K-Adaptability in Two-Stage Distributionally Robust Binary Programming

Grani A. Hanasusanto[1], Daniel Kuhn[1], and Wolfram Wiesemann[2]

[1]Risk Analytics and Optimization Chair, École Polytechnique Fédérale de Lausanne, Switzerland

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

October 26, 2015

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.

1 Introduction

We study two-stage distributionally robust programs of the form

$$\begin{array}{ll} \text{minimize} & \sup\limits_{\mathbb{P}\in\mathcal{P}} \ \mathbb{P}\text{-OCE}_U \left[\tilde{\boldsymbol{\xi}}^\top \boldsymbol{C}\, \boldsymbol{x} + \min\limits_{\boldsymbol{y}\in\mathcal{Y}} \left\{ \tilde{\boldsymbol{\xi}}^\top \boldsymbol{Q}\, \boldsymbol{y} \ : \ \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y} \leq \boldsymbol{H}\tilde{\boldsymbol{\xi}} \right\} \right] \\ \text{subject to} & \boldsymbol{x} \in \mathcal{X}, \end{array} \qquad (\mathcal{DP})$$

where

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

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 $\boldsymbol{C} \in \mathbb{R}^{Q\times N}$,

$\boldsymbol{Q} \in \mathbb{R}^{Q\times M}$, $\boldsymbol{T} \in \mathbb{R}^{L\times N}$, $\boldsymbol{W} \in \mathbb{R}^{L\times M}$ and $\boldsymbol{H} \in \mathbb{R}^{L\times Q}$. The first-stage or *here-and-now* decisions $\boldsymbol{x}$ are selected prior to the observation of the uncertain parameters $\tilde{\boldsymbol{\xi}} \in \mathbb{R}^Q$, and the second-stage or *wait-and-see* decisions $\boldsymbol{y}$ are chosen after $\tilde{\boldsymbol{\xi}}$ has been revealed. We assume that all components of $\boldsymbol{y}$ are binary, while $\boldsymbol{x}$ may 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

$$\begin{array}{ll} \text{minimize} & \sup\limits_{\mathbb{P}\in\mathcal{P}} \ \mathbb{P}\text{-OCE}_U \left[\tilde{\boldsymbol{\xi}}^\top \boldsymbol{C}\, \boldsymbol{x} + \min\limits_{k\in\mathcal{K}} \left\{ \tilde{\boldsymbol{\xi}}^\top \boldsymbol{Q}\, \boldsymbol{y}^k \ : \ \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\tilde{\boldsymbol{\xi}} \right\} \right] \\ \text{subject to} & \boldsymbol{x} \in \mathcal{X}, \ \ \boldsymbol{y}^k \in \mathcal{Y},\ k \in \mathcal{K}, \end{array} \qquad (\mathcal{DP}_K)$$

where $\mathcal{K} = \{1, \ldots, K\}$. Thus, we pre-select exactly K of the $|\mathcal{Y}| \lesssim 2^M$ possible second-stage decisions in the first stage. After the realization $\boldsymbol{\xi}$ of the random parameters $\tilde{\boldsymbol{\xi}}$ has been observed, we implement the best of these K pre-selected candidate decisions. More precisely, among all candidate decisions that are feasible for $\boldsymbol{\xi}$, we implement the one that achieves the lowest second-stage cost in scenario $\boldsymbol{\xi}$. If all K candidate decisions are infeasible for a given $\boldsymbol{\xi}$, 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 $\boldsymbol{\xi} \in \Xi$. 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 $\tilde{\boldsymbol{\xi}}$ impacts only the objective function. We close with numerical results in Section 4. The Appendix 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 $\tilde{\boldsymbol{\xi}}$. 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

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

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 $\tilde{\boldsymbol{\xi}} \in \Xi$ w.p. 1, and it constitutes a nonempty bounded polytope of the form $\Xi = \left\{ \boldsymbol{\xi} \in \mathbb{R}^Q \,:\, \boldsymbol{A}\boldsymbol{\xi} \leq \boldsymbol{b} \right\}$ with $\boldsymbol{A} \in \mathbb{R}^{R \times Q}$ and $\boldsymbol{b} \in \mathbb{R}^R$. We assume that $\boldsymbol{c} \in \mathbb{R}^S$ and that $\boldsymbol{g} : \mathbb{R}^Q \to \mathbb{R}^S$ has convex piecewise linear component functions of the form

$$g_s(\boldsymbol{\xi}) = \max_{t \in \mathcal{T}} \; \boldsymbol{g}_{st}^\top \boldsymbol{\xi} \qquad \forall s \in \mathcal{S} = \{1, \ldots, S\}\,.$$

Without loss of generality, the index t of the linear pieces ranges over the same index set $\mathcal{T} = \{1, \ldots, T\}$ 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 $\mathbb{E}_{\mathbb{P}}[g_s(\tilde{\boldsymbol{\xi}})] < \boldsymbol{c}$ for all $s \in \mathcal{S}$ for which $g_s(\boldsymbol{\xi})$ 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 $\mathbb{P}\text{-OCE}_U[\phi(\tilde{\boldsymbol{\xi}})]$ represents the expected present value of an optimized payment schedule that splits an uncertainty-affected future liability $\phi(\tilde{\boldsymbol{\xi}})$ into a fixed installment θ that is paid today and a remainder $\phi(\tilde{\boldsymbol{\xi}}) - \theta$ that is paid after the realization of $\tilde{\boldsymbol{\xi}}$ 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

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

The choice $I = 1$, $\boldsymbol{s} = 1$ and $\boldsymbol{t} = 0$ corresponds to the worst-case expected value, whereas $I = 2$, $\boldsymbol{s} = ((1-\beta)^{-1}, 0)^\top$ and $\boldsymbol{t} = (0, 0)^\top$ 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 $\tilde{\boldsymbol{\xi}}$ only enter the objective function of the two-stage robust binary program $\mathcal{DP}$:

$$\begin{array}{ll} \text{minimize} & \sup\limits_{\mathbb{P} \in \mathcal{P}} \mathbb{P}\text{-OCE}_U \left[\tilde{\boldsymbol{\xi}}^\top \boldsymbol{C}\, \boldsymbol{x} + \min\limits_{\boldsymbol{y} \in \mathcal{Y}} \left\{ \tilde{\boldsymbol{\xi}}^\top \boldsymbol{Q}\, \boldsymbol{y} \,:\, \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y} \leq \boldsymbol{h} \right\} \right] \\ \text{subject to} & \boldsymbol{x} \in \mathcal{X}, \end{array} \tag{$\mathcal{DPO}$}$$

where $\boldsymbol{h} \in \mathbb{R}^L$. 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 Appendix.

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

$$\begin{array}{ll} \text{minimize} & \sup\limits_{\mathbb{P} \in \mathcal{P}} \mathbb{P}\text{-OCE}_U \left[\tilde{\boldsymbol{\xi}}^\top \boldsymbol{C}\, \boldsymbol{x} + \min\limits_{k \in \mathcal{K}} \left\{ \tilde{\boldsymbol{\xi}}^\top \boldsymbol{Q}\, \boldsymbol{y}^k \,:\, \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{h} \right\} \right] \\ \text{subject to} & \boldsymbol{x} \in \mathcal{X},\ \ \boldsymbol{y}^k \in \mathcal{Y},\ k \in \mathcal{K} \end{array} \tag{3}$$

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

$$\begin{array}{ll} \text{minimize} & \sup\limits_{\mathbb{P}\in\mathcal{P}} \mathbb{P}\text{-OCE}_U\left[\tilde{\boldsymbol{\xi}}^\top \boldsymbol{C}\,\boldsymbol{x} + \min\limits_{k\in\mathcal{K}} \tilde{\boldsymbol{\xi}}^\top \boldsymbol{Q}\,\boldsymbol{y}^k\right] \\ \text{subject to} & \boldsymbol{x}\in\mathcal{X},\ \ \boldsymbol{y}^k\in\mathcal{Y},\ k\in\mathcal{K} \\ & \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{h} \qquad\qquad \forall k\in\mathcal{K}, \end{array} \tag{$\mathcal{DPO}_K$}$$

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.*

$$\begin{array}{ll} \text{minimize} & \alpha + \boldsymbol{c}^\top\boldsymbol{\beta} + \theta \\ \text{subject to} & \boldsymbol{x}\in\mathcal{X},\ \ \boldsymbol{y}^k\in\mathcal{Y},\ k\in\mathcal{K},\ \ \alpha\in\mathbb{R},\ \ \boldsymbol{\beta}\in\mathbb{R}^S_+ \\ & \boldsymbol{\gamma}^i\in\mathbb{R}^R_+,\ \ \boldsymbol{\delta}^i\in\mathbb{R}^K_+,\ \ \boldsymbol{\Lambda}^i\in\mathbb{R}^{S\times T}_+,\ i\in\mathcal{I} \\ & \boldsymbol{z}^{i,k}\in\mathbb{R}^M_+,\ i\in\mathcal{I} \text{ and } k\in\mathcal{K},\ \ \theta\in\mathbb{R} \\ & \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{h} \qquad\qquad \forall k\in\mathcal{K} \\ & \left.\begin{array}{l} \boldsymbol{b}^\top\boldsymbol{\gamma}^i + t_i \leq \alpha + s_i\theta,\ \ \mathbf{e}^\top\boldsymbol{\delta}^i = 1,\ \ \boldsymbol{\Lambda}^i\mathbf{e} = \boldsymbol{\beta} \\ \boldsymbol{A}^\top\boldsymbol{\gamma}^i + \sum\limits_{s\in\mathcal{S}}\sum\limits_{t\in\mathcal{T}} \Lambda^i_{st}\boldsymbol{g}_{st} = s_i\boldsymbol{C}\boldsymbol{x} + \sum\limits_{k\in\mathcal{K}} s_i\boldsymbol{Q}\boldsymbol{z}^{i,k} \\ \left.\begin{array}{l} \boldsymbol{z}^{i,k}\leq \boldsymbol{y}^k,\ \ \boldsymbol{z}^{i,k}\leq \delta^i_k\mathbf{e} \\ \boldsymbol{z}^{i,k}\geq (\delta^i_k-1)\mathbf{e} + \boldsymbol{y}^k \end{array}\right\}\ \forall k\in\mathcal{K} \end{array}\right\}\ \forall i\in\mathcal{I}. \end{array} \tag{4}$$

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 $K \geq I\cdot\min\{\dim\mathcal{Y},\,\mathrm{rk}\,\boldsymbol{Q}\} + I$ policies.*

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

3.1 Persistence

We now study the contribution of each second-stage policy $\boldsymbol{y}^k$ 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 $\boldsymbol{y}^k$ 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 $\boldsymbol{y}^k$ 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 $\tilde{\boldsymbol{\xi}}$.

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

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

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

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

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

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

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

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

Proposition 1. *Let $(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k)$ 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 $\boldsymbol{y}^k$, $k \in \mathcal{K}$, is chosen under any probability distribution $\mathbb{P} \in \mathcal{P}$ is given by the optimal value of the LP*

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

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

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

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

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

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

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

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

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

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

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

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

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

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 $\alpha - \boldsymbol{b}^\top \boldsymbol{\gamma} \geq 1$. A similar reformulation can be derived for the second constraint in (6), which concludes the proof. □

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

Proposition 2. *Let $(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k)$ be a feasible decision in problem $\mathcal{DPO}_K$ 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 $\boldsymbol{\xi}^\top \boldsymbol{Q}\,(\boldsymbol{y}^k - \boldsymbol{y}^{k'}) > 0$. Then, the minimum probability with which policy $\boldsymbol{y}^k$, $k \in \mathcal{K}$, is chosen under any probability distribution $\mathbb{P} \in \mathcal{P}$ is given by the optimal value of the LP*

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

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

$$
\begin{array}{ll}
\text{minimize} & \displaystyle\int_\Xi \mathbb{I}\left[\boldsymbol{\xi}^\top \boldsymbol{Q}\, \boldsymbol{y}^k \leq \boldsymbol{\xi}^\top \boldsymbol{Q}\, \boldsymbol{y}^{k'} \ \ \forall k' \in \mathcal{K}\right] \mathrm{d}\mu(\boldsymbol{\xi}) \\
\text{subject to} & \mu \in \mathcal{M}_+(\mathbb{R}^Q) \\
& \displaystyle\int_\Xi \mathrm{d}\mu(\boldsymbol{\xi}) = 1 \\
& \displaystyle\int_\Xi \boldsymbol{g}(\boldsymbol{\xi})\, \mathrm{d}\mu(\boldsymbol{\xi}) \leq \boldsymbol{c}
\end{array}
$$

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

$$
\begin{array}{ll}
\text{maximize} & \alpha - \boldsymbol{c}^\top \boldsymbol{\beta} \\
\text{subject to} & \alpha \in \mathbb{R},\ \ \boldsymbol{\beta} \in \mathbb{R}^S_+ \\
& \alpha - \boldsymbol{g}(\boldsymbol{\xi})^\top \boldsymbol{\beta} \leq \mathbb{I}\left[\boldsymbol{\xi}^\top \boldsymbol{Q}\, \boldsymbol{y}^k \leq \boldsymbol{\xi}^\top \boldsymbol{Q}\, \boldsymbol{y}^{k'} \ \ \forall k' \in \mathcal{K}\right] \qquad \forall \boldsymbol{\xi} \in \Xi
\end{array}
$$

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

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

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

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

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

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

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

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

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

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 $G = (V, E, \boldsymbol{\xi})$ with nodes $V = \{1, \ldots, N\}$, 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 $\tilde{\boldsymbol{\xi}}$ that is governed by an unknown probability distribution. The goal is to pre-commit to a subset of nodes P_1 before observing $\tilde{\boldsymbol{\xi}}$ and to complete P_1 to a packing P_2 after observing $\tilde{\boldsymbol{\xi}}$ 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:

$$\begin{array}{ll} \text{maximize} & \inf\limits_{\mathbb{P}\in\mathcal{P}} \mathbb{P}\text{-CVaR}_\beta\left[\max\limits_{\boldsymbol{y}\in\{0,1\}^N}\left\{\tilde{\boldsymbol{\xi}}^\top\boldsymbol{y} \,:\, \|\boldsymbol{x}-\boldsymbol{y}\|_1 \leq B,\;\; y_i+y_j\leq 1 \;\; \forall\{i,j\}\in E\right\}\right] \\ \text{subject to} & \boldsymbol{x}\in\{0,1\}^N \end{array}$$

In this formulation, the decisions $\boldsymbol{x}$ and $\boldsymbol{y}$ 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 $N \in \{10, \ldots, 50\}$ 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.

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

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 $\mathcal{O}(KN)$ 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 to determine good feasible solutions. Finally, the table presents

		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.

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 $\mathcal{DPO}$ that results from disregarding the integrality requirement for the second-stage decisions $\boldsymbol{y}$, 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 $\boldsymbol{y} \in \mathcal{Y}$ 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 $\mathcal{DPO}$ 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 $\boldsymbol{y}^k$, $k \in \mathcal{K}$, 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 $\boldsymbol{\mu}$. For this distribution, there is always an optimal solution which selects a single policy with probability 1.

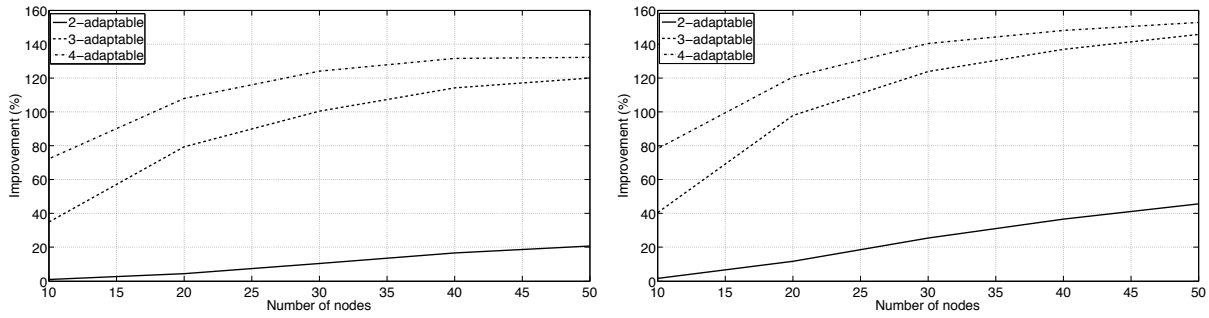

Figure 1. Improvement of the best 2-, 3- and 4-adaptable solutions determined within the set time limit over the static solutions for the vertex packing problem with $B = 0.2N$ (left) and $B = 0.4N$ (right). The figures show the improvements for problems with $N = 10, 20, \ldots, 50$ nodes as averages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

	Number of nodes N		
K	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 $\boldsymbol{y}^k$, $k \in \mathcal{K}$, are chosen in the vertex packing problem. The probabilities are averaged over 100 instances, and the policies are ordered according to increasing maximum probabilities.

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.

References

[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 adaptibility 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.

A The K-Adaptability Problem with Constraint Uncertainty

We study the generic K-adaptability problem $\mathcal{DP}_K$ introduced in Section 1 of the main text. Compared to the K-adaptability problem with objective uncertainty, the objective function in $\mathcal{DP}_K$ can be discontinuous, and the optimal value of the problem may not be attained. Moreover, evaluating the objective function in $\mathcal{DP}_K$ is strongly $\mathcal{NP}$-hard even if $\boldsymbol{g}(\boldsymbol{\xi}) = \mathbf{0}$, $\boldsymbol{c} = \mathbf{0}$, and $U(\phi) = \phi$ [6, Theorem 3], and the problem may attain a strictly higher optimal value than the two-stage robust binary program $\mathcal{DP}$ for any non-trivial number of policies $K < |\mathcal{Y}|$ [6, Theorem 4].

We first derive a MILP reformulation for problem $\mathcal{DP}_K$. Since the optimal value of $\mathcal{DP}_K$ may not be achieved, our MILP formulation constitutes an approximation whose quality is controlled by an approximation parameter ϵ. In the remainder of the section, we use the notation $\mathcal{L} = \{0, \ldots, L\}^K$, $\partial\mathcal{L} = \{\boldsymbol{\ell} \in \mathcal{L} : \boldsymbol{\ell} \not> \mathbf{0}\}$ and $\mathcal{L}_+ = \{\boldsymbol{\ell} \in \mathcal{L} : \boldsymbol{\ell} > \mathbf{0}\}$, where L is the number of second-stage constraints in problem $\mathcal{DP}$. Furthermore, we let $\boldsymbol{h}_k$ be the k-th row of the right-hand side matrix $\boldsymbol{H}$, expressed as a column vector.

Theorem 4. *The following mixed-integer bilinear program provides a lower bound to problem $\mathcal{DP}_K$:*

$$
\begin{aligned}
\text{minimize} \quad & \alpha + \boldsymbol{c}^\top \boldsymbol{\beta} + \theta \\
\text{subject to} \quad & \boldsymbol{x} \in \mathcal{X}, \;\; \boldsymbol{y}^k \in \mathcal{Y},\, k \in \mathcal{K}, \;\; \alpha \in \mathbb{R}, \;\; \boldsymbol{\beta} \in \mathbb{R}^S_+, \;\; \theta \in \mathbb{R} \\
& \left.\begin{array}{l}
\boldsymbol{\lambda}(i,\boldsymbol{\ell}) \in \Delta_K(\boldsymbol{\ell}), \;\; \boldsymbol{\gamma}(i,\boldsymbol{\ell}) \in \mathbb{R}^R_+, \;\; \boldsymbol{\Lambda}(i,\boldsymbol{\ell}) \in \mathbb{R}^{S\times T}_+ \\
\boldsymbol{\chi}^k(i,\boldsymbol{\ell}) \in \mathbb{R}^L_+,\, k \in \mathcal{K}, \;\; \boldsymbol{\psi}(i,\boldsymbol{\ell}) \in \mathbb{R}^K_+ \\
\boldsymbol{b}^\top \boldsymbol{\gamma}(i,\boldsymbol{\ell}) - \sum\limits_{\substack{k\in\mathcal{K}:\\ \ell_k=0}} (\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k)^\top \boldsymbol{\chi}^k(i,\boldsymbol{\ell}) \\
\qquad + \sum\limits_{\substack{k\in\mathcal{K}:\\ \ell_k\neq 0}} \Big([\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k]_{\ell_k} - \epsilon \Big) \psi_k(i,\boldsymbol{\ell}) + t_i \leq \alpha + s_i\theta \\
\boldsymbol{A}^\top \boldsymbol{\gamma}(i,\boldsymbol{\ell}) + \sum\limits_{s\in\mathcal{S}} \sum\limits_{t\in\mathcal{T}} \Lambda_{st}(i,\boldsymbol{\ell}) \boldsymbol{g}_{st} - \sum\limits_{\substack{k\in\mathcal{K}:\\ \ell_k=0}} \boldsymbol{H}^\top \boldsymbol{\chi}^k(i,\boldsymbol{\ell}) \\
\qquad + \sum\limits_{\substack{k\in\mathcal{K}:\\ \ell_k\neq 0}} \boldsymbol{h}_{\ell_k} \psi_k(i,\boldsymbol{\ell}) = s_i \boldsymbol{C}\boldsymbol{x} + \sum\limits_{k\in\mathcal{K}} s_i \lambda_k(i,\boldsymbol{\ell}) \boldsymbol{Q}\boldsymbol{y}^k \\
\boldsymbol{\Lambda}(i,\boldsymbol{\ell})\mathbf{e} = \boldsymbol{\beta}
\end{array}\right\} \begin{array}{l} \forall i \in \mathcal{I}, \\ \forall \boldsymbol{\ell} \in \partial\mathcal{L} \end{array} \qquad (9) \\
& \left.\begin{array}{l}
\boldsymbol{\phi}(\boldsymbol{\ell}) \in \mathbb{R}^R_+, \;\; \boldsymbol{\rho}(\boldsymbol{\ell}) \in \mathbb{R}^K_+ \\
\boldsymbol{b}^\top \boldsymbol{\phi}(\boldsymbol{\ell}) + \sum\limits_{k\in\mathcal{K}} \Big([\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k]_{\ell_k} - \epsilon \Big) \rho_k(\boldsymbol{\ell}) \leq -1 \\
\boldsymbol{A}^\top \boldsymbol{\phi}(\boldsymbol{\ell}) + \sum\limits_{k\in\mathcal{K}} \boldsymbol{h}_{\ell_k} \rho_k(\boldsymbol{\ell}) = \mathbf{0}
\end{array}\right\} \forall \boldsymbol{\ell} \in \mathcal{L}_+,
\end{aligned}
$$

where $\Delta_K(\boldsymbol{\ell}) = \{\boldsymbol{\lambda} \in \mathbb{R}^K_+ \,:\, \mathbf{e}^\top\boldsymbol{\lambda} = 1, \quad \lambda_k = 0 \;\; \forall k \in \mathcal{K} \,:\, \ell_k \neq 0\}$ *and* $\epsilon > 0$ *is a parameter that controls the approximation quality. If* $\mathcal{X} \subseteq \{0,1\}^N$*, then the problem can be reformulated as a MILP using standard Big-M techniques.*

We remark that the size of the MILP reformulation (9) scales with $|\mathcal{L}| = (L+1)^K$, that is, the problem grows exponentially with the number of policies K. This is not surprising as the evaluation of the objective function in problem $\mathcal{DP}_K$, which amounts to fixing the values of all binary variables in (9) and solving the resulting LP, is already strongly $\mathcal{NP}$-hard. We now show that the MILP reformulation (9) converges to the K-adaptability problem $\mathcal{DP}_K$.

Proposition 3. *Let* $\operatorname{dom}(9)$ *and* $\operatorname{dom}(\mathcal{DP}_K)$ *denote the effective domains of the objective functions of problems* (9) *and* $\mathcal{DP}_K$*, respectively, that is, the sets of decisions* $(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k)$ *for which the objective functions in the respective problems are finite (i.e., do not evaluate to* $+\infty$*). Then,*

(i) $\operatorname{dom}(9) = \operatorname{dom}(\mathcal{DP}_K)$ *if* ϵ *is sufficiently small, and*

(ii) over their effective domains, the objective function in (9) *converges uniformly from below to the objective function in* $\mathcal{DP}_K$ *as* ϵ *approaches* 0.

Proposition 3 implies that the optimal value of the problem (9) converges to the optimal value of problem $\mathcal{DP}_K$ as ϵ approaches 0.

A.1 Persistence

In analogy to Section 3.1 of the main text, we now study the contribution of each second-stage policy $\boldsymbol{y}^k$, $k \in \mathcal{K}$, to the objective value in problem $\mathcal{DP}_K$. We first determine the probability with which a particular second-stage decision $\boldsymbol{y}^k$ is chosen under a worst-case distribution for $\mathcal{DP}_K$.

Theorem 5. *Fix a feasible solution* $(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k)$ *to problem* (9) *and let* $(\phi, \boldsymbol{\chi}, \psi, \boldsymbol{\omega})$ *be an optimal solution to the LP*

$$
\begin{array}{lll}
\text{maximize} & \displaystyle\sum_{i\in\mathcal{I}}\sum_{\boldsymbol{\ell}\in\partial\mathcal{L}} t_i\phi(i,\boldsymbol{\ell}) + s_i(\boldsymbol{C}\boldsymbol{x})^\top\boldsymbol{\chi}(i,\boldsymbol{\ell}) + \psi(i,\boldsymbol{\ell}) & \\
\text{subject to} & \phi(i,\boldsymbol{\ell}) \in \mathbb{R}_+,\ \ \boldsymbol{\chi}(i,\boldsymbol{\ell}) \in \mathbb{R}^Q,\ \ \psi(i,\boldsymbol{\ell}) \in \mathbb{R},\ \ \boldsymbol{\omega}(i,\boldsymbol{\ell}) \in \mathbb{R}^S & \left.\begin{array}{l}\\ \\ \\ \\ \\ \\ \end{array}\right\} \begin{array}{l}\forall i \in \mathcal{I},\\ \forall \boldsymbol{\ell} \in \partial\mathcal{L}\end{array} \\
& \boldsymbol{A}\boldsymbol{\chi}(i,\boldsymbol{\ell}) \le \boldsymbol{b}\phi(i,\boldsymbol{\ell}) & \\
& \psi(i,\boldsymbol{\ell}) \le s_i(\boldsymbol{Q}\boldsymbol{y}^k)^\top\boldsymbol{\chi}(i,\boldsymbol{\ell}) \qquad \forall k \in \mathcal{K} : \ell_k = 0 & \\
& \boldsymbol{g}_{st}^\top\boldsymbol{\chi}(i,\boldsymbol{\ell}) \le \omega_s(i,\boldsymbol{\ell}) \qquad \forall s \in \mathcal{S},\ \forall t \in \mathcal{T} & \\
& (\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k)\phi(i,\boldsymbol{\ell}) \le \boldsymbol{H}\boldsymbol{\chi}(i,\boldsymbol{\ell}) \qquad \forall k \in \mathcal{K} : \ell_k = 0 & \\
& \left([\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k]_{\ell_k}\right)\phi(i,\boldsymbol{\ell}) \ge \boldsymbol{h}_{\ell_k}^\top\boldsymbol{\chi}(i,\boldsymbol{\ell}) \qquad \forall k \in \mathcal{K} : \ell_k \neq 0 & \\
& \displaystyle\sum_{i\in\mathcal{I}}\sum_{\boldsymbol{\ell}\in\partial\mathcal{L}} \phi(i,\boldsymbol{\ell}) = 1,\ \ \sum_{i\in\mathcal{I}}\sum_{\boldsymbol{\ell}\in\partial\mathcal{L}} s_i\phi(i,\boldsymbol{\ell}) = 1,\ \ \sum_{i\in\mathcal{I}}\sum_{\boldsymbol{\ell}\in\partial\mathcal{L}} \boldsymbol{\omega}(i,\boldsymbol{\ell}) \le \boldsymbol{c}. &
\end{array}
\tag{10}
$$

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

$$
\mathbb{P}^\star\left[\tilde{\boldsymbol{\xi}} = \frac{\boldsymbol{\chi}(i,\boldsymbol{\ell})}{\phi(i,\boldsymbol{\ell})}\right] = \phi(i,\boldsymbol{\ell}), \quad i \in \mathcal{I},\ \boldsymbol{\ell} \in \partial\mathcal{L} : \phi(i,\boldsymbol{\ell}) > 0,
$$

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

$$
\begin{aligned}
&\mathbb{P}^\star\left[\begin{array}{l}\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \le \boldsymbol{H}\tilde{\boldsymbol{\xi}}, \\ \tilde{\boldsymbol{\xi}}^\top\boldsymbol{Q}\,\boldsymbol{y}^k \le \tilde{\boldsymbol{\xi}}^\top\boldsymbol{Q}\,\boldsymbol{y}^{k'} \quad \forall k' \in \mathcal{K} : \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^{k'} \le \boldsymbol{H}\tilde{\boldsymbol{\xi}}\end{array}\right] \\
= &\sum_{i\in\mathcal{I}}\sum_{\substack{\boldsymbol{\ell}\in\partial\mathcal{L}:\\ \ell_k=0}} \phi(i,\boldsymbol{\ell}) \cdot \mathbb{I}\left[\boldsymbol{\chi}(i,\boldsymbol{\ell})^\top\boldsymbol{Q}\,\boldsymbol{y}^k \le \boldsymbol{\chi}(i,\boldsymbol{\ell})^\top\boldsymbol{Q}\,\boldsymbol{y}^{k'} \quad \forall k' \in \mathcal{K} : \ell_{k'} = 0\right].
\end{aligned}
$$

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

Proposition 4. *For a feasible solution $(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k)$ to problem (9), the maximum probability with which policy $\boldsymbol{y}^k$, $k \in \mathcal{K}$, is chosen under any probability distribution $\mathbb{P} \in \mathcal{P}$ is given by the optimal value of the LP*

$$
\begin{array}{ll}
\text{minimize} & \alpha + \boldsymbol{c}^\top \boldsymbol{\beta} \\
\text{subject to} & \alpha \in \mathbb{R},\ \ \boldsymbol{\beta} \in \mathbb{R}^S_+,\ \ \boldsymbol{\kappa} \in \mathbb{R}^R_+,\ \ \boldsymbol{\Phi} \in \mathbb{R}^{S \times T}_+ \\
& \left.\begin{array}{l}
\boldsymbol{\gamma}(\boldsymbol{\ell}) \in \mathbb{R}^R_+,\ \ \boldsymbol{\delta}(\boldsymbol{\ell}) \in \mathbb{R}^K_+,\ \ \boldsymbol{\Lambda}(\boldsymbol{\ell}) \in \mathbb{R}^{S \times T}_+ \\
\boldsymbol{\chi}(k', \boldsymbol{\ell}) \in \mathbb{R}^L_+,\ k' \in \mathcal{K},\ \ \boldsymbol{\psi}(\boldsymbol{\ell}) \in \mathbb{R}^K_+ \\
\alpha \geq \boldsymbol{b}^\top \boldsymbol{\gamma}(\boldsymbol{\ell}) - \displaystyle\sum_{\substack{k' \in \mathcal{K}: \\ \ell_{k'} = 0}} (\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^{k'})^\top \boldsymbol{\chi}(k', \boldsymbol{\ell}) \\
\qquad\qquad + \displaystyle\sum_{\substack{k' \in \mathcal{K}: \\ \ell_{k'} \neq 0}} \left([\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^{k'}]_{\ell_{k'}} \right) \psi_{k'}(\boldsymbol{\ell}) + 1 \\
\boldsymbol{A}^\top \boldsymbol{\gamma}(\boldsymbol{\ell}) + \displaystyle\sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}} \Lambda_{st}(\boldsymbol{\ell}) \boldsymbol{g}_{st} - \sum_{\substack{k' \in \mathcal{K}: \\ \ell_{k'} = 0}} \boldsymbol{H}^\top \boldsymbol{\chi}(k', \boldsymbol{\ell}) \\
\qquad + \displaystyle\sum_{\substack{k' \in \mathcal{K}: \\ \ell_{k'} \neq 0}} \boldsymbol{h}^\top_{\boldsymbol{\ell}_{k'}} \psi_{k'}(\boldsymbol{\ell}) = \sum_{\substack{k' \in \mathcal{K}: \\ \ell_{k'} = 0}} \boldsymbol{Q}(\boldsymbol{y}^{k'} - \boldsymbol{y}^k) \delta_{k'}(\boldsymbol{\ell}) \\
\boldsymbol{\Lambda}(\boldsymbol{\ell}) \mathbf{e} = \boldsymbol{\beta}
\end{array} \right\} \ \forall \boldsymbol{\ell} \in \partial \mathcal{L},\ \ell_k = 0 \\
& \alpha \geq \boldsymbol{b}^\top \boldsymbol{\kappa},\ \ \boldsymbol{A}^\top \boldsymbol{\kappa} + \displaystyle\sum_{s \in \mathcal{S}} \sum_{t \in \mathcal{T}} \Phi_{st} \boldsymbol{g}_{st} = \mathbf{0},\ \ \boldsymbol{\Phi}\mathbf{e} = \boldsymbol{\beta}.
\end{array}
$$

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

Proposition 5. *For a feasible solution $(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k)$ to problem (9), the minimum probability with which policy $\boldsymbol{y}^k$, $k \in \mathcal{K}$, is chosen under any probability distribution $\mathbb{P} \in \mathcal{P}$ is given by the optimal*

value of the LP

$$
\begin{aligned}
\text{maximize} \quad & \alpha - \boldsymbol{c}^\top \boldsymbol{\beta} \\
\text{subject to} \quad & \alpha \in \mathbb{R},\ \ \boldsymbol{\beta} \in \mathbb{R}^S_+,\ \ \boldsymbol{\delta} \in \mathbb{R}^R_+,\ \ \boldsymbol{\Gamma} \in \mathbb{R}^{S\times T}_+ \\
& \alpha + \boldsymbol{b}^\top \boldsymbol{\delta} \le 1,\ \ \boldsymbol{A}^\top \boldsymbol{\delta} + \sum_{s\in\mathcal{S}} \sum_{t\in\mathcal{T}} \Gamma_{st} \boldsymbol{g}_{st} = \boldsymbol{0},\ \ \boldsymbol{\Gamma}\mathbf{e} = \boldsymbol{\beta} \\
& \left.\begin{array}{l}
\boldsymbol{\gamma}(\boldsymbol{\ell}) \in \mathbb{R}^R_+,\ \ \boldsymbol{\Lambda}(\boldsymbol{\ell}) \in \mathbb{R}^{S\times T}_+,\ \ \boldsymbol{\chi}(k',\boldsymbol{\ell}) \in \mathbb{R}^L_+,\ k' \in \mathcal{K},\ \ \boldsymbol{\psi}(\boldsymbol{\ell}) \in \mathbb{R}^K_+ \\
\alpha + \boldsymbol{b}^\top \boldsymbol{\gamma}(\boldsymbol{\ell}) - \sum\limits_{\substack{k'\in\mathcal{K}:\\ \ell_{k'}=0}} (\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^{k'})^\top \boldsymbol{\chi}(k',\boldsymbol{\ell}) \\
\qquad + \sum\limits_{\substack{k'\in\mathcal{K}:\\ \ell_{k'}\neq 0}} \left([\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^{k'}]_{\ell_{k'}} \right) \psi_{k'}(\boldsymbol{\ell}) \le 0 \\
\boldsymbol{A}^\top \boldsymbol{\gamma}(\boldsymbol{\ell}) + \sum\limits_{s\in\mathcal{S}} \sum\limits_{t\in\mathcal{T}} \Lambda_{st}(\boldsymbol{\ell}) \boldsymbol{g}_{st} - \sum\limits_{\substack{k'\in\mathcal{K}:\\ \ell_{k'}=0}} \boldsymbol{H}^\top \boldsymbol{\chi}(k',\boldsymbol{\ell}) + \sum\limits_{\substack{k'\in\mathcal{K}:\\ \ell_{k'}\neq 0}} \boldsymbol{h}_{\ell_{k'}} \psi_{k'}(\boldsymbol{\ell}) = \boldsymbol{0} \\
\boldsymbol{\Lambda}(\boldsymbol{\ell})\mathbf{e} = \boldsymbol{\beta}
\end{array}\right\} \ \forall \boldsymbol{\ell} \in \partial\mathcal{L} : \ell_k \neq 0 \\
& \left.\begin{array}{l}
\boldsymbol{\kappa}(k',\boldsymbol{\ell}) \in \mathbb{R}^R_+,\ \ \pi(k',\boldsymbol{\ell}) \in \mathbb{R}_+,\ \ \boldsymbol{\Phi}(k',\boldsymbol{\ell}) \in \mathbb{R}^{S\times T}_+ \\
\boldsymbol{\phi}(k',k'',\boldsymbol{\ell}) \in \mathbb{R}^L_+,\ k'' \in \mathcal{K},\ \ \boldsymbol{\omega}(k',\boldsymbol{\ell}) \in \mathbb{R}^K_+ \\
\alpha + \boldsymbol{b}^\top \boldsymbol{\kappa}(k',\boldsymbol{\ell}) - \sum\limits_{\substack{k''\in\mathcal{K}:\\ \ell_{k''}=0}} (\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^{k''})^\top \boldsymbol{\phi}(k',k'',\boldsymbol{\ell}) \\
\qquad + \sum\limits_{\substack{k''\in\mathcal{K}:\\ \ell_{k''}\neq 0}} \left([\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^{k''}]_{\ell_{k''}} \right) \omega_{k''}(k',\boldsymbol{\ell}) \le 0 \\
\boldsymbol{A}^\top \boldsymbol{\kappa}(k',\boldsymbol{\ell}) + \sum\limits_{s\in\mathcal{S}} \sum\limits_{t\in\mathcal{T}} \Phi_{st}(k',\boldsymbol{\ell}) \boldsymbol{g}_{st} - \sum\limits_{\substack{k''\in\mathcal{K}:\\ \ell_{k''}=0}} \boldsymbol{H}^\top \boldsymbol{\phi}(k',k'',\boldsymbol{\ell}) \\
\qquad + \sum\limits_{\substack{k''\in\mathcal{K}:\\ \ell_{k''}\neq 0}} \boldsymbol{h}_{\ell_{k''}}^\top \boldsymbol{\omega}_{k''}(k',\boldsymbol{\ell}) = \boldsymbol{Q}(\boldsymbol{y}^k - \boldsymbol{y}^{k'}) \pi(k',\boldsymbol{\ell}) \\
\boldsymbol{\Phi}(k',\boldsymbol{\ell})\mathbf{e} = \boldsymbol{\beta}
\end{array}\right\} \ \begin{array}{l} \forall \boldsymbol{\ell} \in \partial\mathcal{L} : \ell_k = 0, \\ \forall k' \in \mathcal{K}(k,\boldsymbol{\ell}), \end{array}
\end{aligned}
$$

where the index set $\mathcal{K}(k,\boldsymbol{\ell})$ *is defined as*

$$
\mathcal{K}(k,\boldsymbol{\ell}) = \left\{ k' \in \mathcal{K} \,:\, \ell_{k'} = 0,\ \ \exists \boldsymbol{\xi} \in \Xi \text{ such that } \left[\begin{array}{ll}
\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \le \boldsymbol{H}\boldsymbol{\xi} & \forall k \in \mathcal{K} \,:\, \ell_k = 0 \\
[\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k]_{\boldsymbol{\ell}_k} > [\boldsymbol{H}\boldsymbol{\xi}]_{\ell_k} & \forall k \in \mathcal{K} \,:\, \ell_k \neq 0 \\
\boldsymbol{\xi}^\top \boldsymbol{Q}\,(\boldsymbol{y}^k - \boldsymbol{y}^{k'}) > 0 &
\end{array} \right] \right\}.
$$

We emphasize that the sizes of the LPs in Theorem 5 and Propositions 4 and 5 scale with $|\mathcal{L}| = (L+1)^K$, that is, all three problems grow exponentially with the number of policies K but polynomially in the problem description for any fixed K.

B Omitted Proofs

Proof of Theorem 1. Replacing $\mathbb{P}$-OCE$_U[\cdot]$ with its definition, the objective function in $\mathcal{DPO}_K$ evaluates to

$$\sup_{\mathbb{P}\in\mathcal{P}} \inf_{\theta\in\mathbb{R}} \theta + \mathbb{E}_{\mathbb{P}}\left[U\left(\tilde{\boldsymbol{\xi}}^\top \boldsymbol{C}\,\boldsymbol{x} + \min_{k\in\mathcal{K}} \tilde{\boldsymbol{\xi}}^\top \boldsymbol{Q}\,\boldsymbol{y}^k - \theta\right)\right].$$

Since the expectation is linear in $\mathbb{P}$ and convex in θ, we can use Sion's min-max theorem to reformulate $\mathcal{DPO}_K$ as

$$\begin{array}{ll} \text{minimize} & \sup\limits_{\mathbb{P}\in\mathcal{P}} \theta + \mathbb{E}_{\mathbb{P}}\left[U\left(\tilde{\boldsymbol{\xi}}^\top \boldsymbol{C}\,\boldsymbol{x} + \min\limits_{k\in\mathcal{K}} \tilde{\boldsymbol{\xi}}^\top \boldsymbol{Q}\,\boldsymbol{y}^k - \theta\right)\right] \\ \text{subject to} & \boldsymbol{x}\in\mathcal{X},\ \ \boldsymbol{y}^k\in\mathcal{Y},\, k\in\mathcal{K},\ \ \theta\in\mathbb{R} \\ & \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \le \boldsymbol{h} \qquad\qquad \forall k\in\mathcal{K}. \end{array}$$

We can express the objective function of this problem as the optimal value of the moment problem

$$\begin{array}{ll} \text{maximize} & \theta + \displaystyle\int_\Xi U\left(\boldsymbol{\xi}^\top \boldsymbol{C}\,\boldsymbol{x} + \min_{k\in\mathcal{K}} \boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y}^k - \theta\right) \mathrm{d}\mu(\boldsymbol{\xi}) \\ \text{subject to} & \mu\in\mathcal{M}_+(\mathbb{R}^Q) \\ & \displaystyle\int_\Xi \mathrm{d}\mu(\boldsymbol{\xi}) = 1 \\ & \displaystyle\int_\Xi \boldsymbol{g}(\boldsymbol{\xi})\,\mathrm{d}\mu(\boldsymbol{\xi}) \le \boldsymbol{c}. \end{array}$$

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

$$\begin{array}{ll} \text{minimize} & \alpha + \boldsymbol{c}^\top\boldsymbol{\beta} + \theta \\ \text{subject to} & \alpha\in\mathbb{R},\ \ \boldsymbol{\beta}\in\mathbb{R}^S_+ \\ & \alpha + \boldsymbol{g}(\boldsymbol{\xi})^\top\boldsymbol{\beta} \ge U\left(\boldsymbol{\xi}^\top \boldsymbol{C}\,\boldsymbol{x} + \min\limits_{k\in\mathcal{K}} \boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y}^k - \theta\right) \qquad \forall\boldsymbol{\xi}\in\Xi \end{array}$$

attains the same optimal value. Replacing U with its definition, we see that the semi-infinite constraint in this problem is equivalent to

$$\max_{\boldsymbol{\xi}\in\Xi}\left\{\max_{i\in\mathcal{I}}\left(s_i\cdot\boldsymbol{\xi}^\top \boldsymbol{C}\,\boldsymbol{x} + \min_{k\in\mathcal{K}}\ s_i\cdot\boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y}^k - s_i\theta + t_i\right) - \boldsymbol{\beta}^\top\boldsymbol{g}(\boldsymbol{\xi})\right\} \le \alpha.$$

We can exchange the order of the maximum operators to obtain

$$\max_{\boldsymbol{\xi}\in\Xi}\left\{s_i\cdot\boldsymbol{\xi}^\top \boldsymbol{C}\,\boldsymbol{x} + \min_{k\in\mathcal{K}}\ s_i\cdot\boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y}^k - s_i\theta + t_i - \boldsymbol{\beta}^\top\boldsymbol{g}(\boldsymbol{\xi})\right\} \le \alpha \quad \forall i\in\mathcal{I}. \tag{11}$$

Replacing $\boldsymbol{g}$ with its definition and employing an epigraph formulation, we can express the maximization embedded in the i-th constraint as the optimal value of the following linear program.

$$\begin{array}{lll} \text{maximize} & s_i \cdot \boldsymbol{\xi}^\top \boldsymbol{C}\,\boldsymbol{x} + \zeta - s_i\theta + t_i + \boldsymbol{\beta}^\top \boldsymbol{\eta} & \\ \text{subject to} & \boldsymbol{\xi} \in \mathbb{R}^Q,\ \ \zeta \in \mathbb{R},\ \ \boldsymbol{\eta} \in \mathbb{R}^S & \\ & \boldsymbol{A}\boldsymbol{\xi} \leq \boldsymbol{b} & \\ & \zeta \leq s_i \cdot \boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y}^k & \forall k \in \mathcal{K} \\ & \eta_s \leq -\boldsymbol{g}_{st}^\top \boldsymbol{\xi} & \forall s \in \mathcal{S},\, \forall t \in \mathcal{T} \end{array} \tag{12}$$

This reformulation exploits the fact that $\boldsymbol{\beta} \geq \mathbf{0}$. Strong linear programming duality, which applies since problem (12) is feasible, implies that (12) has the same optimal value as its dual problem,

$$\begin{array}{ll} \text{minimize} & \boldsymbol{b}^\top \boldsymbol{\gamma} - s_i\theta + t_i \\ \text{subject to} & \boldsymbol{\gamma} \in \mathbb{R}^R_+,\ \ \boldsymbol{\delta} \in \mathbb{R}^K_+,\ \ \boldsymbol{\Lambda} \in \mathbb{R}^{S\times T}_+ \\ & \boldsymbol{A}^\top \boldsymbol{\gamma} + \displaystyle\sum_{s\in\mathcal{S}} \sum_{t\in\mathcal{T}} \Lambda_{st} \boldsymbol{g}_{st} = s_i \boldsymbol{C}\boldsymbol{x} + \sum_{k\in\mathcal{K}} s_i \delta_k \boldsymbol{Q}\boldsymbol{y}^k \\ & \mathbf{e}^\top \boldsymbol{\delta} = 1,\ \ \boldsymbol{\Lambda}\mathbf{e} = \boldsymbol{\beta}. \end{array}$$

The result now follows if we replace the bilinear terms $\delta_k \boldsymbol{y}^k$ with auxiliary variables $\boldsymbol{z}^k \in \mathbb{R}^M_+$, $k \in \mathcal{K}$, subject to the constraints that

$$\boldsymbol{z}^k = \delta_k \boldsymbol{y}^k \quad \Longleftrightarrow \quad \boldsymbol{z}^k \leq \boldsymbol{y}^k,\ \ \boldsymbol{z}^k \leq \delta_k \mathbf{e},\ \ \boldsymbol{z}^k \geq (\delta_k - 1)\mathbf{e} + \boldsymbol{y}^k.$$

This reformulation exploits the fact that $\mathbf{0} \leq \boldsymbol{\delta}$, $\boldsymbol{y}^k \leq \mathbf{e}$ and that $\boldsymbol{y}^k$ is binary. □

Proof of Theorem 2. Consider the $\overline{K}$-adaptability problem with $\overline{K} = |\mathcal{Y}|$, which by construction has the same optimal value as problem $\mathcal{DP}$. By Theorem 1, this problem can be formulated as

$$\begin{array}{lll} \text{minimize} & \alpha + \boldsymbol{c}^\top \boldsymbol{\beta} + \theta & \\ \text{subject to} & \boldsymbol{x} \in \mathcal{X},\ \ \boldsymbol{y}^k \in \mathcal{Y},\, k \in \overline{\mathcal{K}},\ \ \alpha \in \mathbb{R},\ \ \boldsymbol{\beta} \in \mathbb{R}^S_+ & \\ & \boldsymbol{\gamma}^i \in \mathbb{R}^R_+,\ \ \boldsymbol{\delta}^i \in \mathbb{R}^{\overline{K}}_+,\ \ \boldsymbol{\Lambda}^i \in \mathbb{R}^{S\times T}_+,\, i \in \mathcal{I},\ \ \theta \in \mathbb{R} & \\ & \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{h} & \forall k \in \overline{\mathcal{K}} \\ & \left.\begin{array}{l} \boldsymbol{b}^\top \boldsymbol{\gamma}^i + t_i \leq \alpha + s_i\theta,\ \ \mathbf{e}^\top \boldsymbol{\delta}^i = 1,\ \ \boldsymbol{\Lambda}^i \mathbf{e} = \boldsymbol{\beta} \\ \boldsymbol{A}^\top \boldsymbol{\gamma}^i + \displaystyle\sum_{s\in\mathcal{S}} \sum_{t\in\mathcal{T}} \Lambda^i_{st} \boldsymbol{g}_{st} = s_i \boldsymbol{C}\boldsymbol{x} + \sum_{k\in\overline{\mathcal{K}}} s_i \delta^i_k \boldsymbol{Q}\boldsymbol{y}^k \end{array}\right\} & \forall i \in \mathcal{I}, \end{array} \tag{13}$$

where $\overline{\mathcal{K}} = \{1, \ldots, \overline{K}\}$ and where we have replaced the auxiliary variables $\boldsymbol{z}^{i,k}$ in problem (4) with $\delta_k^i \boldsymbol{y}^k$. Fix any optimal solution $(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k, \alpha, \boldsymbol{\beta}, \{\boldsymbol{\gamma}^i\}_i, \{\boldsymbol{\delta}^i\}_i, \{\boldsymbol{\Lambda}^i\}_i, \theta)$ to this problem. Since $\mathbf{e}^\top \boldsymbol{\delta}^i = 1$, the expression $\sum_{k \in \overline{\mathcal{K}}} s_i \delta_k^i \boldsymbol{Q}\boldsymbol{y}^k$ in the last constraint of (13) can be viewed either as a convex combination of the terms $s_i \boldsymbol{Q}\boldsymbol{y}^1, \ldots, s_i \boldsymbol{Q}\boldsymbol{y}^{\overline{K}} \in \mathbb{R}^Q$ or as a convex combination of the terms $s_i \boldsymbol{y}^1, \ldots, s_i \boldsymbol{y}^{\overline{K}} \in \mathbb{R}^M$, multiplied from the left with the matrix $\boldsymbol{Q}$. Since $\boldsymbol{\delta}^i$ does not participate in any other equations, Carathéodory's Theorem allows us to assume that at most $\min\{\dim \mathcal{Y}, \operatorname{rk} \boldsymbol{Q}\} + 1$ components of $\boldsymbol{\delta}^i$ are nonzero. If we further relax problem (13) by removing those constraints

$$\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{h}, \quad k \in \mathcal{K}$$

for which $\delta_k^i = 0$ for all $i \in \mathcal{I}$, then we recover the K-adaptability problem with

$$K \leq I \cdot (\min\{\dim \mathcal{Y}, \operatorname{rk} \boldsymbol{Q}\} + 1).$$

By construction, this problem still attains the same optimal value as the problem $\mathcal{DP}$. □

The proof of Theorem 3 requires the following auxiliary result, which we prove first.

Lemma 1. *If a feasible solution $(\boldsymbol{\phi}, \{\boldsymbol{\chi}^i\}_i, \boldsymbol{\psi}, \{\boldsymbol{\omega}^i\}_i)$ to problem (5) satisfies $\phi_i = 0$ for some $i \in \mathcal{I}$, then it also satisfies $\boldsymbol{\chi}^i = \mathbf{0}$.*

Proof. Assume to the contrary that there is a feasible solution $(\boldsymbol{\phi}, \{\boldsymbol{\chi}^i\}_i, \boldsymbol{\psi}, \{\boldsymbol{\omega}^i\}_i)$ to problem (5) such that $\phi_i = 0$ and $\boldsymbol{\chi}^i \neq \mathbf{0}$ for some $i \in \mathcal{I}$. From the constraints in (5) we then conclude that $\boldsymbol{A}\boldsymbol{\chi}^i \leq \mathbf{0}$. Choose any $\boldsymbol{\xi} \in \Xi$, which exists since Ξ is nonempty. We have $\boldsymbol{A}(\boldsymbol{\xi} + \lambda \boldsymbol{\chi}^i) \leq \boldsymbol{b}$, that is, $\boldsymbol{\xi} + \lambda \boldsymbol{\chi}^i \in \Xi$, for all $\lambda \in \mathbb{R}_+$. This contradicts the fact that Ξ is bounded, and we thus conclude that our assumption is wrong, that is, the assertion of the lemma indeed holds true. □

Proof of Theorem 3. We first show that $\mathbb{P}^\star$ is an element of the ambiguity set $\mathcal{P}$. To this end, we note that $\boldsymbol{\chi}^i/\phi_i \in \Xi$ for all $i \in \mathcal{I}$ that satisfy $\phi_i > 0$ since the constraints in (5) imply that

$$\boldsymbol{A}\boldsymbol{\chi}^i \leq \boldsymbol{b}\phi_i \quad \Longleftrightarrow \quad \boldsymbol{A}\frac{\boldsymbol{\chi}^i}{\phi_i} \leq \boldsymbol{b} \qquad \forall i \in \mathcal{I} : \phi_i > 0.$$

We furthermore observe that for all $s \in \mathcal{S}$,

$$\mathbb{E}_{\mathbb{P}^\star}\left[\max_{t \in \mathcal{T}} \boldsymbol{g}_{st}^\top \tilde{\boldsymbol{\xi}}\right] = \sum_{\substack{i \in \mathcal{I}: \\ \phi_i > 0}} \phi_i \max_{t \in \mathcal{T}} \frac{\boldsymbol{g}_{st}^\top \boldsymbol{\chi}^i}{\phi_i} = \sum_{i \in \mathcal{I}} \max_{t \in \mathcal{T}} \boldsymbol{g}_{st}^\top \boldsymbol{\chi}^i \leq \sum_{i \in \mathcal{I}} \omega_s^i \leq c_s.$$

Here, the first identity follows from the definition of $\mathbb{P}^\star$, the second one from Lemma 1, and the last two inequalities are due to the constraints in (5). Hence, we have $\mathbb{P}^\star \in \mathcal{P}$.

To show that $\mathbb{P}^\star$ is in fact a worst-case distribution for $(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k)$ in problem $\mathcal{DPO}_K$, we first show that the optimized certainty equivalent of $(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k)$ under $\mathbb{P}^\star$ is bounded below by the optimal value of problem (5). Afterwards, we prove that the optimal value of problem (5) is identical to the worst-case (maximum) optimized certainty equivalent of $(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k)$ in $\mathcal{DPO}_K$.

In view of the first step, we observe that

$$
\begin{array}{rl}
& \mathbb{P}^\star\text{-OCE}_U \left[\tilde{\boldsymbol{\xi}}^\top \boldsymbol{C}\, \boldsymbol{x} + \min\limits_{k\in\mathcal{K}} \ \tilde{\boldsymbol{\xi}}^\top \boldsymbol{Q}\, \boldsymbol{y}^k \right] \\
= & \inf\limits_{\theta\in\mathbb{R}} \ \theta + \mathbb{E}_{\mathbb{P}^\star} \left[\max\limits_{i'\in\mathcal{I}} \left\{ s_{i'} \tilde{\boldsymbol{\xi}}^\top \boldsymbol{C}\, \boldsymbol{x} + \min\limits_{k\in\mathcal{K}} \ s_{i'} \tilde{\boldsymbol{\xi}}^\top \boldsymbol{Q}\, \boldsymbol{y}^k - s_{i'}\theta + t_{i'} \right\} \right] \\
= & \inf\limits_{\theta\in\mathbb{R}} \ \theta + \sum\limits_{\substack{i\in\mathcal{I}: \\ \phi_i>0}} \phi_i \max\limits_{i'\in\mathcal{I}} \left\{ \dfrac{s_{i'} (\boldsymbol{\chi}^i)^\top \boldsymbol{C}\, \boldsymbol{x}}{\phi_i} + \min\limits_{k\in\mathcal{K}} \ \dfrac{s_{i'} (\boldsymbol{\chi}^i)^\top \boldsymbol{Q}\, \boldsymbol{y}^k}{\phi_i} - s_{i'}\theta + t_{i'} \right\} \\
\geq & \inf\limits_{\theta\in\mathbb{R}} \ \theta + \sum\limits_{\substack{i\in\mathcal{I}: \\ \phi_i>0}} \phi_i \dfrac{s_i (\boldsymbol{\chi}^i)^\top \boldsymbol{C}\, \boldsymbol{x}}{\phi_i} + \min\limits_{k\in\mathcal{K}} \ \phi_i \dfrac{s_i (\boldsymbol{\chi}^i)^\top \boldsymbol{Q}\, \boldsymbol{y}^k}{\phi_i} - s_i\phi_i\theta + \phi_i t_i \\
= & \inf\limits_{\theta\in\mathbb{R}} \ \theta + \sum\limits_{i\in\mathcal{I}} s_i (\boldsymbol{\chi}^i)^\top \boldsymbol{C}\, \boldsymbol{x} + \min\limits_{k\in\mathcal{K}} \ s_i (\boldsymbol{\chi}^i)^\top \boldsymbol{Q}\, \boldsymbol{y}^k - s_i\phi_i\theta + \phi_i t_i \\
\geq & \inf\limits_{\theta\in\mathbb{R}} \ \theta + \sum\limits_{i\in\mathcal{I}} s_i (\boldsymbol{\chi}^i)^\top \boldsymbol{C}\, \boldsymbol{x} + \psi_i + \phi_i t_i - \theta \ = \ \sum\limits_{i\in\mathcal{I}} s_i (\boldsymbol{\chi}^i)^\top \boldsymbol{C}\, \boldsymbol{x} + \psi_i + \phi_i t_i,
\end{array}
$$

where the first identity holds because of the definition of the optimized certainty equivalent and the disutility function U, the second one is due to the definition of $\mathbb{P}^\star$, the third one follows from Lemma 1, and the last inequality is implied by the constraints in (5). By definition of $\boldsymbol{\phi}$, $\boldsymbol{\chi}^i$, $\boldsymbol{\psi}$ and $\boldsymbol{\omega}^i$, the last expression is the optimal value of problem (5), which proves the first statement.

To prove that the optimal value of problem (5) is identical to the worst-case expected disutility of $(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k)$ in $\mathcal{DPO}_K$, we consider the dual of (5):

$$
\begin{array}{ll}
\text{minimize} & \alpha + \boldsymbol{c}^\top \boldsymbol{\beta} + \theta \\
\text{subject to} & \alpha \in \mathbb{R}, \ \ \boldsymbol{\beta} \in \mathbb{R}^S_+, \ \ \theta \in \mathbb{R} \\
& \left.
\begin{array}{l}
\boldsymbol{\gamma}^i \in \mathbb{R}^R_+, \ \ \boldsymbol{\delta}^i \in \mathbb{R}^K_+, \ \ \boldsymbol{\Lambda}^i \in \mathbb{R}^{S\times T}_+ \\
\boldsymbol{b}^\top \boldsymbol{\gamma}^i + t_i \leq \alpha + s_i\theta, \ \ \mathbf{e}^\top \boldsymbol{\delta}^i = 1, \ \ \boldsymbol{\Lambda}^i \mathbf{e} = \boldsymbol{\beta} \\
\boldsymbol{A}^\top \boldsymbol{\gamma}^i + \sum\limits_{s\in\mathcal{S}} \sum\limits_{t\in\mathcal{T}} \Lambda^i_{st} \boldsymbol{g}_{st} = s_i \boldsymbol{C}\boldsymbol{x} + \sum\limits_{k\in\mathcal{K}} s_i \boldsymbol{Q} \delta^i_k \boldsymbol{y}^k
\end{array}
\right\} \ \forall i \in \mathcal{I}
\end{array}
\tag{14}
$$

Strong linear programming duality, which applies since problem (14) is feasible, implies that the optimal value of this problem is identical to the optimal value of (5). An inspection of the proof of Theorem 1 reveals, however, that the optimal value of problem (14) is the value of the objective function of problem $\mathcal{DPO}_K$ at $(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k)$.

To compute the probability that policy $\boldsymbol{y}^k$, $k \in \mathcal{K}$, is chosen under $\mathbb{P}^\star$, we observe that

$$\begin{aligned}
& \mathbb{P}^\star \left[\tilde{\boldsymbol{\xi}}^\top \boldsymbol{Q}\, \boldsymbol{y}^k \leq \tilde{\boldsymbol{\xi}}^\top \boldsymbol{Q}\, \boldsymbol{y}^{k'} \;\; \forall k' \in \mathcal{K} \right] \\
= & \sum_{\substack{i \in \mathcal{I}: \\ \phi_i > 0}} \phi_i \cdot \mathbb{I} \left[\frac{(\boldsymbol{\chi}^i)^\top \boldsymbol{Q}\, \boldsymbol{y}^k}{\phi_i} \leq \frac{(\boldsymbol{\chi}^i)^\top \boldsymbol{Q}\, \boldsymbol{y}^{k'}}{\phi_i} \;\; \forall k' \in \mathcal{K} \right] \\
= & \sum_{i \in \mathcal{I}} \phi_i \cdot \mathbb{I} \left[(\boldsymbol{\chi}^i)^\top \boldsymbol{Q}\, \boldsymbol{y}^k \leq (\boldsymbol{\chi}^i)^\top \boldsymbol{Q}\, \boldsymbol{y}^{k'} \;\; \forall k' \in \mathcal{K} \right].
\end{aligned}$$

This completes the proof. □

For the proofs of the statements in Section A of this Appendix, we introduce the sets

$$\Xi(\boldsymbol{\ell}) = \left\{ \boldsymbol{\xi} \in \Xi : \begin{array}{ll} \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi} & \forall k \in \mathcal{K} : \ell_k = 0 \\ [\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k]_{\ell_k} > [\boldsymbol{H}\boldsymbol{\xi}]_{\ell_k} & \forall k \in \mathcal{K} : \ell_k \neq 0 \end{array} \right\} \quad \text{for } \boldsymbol{\ell} \in \mathcal{L},$$

as well as their parameterized closed inner approximations

$$\Xi_\epsilon(\boldsymbol{\ell}) = \left\{ \boldsymbol{\xi} \in \Xi : \begin{array}{ll} \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi} & \forall k \in \mathcal{K} : \ell_k = 0 \\ [\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k]_{\ell_k} \geq [\boldsymbol{H}\boldsymbol{\xi}]_{\ell_k} + \epsilon & \forall k \in \mathcal{K} : \ell_k \neq 0 \end{array} \right\} \quad \text{for } \boldsymbol{\ell} \in \mathcal{L},\ \epsilon > 0.$$

Proof of Theorem 4. By construction, for fixed $(\boldsymbol{x}, \{\boldsymbol{y}\}_k)$ the objective value of $\mathcal{DP}_K$ is finite if and only if $\Xi(\boldsymbol{\ell}) = \emptyset$ for all $\boldsymbol{\ell} \in \mathcal{L}_+$. In that case, we can proceed as in the proof of Theorem 1 to conclude that the objective function of $\mathcal{DP}_K$ equals the optimal value of the problem

$$\begin{array}{ll}
\text{minimize} & \alpha + \boldsymbol{c}^\top \boldsymbol{\beta} + \theta \\
\text{subject to} & \alpha \in \mathbb{R},\ \ \boldsymbol{\beta} \in \mathbb{R}^S_+ \\
& \alpha + \boldsymbol{g}(\boldsymbol{\xi})^\top \boldsymbol{\beta} \geq U \left(\boldsymbol{\xi}^\top \boldsymbol{C}\, \boldsymbol{x} + \min_{k \in \mathcal{K}} \left\{ \boldsymbol{\xi}^\top \boldsymbol{Q}\, \boldsymbol{y}^k \,:\, \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi} \right\} - \theta \right) \qquad \forall \boldsymbol{\xi} \in \Xi.
\end{array} \tag{15}$$

The semi-infinite constraint in this problem can be reformulated as follows:

$$\max_{\boldsymbol{\xi}\in\Xi}\left\{U\left(\boldsymbol{\xi}^\top\boldsymbol{C}\,\boldsymbol{x}+\min_{k\in\mathcal{K}}\left\{\boldsymbol{\xi}^\top\boldsymbol{Q}\,\boldsymbol{y}^k\;:\;\boldsymbol{T}\boldsymbol{x}+\boldsymbol{W}\boldsymbol{y}^k\leq\boldsymbol{H}\boldsymbol{\xi}\right\}-\theta\right)-\boldsymbol{g}(\boldsymbol{\xi})^\top\boldsymbol{\beta}\right\}\leq\alpha$$

$$\iff\quad\max_{\boldsymbol{\ell}\in\partial\mathcal{L}}\max_{\boldsymbol{\xi}\in\Xi(\boldsymbol{\ell})}\left\{U\left(\boldsymbol{\xi}^\top\boldsymbol{C}\,\boldsymbol{x}+\min_{k\in\mathcal{K}}\left\{\boldsymbol{\xi}^\top\boldsymbol{Q}\,\boldsymbol{y}^k\;:\;\boldsymbol{T}\boldsymbol{x}+\boldsymbol{W}\boldsymbol{y}^k\leq\boldsymbol{H}\boldsymbol{\xi}\right\}-\theta\right)-\boldsymbol{g}(\boldsymbol{\xi})^\top\boldsymbol{\beta}\right\}\leq\alpha$$

$$\iff\max_{\boldsymbol{\ell}\in\partial\mathcal{L}}\max_{\boldsymbol{\xi}\in\Xi(\boldsymbol{\ell})}\max_{i\in\mathcal{I}}\left\{s_i\cdot\boldsymbol{\xi}^\top\boldsymbol{C}\,\boldsymbol{x}+s_i\cdot\min_{k\in\mathcal{K}}\left\{\boldsymbol{\xi}^\top\boldsymbol{Q}\,\boldsymbol{y}^k\;:\;\boldsymbol{T}\boldsymbol{x}+\boldsymbol{W}\boldsymbol{y}^k\leq\boldsymbol{H}\boldsymbol{\xi}\right\}-s_i\theta+t_i-\boldsymbol{g}(\boldsymbol{\xi})^\top\boldsymbol{\beta}\right\}\leq\alpha$$

$$\iff\max_{\boldsymbol{\ell}\in\partial\mathcal{L}}\max_{\boldsymbol{\xi}\in\Xi(\boldsymbol{\ell})}\max_{i\in\mathcal{I}}\left\{s_i\cdot\boldsymbol{\xi}^\top\boldsymbol{C}\,\boldsymbol{x}+s_i\cdot\min_{k\in\mathcal{K}:\ell_k=0}\boldsymbol{\xi}^\top\boldsymbol{Q}\,\boldsymbol{y}^k-s_i\theta+t_i-\boldsymbol{g}(\boldsymbol{\xi})^\top\boldsymbol{\beta}\right\}\leq\alpha\tag{16}$$

$$\iff\max_{\boldsymbol{\ell}\in\partial\mathcal{L}}\max_{\boldsymbol{\xi}\in\Xi(\boldsymbol{\ell})}\max_{i\in\mathcal{I}}\left\{s_i\cdot\boldsymbol{\xi}^\top\boldsymbol{C}\,\boldsymbol{x}+s_i\cdot\min_{\boldsymbol{\lambda}\in\Delta_K(\boldsymbol{\ell})}\sum_{k\in\mathcal{K}}\lambda_k\cdot\boldsymbol{\xi}^\top\boldsymbol{Q}\,\boldsymbol{y}^k-s_i\theta+t_i-\boldsymbol{g}(\boldsymbol{\xi})^\top\boldsymbol{\beta}\right\}\leq\alpha\tag{17}$$

Here, the first equivalence follows from the fact that $\Xi=\bigcup_{\boldsymbol{\ell}\in\mathcal{L}}\Xi(\boldsymbol{\ell})$ and the assumption that $\Xi(\boldsymbol{\ell})=\emptyset$ for all $\boldsymbol{\ell}\in\mathcal{L}_+$, the second one is due to the definition of the disutility function U and the last one follows from the definition of $\Xi(\boldsymbol{\ell})$ and $\Delta_K(\boldsymbol{\ell})$. The expression on the right-hand side of the last equivalence, however, is satisfied if and only if there exist $\boldsymbol{\lambda}(\boldsymbol{\ell})\in\Delta_K(\boldsymbol{\ell})$, $\boldsymbol{\ell}\in\partial\mathcal{L}$, such that

$$s_i\cdot\boldsymbol{\xi}^\top\boldsymbol{C}\,\boldsymbol{x}+s_i\sum_{k\in\mathcal{K}}\lambda_k(\boldsymbol{\ell})\cdot\boldsymbol{\xi}^\top\boldsymbol{Q}\,\boldsymbol{y}^k-s_i\theta+t_i\leq\alpha+\boldsymbol{g}(\boldsymbol{\xi})^\top\boldsymbol{\beta}\qquad\forall i\in\mathcal{I},\;\forall\boldsymbol{\ell}\in\partial\mathcal{L},\;\forall\boldsymbol{\xi}\in\Xi(\boldsymbol{\ell}).$$

Our previous derivations thus imply that the problem $\mathcal{DP}_K$ is equivalent to

$$\begin{array}{ll}\text{minimize} & \alpha+\boldsymbol{c}^\top\boldsymbol{\beta}+\theta\\[2ex]\text{subject to} & \boldsymbol{x}\in\mathcal{X},\;\;\boldsymbol{y}^k\in\mathcal{Y},\,k\in\mathcal{K},\;\;\alpha\in\mathbb{R},\;\;\boldsymbol{\beta}\in\mathbb{R}^S_+\\ & s_i\cdot\boldsymbol{\xi}^\top\boldsymbol{C}\,\boldsymbol{x}+s_i\displaystyle\sum_{k\in\mathcal{K}}\lambda_k(\boldsymbol{\ell})\cdot\boldsymbol{\xi}^\top\boldsymbol{Q}\,\boldsymbol{y}^k-s_i\theta+t_i\leq\alpha+\boldsymbol{g}(\boldsymbol{\xi})^\top\boldsymbol{\beta}\quad\forall i\in\mathcal{I},\;\forall\boldsymbol{\ell}\in\partial\mathcal{L},\;\forall\boldsymbol{\xi}\in\Xi(\boldsymbol{\ell})\\ & \Xi(\boldsymbol{\ell})=\emptyset\qquad\qquad\qquad\qquad\qquad\qquad\qquad\qquad\forall\boldsymbol{\ell}\in\mathcal{L}_+.\end{array}$$

The assertion now follows if we replace the sets $\Xi(\boldsymbol{\ell})$ with their inner approximations $\Xi_\epsilon(\boldsymbol{\ell})$, $\boldsymbol{\ell}\in\mathcal{L}$, and reformulate the constraints as in [6, Theorem 5]. □

Proof of Proposition 3. As for assertion (i), it follows from [6, Lemma 1] that for small $\epsilon>0$, the inner approximations $\Xi_\epsilon(\boldsymbol{\ell})$, $\boldsymbol{\ell}\in\mathcal{L}$, are nonempty if and only if the exact uncertainty sets $\Xi(\boldsymbol{\ell})$ are. Thus, the proof of Theorem 4 implies that $\operatorname{dom}(9)=\operatorname{dom}(\mathcal{DP}_K)$ for sufficiently small $\epsilon>0$.

As for assertion (ii), fix any $(\boldsymbol{x},\{\boldsymbol{y}^k\}_k)\in\operatorname{dom}(\mathcal{DP}_K)$ and $\epsilon>0$ such that $\operatorname{dom}(9)=\operatorname{dom}(\mathcal{DP}_K)$. It then follows from (15) and (17) in the proof of Theorem 4 that the objective function value of

$(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k)$ in problem $\mathcal{DP}_K$ is representable as

$$\begin{array}{lll} \varphi = \text{minimize} & \boldsymbol{c}^\top \boldsymbol{\beta} + \theta + \max\limits_{\boldsymbol{\ell}\in\mathcal{L}} \max\limits_{\boldsymbol{\xi}\in\Xi(\boldsymbol{\ell})} \max\limits_{i\in\mathcal{I}} -\boldsymbol{g}(\boldsymbol{\xi})^\top \boldsymbol{\beta} + s_i \cdot \boldsymbol{\xi}^\top \boldsymbol{C}\, \boldsymbol{x} + s_i \min\limits_{\substack{k\in\mathcal{K}:\\ \boldsymbol{\ell}_k = 0}} \boldsymbol{\xi}^\top \boldsymbol{Q}\, \boldsymbol{y}^k - s_i\theta + t_i \\ \text{subject to} & \boldsymbol{\beta} \in \mathbb{R}^S_+,\ \theta \in \mathbb{R}. \end{array} \tag{18}$$

Likewise, Theorem 4 implies that the objective value of $(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k)$ in the ϵ-approximation (9) can be expressed as

$$\begin{array}{lll} \varphi_\epsilon = \text{minimize} & \boldsymbol{c}^\top \boldsymbol{\beta} + \theta + \max\limits_{\boldsymbol{\ell}\in\mathcal{L}} \max\limits_{\boldsymbol{\xi}\in\Xi_\epsilon(\boldsymbol{\ell})} \max\limits_{i\in\mathcal{I}} -\boldsymbol{g}(\boldsymbol{\xi})^\top \boldsymbol{\beta} + s_i \cdot \boldsymbol{\xi}^\top \boldsymbol{C}\, \boldsymbol{x} + s_i \min\limits_{\substack{k\in\mathcal{K}:\\ \boldsymbol{\ell}_k = 0}} \boldsymbol{\xi}^\top \boldsymbol{Q}\, \boldsymbol{y}^k - s_i\theta + t_i \\ \text{subject to} & \boldsymbol{\beta} \in \mathbb{R}^S_+,\ \theta \in \mathbb{R}. \end{array} \tag{18$_\epsilon$}$$

Let $(\boldsymbol{\beta}^\epsilon, \theta^\epsilon)$ be a minimizer of (18$_\epsilon$). Similar derivations as in [6, Proposition 2] imply that

$$\begin{aligned} 0 \leq \varphi - \varphi_\epsilon &\leq \max_{i\in\mathcal{I}} \max_{\boldsymbol{\ell}\in\mathcal{L}} \max_{\boldsymbol{\xi}\in\Xi(\boldsymbol{\ell})} \min_{\boldsymbol{\xi}'\in\Xi(\boldsymbol{\ell})} \max_{\substack{k\in\mathcal{K}:\\ \ell_k = 0}} \left\{ \sum_{s\in\mathcal{S}} \beta_s^\epsilon \max_{t\in\mathcal{T}} \boldsymbol{g}_{st}^\top (\boldsymbol{\xi}' - \boldsymbol{\xi}) + s_i (\boldsymbol{\xi} - \boldsymbol{\xi}')^\top (\boldsymbol{C}\, \boldsymbol{x} + \boldsymbol{Q}\, \boldsymbol{y}^k) \right\} \\ &\leq \left(\max_{\boldsymbol{\ell}\in\mathcal{L}} \max_{\boldsymbol{\xi}\in\Xi(\boldsymbol{\ell})} \min_{\boldsymbol{\xi}'\in\Xi(\boldsymbol{\ell})} \|\boldsymbol{\xi}' - \boldsymbol{\xi}\| \right) \cdot \left(\max_{i\in\mathcal{I}} \max_{\substack{k\in\mathcal{K}:\\ \ell_k = 0}} \left\{ \sum_{s\in\mathcal{S}} \beta_s^\epsilon \max_{t\in\mathcal{T}} \|\boldsymbol{g}_{st}\| + s_i \|\boldsymbol{C}\, \boldsymbol{x} + \boldsymbol{Q}\, \boldsymbol{y}^k\| \right\} \right), \end{aligned}$$

where the last inequality invokes weak duality and the Cauchy-Schwarz inequality. It follows from [6, Lemma 1] that the first of the two factors in the last expression can be made arbitrarily small by choosing $\epsilon > 0$ appropriately. The second product term, on the other hand, is bounded over $(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k) \in \operatorname{dom}(\mathcal{DP}_K)$ since $\mathcal{X}$ and $\mathcal{Y}$ are bounded, while the set of optimal solutions $(\boldsymbol{\beta}^\epsilon, \theta^\epsilon)$ to problem (18$_\epsilon$) can without loss of generality be bounded uniformly over $(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k) \in \operatorname{dom}(\mathcal{DP}_K)$. For sufficiently small $\epsilon > 0$, we can thus upper bound the difference $\varphi - \varphi_\epsilon$ uniformly over $(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k) \in \operatorname{dom}(\mathcal{DP}_K)$ by an arbitrarily small constant, which concludes the proof. □

Proof of Theorem 5. The proof is analogous to the proof of Theorem 3. Indeed, using the same reasoning as in that proof, we can show that $\mathbb{P}^\star$ is an element of the ambiguity set $\mathcal{P}$.

To prove that $\mathbb{P}^\star$ is a worst-case distribution for $(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k)$ in problem $\mathcal{DP}_K$, we first show that the optimized certainty equivalent of $(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k)$ under $\mathbb{P}^\star$ is bounded below by the optimal value of problem (10). Indeed, a similar argument as in the proof of Theorem 3 shows that

$$\mathbb{P}^\star\text{-OCE}_U \left[\tilde{\boldsymbol{\xi}}^\top \boldsymbol{C}\, \boldsymbol{x} + \min_{k\in\mathcal{K}} \left\{ \tilde{\boldsymbol{\xi}}^\top \boldsymbol{Q}\, \boldsymbol{y}^k \,:\, \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\tilde{\boldsymbol{\xi}} \right\} \right] \;\geq\; \sum_{i\in\mathcal{I},\, \boldsymbol{\ell}\in\partial\mathcal{L}} t_i \phi(i,\boldsymbol{\ell}) + s_i (\boldsymbol{C}\, \boldsymbol{x})^\top \boldsymbol{\chi}(i,\boldsymbol{\ell}) + \psi(i,\boldsymbol{\ell}),$$

and the last expression is the optimal value of problem (10) by definition of ϕ, $\boldsymbol{\chi}$ and ψ.

We now prove that the optimal value of problem (10) is identical to the worst-case (maximum) optimized certainty equivalent of $(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k)$ in $\mathcal{DP}_K$. To this end, we consider the dual of (10):

$$
\begin{array}{ll}
\text{minimize} & \alpha + \boldsymbol{c}^\top \boldsymbol{\beta} + \theta \\
\text{subject to} & \alpha \in \mathbb{R},\ \ \boldsymbol{\beta} \in \mathbb{R}^S_+,\ \ \theta \in \mathbb{R} \\
& \left.\begin{array}{l}
\boldsymbol{\lambda}(i,\boldsymbol{\ell}) \in \Delta_K(\boldsymbol{\ell}),\ \ \boldsymbol{\gamma}(i,\boldsymbol{\ell}) \in \mathbb{R}^R_+,\ \ \boldsymbol{\Lambda}(i,\boldsymbol{\ell}) \in \mathbb{R}^{S\times T}_+ \\
\boldsymbol{\chi}^k(i,\boldsymbol{\ell}) \in \mathbb{R}^L_+,\ k \in \mathcal{K},\ \ \boldsymbol{\psi}(i,\boldsymbol{\ell}) \in \mathbb{R}^K_+ \\
\boldsymbol{b}^\top \boldsymbol{\gamma}(i,\boldsymbol{\ell}) - \displaystyle\sum_{\substack{k\in\mathcal{K}:\\ \ell_k = 0}} (\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k)^\top \boldsymbol{\chi}^k(i,\boldsymbol{\ell}) \\
\qquad + \displaystyle\sum_{\substack{k\in\mathcal{K}:\\ \ell_k \neq 0}} \Big([\boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k]_{\ell_k} \Big) \psi_k(i,\boldsymbol{\ell}) + t_i \leq \alpha + s_i \theta \\
\boldsymbol{A}^\top \boldsymbol{\gamma}(i,\boldsymbol{\ell}) + \displaystyle\sum_{s\in\mathcal{S}} \sum_{t\in\mathcal{T}} \Lambda_{st}(i,\boldsymbol{\ell}) \boldsymbol{g}_{st} - \sum_{\substack{k\in\mathcal{K}:\\ \ell_k = 0}} \boldsymbol{H}^\top \boldsymbol{\chi}^k(i,\boldsymbol{\ell}) \\
\qquad + \displaystyle\sum_{\substack{k\in\mathcal{K}:\\ \ell_k \neq 0}} \boldsymbol{h}_{\ell_k} \psi_k(i,\boldsymbol{\ell}) = s_i \boldsymbol{C}\boldsymbol{x} + \sum_{k\in\mathcal{K}} s_i \lambda_k(i,\boldsymbol{\ell}) \boldsymbol{Q}\boldsymbol{y}^k \\
\boldsymbol{\Lambda}(i,\boldsymbol{\ell})\mathbf{e} = \boldsymbol{\beta}
\end{array}\right\} \begin{array}{l} \forall i \in \mathcal{I}, \\ \forall \boldsymbol{\ell} \in \partial\mathcal{L}. \end{array}
\end{array}
\tag{19}
$$

Strong linear programming duality, which holds since problem (19) can be shown to be feasible, implies that the optimal value of this problem is identical to the optimal value of (10). An inspection of the proof of Theorem 4 reveals, however, that the optimal value of problem (19) is the value of the objective function of problem $\mathcal{DP}_K$ at $(\boldsymbol{x}, \{\boldsymbol{y}^k\}_k)$.

Finally, we can compute the probability with which policy $\boldsymbol{y}^k$, $k \in \mathcal{K}$, is chosen under $\mathbb{P}^\star$ in the same way as in the proof of Theorem 3. This concludes the proof. □

Proof of Proposition 4. The maximum probability with which policy $\boldsymbol{y}^k$ is chosen under any probability distribution $\mathbb{P} \in \mathcal{P}$ is given by the optimal value of the following moment problem.

$$
\begin{array}{ll}
\text{maximize} & \displaystyle\int_\Xi \mathbb{I} \left[\begin{array}{l} \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\tilde{\boldsymbol{\xi}}, \\ \tilde{\boldsymbol{\xi}}^\top \boldsymbol{Q}\, \boldsymbol{y}^k \leq \tilde{\boldsymbol{\xi}}^\top \boldsymbol{Q}\, \boldsymbol{y}^{k'} \quad \forall k' \in \mathcal{K} : \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^{k'} \leq \boldsymbol{H}\tilde{\boldsymbol{\xi}} \end{array} \right] \mathrm{d}\mu(\boldsymbol{\xi}) \\
\text{subject to} & \mu \in \mathcal{M}_+(\mathbb{R}^Q) \\
& \displaystyle\int_\Xi \mathrm{d}\mu(\boldsymbol{\xi}) = 1 \\
& \displaystyle\int_\Xi \boldsymbol{g}(\boldsymbol{\xi})\, \mathrm{d}\mu(\boldsymbol{\xi}) \leq \boldsymbol{c}
\end{array}
$$

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

$$\begin{array}{ll} \text{minimize} & \alpha + \boldsymbol{c}^\top \boldsymbol{\beta} \\ \text{subject to} & \alpha \in \mathbb{R}, \;\; \boldsymbol{\beta} \in \mathbb{R}^S_+ \\ & \alpha + \boldsymbol{g}(\boldsymbol{\xi})^\top \boldsymbol{\beta} \geq \mathbb{I}\left[\begin{array}{l} \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\tilde{\boldsymbol{\xi}}, \\ \tilde{\boldsymbol{\xi}}^\top \boldsymbol{Q}\,\boldsymbol{y}^k \leq \tilde{\boldsymbol{\xi}}^\top \boldsymbol{Q}\,\boldsymbol{y}^{k'} \quad \forall k' \in \mathcal{K} : \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^{k'} \leq \boldsymbol{H}\tilde{\boldsymbol{\xi}} \end{array}\right] \qquad \forall \boldsymbol{\xi} \in \Xi \end{array}$$

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

$$\begin{array}{l} \alpha + \boldsymbol{g}(\boldsymbol{\xi})^\top \boldsymbol{\beta} \geq 1 \quad \forall \boldsymbol{\xi} \in \Xi : \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\boldsymbol{\xi},\; \boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y}^k \leq \boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y}^{k'} \;\; \forall k' \in \mathcal{K} : \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^{k'} \leq \boldsymbol{H}\boldsymbol{\xi} \\ \alpha + \boldsymbol{g}(\boldsymbol{\xi})^\top \boldsymbol{\beta} \geq 0 \quad \forall \boldsymbol{\xi} \in \Xi. \end{array} \tag{20}$$

In the following, we focus on the reformulation of the first constraint; the second constraint can be dealt with analogously. The first constraint is equivalent to the system of semi-infinite constraints

$$\alpha + \boldsymbol{g}(\boldsymbol{\xi})^\top \boldsymbol{\beta} \geq 1 \quad \forall \boldsymbol{\ell} \in \partial\mathcal{L} : \ell_k = 0,\; \forall \boldsymbol{\xi} \in \Xi(\boldsymbol{\ell}) : \boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y}^k \leq \boldsymbol{\xi}^\top \boldsymbol{Q}\,\boldsymbol{y}^{k'} \;\; \forall k' \in \mathcal{K} : \ell_{k'} = 0. \tag{21}$$

By dualizing the resulting semi-infinite constraints as in the proof of Proposition 1, we arrive at the linear program in the statement of the proposition. □

Proof of Proposition 5. The minimum probability with which policy $\boldsymbol{y}^k$ is chosen under any probability distribution $\mathbb{P} \in \mathcal{P}$ is given by the optimal value of the following moment problem.

$$\begin{array}{ll} \text{minimize} & \displaystyle\int_\Xi \mathbb{I}\left[\begin{array}{l} \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\tilde{\boldsymbol{\xi}}, \\ \tilde{\boldsymbol{\xi}}^\top \boldsymbol{Q}\,\boldsymbol{y}^k \leq \tilde{\boldsymbol{\xi}}^\top \boldsymbol{Q}\,\boldsymbol{y}^{k'} \quad \forall k' \in \mathcal{K} : \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^{k'} \leq \boldsymbol{H}\tilde{\boldsymbol{\xi}} \end{array}\right] \mathrm{d}\mu(\boldsymbol{\xi}) \\ \text{subject to} & \mu \in \mathcal{M}_+(\mathbb{R}^Q) \\ & \displaystyle\int_\Xi \mathrm{d}\mu(\boldsymbol{\xi}) = 1 \\ & \displaystyle\int_\Xi \boldsymbol{g}(\boldsymbol{\xi})\,\mathrm{d}\mu(\boldsymbol{\xi}) \leq \boldsymbol{c} \end{array}$$

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

$$\begin{aligned}
&\text{maximize} && \alpha - \boldsymbol{c}^\top \boldsymbol{\beta} \\
&\text{subject to} && \alpha \in \mathbb{R},\ \ \boldsymbol{\beta} \in \mathbb{R}^S_+ \\
& && \alpha - \boldsymbol{g}(\boldsymbol{\xi})^\top \boldsymbol{\beta} \leq \mathbb{I} \left[\begin{array}{l} \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^k \leq \boldsymbol{H}\tilde{\boldsymbol{\xi}}, \\ \tilde{\boldsymbol{\xi}}^\top \boldsymbol{Q}\, \boldsymbol{y}^k \leq \tilde{\boldsymbol{\xi}}^\top \boldsymbol{Q}\, \boldsymbol{y}^{k'} \quad \forall k' \in \mathcal{K} : \boldsymbol{T}\boldsymbol{x} + \boldsymbol{W}\boldsymbol{y}^{k'} \leq \boldsymbol{H}\tilde{\boldsymbol{\xi}} \end{array} \right] \qquad \forall \boldsymbol{\xi} \in \Xi
\end{aligned}$$

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

$$\begin{aligned}
&\alpha - \boldsymbol{g}(\boldsymbol{\xi})^\top \boldsymbol{\beta} \leq 1 \quad \forall \boldsymbol{\xi} \in \Xi \\
&\alpha - \boldsymbol{g}(\boldsymbol{\xi})^\top \boldsymbol{\beta} \leq 0 \quad \forall \boldsymbol{\ell} \in \mathcal{L} : \ell_k \neq 0,\ \forall \boldsymbol{\xi} \in \Xi(\boldsymbol{\ell}) \\
&\alpha - \boldsymbol{g}(\boldsymbol{\xi})^\top \boldsymbol{\beta} \leq 0 \quad \forall \boldsymbol{\ell} \in \mathcal{L} : \ell_k = 0, \forall k' \in \mathcal{K} \setminus k : \ell'_k = 0,\ \forall \boldsymbol{\xi} \in \Xi(\boldsymbol{\ell}) : \boldsymbol{\xi}^\top \boldsymbol{Q}\boldsymbol{y}^k > \boldsymbol{\xi}^\top \boldsymbol{Q}\boldsymbol{y}^{k'}.
\end{aligned}$$

Since the function $\alpha - \boldsymbol{g}(\boldsymbol{\xi})^\top \boldsymbol{\beta}$ is continuous in $\boldsymbol{\xi}$, we can replace the strict inequality in the third constraint with a weak one whenever the set $\{\boldsymbol{\xi} \in \Xi(\boldsymbol{\ell}) : \boldsymbol{\xi}^\top \boldsymbol{Q}\boldsymbol{y}^k > \boldsymbol{\xi}^\top \boldsymbol{Q}\boldsymbol{y}^{k'}\}$ is non-empty. The result now follows if we apply similar reformulations to the resulting semi-infinite constraints as in the proof of Proposition 2. □