \mathcal{Y} is convex.*
- (ii) *The problem $Q(\mathbf{x})$ is convex, if Ξ is convex and only the objective function is uncertain, irrespective of \mathcal{Y} . Otherwise, $Q(\mathbf{x})$ is typically not convex, even if Ξ and \mathcal{Y} are convex and only the constraint right-hand sides are uncertain.*
- (iii) *The problem (1') is convex, if \mathcal{X} and \mathcal{Y} are convex, irrespective of Ξ .*

Proof. The first statement directly follows from the linearity of the objective function and the convexity of the feasible set.

In view of the second statement, assume first that Ξ is convex and only the objective function in problem (1') is uncertain; that is, $\mathbf{T}(\boldsymbol{\xi}) = \mathbf{T}$, $\mathbf{W}(\boldsymbol{\xi}) = \mathbf{W}$ and $\mathbf{h}(\boldsymbol{\xi}) = \mathbf{h}$ for all $\boldsymbol{\xi} \in \Xi$. Then the inner problem $Q(\mathbf{x}, \boldsymbol{\xi})$ simplifies to $Q(\mathbf{x}, \boldsymbol{\xi}) = \inf\{\mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y} : \mathbf{y} \in \mathcal{Y}_{\mathbf{W}}\}$ with $\mathcal{Y}_{\mathbf{W}} = \{\mathbf{y} \in \mathcal{Y} : \mathbf{W}\mathbf{y} \leq \mathbf{h} - \mathbf{T}\mathbf{x}\}$. Assume that $\mathcal{Y}_{\mathbf{W}} \neq \emptyset$; the other case is trivial. Then we have $Q(\mathbf{x}, \boldsymbol{\xi}) =$

$\inf\{\mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y} : \mathbf{y} \in \text{ext conv } \mathcal{Y}_{\mathbf{W}}\}$, where $\text{ext conv } \mathcal{Y}_{\mathbf{W}}$ denotes the set of (finitely many) extreme points of the convex hull of $\mathcal{Y}_{\mathbf{W}}$. We thus conclude from [33, Proposition 2.9] that $Q(\mathbf{x}, \boldsymbol{\xi})$ is concave in $\boldsymbol{\xi}$ for every fixed \mathbf{x} , which implies that $Q(\mathbf{x})$ is a convex optimization problem.

Assume now that the constraint right-hand sides in problem (1') are allowed to be uncertain, and consider the following instance of problem $Q(\mathbf{x})$:

$$\sup_{\xi \in [-1, 1]} \inf_{y \in \mathbb{R}} \{y : y \geq \xi, y \geq -\xi\}$$

Since the inner minimization problem is optimized by $y^* = |\xi|$, the problem maximizes the (convex) 1-norm of ξ over the interval $[-1, 1]$, which amounts to a non-convex optimization problem.

As for the third statement, assume that \mathcal{X} and \mathcal{Y} are convex, and consider the function

$$Q(\mathbf{x}, \boldsymbol{\xi}, \mathbf{y}) = \begin{cases} \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y} & \text{if } \mathbf{y} \in \mathcal{Y} \text{ and } \mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}(\boldsymbol{\xi})\mathbf{y} \leq \mathbf{h}(\boldsymbol{\xi}), \\ +\infty & \text{otherwise.} \end{cases}$$

The function $Q(\mathbf{x}, \boldsymbol{\xi}, \mathbf{y})$ is convex in (\mathbf{x}, \mathbf{y}) for every fixed $\boldsymbol{\xi} \in \Xi$. From [33, Proposition 2.22] and [33, Proposition 2.9] we conclude that the convexity is preserved by the partial minimization over \mathbf{y} and the partial maximization over $\boldsymbol{\xi}$. The problem (1') thus minimizes the sum of two convex functions $\mathbf{c}^\top \mathbf{x}$ and $Q(\mathbf{x})$ over the convex set \mathcal{X} , which is a convex optimization problem. \square

One readily verifies that the third statement in Proposition A.2 does not hold if \mathcal{Y} is not convex. We now investigate under which conditions we can solve problem (1') in polynomial time.

Proposition A.3 (Tractability). *Problem (1') satisfies the following properties.*

- (i) *The problem $Q(\mathbf{x}, \boldsymbol{\xi})$ can be solved in polynomial time, if \mathcal{Y} is convex.*
- (ii) *The problem $Q(\mathbf{x})$ can be solved in polynomial time, if Ξ and \mathcal{Y} are convex and only the objective function is uncertain. Otherwise, $Q(\mathbf{x})$ is strongly NP-hard, even if Ξ and \mathcal{Y} are convex and only the constraint right-hand sides are uncertain.*
- (iii) *The problem (1') can be solved in polynomial time, if \mathcal{X} , \mathcal{Y} and Ξ are convex and only the objective function is uncertain. Otherwise, the problem is strongly NP-hard, even if \mathcal{X} , \mathcal{Y} and Ξ are convex and only the constraint right-hand sides are uncertain.*

Proof. If \mathcal{Y} is convex, then the problem $Q(\mathbf{x}, \boldsymbol{\xi})$ amounts to a linear program that can be solved in polynomial time. This shows the first statement.

The first part of the second statement follows from the proof of the first part of the third statement below if we fix $\mathcal{X} = \{\mathbf{x}\}$ and $\mathbf{c} = \mathbf{0}$ in problem (1'). For the second part of the second statement, we recall the strongly NP-hard 0/1 Integer Programming (IP) feasibility problem [17]:

0/1 INTEGER PROGRAMMING FEASIBILITY.

Instance. Given are $\mathbf{A} \in \mathbb{Z}^{R \times N_p}$ and $\mathbf{b} \in \mathbb{Z}^R$.

Question. Is there a vector $\boldsymbol{\xi} \in \{0, 1\}^{N_p}$ such that $\mathbf{A}\boldsymbol{\xi} \leq \mathbf{b}$?

We show that the IP feasibility problem has an affirmative answer if and only if the problem

$$\sup \left\{ \inf \left\{ \mathbf{e}^\top \mathbf{y} : \mathbf{y} \in \mathbb{R}^{N_p}, \mathbf{y} \geq \boldsymbol{\xi} - \frac{1}{2}\mathbf{e}, \mathbf{y} \geq \frac{1}{2}\mathbf{e} - \boldsymbol{\xi} \right\} : \boldsymbol{\xi} \in [0, 1]^{N_p}, \mathbf{A}\boldsymbol{\xi} \leq \mathbf{b} \right\} \quad (9)$$

has an optimal value of $N_p/2$. Note that (9) can be interpreted as an instance of $Q(\mathbf{x})$. For any fixed $\boldsymbol{\xi}$, the infimum in (9) evaluates to $\|\boldsymbol{\xi} - \mathbf{e}/2\|_1$. Thus, the optimal value of (9) is equal to $N_p/2$ if and only if there is $\boldsymbol{\xi} \in [0, 1]^{N_p}$ satisfying $\|\boldsymbol{\xi} - \mathbf{e}/2\|_1 = N_p/2$ and $\mathbf{A}\boldsymbol{\xi} \leq \mathbf{b}$. The statement now follows from the fact that $\|\boldsymbol{\xi} - \mathbf{e}/2\|_1 = N_p/2$ if and only if $\boldsymbol{\xi} \in \{0, 1\}^{N_p}$.

In view of the first part of the third statement, the proof of Proposition A.4 below shows that, under the stated assumptions, problem (1') reduces to a single-stage robust optimization problem. Dualizing the inner maximization problem in that single-stage reformulation allows us to solve the overall problem in polynomial time as a linear program, see [2]. Finally, the second part of the third statement follows from the proof of the second part of the second statement if we amend problem (9) by appending an outer infimum over the singleton set $\mathcal{X} = \{1\}$. \square

The NP-hardness of the two-stage robust optimization problem (1') has previously been established for the case where \mathcal{X} , \mathcal{Y} and Ξ are convex and only the constraint right-hand sides are uncertain [24, Theorem 3.5]. We provide an alternative proof in Proposition A.3 (iii) to facilitate a self-contained comparison with the K -adaptability problem (2) in Proposition B.3 (ii) below. Moreover, it is shown in [24, Theorem 3.3] that problem (1') can be solved in polynomial time whenever \mathcal{X} and \mathcal{Y} are convex, \mathbf{d} and \mathbf{W} are deterministic and Ξ is described in terms of its extreme points. We close our analysis of the two-stage robust optimization problem (1') with a special case where it reduces to a single-stage robust optimization problem.

Proposition A.4 (Reduction to Static Problem). *The problem (1') reduces to a static robust optimization problem where \mathcal{Y} is replaced with its convex hull, if Ξ is convex and only the objective function is uncertain, irrespective of \mathcal{X} and \mathcal{Y} .*

Proof. Assume that Ξ is convex and that only the objective function in problem (1') is uncertain; that is, $\mathbf{T}(\boldsymbol{\xi}) = \mathbf{T}$, $\mathbf{W}(\boldsymbol{\xi}) = \mathbf{W}$ and $\mathbf{h}(\boldsymbol{\xi}) = \mathbf{h}$ for all $\boldsymbol{\xi} \in \Xi$. Problem (1') then simplifies to

$$\inf_{\mathbf{x} \in \mathcal{X}} \sup_{\boldsymbol{\xi} \in \Xi} \inf_{\mathbf{y} \in \text{conv } \mathcal{Y}} \left\{ \mathbf{c}^\top \mathbf{x} + \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y} : \mathbf{T}\mathbf{x} + \mathbf{W}\mathbf{y} \leq \mathbf{h} \right\},$$

where the replacement of the second-stage feasible region \mathcal{Y} with its convex hull, $\text{conv } \mathcal{Y}$, is justified since the inner minimization has a linear objective function. The classical minimax theorem now allows us to exchange the order of the inner two operators:

$$\inf_{\substack{\mathbf{x} \in \mathcal{X}, \\ \mathbf{y} \in \text{conv } \mathcal{Y}}} \left\{ \sup_{\boldsymbol{\xi} \in \Xi} \left\{ \mathbf{c}^\top \mathbf{x} + \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y} \right\} : \mathbf{T}\mathbf{x} + \mathbf{W}\mathbf{y} \leq \mathbf{h} \right\}.$$

This problem is readily recognized as a single-stage robust optimization problem.

We emphasize that the previous argument requires Ξ to be convex. Indeed, we have

$$-1 = \sup_{\xi \in \{-1,1\}} \inf_{y \in [-1,1]} \xi y \neq \inf_{y \in [-1,1]} \sup_{\xi \in \{-1,1\}} \xi y = 0,$$

and we cannot establish equivalence by replacing Ξ with $\text{conv } \Xi = [-1, 1]$ in the second optimization problem either. \square

A related result was established in [3, Theorem 2.1], where it is shown that the two-stage robust optimization problem (1') reduces to a single-stage robust optimization problem if \mathcal{X} , \mathcal{Y} and Ξ are convex and the uncertain parameters can be partitioned into subsets such that each constraint is only affected by the parameters of one such subset, and no two constraints are affected by parameters of the same subset.

B Analysis of the K -Adaptability Problem

To analyze problem (2), we equivalently rewrite it as

$$\begin{aligned} \inf_{\substack{\mathbf{x} \in \mathcal{X}, \\ \mathbf{y} \in \mathcal{Y}^k}} \mathbf{c}^\top \mathbf{x} + \mathcal{Q}(\mathbf{x}, \mathbf{y}) \quad \text{with} \quad \mathcal{Q}(\mathbf{x}, \mathbf{y}) = \sup_{\boldsymbol{\xi} \in \Xi} \mathcal{Q}(\mathbf{x}, \mathbf{y}, \boldsymbol{\xi}), \\ \text{where} \quad \mathcal{Q}(\mathbf{x}, \mathbf{y}, \boldsymbol{\xi}) = \inf_{k \in \mathcal{K}} \left\{ \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_k : \mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}(\boldsymbol{\xi})\mathbf{y}_k \leq \mathbf{h}(\boldsymbol{\xi}) \right\} \end{aligned} \tag{2'}$$

for $\mathcal{Q} : \mathcal{X} \times \mathcal{Y}^K \mapsto \mathbb{R} \cup \{+\infty\}$ and $Q : \mathcal{X} \times \mathcal{Y}^K \times \Xi \mapsto \mathbb{R} \cup \{+\infty\}$.

In analogy to Proposition A.1, we first investigate whether (2') attains its infima and supremum.

Proposition B.1 (Continuity). *Problem (2') satisfies the following properties.*

(i) *The problem $\mathcal{Q}(\mathbf{x}, \mathbf{y})$ attains its supremum, if only the objective function or only the constraints are uncertain. Otherwise, $\mathcal{Q}(\mathbf{x}, \mathbf{y})$ does not necessarily attain its supremum, even if the constraint left-hand sides are not uncertain.*

(ii) *The problem (2') attains its infimum, if it is feasible.*

Proof. In view of the first statement, assume that only the objective function in problem (2') is uncertain; that is, $\mathbf{T}(\boldsymbol{\xi}) = \mathbf{T}$, $\mathbf{W}(\boldsymbol{\xi}) = \mathbf{W}$ and $\mathbf{h}(\boldsymbol{\xi}) = \mathbf{h}$ for all $\boldsymbol{\xi} \in \Xi$. Then the inner problem $Q(\mathbf{x}, \mathbf{y}, \boldsymbol{\xi})$ simplifies to $Q(\mathbf{x}, \mathbf{y}, \boldsymbol{\xi}) = \inf\{\mathbf{y}_k^\top \mathbf{d}(\boldsymbol{\xi}) : k \in \mathcal{K}_{\mathbf{W}}\}$, where $\mathcal{K}_{\mathbf{W}} = \{k \in \mathcal{K} : \mathbf{W}\mathbf{y}_k \leq \mathbf{h} - \mathbf{T}\mathbf{x}\}$. If $\mathcal{K}_{\mathbf{W}} = \emptyset$, then $Q(\mathbf{x}, \mathbf{y}, \boldsymbol{\xi}) = +\infty$ for all $\boldsymbol{\xi} \in \Xi$, and any $\boldsymbol{\xi} \in \Xi$ attains the supremum. Otherwise, $Q(\mathbf{x}, \mathbf{y}, \boldsymbol{\xi})$ is readily verified to be continuous in $\boldsymbol{\xi}$ for every fixed \mathbf{x} and \mathbf{y} , and [33, Theorem 1.9] then implies that $\mathcal{Q}(\mathbf{x}, \mathbf{y})$ attains its supremum since Ξ is a compact set.

Assume now that only the constraints in problem (2') are uncertain. Then $Q(\mathbf{x}, \mathbf{y}, \boldsymbol{\xi})$ simplifies to $Q(\mathbf{x}, \mathbf{y}, \boldsymbol{\xi}) = \inf\{\mathbf{d}^\top \mathbf{y}_k : k \in \mathcal{K}, \boldsymbol{\xi} \in \Xi_k\}$, where $\Xi_k = \{\boldsymbol{\xi} \in \Xi : \mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}(\boldsymbol{\xi})\mathbf{y}_k \leq \mathbf{h}(\boldsymbol{\xi})\}$. If $\bigcup_{k \in \mathcal{K}} \Xi_k \neq \Xi$, then $Q(\mathbf{x}, \mathbf{y}, \boldsymbol{\xi}) = +\infty$ for any $\boldsymbol{\xi} \in \Xi \setminus \bigcup_{k \in \mathcal{K}} \Xi_k$ and $\mathcal{Q}(\mathbf{x}, \mathbf{y})$ attains its supremum. Otherwise, $\bigcup_{k \in \mathcal{K}} \Xi_k = \Xi$, and the supremum in $\mathcal{Q}(\mathbf{x}, \mathbf{y})$ is attained by any $\boldsymbol{\xi} \in \bigcap_{k \in \mathcal{K}^*} \Xi_k$, where

$$\mathcal{K}^* \in \arg \max_{\mathcal{K}' \subseteq \mathcal{K}} \left\{ \min_{k \in \mathcal{K}'} \mathbf{d}^\top \mathbf{y}_k : \bigcap_{k \in \mathcal{K}'} \Xi_k \neq \emptyset \right\}.$$

Finally, assume that both the objective function and the constraint right-hand sides in problem (2') are allowed to be uncertain, and consider the following instance of problem $\mathcal{Q}(\mathbf{x}, \mathbf{y})$:

$$\sup_{\xi \in [0,1]} \inf_{k \in \{1,2\}} \{\xi - y_k : y_k \leq \xi\} \quad \text{with } y_1 = 1 \text{ and } y_2 = 0.$$

The inner minimization problem is optimized by $k^* = 1$ if $\xi = 1$ and $k^* = 2$ otherwise. The supremum is 1, and it is approached by the sequence of feasible solutions $\xi \uparrow 1$. Note, however, that the limit point $\xi^* = 1$ of this sequence results in an objective function value of 0.

As for the second statement, we first note that each extended real-valued function

$$Q_k(\mathbf{x}, \mathbf{y}, \boldsymbol{\xi}) = \begin{cases} \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_k & \text{if } \mathbf{y}_k \in \mathcal{Y} \text{ and } \mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}(\boldsymbol{\xi})\mathbf{y}_k \leq \mathbf{h}(\boldsymbol{\xi}), \\ +\infty & \text{otherwise,} \end{cases} \quad k \in \mathcal{K},$$

is lower semicontinuous in $(\mathbf{x}, \mathbf{y}, \boldsymbol{\xi})$ since the sets $\{(\mathbf{x}, \mathbf{y}, \boldsymbol{\xi}) \in \mathbb{R}^{N_1} \times (\mathbb{R}^{N_2})^K \times \mathbb{R}^{N_p} : \mathbf{y}_k \in \mathcal{Y}, \mathbf{T}(\boldsymbol{\xi})\mathbf{x} + \mathbf{W}(\boldsymbol{\xi})\mathbf{y}_k \leq \mathbf{h}(\boldsymbol{\xi})\}$ are closed. From [33, Proposition 1.26] we conclude that the lower semicontinuity is preserved by the partial minimization over k and the partial maximization over $\boldsymbol{\xi}$. Due to [33, Theorem 1.9] the problem (2') then attains its infimum whenever it is feasible. \square

Similar to Proposition A.2, we now consider the convexity properties of problem (2').

Proposition B.2 (Convexity). *Problem (2') satisfies the following properties.*

- (i) *The problem $\mathcal{Q}(\mathbf{x}, \mathbf{y})$ is convex, if Ξ is convex and only the objective function is uncertain. Otherwise, $\mathcal{Q}(\mathbf{x}, \mathbf{y})$ is typically not convex, even if Ξ is convex and only the constraint right-hand sides are uncertain.*
- (ii) *Problem (2') is typically not convex, even if \mathcal{X} and \mathcal{Y} are convex and Ξ is a singleton.*

Proof. In view of the first statement, assume that Ξ is convex and that only the objective function in problem (2') is uncertain; that is, $\mathbf{T}(\boldsymbol{\xi}) = \mathbf{T}$, $\mathbf{W}(\boldsymbol{\xi}) = \mathbf{W}$ and $\mathbf{h}(\boldsymbol{\xi}) = \mathbf{h}$ for all $\boldsymbol{\xi} \in \Xi$. Then the inner problem $Q(\mathbf{x}, \mathbf{y}, \boldsymbol{\xi})$ simplifies to $Q(\mathbf{x}, \mathbf{y}, \boldsymbol{\xi}) = \inf\{\mathbf{y}_k^\top \mathbf{d}(\boldsymbol{\xi}) : k \in \mathcal{K}_{\mathbf{W}}\}$, where $\mathcal{K}_{\mathbf{W}} = \{k \in \mathcal{K} : \mathbf{W}\mathbf{y}_k \leq \mathbf{h} - \mathbf{T}\mathbf{x}\}$. Assume that $\mathcal{K}_{\mathbf{W}} \neq \emptyset$; the other case is trivial. Then $Q(\mathbf{x}, \mathbf{y}, \boldsymbol{\xi})$ is a piecewise-affine concave function in $\boldsymbol{\xi}$ for every fixed \mathbf{x} and \mathbf{y} . We thus conclude that $\mathcal{Q}(\mathbf{x}, \mathbf{y})$ is a convex optimization problem as it maximizes a concave function over a convex set.

Assume now that the constraint right-hand sides in problem (2') are allowed to be uncertain, and consider the following instance of problem $\mathcal{Q}(\mathbf{x}, \mathbf{y})$:

$$\sup_{\xi \in [-1, 1]} \inf_{k \in \mathcal{K}} \{y_k : y_k \geq \xi\} \quad \text{with } y_1 = 1 \text{ and } y_2 = 0.$$

The inner minimization problem is optimized by $k^* = 1$, if $\xi > 0$, and $k^* = 2$, otherwise. The outer maximization problem thus maximizes the non-convex function $\mathbb{I}[\xi > 0]$ over the interval $[-1, 1]$, which amounts to a non-convex optimization problem.

In view of the second statement, consider the following problem instance:

$$\inf_{y_1, y_2 \in [-1, 1]} \sup_{\xi \in \{1\}} \inf_{k \in \{1, 2\}} \xi y_k$$

The problem attains the objective value -1 for $(y_1, y_2) \in \{(-1, 1), (1, -1)\}$, but it attains the larger objective value of 0 for $(y_1, y_2) = 1/2 \cdot (-1, 1) + 1/2 \cdot (1, -1) = (0, 0)$. \square

In analogy to Proposition A.3, we now investigate under which conditions problem (2') is tractable.

Proposition B.3 (Tractability). *Problem (2') satisfies the following properties.*

- (i) *The problem $\mathcal{Q}(\mathbf{x}, \mathbf{y})$ can be solved in polynomial time, if Ξ is convex and only the objective function is uncertain. Otherwise, $\mathcal{Q}(\mathbf{x}, \mathbf{y})$ is strongly NP-hard, even if Ξ is convex and only the constraint right-hand sides are uncertain.*
- (ii) *The problem (2') can be solved in polynomial time, if \mathcal{X} , \mathcal{Y} and Ξ are convex and only the objective function is uncertain. Otherwise, the problem is strongly NP-hard, even if \mathcal{X} , \mathcal{Y} and Ξ are convex and only the constraint right-hand sides are uncertain.*

Proof. The first part of the first statement follows directly from [25, Observation 2].

In view of the second part of the first statement, we consider the following variant of the IP feasibility problem:

<p>APPROXIMATE 0/1 INTEGER PROGRAMMING FEASIBILITY.</p> <p>Instance. Given are $\mathbf{A} \in \mathbb{Z}^{R \times N_p}$, $\mathbf{b} \in \mathbb{Z}^R$.</p> <p>Question. Is there $\boldsymbol{\xi} \in ([0, \epsilon] \cup (1 - \epsilon, 1])^{N_p}$, $\epsilon = (\min_r \sum_q A_{rq})^{-1}$, such that $\mathbf{A}\boldsymbol{\xi} \leq \mathbf{b}$?</p>

It follows from [38, Lemma 2] that the approximate IP feasibility problem is strongly NP-hard. We claim that the approximate IP feasibility problem has an affirmative answer if and only if

$$\mathcal{Q}(\mathbf{x}, \mathbf{y}) = \sup_{\boldsymbol{\xi} \in \Xi} \inf_{k \in K} \left\{ y_k^0 : \underline{\mathbf{y}}_k \leq \boldsymbol{\xi} \leq \bar{\mathbf{y}}_k \right\} = 1, \quad (10)$$

where $\Xi = \{\boldsymbol{\xi} \in [0, 1]^{N_p} : \mathbf{A}\boldsymbol{\xi} \leq \mathbf{b}\}$ for (\mathbf{A}, \mathbf{b}) from the approximate IP feasibility instance, $K = N_p + 1$, $\mathbf{y}_k = (y_k^0, \underline{\mathbf{y}}_k, \bar{\mathbf{y}}_k) = (0, \epsilon \mathbf{e}_k, \mathbf{e} - \epsilon \mathbf{e}_k)$ for $k = 1, \dots, N_p$ and $\mathbf{y}_{N_p+1} = (y_{N_p+1}^0, \underline{\mathbf{y}}_{N_p+1}, \bar{\mathbf{y}}_{N_p+1}) = (1, \mathbf{0}, \mathbf{e})$. Note that we do not specify the first-stage decision \mathbf{x} as it is not required for the argument.

Assume first that the approximate IP feasibility problem has an affirmative answer, i.e., there is $\boldsymbol{\xi}^* \in ([0, \epsilon] \cup (1 - \epsilon, 1])^{N_p}$ such that $\mathbf{A}\boldsymbol{\xi}^* \leq \mathbf{b}$. Note that $\boldsymbol{\xi}^* \in \Xi$, but for every $k = 1, \dots, N_p$ we have $\boldsymbol{\xi}^* \notin [\underline{\mathbf{y}}_k, \bar{\mathbf{y}}_k]$. Since $\boldsymbol{\xi}^* \in [\mathbf{0}, \mathbf{e}]$, \mathbf{y}_{N_p+1} is the only feasible candidate policy, and the optimal value of problem (10) is indeed 1.

Assume now that the optimal value of problem (10) is 1. In that case, there is $\boldsymbol{\xi}^* \in \Xi$ such that $\boldsymbol{\xi}^* \in [\mathbf{0}, \mathbf{e}] \setminus \bigcup_{k=1}^{N_p} [\epsilon \mathbf{e}_k, \mathbf{e} - \epsilon \mathbf{e}_k]$, i.e., $\boldsymbol{\xi}^* \in ([0, \epsilon] \cup (1 - \epsilon, 1])^{N_p}$. Since $\mathbf{A}\boldsymbol{\xi}^* \leq \mathbf{b}$ by construction of Ξ , we conclude that the approximate IP feasibility problem has an affirmative answer.

As for the first part of the second statement, assume first that \mathcal{X} , \mathcal{Y} and Ξ are convex and that only the objective function in problem (2') is uncertain; that is, $\mathbf{T}(\boldsymbol{\xi}) = \mathbf{T}$, $\mathbf{W}(\boldsymbol{\xi}) = \mathbf{W}$ and $\mathbf{h}(\boldsymbol{\xi}) = \mathbf{h}$ for all $\boldsymbol{\xi} \in \Xi$. The classical minimax theorem then implies that

$$\inf_{\substack{\mathbf{x} \in \mathcal{X}, \\ \mathbf{y}_1 \in \mathcal{Y}}} \left\{ \sup_{\boldsymbol{\xi} \in \Xi} \left\{ \mathbf{c}^\top \mathbf{x} + \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_1 \right\} : \mathbf{T}\mathbf{x} + \mathbf{W}\mathbf{y}_1 \leq \mathbf{h} \right\} = \inf_{\mathbf{x} \in \mathcal{X}} \sup_{\boldsymbol{\xi} \in \Xi} \inf_{\mathbf{y} \in \mathcal{Y}} \left\{ \mathbf{c}^\top \mathbf{x} + \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y} : \mathbf{T}\mathbf{x} + \mathbf{W}\mathbf{y} \leq \mathbf{h} \right\},$$

i.e., the 1-adaptability problem achieves the same objective value as the two-stage robust optimization problem (1). Since the K -adaptability problem is bounded from above by the 1-adaptability problem and from below by the two-stage robust optimization problem, we thus conclude that, under the stated assumptions, the 1-adaptability problem coincides with the K -adaptability problem. The statement now follows from the fact that we can reformulate the 1-adaptability problem as a linear program by dualizing the inner maximization problem, see [2]. Note that the convexity of Ξ is needed in this argument since

$$-1 = \max_{\xi \in \{-1,1\}} \min_{y \in [-1,1]} \xi y \neq \min_{y \in [-1,1]} \max_{\xi \in \{-1,1\}} \xi y = 0.$$

A similar argument with $\Xi = [-1, 1]$ and $\mathcal{Y} = \{-1, 1\}$ shows that the convexity of \mathcal{Y} is needed, too.

Finally, in view of the second part of the second statement, consider the following problem:

$$\inf_{\mathbf{y} \in \mathcal{Y}^K} \sup_{\boldsymbol{\xi} \in \Xi} \inf_{k \in \mathcal{K}} \{0 : \mathbf{y}_k \leq \boldsymbol{\xi} \leq \mathbf{e} - \mathbf{y}_k\}, \quad (11)$$

where $K = N_p$, $\mathcal{Y} = \{\mathbf{y} \in \mathbb{R}_+^{N_p} : \mathbf{e}^\top \mathbf{y} = \epsilon\}$, $\epsilon < 1/2$, and $\Xi = \{\boldsymbol{\xi} \in [0, 1]^{N_p} : \mathbf{A}\boldsymbol{\xi} \leq \mathbf{b}\}$ for (\mathbf{A}, \mathbf{b}) from an approximate IP feasibility instance. We claim that the optimal value of this problem is $+\infty$, i.e., there is $\boldsymbol{\xi} \in \Xi$ for which no decision $\mathbf{y} \in \mathcal{Y}^K$ is feasible in the second stage, if and only if the approximate IP feasibility instance has an affirmative answer.

Assume first that the approximate IP feasibility problem has an affirmative answer. It then follows from [38, Lemma 2] that the *exact* IP feasibility problem also has an affirmative answer; that is, there is $\boldsymbol{\xi}^* \in [0, 1]^{N_p}$ such that $\mathbf{A}\boldsymbol{\xi}^* \leq \mathbf{b}$. Fix any feasible candidate policy $\mathbf{y}_k \in \mathcal{Y}$. By construction, there is $i \in \{1, \dots, N_p\}$ such that $y_{ki} > 0$. Thus, the candidate policy is feasible in the second stage under the parameter realization $\boldsymbol{\xi}^*$ only if $\xi_i^* \in [y_{ki}, 1 - y_{ki}] \subseteq (0, 1)$, which is not the case since $\xi_i^* \in \{0, 1\}$. We thus conclude that the optimal value of problem (11) is indeed $+\infty$ whenever the approximate IP feasibility problem has an affirmative answer.

Assume now that the optimal value of problem (11) is $+\infty$ and consider the K candidate policies $\mathbf{y}_k = \epsilon \mathbf{e}_k$, $k = 1, \dots, N_p$. Since $\mathbf{y} = (\mathbf{y}_1, \dots, \mathbf{y}_K) \in \mathcal{Y}^K$ and problem (11) evaluates to $+\infty$

under this choice of \mathbf{y} , there must be $\boldsymbol{\xi}^* \in \Xi$ such that $\boldsymbol{\xi}^* \in [\mathbf{0}, \mathbf{e}] \setminus \bigcup_{k=1}^{N_p} [\epsilon \mathbf{e}_k, \mathbf{e} - \epsilon \mathbf{e}_k]$; that is, $\boldsymbol{\xi}^* \in ([0, \epsilon) \cup (1 - \epsilon, 1])^{N_p}$. Since $\mathbf{A}\boldsymbol{\xi}^* \leq \mathbf{b}$ by construction of Ξ , we conclude that the approximate IP feasibility problem has an affirmative answer. \square

We note that $\mathcal{Q}(\mathbf{x}, \mathbf{y})$ can be solved in polynomial time if Ξ is convex and the number of policies K is fixed (even when the objective function, the constraint coefficients and the right-hand sides are uncertain), see [25, Corollary 1]. The NP-hardness of $\mathcal{Q}(\mathbf{x}, \mathbf{y})$ has previously been established under the more restrictive assumption that both the objective function and the constraint right-hand sides in problem (2') are uncertain, see [25, Theorem 3]. For the special case where $K = 2$, the K -adaptability problem with objective and constraint uncertainty can be solved in polynomial time if any of N_p , $\max\{N_1, N_2\}$ or L is fixed [5, Proposition 5], while the problem becomes NP-hard otherwise [5, Proposition 6]. Proposition B.3 (ii) provides an alternative proof of the NP-hardness of problem (2'), which facilitates a direct comparison to the two-stage robust optimization problem (1'), see Proposition A.3 (iii).

We also note that the K -adaptability problem (2') simplifies in the absence of first-stage decisions \mathbf{x} . For this case, it has been shown in [13, Theorem 2] that the problem can be solved in polynomial time, if only the objective function is uncertain, the deterministic second-stage problem is polynomial time solvable, the uncertainty set Ξ has a tractable representation, and if any number of policies $K > N_2$ is acceptable. The same problem becomes NP-hard, however, when the number of policies K is fixed [13, Corollary 3] or when the uncertainty sets are discrete [14, Corollaries 1–4].

Similar to Proposition A.4, we next consider when the K -adaptability problem (2') reduces to a single-stage robust optimization problem.

Proposition B.4 (Reduction to Static Problem). *The problem (2') reduces to a static robust optimization problem where \mathcal{Y} is replaced with its convex hull, if Ξ is convex, only the objective function is uncertain and $K > \min\{N_2, N_p\}$, irrespective of \mathcal{X} and \mathcal{Y} .*

Proof. Assume that Ξ is convex and that only the objective function in problem (2') is uncertain; that is, $\mathbf{T}(\boldsymbol{\xi}) = \mathbf{T}$, $\mathbf{W}(\boldsymbol{\xi}) = \mathbf{W}$ and $\mathbf{h}(\boldsymbol{\xi}) = \mathbf{h}$ for all $\boldsymbol{\xi} \in \Xi$. Problem (2') then simplifies to

$$\inf_{\substack{\mathbf{x} \in \mathcal{X}, \\ \mathbf{y} \in \mathcal{Y}^K}} \sup_{\boldsymbol{\xi} \in \Xi} \inf_{k \in \mathcal{K}} \left\{ \mathbf{c}^\top \mathbf{x} + \mathbf{d}(\boldsymbol{\xi})^\top \mathbf{y}_k : \mathbf{T}\mathbf{x} + \mathbf{W}\mathbf{y}_k \leq \mathbf{h} \right\}.$$

The inner discrete minimization in this problem can be replaced by a continuous minimization over

all convex combinations $\boldsymbol{\lambda} \in \Delta(\boldsymbol{x}, \boldsymbol{y}) = \{\boldsymbol{\lambda} \in \mathbb{R}_+^K : \mathbf{e}^\top \boldsymbol{\lambda} = 1, \lambda_k = 0 \text{ if } \mathbf{T}\boldsymbol{x} + \mathbf{W}\boldsymbol{y}_k \not\leq \boldsymbol{h}\}$:

$$\inf_{\substack{\boldsymbol{x} \in \mathcal{X}, \\ \boldsymbol{y} \in \mathcal{Y}^K}} \sup_{\boldsymbol{\xi} \in \Xi} \inf_{\boldsymbol{\lambda} \in \Delta(\boldsymbol{x}, \boldsymbol{y})} \left\{ \mathbf{c}^\top \boldsymbol{x} + \sum_{k \in \mathcal{K}} \lambda_k \mathbf{d}(\boldsymbol{\xi})^\top \boldsymbol{y}_k \right\}$$

The classical minimax theorem then allows us to exchange the order of the inner two operators:

$$\inf_{\substack{\boldsymbol{x} \in \mathcal{X}, \\ \boldsymbol{y} \in \mathcal{Y}^K}} \inf_{\boldsymbol{\lambda} \in \Delta(\boldsymbol{x}, \boldsymbol{y})} \sup_{\boldsymbol{\xi} \in \Xi} \left\{ \mathbf{c}^\top \boldsymbol{x} + \mathbf{d}(\boldsymbol{\xi})^\top \left[\sum_{k \in \mathcal{K}} \lambda_k \boldsymbol{y}_k \right] \right\} \quad (12)$$

If $K \geq N_2 + 1$, then we can use Carathéodory's theorem to replace the convex combinations of K candidate decisions $\boldsymbol{y}_k \in \mathcal{Y}$ with a single decision from the convex hull, $\text{conv } \mathcal{Y}$, resulting in

$$\inf_{\substack{\boldsymbol{x} \in \mathcal{X}, \\ \boldsymbol{y} \in \text{conv } \mathcal{Y}}} \left\{ \sup_{\boldsymbol{\xi} \in \Xi} \left\{ \mathbf{c}^\top \boldsymbol{x} + \mathbf{d}(\boldsymbol{\xi})^\top \boldsymbol{y} \right\} : \mathbf{T}\boldsymbol{x} + \mathbf{W}\boldsymbol{y} \leq \boldsymbol{h} \right\}, \quad (13)$$

which is readily recognized as a single-stage robust optimization problem. Similarly, problem (12) can be rewritten as

$$\inf_{\substack{\boldsymbol{x} \in \mathcal{X}, \\ \boldsymbol{y} \in \mathcal{Y}^K}} \inf_{\boldsymbol{\lambda} \in \Delta(\boldsymbol{x}, \boldsymbol{y})} \sup_{\boldsymbol{\xi} \in \Xi} \left\{ \mathbf{c}^\top \boldsymbol{x} + \sum_{k \in \mathcal{K}} \lambda_k \left[\mathbf{d}(\boldsymbol{\xi})^\top \boldsymbol{y}_k \right] \right\},$$

and if $K \geq N_p + 1$, then we can use Carathéodory's theorem to replace the convex combinations of K candidate decisions $\mathbf{d}(\boldsymbol{\xi})^\top \boldsymbol{y}_k \in \{\mathbf{d}(\boldsymbol{\xi})^\top \boldsymbol{y} : \boldsymbol{y} \in \mathcal{Y}\}$ with a single decision from the convex hull of their domain, $\text{conv } \{\mathbf{d}(\boldsymbol{\xi})^\top \boldsymbol{y} : \boldsymbol{y} \in \mathcal{Y}\} = \{\mathbf{d}(\boldsymbol{\xi})^\top \boldsymbol{y} : \boldsymbol{y} \in \text{conv } \mathcal{Y}\}$, resulting again in the single-stage robust optimization problem (13).

As in Proposition A.4, the previous arguments require Ξ to be convex. Indeed, we have

$$-1 = \sup_{\xi \in \{-1, 1\}} \inf_{k \in \{1, 2\}} \xi y_k \neq \inf_{y \in [-1, 1]} \sup_{\xi \in \{-1, 1\}} \xi y = 0 \quad \text{with } y_1 = -1 \text{ and } y_2 = 1,$$

and we cannot establish equivalence by replacing Ξ with $\text{conv } \Xi = [-1, 1]$ in the second optimization problem either. To see that the number of policies K must exceed $\min\{N_2, N_p\}$ in general in the previous arguments, we compare the two problems

$$\inf_{\boldsymbol{y} \in \{[-1, 1]^{N_2}\}^K} \sup_{\boldsymbol{\xi} \in [-1, 1]^{N_p}} \inf_{k \in \mathcal{K}} \xi_1 y_{k1} \quad \text{and} \quad \inf_{\boldsymbol{y} \in [-1, 1]^{N_2}} \sup_{\boldsymbol{\xi} \in [-1, 1]^{N_p}} \xi_1 y_1.$$

For every $N_2, N_p \geq 1$, the problem on the right attains its optimal value 0 at any $\boldsymbol{y} \in [-1, 1]^{N_2}$ with $y_1 = 0$. Similarly, for every $N_2, N_p \geq 1$, the problem on the left attains an optimal value of 0, if $K \geq 2$, and an optimal value of 1, if $K = 1$. We thus verify that $K > \min\{N_2, N_p\}$ is required for the optimal values to coincide by choosing $(N_2, N_p) \in \{(1, 2), (2, 1)\}$. \square

We close this section with several special cases where the optimal value of the K -adaptability problem (2') coincides with the optimal value of the two-stage robust optimization problem (1').

Proposition B.5 (Optimality). *The optimal values of problems (1') and (2') coincide if (i) \mathcal{Y} and Ξ are convex and only the objective function is uncertain, irrespective of \mathcal{X} and K , or if (ii) Ξ is convex, only the objective function is uncertain and $K > \min\{N_2, N_p\}$, irrespective of \mathcal{X} and \mathcal{Y} , or if (iii) \mathcal{Y} has a finite cardinality and $K \geq |\mathcal{Y}|$, irrespective of \mathcal{X} and Ξ . Otherwise, their optimal values may differ for any finite K , even if \mathcal{X} , \mathcal{Y} and Ξ are convex and only the constraint right-hand sides are uncertain.*

Proof. The fact that the two optimal values coincide under the first set of conditions follows from the proof of Proposition B.3 (ii), where we have shown that, under the stated assumptions, the optimal values of the 1-adaptability problem and the two-stage robust optimization problem (1') coincide. We have also shown there that the convexity of \mathcal{Y} and Ξ is crucial for the proof to hold.

The fact that the two optimal values coincide under the second set of conditions follows from Proposition B.4, which shows that, under the stated assumptions, the K -adaptability problem (2') is equivalent to

$$\inf_{\substack{\mathbf{x} \in \mathcal{X}, \\ \mathbf{y} \in \text{conv} \mathcal{Y}}} \left\{ \sup_{\xi \in \Xi} \left\{ \mathbf{c}^\top \mathbf{x} + \mathbf{d}(\xi)^\top \mathbf{y} \right\} : \mathbf{T}\mathbf{x} + \mathbf{W}\mathbf{y} \leq \mathbf{h} \right\}.$$

The classical minimax theorem, which is applicable since $\text{conv} \mathcal{Y}$ and Ξ are convex, then implies that this problem is equivalent to

$$\inf_{\mathbf{x} \in \mathcal{X}} \sup_{\xi \in \Xi} \inf_{\mathbf{y} \in \text{conv} \mathcal{Y}} \left\{ \mathbf{c}^\top \mathbf{x} + \mathbf{d}(\xi)^\top \mathbf{y} : \mathbf{T}\mathbf{x} + \mathbf{W}\mathbf{y} \leq \mathbf{h} \right\},$$

which provides a lower bound to the two-stage robust optimization problem (1') since $\text{conv} \mathcal{Y} \supseteq \mathcal{Y}$. On the other hand, we know that the K -adaptability problem (2') by construction bounds (1') from above. We thus conclude that the optimal values of both problems must coincide.

In view of the third set of conditions, it is clear that $K \geq |\mathcal{Y}|$ policies are sufficient for the optimal values of problems (1') and (2') to coincide. Moreover, [25, Theorem 4] presents a problem where the optimal values of problems (1') and (2') differ for every $K < |\mathcal{Y}|$.

As for the second part of the statement, consider the problem

$$\sup_{\xi \in [-1,1]} \inf_{\substack{\tau \in \mathbb{R}, \\ y \in [-1,1]}} \{ \tau : \tau \geq \xi - y, \tau \geq y - \xi \}.$$

The optimal second-stage decision to this two-stage robust optimization problem satisfies $\tau(\xi) = 0$ and $y(\xi) = \xi$, and any $\xi \in [-1, 1]$ attains the optimal objective value 0. Consider now the K -adaptability problem

$$\inf_{\substack{\boldsymbol{\tau} \in \mathbb{R}^K, \\ \boldsymbol{y} \in [-1, 1]^K}} \sup_{\xi \in [-1, 1]} \inf_{k \in \mathcal{K}} \{ \tau_k : \tau_k \geq \xi - y_k, \tau_k \geq y_k - \xi \}.$$

By construction, the objective value of this problem evaluates to at least $\max_{\xi \in [-1, 1]} \min_{k \in \mathcal{K}} |\xi - y_k| > 0$ for any finite K and any feasible solution $(\boldsymbol{\tau}, \boldsymbol{y}) \in \mathbb{R}^K \times [-1, 1]^K$. \square

The validity of the first part of the statement under the second set of conditions in Proposition B.5 was established for the subclass of purely binary K -adaptability problems in [25, Theorem 1]. We also mention that the optimal value of the K -adaptability problem (2') approaches the optimal value of the two-stage robust optimization problem (1') if $K \rightarrow \infty$ and a continuity assumption is satisfied, see [5, Proposition 1].