

Conic Programming Reformulations of Two-Stage Distributionally Robust Linear Programs over Wasserstein Balls

Grani A. Hanasusanto¹ and Daniel Kuhn²

¹Graduate Program in Operations Research and Industrial Engineering, The University of Texas at Austin, USA

²Risk Analytics and Optimization Chair, École Polytechnique Fédérale de Lausanne, Switzerland

July 14, 2017

Abstract

Adaptive robust optimization problems are usually solved *approximately* by restricting the adaptive decisions to simple parametric decision rules. However, the corresponding approximation error can be substantial. In this paper we show that two-stage robust and distributionally robust linear programs can often be reformulated *exactly* as conic programs that scale polynomially with the problem dimensions. Specifically, when the ambiguity set constitutes a 2-Wasserstein ball centered at a discrete distribution, then the distributionally robust linear program is equivalent to a copositive program (if the problem has complete recourse) or can be approximated arbitrarily closely by a sequence of copositive programs (if the problem has sufficiently expensive recourse). These results directly extend to the classical robust setting and motivate strong tractable approximations of two-stage problems based on semidefinite approximations of the copositive cone. We also demonstrate that the two-stage distributionally robust optimization problem is equivalent to a tractable linear program when the ambiguity set constitutes a 1-Wasserstein ball centered at a discrete distribution and there are no support constraints.

1 Introduction

In two-stage optimization under uncertainty an agent selects a here-and-now decision before observing the realization of some decision-relevant random vector. Once the uncertainty has been revealed, a wait-and-see decision is taken in order to correct any undesired effects of the here-and-now decision in the realized scenario. Classical stochastic programming seeks a single here-and-now decision and a family of (possibly infinitely many) wait-and-see decisions—one for each possible uncertainty realization—with the goal to minimize the sum of a deterministic here-and-now cost and the expectation of an uncertain wait-and-see cost [53].

Classical robust optimization, in contrast, seeks decisions that minimize the worst case of the total cost across all possible uncertainty realizations [3]. While stochastic programming assumes full knowledge of the distribution governing the uncertain problem parameters, which is needed to evaluate the expectation of the total costs, robust optimization denies (or ignores) any knowledge of this distribution except for its support.

Distributionally robust optimization is an alternative modeling paradigm pioneered in [17, 48, 54]. It has gained new thrust over the last decade and challenges the black-and-white view of stochastic and robust optimization. Specifically, it assumes that the decision maker has access to some limited probabilistic information (*e.g.*, in the form of the distribution’s moments, its structural properties or its distance to a reference distribution); but not enough to pin down the true distribution precisely. In this setting, a meaningful objective is to minimize the worst-case expected total cost, where the worst case is evaluated across an ambiguity set that contains all distributions consistent with the available probabilistic information. Distributionally robust models enjoy strong theoretical justification from decision theory [22], and there is growing evidence that they provide high-quality decisions at a moderate computational cost [15, 23, 61].

Two-stage decision problems under uncertainty—whether stochastic, robust or distributionally robust—typically involve a continuum of wait-and-see decisions and thus constitute infinite-dimensional functional optimization problems. Therefore, they can only be solved approximately, except in contrived circumstances. The existing approximation methods can roughly be subdivided into *discretization schemes* [25, 32, 52] and *decision rule methods* [4, 21, 23]. Discretization schemes approximate the support of the uncertain parameters with a finite subset, which entails a relaxation of the original problem and encourages optimistically biased solutions. Decision rule methods, on the other hand, approximate the infinite-dimensional space of all wait-and-see decisions with a finite-dimensional subspace of linearly parameterized decision rules, which entails a restriction of the original problem and leads to pessimistically biased solutions. In this paper, we introduce a new method for approximating two-stage distributionally robust linear programs, which can neither be classified as a discretization scheme nor as a decision rule method: We first reformulate the original infinite-dimensional optimization problem as an equivalent finite-dimensional conic program of polynomial size, which absorbs all the complexity in its cones, and then replace the cones with tractable inner approximations.

Our exposition focuses on distributionally robust linear programs whose ambiguity sets contain all discrete and continuous distributions supported on a polytope that have a Wasserstein distance of at most ϵ from a discrete reference distribution (such as the empirical distribution corresponding to finitely many samples from the unknown true distribution). This problem class encapsulates the two-stage stochastic linear programs with discrete distributions (for $\epsilon = 0$) and the two-stage robust optimization problems with bounded polyhedral uncertainty sets (for $\epsilon = \infty$) as special cases. Wasserstein ambiguity sets have first been used in the context of portfolio optimization [45]. The corresponding distributionally robust optimization models were initially perceived as difficult and thus tackled with methods from global optimization [62]; see also [44,

Chapter 7]. Recently it has been discovered, however, that distributionally robust optimization problems with Wasserstein ambiguity sets can often be reformulated as finite convex programs [39, 66]. Single-stage problems with piecewise linear cost functions, for instance, are tractable and admit convex reformulations of polynomial sizes [39]. Two-stage problems, on the other hand, are generically NP-hard. Their convex reformulations have exponential size but are amenable to Benders-type decomposition algorithms [66]. Alternatively, two-stage problems can be converted to single-stage problems via a decision rule approximation, in which case they admit again a convex reformulation of polynomial size and thus regain tractability [20].

This paper extends the state-of-the-art in two-stage distributionally robust linear programming along several dimensions. We highlight the following main contributions:

- (i) We prove that any two-stage distributionally robust linear program with *complete recourse* is equivalent to a copositive program of polynomial size if the ambiguity set constitutes a 2-Wasserstein ball centered at a discrete distribution.
- (ii) We prove that any two-stage distributionally robust linear program with *sufficiently expensive recourse* can be approximated arbitrarily closely by a sequence of copositive programs of a fixed polynomial size if the ambiguity set constitutes a 2-Wasserstein ball centered at a discrete distribution.
- (iii) By using nested hierarchies of tractable convex cones to approximate the (intractable) copositive cones from the inside [10, 14, 43], we obtain sequences of tractable conservative approximations for the two-stage distributionally robust linear programs described in (i) and (ii). These approximations can be made arbitrarily accurate. However, numerical tests suggest that even the coarsest of these approximations distinctly outperform the state-of-the-art decision rule approximations in terms of accuracy.
- (iv) We prove that any two-stage distributionally robust linear program with *fixed costs* is equivalent to a tractable linear program if the ambiguity set constitutes a 1-Wasserstein ball centered at a discrete distribution and if there are no support constraints. We also show that this tractability result is sharp.
- (v) We demonstrate that all of the above results carry directly over to classical two-stage robust optimization problems with bounded polyhedral uncertainty sets. To our best knowledge, we provide the first (polynomially-sized) conic programming reformulations for generic problem instances in this class.

Two-stage distributionally robust linear programs with objective uncertainty are studied in [7]. Assuming that only the first- and second-order moments of the uncertain cost coefficients are known, these problems can be reformulated as tractable semidefinite programs. In the presence of constraint uncertainty, however, these problems become intractable. Two-stage distributionally robust binary programs with polyhedral moment information are studied in [29]. If only the cost coefficients are uncertain, these problems can be

reformulated as explicit mixed-integer linear programs of polynomial sizes. While two-stage distributionally robust optimization endeavors to minimize the worst-case (maximal) expected wait-and-see cost, a parallel stream of research investigates the best-case (minimal) expectations of the minima of mixed zero-one linear programs with objective uncertainty. Under first- and second-order moment information, any such best-case expectation can be reformulated as the optimal value of a completely positive program [42]. In fact, this best-case expectation even reduces to the optimal value of a tractable semidefinite program whenever the convex hull of all rank-1 outer products of feasible wait-and-see decisions with themselves is semidefinite-representable [41]. These deep theoretical results have recently opened up new avenues for modeling and solving stochastic appointment scheduling problems [33] and have also ramifications for computing best-worst choice probabilities in discrete choice models [41]. A comprehensive survey of recent results at the interface of distributionally robust optimization and completely positive programming is provided in [36].

In contrast to the existing literature, here we develop copositive programming reformulations for generic two-stage distributionally robust linear programs where both the objective function and the constraints may be affected by the uncertainty. We also present new linear programming reformulations for two-stage distributionally robust linear programs where the uncertainty affects only the constraints. These exact reformulations are reminiscent of the conservative approximation models for two-stage robust optimization models derived in [1] by leveraging popular reformulation-linearization techniques from bilinear programming. Another main difference to the existing literature is our focus on Wasserstein balls instead of moment ambiguity sets to capture distributional uncertainty. This has the advantage that the degree of ambiguity aversion can be controlled by tuning the radius of the Wasserstein ball.

A key benefit of Wasserstein balls is that they provide natural confidence sets for the unknown distribution of the uncertain problem parameters. Specifically, the Wasserstein ball around the empirical distribution on I independent historical samples contains the unknown true distribution with confidence $1 - \beta$ if its radius exceeds an explicit threshold $\epsilon_I(\beta)$ that is known in closed form [39, 66]. Thus, the corresponding distributionally robust optimization problem offers a $1 - \beta$ upper confidence bound on the optimal value of the true stochastic program. One can also show that this data-driven distributionally robust optimization problem converges to the corresponding true stochastic program as the sample size I tends to infinity [39, 66]. Other data-driven distributionally robust optimization models that offer finite sample and asymptotic guarantees are discussed in [8] based on goodness-of-fit ambiguity sets, in [31] based on L^1 -norm ball ambiguity sets and in [38] based on Φ -divergence ambiguity sets.

While this paper was under review, we became aware of the paper [63] by Xu and Burer, which was submitted simultaneously. It turns out that our Corollary 1 is equivalent to Theorem 1 in [63], and so we mention it here for the readers reference. While [63] focuses on two-stage *robust* linear programs with right hand side uncertainty, we develop copositive programming reformulations for *distributionally robust*

two-stage linear programs with objective *and* constraint uncertainty.

The rest of the paper is structured as follows. Section 2 provides a formal problem statement and reviews some fundamental results from [39, 66]. In Section 3 we derive copositive programming reformulations for two-stage distributionally robust linear programs over 2-Wasserstein balls and discuss tractable approximations. Exact tractable linear programming reformulations for two-stage distributionally robust linear programs over 1-Wasserstein balls are described in Section 4. Section 5 reports on numerical results.

Notation: For any $I \in \mathbb{N}$, we define $[I]$ as the index set $\{1, \dots, I\}$. We denote by \mathbb{I} the identity matrix and by \mathbf{e} the vector of all ones. Their dimensions will be clear from the context. The trace of a square matrix \mathbf{M} is denoted as $\text{tr}(\mathbf{M})$. We define $\text{diag}(\mathbf{v})$ as the diagonal matrix with the vector \mathbf{v} on its main diagonal. The set of non-negative (positive) reals is denoted as \mathbb{R}_+ (\mathbb{R}_{++}). The set of all symmetric matrices in $\mathbb{R}^{K \times K}$ is denoted as \mathbb{S}^K , while the cone of positive semidefinite matrices in $\mathbb{R}^{K \times K}$ is denoted as \mathbb{S}_+^K . We define the cone of copositive matrices as $\mathcal{C} = \{\mathbf{M} \in \mathbb{S}^K : \boldsymbol{\xi}^\top \mathbf{M} \boldsymbol{\xi} \geq 0 \forall \boldsymbol{\xi} \geq \mathbf{0}\}$ and the cone of completely positive matrices as $\mathcal{C}^* = \{\mathbf{M} \in \mathbb{S}^K : \mathbf{M} = \mathbf{B}\mathbf{B}^\top \text{ for some } \mathbf{B} \in \mathbb{R}_+^{K \times g(K)}\}$, where $g(K) = \max\{\binom{K+1}{2} - 4, K\}$ [50]. For any $\mathbf{Q}, \mathbf{R} \in \mathbb{S}^K$, the relations $\mathbf{Q} \succeq \mathbf{R}$, $\mathbf{Q} \succeq_{\mathcal{C}} \mathbf{R}$, and $\mathbf{Q} \succeq_{\mathcal{C}^*} \mathbf{R}$ mean that $\mathbf{Q} - \mathbf{R}$ is an element of \mathbb{S}_+^K , \mathcal{C} , and \mathcal{C}^* , respectively. We denote the j -th row (j -th column) of a matrix \mathbf{M} as \mathbf{M}_j : ($\mathbf{M}_{:j}$). All random variables are designated by tilde signs (*e.g.*, $\tilde{\boldsymbol{\xi}}$), while their realizations are denoted without tildes (*e.g.*, $\boldsymbol{\xi}$). The characteristic function of a set \mathcal{S} is defined as $\chi_{\mathcal{S}}(\boldsymbol{\xi}) = 0$ if $\boldsymbol{\xi} \in \mathcal{S}$; $= \infty$ otherwise.

2 Problem Formulation

We study two-stage distributionally robust linear programs of the form

$$\begin{aligned} & \text{minimize} && \mathbf{c}^\top \mathbf{x} + \mathcal{Z}(\mathbf{x}) \\ & \text{subject to} && \mathbf{x} \in \mathcal{X}, \end{aligned} \tag{1}$$

where $\mathcal{X} \subseteq \mathbb{R}^{N_1}$ is the feasible set of the here-and-now decisions, $\mathbf{c}^\top \mathbf{x}$ is the here-and-now cost, and $\mathcal{Z}(\mathbf{x})$ is the worst-case expected wait-and-see cost. Formally, we set

$$\mathcal{Z}(\mathbf{x}) = \sup_{\mathbb{P} \in \hat{\mathcal{P}}} \mathbb{E}_{\mathbb{P}} \left[Z(\mathbf{x}, \tilde{\boldsymbol{\xi}}) \right], \tag{2}$$

where $\tilde{\boldsymbol{\xi}} \in \Xi \subseteq \mathbb{R}^K$ is a random vector comprising the uncertain problem parameters, and $\hat{\mathcal{P}}$ is an ambiguity set that contains the possible distributions of $\tilde{\boldsymbol{\xi}}$. The recourse function $Z(\mathbf{x}, \boldsymbol{\xi})$ in (2) constitutes the optimal value of the recourse problem, that is,

$$\begin{aligned} Z(\mathbf{x}, \boldsymbol{\xi}) = & \inf && (\mathbf{Q}\boldsymbol{\xi} + \mathbf{q})^\top \mathbf{y} \\ \text{s.t.} &&& \mathbf{y} \in \mathbb{R}^{N_2} \\ &&& \mathbf{T}(\mathbf{x})\boldsymbol{\xi} + \mathbf{h}(\mathbf{x}) \leq \mathbf{W}\mathbf{y}, \end{aligned} \tag{3}$$

where $\mathbf{T}(\mathbf{x}) \in \mathbb{R}^{M \times K}$ and $\mathbf{h}(\mathbf{x}) \in \mathbb{R}^M$ are matrix- and vector-valued affine functions, respectively. The dual of the recourse problem is given by

$$\begin{aligned} Z_d(\mathbf{x}, \boldsymbol{\xi}) = & \sup (\mathbf{T}(\mathbf{x})\boldsymbol{\xi} + \mathbf{h}(\mathbf{x}))^\top \mathbf{p} \\ \text{s.t. } & \mathbf{p} \in \mathbb{R}_+^M \\ & \mathbf{Q}\boldsymbol{\xi} + \mathbf{q} = \mathbf{W}^\top \mathbf{p}. \end{aligned} \quad (4)$$

We introduce the following standard terminology that will be used throughout the paper.

Definition 1 (Complete Recourse). *We say that the two-stage distributionally robust linear program (1) has complete recourse if there exists $\mathbf{y}^+ \in \mathbb{R}^{N_2}$ with $\mathbf{W}\mathbf{y}^+ > \mathbf{0}$.*

Complete recourse implies that problem (3) is feasible for every $\mathbf{x} \in \mathbb{R}^{N_1}$ and $\boldsymbol{\xi} \in \mathbb{R}^K$. Indeed, it implies that there is always a $\lambda > 0$ such that $\mathbf{y} = \lambda \mathbf{y}^+$ exceeds $\mathbf{T}(\mathbf{x})\boldsymbol{\xi} + \mathbf{h}(\mathbf{x})$.

Definition 2 (Sufficiently Expensive Recourse). *We say that the two-stage distributionally robust linear program (1) has sufficiently expensive recourse if for any fixed $\boldsymbol{\xi} \in \Xi$ the dual problem (4) is feasible.*

If problem (1) has complete recourse, then $Z(\mathbf{x}, \boldsymbol{\xi}) < +\infty$ for every $\mathbf{x} \in \mathbb{R}^{N_1}$ and $\boldsymbol{\xi} \in \mathbb{R}^K$. On the other hand, if problem (1) has sufficiently expensive recourse, then $Z(\mathbf{x}, \boldsymbol{\xi}) > -\infty$ for every $\mathbf{x} \in \mathcal{X}$ and $\boldsymbol{\xi} \in \Xi$. If both conditions are satisfied, then $Z(\mathbf{x}, \boldsymbol{\xi})$ is finite. Each condition on itself implies that strong duality holds between the primal and dual linear programs (3) and (4), respectively. Throughout this paper, we will always assume that problem (1) has sufficiently expensive recourse. This is a weak condition that is even satisfied by many problems with induced constraints. The complete recourse assumption, which rules out induced constraints, will only be imposed occasionally to obtain stronger results.

Following [39, 66], we assume henceforth that the true distribution of $\tilde{\boldsymbol{\xi}}$ is unknown but that we have access to I samples $\hat{\boldsymbol{\xi}}_1, \dots, \hat{\boldsymbol{\xi}}_I$ from this distribution. In this case, we can define the empirical distribution $\hat{\mathbb{P}}_I = \frac{1}{I} \sum_{i \in [I]} \delta_{\hat{\boldsymbol{\xi}}_i}$, that is, the uniform distribution on the samples. The ambiguity set $\hat{\mathcal{P}}$ in (1) can then be defined as the family of all distributions that are close to the empirical distribution $\hat{\mathbb{P}}_I$ with respect to the Wasserstein metric.

Definition 3 (Wasserstein Metric). *For any $r \geq 1$, let $\mathcal{M}^r(\Xi)$ be the set of all probability distributions \mathbb{P} supported on Ξ satisfying $\mathbb{E}_{\mathbb{P}}[d(\tilde{\boldsymbol{\xi}}, \boldsymbol{\xi}_0)^r] = \int_{\Xi} d(\boldsymbol{\xi}, \boldsymbol{\xi}_0)^r \mathbb{P}(d\boldsymbol{\xi}) < \infty$ where $\boldsymbol{\xi}_0 \in \Xi$ is some reference point, and $d(\boldsymbol{\xi}, \boldsymbol{\xi}_0)$ is a continuous reference metric on Ξ . For any $r \geq 1$, the r -Wasserstein distance between two distributions $\mathbb{P}_1, \mathbb{P}_2 \in \mathcal{M}^r(\Xi)$ is defined as*

$$W^r(\mathbb{P}_1, \mathbb{P}_2) = \inf \left\{ \left(\int_{\Xi^2} d(\boldsymbol{\xi}_1, \boldsymbol{\xi}_2)^r \mathbb{Q}(d\boldsymbol{\xi}_1, d\boldsymbol{\xi}_2) \right)^{\frac{1}{r}} : \begin{array}{l} \mathbb{Q} \text{ is a joint distribution of } \boldsymbol{\xi}_1 \text{ and } \boldsymbol{\xi}_2 \\ \text{with marginals } \mathbb{P}_1 \text{ and } \mathbb{P}_2, \text{ respectively} \end{array} \right\}.$$

We denote the Wasserstein ball of radius ϵ centered at the empirical distribution by

$$\mathcal{B}_\epsilon^r(\hat{\mathbb{P}}_I) = \left\{ \mathbb{P} \in \mathcal{M}^r(\Xi) : W^r(\mathbb{P}, \hat{\mathbb{P}}_I) \leq \epsilon \right\}.$$

The following theorem, which is adapted from [39, 20] and relies on the Knothe-Rosenblatt rearrangement [59], establishes that the worst-case expectation (2) over a Wasserstein ambiguity set $\hat{\mathcal{P}} = \mathcal{B}_\epsilon^r(\hat{\mathbb{P}}_I)$ can be reformulated in terms of a generalized moment problem and the corresponding dual robust optimization problem. To keep this paper self-contained, we prove this theorem in the appendix.

Theorem 1. *If $\hat{\mathcal{P}} = \mathcal{B}_\epsilon^r(\hat{\mathbb{P}}_I)$, the worst-case expectation (2) coincides with the optimal value of the generalized moment problem*

$$\begin{aligned} \mathcal{Z}(\mathbf{x}) = & \sup \frac{1}{I} \sum_{i \in [I]} \int_{\Xi} Z(\mathbf{x}, \boldsymbol{\xi}) \mathbb{P}_i(d\boldsymbol{\xi}) \\ \text{s.t. } & \mathbb{P}_i \in \mathcal{M}^r(\Xi) \quad \forall i \in [I] \\ & \frac{1}{I} \sum_{i \in [I]} \int_{\Xi} d(\boldsymbol{\xi}, \hat{\boldsymbol{\xi}}_i)^r \mathbb{P}_i(d\boldsymbol{\xi}) \leq \epsilon^r. \end{aligned} \quad (5)$$

Furthermore, for $\epsilon > 0$ this problem admits the strong dual robust optimization problem

$$\mathcal{Z}(\mathbf{x}) = \inf_{\lambda \in \mathbb{R}_+} \epsilon^r \lambda + \frac{1}{I} \sum_{i \in [I]} \sup_{\boldsymbol{\xi} \in \Xi} Z(\mathbf{x}, \boldsymbol{\xi}) - \lambda d(\boldsymbol{\xi}, \hat{\boldsymbol{\xi}}_i)^r. \quad (6)$$

All results of this paper directly extend to the class of two-stage *robust* optimization problems, which model the uncertainties only through their uncertainty set Ξ .

Remark 1 (Two-Stage Robust Optimization). *If Ξ is compact and the radius ϵ of the Wasserstein ball is larger than the diameter of Ξ , then the two-stage distributionally robust linear program (1) simplifies to the two-stage robust optimization problem*

$$\begin{aligned} \text{minimize } & \mathbf{c}^\top \mathbf{x} + \max_{\boldsymbol{\xi} \in \Xi} Z(\mathbf{x}, \boldsymbol{\xi}) \\ \text{subject to } & \mathbf{x} \in \mathcal{X}. \end{aligned} \quad (7)$$

Indeed, if we set $\epsilon \geq \max_{\boldsymbol{\xi}, \boldsymbol{\xi}' \in \Xi} d(\boldsymbol{\xi}, \boldsymbol{\xi}')$, then the Wasserstein ball $\mathcal{B}_\epsilon^r(\hat{\mathbb{P}}_I)$ contains all Dirac distributions $\delta_{\boldsymbol{\xi}}$, $\boldsymbol{\xi} \in \Xi$. This implies that the worst-case expected cost (2) reduces to the worst-case cost $\max_{\boldsymbol{\xi} \in \Xi} Z(\mathbf{x}, \boldsymbol{\xi})$, irrespective of the number and positions of the samples $\hat{\boldsymbol{\xi}}_i$, $i \in [I]$.

3 Copositive Programming Reformulation

Throughout this section we work with the 2-Wasserstein metric and the 2-norm reference distance $d(\boldsymbol{\xi}_1, \boldsymbol{\xi}_2) = \|\boldsymbol{\xi}_1 - \boldsymbol{\xi}_2\|_2$. We further assume that the support set Ξ is a non-empty polyhedron of the form

$$\Xi = \{\boldsymbol{\xi} \in \mathbb{R}_+^K : \mathbf{S}\boldsymbol{\xi} \leq \mathbf{t}\} \quad (8)$$

for some $\mathbf{S} \in \mathbb{R}^{J \times K}$ and $\mathbf{t} \in \mathbb{R}^J$. Note that we assume without loss of generality that Ξ is a (possibly unbounded) subset of the non-negative orthant. Finally, we also assume that problem (1) has sufficiently expensive recourse. We will show that under these assumptions the two-stage distributionally robust linear program (1) admits an equivalent reformulation as a copositive program.

3.1 A Copositive Upper Bound on $\mathcal{Z}(\mathbf{x})$

To derive a copositive programming-based upper bound on $\mathcal{Z}(\mathbf{x})$, we will need the following technical lemma.

Lemma 1. *For any symmetric matrix $\mathbf{M} \in \mathbb{S}^K$, we have $\mathbf{M} \succeq_c \mathbf{0}$ if and only if*

$$[\mathbf{z}^\top \ 1] \mathbf{M} [\mathbf{z}^\top \ 1]^\top \geq 0 \quad \forall \mathbf{z} \in \mathbb{R}_+^{K-1}. \quad (9)$$

Proof. To prove sufficiency, we recall that $\mathbf{M} \succeq_c \mathbf{0}$ if and only if $\boldsymbol{\xi}^\top \mathbf{M} \boldsymbol{\xi} \geq 0$ for all $\boldsymbol{\xi} \in \mathbb{R}_+^K$. Thus, (9) follows by focusing on those $\boldsymbol{\xi} \in \mathbb{R}^K$ with $\xi_K = 1$.

To prove the converse implication, assume that (9) holds. Hence, we have

$$\begin{aligned} [\mathbf{z}^\top \ 1] \mathbf{M} [\mathbf{z}^\top \ 1]^\top \geq 0 \quad \forall \mathbf{z} \in \mathbb{R}_+^{K-1} &\implies [t\mathbf{z}^\top \ t] \mathbf{M} [t\mathbf{z}^\top \ t]^\top \geq 0 \quad \forall \mathbf{z} \in \mathbb{R}_+^{K-1} \ \forall t \in \mathbb{R}_{++} \\ &\implies [\mathbf{y}^\top \ t] \mathbf{M} [\mathbf{y}^\top \ t]^\top \geq 0 \quad \forall \mathbf{y} \in \mathbb{R}_+^{K-1} \ \forall t \in \mathbb{R}_{++}, \end{aligned}$$

that is, for any fixed $\mathbf{y} \in \mathbb{R}_+^{K-1}$, the univariate quadratic function $[\mathbf{y}^\top \ t] \mathbf{M} [\mathbf{y}^\top \ t]^\top$ is non-negative for all $t > 0$. As this function is continuous, it must in fact be non-negative for all $t \geq 0$. Thus, $\mathbf{M} \succeq_c \mathbf{0}$. \square

We are now ready to derive an upper bound on the worst-case expectation $\mathcal{Z}(\mathbf{x})$. This bound is expressed as the optimal value of a copositive minimization problem and can thus be used to conservatively approximate (1) with a finite-dimensional minimization problem that is principally amenable to numerical solution.

Theorem 2 (Copositive Upper Bound). *For any fixed first-stage decision $\mathbf{x} \in \mathcal{X}$, the worst-case expectation $\mathcal{Z}(\mathbf{x})$ in (2) is bounded above by the optimal value of the copositive program*

$$\begin{aligned} &\overline{\mathcal{Z}}(\mathbf{x}) \\ = &\inf \quad \epsilon^2 \lambda + \frac{1}{I} \sum_{i \in [I]} \left[s_i + \mathbf{q}^\top \boldsymbol{\psi}_i - \lambda \|\hat{\boldsymbol{\xi}}_i\|_2^2 + \sum_{j \in [N_2+J]} \phi_{ij} q_j^2 \right] \\ \text{s.t.} &\quad \lambda \in \mathbb{R}_+, \ s_i \in \mathbb{R}, \ \boldsymbol{\psi}_i, \ \phi_i \in \mathbb{R}^{N_2+J} \quad \forall i \in [I] \\ &\quad \begin{bmatrix} \lambda \mathbb{I} + \mathbf{Q}^\top \text{diag}(\phi_i) \mathbf{Q} & -\frac{1}{2} \boldsymbol{\mathcal{T}}(\mathbf{x})^\top - \mathbf{Q}^\top \text{diag}(\phi_i) \boldsymbol{\mathcal{W}}^\top & -\lambda \hat{\boldsymbol{\xi}}_i - \frac{1}{2} \mathbf{Q}^\top \boldsymbol{\psi}_i \\ -\frac{1}{2} \boldsymbol{\mathcal{T}}(\mathbf{x}) - \boldsymbol{\mathcal{W}} \text{diag}(\phi_i) \mathbf{Q} & \boldsymbol{\mathcal{W}} \text{diag}(\phi_i) \boldsymbol{\mathcal{W}}^\top & \frac{1}{2} (\boldsymbol{\mathcal{W}} \boldsymbol{\psi}_i - \mathbf{h}(\mathbf{x})) \\ (-\lambda \hat{\boldsymbol{\xi}}_i - \frac{1}{2} \mathbf{Q}^\top \boldsymbol{\psi}_i)^\top & \frac{1}{2} (\boldsymbol{\mathcal{W}} \boldsymbol{\psi}_i - \mathbf{h}(\mathbf{x}))^\top & s_i \end{bmatrix} \succeq_c \mathbf{0} \quad \forall i \in [I], \end{bmatrix} \quad (10) \end{aligned}$$

where

$$\mathbf{Q} = \begin{bmatrix} \mathbf{Q} \\ \mathbf{S} \end{bmatrix}, \quad \mathbf{q} = \begin{bmatrix} \mathbf{q} \\ -t \end{bmatrix}, \quad \boldsymbol{\mathcal{T}}(\mathbf{x}) = \begin{bmatrix} \boldsymbol{\mathcal{T}}(\mathbf{x}) \\ \mathbf{0} \end{bmatrix}, \quad \mathbf{h}(\mathbf{x}) = \begin{bmatrix} \mathbf{h}(\mathbf{x}) \\ \mathbf{0} \end{bmatrix}, \quad \text{and} \quad \boldsymbol{\mathcal{W}} = \begin{bmatrix} \boldsymbol{\mathcal{W}} & \mathbf{0} \\ \mathbf{0} & -\mathbb{I} \end{bmatrix}. \quad (11)$$

Remark 2. *Note that the extended recourse parameters defined in (11) combine the input data of the recourse problem (3) with the parameters characterizing the support set (8).*

Proof of Theorem 2. By strong linear programming duality, which holds because problem (1) has sufficiently expensive recourse, we have $Z(\mathbf{x}, \boldsymbol{\xi}) = Z_d(\mathbf{x}, \boldsymbol{\xi})$ for every $\mathbf{x} \in \mathcal{X}$ and $\boldsymbol{\xi} \in \Xi$. Recalling that $r = 2$ and $d(\boldsymbol{\xi}_1, \boldsymbol{\xi}_2) = \|\boldsymbol{\xi}_1 - \boldsymbol{\xi}_2\|_2$, the explicit formula (4) for the optimal value $Z_d(\mathbf{x}, \boldsymbol{\xi})$ of the dual recourse problem and the polyhedral representation (8) for Ξ allow us to reformulate (6) as

$$\begin{aligned} \mathcal{Z}(\mathbf{x}) &= \inf_{\lambda \geq 0} \epsilon^2 \lambda + \frac{1}{I} \sum_{i \in [I]} \sup_{\substack{\boldsymbol{\xi} \geq \mathbf{0} \\ \mathbf{S}\boldsymbol{\xi} \leq \mathbf{t}}} \sup_{\substack{\mathbf{p} \geq \mathbf{0} \\ \mathbf{Q}\boldsymbol{\xi} + \mathbf{q} = \mathbf{W}^\top \mathbf{p}}} (\mathcal{T}(\mathbf{x})\boldsymbol{\xi} + \mathbf{h}(\mathbf{x}))^\top \mathbf{p} - \lambda \|\boldsymbol{\xi} - \hat{\boldsymbol{\xi}}_i\|_2^2 \\ &= \inf_{\lambda \geq 0} \epsilon^2 \lambda + \frac{1}{I} \sum_{i \in [I]} \sup_{\substack{\boldsymbol{\xi}, \boldsymbol{\pi} \geq \mathbf{0} \\ \mathbf{Q}\boldsymbol{\xi} + \mathbf{q} = \mathbf{W}^\top \boldsymbol{\pi}}} (\mathcal{T}(\mathbf{x})\boldsymbol{\xi} + \mathbf{h}(\mathbf{x}))^\top \boldsymbol{\pi} - \lambda \|\boldsymbol{\xi} - \hat{\boldsymbol{\xi}}_i\|_2^2, \end{aligned} \quad (12)$$

where the second equality uses the definitions in (11). Note that the first M components of the new decision variable $\boldsymbol{\pi} \in \mathbb{R}^{M+J}$ correspond to the dual variable \mathbf{p} , while the remaining J components represent slack variables for the support constraints $\mathbf{S}\boldsymbol{\xi} \leq \mathbf{t}$. Next, we add the following non-convex constraints to each of the I inner maximization problems in (12).

$$q_j^2 = (\mathbf{W}_{:j}^\top \boldsymbol{\pi} - \mathbf{Q}_{j:}^\top \boldsymbol{\xi})^2 = (\mathbf{Q}_{j:}^\top \boldsymbol{\xi})^2 - 2\mathbf{Q}_{j:}^\top \boldsymbol{\xi} \boldsymbol{\pi}^\top \mathbf{W}_{:j} + (\mathbf{W}_{:j}^\top \boldsymbol{\pi})^2 \quad \forall j \in [N_2 + J]$$

Note that these constraints are redundant as they follow from $\mathbf{Q}\boldsymbol{\xi} + \mathbf{q} = \mathbf{W}^\top \boldsymbol{\pi}$. As will be revealed later, however, these constraints ensure that the optimal value of the completely positive program dual to (10) coincides with $\mathcal{Z}(\mathbf{x})$. Thanks to a recent result from the theory of quadratic programming [11], each of the emerging (nonconvex) quadratically constrained quadratic subproblems in (12) can be reformulated as a completely positive maximization problem. Thus, we could apply standard dualization techniques to reformulate (12) as a finite copositive minimization problem. Here, we instead pursue a more direct approach which leverages Lemma 1. By expressing all linear and quadratic constraints of the subproblems in Lagrangian form, we can reformulate (12) as

$$\begin{aligned} \mathcal{Z}(\mathbf{x}) &= \inf_{\lambda \in \mathbb{R}_+} \epsilon^2 \lambda + \frac{1}{I} \sum_{i \in [I]} \sup_{\boldsymbol{\xi}, \boldsymbol{\pi} \geq \mathbf{0}} \inf_{\boldsymbol{\psi}_i, \phi_i} \left[(\mathcal{T}(\mathbf{x})\boldsymbol{\xi} + \mathbf{h}(\mathbf{x}))^\top \boldsymbol{\pi} - \lambda \|\boldsymbol{\xi} - \hat{\boldsymbol{\xi}}_i\|_2^2 + \boldsymbol{\psi}_i^\top (\mathbf{Q}\boldsymbol{\xi} + \mathbf{q} - \mathbf{W}^\top \boldsymbol{\pi}) \right. \\ &\quad \left. + \sum_{j \in [N_2 + J]} \phi_{ij} \left(q_j^2 - (\mathbf{Q}_{j:}^\top \boldsymbol{\xi})^2 + 2\mathbf{Q}_{j:}^\top \boldsymbol{\xi} \boldsymbol{\pi}^\top \mathbf{W}_{:j} - (\mathbf{W}_{:j}^\top \boldsymbol{\pi})^2 \right) \right] \\ &\leq \inf_{\lambda \geq 0, \boldsymbol{\psi}_i, \phi_i} \epsilon^2 \lambda + \frac{1}{I} \sum_{i \in [I]} \sup_{\boldsymbol{\xi}, \boldsymbol{\pi} \geq \mathbf{0}} \left[(\mathcal{T}(\mathbf{x})\boldsymbol{\xi} + \mathbf{h}(\mathbf{x}))^\top \boldsymbol{\pi} - \lambda \|\boldsymbol{\xi} - \hat{\boldsymbol{\xi}}_i\|_2^2 + \boldsymbol{\psi}_i^\top (\mathbf{Q}\boldsymbol{\xi} + \mathbf{q} - \mathbf{W}^\top \boldsymbol{\pi}) \right. \\ &\quad \left. + \sum_{j \in [N_2 + J]} \phi_{ij} \left(q_j^2 - (\mathbf{Q}_{j:}^\top \boldsymbol{\xi})^2 + 2\mathbf{Q}_{j:}^\top \boldsymbol{\xi} \boldsymbol{\pi}^\top \mathbf{W}_{:j} - (\mathbf{W}_{:j}^\top \boldsymbol{\pi})^2 \right) \right], \end{aligned} \quad (13)$$

where ϕ_{ij} denotes the j -th entry of $\boldsymbol{\phi}_i$. Here, the inequality follows from interchanging the order of the supremum and the infimum operators. We observe now that the terms in square brackets constitute quadratic forms in $\boldsymbol{\xi}$ and $\boldsymbol{\pi}$. Thus, by introducing auxiliary epigraphical variables s_i , $i \in [I]$, to eliminate the suprema

over ξ and π , we can reformulate (13) as the quadratically parameterized semi-infinite linear program

$$\begin{aligned}
& \inf \quad \epsilon^2 \lambda + \frac{1}{I} \sum_{i \in [I]} \left[s_i + \mathbf{q}^\top \boldsymbol{\psi}_i - \lambda \|\hat{\xi}_i\|_2^2 + \sum_{j \in [N_2+J]} \phi_{ij} \mathbf{q}_j^2 \right] \\
& \text{s.t.} \quad \lambda \in \mathbb{R}_+, s_i \in \mathbb{R}, \boldsymbol{\psi}_i, \phi_i \in \mathbb{R}^{N_2+J} \quad \forall i \in [I] \\
& \quad \begin{bmatrix} \xi \\ \pi \\ 1 \end{bmatrix}^\top \begin{bmatrix} \lambda \mathbb{I} + \mathcal{Q}^\top \text{diag}(\phi_i) \mathcal{Q} & -\frac{1}{2} \mathcal{T}(\mathbf{x})^\top - \mathcal{Q}^\top \text{diag}(\phi_i) \mathcal{W}^\top & -\lambda \hat{\xi}_i - \frac{1}{2} \mathcal{Q}^\top \boldsymbol{\psi}_i \\ -\frac{1}{2} \mathcal{T}(\mathbf{x}) - \mathcal{W} \text{diag}(\phi_i) \mathcal{Q} & \mathcal{W} \text{diag}(\phi_i) \mathcal{W}^\top & \frac{1}{2} (\mathcal{W} \boldsymbol{\psi}_i - \mathbf{h}(\mathbf{x})) \\ (-\lambda \hat{\xi}_i - \frac{1}{2} \mathcal{Q}^\top \boldsymbol{\psi}_i)^\top & \frac{1}{2} (\mathcal{W} \boldsymbol{\psi}_i - \mathbf{h}(\mathbf{x}))^\top & s_i \end{bmatrix} \begin{bmatrix} \xi \\ \pi \\ 1 \end{bmatrix} \geq 0 \\
& \quad \forall i \in [I] \quad \forall (\xi, \pi) \in \mathbb{R}_+^{K+M+J},
\end{aligned}$$

which is equivalent to the copositive program (10) by virtue of Lemma 1. \square

3.2 A Completely Positive Reformulation of $\mathcal{Z}(\mathbf{x})$

We now derive the dual of the copositive program (10). As we will see later, even though (10) provides an upper bound on (2), the optimal value of its dual problem coincides with the worst-case expectation (2).

Proposition 1. *The copositive program (10) is dual to the following completely positive program.*

$$\begin{aligned}
\underline{\mathcal{Z}}(\mathbf{x}) = & \sup \quad \frac{1}{I} \sum_{i \in [I]} \text{tr}(\mathcal{T}(\mathbf{x}) \mathbf{Y}_i) + \mathbf{h}(\mathbf{x})^\top \boldsymbol{\gamma}_i \\
& \text{s.t.} \quad \boldsymbol{\gamma}_i \in \mathbb{R}_+^{M+J}, \boldsymbol{\mu}_i \in \mathbb{R}_+^K, \boldsymbol{\Gamma}_i \in \mathbb{S}_+^{M+J}, \boldsymbol{\Omega}_i \in \mathbb{S}_+^K, \mathbf{Y}_i \in \mathbb{R}^{K \times (M+J)} \quad \forall i \in [I] \\
& \quad \mathcal{Q} \boldsymbol{\mu}_i + \mathbf{q} = \mathcal{W}^\top \boldsymbol{\gamma}_i \quad \forall i \in [I] \\
& \quad \mathcal{Q}_j^\top \boldsymbol{\Omega}_i \mathcal{Q}_j - 2 \mathcal{Q}_j^\top \mathbf{Y}_i \mathcal{W}_{:,j} + \mathcal{W}_{:,j}^\top \boldsymbol{\Gamma}_i \mathcal{W}_{:,j} = \mathbf{q}_j^2 \quad \forall i \in [I] \quad \forall j \in [N_2+J] \\
& \quad \frac{1}{I} \sum_{i \in [I]} \text{tr}(\boldsymbol{\Omega}_i) - 2 \hat{\xi}_i^\top \boldsymbol{\mu}_i + \hat{\xi}_i^\top \hat{\xi}_i \leq \epsilon^2 \\
& \quad \begin{bmatrix} \boldsymbol{\Omega}_i & \mathbf{Y}_i & \boldsymbol{\mu}_i \\ \mathbf{Y}_i^\top & \boldsymbol{\Gamma}_i & \boldsymbol{\gamma}_i \\ \boldsymbol{\mu}_i^\top & \boldsymbol{\gamma}_i^\top & 1 \end{bmatrix} \succeq_{\mathcal{C}^*} \mathbf{0} \quad \forall i \in [I]
\end{aligned} \tag{14}$$

Proof. The claim follows from standard conic duality theory. Details are omitted for brevity. \square

In the following, we show that the worst-case expectation $\mathcal{Z}(\mathbf{x})$ is in fact equal to the optimal value $\underline{\mathcal{Z}}(\mathbf{x})$ of the completely positive program (14).

Theorem 3 (Completely Positive Reformulation). *For any fixed $\mathbf{x} \in \mathcal{X}$ we have $\mathcal{Z}(\mathbf{x}) = \underline{\mathcal{Z}}(\mathbf{x})$.*

Proof. Recall from Theorem 1 that $\mathcal{Z}(\mathbf{x})$ coincides with the optimal value of the moment problem (5). We first prove that $\mathcal{Z}(\mathbf{x}) \leq \underline{\mathcal{Z}}(\mathbf{x})$. To this end, we show that any feasible solution $\{\mathbb{P}_i\}_{i \in [I]}$ of (5) gives rise to a feasible solution to (14) with the same objective function value. Let $\mathbf{p}(\xi)$ be a measurable selector of

the dual feasible set mapping $\boldsymbol{\xi} \Rightarrow \{\mathbf{p} \in \mathbb{R}_+^M : \mathbf{Q}\boldsymbol{\xi} + \mathbf{q} = \mathbf{W}^\top \mathbf{p}\}$, $\boldsymbol{\xi} \in \Xi$, which exists due to [47, Corollary 14.6] and because problem (1) has sufficiently expensive recourse. Next, define $\boldsymbol{\pi}(\boldsymbol{\xi}) = (\mathbf{p}(\boldsymbol{\xi}), \mathbf{t} - \mathbf{S}\boldsymbol{\xi})$. By construction we have $\mathbf{Q}\boldsymbol{\xi} + \mathbf{q} = \mathbf{W}^\top \boldsymbol{\pi}(\boldsymbol{\xi})$ for all $\boldsymbol{\xi} \in \Xi$. Next, define the following candidate solution for (14):

$$\left. \begin{aligned} \boldsymbol{\mu}_i &= \int_{\Xi} \boldsymbol{\xi} \mathbb{P}_i(d\boldsymbol{\xi}), & \boldsymbol{\Omega}_i &= \int_{\Xi} \boldsymbol{\xi} \boldsymbol{\xi}^\top \mathbb{P}_i(d\boldsymbol{\xi}), \\ \boldsymbol{\gamma}_i &= \int_{\Xi} \boldsymbol{\pi}(\boldsymbol{\xi}) \mathbb{P}_i(d\boldsymbol{\xi}), & \boldsymbol{\Gamma}_i &= \int_{\Xi} \boldsymbol{\pi}(\boldsymbol{\xi}) \boldsymbol{\pi}(\boldsymbol{\xi})^\top \mathbb{P}_i(d\boldsymbol{\xi}), & \mathbf{Y}_i &= \int_{\Xi} \boldsymbol{\xi} \boldsymbol{\pi}(\boldsymbol{\xi})^\top \mathbb{P}_i(d\boldsymbol{\xi}) \end{aligned} \right\} \forall i \in [I]. \quad (15)$$

Since $\tilde{\boldsymbol{\xi}}$ and $\boldsymbol{\pi}(\tilde{\boldsymbol{\xi}})$ are non-negative random vectors, the matrix of their moments of degree ≤ 2 is completely positive. Thus, the candidate solution (15) satisfies the last constraint in (14). We further have

$$\mathbf{Q}\boldsymbol{\xi} + \mathbf{q} = \mathbf{W}^\top \boldsymbol{\pi}(\boldsymbol{\xi}) \quad \forall \boldsymbol{\xi} \in \Xi \quad \implies \quad \mathbf{Q}\boldsymbol{\mu}_i + \mathbf{q} = \mathbf{W}^\top \boldsymbol{\gamma}_i$$

and

$$(\mathbf{Q}_{j\cdot}^\top \boldsymbol{\xi} + q_j)^2 = (\mathbf{W}_{j\cdot}^\top \boldsymbol{\pi}(\boldsymbol{\xi}))^2 \quad \forall \boldsymbol{\xi} \in \Xi \quad \implies \quad \mathbf{Q}_{j\cdot}^\top \boldsymbol{\Omega}_i \mathbf{Q}_{j\cdot} - 2\mathbf{Q}_{j\cdot}^\top \mathbf{Y}_i \mathbf{W}_{j\cdot} + \mathbf{W}_{j\cdot}^\top \boldsymbol{\Gamma}_i \mathbf{W}_{j\cdot} = q_j^2$$

for all $j \in [N_2 + J]$. Here, the implications follow from taking expectations with respect to \mathbb{P}_i on both sides of the semi-infinite constraints. Thus, the candidate solution (15) also satisfies the first and the second constraint systems in (14). Next, the feasibility of $\{\mathbb{P}_i\}_{i \in [I]}$ in the generalized moment problem (5) implies that $\frac{1}{I} \sum_{i \in [I]} \int_{\Xi} \|\boldsymbol{\xi} - \hat{\boldsymbol{\xi}}_i\|_2^2 \mathbb{P}_i(d\boldsymbol{\xi}) \leq \epsilon^2$. Expanding the squared norm term, we obtain

$$\epsilon^2 \geq \frac{1}{I} \sum_{i \in [I]} \int_{\Xi} \boldsymbol{\xi}^\top \boldsymbol{\xi} - 2\hat{\boldsymbol{\xi}}_i^\top \boldsymbol{\xi} + \hat{\boldsymbol{\xi}}_i^\top \hat{\boldsymbol{\xi}}_i \mathbb{P}_i(d\boldsymbol{\xi}) = \frac{1}{I} \sum_{i \in [I]} \text{tr}(\boldsymbol{\Omega}_i) - 2\hat{\boldsymbol{\xi}}_i^\top \boldsymbol{\mu}_i + \hat{\boldsymbol{\xi}}_i^\top \hat{\boldsymbol{\xi}}_i,$$

and thus the candidate solution (15) also satisfies the penultimate constraint in (14). Lastly, the objective function of (5) can be reformulated as

$$\frac{1}{I} \sum_{i \in [I]} \int_{\Xi} Z(\mathbf{x}, \boldsymbol{\xi}) \mathbb{P}_i(d\boldsymbol{\xi}) = \frac{1}{I} \sum_{i \in [I]} \int_{\Xi} (\mathcal{T}(\mathbf{x})\boldsymbol{\xi} + \mathbf{h}(\mathbf{x}))^\top \boldsymbol{\pi}(\boldsymbol{\xi}) \mathbb{P}_i(d\boldsymbol{\xi}) = \frac{1}{I} \sum_{i \in [I]} \text{tr}(\mathcal{T}(\mathbf{x})\mathbf{Y}_i) + \mathbf{h}(\mathbf{x})^\top \boldsymbol{\gamma}_i,$$

where the first equality follows from the observation that $(\mathcal{T}(\mathbf{x})\boldsymbol{\xi} + \mathbf{h}(\mathbf{x}))^\top \boldsymbol{\pi}(\boldsymbol{\xi}) = (\mathbf{T}(\mathbf{x})\boldsymbol{\xi} + \mathbf{h}(\mathbf{x}))^\top \mathbf{p}(\boldsymbol{\xi})$ for all $\boldsymbol{\xi} \in \Xi$. Note that the rightmost term in the above equation corresponds to the objective value of the candidate solution (15) in (14). We have thus shown that from any feasible solution $\{\mathbb{P}_i\}_{i \in [I]}$ to the moment problem (5) we can construct a feasible solution $\{(\boldsymbol{\mu}_i, \boldsymbol{\gamma}_i, \boldsymbol{\Omega}_i, \boldsymbol{\Gamma}_i, \mathbf{Y}_i)\}_{i \in [I]}$ to the completely positive program (14) that attains the same objective value. This demonstrates that $\mathcal{Z}(\mathbf{x}) \leq \underline{\mathcal{Z}}(\mathbf{x})$.

To prove the converse inequality, consider any feasible solution $\{(\boldsymbol{\mu}_i, \boldsymbol{\gamma}_i, \boldsymbol{\Omega}_i, \boldsymbol{\Gamma}_i, \mathbf{Y}_i)\}_{i \in [I]}$ to (14), which gives rise to a moment matrix with the following completely positive decomposition,

$$\begin{bmatrix} \boldsymbol{\Omega}_i & \mathbf{Y}_i & \boldsymbol{\mu}_i \\ \mathbf{Y}_i^\top & \boldsymbol{\Gamma}_i & \boldsymbol{\gamma}_i \\ \boldsymbol{\mu}_i^\top & \boldsymbol{\gamma}_i^\top & 1 \end{bmatrix} = \sum_{\ell \in \mathcal{L}_i} \begin{bmatrix} \chi_{i\ell} \\ \boldsymbol{\eta}_{i\ell} \\ \alpha_{i\ell} \end{bmatrix} \begin{bmatrix} \chi_{i\ell} \\ \boldsymbol{\eta}_{i\ell} \\ \alpha_{i\ell} \end{bmatrix}^\top, \quad (16)$$

where \mathcal{L}_i is a finite index set, while $\boldsymbol{\chi}_{i\ell} \in \mathbb{R}_+^K$, $\boldsymbol{\eta}_{i\ell} \in \mathbb{R}_+^{M+J}$ and $\alpha_{i\ell} \in \mathbb{R}_+$ for every $\ell \in \mathcal{L}_i$. Partitioning \mathcal{L}_i into $\mathcal{L}_i^+ = \{\ell \in \mathcal{L} : \alpha_{i\ell} > 0\}$ and $\mathcal{L}_i^0 = \{\ell \in \mathcal{L} : \alpha_{i\ell} = 0\}$, the decomposition (16) reduces to

$$\begin{bmatrix} \boldsymbol{\Omega}_i & \mathbf{Y}_i & \boldsymbol{\mu}_i \\ \mathbf{Y}_i^\top & \boldsymbol{\Gamma}_i & \boldsymbol{\gamma}_i \\ \boldsymbol{\mu}_i^\top & \boldsymbol{\gamma}_i^\top & 1 \end{bmatrix} = \sum_{\ell \in \mathcal{L}_i^+} \begin{bmatrix} \boldsymbol{\chi}_{i\ell} \boldsymbol{\chi}_{i\ell}^\top & \boldsymbol{\chi}_{i\ell} \boldsymbol{\eta}_{i\ell}^\top & \alpha_{i\ell} \boldsymbol{\chi}_{i\ell} \\ \boldsymbol{\eta}_{i\ell} \boldsymbol{\chi}_{i\ell}^\top & \boldsymbol{\eta}_{i\ell} \boldsymbol{\eta}_{i\ell}^\top & \alpha_{i\ell} \boldsymbol{\eta}_{i\ell} \\ \alpha_{i\ell} \boldsymbol{\chi}_{i\ell}^\top & \alpha_{i\ell} \boldsymbol{\eta}_{i\ell}^\top & \alpha_{i\ell}^2 \end{bmatrix} + \sum_{\ell \in \mathcal{L}_i^0} \begin{bmatrix} \boldsymbol{\chi}_{i\ell} \boldsymbol{\chi}_{i\ell}^\top & \boldsymbol{\chi}_{i\ell} \boldsymbol{\eta}_{i\ell}^\top & \mathbf{0} \\ \boldsymbol{\eta}_{i\ell} \boldsymbol{\chi}_{i\ell}^\top & \boldsymbol{\eta}_{i\ell} \boldsymbol{\eta}_{i\ell}^\top & \mathbf{0} \\ \mathbf{0}^\top & \mathbf{0}^\top & 0 \end{bmatrix}. \quad (17)$$

Next, we construct a sequence of discrete distributions \mathbb{P}_i^κ , $i \in [I]$, parametrized by $\kappa \in [0, 1]$, that satisfy

$$\left. \begin{aligned} \mathbb{P}_i^\kappa \left(\tilde{\boldsymbol{\xi}} = \frac{\boldsymbol{\chi}_{i\ell}}{\alpha_{i\ell}} \right) &= (1 - \kappa^2) \alpha_{i\ell}^2 & \forall \ell \in \mathcal{L}_i^+ \\ \mathbb{P}_i^\kappa \left(\tilde{\boldsymbol{\xi}} = \hat{\boldsymbol{\xi}}_i + \frac{1}{\kappa} \sqrt{|\mathcal{L}_i^0|} \boldsymbol{\chi}_{i\ell} \right) &= \frac{\kappa^2}{|\mathcal{L}_i^0|} & \forall \ell \in \mathcal{L}_i^0 \end{aligned} \right\} \forall i \in [I].$$

Observe that each \mathbb{P}_i^κ is indeed a probability distribution since $\sum_{\ell \in \mathcal{L}_i^+} \alpha_{i\ell}^2 = 1$ due to (17). Lemma 2 below implies that $\boldsymbol{\chi}_{i\ell}/\alpha_{i\ell} \in \Xi$ for every $\ell \in \mathcal{L}_i^+$ and $\hat{\boldsymbol{\xi}}_i + \frac{1}{\kappa} \sqrt{|\mathcal{L}_i^0|} \boldsymbol{\chi}_{i\ell} \in \Xi$ for every $\ell \in \mathcal{L}_i^0$. Thus, \mathbb{P}_i^κ is supported on Ξ . We further have

$$\begin{aligned} \frac{1}{I} \sum_{i \in [I]} \mathbb{E}_{\mathbb{P}_i^\kappa} \left[\|\tilde{\boldsymbol{\xi}} - \hat{\boldsymbol{\xi}}_i\|_2^2 \right] &= \frac{1}{I} \sum_{i \in [I]} \left[\sum_{\ell \in \mathcal{L}_i^+} (1 - \kappa^2) \alpha_{i\ell}^2 \|\boldsymbol{\chi}_{i\ell}/\alpha_{i\ell} - \hat{\boldsymbol{\xi}}_i\|_2^2 + \sum_{\ell \in \mathcal{L}_i^0} \frac{\kappa^2}{|\mathcal{L}_i^0|} \left\| \hat{\boldsymbol{\xi}}_i + \frac{1}{\kappa} \sqrt{|\mathcal{L}_i^0|} \boldsymbol{\chi}_{i\ell} - \hat{\boldsymbol{\xi}}_i \right\|_2^2 \right] \\ &\leq \frac{1}{I} \sum_{i \in [I]} \left[\sum_{\ell \in \mathcal{L}_i^+} \left(\boldsymbol{\chi}_{i\ell}^\top \boldsymbol{\chi}_{i\ell} - 2 \hat{\boldsymbol{\xi}}_i^\top (\alpha_{i\ell} \boldsymbol{\chi}_{i\ell}) \right) + \hat{\boldsymbol{\xi}}_i^\top \hat{\boldsymbol{\xi}}_i + \sum_{\ell \in \mathcal{L}_i^0} \boldsymbol{\chi}_{i\ell}^\top \boldsymbol{\chi}_{i\ell} \right] \\ &= \frac{1}{I} \sum_{i \in [I]} \text{tr}(\boldsymbol{\Omega}_i) - 2 \hat{\boldsymbol{\xi}}_i^\top \boldsymbol{\mu}_i + \hat{\boldsymbol{\xi}}_i^\top \hat{\boldsymbol{\xi}}_i \leq \epsilon^2, \end{aligned}$$

where the first inequality holds since $(1 - \kappa^2) \leq 1$, the second equality follows from the decomposition (17), and the last inequality follows from the penultimate constraint in (14). Thus, the distributions \mathbb{P}_i^κ , $i \in [I]$, are feasible in the generalized moment problem (5). We next construct feasible solutions for the dual recourse problem (4). For any $i \in [I]$ and $\ell \in \mathcal{L}_i$, we define $\boldsymbol{\rho}_{i\ell}$ as the vector of the first M elements of $\boldsymbol{\eta}_{i\ell} \in \mathbb{R}_+^{M+J}$. Lemma 2 implies that $\mathbf{Q} \boldsymbol{\chi}_{i\ell}/\alpha_{i\ell} + \mathbf{q} = \mathbf{W}^\top \boldsymbol{\rho}_{i\ell}/\alpha_{i\ell}$ for every $\ell \in \mathcal{L}_i^+$. Thus, $\mathbf{p}_{i\ell} = \boldsymbol{\rho}_{i\ell}/\alpha_{i\ell}$ is feasible in the dual recourse problem (4) at $\boldsymbol{\xi} = \boldsymbol{\chi}_{i\ell}/\alpha_{i\ell}$ for $\ell \in \mathcal{L}_i^+$. Next, for any $i \in [I]$, let $\hat{\mathbf{p}}_i$ be a feasible solution to the dual recourse problem (4) at $\boldsymbol{\xi} = \hat{\boldsymbol{\xi}}_i$, which exists because problem (1) has sufficiently expensive recourse. Hence, we have $\mathbf{Q} \hat{\boldsymbol{\xi}}_i + \mathbf{q} = \mathbf{W}^\top \hat{\mathbf{p}}_i$. Lemma 2 further implies that $\mathbf{Q} \boldsymbol{\chi}_{i\ell} = \mathbf{W}^\top \boldsymbol{\rho}_{i\ell}$ for every $\ell \in \mathcal{L}_i^0$. Combining the last two equalities yields

$$\mathbf{Q} \left(\hat{\boldsymbol{\xi}}_i + \frac{1}{\kappa} \sqrt{|\mathcal{L}_i^0|} \boldsymbol{\chi}_{i\ell} \right) + \mathbf{q} = \mathbf{W}^\top \left(\hat{\mathbf{p}}_i + \frac{1}{\kappa} \sqrt{|\mathcal{L}_i^0|} \boldsymbol{\rho}_{i\ell} \right) \quad \forall \ell \in \mathcal{L}_i^0.$$

Thus, $\mathbf{p}_{i\ell} = \hat{\mathbf{p}}_i + \frac{1}{\kappa} \sqrt{|\mathcal{L}_i^0|} \boldsymbol{\rho}_{i\ell}$ constitutes a feasible solution in the dual recourse problem (4) at $\boldsymbol{\xi} = \hat{\boldsymbol{\xi}}_i + \frac{1}{\kappa} \sqrt{|\mathcal{L}_i^0|} \boldsymbol{\chi}_{i\ell}$ for $\ell \in \mathcal{L}_i^0$. In summary, we have

$$Z \left(\mathbf{x}, \frac{\boldsymbol{\chi}_{i\ell}}{\alpha_{i\ell}} \right) = Z_d \left(\mathbf{x}, \frac{\boldsymbol{\chi}_{i\ell}}{\alpha_{i\ell}} \right) \geq \left(\mathbf{T}(\mathbf{x}) \frac{\boldsymbol{\chi}_{i\ell}}{\alpha_{i\ell}} + \mathbf{h}(\mathbf{x}) \right)^\top \mathbf{p}_{i\ell} \quad \forall \ell \in \mathcal{L}_i^+$$

and

$$Z\left(\mathbf{x}, \hat{\boldsymbol{\xi}}_i + \frac{1}{\kappa} \sqrt{|\mathcal{L}_i^0|} \boldsymbol{\chi}_{il}\right) = Z_d\left(\mathbf{x}, \hat{\boldsymbol{\xi}}_i + \frac{1}{\kappa} \sqrt{|\mathcal{L}_i^0|} \boldsymbol{\chi}_{il}\right) \geq \left(\mathbf{T}(\mathbf{x}) \left(\hat{\boldsymbol{\xi}}_i + \frac{1}{\kappa} \sqrt{|\mathcal{L}_i^0|} \boldsymbol{\chi}_{il}\right) + \mathbf{h}(\mathbf{x})\right)^\top \mathbf{p}_{il} \quad \forall \ell \in \mathcal{L}_i^0.$$

Using these estimates, we can now bound the objective value of the discrete distributions \mathbb{P}_i^κ , $i \in [I]$, in (5).

Specifically, we obtain

$$\begin{aligned} & \frac{1}{I} \sum_{i \in [I]} \mathbb{E}_{\mathbb{P}_i^\kappa} \left[Z(\mathbf{x}, \tilde{\boldsymbol{\xi}}) \right] \\ &= \frac{1}{I} \sum_{i \in [I]} \left[\sum_{\ell \in \mathcal{L}_i^+} (1 - \kappa^2) \alpha_{i\ell}^2 Z\left(\mathbf{x}, \frac{\boldsymbol{\chi}_{i\ell}}{\alpha_{i\ell}}\right) + \sum_{\ell \in \mathcal{L}_i^0} \frac{\kappa^2}{|\mathcal{L}_i^0|} Z\left(\mathbf{x}, \hat{\boldsymbol{\xi}}_i + \frac{\sqrt{|\mathcal{L}_i^0|}}{\kappa} \boldsymbol{\chi}_{i\ell}\right) \right] \\ &\geq \frac{1}{I} \sum_{i \in [I]} \left[\sum_{\ell \in \mathcal{L}_i^+} (1 - \kappa^2) \alpha_{i\ell}^2 \left(\mathbf{T}(\mathbf{x}) \frac{\boldsymbol{\chi}_{i\ell}}{\alpha_{i\ell}} + \mathbf{h}(\mathbf{x})\right)^\top \mathbf{p}_{i\ell} + \sum_{\ell \in \mathcal{L}_i^0} \frac{\kappa^2}{|\mathcal{L}_i^0|} \left(\mathbf{T}(\mathbf{x}) \left(\hat{\boldsymbol{\xi}}_i + \frac{\sqrt{|\mathcal{L}_i^0|}}{\kappa} \boldsymbol{\chi}_{i\ell}\right) + \mathbf{h}(\mathbf{x})\right)^\top \mathbf{p}_{i\ell} \right] \\ &= \frac{1}{I} \sum_{i \in [I]} \left[\sum_{\ell \in \mathcal{L}_i^+} (1 - \kappa^2) \alpha_{i\ell}^2 \left(\mathcal{T}(\mathbf{x}) \frac{\boldsymbol{\chi}_{i\ell}}{\alpha_{i\ell}} + \mathbf{h}(\mathbf{x})\right)^\top \frac{\boldsymbol{\eta}_{i\ell}}{\alpha_{i\ell}} \right. \\ &\quad \left. + \sum_{\ell \in \mathcal{L}_i^0} \frac{\kappa^2}{|\mathcal{L}_i^0|} \left(\mathcal{T}(\mathbf{x}) \left(\hat{\boldsymbol{\xi}}_i + \frac{\sqrt{|\mathcal{L}_i^0|}}{\kappa} \boldsymbol{\chi}_{i\ell}\right) + \mathbf{h}(\mathbf{x})\right)^\top \left([\hat{\mathbf{p}}_i^\top \mathbf{0}]^\top + \frac{\sqrt{|\mathcal{L}_i^0|}}{\kappa} \boldsymbol{\eta}_{i\ell}\right) \right] \\ &= \frac{1}{I} \sum_{i \in [I]} \left[\sum_{\ell \in \mathcal{L}_i^+} (1 - \kappa^2) [\text{tr}(\mathcal{T}(\mathbf{x}) \boldsymbol{\chi}_{i\ell} \boldsymbol{\eta}_{i\ell}^\top) + \mathbf{h}(\mathbf{x})^\top (\alpha_{i\ell} \boldsymbol{\eta}_{i\ell})] + \sum_{\ell \in \mathcal{L}_i^0} \text{tr}(\mathcal{T}(\mathbf{x}) \boldsymbol{\chi}_{i\ell} \boldsymbol{\eta}_{i\ell}^\top) \right] \\ &\quad + \frac{1}{I} \sum_{i \in [I]} \left[\sum_{\ell \in \mathcal{L}_i^0} \frac{\kappa^2}{|\mathcal{L}_i^0|} \left(\left(\mathbf{T}(\mathbf{x}) \left(\hat{\boldsymbol{\xi}}_i + \frac{\sqrt{|\mathcal{L}_i^0|}}{\kappa} \boldsymbol{\chi}_{i\ell}\right) + \mathbf{h}(\mathbf{x})\right)^\top \hat{\mathbf{p}}_i + \left(\mathcal{T}(\mathbf{x}) \hat{\boldsymbol{\xi}}_i + \mathbf{h}(\mathbf{x})\right)^\top \frac{\sqrt{|\mathcal{L}_i^0|}}{\kappa} \boldsymbol{\eta}_{i\ell} \right) \right]. \end{aligned}$$

Together with the decomposition (17), the above estimate implies that

$$\lim_{\kappa \downarrow 0} \frac{1}{I} \sum_{i \in [I]} \mathbb{E}_{\mathbb{P}_i^\kappa} \left[Z(\mathbf{x}, \tilde{\boldsymbol{\xi}}) \right] \geq \frac{1}{I} \sum_{i \in [I]} \text{tr}(\mathcal{T}(\mathbf{x}) \mathbf{Y}_i) + \mathbf{h}(\mathbf{x})^\top \boldsymbol{\gamma}_i.$$

We have therefore shown that any feasible solution to (14) can be used to construct a sequence of feasible solutions to the generalized moment problem (5) that asymptotically attain a (weakly) larger objective value.

This demonstrates that $\mathcal{Z}(\mathbf{x}) \geq \underline{\mathcal{Z}}(\mathbf{x})$. Thus the claim follows. \square

The proof of Theorem 3 relies on the following lemma, which is inspired by Lemma 2.2 in [11] and Proposition 3.1 in [42].

Lemma 2. *If $\{(\boldsymbol{\mu}_i, \boldsymbol{\gamma}_i, \boldsymbol{\Omega}_i, \boldsymbol{\Gamma}_i, \mathbf{Y}_i)\}_{i \in [I]}$ is feasible in (14) and if $(\boldsymbol{\chi}_{i\ell}, \boldsymbol{\eta}_{i\ell}, \alpha_{i\ell}) \in \mathbb{R}_+^K \times \mathbb{R}_+^{M+J} \times \mathbb{R}_+$, $\ell \in \mathcal{L}_i$, satisfies (16) for some $i \in [I]$, then*

$$\boldsymbol{\chi}_{i\ell} / \alpha_{i\ell} \in \Xi \quad \text{and} \quad \mathbf{Q}(\boldsymbol{\chi}_{i\ell} / \alpha_{i\ell}) + \mathbf{q} = \mathbf{W}^\top (\boldsymbol{\rho}_{i\ell} / \alpha_{i\ell}) \quad \forall \ell \in \mathcal{L}_i^+,$$

while

$$\boldsymbol{\chi}_{i\ell} \in \text{recc}(\Xi) \quad \text{and} \quad \mathbf{Q}\boldsymbol{\chi}_{i\ell} = \mathbf{W}^\top \boldsymbol{\rho}_{i\ell} \quad \forall \ell \in \mathcal{L}_i^0,$$

where $\text{recc}(\Xi) = \{\boldsymbol{\xi} \in \mathbb{R}_+^K : \mathbf{S}\boldsymbol{\xi} \leq \mathbf{0}\}$ is the recession cone of Ξ , and $\boldsymbol{\rho}_{i\ell}$ is the vector of the first M elements of $\boldsymbol{\eta}_{i\ell}$.

Proof. We substitute the decomposition (16) into the constraints of problem (14) to obtain

$$\sum_{\ell \in \mathcal{L}_i} \alpha_{i\ell} \left(\boldsymbol{w}_{:j}^\top \boldsymbol{\eta}_{i\ell} - \boldsymbol{q}_j^\top \boldsymbol{\chi}_{i\ell} \right) = q_j \quad (18)$$

and

$$\sum_{\ell \in \mathcal{L}_i} \left(\boldsymbol{w}_{:j}^\top \boldsymbol{\eta}_{i\ell} - \boldsymbol{q}_j^\top \boldsymbol{\chi}_{i\ell} \right)^2 = q_j^2 \quad (19)$$

for every $j \in [N_2 + J]$. Squaring (18) and eliminating q_j^2 by using (19) yields

$$\left(\sum_{\ell \in \mathcal{L}_i} \alpha_{i\ell} \left(\boldsymbol{w}_{:j}^\top \boldsymbol{\eta}_{i\ell} - \boldsymbol{q}_j^\top \boldsymbol{\chi}_{i\ell} \right) \right)^2 = \sum_{\ell \in \mathcal{L}_i} \left(\boldsymbol{w}_{:j}^\top \boldsymbol{\eta}_{i\ell} - \boldsymbol{q}_j^\top \boldsymbol{\chi}_{i\ell} \right)^2 = \left(\sum_{\ell \in \mathcal{L}_i} \alpha_{i\ell}^2 \right) \sum_{\ell \in \mathcal{L}_i} \left(\boldsymbol{w}_{:j}^\top \boldsymbol{\eta}_{i\ell} - \boldsymbol{q}_j^\top \boldsymbol{\chi}_{i\ell} \right)^2.$$

Here, the second equality follows from the fact that $\sum_{\ell \in \mathcal{L}} \alpha_{i\ell}^2 = 1$. The tightness condition of the Cauchy-Schwartz inequality therefore implies that there exists $\tau \in \mathbb{R}$ with

$$\boldsymbol{w}_{:j}^\top \boldsymbol{\eta}_{i\ell} - \boldsymbol{q}_j^\top \boldsymbol{\chi}_{i\ell} = \tau \alpha_{i\ell} \quad \forall \ell \in \mathcal{L}_i. \quad (20)$$

Thus, we have

$$\boldsymbol{w}_{:j}^\top \boldsymbol{\eta}_{i\ell} - \boldsymbol{q}_j^\top \boldsymbol{\chi}_{i\ell} = 0 \quad \forall \ell \in \mathcal{L}_i^0,$$

which confirms the second claim. Next, we observe from (18) and (20) that

$$q_j = \sum_{\ell \in \mathcal{L}_i} \alpha_{i\ell} \left(\boldsymbol{w}_{:j}^\top \boldsymbol{\eta}_{i\ell} - \boldsymbol{q}_j^\top \boldsymbol{\chi}_{i\ell} \right) = \sum_{\ell \in \mathcal{L}_i} \tau \alpha_{i\ell}^2 = \tau. \quad (21)$$

Replacing τ with q_j in (20) and using the fact that $\alpha_{i\ell} > 0$ for $\ell \in \mathcal{L}_i^+$ yields

$$\boldsymbol{w}_{:j}^\top (\boldsymbol{\eta}_{i\ell} / \alpha_{i\ell}) - \boldsymbol{q}_j^\top (\boldsymbol{\chi}_{i\ell} / \alpha_{i\ell}) = q_j \quad \forall \ell \in \mathcal{L}_i^+,$$

which establishes the first claim. □

3.3 A Copositive Reformulation of Problem (1)

So far, we have seen that $\mathcal{Z}(\boldsymbol{x}) = \underline{\mathcal{Z}}(\boldsymbol{x}) \leq \overline{\mathcal{Z}}(\boldsymbol{x})$. Unfortunately, as we exemplify below, the duality gap between $\underline{\mathcal{Z}}(\boldsymbol{x})$ and $\overline{\mathcal{Z}}(\boldsymbol{x})$ can be strictly positive.

Example 1 (Infinite Duality Gap). *Consider the following pair of primal and dual recourse problems*

$$Z(x, \xi) = \inf_{y \in \mathbb{R}} \{(\xi - 1)y : \xi - 1 \leq 0y\} \quad \text{and} \quad Z_d(x, \xi) = \sup_{p \in \mathbb{R}_+} \{(\xi - 1)p : \xi - 1 = 0p\},$$

respectively, and set $\Xi = \{\xi \in \mathbb{R}_+ : \xi \leq 1, -\xi \leq -1\} = \{1\}$. Assume that there is only one sample $\hat{\xi}_1 = 1$ and that the Wasserstein radius is set to $\epsilon = 1$. Note that both linear programs are feasible with the same optimal value 0 for $\xi = 1$. Theorem 3 therefore implies that $\underline{\mathcal{Z}}(\mathbf{x}) = \mathcal{Z}(\mathbf{x}) = 0$.

Under the current setting, the extended recourse parameters defined in (11) are given by

$$\mathcal{Q} = \begin{bmatrix} 1 \\ 1 \\ -1 \end{bmatrix}, \mathbf{q} = \begin{bmatrix} -1 \\ -1 \\ 1 \end{bmatrix}, \mathcal{T}(\mathbf{x}) = \begin{bmatrix} 1 \\ 0 \\ 0 \end{bmatrix}, \mathbf{W} = \begin{bmatrix} 0 & 0 & 0 \\ 0 & -1 & 0 \\ 0 & 0 & -1 \end{bmatrix}, \text{ and } \mathbf{h}(\mathbf{x}) = \begin{bmatrix} -1 \\ 0 \\ 0 \end{bmatrix}.$$

Hence, the copositive program (10) simplifies to

$$\begin{aligned} \overline{\mathcal{Z}}(\mathbf{x}) = & \inf \quad s - \psi_1 - \psi_2 + \psi_3 + \phi_1 + \phi_2 + \phi_3 \\ \text{s.t. } & \lambda \in \mathbb{R}_+, s \in \mathbb{R}, \boldsymbol{\psi}, \boldsymbol{\phi} \in \mathbb{R}^3 \\ & \begin{bmatrix} \lambda + \phi_1 + \phi_2 + \phi_3 & -\frac{1}{2} & \phi_2 & -\phi_3 & -\lambda - \frac{\psi_1 + \psi_2 - \psi_3}{2} \\ -\frac{1}{2} & 0 & 0 & 0 & \frac{1}{2} \\ \phi_2 & 0 & \phi_2 & 0 & -\frac{\psi_2}{2} \\ -\phi_3 & 0 & 0 & \phi_3 & -\frac{\psi_3}{2} \\ -\lambda - \frac{\psi_1 + \psi_2 - \psi_3}{2} & \frac{1}{2} & -\frac{\psi_2}{2} & -\frac{\psi_3}{2} & s \end{bmatrix} \succeq_{\mathcal{C}} \mathbf{0}. \end{aligned} \quad (22)$$

However, multiplying the copositive constraint from both sides with the vector $[\theta \ 1 \ 0 \ 0 \ 0]^\top$, $\theta \geq 0$, implies that $(\lambda + \phi_1 + \phi_2 + \phi_3)\theta^2 - \theta \geq 0$ for all $\theta \geq 0$. As no values of λ , ϕ_1 , ϕ_2 and ϕ_3 can satisfy this inequality, problem (22) is infeasible, i.e., $\overline{\mathcal{Z}}(\mathbf{x}) = +\infty$. Thus, there is an infinite duality gap between $\underline{\mathcal{Z}}(\mathbf{x})$ and $\overline{\mathcal{Z}}(\mathbf{x})$.

Even though $\underline{\mathcal{Z}}(\mathbf{x})$ and $\overline{\mathcal{Z}}(\mathbf{x})$ may differ, one can prove that $\underline{\mathcal{Z}}(\mathbf{x}) = \mathcal{Z}(\mathbf{x}) = \overline{\mathcal{Z}}(\mathbf{x})$ if the two-stage distributionally robust linear program (1) has complete recourse. To show this, we first prove two lemmas.

Lemma 3. *If Problem (1) has complete recourse, then $\mathbf{W}\mathbf{W}^\top \succ_{\mathcal{C}} \mathbf{0}$.*

Proof. The complete recourse property is equivalent to the unboundedness of the linear program

$$\begin{aligned} & \text{maximize} \quad z \\ & \text{subject to} \quad \mathbf{y} \in \mathbb{R}^{N_2}, z \in \mathbb{R} \\ & \quad \quad \quad \mathbf{W}\mathbf{y} \geq z\mathbf{e}, \end{aligned}$$

whose dual linear program is given by

$$\begin{aligned} & \text{minimize} \quad 0 \\ & \text{subject to} \quad \boldsymbol{\lambda} \in \mathbb{R}_+^M \\ & \quad \quad \quad \mathbf{e}^\top \boldsymbol{\lambda} = 1, \mathbf{W}^\top \boldsymbol{\lambda} = \mathbf{0}. \end{aligned}$$

As the primal problem is unbounded, the dual problem is infeasible by weak duality, implying that $\mathbf{W}^\top \boldsymbol{\lambda} \neq \mathbf{0}$ for all $\boldsymbol{\lambda} \in \mathbb{R}_+^M$ such that $\mathbf{e}^\top \boldsymbol{\lambda} = 1$. By rescaling, we thus have $\boldsymbol{\lambda}^\top \mathbf{W}\mathbf{W}^\top \boldsymbol{\lambda} > 0$ for all $\boldsymbol{\lambda} \in \mathbb{R}_+^M$ such that $\boldsymbol{\lambda} \neq \mathbf{0}$. This implies that $\mathbf{W}\mathbf{W}^\top$ lies in the interior of the copositive cone \mathcal{C} . \square

Lemma 4 (Copositive Schur Complements). *Consider the symmetric matrix*

$$M = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B}^\top & \mathbf{C} \end{bmatrix},$$

with $\mathbf{A} \succ \mathbf{0}$. We then have $M \succ_c \mathbf{0}$ if $\mathbf{C} - \mathbf{B}^\top \mathbf{A}^{-1} \mathbf{B} \succ_c \mathbf{0}$.

Proof. Multiplying the matrix M from both sides with a non-negative vector $[\boldsymbol{\xi}^\top \ \boldsymbol{\rho}^\top]^\top \in \mathbb{R}_+^{K+M}$ satisfying $\mathbf{e}^\top \boldsymbol{\xi} + \mathbf{e}^\top \boldsymbol{\rho} = 1$, we obtain

$$\begin{aligned} [\boldsymbol{\xi}^\top \ \boldsymbol{\rho}^\top] M [\boldsymbol{\xi}^\top \ \boldsymbol{\rho}^\top]^\top &= \boldsymbol{\xi}^\top \mathbf{A} \boldsymbol{\xi} + 2 \boldsymbol{\xi}^\top \mathbf{B} \boldsymbol{\rho} + \boldsymbol{\rho}^\top \mathbf{C} \boldsymbol{\rho} \\ &= (\boldsymbol{\xi} + \mathbf{A}^{-1} \mathbf{B} \boldsymbol{\rho})^\top \mathbf{A} (\boldsymbol{\xi} + \mathbf{A}^{-1} \mathbf{B} \boldsymbol{\rho}) + \boldsymbol{\rho}^\top (\mathbf{C} - \mathbf{B}^\top \mathbf{A}^{-1} \mathbf{B}) \boldsymbol{\rho}. \end{aligned}$$

Since $\mathbf{A} \succ \mathbf{0}$, the term $(\boldsymbol{\xi} + \mathbf{A}^{-1} \mathbf{B} \boldsymbol{\rho})^\top \mathbf{A} (\boldsymbol{\xi} + \mathbf{A}^{-1} \mathbf{B} \boldsymbol{\rho})$ is non-negative. If $\boldsymbol{\rho} = \mathbf{0}$, then we have $\mathbf{e}^\top \boldsymbol{\xi} = 1$, which implies that this term is positive. If $\boldsymbol{\rho} \neq \mathbf{0}$, then the assumption $\mathbf{C} - \mathbf{B}^\top \mathbf{A}^{-1} \mathbf{B} \succ_c \mathbf{0}$ implies that the term $\boldsymbol{\rho}^\top (\mathbf{C} - \mathbf{B}^\top \mathbf{A}^{-1} \mathbf{B}) \boldsymbol{\rho}$ is positive. In both cases, by rescaling we find that $[\boldsymbol{\xi}^\top \ \boldsymbol{\rho}^\top] M [\boldsymbol{\xi}^\top \ \boldsymbol{\rho}^\top]^\top$ for all $\boldsymbol{\xi} \in \mathbb{R}_+^K$ and all $\boldsymbol{\rho} \in \mathbb{R}_+^M$ such that $[\boldsymbol{\xi}^\top \ \boldsymbol{\rho}^\top]^\top \neq \mathbf{0}$. Hence, $M \succ_c \mathbf{0}$. \square

Lemmas 3 and 4 enable us to prove the following exactness result.

Theorem 4. *If problem (1) has complete recourse, then $\mathcal{Z}(\mathbf{x}) = \underline{\mathcal{Z}}(\mathbf{x}) = \overline{\mathcal{Z}}(\mathbf{x})$ for any fixed $\mathbf{x} \in \mathcal{X}$.*

Proof. We already know from Theorem 3 that $\mathcal{Z}(\mathbf{x}) = \underline{\mathcal{Z}}(\mathbf{x})$. To show that $\overline{\mathcal{Z}}(\mathbf{x}) = \underline{\mathcal{Z}}(\mathbf{x})$, it suffices to prove strong duality between problems (10) and (14). Specifically, we will construct a Slater point $(\lambda^s, (s_i^s, \boldsymbol{\phi}_i^s, \boldsymbol{\psi}_i^s)_{i \in [I]})$ for problem (10). To this end, we first set $\boldsymbol{\phi}_i^s = \mathbf{e}$ and $\boldsymbol{\psi}_i^s = \mathbf{0}$ for all $i \in [I]$. Using this solution, the i -th constraint matrix in (10) can be decomposed as

$$\begin{bmatrix} \lambda^s \mathbb{I} & -\frac{1}{2} \mathbf{T}(\mathbf{x})^\top - \mathbf{Q}^\top \mathbf{W}^\top & \mathbf{S}^\top & \mathbf{0} \\ -\frac{1}{2} \mathbf{T}(\mathbf{x}) - \mathbf{W} \mathbf{Q} & \mathbf{W} \mathbf{W}^\top & \mathbf{0} & -\frac{1}{2} \mathbf{h}(\mathbf{x}) \\ \mathbf{S} & \mathbf{0} & \mathbb{I} & \mathbf{0} \\ \mathbf{0}^\top & -\frac{1}{2} \mathbf{h}(\mathbf{x})^\top & \mathbf{0}^\top & \lambda^s \end{bmatrix} + \begin{bmatrix} \mathbf{Q}^\top \mathbf{Q} & \mathbf{0} & \mathbf{0} & -\lambda^s \hat{\boldsymbol{\xi}}_i \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ -\lambda^s \hat{\boldsymbol{\xi}}_i^\top & \mathbf{0}^\top & \mathbf{0}^\top & s_i^s - \lambda^s \end{bmatrix}. \quad (23)$$

Next, we select s_i^s large enough to ensure that the right matrix in (23) is positive semidefinite. In this case, a Slater point can be obtained by ensuring that the left matrix is strictly copositive. As problem (1) has complete recourse, Lemma 3 is applicable and implies that

$$\begin{bmatrix} \mathbf{W} \mathbf{W}^\top & \mathbf{0} \\ \mathbf{0} & \mathbb{I} \end{bmatrix} \succ_c \mathbf{0}.$$

Moreover Lemma 4 implies that the left matrix in (23) is strictly copositive if

$$\begin{bmatrix} \mathbf{W} \mathbf{W}^\top & \mathbf{0} \\ \mathbf{0} & \mathbb{I} \end{bmatrix} \succ_c \frac{1}{\lambda^s} \begin{bmatrix} -\frac{1}{2} \mathbf{T}(\mathbf{x})^\top + \mathbf{Q}^\top \mathbf{W}^\top & \mathbf{S}^\top \\ -\frac{1}{2} \mathbf{h}(\mathbf{x})^\top & \mathbf{0}^\top \end{bmatrix}^\top \begin{bmatrix} -\frac{1}{2} \mathbf{T}(\mathbf{x})^\top + \mathbf{Q}^\top \mathbf{W}^\top & \mathbf{S}^\top \\ -\frac{1}{2} \mathbf{h}(\mathbf{x})^\top & \mathbf{0}^\top \end{bmatrix},$$

which is true whenever λ^s is sufficiently large. We have therefore constructed a strictly feasible solution to (10). Thus, $\overline{\mathcal{Z}}(\mathbf{x}) = \underline{\mathcal{Z}}(\mathbf{x}) = \mathcal{Z}(\mathbf{x})$ by strong conic duality. \square

Theorem 4 implies that if problem (1) has complete recourse then it is equivalent to the copositive minimization problem obtained by replacing $\mathcal{Z}(\mathbf{x})$ in (1) with $\overline{\mathcal{Z}}(\mathbf{x})$. Conversely, if problem (1) fails to have complete recourse, it may only be possible to approximate $\mathcal{Z}(\mathbf{x})$ by the optimal value of a copositive minimization problem. To show this, we construct a relaxation of problem (10) parameterized by $\delta \geq 0$.

$$\begin{aligned}
& \overline{\mathcal{Z}}_\delta(\mathbf{x}) \\
= & \inf \quad \epsilon^2 \lambda + \frac{1}{I} \sum_{i \in [I]} \left[s_i + \mathbf{q}^\top \boldsymbol{\psi}_i - \lambda \|\hat{\boldsymbol{\xi}}_i\|_2^2 + \sum_{j \in [N_2+J]} \phi_{ij} q_j^2 \right] \\
& \text{s.t.} \quad \lambda \in \mathbb{R}_+, s_i \in \mathbb{R}, \boldsymbol{\psi}_i, \phi_i \in \mathbb{R}^{N_2+J} \quad \forall i \in [I] \\
& \quad \begin{bmatrix} \lambda \mathbb{I} + \mathcal{Q}^\top \text{diag}(\phi_i) \mathcal{Q} & -\frac{1}{2} \mathcal{T}(\mathbf{x})^\top - \mathcal{Q}^\top \text{diag}(\phi_i) \mathcal{W}^\top & -\lambda \hat{\boldsymbol{\xi}}_i - \frac{1}{2} \mathcal{Q}^\top \boldsymbol{\psi}_i \\ -\frac{1}{2} \mathcal{T}(\mathbf{x}) - \mathcal{W} \text{diag}(\phi_i) \mathcal{Q} & \mathcal{W} \text{diag}(\phi_i) \mathcal{W}^\top + \delta \mathbb{I} & \frac{1}{2} (\mathcal{W} \boldsymbol{\psi}_i - \mathbf{h}(\mathbf{x})) \\ (-\lambda \hat{\boldsymbol{\xi}}_i - \frac{1}{2} \mathcal{Q}^\top \boldsymbol{\psi}_i)^\top & \frac{1}{2} (\mathcal{W} \boldsymbol{\psi}_i - \mathbf{h}(\mathbf{x}))^\top & s_i \end{bmatrix} \succeq_{\mathcal{C}} \mathbf{0} \quad \forall i \in [I]
\end{aligned} \tag{24}$$

Note that δ only affects the middle block of the copositive matrix. One can further show that the completely positive program dual to (24) constitutes a restriction of problem (14) with a perturbed objective function.

$$\begin{aligned}
\underline{\mathcal{Z}}_\delta(\mathbf{x}) = & \sup \quad \frac{1}{I} \sum_{i \in [I]} [\text{tr}(\mathcal{T}(\mathbf{x}) \mathbf{Y}_i) + \mathbf{h}(\mathbf{x})^\top \boldsymbol{\gamma}_i - \delta \text{tr}(\boldsymbol{\Gamma}_i)] \\
& \text{s.t.} \quad \boldsymbol{\gamma}_i \in \mathbb{R}_+^{M+J}, \boldsymbol{\mu}_i \in \mathbb{R}_+^K, \boldsymbol{\Gamma}_i \in \mathbb{S}_+^{M+J}, \boldsymbol{\Omega}_i \in \mathbb{S}_+^K, \mathbf{Y}_i \in \mathbb{R}^{K \times (M+J)} \quad \forall i \in [I] \\
& \quad \mathcal{Q} \boldsymbol{\mu}_i + \mathbf{q} = \mathcal{W}^\top \boldsymbol{\gamma}_i \quad \forall i \in [I] \\
& \quad \mathcal{Q}_j^\top \boldsymbol{\Omega}_i \mathcal{Q}_j - 2 \mathcal{Q}_j^\top \mathbf{Y}_i \mathcal{W}_j + \mathcal{W}_j^\top \boldsymbol{\Gamma}_i \mathcal{W}_j = q_j^2 \quad \forall i \in [I] \forall j \in [N_2+J] \\
& \quad \frac{1}{I} \sum_{i \in [I]} \text{tr}(\boldsymbol{\Omega}_i) - 2 \hat{\boldsymbol{\xi}}_i^\top \boldsymbol{\mu}_i + \hat{\boldsymbol{\xi}}_i^\top \hat{\boldsymbol{\xi}}_i \leq \epsilon^2 \\
& \quad \begin{bmatrix} \boldsymbol{\Omega}_i & \mathbf{Y}_i & \boldsymbol{\mu}_i \\ \mathbf{Y}_i^\top & \boldsymbol{\Gamma}_i & \boldsymbol{\gamma}_i \\ \boldsymbol{\mu}_i^\top & \boldsymbol{\gamma}_i^\top & 1 \end{bmatrix} \succeq_{\mathcal{C}^*} \mathbf{0} \quad \forall i \in [I]
\end{aligned} \tag{25}$$

Observe that δ only affects the objective function of this dual problem.

Proposition 2. *For any fixed $\mathbf{x} \in \mathcal{X}$, $\overline{\mathcal{Z}}_\delta(\mathbf{x}) = \underline{\mathcal{Z}}_\delta(\mathbf{x})$ is finite for all $\delta > 0$, and $\lim_{\delta \downarrow 0} \overline{\mathcal{Z}}_\delta(\mathbf{x}) = \mathcal{Z}(\mathbf{x})$.*

Proof. We first show that $\overline{\mathcal{Z}}_\delta(\mathbf{x}) = \underline{\mathcal{Z}}_\delta(\mathbf{x})$ by proving strong duality between problems (24) and (25). To this end, we first construct a Slater point $(\lambda^s, (s_i^s, \phi_i^s, \boldsymbol{\psi}_i^s)_{i \in [I]})$ for problem (24). Specifically, we set $\boldsymbol{\psi}_i^s = \mathbf{0}$ and $\phi_i^s = 0$ for all $i \in [I]$, and we select λ^s satisfying $\lambda^s \mathbb{I} \succ \frac{1}{4\delta} \mathcal{T}(\mathbf{x}) \mathcal{T}(\mathbf{x})^\top$. This is possible because $\delta > 0$.

A standard Schur complement argument then implies that for all sufficiently large $s_i^s > 0$, $i \in [I]$, we have

$$\begin{bmatrix} \lambda^s \mathbb{I} & -\frac{1}{2} \mathcal{T}(\mathbf{x})^\top \\ -\frac{1}{2} \mathcal{T}(\mathbf{x}) & \delta \mathbb{I} \end{bmatrix} \succ \mathbf{0} \implies \begin{bmatrix} \lambda^s \mathbb{I} & -\frac{1}{2} \mathcal{T}(\mathbf{x})^\top \\ -\frac{1}{2} \mathcal{T}(\mathbf{x}) & \delta \mathbb{I} \end{bmatrix} \succ \frac{1}{s_i^s} \begin{bmatrix} \lambda^s \hat{\boldsymbol{\xi}}_i \\ \frac{1}{2} \mathbf{h}(\mathbf{x}) \end{bmatrix} \begin{bmatrix} \lambda^s \hat{\boldsymbol{\xi}}_i \\ \frac{1}{2} \mathbf{h}(\mathbf{x}) \end{bmatrix}^\top \quad \forall i \in [I].$$

A second Schur complement argument then ensures that

$$\begin{bmatrix} \lambda^s \mathbb{I} & -\frac{1}{2} \mathcal{T}(\mathbf{x})^\top & -\lambda^s \hat{\boldsymbol{\xi}}_i \\ -\frac{1}{2} \mathcal{T}(\mathbf{x}) & \delta \mathbb{I} & -\frac{1}{2} \mathbf{h}(\mathbf{x}) \\ -\lambda^s \hat{\boldsymbol{\xi}}_i^\top & -\frac{1}{2} \mathbf{h}(\mathbf{x})^\top & s_i^s \end{bmatrix} \succ \mathbf{0} \quad \forall i \in [I]. \quad (26)$$

Consequently, the matrix on the left-hand side of (26) is in the interior of the copositive cone \mathcal{C} . This proves that $(\lambda^s, (s_i^s, \phi_i^s, \psi_i^s)_{i \in [I]})$ is indeed a Slater point for (24). Hence, $\bar{\mathcal{Z}}_\delta(\mathbf{x}) = \underline{\mathcal{Z}}_\delta(\mathbf{x})$ by strong conic duality. As problem (24) is feasible, we have $\bar{\mathcal{Z}}_\delta(\mathbf{x}) < +\infty$ for any fixed $\delta > 0$. Moreover, as problem (1) has sufficiently expensive recourse, we have $Z(\mathbf{x}, \boldsymbol{\xi}) > -\infty$ for any fixed $\mathbf{x} \in \mathcal{X}$ and $\boldsymbol{\xi} \in \Xi$. Since $\hat{\mathcal{P}}$ is non-empty, evaluating the worst-case expectation in (2) yields $\underline{\mathcal{Z}}(\mathbf{x}) = \mathcal{Z}(\mathbf{x}) > -\infty$, where the equality follows from Theorem 3. Thus, the completely positive program (14) and its restriction (25) are both feasible, implying that $\bar{\mathcal{Z}}_\delta(\mathbf{x}) = \underline{\mathcal{Z}}_\delta(\mathbf{x}) > -\infty$. This proves finiteness of $\bar{\mathcal{Z}}_\delta(\mathbf{x})$.

To prove the second claim, we observe that $\underline{\mathcal{Z}}_\delta(\mathbf{x})$ constitutes a pointwise supremum of a family of affine functions in δ . Thus, $\underline{\mathcal{Z}}_\delta(\mathbf{x})$ is convex and lower-semicontinuous in δ for every fixed $\mathbf{x} \in \mathcal{X}$. Since $\underline{\mathcal{Z}}_\delta(\mathbf{x})$ is also non-increasing in δ by construction, it is indeed right-continuous. Thus, we have

$$\lim_{\delta \downarrow 0} \bar{\mathcal{Z}}_\delta(\mathbf{x}) = \lim_{\delta \downarrow 0} \underline{\mathcal{Z}}_\delta(\mathbf{x}) = \underline{\mathcal{Z}}_0(\mathbf{x}) = \mathcal{Z}(\mathbf{x}). \quad (27)$$

Here, the first equality holds because $\bar{\mathcal{Z}}_\delta(\mathbf{x}) = \underline{\mathcal{Z}}_\delta(\mathbf{x})$ for $\delta > 0$, while the second equality follows from the right-continuity of $\underline{\mathcal{Z}}_\delta(\mathbf{x})$. The last equality is due to Theorem 3. This completes the proof. \square

The findings of this section culminate in the following main theorem.

Theorem 5. *Consider the following family of copositive programs parametrized in δ .*

$$\begin{aligned} & \text{minimize} && \mathbf{c}^\top \mathbf{x} + \epsilon^2 \lambda + \frac{1}{I} \sum_{i \in [I]} \left[s_i + \mathbf{q}^\top \boldsymbol{\psi}_i - \lambda \|\hat{\boldsymbol{\xi}}_i\|_2^2 + \sum_{j \in [N_2+J]} \phi_{ij} q_j^2 \right] \\ & \text{subject to} && \mathbf{x} \in \mathcal{X}, \lambda \in \mathbb{R}_+, s_i \in \mathbb{R}, \boldsymbol{\psi}_i, \phi_i \in \mathbb{R}^{N_2+J} \quad \forall i \in [I] \\ & && \begin{bmatrix} \lambda \mathbb{I} + \mathcal{Q}^\top \text{diag}(\phi_i) \mathcal{Q} & -\frac{1}{2} \mathcal{T}(\mathbf{x})^\top - \mathcal{Q}^\top \text{diag}(\phi_i) \mathcal{W}^\top & -\lambda \hat{\boldsymbol{\xi}}_i - \frac{1}{2} \mathcal{Q}^\top \boldsymbol{\psi}_i \\ -\frac{1}{2} \mathcal{T}(\mathbf{x}) - \mathcal{W} \text{diag}(\phi_i) \mathcal{Q} & \mathcal{W} \text{diag}(\phi_i) \mathcal{W}^\top + \delta \mathbb{I} & \frac{1}{2} (\mathcal{W} \boldsymbol{\psi}_i - \mathbf{h}(\mathbf{x})) \\ (-\lambda \hat{\boldsymbol{\xi}}_i - \frac{1}{2} \mathcal{Q}^\top \boldsymbol{\psi}_i)^\top & \frac{1}{2} (\mathcal{W} \boldsymbol{\psi}_i - \mathbf{h}(\mathbf{x}))^\top & s_i \end{bmatrix} \succeq_{\mathcal{C}} \mathbf{0} \quad \forall i \in [I] \end{bmatrix} \end{aligned} \quad (28)$$

Then, the following statements hold.

(i) If $\delta = 0$ and (1) has complete recourse, then (28) is equivalent to (1).

(ii) If $\delta = 0$ and (1) fails to have complete recourse, then (28) provides an upper bound on (1).

(iii) If $\delta > 0$, then (28) provides a lower bound on (1).

(iv) If \mathcal{X} is compact, then the optimal value of (28) converges to that of (1) for $\delta \downarrow 0$. Moreover, every cluster point \mathbf{x}^* of a sequence $\{\mathbf{x}_\delta^*\}_{\delta \downarrow 0}$ of minimizers for (28) is a minimizer for (1).

Proof. Replacing $\mathcal{Z}(\mathbf{x})$ in (1) with $\overline{\mathcal{Z}}_\delta(\mathbf{x})$ yields (28). Assertion (i) thus follows from Theorem 4, while assertion (ii) follows from Theorem 2, which implies that $\mathcal{Z}(\mathbf{x}) \leq \overline{\mathcal{Z}}(\mathbf{x}) = \overline{\mathcal{Z}}_0(\mathbf{x})$ for every $\mathbf{x} \in \mathcal{X}$. Assertion (iii) holds because $\overline{\mathcal{Z}}_\delta(\mathbf{x}) = \underline{\mathcal{Z}}_\delta(\mathbf{x}) \leq \underline{\mathcal{Z}}(\mathbf{x}) = \mathcal{Z}(\mathbf{x})$, where the first equality follows from Proposition 2, the inequality holds because problem (25) constitutes a relaxation of (10), and the second equality is due to Theorem 3. As for assertion (iv), recall that $\underline{\mathcal{Z}}_\delta(\mathbf{x})$ is the optimal value of problem (25) and thus constitutes a pointwise supremum of affine functions in \mathbf{x} . Therefore, $\overline{\mathcal{Z}}_\delta(\mathbf{x}) = \underline{\mathcal{Z}}_\delta(\mathbf{x})$ is convex and lower semicontinuous in \mathbf{x} for every fixed $\delta > 0$. As \mathcal{X} is compact, we may thus conclude that there exists a minimizer $\mathbf{x}_\delta^* \in \arg \min_{\mathbf{x} \in \mathcal{X}} \mathbf{c}^\top \mathbf{x} + \overline{\mathcal{Z}}_\delta(\mathbf{x})$ for every $\delta > 0$. Next, note that the sequence of functions $\{\overline{\mathcal{Z}}_\delta(\mathbf{x})\}_{\delta \downarrow 0}$ is non-decreasing and thus epi-converges to $\mathcal{Z}(\mathbf{x})$ by [47, Proposition 7.4(d)]. Moreover, the sequence of minimizers $\{\mathbf{x}_\delta^*\}_{\delta \downarrow 0}$ admits at least one cluster point $\mathbf{x}^* \in \mathcal{X}$. By [53, Proposition 7.30], \mathbf{x}^* constitutes a minimizer for (1), and we have

$$\lim_{\delta \downarrow 0} \min_{\mathbf{x} \in \mathcal{X}} \mathbf{c}^\top \mathbf{x} + \overline{\mathcal{Z}}_\delta(\mathbf{x}) = \min_{\mathbf{x} \in \mathcal{X}} \mathbf{c}^\top \mathbf{x} + \mathcal{Z}(\mathbf{x}).$$

This completes the proof. \square

Theorem 5 immediately extends to two-stage robust optimization problems of the form (7).

Corollary 1 (Two-Stage Robust Optimization). *Assume that Ξ is bounded and set $I = 1$. Moreover, choose $\epsilon \geq 0$ and $\hat{\xi}_1 \in \Xi$ such that $d(\xi, \hat{\xi}_1) \leq \epsilon \forall \xi \in \Xi$. Then, the following statements hold.*

(i) If $\delta = 0$ and (7) has complete recourse, then (28) is equivalent to (7).

(ii) If $\delta = 0$ and (7) fails to have complete recourse, then (28) provides an upper bound on (7).

(iii) If $\delta > 0$, then (28) provides a lower bound on (7).

(iv) If \mathcal{X} is compact, then the optimal value of (28) converges to that of (7) for $\delta \downarrow 0$. Moreover, every cluster point \mathbf{x}^* of a sequence $\{\mathbf{x}_\delta^*\}_{\delta \downarrow 0}$ of minimizers for (28) is a minimizer for (7).

Proof. By construction of ϵ and $\hat{\xi}_1$, the Wasserstein ball $\mathcal{B}_\epsilon^2(\hat{\mathbb{P}}_I)$ contains all Dirac distributions δ_ξ , $\xi \in \Xi$. Therefore, the worst-case expected cost (2) reduces to $\max_{\xi \in \Xi} Z(\mathbf{x}, \xi)$. The claim thus follows immediately from Theorem 5. \square

To our best knowledge, Corollary 1 provides the first exact finite conic programming reformulation for the *generic* two-stage robust optimization problem (7). When the uncertainty appears only in the constraints of the recourse problem (3), approximation schemes based on the cutting plane method are available [57, 64]. These approximations construct increasingly tight lower bounds for the wait-and-see cost $\max_{\boldsymbol{\xi} \in \Xi} Z(\boldsymbol{x}, \boldsymbol{\xi})$. Each iteration is costly, however, as it involves the solution of a bilinear maximization problem. Thus, additional assumptions on the uncertainty set Ξ are needed to alleviate the computational burden of generating the cuts. For example, when Ξ is a finite or a budgeted uncertainty set, then one can reformulate the corresponding bilinear maximization problems as mixed integer linear programs of moderate sizes. If problem (7) fails to have relatively complete recourse, however, then many expensive iterations may be required to obtain a first feasible solution. The results portrayed in Theorem 5 and Corollary 1 motivate an alternative conservative approximation scheme to solve problem (7) that can immediately produce a feasible solution. This is achieved by employing a tractable inner approximation for the copositive cone \mathcal{C} . We discuss this approach in the next section.

3.4 A Hierarchy of Semidefinite Programming Approximations

To justify the use of an approximation scheme, we first establish that two-stage distributionally robust linear programs of the form (1) are intractable.

Proposition 3. *The two-stage distributionally robust linear program (1) is NP-hard even if $\mathbf{Q} = \mathbf{0}$ and \mathcal{X} is a polyhedron specified by a list of linear inequalities.*

Proof. Recall from Remark 1 that the distributionally robust linear programs (1) encapsulate the class of all two-stage robust optimization problems of the form (7). The claim thus follows immediately from [24, Theorem 3.5], which asserts that two-stage robust optimization problems with fixed recourse are NP-hard. \square

The complexity result of Proposition 3 is also plausible in view of Theorem 5 and the known fact that linear programs over copositive cones are generically intractable [40]. A tractable conservative approximation for (28) can be obtained by replacing the copositive cone \mathcal{C} with

$$\mathcal{C}^0 = \{ \mathbf{M} \in \mathbb{S}^K : \mathbf{M} = \mathbf{P} + \mathbf{N}, \mathbf{P} \succeq \mathbf{0}, \mathbf{N} \geq \mathbf{0} \}.$$

By construction, we have $\mathcal{C}^0 \subseteq \mathcal{C}$, but for dimensions $K \leq 4$ one can prove that $\mathcal{C}^0 = \mathcal{C}$ [16]. For $K > 4$, \mathcal{C}^0 is a strict subset of \mathcal{C} . In this case, there exists a hierarchy of semidefinite representable cones $\{\mathcal{C}^\ell\}_{\ell \geq 1}$ that provide increasingly tight inner approximations for \mathcal{C} and converge in finitely many iterations to \mathcal{C} [43, 10, 14, 34]. If these tractable cones are used to replace \mathcal{C} in (28), then the sizes of the resulting approximate problems can, however, become prohibitively large for $\ell > 0$. In practice, we find that replacing the cone \mathcal{C} with \mathcal{C}^0 is sufficient to generate solutions that enjoy an acceptable accuracy.

If $\mathcal{Z}(\mathbf{x})$ denotes the worst-case expected disutility function (29), then the following statements hold.

(i) If $\delta = 0$ and (1) has complete recourse, then (30) is equivalent to (1).

(ii) If $\delta = 0$ and (1) fails to have complete recourse, then (30) provides an upper bound on (1).

(iii) If $\delta > 0$, then (30) provides a lower bound on (1).

(iv) If \mathcal{X} is compact, then the optimal value of (30) converges to that of (1) for $\delta \downarrow 0$. Moreover, every cluster point \mathbf{x}^* of a sequence $\{\mathbf{x}_\delta^*\}_{\delta \downarrow 0}$ of minimizers for (30) is a minimizer for (1).

Proof. This is an immediate generalization of Theorem 5. Details are omitted for brevity. \square

4 Linear Programming Reformulation for $\mathbf{Q} = \mathbf{0}$

We assume now that the uncertainty affects only the constraints of the recourse problem (3), that is, we assume that $\mathbf{Q} = \mathbf{0}$. Unless stated otherwise, we further assume throughout this section that $\Xi = \mathbb{R}^K$, and that the ambiguity set is constructed using the 1-Wasserstein metric with reference distance $d(\xi_1, \xi_2) = \|\xi_1 - \xi_2\|$, where the norm $\|\cdot\|$ is defined through

$$\|\xi\| = \mathbf{e}^\top \max \{w_+ \cdot \xi, -w_- \cdot \xi\} \quad (31)$$

for some positive scaling parameters w_+ and w_- . Note, that (31) reduces to the 1-norm if $w_+ = w_- = 1$.

Finally, we always assume that (1) has sufficiently expensive recourse. Under these assumptions, the two-stage distributionally robust linear program (1) admits an equivalent reformulation as a tractable linear program.

4.1 Tractable Formulation

We first establish the main tractability result for the two-stage distributionally robust linear program (1).

Theorem 6. *The two-stage distributionally robust optimization problem (1) is equivalent to the tractable linear program*

$$\begin{aligned} \text{minimize} \quad & \epsilon\lambda + \frac{1}{I} \sum_{i \in [I]} \mathbf{q}^\top \mathbf{y}_i \\ \text{subject to} \quad & \mathbf{x} \in \mathcal{X}, \lambda \in \mathbb{R}_+, \mathbf{y}_i \in \mathbb{R}^{N_2} && \forall i \in [I] \\ & \phi_k, \psi_k \in \mathbb{R}^{N_2} && \forall k \in [K] \\ & \mathbf{T}(\mathbf{x})\hat{\xi}_i + \mathbf{h}(\mathbf{x}) \leq \mathbf{W}\mathbf{y}_i && \forall i \in [I] \\ & \left. \begin{aligned} \mathbf{q}^\top \phi_k \leq \lambda, \mathbf{q}^\top \psi_k \leq \lambda \\ \mathbf{T}(\mathbf{x})\mathbf{e}_k/w_+ \leq \mathbf{W}\phi_k, -\mathbf{T}(\mathbf{x})\mathbf{e}_k/w_- \leq \mathbf{W}\psi_k \end{aligned} \right\} && \forall k \in [K]. \end{aligned} \quad (32)$$

Proof. By strong linear programming duality, which holds because problem (1) has sufficiently expensive recourse, we have $Z(\mathbf{x}, \boldsymbol{\xi}) = Z_d(\mathbf{x}, \boldsymbol{\xi})$ for every $\mathbf{x} \in \mathcal{X}$ and $\boldsymbol{\xi} \in \mathbb{R}^K$. Theorem 1 thus implies that

$$\mathcal{Z}(\mathbf{x}) = \inf_{\lambda \geq 0} \epsilon \lambda + \frac{1}{N} \sum_{i \in [I]} \sup_{\boldsymbol{\xi}} \sup_{\substack{\mathbf{p} \geq \mathbf{0} \\ \mathbf{W}^\top \mathbf{p} = \mathbf{q}}} (\mathbf{T}(\mathbf{x})\boldsymbol{\xi} + \mathbf{h}(\mathbf{x}))^\top \mathbf{p} - \lambda \|\boldsymbol{\xi} - \hat{\boldsymbol{\xi}}_i\|.$$

Invoking the definition of the dual norm and interchanging the order of the supremum operators over $\boldsymbol{\xi}$ and \mathbf{p} , we further obtain

$$\mathcal{Z}(\mathbf{x}) = \inf_{\lambda \geq 0} \epsilon \lambda + \frac{1}{N} \sum_{i \in [I]} \sup_{\substack{\mathbf{p} \geq \mathbf{0} \\ \mathbf{W}^\top \mathbf{p} = \mathbf{q}}} \sup_{\boldsymbol{\xi}} \inf_{\|\boldsymbol{\gamma}\|_* \leq \lambda} (\mathbf{T}(\mathbf{x})\boldsymbol{\xi} + \mathbf{h}(\mathbf{x}))^\top \mathbf{p} - \boldsymbol{\gamma}^\top \boldsymbol{\xi} + \boldsymbol{\gamma}^\top \hat{\boldsymbol{\xi}}_i.$$

Next, we interchange the order of the innermost supremum over $\boldsymbol{\xi}$ and the infimum over $\boldsymbol{\gamma}$, which is allowed by the classical minimax theorem [6, Proposition 5.5.4] since $\boldsymbol{\gamma}$ ranges over a compact set. This yields

$$\mathcal{Z}(\mathbf{x}) = \inf_{\lambda \geq 0} \epsilon \lambda + \frac{1}{I} \sum_{i \in [I]} \sup_{\substack{\mathbf{p} \geq \mathbf{0} \\ \mathbf{W}^\top \mathbf{p} = \mathbf{q}}} \inf_{\|\boldsymbol{\gamma}\|_* \leq \lambda} \sup_{\boldsymbol{\xi}} (\mathbf{T}(\mathbf{x})\boldsymbol{\xi} + \mathbf{h}(\mathbf{x}))^\top \mathbf{p} - \boldsymbol{\gamma}^\top \boldsymbol{\xi} + \boldsymbol{\gamma}^\top \hat{\boldsymbol{\xi}}_i.$$

Evaluating the inner maximization over $\boldsymbol{\xi}$ analytically further yields

$$\begin{aligned} \mathcal{Z}(\mathbf{x}) &= \inf_{\lambda \geq 0} \epsilon \lambda + \frac{1}{I} \sum_{i \in [I]} \sup_{\substack{\mathbf{p} \geq \mathbf{0} \\ \mathbf{W}^\top \mathbf{p} = \mathbf{q}}} \inf_{\|\boldsymbol{\gamma}\|_* \leq \lambda} \mathbf{h}(\mathbf{x})^\top \mathbf{p} + \boldsymbol{\gamma}^\top \hat{\boldsymbol{\xi}}_i + \chi_{\{\boldsymbol{\gamma} = \mathbf{T}(\mathbf{x})^\top \mathbf{p}\}}(\boldsymbol{\gamma}; \mathbf{p}) \\ &= \inf_{\lambda \geq 0} \epsilon \lambda + \frac{1}{I} \sum_{i \in [I]} \sup_{\substack{\mathbf{p} \geq \mathbf{0} \\ \mathbf{W}^\top \mathbf{p} = \mathbf{q}}} \mathbf{h}(\mathbf{x})^\top \mathbf{p} + (\mathbf{T}(\mathbf{x})^\top \mathbf{p})^\top \hat{\boldsymbol{\xi}}_i + \chi_{\{\|\mathbf{T}(\mathbf{x})^\top \mathbf{p}\|_* \leq \lambda\}}(\mathbf{p}). \end{aligned}$$

The minimization over λ in the last problem can also be evaluated analytically. In fact, the unique optimal solution is $\lambda^* = \sup \{\|\mathbf{T}(\mathbf{x})^\top \mathbf{p}\|_* : \mathbf{p} \in \mathbb{R}_+^M, \mathbf{W}^\top \mathbf{p} = \mathbf{q}\}$. Note that for any $\lambda < \lambda^*$, the supremum over \mathbf{p} would be unbounded, and any $\lambda > \lambda^*$ would incur an unnecessarily high cost as λ is penalized by ϵ in the objective function. We thus obtain

$$\begin{aligned} \mathcal{Z}(\mathbf{x}) &= \inf_{\lambda \in \mathbb{R}_+} \epsilon \lambda + \frac{1}{I} \sum_{i \in [I]} \sup_{\substack{\mathbf{p} \geq \mathbf{0} \\ \mathbf{W}^\top \mathbf{p} = \mathbf{q}}} \mathbf{h}(\mathbf{x})^\top \mathbf{p} + (\mathbf{T}(\mathbf{x})^\top \mathbf{p})^\top \hat{\boldsymbol{\xi}}_i \\ \text{s.t. } &\lambda \in \mathbb{R}_+ \\ &\|\mathbf{T}(\mathbf{x})^\top \mathbf{p}\|_* \leq \lambda \quad \forall \mathbf{p} \in \mathbb{R}_+^M : \mathbf{W}^\top \mathbf{p} = \mathbf{q}. \end{aligned} \tag{33}$$

Next, the norm dual to (31) is given by

$$\|\mathbf{z}\|_* = \max_{k \in [K]} \left[\max \left\{ \frac{z_k}{w_+}, -\frac{z_k}{w_-} \right\} \right].$$

Thus, the last constraint in (33) can be decomposed into a system of $\mathcal{O}(K)$ linear constraints as follows.

$$\begin{aligned} &\|\mathbf{T}(\mathbf{x})^\top \mathbf{p}\|_* \leq \lambda && \forall \mathbf{p} \in \mathbb{R}_+^M : \mathbf{W}^\top \mathbf{p} = \mathbf{q} \\ \iff &\sup_{\substack{\mathbf{p} \geq \mathbf{0} \\ \mathbf{W}^\top \mathbf{p} = \mathbf{q}}} \mathbf{e}_k^\top \mathbf{T}(\mathbf{x})^\top \mathbf{p} / w_+ \leq \lambda, \quad \sup_{\substack{\mathbf{p} \geq \mathbf{0} \\ \mathbf{W}^\top \mathbf{p} = \mathbf{q}}} -\mathbf{e}_k^\top \mathbf{T}(\mathbf{x})^\top \mathbf{p} / w_- \leq \lambda && \forall k \in [K] \\ \iff &\exists \boldsymbol{\phi}_k, \boldsymbol{\psi}_k \in \mathbb{R}^{N_2} : \left. \begin{aligned} &\mathbf{q}^\top \boldsymbol{\phi}_k \leq \lambda, \quad \mathbf{q}^\top \boldsymbol{\psi}_k \leq \lambda \\ &\mathbf{T}(\mathbf{x})\mathbf{e}_k / w_+ \leq \mathbf{W}\boldsymbol{\phi}_k, \quad -\mathbf{T}(\mathbf{x})\mathbf{e}_k / w_- \leq \mathbf{W}\boldsymbol{\psi}_k \end{aligned} \right\} && \forall k \in [K] \end{aligned}$$

Here, the second equivalence follows from dualizing the linear programs over \mathbf{p} , all of which are feasible because problem (1) has sufficiently expensive recourse. The claim then follows from substituting the last constraint system into (33). \square

Example 2 (Regression). *Consider the least absolute deviations (LAD) regression problem*

$$\begin{aligned} & \text{minimize} && \mathbb{E}_{\mathbb{P}} \left[|\mathbf{x}^\top \tilde{\boldsymbol{\xi}} + x_0 - \tilde{\chi}| \right] \\ & \text{subject to} && (\mathbf{x}, x_0) \in \mathcal{X}. \end{aligned}$$

The objective of this problem is to find the slope \mathbf{x} and intercept x_0 of an affine function of the explanatory random variables $\tilde{\boldsymbol{\xi}}$ that tightly approximates the independent variable $\tilde{\chi}$ in terms of the mean absolute deviation. In statistics, however, the data-generating distribution \mathbb{P} of $(\tilde{\boldsymbol{\xi}}, \tilde{\chi})$ is never known. Only the empirical distribution $\hat{\mathbb{P}}_I$ corresponding to a set of I training samples is given. In this case \mathbb{P} is ambiguous, and it may make sense to solve the distributionally robust LAD problem

$$\begin{aligned} & \text{minimize} && \sup_{\mathbb{P} \in \hat{\mathcal{P}}} \mathbb{E}_{\mathbb{P}} \left[|\mathbf{x}^\top \tilde{\boldsymbol{\xi}} + x_0 - \tilde{\chi}| \right] \\ & \text{subject to} && (\mathbf{x}, x_0) \in \mathcal{X}, \end{aligned}$$

which can be identified as an instance of the two-stage distributionally robust optimization problem (1) with recourse function

$$Z((\mathbf{x}, x_0), (\boldsymbol{\xi}, \chi)) = \min\{y : y \in \mathbb{R}, y \geq \mathbf{x}^\top \boldsymbol{\xi} + x_0 - \chi, y \geq \chi - x_0 - \mathbf{x}^\top \boldsymbol{\xi}\}.$$

From equation (33) in the proof of Theorem 6 it is evident that the distributionally robust LAD problem is equivalent to

$$\begin{aligned} & \text{minimize} && \epsilon \|\mathbf{x}\|_* + \frac{1}{N} \sum_{i \in [I]} |\mathbf{x}^\top \hat{\boldsymbol{\xi}}_i + x_0 - \hat{\chi}_i| \\ & \text{subject to} && (\mathbf{x}, x_0) \in \mathcal{X}. \end{aligned}$$

Note that the above formulation holds for arbitrary norms (not just the one defined in (32)). The second term in the objective function represents the empirical LAD loss, while the first term acts as a regularizer for the regression coefficient \mathbf{x} . If the reference distance is set to the infinity norm, then we recover the celebrated LASSO regularizer [58, 60]. On the other hand, if the reference distance is set to the 1-norm, then we obtain an infinity norm regularizer which has recently been employed in the context of logistic regression [49].

Example 3 (Multi-Task Learning). *We can extend Example 2 to a distributionally robust multi-task learning problem [12, 2] where several regression problems are to be solved simultaneously.*

$$\begin{aligned} & \text{minimize} && \sup_{\mathbb{P} \in \hat{\mathcal{P}}} \mathbb{E}_{\mathbb{P}} [\|\mathbf{X} \tilde{\boldsymbol{\xi}} + \mathbf{x} - \tilde{\boldsymbol{\chi}}\|_1] \\ & \text{subject to} && (\mathbf{X}, \mathbf{x}) \in \mathcal{X} \end{aligned}$$

This model has many applications in marketing [35], healthcare [65], natural language processing [13], etc. The distributionally robust multi-task learning model still constitutes an instance of problem (1), where the recourse function is now given by

$$Z((\mathbf{X}, \mathbf{x}), (\boldsymbol{\xi}, \boldsymbol{\chi})) = \min\{\mathbf{e}^\top \mathbf{y} : \mathbf{y} \in \mathbb{R}^L, \mathbf{y} \geq \mathbf{X}\boldsymbol{\xi} + \mathbf{x} - \boldsymbol{\chi}, \mathbf{y} \geq \boldsymbol{\chi} - \mathbf{x} - \mathbf{X}\boldsymbol{\xi}\}.$$

By Theorem 6, this problem is equivalent to the linear program

$$\begin{aligned} & \text{minimize} && \frac{\epsilon}{\min\{w_+, w_-\}} \max_{k \in [K]} \|\mathbf{X}_{:k}\|_1 + \frac{1}{N} \sum_{i \in [I]} \|\mathbf{X}\hat{\boldsymbol{\xi}}_i + \mathbf{x} - \hat{\boldsymbol{\chi}}_i\|_1 \\ & \text{subject to} && (\mathbf{X}, \mathbf{x}) \in \mathcal{X}. \end{aligned}$$

Here, the first term in the objective function acts again as a regularizer for the regression coefficient \mathbf{X} , while the second term represents the empirical LAD loss.

4.2 Complexity Analysis

Unfortunately, tractability of the distributionally robust linear program (1) is lost when the reference distance is defined via a p -norm with $p > 1$ even if all other conditions of Theorem 6 remain valid. This can be shown by using a reduction from the NP-hard MATRIX NORM MAXIMIZATION problem [56].

MATRIX NORM MAXIMIZATION

Instance. Given a positive semidefinite matrix $\mathbf{M} \in \mathbb{S}_+^K$.

Question. For a fixed $q \in [1, \infty)$, compute the matrix norm $\|\mathbf{M}\|_{\infty, q} = \max_{\|\mathbf{z}\|_\infty \leq 1} \|\mathbf{M}\mathbf{z}\|_q$.

Theorem 7. *Computing the optimal value of (1) is NP-hard whenever the reference distance is set to $d(\boldsymbol{\xi}_1, \boldsymbol{\xi}_2) = \|\boldsymbol{\xi}_1 - \boldsymbol{\xi}_2\|_p$ for any $p > 1$, even if $\mathbf{Q} = \mathbf{0}$, $r = 1$, $\Xi = \mathbb{R}^K$ and there are no first-stage decisions.*

Proof. Fix $p \in (1, \infty]$ and set $q = \frac{p}{p-1} \in [1, \infty)$. For any instance $\mathbf{M} \in \mathbb{S}_+^K$ of the MATRIX NORM MAXIMIZATION problem, construct an instance of the distributionally robust linear program (1) as follows. Set the parameters of the recourse problem (3) to $\mathbf{Q} = \mathbf{0}$, $\mathbf{q} = \mathbf{0}$, $\mathbf{T}(\mathbf{x}) = [\mathbf{M} \quad -\mathbf{M}]^\top$, $\mathbf{h}(\mathbf{x}) = \mathbf{0}$ and $\mathbf{W} = [\mathbb{I} \quad \mathbb{I}]^\top$. Moreover, assume that there is only one sample $\hat{\boldsymbol{\xi}}_1 = \mathbf{0}$, and set $\epsilon = 1$. Equation (33) in the proof of Theorem 6 implies that problem (1) is equivalent to

$$\begin{aligned} & \text{minimize} && \lambda + \mathbf{e}^\top \mathbf{y}_1 \\ & \text{subject to} && \lambda \in \mathbb{R}_+, \mathbf{y}_1 \in \mathbb{R}^{N_2} \\ & && \mathbf{0} \leq \mathbf{y}_1 \\ & && \|\mathbf{M}(\mathbf{p}_+ - \mathbf{p}_-)\|_q \leq \lambda \quad \forall \mathbf{p}_+, \mathbf{p}_- \in \mathbb{R}_+^K : \mathbf{p}_+ + \mathbf{p}_- = \mathbf{e}. \end{aligned}$$

Note that $\mathbf{y}_1 = \mathbf{0}$ at optimality irrespective of λ , and thus the optimal value of this problem coincides with

$$\max \left\{ \|\mathbf{M}(\mathbf{p}_+ - \mathbf{p}_-)\|_q : \mathbf{p}_+, \mathbf{p}_- \in \mathbb{R}_+^K, \mathbf{p}_+ + \mathbf{p}_- = \mathbf{e} \right\} = \max_{\|\mathbf{z}\|_\infty \leq 1} \|\mathbf{M}\mathbf{z}\|_q.$$

We conclude that computing the optimal value of (1) is at least as hard as solving the NP-hard MATRIX NORM MAXIMIZATION problem. \square

5 Numerical Results

We now assess the computational and statistical properties of the two-stage distributionally robust linear programs over 2-Wasserstein balls studied in Section 3. All optimization problems are solved with MOSEK v7 using the YALMIP interface [37] on an 8-core 3.4 GHz computer with 16 GB RAM.

5.1 Approximation Quality

We first assess the error introduced by approximating the copositive cone \mathcal{C} in (28) with its semidefinite inner approximation \mathcal{C}^0 . To this end, we study recourse problems of the form

$$Z(\mathbf{x}, \boldsymbol{\xi}) = \inf_{\mathbf{y} \in \mathbb{R}_+^{N_2}} \{ \mathbf{e}^\top \mathbf{y} : \mathbf{A}\boldsymbol{\xi} - \mathbf{b} \leq \mathbf{y} \} = \sum_{n \in [N_2]} \max\{ \mathbf{A}_{n\cdot}^\top \boldsymbol{\xi} - b_n, 0 \} = \max_{\boldsymbol{\ell} \in \{0,1\}^{N_2}} (\mathbf{A}\boldsymbol{\xi} - \mathbf{b})^\top \boldsymbol{\ell}, \quad (34)$$

where $\mathbf{A} \in [0, 1]^{N_2 \times K}$, $\mathbf{b} \in [0, 1]^K$ and the random vector $\tilde{\boldsymbol{\xi}} \in \mathbb{R}^K$ is supported on $\Xi = [0, 1]^K$. Note that the wait-and-see cost is independent of \mathbf{x} and representable as a sum of N_2 max functions, which can be expressed as the pointwise maximum of 2^{N_2} affine functions in $\boldsymbol{\xi}$. Recourse problems of this type are hard both in the stochastic as well as in the robust setting. Indeed, evaluating the expectation of $Z(\mathbf{x}, \tilde{\boldsymbol{\xi}})$ is #P-hard even if $N_2 = 1$ and $\tilde{\boldsymbol{\xi}}$ follows the uniform distribution on Ξ [28, Corollary 1]. Similarly, evaluating the worst case of $Z(\mathbf{x}, \boldsymbol{\xi})$ over all $\boldsymbol{\xi} \in \Xi$ is strongly NP-hard [26, Example 1.1.9]. The following proposition shows that the worst-case expectation of $Z(\mathbf{x}, \tilde{\boldsymbol{\xi}})$ over all distributions of $\tilde{\boldsymbol{\xi}} \in \Xi$ within a given 2-Wasserstein ball can be expressed as the optimal value of a second-order cone program (SOCP) with $\mathcal{O}(2^{N_2})$ constraints.

Proposition 4. *If $\hat{\mathcal{P}} = \mathcal{B}_\epsilon^2(\hat{\mathbb{P}}_I)$, then the worst-case expectation (2) of the wait-and-see cost (34) amounts to*

$$\begin{aligned} \mathcal{Z}(\mathbf{x}) = \inf \quad & \epsilon^2 \lambda + \frac{1}{I} \sum_{i \in [I]} s_i \\ \text{s.t.} \quad & \lambda \in \mathbb{R}_+, s_i \in \mathbb{R}_+, \boldsymbol{\theta}_{i\ell}, \boldsymbol{\eta}_{i\ell} \in \mathbb{R}_+^K & \forall i \in [I] \forall \ell \in \{0, 1\}^{N_2} \\ & s_i + b + \lambda \|\hat{\boldsymbol{\xi}}_i\|_2^2 - \mathbf{e}^\top \boldsymbol{\eta}_{i\ell} \geq 0 & \forall i \in [I] \forall \ell \in \{0, 1\}^{N_2} \\ & \left\| \begin{bmatrix} \mathbf{A}^\top \boldsymbol{\ell} + 2\lambda \hat{\boldsymbol{\xi}}_i + \boldsymbol{\theta}_{i\ell} - \boldsymbol{\eta}_{i\ell} \\ s_i + \mathbf{b}^\top \boldsymbol{\ell} + \lambda \|\hat{\boldsymbol{\xi}}_i\|_2^2 - \mathbf{e}^\top \boldsymbol{\eta}_{i\ell} - \lambda \end{bmatrix} \right\|_2 \leq s_i + \mathbf{b}^\top \boldsymbol{\ell} + \lambda \|\hat{\boldsymbol{\xi}}_i\|_2^2 - \mathbf{e}^\top \boldsymbol{\eta}_{i\ell} + \lambda & \forall i \in [I] \forall \ell \in \{0, 1\}^{N_2}. \end{aligned} \quad (35)$$

Proof. By Theorem 1, the worst-case expectation (2) is representable as

$$\mathcal{Z}(\mathbf{x}) = \inf_{\lambda \in \mathbb{R}_+} \epsilon^2 \lambda + \frac{1}{I} \sum_{i \in [I]} \max_{\boldsymbol{\xi} \in \Xi} \max_{\boldsymbol{\ell} \in \{0,1\}^{N_2}} \boldsymbol{\ell}^\top \mathbf{A}\boldsymbol{\xi} - \mathbf{b}^\top \boldsymbol{\ell} - \|\boldsymbol{\xi} - \hat{\boldsymbol{\xi}}_i\|_2^2.$$

Thus, by introducing auxiliary epigraphical variables s_i , $i \in [I]$, we can reformulate the above optimization problem as the semi-infinite linear program

$$\begin{aligned} \mathcal{Z}(\mathbf{x}) = \quad & \inf \quad \epsilon^2 \lambda + \frac{1}{I} \sum_{i \in [I]} s_i \\ \text{s.t.} \quad & \lambda \in \mathbb{R}_+, s_i \in \mathbb{R}_+ \quad \forall i \in [I] \\ & \max_{\xi \in \Xi} \ell^\top \mathbf{A} \xi - \mathbf{b}^\top \ell - \|\xi - \hat{\xi}_i\|_2^2 \leq s_i \quad \forall i \in [I] \forall \ell \in \{0, 1\}^{N_2}. \end{aligned}$$

Strong quadratic programming duality implies that the (i, ℓ) -th semi-infinite constraint is satisfied if and only if there exist $\theta_{i\ell}, \eta_{i\ell} \in \mathbb{R}_+^K$ that satisfy the hyperbolic constraint

$$\frac{1}{4} \left\| \mathbf{A}^\top \ell + 2\lambda \hat{\xi}_i + \theta_{i\ell} - \eta_{i\ell} \right\|_2^2 \leq \lambda \left(s_i + \mathbf{b}^\top \ell + \lambda \|\hat{\xi}_i\|_2^2 - \mathbf{e}^\top \eta_{i\ell} \right),$$

which is equivalent to the standard SOCP constraints

$$\begin{aligned} \lambda \geq 0, s_i + b + \lambda \|\hat{\xi}_i\|_2^2 - \mathbf{e}^\top \eta_{i\ell} \geq 0, \\ \left\| \begin{bmatrix} \mathbf{A}^\top \ell + 2\lambda \hat{\xi}_i + \theta_{i\ell} - \eta_{i\ell} \\ s_i + \mathbf{b}^\top \ell + \lambda \|\hat{\xi}_i\|_2^2 - \mathbf{e}^\top \eta_{i\ell} - \lambda \end{bmatrix} \right\|_2 \leq s_i + \mathbf{b}^\top \ell + \lambda \|\hat{\xi}_i\|_2^2 - \mathbf{e}^\top \eta_{i\ell} + \lambda. \end{aligned}$$

Thus, the worst-case expectation (2) indeed coincides with the optimal value of (35). \square

As it is hard to evaluate $\mathcal{Z}(\mathbf{x})$ exactly—as reflected by the exponential size of the SOCP (35)—we now investigate two efficient methods for evaluating $\mathcal{Z}(\mathbf{x})$ approximately: the \mathcal{C}_0 approximation of the equivalent copositive program (10) and a state-of-the-art quadratic decision rule approximation. The \mathcal{C}_0 approximation is obtained by solving (10) with inputs $\mathbf{S} = \mathbb{I}$, $\mathbf{t} = \mathbf{e}$, $\mathbf{Q} = \mathbf{0}$, $\mathbf{q} = \mathbf{e}$, $\mathbf{T}(\mathbf{x}) = [\mathbf{A}^\top \mathbf{0}]^\top$, $\mathbf{h}(\mathbf{x}) = [-\mathbf{b}^\top \mathbf{0}^\top]^\top$, and $\mathbf{W} = [\mathbb{I} \mathbb{I}]^\top$, while approximating the copositive cone \mathcal{C} with \mathcal{C}_0 .

In order to develop decision rule approximation, we first use [47, Theorem 14.60] to reformulate (2) as

$$\mathcal{Z}(\mathbf{x}) = \inf_{\mathbf{y} \in \mathcal{L}_{K, N_2}} \left\{ \sup_{\mathbb{P} \in \tilde{\mathcal{P}}} \mathbb{E}_{\mathbb{P}} \left[\mathbf{e}^\top \mathbf{y}(\tilde{\xi}) \right] : \mathbf{A} \xi - \mathbf{b} \leq \mathbf{y}(\xi) \forall \xi \in \Xi, \mathbf{y}(\xi) \geq \mathbf{0} \forall \xi \in \Xi \right\}, \quad (36)$$

where \mathcal{L}_{K, N_2} denotes the linear space of all measurable functions from \mathbb{R}^K to \mathbb{R}^{N_2} . A tractable upper bound on $\mathcal{Z}(\mathbf{x})$ is obtained by restricting \mathcal{L}_{K, N_2} to the subspace of all affine functions; see *e.g.* [20]. A tighter tractable upper bound can be obtained, however, by restricting \mathcal{L}_{K, N_2} to the subspace of all quadratic functions and by conservatively approximating the emerging semi-infinite constraints by semidefinite constraints using the approximate \mathcal{S} -lemma. Quadratic decision rule approximations of this type are also studied in [27].

We run numerical experiments for different values of the uncertainty dimension K and the sample size I and set the Wasserstein radius to $\epsilon = 1/\sqrt{I}$, thus enforcing the scaling rule advocated in [9] and [66].¹ All results are averaged over 100 instances generated randomly as follows. We sample the dimension N_2

¹We also ran all experiments with $\epsilon = 1$ and $\epsilon = 1/I$ but did not observe any qualitative changes in the results.

of the wait-and-see decision uniformly at random from $\{1, 2, \dots, \lceil \log(K + 1) \rceil\}$, which guarantees that the SOCP (35) grows at most polynomially with K and I . Next, we sample \mathbf{A} uniformly from $[0, 1]^{N_2 \times K}$ and \mathbf{b} uniformly from $[0, \mathbf{e}^\top \mathbf{A}_{1:}] \times \dots \times [0, \mathbf{e}^\top \mathbf{A}_{N_2:}]$. We then generate independent training samples $\{\hat{\xi}_i\}_{i \in [I]}$ from the uniform distribution on $[0, 1]^K$. Lastly, we evaluate the worst-case expectation (2) exactly by solving the SOCP (35), and also approximately by computing the \mathcal{C}^0 and the quadratic decision rule approximations.

Table 1 reports the optimality gaps of the two approximations relative to the exact worst-case expectation, averaged across all solvable instances. While the optimality gaps of the \mathcal{C}^0 approximation remain consistently below 2.3%, the state-of-the-art quadratic decision rule approximation can incur alarmingly large optimality gaps of more than 100%. On the other hand, while MOSEK is able to solve all instances of the quadratic decision rule approximation, it encounters numerical difficulties when solving some of the larger instances of the \mathcal{C}^0 approximation. For $K = 64$, for instance, the underlying semidefinite constraints involve blocks of the size 139×139 , which pose a distinct challenge for state-of-the-art interior point solvers. The percentages of all instances of the \mathcal{C}^0 approximation that could be solved to global optimality are reported in Table 2. The average runtimes of both approximations are presented in Table 3.

		K													
		1		2		4		8		16		32		64	
I	5	0.0	0.8	0.0	2.8	0.0	5.2	0.0	3.9	0.5	6.7	0.3	5.5	0.5	3.0
	10	0.0	2.2	0.0	6.1	0.0	11.0	0.0	13.0	0.5	15.0	0.4	16.1	0.8	10.2
	20	0.0	4.9	0.0	10.6	0.0	17.6	0.0	20.0	0.5	27.6	0.5	33.1	1.9	19.0
	40	0.0	8.6	0.0	14.9	0.0	25.4	0.0	28.2	0.8	55.0	0.6	65.2	1.8	54.0
	80	0.0	12.4	0.0	19.9	0.0	33.2	0.0	42.0	0.9	77.1	0.8	117.1	0.0	121.1
	160	0.0	18.2	0.0	26.8	0.0	43.1	0.0	51.2	1.3	129.6	1.3	201.2	-	220.8
	320	0.0	25.3	0.0	34.4	0.0	58.9	0.0	74.3	2.3	237.7	3.7	254.3	-	416.6
	640	0.0	33.4	0.0	43.7	0.0	79.1	0.0	102.3	2.3	343.9	0.2	498.3	-	1137.1

Table 1. Optimality gaps (in %) of the \mathcal{C}^0 approximation (left) and the quadratic decision rule approximation (right).

(I, K)										
(10,64)	(20,64)	(40,64)	(80,64)	(160,32)	(160,64)	(320,32)	(320,64)	(640,32)	(640,64)	
95	48	6	1	88	0	19	0	2	0	

Table 2. Percentage of solvable instances from using the \mathcal{C}^0 approximation. We report only those (I, K) pairs for which fewer than 100% of all instances were solved to optimality.

		K													
		1		2		4		8		16		32		64	
I	5	<0.1	<0.1	<0.1	<0.1	0.1	0.1	0.5	0.3	2.3	0.5	41.5	6.2	1375.9	230.4
	10	<0.1	<0.1	0.2	0.1	0.4	0.1	0.8	0.5	5.2	1.0	93.9	12.1	3091.6	355.1
	20	0.1	0.1	0.5	0.1	0.7	0.2	1.9	0.9	11.6	1.4	194.2	26.6	5799.2	490.2
	40	0.3	0.2	1.6	0.3	1.2	0.3	3.0	0.7	25.0	3.0	398.7	52.2	11276.2	1617.3
	80	0.6	0.3	3.3	0.5	2.2	0.5	8.1	1.1	53.5	6.7	905.9	116.9	19887.3	3134.0
	160	1.3	0.8	3.2	1.1	2.6	0.5	19.4	2.4	109.9	17.7	1956.5	253.7	-	7586.9
	320	0.8	2.0	2.1	0.5	12.7	1.0	41.5	5.2	251.4	40.8	3715.2	415.4	-	16511.3
	640	1.7	0.7	31.3	1.1	15.5	2.1	86.8	15.1	514.7	79.1	9777.1	1441.9	-	22365.2

Table 3. Solution times in seconds of the C^0 approximation (left) and the quadratic decision rule approximation (right).

5.2 Out-of-Sample Performance

Next, we assess the out-of-sample performance of different data-driven policies in the context of a multi-item newsvendor problem, where an inventory planner has to select a vector $\mathbf{x} \in \mathbb{R}_+^K$ of order quantities for K different products at the beginning of a sales period. We assume that the total order quantity $\mathbf{e}^\top \mathbf{x}$ may not exceed a given budget B . The demands of the products can be described by a random vector $\tilde{\boldsymbol{\xi}} \in \mathbb{R}_+^K$ that follows an unknown multivariate distribution \mathbb{P}^* . We also assume that there are no ordering costs but that excess inventory of the k -th product incurs a per-unit holding cost of b_k , while unmet demand incurs a per-unit stock-out cost of s_k . The total cost of an order \mathbf{x} incurred in scenario $\boldsymbol{\xi}$ thus amounts to

$$Z(\mathbf{x}, \boldsymbol{\xi}) = \inf_{\mathbf{y} \in \mathbb{R}^K} \{ \mathbf{e}^\top \mathbf{y} : \text{diag}(\mathbf{b})(\mathbf{x} - \boldsymbol{\xi}) \leq \mathbf{y}, \text{diag}(\mathbf{s})(\boldsymbol{\xi} - \mathbf{x}) \leq \mathbf{y} \},$$

where $\mathbf{b} = (b_1, \dots, b_K)^\top$ and $\mathbf{s} = (s_1, \dots, s_K)^\top$. By construction, this recourse problem has sufficiently expensive recourse as well as complete recourse. We assume that the inventory planner is both risk-averse and ambiguity-averse and thus solves the two-stage distributionally robust linear program

$$\begin{aligned} & \text{minimize} && \sup_{\mathbb{P} \in \hat{\mathcal{P}}} \mathbb{P}\text{-CVaR}_\rho \left[Z(\mathbf{x}, \tilde{\boldsymbol{\xi}}) \right] \\ & \text{subject to} && \mathbf{x} \in \mathbb{R}_+^K \\ & && \mathbf{e}^\top \mathbf{x} \leq B, \end{aligned} \tag{37}$$

which minimizes the worst-case conditional value-at-risk (CVaR) of $Z(\mathbf{x}, \tilde{\boldsymbol{\xi}})$ at level $\rho \in (0, 1]$, where the worst case is taken over all distributions within some ambiguity set $\hat{\mathcal{P}}$. For any fixed $\mathbb{P} \in \hat{\mathcal{P}}$, the \mathbb{P} -CVaR of $Z(\mathbf{x}, \tilde{\boldsymbol{\xi}})$ at level ρ is defined through

$$\mathbb{P}\text{-CVaR}_\rho \left[Z(\mathbf{x}, \tilde{\boldsymbol{\xi}}) \right] = \inf_{\theta \in \mathbb{R}} \theta + \frac{1}{\rho} \mathbb{E}_{\mathbb{P}} \left[\max\{Z(\mathbf{x}, \tilde{\boldsymbol{\xi}}) - \theta, 0\} \right]$$

and can be viewed as the conditional expectation of $Z(\mathbf{x}, \tilde{\boldsymbol{\xi}})$ above its $(1 - \rho)$ -percentile under \mathbb{P} [46]. Note that the worst-case CVaR can be viewed as an instance of the worst-case optimized certainty equivalent (29) corresponding to the disutility function $U(y) = \max\{y, 0\}$.

In the following we review different approaches to construct $\hat{\mathcal{P}}$ from I demand samples $\hat{\boldsymbol{\xi}}_1, \dots, \hat{\boldsymbol{\xi}}_I$, and we show that in each case problem (37) can be reformulated as a conic program. The first possibility is to set $\hat{\mathcal{P}}$ to the 2-Wasserstein ball $\mathcal{B}_\epsilon^2(\hat{\mathbb{P}}_I)$ around the empirical distribution on the demand samples as in Section 3.

Proposition 5. *If $\hat{\mathcal{P}} = \mathcal{B}_\epsilon^2(\hat{\mathbb{P}}_I)$, then (37) is equivalent to the copositive program*

$$\begin{aligned}
& \text{minimize} && \theta + \frac{1}{\rho} \left(\epsilon^2 \lambda + \frac{1}{I} \sum_{i \in [I]} s_i \right) \\
& \text{subject to} && \mathbf{x} \in \mathbb{R}_+^K, \lambda, s_i \in \mathbb{R}_+, \theta \in \mathbb{R}, \boldsymbol{\psi}_i, \boldsymbol{\phi}_i \in \mathbb{R}^{N_2+J} && \forall i \in [I] \\
& && \mathbf{e}^\top \mathbf{x} \leq B && (38) \\
& && \begin{bmatrix} \lambda \mathbb{I} & -\frac{1}{2} \mathbf{T}(\mathbf{x})^\top & -\lambda \hat{\boldsymbol{\xi}}_i \\ -\frac{1}{2} \mathbf{T}(\mathbf{x}) & \mathbf{W} \text{diag}(\boldsymbol{\phi}_i) \mathbf{W}^\top & \frac{1}{2} (\mathbf{W} \boldsymbol{\psi}_i - \mathbf{h}(\mathbf{x})) \\ -\lambda \hat{\boldsymbol{\xi}}_i^\top & \frac{1}{2} (\mathbf{W} \boldsymbol{\psi}_i - \mathbf{h}(\mathbf{x}))^\top & s_i + \theta - \mathbf{e}^\top (\boldsymbol{\psi}_i + \boldsymbol{\phi}_i) + \lambda \|\hat{\boldsymbol{\xi}}_i\|_2^2 \end{bmatrix} \succeq_c \mathbf{0} && \forall i \in [I],
\end{aligned}$$

where $\mathbf{T}(\mathbf{x}) = [-\text{diag}(\mathbf{b}) \text{diag}(\mathbf{s})]^\top$, $\mathbf{h}(\mathbf{x}) = [\mathbf{x}^\top \text{diag}(\mathbf{b}), -\mathbf{x}^\top \text{diag}(\mathbf{s})]^\top$, and $\mathbf{W} = [\mathbb{I} \ \mathbb{I}]^\top$.

Proof. The claim follows immediately from Corollary 2 and the observation that (37) has both sufficiently expensive as well as complete recourse. \square

As a second possibility, we can use the I demand samples to estimate the sample mean $\hat{\boldsymbol{\mu}} = \frac{1}{I} \sum_{i \in [I]} \hat{\boldsymbol{\xi}}_i$ and the sample covariance matrix $\hat{\boldsymbol{\Sigma}} = \frac{1}{I} \sum_{i \in [I]} (\hat{\boldsymbol{\xi}}_i - \hat{\boldsymbol{\mu}})(\hat{\boldsymbol{\xi}}_i - \hat{\boldsymbol{\mu}})^\top$ of $\tilde{\boldsymbol{\xi}}$, which can in turn be used to construct a Chebyshev ambiguity set of the form

$$\mathcal{P}(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}}, \gamma_1, \gamma_2) = \left\{ \mathbb{P} \in \mathcal{M}^2(\mathbb{R}_+^K) : \begin{array}{l} (\mathbb{E}_{\mathbb{P}}[\tilde{\boldsymbol{\xi}}] - \hat{\boldsymbol{\mu}})^\top \hat{\boldsymbol{\Sigma}}^{-1} (\mathbb{E}_{\mathbb{P}}[\tilde{\boldsymbol{\xi}}] - \hat{\boldsymbol{\mu}}) \leq \gamma_1 \\ \mathbb{E}_{\mathbb{P}}[(\tilde{\boldsymbol{\xi}} - \hat{\boldsymbol{\mu}})(\tilde{\boldsymbol{\xi}} - \hat{\boldsymbol{\mu}})^\top] \preceq (1 + \gamma_2) \hat{\boldsymbol{\Sigma}} \end{array} \right\}, \quad (39)$$

where $\gamma_1, \gamma_2 \in \mathbb{R}_+$ represent two confidence parameters. This ambiguity set has been proposed in [15].

Proposition 6. *If $\hat{\mathcal{P}} = \mathcal{P}(\hat{\boldsymbol{\mu}}, \hat{\boldsymbol{\Sigma}}, \gamma_1, \gamma_2)$, then (37) is equivalent to the copositive program*

$$\begin{aligned}
& \text{minimize} && \theta + \frac{1}{\rho} \left[s + \text{tr} \left((\gamma_2 \hat{\boldsymbol{\Sigma}} + \hat{\boldsymbol{\mu}} \hat{\boldsymbol{\mu}}^\top) \mathbf{M} \right) + \hat{\boldsymbol{\mu}}^\top \mathbf{m} + \sqrt{\gamma_1} \|\hat{\boldsymbol{\Sigma}}^{\frac{1}{2}} (\mathbf{m} + 2\mathbf{M} \hat{\boldsymbol{\mu}})\|_2 \right] \\
& \text{subject to} && \mathbf{x} \in \mathbb{R}_+^K, \theta, s \in \mathbb{R}, \mathbf{m}, \boldsymbol{\psi}, \boldsymbol{\phi} \in \mathbb{R}^K, \mathbf{M} \in \mathbb{S}_+^K \\
& && \mathbf{e}^\top \mathbf{x} \leq B \\
& && \begin{bmatrix} \mathbf{M} & -\frac{1}{2} \mathbf{T}(\mathbf{x})^\top & \frac{1}{2} \mathbf{m} \\ -\frac{1}{2} \mathbf{T}(\mathbf{x}) & \mathbf{W} \text{diag}(\boldsymbol{\phi}) \mathbf{W}^\top & \frac{1}{2} (\mathbf{W} \boldsymbol{\psi} - \mathbf{h}(\mathbf{x})) \\ \frac{1}{2} \mathbf{m}^\top & \frac{1}{2} (\mathbf{W} \boldsymbol{\psi} - \mathbf{h}(\mathbf{x}))^\top & s + \theta - \mathbf{e}^\top (\boldsymbol{\psi} + \boldsymbol{\phi}) \end{bmatrix} \succeq_c \mathbf{0}, \quad \begin{bmatrix} \mathbf{M} & \frac{1}{2} \mathbf{m} \\ \frac{1}{2} \mathbf{m}^\top & s \end{bmatrix} \succeq_c \mathbf{0},
\end{aligned} \quad (40)$$

where $\mathbf{T}(\mathbf{x}) = [-\text{diag}(\mathbf{b}) \text{diag}(\mathbf{s})]^\top$, $\mathbf{h}(\mathbf{x}) = [\mathbf{x}^\top \text{diag}(\mathbf{b}), -\mathbf{x}^\top \text{diag}(\mathbf{s})]^\top$, and $\mathbf{W} = [\mathbb{I} \ \mathbb{I}]^\top$.

Proof. As in the proof of Proposition 5, we may transform (37) to a worst-case expected disutility minimization problem by using Sion’s minimax theorem [55]. The worst-case expectation in the resulting objective function can then be re-expressed as a maximization problem over completely positive cones by leveraging ideas from [42, Section 4.4]. Finally, problem (40) is obtained via strong conic duality, which allows us to convert the completely positive maximization problem to an equivalent copositive minimization problem. \square

Distributionally robust multi-item newsvendor problems with Chebyshev ambiguity sets are known to be NP-hard even if $\gamma_1 = \gamma_2 = 0$ and $\tilde{\boldsymbol{\xi}}$ is supported on \mathbb{R}^K , but they admit tractable conservative approximations based on quadratic decision rules [27].

A third possibility is to set $\hat{\mathcal{P}} = \{\hat{\mathbb{P}}_I\}$. This singleton ambiguity set corresponds to a Wasserstein ball around the empirical distribution with radius $\epsilon = 0$. Problem (37) then simply reduces to the corresponding sample average approximation (SAA) problem, which is equivalent to a tractable linear program.

In order to assess the performance of the Wasserstein, Chebyshev and SAA policies obtained from the respective distributionally robust optimization models, we conduct out-of-sample experiments for the K -item newsvendor problem with $K = 3$ and training datasets containing $I = 10, 20, 40, 80, 160, 320, 640$ and $1,028$ independent samples. In all experiments, we replace each copositive cone \mathcal{C} appearing in (38) and (40) with its first inner approximation \mathcal{C}^0 . We fix the vectors of holding and stock-out costs to $\mathbf{b} = \mathbf{e}$ and $\mathbf{s} = 10\mathbf{e}$, respectively, and we set the ordering budget to $B = 30$. We further fix the risk level of the CVaR to $\rho = 10\%$.

The results of all experiments are averaged over 100 random trials generated in the following manner. The true demand distribution \mathbb{P}^* of $\tilde{\boldsymbol{\xi}}$ is assumed to be lognormal, that is, $\tilde{\xi}_k = \exp(\tilde{\chi}_k)$, $k \in [K]$, where the $\tilde{\chi}_k$, $k \in [K]$, represent jointly normally distributed random variables with first- and second-order moments given by $\boldsymbol{\nu} \in \mathbb{R}_+^K$ and $\boldsymbol{\Sigma} \in \mathbb{S}_+^K$, respectively. In each trial, we sample $\boldsymbol{\nu}$ uniformly at random from $[0, 2]^K$ while the matrix $\boldsymbol{\Sigma}$ is generated randomly using the following procedure. We set the vector of standard deviations to $\boldsymbol{\sigma} = 1/4\mathbf{e}$, sample a random correlation matrix $\mathbf{C} \in \mathbb{S}_+^K$ using the MATLAB command ‘gallery(‘randcorr’, 3)’, and set $\boldsymbol{\Sigma} = \text{diag}(\boldsymbol{\sigma})\mathbf{C}\text{diag}(\boldsymbol{\sigma}) + \boldsymbol{\nu}\boldsymbol{\nu}^\top$. Next, we sample I independent training samples $\{\hat{\boldsymbol{\xi}}_i\}_{i \in [I]}$ from \mathbb{P}^* . We then compute the Wasserstein policy $\mathbf{x}_{\text{Wass}}^*$ by solving (38) with \mathcal{C}^0 instead of \mathcal{C} and where the Wasserstein radius ϵ is chosen by 5-fold cross-validation so as to minimize the out-of-sample risk [18, 19]. Similarly, the Chebyshev policy $\mathbf{x}_{\text{Cheb}}^*$ is obtained by solving (40) with \mathcal{C}^0 instead of \mathcal{C} and where the confidence parameters are again determined via 5-fold cross-validation. Finally, we compute the SAA policy $\mathbf{x}_{\text{SAA}}^*$ by solving (37) with $\hat{\mathcal{P}} = \{\hat{\mathbb{P}}_I\}$. The out-of-sample risk $\mathbb{P}^*\text{-CVaR}_\rho[Z(\mathbf{x}^*, \tilde{\boldsymbol{\xi}})]$ of each of the three data-driven strategies $\mathbf{x}_{\text{Wass}}^*$, $\mathbf{x}_{\text{Cheb}}^*$ and $\mathbf{x}_{\text{SAA}}^*$ is then estimated at high accuracy using 20,000 test samples from \mathbb{P}^* .

Figure 1 visualizes the out-of-sample risk of the Wasserstein and Chebyshev policies relative to the SAA policy as a function of the training sample size I . Observe that the Wasserstein policy dominates the SAA policy with high confidence uniformly across all sample sizes. Moreover, for training datasets of size $I \leq 20$,

both the Wasserstein and Chebyshev policies outperform the SAA policy with high confidence by more than 20%. This suggests that the distributionally robust policies are preferable whenever there is significant ambiguity about the true distribution \mathbb{P}^* . While the Wasserstein policy consistently outperforms the SAA policy, the quality of the Chebyshev policy starts to deteriorate for $I \geq 30$.




Figure 1. Improvement of the Wasserstein (left) and Chebyshev (right) policy relative to the SAA policy in terms of out-of-sample CVaR. The solid blue lines represent the mean, and the error bars visualize the 20% and 80% quantiles of the relative improvement, respectively.

Figure 2 depicts the optimality gaps of the three policies with respect to the true optimal policy, which we estimate by solving another SAA problem using 20,000 samples from \mathbb{P}^* . For $I = 10$, we find that the SAA policy is $\sim 70\%$ suboptimal, while both the Chebyshev and Wasserstein policies are only $\sim 25\%$ suboptimal on average. For $I = 20$, the SAA policy remains $\sim 40\%$ suboptimal while the Wasserstein policy exhibits a marginally better suboptimality of $\sim 20\%$ on average. The Chebyshev policy, on the other hand, reaches a steady state already for $I = 10$ with an average suboptimality of about 25%. This suggests that the empirical estimates of the first- and second-order moments are already accurate enough for small training datasets. Unfortunately, as the sample size grows, the Chebyshev policy cannot improve as the first two moments are insufficient to describe the entire shape of the true distribution \mathbb{P}^* . The Wasserstein policy, on the other hand, strikes a good balance between robustness and asymptotic consistency. In particular, we find that it converges quickly to the true optimal policy as the size of the training dataset grows.




Figure 2. Optimality gaps of the Wasserstein (left), Chebyshev (middle), and SAA (right) policies. The solid blue lines represent the mean, and the error bars visualize the 20% and 80% quantiles of the optimality gaps, respectively.

Acknowledgements

This research was supported by the Swiss National Science Foundation grant BSCGIO_157733.

References

- [1] A. Ardestani-Jaafari and E. Delage. Linearized robust counterparts of two-stage robust optimization problems with applications in operations management. *Available on Optimization Online*, 2016.
- [2] J. Baxter. A model of inductive bias learning. *Journal of Artificial Intelligence Research*, 12(3):149–198, 2000.
- [3] A. Ben-Tal, L. El Ghaoui, and A. Nemirovski. *Robust Optimization*. Princeton University Press, 2009.
- [4] A. Ben-Tal, A. Goryashko, E. Guslitzer, and A. Nemirovski. Adjustable robust solutions of uncertain linear programs. *Mathematical Programming A*, 99(2):351–376, 2004.
- [5] 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.
- [6] D. P. Bertsekas. *Convex Optimization Theory*. Athena Scientific, 2009.
- [7] D. Bertsimas, X. V. Doan, K. Natarajan, and C.-P. Teo. Models for minimax stochastic linear optimization problems with risk aversion. *Mathematics of Operations Research*, 35(3):580–602, 2010.
- [8] D. Bertsimas, V. Gupta, and N. Kallus. Robust sample average approximation. *Available on arXiv.org*, 2014.
- [9] J. Blanchet and Y. Kang. Sample out-of-sample inference based on Wasserstein distance. *Available on ArXiv*, 2016.
- [10] I. M. Bomze and E. de Klerk. Solving standard quadratic optimization problems via linear, semidefinite and copositive programming. *Journal of Global Optimization*, 24(2):163–185, 2002.
- [11] S. Burer. On the copositive representation of binary and continuous nonconvex quadratic programs. *Mathematical Programming A*, 120(2):479–495, 2009.
- [12] R. Caruana. Multi-task learning. *Machine Learning*, 28(1):41–75, 1996.
- [13] R. Collobert and J. Weston. A unified architecture for natural language processing: Deep neural networks with multitask learning. In *International Conference on Machine Learning*, pages 160–167, 2008.

- [14] E. de Klerk and D. V. Pasechnik. Approximation of the stability number of a graph via copositive programming. *SIAM Journal on Optimization*, 12(4):875–892, 2002.
- [15] E. Delage and Y. Ye. Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research*, 58(3):595–612, 2010.
- [16] P. H. Diananda. On non-negative forms in real variables some or all of which are non-negative. *Mathematical Proceedings of the Cambridge Philosophical Society*, 58(1):17–25, 1962.
- [17] J. Dupačová (as Žáčková). On minimax solutions of stochastic linear programming problems. *Časopis pro pěstování matematiky*, 91(4):423–430, 1966.
- [18] B. Efron and R. J. Tibshirani. *An Introduction to the Bootstrap*. CRC Press, 1994.
- [19] J. Friedman, T. Hastie, and R. Tibshirani. *The Elements of Statistical Learning*. Springer, 2001.
- [20] R. Gao and A. J. Kleywegt. Distributionally robust stochastic optimization with Wasserstein distance. *Available on Optimization Online*, 2016.
- [21] A. Georghiou, W. Wiesemann, and D. Kuhn. Generalized decision rule approximations for stochastic programming via liftings. *Mathematical Programming A*, 152(1-2):301–338, 2015.
- [22] I. Gilboa and D. Schmeidler. Maxmin expected utility with non-unique prior. *Journal of Mathematical Economics*, 18(2):141–153, 1989.
- [23] J. Goh and M. Sim. Distributionally robust optimization and its tractable approximations. *Operations research*, 58(4-part-1):902–917, 2010.
- [24] E. Guslitser. Uncertainty-Immunized Solutions in Linear Programming. Master’s thesis, Technion – Israel Institute of Technology, 2002.
- [25] M. J. Hadjiyiannis, P. J. Goulart, and D. Kuhn. A scenario approach for estimating the suboptimality of linear decision rules in two-stage robust optimization. In *IEEE Conference on Decision and Control and European Control Conference*, pages 7386–7391, 2011.
- [26] G. A. Hanasusanto. *Decision Making under Uncertainty: Robust and Data-Driven Approaches*. PhD thesis, Imperial College London, 2015.
- [27] G. A. Hanasusanto, D. Kuhn, S. W. Wallace, and S. Zymler. Distributionally robust multi-item news vendor problems with multimodal demand distributions. *Mathematical Programming A*, 152(1-2):1–32, 2015.

- [28] G. A. Hanasusanto, D. Kuhn, and W. Wiesemann. A comment on “Computational complexity of stochastic programming problems”. *Mathematical Programming A*, 159(1):557–569, 2016.
- [29] G. A. Hanasusanto, D. Kuhn, and W. Wiesemann. K -adaptability in two-stage distributionally robust binary programming. *Operations Research Letters*, 44(1):6–11, 2016.
- [30] G. A. Hanasusanto, V. Roitch, D. Kuhn, and W. Wiesemann. Ambiguous joint chance constraints under mean and dispersion information. *Available on Optimization Online*, 2015.
- [31] R. Jiang and Y. Guan. Risk-averse two-stage stochastic program with distributional ambiguity. *Available on Optimization Online*, 2015.
- [32] A. J. Kleywegt, A. Shapiro, and T. Homem de Mello. The sample average approximation method for stochastic discrete optimization. *SIAM Journal on Optimization*, 12(2):479–502, 2002.
- [33] Q. Kong, C.-Y. Lee, C.-P. Teo, and Z. Zheng. Scheduling arrivals to a stochastic service delivery system using copositive cones. *Operations Research*, 61(3):711–726, 2013.
- [34] J. B. Lasserre. Convexity in semialgebraic geometry and polynomial optimization. *SIAM Journal on Optimization*, 19(4):1995–2014, 2009.
- [35] P. J. Lenk, W. S. DeSarbo, P. E. Green, and R. M. Young. Hierarchical Bayes conjoint analysis: Recovery of partworth heterogeneity from reduced experimental designs. *Marketing Science*, 15(2):173–191, 1996.
- [36] X. Li, K. Natarajan, C.-P. Teo, and Z. Zheng. Distributionally robust mixed integer linear programs: Persistency models with applications. *European Journal of Operational Research*, 233(3):459–473, 2014.
- [37] 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.
- [38] D. Love and G. Bayraksan. Phi-divergence constrained ambiguous stochastic programs for data-driven optimization. *Available on Optimization Online*, 2016.
- [39] P. Mohajerin Esfahani and D. Kuhn. Data-driven distributionally robust optimization using the Wasserstein metric: Performance guarantees and tractable reformulations. *Available on Optimization Online*, 2015.
- [40] K. G. Murty and S. N. Kabadi. Some NP-complete problems in quadratic and nonlinear programming. *Mathematical Programming*, 39(2):117–129, 1987.
- [41] K. Natarajan and C.-P. Teo. On reduced semidefinite programs for second order moment bounds with applications. *Mathematical Programming A*, 161(1-2):487–518, 2017.

- [42] K. Natarajan, C.-P. Teo, and Z. Zheng. Mixed 0-1 linear programs under objective uncertainty: A completely positive representation. *Operations Research*, 59(3):713–728, 2011.
- [43] P. A. Parrilo. *Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization*. PhD thesis, California Institute of Technology, 2000.
- [44] G. Pflug and A. Pichler. *Multistage Stochastic Optimization*. Springer, 2014.
- [45] G. Pflug and D. Wozabal. Ambiguity in portfolio selection. *Quantitative Finance*, 7(4):435–442, 2007.
- [46] R. T. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2:21–42, 2000.
- [47] R. T. Rockafellar and R. J.-B. Wets. *Variational Analysis*. Springer Science & Business Media, 2009.
- [48] H. E. Scarf. A min-max solution to an inventory problem. In K. J. Arrow, S. Karlin, and H. E. Scarf, editors, *Studies in Mathematical Theory of Inventory and Production*, pages 201–209. Stanford University Press, Stanford, CA, 1958.
- [49] S. Shafieezadeh Abadeh, P. Mohajerin Esfahani, and D. Kuhn. Distributionally robust logistic regression. In *Advances in Neural Information Processing Systems*, pages 1576–1584, 2015.
- [50] N. Shaked-Monderer, A. Berman, I. M. Bomze, F. Jarre, and W. Schachinger. New results on the cp-rank and related properties of co(mpletely)positive matrices. *Linear and Multilinear Algebra*, 63(2):384–396, 2015.
- [51] A. Shapiro. On duality theory of conic linear problems. In M. A. Goberna and M. A. López, editors, *Semi-Infinite Programming*, pages 135–165. Kluwer Academic Publishers, 2001.
- [52] A. Shapiro. Monte Carlo sampling methods. In A. Ruszczyński and A. Shapiro, editors, *Handbook in Operations Research and Management Science*, volume 10, pages 353–425. Elsevier, 2003.
- [53] A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on Stochastic Programming: Modeling and Theory*. SIAM, 2014.
- [54] A. Shapiro and A. Kleywegt. Minimax analysis of stochastic problems. *Optimization Methods and Software*, 17(3):523–542, 2002.
- [55] M. Sion. On general minimax theorems. *Pacific Journal of Mathematics*, 8(1):171–176, 1958.
- [56] D. Steinberg. Computation of Matrix Norms with Applications to Robust Optimization. Master’s thesis, Technion – Israel Institute of Technology, 2005.

- [57] A. Thiele, T. Terry, and M. Epelman. Robust linear optimization with recourse. *Available on Optimization Online*, 2009.
- [58] R. Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society Series B*, 58(1):267–288, 1996.
- [59] C. Villani. *Optimal Transport: Old and New*. Springer Science & Business Media, 2008.
- [60] L. Wang, M. D. Gordon, and J. Zhu. Regularized least absolute deviations regression and an efficient algorithm for parameter tuning. In *IEEE International Conference on Data Mining*, pages 690–700, 2006.
- [61] W. Wiesemann, D. Kuhn, and M. Sim. Distributionally robust convex optimization. *Operations Research*, 62(6):1358–1376, 2014.
- [62] D. Wozabal. A framework for optimization under ambiguity. *Annals of Operations Research*, 193(1):21–47, 2012.
- [63] G. Xu and S. Burer. A copositive approach for two-stage adjustable robust optimization with uncertain right-hand sides. *Available on arXiv*, 2016.
- [64] B. Zeng and L. Zhao. Solving two-stage robust optimization problems using a column-and-constraint generation method. *Operations Research Letters*, 41(5):457–461, 2013.
- [65] D. Zhang, D. Shen, Alzheimer’s Disease Neuroimaging Initiative, and others. Multi-modal multi-task learning for joint prediction of multiple regression and classification variables in Alzheimer’s disease. *Neuroimage*, 59(2):895–907, 2012.
- [66] C. Zhao and Y. Guan. Data-driven risk-averse stochastic optimization with Wasserstein metric. *Available on Optimization Online*, 2015.

A Appendix: Proof of Theorem 1

Proof of Theorem 1. By the definition of the Wasserstein metric, the worst-case expected wait-and-see cost over the ambiguity set $\mathcal{B}_\epsilon^r(\hat{\mathbb{P}}_I)$ can be expressed as

$$\begin{aligned}
 \mathcal{Z}(\mathbf{x}) &= \sup_{\mathbb{P} \in \mathcal{M}^r(\Xi)} \left\{ \mathbb{E}_{\mathbb{P}} [Z(\mathbf{x}, \tilde{\boldsymbol{\xi}})] : W^r(\mathbb{P}, \hat{\mathbb{P}}_I) \leq \epsilon \right\} \\
 &= \sup_{\mathbb{P} \in \mathcal{M}^r(\Xi)} \int_{\Xi} Z(\mathbf{x}, \boldsymbol{\xi}) \mathbb{P}(d\boldsymbol{\xi}) \\
 &\quad \text{s.t. } \mathbb{P} \in \mathcal{M}^r(\Xi), \Pi \in \mathcal{M}^r(\Xi \times \Xi) \\
 &\quad \int_{\Xi \times \Xi} d(\boldsymbol{\xi}, \boldsymbol{\xi}')^r \Pi(d\boldsymbol{\xi}, d\boldsymbol{\xi}') \leq \epsilon^r \\
 &\quad \Pi \text{ is a joint distribution of } \tilde{\boldsymbol{\xi}} \text{ and } \tilde{\boldsymbol{\xi}}' \text{ with marginals } \mathbb{P} \text{ and } \hat{\mathbb{P}}_I, \text{ respectively.}
 \end{aligned} \tag{41}$$

By the law of total probability we can decompose the transportation plan as $\Pi = \frac{1}{I} \sum_{i \in [I]} \mathbb{P}_i$, where \mathbb{P}_i represents the distribution of $\tilde{\boldsymbol{\xi}}$ conditional on $\tilde{\boldsymbol{\xi}}' = \hat{\boldsymbol{\xi}}_i$. Thus, $\mathcal{Z}(\mathbf{x})$ coincides with the optimal value of the generalized moment problem (5). This establishes the first claim.

The second claim follows from the observation that the semi-infinite linear program (6) is dual to the generalized moment problem (5). Indeed, if $Z(\mathbf{x}, \boldsymbol{\xi})$ is finite for all $\boldsymbol{\xi} \in \Xi$, then strong duality between (5) and (6) holds for all $\epsilon > 0$ due to a straightforward generalization of [51, Proposition 3.4] (see also [30, Lemma 7]). If there exists $\bar{\boldsymbol{\xi}} \in \Xi$ with $Z(\mathbf{x}, \bar{\boldsymbol{\xi}}) = \infty$, on the other hand, then the dual problem (6) is infeasible as all its inner maximization problems are unbounded. In this case, however, the primal problem (5) is unbounded. Indeed, since $W^r(\hat{\mathbb{P}}_I, \hat{\mathbb{P}}_I) = 0$, the continuity of the reference metric d implies that the mixture distribution $\mathbb{P}^* = (1 - \tau)\hat{\mathbb{P}}_I + \tau\delta_{\bar{\boldsymbol{\xi}}}$ is feasible in (41) for some $\tau \in (0, 1]$, which implies that $\mathcal{Z}(\mathbf{x}) \geq \mathbb{E}_{\mathbb{P}^*}[Z(\mathbf{x}, \tilde{\boldsymbol{\xi}})] = \infty$. Thus, the optimal value of the dual problem (6) always coincides with $\mathcal{Z}(\mathbf{x})$. This completes the proof. \square