K -Adaptability in Two-Stage Distributionally Robust Binary Programming

Grani A. Hanasusanto^{a,*}, Daniel Kuhn^a, Wolfram Wiesemann^b

^aRisk Analytics and Optimization Chair, École Polytechnique Fédérale de Lausanne, Switzerland

^bImperial College Business School, Imperial College London, United Kingdom

Abstract

We propose to approximate two-stage distributionally robust programs with binary recourse decisions by their associated K -adaptability problems, which pre-select K candidate second-stage policies here-and-now and implement the best of these policies once the uncertain parameters have been observed. We analyze the approximation quality and the computational complexity of the K -adaptability problem, and we derive explicit mixed-integer linear programming reformulations. We also provide efficient procedures for bounding the probabilities with which each of the K second-stage policies is selected.

Keywords: distributionally robust optimization, integer programming, two-stage problems

1. Introduction

We study two-stage distributionally robust programs of the form

$$\begin{aligned} & \text{minimize} \quad \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}\text{-OCE}_U \left[\tilde{\xi}^\top C \mathbf{x} + \min_{\mathbf{y} \in \mathcal{Y}} \left\{ \tilde{\xi}^\top Q \mathbf{y} : T \mathbf{x} + W \mathbf{y} \leq H \tilde{\xi} \right\} \right] \\ & \text{subject to} \quad \mathbf{x} \in \mathcal{X}, \end{aligned} \quad (\mathcal{DP})$$

where

$$\mathbb{P}\text{-OCE}_U [\phi(\tilde{\xi})] = \inf_{\theta \in \mathbb{R}} \theta + \mathbb{E}_{\mathbb{P}} [U(\phi(\tilde{\xi}) - \theta)]$$

denotes the *optimized certainty equivalent* of a disutility function U under the probability distribution $\mathbb{P}$. Here, $\mathcal{X} \subseteq \mathbb{R}^N$ and $\mathcal{Y} \subseteq \{0, 1\}^M$ are bounded mixed-integer linear sets, while $C \in \mathbb{R}^{Q \times N}$, $Q \in \mathbb{R}^{Q \times M}$, $T \in \mathbb{R}^{L \times N}$, $W \in \mathbb{R}^{L \times M}$ and $H \in \mathbb{R}^{L \times Q}$. The first-stage or *here-and-now* decisions $\mathbf{x}$ are selected prior to the observation of the uncertain parameters $\tilde{\xi} \in \mathbb{R}^Q$, and the second-stage or *wait-and-see* decisions $\mathbf{y}$ are chosen after $\tilde{\xi}$ has been revealed. We assume that all components of $\mathbf{y}$ are binary, while $\mathbf{x}$ may

^{*}An extended version of the paper with all the proofs can be accessed at http://www.optimization-online.org/DB_HTML/2015/04/4878.html.

^{*}Corresponding author: grani.hanasusanto@epfl.ch

have continuous as well as binary components. Problem $\mathcal{DP}$ minimizes the worst-case optimized certainty equivalent over all distributions $\mathbb{P}$ from an ambiguity set $\mathcal{P}$. Problems of the type $\mathcal{DP}$ have many applications, for example in facility location, vehicle routing, unit commitment, layout planning, project scheduling, portfolio selection and game theory.

Despite the broad applicability of problem $\mathcal{DP}$, its numerical solution is extremely challenging. Instead of solving $\mathcal{DP}$ exactly, we propose to solve its associated K -adaptability problem

$$\begin{aligned} & \text{minimize} && \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}\text{-OCE}_U \left[\tilde{\xi}^\top \mathbf{C} \mathbf{x} + \min_{k \in \mathcal{K}} \left\{ \tilde{\xi}^\top \mathbf{Q} \mathbf{y}^k : T\mathbf{x} + \mathbf{W}\mathbf{y}^k \leq \mathbf{H}\tilde{\xi} \right\} \right] \\ & \text{subject to} && \mathbf{x} \in \mathcal{X}, \quad \mathbf{y}^k \in \mathcal{Y}, \quad k \in \mathcal{K}, \end{aligned} \quad (\mathcal{DP}_K)$$

where $\mathcal{K} = \{1, \dots, K\}$. Thus, we pre-select exactly K of the $|\mathcal{Y}| \lesssim 2^M$ possible second-stage decisions in the first stage. After the realization ξ of the random parameters $\tilde{\xi}$ has been observed, we implement the best of these K pre-selected candidate decisions. More precisely, among all candidate decisions that are feasible for ξ , we implement the one that achieves the lowest second-stage cost in scenario ξ . If all K candidate decisions are infeasible for a given ξ , then the innermost minimum in $\mathcal{DP}_K$ is interpreted as an infimum that evaluates to $+\infty$. By construction, the K -adaptability problem may provide a strict upper bound on the optimal value of problem $\mathcal{DP}$ unless $K = |\mathcal{Y}|$, but the hope is that the corresponding optimality gap is small already for some $K \ll |\mathcal{Y}|$. The K -adaptability problem has been studied in [3, 6] in the context of robust two-stage integer programming, where the objective is to minimize the worst-case cost over all uncertainty realizations $\xi \in \Xi$. We refer to [6] for a review of the literature on robust two-stage integer programming.

The rest of the paper develops as follows. We provide a detailed formulation of the K -adaptability problem $\mathcal{DP}_K$ in Section 2. Section 3 studies the approximation quality as well as an explicit MILP reformulation of the distributionally robust K -adaptability problem with objective uncertainty, that is, for instances of $\mathcal{DP}_K$ where $\tilde{\xi}$ impacts only the objective function. We close with numerical results in Section 4. The Electronic Companion contains a treatment of generic instances of $\mathcal{DP}_K$ where both the objective function and the constraints are affected by uncertainty, as well as most of the proofs.

Notation. Variables with tilde signs represent random objects. We denote by $\mathbf{e}$ the vector of ones and by $\mathbf{e}_i$ the vector whose i -th entry is 1 while all other entries are 0. The indicator function $\mathbb{I}[\mathcal{E}]$ of a logical expression $\mathcal{E}$ is defined through $\mathbb{I}[\mathcal{E}] = 1$ if $\mathcal{E}$ is true; $= 0$ otherwise. To avoid tedious case distinctions, we define the minimum (maximum) of an empty set as $+\infty$ ($-\infty$).

2. Problem Formulation

Problem $\mathcal{DP}$ accounts for both ambiguity aversion (through the ambiguity set) and risk aversion (through the optimized certainty equivalent). We discuss both of these components in turn.

Ambiguity Set. By construction, the ambiguity set $\mathcal{P}$ contains all probability distributions that share certain known properties of the unknown true distribution $\mathbb{P}^0$ of $\tilde{\xi}$. Hedging against the worst probability distribution within the ambiguity set $\mathcal{P}$ reflects an aversion against distributional ambiguity, which enjoys strong justification from decision theory [4, 5].

We henceforth assume that the ambiguity set $\mathcal{P}$ is of the form

$$\mathcal{P} = \left\{ \mathbb{P} \in \mathcal{M}_+(\mathbb{R}^Q) : \mathbb{P}[\tilde{\xi} \in \Xi] = 1, \mathbb{E}_{\mathbb{P}}[g(\tilde{\xi})] \leq c \right\}, \quad (1)$$

where $\mathcal{M}_+(\mathbb{R}^Q)$ denotes the cone of nonnegative Borel measures supported on $\mathbb{R}^Q$. The support set Ξ is defined as the smallest set that is known to satisfy $\tilde{\xi} \in \Xi$ w.p. 1, and it constitutes a nonempty bounded polytope of the form $\Xi = \{\xi \in \mathbb{R}^Q : A\xi \leq b\}$ with $A \in \mathbb{R}^{R \times Q}$ and $b \in \mathbb{R}^R$. We assume that $c \in \mathbb{R}^S$ and that $g : \mathbb{R}^Q \rightarrow \mathbb{R}^S$ has convex piecewise linear component functions of the form

$$g_s(\xi) = \max_{t \in \mathcal{T}} g_{st}^\top \xi \quad \forall s \in S = \{1, \dots, S\}.$$

Without loss of generality, the index t of the linear pieces ranges over the same index set $\mathcal{T} = \{1, \dots, T\}$ for each component g_s . To ensure the applicability of strong semi-infinite duality results, we finally assume that the ambiguity set $\mathcal{P}$ contains a Slater point in the sense that there is a distribution $\mathbb{P} \in \mathcal{P}$ such that $\mathbb{E}_{\mathbb{P}}[g_s(\tilde{\xi})] < c$ for all $s \in S$ for which $g_s(\xi)$ is nonlinear.

Ambiguity sets of the form (1) are flexible enough to encode bounds on the expected value and the mean absolute deviation of the unknown true distribution $\mathbb{P}^0$, and they can also be used to approximate nonlinear dispersion measures such as the variance or the standard deviation [11].

Optimized Certainty Equivalent. Intuitively, the optimized certainty equivalent $\mathbb{P}\text{-OCE}_U[\phi(\tilde{\xi})]$ represents the expected present value of an optimized payment schedule that splits an uncertainty-affected future liability $\phi(\tilde{\xi})$ into a fixed installment θ that is paid today and a remainder $\phi(\tilde{\xi}) - \theta$ that is paid after the realization of $\tilde{\xi}$ has been observed [2].

The decision maker's attitude towards risk is controlled by the choice of the disutility function U . We consider increasing, convex and piecewise affine disutility functions of the form

$$U(x) = \max_{i \in I} \{s_i x + t_i\}, \quad \text{where } I = \{1, \dots, I\} \text{ and } s \geq \mathbf{0}, s \neq \mathbf{0}. \quad (2)$$

The choice $I = 1$, $s = 1$ and $t = 0$ corresponds to the worst-case expected value, whereas $I = 2$, $s = ((1 - \beta)^{-1}, 0)^\top$ and $t = (0, 0)^\top$ recovers the worst-case conditional value-at-risk at level β , that is, the worst-case expected value of the $(1 - \beta) \cdot 100\%$ largest outcomes.

3. The K -Adaptability Problem with Objective Uncertainty

In this section, we assume that the random parameters $\tilde{\xi}$ only enter the objective function of the two-stage robust binary program $\mathcal{DP}$:

$$\begin{aligned} & \text{minimize} \quad \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}\text{-OCE}_U \left[\tilde{\xi}^\top C \mathbf{x} + \min_{\mathbf{y} \in \mathcal{Y}} \{ \tilde{\xi}^\top Q \mathbf{y} : T \mathbf{x} + W \mathbf{y} \leq \mathbf{h} \} \right] \\ & \text{subject to} \quad \mathbf{x} \in \mathcal{X}, \end{aligned} \quad (\mathcal{DPO})$$

where $\mathbf{h} \in \mathbb{R}^L$. Problem $\mathcal{DPO}$ arises naturally in a number of application domains, such as traveling salesman and vehicle routing problems with uncertain travel times, network expansion problems with uncertain costs, facility location problems with uncertain customer demands and layout planning problems with uncertain production quantities. A discussion of generic instances of $\mathcal{DP}$, where both the objective function and the constraints are affected by uncertainty, is relegated to the Electronic Companion.

In the following, we study the K -adaptability problem associated with $\mathcal{DPO}$:

$$\begin{aligned} & \text{minimize} \quad \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}\text{-OCE}_U \left[\tilde{\xi}^\top C \mathbf{x} + \min_{k \in \mathcal{K}} \{ \tilde{\xi}^\top Q \mathbf{y}^k : T \mathbf{x} + W \mathbf{y}^k \leq \mathbf{h} \} \right] \\ & \text{subject to} \quad \mathbf{x} \in \mathcal{X}, \mathbf{y}^k \in \mathcal{Y}, k \in \mathcal{K} \end{aligned} \quad (3)$$

We note that the K -adaptability problem (3) is equivalent to the problem

$$\begin{aligned} & \text{minimize} \quad \sup_{\mathbb{P} \in \mathcal{P}} \mathbb{P}\text{-OCE}_U \left[\tilde{\xi}^\top C \mathbf{x} + \min_{k \in \mathcal{K}} \tilde{\xi}^\top Q \mathbf{y}^k \right] \\ & \text{subject to} \quad \mathbf{x} \in \mathcal{X}, \mathbf{y}^k \in \mathcal{Y}, k \in \mathcal{K} \\ & \quad \quad \quad T \mathbf{x} + W \mathbf{y}^k \leq \mathbf{h} \quad \forall k \in \mathcal{K}, \end{aligned} \quad (\mathcal{DPO}_K)$$

where we have shifted the second-stage constraints to the first stage, see [6, Observation 1].

We now show that the problem $\mathcal{DPO}_K$ has an equivalent reformulation as a MILP. To this end, we interpret the objective function in $\mathcal{DPO}_K$ as a moment problem. The dual of this moment problem constitutes a semi-infinite program which we can simplify using a standard LP dualization.

Theorem 1. *For the ambiguity set $\mathcal{P}$ defined in (1) and the disutility function U defined in (2), problem $\mathcal{DPO}_K$ is equivalent to the following MILP.*

$$\begin{aligned} & \text{minimize} \quad \alpha + \mathbf{c}^\top \boldsymbol{\beta} + \theta \\ & \text{subject to} \quad \left. \begin{aligned} & \mathbf{x} \in \mathcal{X}, \mathbf{y}^k \in \mathcal{Y}, k \in \mathcal{K}, \alpha \in \mathbb{R}, \boldsymbol{\beta} \in \mathbb{R}_+^S \\ & \boldsymbol{\gamma}^i \in \mathbb{R}_+^R, \boldsymbol{\delta}^i \in \mathbb{R}_+^K, \boldsymbol{\Lambda}^i \in \mathbb{R}_+^{S \times T}, i \in \mathcal{I} \\ & \mathbf{z}^{i,k} \in \mathbb{R}_+^M, i \in \mathcal{I} \text{ and } k \in \mathcal{K}, \theta \in \mathbb{R} \\ & T \mathbf{x} + W \mathbf{y}^k \leq \mathbf{h} \quad \forall k \in \mathcal{K} \\ & \mathbf{b}^\top \boldsymbol{\gamma}^i + t_i \leq \alpha + s_i \theta, \mathbf{e}^\top \boldsymbol{\delta}^i = 1, \boldsymbol{\Lambda}^i \mathbf{e} = \boldsymbol{\beta} \\ & A^\top \boldsymbol{\gamma}^i + \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}} \Lambda_{st}^i \mathbf{g}_{st} = s_i C \mathbf{x} + \sum_{k \in \mathcal{K}} s_i Q \mathbf{z}^{i,k} \\ & \mathbf{z}^{i,k} \leq \mathbf{y}^k, \mathbf{z}^{i,k} \leq \delta_k^i \mathbf{e} \\ & \mathbf{z}^{i,k} \geq (\delta_k^i - 1) \mathbf{e} + \mathbf{y}^k \end{aligned} \right\} \quad \forall i \in \mathcal{I}. \end{aligned} \quad (4)$$

By construction, the optimal value of problem $\mathcal{DPO}_K$ constitutes an upper bound on the optimal value of problem $\mathcal{DPO}$ as we restrict our flexibility in the second stage. For classical two-stage robust binary programs, it has been shown that the approximation provided by the K -adaptability formulation is tight whenever K exceeds the affine dimension of either the uncertainty set or the second-stage feasible region, see [6, Theorem 1]. We now show that this favorable approximation behavior generalizes to distributionally robust two-stage binary programs.

Theorem 2. *Problem $\mathcal{DPO}_K$ has the same optimal value as problem $\mathcal{DPO}$ if we choose $K \geq I \cdot \min \{\dim \mathcal{Y}, \text{rk } \mathbf{Q}\} + I$ policies.*

We note that $\dim \mathcal{Y} \leq M$ and $\text{rk } \mathbf{Q} \leq Q$ by construction. Without loss of generality, we can further assume that $\text{rk } \mathbf{Q} \leq \dim \Xi + 1$. Indeed, if this is not the case, then there is a matrix $\mathbf{Q}'$ such that $\text{rk } \mathbf{Q}' \leq \dim \Xi + 1$ and the optimal value and the optimal solutions to $\mathcal{DPO}_K$ do not change if we replace $\mathbf{Q}$ with $\mathbf{Q}'$, see also [6, Remark 3].

3.1. Persistence

We now study the contribution of each second-stage policy $\mathbf{y}^k$ to the objective value in problem $\mathcal{DPO}_K$. This analysis can provide important insights for practical decision-making. Amongst others, it can help to determine how much adaptability is needed (*i.e.*, it can inform the choice of K), and it can elucidate the relative ‘criticality’ of each second-stage policy.

We develop two alternative measures for the importance of a second-stage policy. We first determine the probability with which a particular policy $\mathbf{y}^k$ is chosen under a worst-case distribution in problem $\mathcal{DPO}_K$. This approach is reminiscent of the persistency analysis in [8], and it is justified by the assumption that the decision maker optimizes in view of the *worst* distribution from within $\mathcal{P}$. Afterwards, we determine the minimum and maximum probability that a policy $\mathbf{y}^k$ is chosen if the unknown true probability distribution $\mathbb{P}^0$ can be *any* distribution within the ambiguity set $\mathcal{P}$. These probability bounds characterize the ambiguity of the persistence inherited by the ambiguity of $\mathbb{P}^0$ that governs $\tilde{\xi}$.

We now study the probability that a particular second-stage policy $\mathbf{y}^k$ is chosen under a worst-case distribution for $\mathcal{DPO}_K$.

Theorem 3. *For a feasible decision $(\mathbf{x}, \{\mathbf{y}^k\}_k)$ in problem $\mathcal{DPO}_K$ with the ambiguity set (1) and the disutility function (2), and let $(\phi, \{\chi^i\}_i, \psi, \{\omega^i\}_i)$ be an optimal solution to the LP*

$$\begin{aligned} & \text{maximize} && \sum_{i \in \mathcal{I}} t_i \phi_i + s_i (\mathbf{C}\mathbf{x})^\top \chi^i + \psi_i \\ & \text{subject to} && \left. \begin{aligned} & \phi_i \in \mathbb{R}_+, \chi^i \in \mathbb{R}^Q, \psi_i \in \mathbb{R}, \omega^i \in \mathbb{R}^S \\ & \mathbf{A}\chi^i \leq \mathbf{b}\phi_i \\ & \psi_i \leq s_i (\mathbf{Q}\mathbf{y}^k)^\top \chi^i \quad \forall k \in \mathcal{K} \\ & \mathbf{g}_{st}^\top \chi^i \leq \omega_s^i \quad \forall s \in \mathcal{S}, \forall t \in \mathcal{T} \\ & \mathbf{e}^\top \phi = 1, \mathbf{s}^\top \phi = 1, \sum_{i \in \mathcal{I}} \omega^i \leq \mathbf{c}. \end{aligned} \right\} \quad \forall i \in \mathcal{I} \end{aligned} \quad (5)$$

Then, a worst-case distribution $\mathbb{P}^\star$ for $(\mathbf{x}, \{\mathbf{y}^k\}_k)$ in problem $\mathcal{DPO}_K$ is defined through

$$\mathbb{P}^\star \left[\tilde{\xi} = \frac{\chi^i}{\phi_i} \right] = \phi_i, \quad i \in \mathcal{I} : \phi_i > 0,$$

and the probability with which a policy $\mathbf{y}^k$, $k \in \mathcal{K}$, is chosen under $\mathbb{P}^\star$ is

$$\mathbb{P}^\star \left[\tilde{\xi}^\top \mathbf{Q} \mathbf{y}^k \leq \tilde{\xi}^\top \mathbf{Q} \mathbf{y}^{k'} \quad \forall k' \in \mathcal{K} \right] = \sum_{i \in \mathcal{I}} \phi_i \cdot \mathbb{I} \left[(\chi^i)^\top \mathbf{Q} \mathbf{y}^k \leq (\chi^i)^\top \mathbf{Q} \mathbf{y}^{k'} \quad \forall k' \in \mathcal{K} \right].$$

We now evaluate the maximum probability with which a second-stage policy $\mathbf{y}^k$ is chosen under any distribution $\mathbb{P} \in \mathcal{P}$.

Proposition 1. *Let $(\mathbf{x}, \{\mathbf{y}^k\}_k)$ be a feasible decision in problem $\mathcal{DPO}_K$ with the ambiguity set (1) and the disutility function (2). The maximum probability with which policy $\mathbf{y}^k$, $k \in \mathcal{K}$, is chosen under any probability distribution $\mathbb{P} \in \mathcal{P}$ is given by the optimal value of the LP*

$$\begin{aligned} & \text{minimize} \quad \alpha + \mathbf{c}^\top \boldsymbol{\beta} \\ & \text{subject to} \quad \alpha \in \mathbb{R}, \quad \boldsymbol{\beta} \in \mathbb{R}_+^S, \quad \boldsymbol{\gamma} \in \mathbb{R}_+^R, \quad \boldsymbol{\delta} \in \mathbb{R}_+^K, \quad \boldsymbol{\kappa} \in \mathbb{R}_+^R \\ & \quad \quad \quad \boldsymbol{\Lambda} \in \mathbb{R}_+^{S \times T}, \quad \boldsymbol{\Phi} \in \mathbb{R}_+^{S \times T} \\ & \quad \quad \quad \alpha \geq \mathbf{b}^\top \boldsymbol{\gamma} + 1, \quad \alpha \geq \mathbf{b}^\top \boldsymbol{\kappa}, \quad \boldsymbol{\Lambda} \mathbf{e} = \boldsymbol{\beta}, \quad \boldsymbol{\Phi} \mathbf{e} = \boldsymbol{\beta} \\ & \quad \quad \quad \mathbf{A}^\top \boldsymbol{\gamma} + \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}} \Lambda_{st} \mathbf{g}_{st} = \sum_{k' \in \mathcal{K}} \mathbf{Q}(\mathbf{y}^{k'} - \mathbf{y}^k) \delta_{k'} \\ & \quad \quad \quad \mathbf{A}^\top \boldsymbol{\kappa} + \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}} \Phi_{st} \mathbf{g}_{st} = \mathbf{0}. \end{aligned}$$

Proof. The maximum probability with which policy $\mathbf{y}^k$ is chosen under any probability distribution $\mathbb{P} \in \mathcal{P}$ is given by the optimal value of the following moment problem.

$$\begin{aligned} & \text{maximize} \quad \int_{\Xi} \mathbb{I} \left[\xi^\top \mathbf{Q} \mathbf{y}^k \leq \xi^\top \mathbf{Q} \mathbf{y}^{k'} \quad \forall k' \in \mathcal{K} \right] d\mu(\xi) \\ & \text{subject to} \quad \mu \in \mathcal{M}_+(\mathbb{R}^Q) \\ & \quad \quad \quad \int_{\Xi} d\mu(\xi) = 1 \\ & \quad \quad \quad \int_{\Xi} \mathbf{g}(\xi) d\mu(\xi) \leq \mathbf{c} \end{aligned}$$

Strong duality is guaranteed by Proposition 3.4 in [10], which is applicable since the ambiguity set $\mathcal{P}$ contains a Slater point. Thus, the dual problem

$$\begin{aligned} & \text{minimize} \quad \alpha + \mathbf{c}^\top \boldsymbol{\beta} \\ & \text{subject to} \quad \alpha \in \mathbb{R}, \quad \boldsymbol{\beta} \in \mathbb{R}_+^S \\ & \quad \quad \quad \alpha + \mathbf{g}(\xi)^\top \boldsymbol{\beta} \geq \mathbb{I} \left[\xi^\top \mathbf{Q} \mathbf{y}^k \leq \xi^\top \mathbf{Q} \mathbf{y}^{k'} \quad \forall k' \in \mathcal{K} \right] \quad \forall \xi \in \Xi \end{aligned}$$

attains the same optimal value. We can replace the semi-infinite constraint in this problem with the two constraints

$$\begin{aligned} \alpha + \mathbf{g}(\xi)^\top \boldsymbol{\beta} & \geq 1 \quad \forall \xi \in \Xi : \xi^\top \mathbf{Q} \mathbf{y}^k \leq \xi^\top \mathbf{Q} \mathbf{y}^{k'} \quad \forall k' \in \mathcal{K} \\ \alpha + \mathbf{g}(\xi)^\top \boldsymbol{\beta} & \geq 0 \quad \forall \xi \in \Xi. \end{aligned} \tag{6}$$

In the following, we focus on the reformulation of the first constraint; the second constraint can be dealt with analogously. The first constraint is satisfied if and only if the optimal value of

$$\begin{aligned} & \text{minimize} && \alpha + \boldsymbol{\beta}^\top \mathbf{g}(\boldsymbol{\xi}) \\ & \text{subject to} && \boldsymbol{\xi} \in \mathbb{R}^Q \\ & && \mathbf{A}\boldsymbol{\xi} \leq \mathbf{b} \\ & && \boldsymbol{\xi}^\top \mathbf{Q}\mathbf{y}^k \leq \boldsymbol{\xi}^\top \mathbf{Q}\mathbf{y}^{k'} \quad \forall k' \in \mathcal{K} \end{aligned}$$

is greater than or equal to 1. By employing an epigraph reformulation, we can replace $\mathbf{g}$ with its definition to obtain the following equivalent problem:

$$\begin{aligned} & \text{minimize} && \alpha + \boldsymbol{\beta}^\top \boldsymbol{\eta} \\ & \text{subject to} && \boldsymbol{\xi} \in \mathbb{R}^Q, \quad \boldsymbol{\eta} \in \mathbb{R}^S \\ & && \mathbf{A}\boldsymbol{\xi} \leq \mathbf{b} \\ & && \boldsymbol{\xi}^\top \mathbf{Q}\mathbf{y}^k \leq \boldsymbol{\xi}^\top \mathbf{Q}\mathbf{y}^{k'} \quad \forall k' \in \mathcal{K} \\ & && \eta_s \geq \mathbf{g}_{st}^\top \boldsymbol{\xi} \quad \forall s \in \mathcal{S}, \forall t \in \mathcal{T} \end{aligned}$$

This reformulation exploits the fact that $\boldsymbol{\beta} \geq \mathbf{0}$. The dual problem

$$\begin{aligned} & \text{maximize} && \alpha - \mathbf{b}^\top \boldsymbol{\gamma} \\ & \text{subject to} && \boldsymbol{\gamma} \in \mathbb{R}_+^R, \quad \boldsymbol{\delta} \in \mathbb{R}_+^K, \quad \boldsymbol{\Lambda} \in \mathbb{R}_+^{S \times T} \\ & && \mathbf{A}^\top \boldsymbol{\gamma} + \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}} \Lambda_{st} \mathbf{g}_{st} = \sum_{k' \in \mathcal{K}} \mathbf{Q}(\mathbf{y}^{k'} - \mathbf{y}^k) \delta_{k'} \\ & && \boldsymbol{\Lambda} \mathbf{e} = \boldsymbol{\beta} \end{aligned} \tag{7}$$

is feasible, which implies that strong duality holds. We conclude that the first constraint in (6) is satisfied if and only if there exists $(\boldsymbol{\gamma}, \boldsymbol{\delta}, \boldsymbol{\Lambda})$ feasible in (7) for which $\alpha - \mathbf{b}^\top \boldsymbol{\gamma} \geq 1$. A similar reformulation can be derived for the second constraint in (6), which concludes the proof. $\square$

We close the section with the following result about the minimum probability with which a second-stage policy $\mathbf{y}^k$ is chosen under any distribution $\mathbb{P} \in \mathcal{P}$.

Proposition 2. *Let $(\mathbf{x}, \{\mathbf{y}^k\}_k)$ be a feasible decision in problem $\mathcal{DP}\mathcal{O}_K$ with the ambiguity set (1) and the disutility function (2). For $k \in \mathcal{K}$, let $\mathcal{K}(k)$ be the subset of indices $k' \in \mathcal{K}$ for which there is a parameter realization $\boldsymbol{\xi} \in \Xi$ such that $\boldsymbol{\xi}^\top \mathbf{Q}(\mathbf{y}^k - \mathbf{y}^{k'}) > 0$. Then, the minimum probability with which policy $\mathbf{y}^k$, $k \in \mathcal{K}$, is chosen under any probability distribution $\mathbb{P} \in \mathcal{P}$ is given by the optimal value of the LP*

$$\begin{aligned} & \text{maximize} && \alpha - \mathbf{c}^\top \boldsymbol{\beta} \\ & \text{subject to} && \alpha \in \mathbb{R}, \quad \boldsymbol{\beta} \in \mathbb{R}_+^S, \quad \boldsymbol{\gamma} \in \mathbb{R}_+^R, \quad \boldsymbol{\Lambda} \in \mathbb{R}_+^{S \times T} \\ & && \boldsymbol{\kappa}^{k'} \in \mathbb{R}_+^R, \quad \pi_{k'} \in \mathbb{R}_+, \quad \boldsymbol{\Phi}^{k'} \in \mathbb{R}_+^{S \times T} \\ & && \left. \begin{aligned} & \mathbf{b}^\top \boldsymbol{\gamma} + \alpha \leq 1, \quad \mathbf{b}^\top \boldsymbol{\kappa}^{k'} + \alpha \leq 0, \quad \boldsymbol{\Lambda} \mathbf{e} = \boldsymbol{\beta}, \quad \boldsymbol{\Phi}^{k'} \mathbf{e} = \boldsymbol{\beta} \\ & \mathbf{A}^\top \boldsymbol{\kappa}^{k'} + \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}} \Phi_{st}^{k'} \mathbf{g}_{st} = \mathbf{Q}(\mathbf{y}^k - \mathbf{y}^{k'}) \pi_{k'} \end{aligned} \right\} \quad \forall k' \in \mathcal{K}(k) \\ & && \mathbf{A}^\top \boldsymbol{\gamma} + \sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}} \Lambda_{st} \mathbf{g}_{st} = \mathbf{0}. \end{aligned}$$

Proof. The minimum probability with which policy $\mathbf{y}^k$ is chosen under any probability distribution $\mathbb{P} \in \mathcal{P}$ is given by the optimal value of the following moment problem.

$$\begin{aligned} & \text{minimize} && \int_{\Xi} \mathbb{I} [\xi^\top \mathbf{Q} \mathbf{y}^k \leq \xi^\top \mathbf{Q} \mathbf{y}^{k'} \quad \forall k' \in \mathcal{K}] d\mu(\xi) \\ & \text{subject to} && \mu \in \mathcal{M}_+(\mathbb{R}^Q) \\ & && \int_{\Xi} d\mu(\xi) = 1 \\ & && \int_{\Xi} \mathbf{g}(\xi) d\mu(\xi) \leq \mathbf{c} \end{aligned}$$

Strong duality is guaranteed by Proposition 3.4 in [10], which is applicable since the ambiguity set $\mathcal{P}$ contains a Slater point. Thus, the dual problem

$$\begin{aligned} & \text{maximize} && \alpha - \mathbf{c}^\top \boldsymbol{\beta} \\ & \text{subject to} && \alpha \in \mathbb{R}, \quad \boldsymbol{\beta} \in \mathbb{R}_+^S \\ & && \alpha - \mathbf{g}(\xi)^\top \boldsymbol{\beta} \leq \mathbb{I} [\xi^\top \mathbf{Q} \mathbf{y}^k \leq \xi^\top \mathbf{Q} \mathbf{y}^{k'} \quad \forall k' \in \mathcal{K}] \quad \forall \xi \in \Xi \end{aligned}$$

attains the same optimal value. We can replace the semi-infinite constraint in this problem with the two constraints

$$\begin{aligned} \alpha - \mathbf{g}(\xi)^\top \boldsymbol{\beta} &\leq 1 & \forall \xi \in \Xi \\ \alpha - \mathbf{g}(\xi)^\top \boldsymbol{\beta} &\leq 0 & \forall \xi \in \Xi \setminus \Xi_k, \end{aligned} \tag{8}$$

where $\Xi_k = \{\xi \in \Xi : \xi^\top \mathbf{Q} \mathbf{y}^k \leq \xi^\top \mathbf{Q} \mathbf{y}^{k'} \quad \forall k' \in \mathcal{K}\}$. Note that

$$\Xi \setminus \Xi_k = \bigcup_{k' \in \mathcal{K}} \{\xi \in \Xi : \xi^\top \mathbf{Q} (\mathbf{y}^k - \mathbf{y}^{k'}) > 0\}.$$

Since the constraints in (8) are continuous in ξ , we can replace the set $\Xi \setminus \Xi_k$ in (8) with

$$\begin{aligned} \text{cl}(\Xi \setminus \Xi_k) &= \text{cl} \bigcup_{k' \in \mathcal{K}} \{\xi \in \Xi : \xi^\top \mathbf{Q} (\mathbf{y}^k - \mathbf{y}^{k'}) > 0\} \\ &= \bigcup_{k' \in \mathcal{K}} \text{cl} \{\xi \in \Xi : \xi^\top \mathbf{Q} (\mathbf{y}^k - \mathbf{y}^{k'}) > 0\} \\ &= \bigcup_{k' \in \mathcal{K}(k)} \{\xi \in \Xi : \xi^\top \mathbf{Q} (\mathbf{y}^k - \mathbf{y}^{k'}) \geq 0\}. \end{aligned}$$

Here, the second identity holds because $\mathcal{K}$ is finite, which implies that the union and the closure operators commute. The last identity follows from the fact that the k' -th set in the union is non-empty if and only if $k' \in \mathcal{K}(k)$. Hence, the semi-infinite constraints in (8) are satisfied if and only if the constraints

$$\begin{aligned} \alpha - \mathbf{g}(\xi)^\top \boldsymbol{\beta} &\leq 1 & \forall \xi \in \Xi \\ \alpha - \mathbf{g}(\xi)^\top \boldsymbol{\beta} &\leq 0 & \forall k' \in \mathcal{K}(k), \quad \forall \xi \in \Xi : \xi^\top \mathbf{Q} \mathbf{y}^k \geq \xi^\top \mathbf{Q} \mathbf{y}^{k'} \end{aligned}$$

are satisfied. The result now follows if we apply similar reformulations to these semi-infinite constraints as in the proof of Proposition 1. $\square$

4. Numerical Experiments

We apply the K -adaptability approximation to a stylized two-stage version of the vertex packing problem. The optimization problems in this section are solved using the YALMIP modeling language [9] and Gurobi Optimizer 5.6 [7] with the default settings and a time limit of 7,200 seconds.

Consider an undirected, node-weighted graph $G = (V, E, \xi)$ with nodes $V = \{1, \dots, N\}$, edges $E \subseteq \{\{i, j\} : i, j \in V\}$ and weights $\xi_i \in \mathbb{R}_+$, $i \in V$. The vertex packing problem asks for a packing P , that is, a subset of nodes $P \subseteq V$ such that no pair of nodes in P is connected by an edge, that maximizes the sum of node weights $\sum_{i \in P} \xi_i$. We consider a two-stage distributionally robust variant of the problem where the node weights are modeled as a random vector $\tilde{\xi}$ that is governed by an unknown probability distribution. The goal is to pre-commit to a subset of nodes P_1 before observing $\tilde{\xi}$ and to complete P_1 to a packing P_2 after observing $\tilde{\xi}$ so that $|P_1 \setminus P_2| + |P_2 \setminus P_1| \leq B$, where B denotes the budget of change, and the packing P_2 maximizes the sum of node weights.

Using the worst-case conditional value-at-risk as a risk measure, the problem can be formulated as a distributionally robust two-stage binary program with objective uncertainty:

$$\begin{aligned} \text{maximize} \quad & \inf_{\mathbb{P} \in \mathcal{P}} \mathbb{P}\text{-CVaR}_\beta \left[\max_{\mathbf{y} \in \{0,1\}^N} \left\{ \tilde{\xi}^\top \mathbf{y} : \|\mathbf{x} - \mathbf{y}\|_1 \leq B, \quad y_i + y_j \leq 1 \quad \forall \{i, j\} \in E \right\} \right] \\ \text{subject to} \quad & \mathbf{x} \in \{0, 1\}^N \end{aligned}$$

In this formulation, the decisions $\mathbf{x}$ and $\mathbf{y}$ represent indicator functions for the sets P_1 and P_2 , respectively, that is, we have $x_i = \mathbb{I}[i \in P_1]$ and $y_i = \mathbb{I}[i \in P_2]$, $i \in V$, and $\mathbb{P}\text{-CVaR}_\beta$ denotes the conditional value-at-risk at level β under the distribution $\mathbb{P} \in \mathcal{P}$.

For our numerical experiments, we generate random graphs with $N \in \{10, \dots, 50\}$ nodes and $|E| = 5|V|$ edges. Note that the number of edges scales linearly with the number of nodes, which implies that the expected size of a maximum cardinality packing (*i.e.*, the maximum packing for unit node weights) grows proportionally with N . This ensures that the optimal packings do not degenerate to a trivial solution (*i.e.*, none or almost all nodes) as N increases. We choose the following ambiguity set for the node weights.

$$\mathcal{P} = \left\{ \mathbb{P} \in \mathcal{M}_+(\mathbb{R}^N) : \mathbb{E}_{\mathbb{P}}[\tilde{\xi}] = \boldsymbol{\mu}, \quad \mathbb{E}_{\mathbb{P}}[|\tilde{\xi} - \boldsymbol{\mu}|] \leq 0.15\boldsymbol{\mu}, \quad \mathbb{E}_{\mathbb{P}}[|\mathbf{e}^\top(\tilde{\xi} - \boldsymbol{\mu})|] \leq 0.15N^{-1/2}\mathbf{e}^\top\boldsymbol{\mu} \right\}$$

Here, the average node weights $\boldsymbol{\mu}$ are chosen uniformly at random from the interval $[0, 10]^N$. The second condition in $\mathcal{P}$ imposes upper bounds on the mean absolute deviations of the individual node weights. The third condition is inspired by the central limit theorem, and it imposes an upper bound on the cumulative deviation of the node weights from their expected values [1, §2].

The K -adaptability formulation (4) corresponding to the vertex packing problem has $O(KN)$ binary variables. Table 1 shows that most of the problem instances can be solved to optimality within the set time limit. The table also shows that problems with a smaller budget of change B are harder to solve. An investigation of the reports generated by Gurobi reveals that for small values of B , the solver requires a long time

		Number of nodes N				
K		10	20	30	40	50
$B = 0.2N$	2	100%/ $<1s$ /0%	100%/ $<1s$ /0%	100%/ $1s$ /0%	100%/ $17s$ /0%	100%/ $4m:12s$ /0%
	3	100%/ $<1s$ /0%	100%/ $<1s$ /0%	100%/ $12s$ /0%	100%/ $3m:14s$ /0%	74%/ $44m:39s$ /22.35%
	4	100%/ $<1s$ /0%	100%/ $2s$ /0%	100%/ $2m:11s$ /0%	65%/ $54m:27s$ /13.52%	0%/ $-$ /22.38%
$B = 0.4N$	2	100%/ $<1s$ /0%	100%/ $<1s$ /0%	100%/ $1s$ /0%	100%/ $9s$ /0%	100%/ $1m:13s$ /0%
	3	100%/ $<1s$ /0%	100%/ $<1s$ /0%	100%/ $3s$ /0%	100%/ $49s$ /0%	100%/ $12m:24s$ /0%
	4	100%/ $<1s$ /0%	100%/ $1s$ /0%	100%/ $36s$ /0%	100%/ $17m:33s$ /0%	10%/ $29m:05s$ /21.68%

Table 1: Summary of the results for the vertex packing problem. Each entry in the table documents the percentage of instances solved within the time limit, the average solution time for the instances solved within the time limit and the average optimality gap for the instances not solved to optimality. All results are averaged over 100 instances.

Figure 1: Improvement of the best 2-, 3- and 4-adaptable solutions determined within the set time limit over the static solutions for the vertex packing problem with $B = 0.2N$ (left) and $B = 0.4N$ (right). The figures show the improvements for problems with $N = 10, 20, \dots, 50$ nodes as averages over 100 instances.

to determine good feasible solutions. Finally, the table presents estimates for the optimality gaps of those instances that could not be solved to optimality. These estimates are derived from a progressive (upper) bound on the optimal value of problem DPO that results from disregarding the integrality requirement for the second-stage decisions y , applying the classical min-max theorem to exchange the order of the maximization problem over $\mathbb{P} \in \mathcal{P}$ and the minimization problem over $y \in \mathcal{Y}$ and subsequently dualizing the maximization problem.

Figure 1 shows the improvement of the 2-, 3- and 4-adaptable solutions over the static solutions that take all decisions here-and-now. The figure reveals that the improvement increases with the number of nodes N , but it saturates as N increases. Indeed, one can show that for the considered problem class, the outperformance of the fully adaptable solutions to problem DPO over the static solutions is bounded by a constant. The figure also shows that the improvement of the adaptable solutions increases with the budget of change B . This is intuitive as higher values of B give more flexibility to modify the vertex packing in the second stage when the node weights are known.

Table 2 shows the minimum and maximum probabilities with which the policies y^k , $k \in \mathcal{K}$, are chosen. Note that in all instances, one policy has a maximum probability of 1 while all other policies have a minimum probability of 0. This is due to the fact that the ambiguity set $\mathcal{P}$ contains the Dirac distribution that places all probability mass on the expected value μ . For this distribution, there is always an optimal solution which

K	Number of nodes N		
	10	30	50
2	[0 0.32], [0.62 1]	[0 0.65], [0.32 1]	[0 0.92], [0.07 1]
3	[0 0.36], [0 0.66], [0.21 1]	[0 0.60], [0 0.88], [0.05 1]	[0 0.75], [0 0.91], [0.02 1]
4	[0 0.27], [0 0.46], [0 0.74], [0.11 1]	[0 0.34], [0 0.62], [0 0.86], [0.04 1]	[0 0.35], [0 0.75], [0 0.91], [0.02 1]

Table 2: Minimum and maximum probabilities with which the policies y^k , $k \in \mathcal{K}$, are chosen in the vertex packing problem. The probabilities are averaged over 100 instances, and the policies are ordered according to increasing maximum probabilities.

selects a single policy with probability 1.

Acknowledgements. The authors are grateful to Rüdiger Schultz and the anonymous referee for their thoughtful comments which substantially improved the paper. This research was supported by the Swiss National Science Foundation grant BSCGI0_157733 and the EPSRC grants EP/M028240/1 and EP/M027856/1.

- [1] C. Bandi and D. Bertsimas. Tractable stochastic analysis in high dimensions via robust optimization. *Mathematical Programming B*, 134(1):23–70, 2012.
- [2] A. Ben-Tal and M. Teboulle. An old-new concept of convex risk measures: The optimized certainty equivalent. *Mathematical Finance*, 17(3):449–476, 2007.
- [3] D. Bertsimas and C. Caramanis. Finite adaptability in multistage linear optimization. *IEEE Transactions on Automatic Control*, 55(12):2751–2766, 2010.
- [4] L. G. Epstein. A definition of uncertainty aversion. *Review of Economic Studies*, 66(3):579–608, 1999.
- [5] I. Gilboa and D. Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989.
- [6] G. A. Hanasusanto, D. Kuhn, and W. Wiesemann. K -adaptability in two-stage robust binary programming. *Operations Research*, 63(4):877–891, 2015.
- [7] Gurobi Optimization, Inc. Gurobi optimizer reference manual, 2014.
- [8] X. Li, K. Natarajan, C.-P. Teo, and Z. Zheng. Distributionally robust mixed integer linear programs: Persistency model with applications. *European Journal of Operational Research*, 233(3):459–473, 2014.
- [9] J. Löfberg. YALMIP : A toolbox for modeling and optimization in MATLAB. In *IEEE International Symposium on Computer Aided Control Systems Design*, pages 284–289, 2004.
- [10] A. Shapiro. On duality theory of conic linear problems. In *Semi-Infinite Programming*, chapter 7, pages 135–165. Kluwer Academic Publishers, 2001.
- [11] W. Wiesemann, D. Kuhn, and M. Sim. Distributionally robust convex optimization. *Operations Research*, 62(6):1358–1376, 2014.