Les Cahiers du GERAD

The value of randomized solutions in mixed-integer distributionally robust optimization problems

E. Delage,
A. Saif

G–2018–45

June 2018
Revised: April 2019

The value of randomized solutions in mixed-integer distributionally robust optimization problems

Erick Delage [a]

Ahmed Saif [b]

[a] *GERAD & Department of Decision Sciences, HEC Montréal, Montréal (Québec), Canada, H3T 2A7*

[b] *Department of Industrial Engineering, Dalhousie University, Halifax (Nova Scotia), Canada, B3J 1B6*

erick.delage@hec.ca
ahmed.saif@dal.ca

Abstract: *Randomized decision making* refers to the process of taking decisions randomly according to the outcome of an independent randomization device such as a dice or a coin flip. The concept is unconventional, and somehow counterintuitive, in the domain of mathematical programming, where deterministic decisions are usually sought even when the problem parameters are uncertain. However, it has recently been shown that using a *randomized*, rather than a *deterministic*, strategy in non-convex distributionally robust optimization (DRO) problems can lead to improvement in their objective values. It is still unknown, though, what is the magnitude of improvement that can be attained through *randomization* or how to numerically find the optimal *randomized strategy.* In this paper, we study the value of randomization in mixed-integer DRO problems and show that it is bounded by the improvement achievable through its convex relaxation. Furthermore, we identify conditions under which the bound is tight. We then develop an algorithmic procedure, based on column generation, for solving two-stage linear DRO problems with randomization that can be used with both *moment-based* and *Wasserstein* ambiguity sets. Finally, we apply the proposed algorithm to solve three classical discrete DRO problems: the assignment problem, the uncapacitated facility location problem, and the capacitated facility location problem, and report numerical results that show the quality of our bounds, the computational efficiency of the proposed solution method, and the magnitude of performance improvement achieved by randomized decisions.

Keywords: Distributionally robust optimization, randomized strategies, mixed-integer programming

1   Introduction

Distributionally robust optimization (DRO) is a relatively new paradigm in decision making under uncertainty that has attracted considerable attention due to its favorable characteristics (Van Parys et al. 2017). In DRO, we minimize the worst-case expected value of a random cost, *i.e.,* taken with respect to a probability distribution that belongs to a *distributional ambiguity set*. In fact, DRO can be considered both a unifying framework and a viable alternative to two classical approaches for dealing with uncertainty in decision problems: stochastic programming (SP) and robust optimization (RO). Unlike SP, it alleviates the optimistic, and often unrealistic, assumption of the decision maker's complete knowledge of the probability distribution governing the uncertain parameters. Hence, it can prevent the *ex-post* performance disappointment often referred to as the *optimizer's curse* that is common in SP models (Smith and Winkler 2006). Moreover, the DRO counterparts of many decision problems are more computationally tractable than their SP formulations. On the other hand, DRO avoids the inherent over-conservatism of RO that usually leads to poor expected performances, and allows for a better utilization of the available data. In this sense, it can be considered a *data-driven* approach. With a careful design of the ambiguity set, one can usually obtain a statistical guarantee on the *out-of-sample* performance of the DRO problem's solution.

Recently, Delage et al. (2018) introduced the idea of exploiting *randomized strategies* in DRO problems that arise when using ambiguity averse risk measures, *e.g.,* worst-case expected value, worst-case conditional value-at-risk, etc. A randomized strategy describes the process of implementing an action that depends on the outcome of an independent randomization device, such as a dice-roll or a coin-flip. The concept is somewhat counterintuitive (and at first sight computationally unattractive) in the domain of mathematical programming, where the optimal decisions sought are usually deterministic ones even when the problem parameters are uncertain. In particular, Delage et al. (2018, Theorem 13) showed that when the feasible set of a DRO problem is nonconvex, deterministic decisions can be sub-optimal. More precisely, there might exist a randomized strategy that exposes the decision maker to a strictly lower risk (measured using an ambiguity averse risk measure) than what can be achieved by a deterministic one. Despite its significance, this result is still more theoretical than practical given that it is still unclear how much improvement can be obtained in real application problems and whether optimal randomized strategies can be found efficiently.

In this paper, we focus on studying the value of randomized solutions in DRO problems with a mixed-integer linear representable decision space. The contribution is three-fold.

- On the theory side, we prove for the first time that the value of randomization in mixed-integer DRO problems with convex cost functions and convex risk measure is bounded by the difference between the optimal values of the nominal DRO problem and that of its convex relaxation, which is typically straightforward to compute. Furthermore, we show that when the risk measure is an expected value and the cost function is affine with respect to the decisions, this bound becomes tight and can be used to design an efficient solution scheme. Finally, we demonstrate, for the first time, how a finitely supported optimal randomized strategy always exists for this class of problems.

- On the algorithmic side, we devise a finitely convergent column generation algorithm for solving single- and two-stage linear integer DRO problems with randomization. The algorithm iterates between solving a restricted primal problem to generate new vertices in the lifted support set and solving a restricted dual problem to generate new feasible deterministic solutions. Unlike the scheme proposed in Zeng and Zhao (2013), our formulation of the primal subproblem ensures that the number of integer variables does not depend on the size of the support set of the randomized strategy. We also show how the algorithm can be extended to the general mixed-integer case through projection while preserving its finite convergence property, and discuss the cases in which the linearity assumptions do not hold, showing the generic nature of the proposed approach. Despite the theoretical complexity of the problem, the solution algorithm shows surprisingly good performance relative to the deterministic strategy case.

- We provide some empirical evidence that randomization can indeed significantly improve the performance of decisions. This is done using synthetic, yet realistic, instances of three popular stochastic integer programming problems: an assignment problem, and both an uncapacitated and capacitated facility location problem. For some of the assignment problem test instances, a relative improvement of up to 47% in the worst-case expected cost was achievable by a randomized strategy compared to the best deterministic one. In comparison, the improvement achieved in facility location problems appears to be more modest.

The rest of this paper is organized as follows. In the next section, we review the literature that is related to our work. We then motivate our work by solving a small example in Section 3 to illustrate how risk in mixed-integer DRO problems can be reduced through randomization. In Section 4, we study the relationship between randomization and convex relaxation and the structure of optimal randomized strategies. Section 5 encompasses the algorithmic part of the paper. We devise a two-layer column generation algorithm for solving two-stage distributionally robust mixed-integer linear programming problems and explain how it can be modified to solve single-stage problems. This algorithm is used to solve three classical discrete problems in Section 6: the capacitated and uncapacitated facility location problem, and the assignment problem. Section 7 presents numerical results for the aforementioned problems that demonstrate the value of randomized solutions and the performance of the proposed solution algorithms. Conclusions are drawn and directions for future research are proposed in Section 8. Finally, we note that all the proofs of our theorems are deferred to Appendix A.

Notation. We use lower case letters for scalars and vectors and upper case letters for matrices. However, depending on the context, upper case letters are also used to denote random variables (*e.g.,* X) or randomized strategies/distributions (*e.g.,* F_x). Special matrices and vectors used include I, identity matrix of appropriate size, e, *all-ones* vector of appropriate size and e_i, vector of all zeros except for 1 at position i. The operation $[x_1^\intercal \ x_2^\intercal]^\intercal$ is used to denote the concatenation of the vectors x_1 and x_2. We denote by $\mathbb{R}_+$ and $\mathbb{Z}$ the non-negative reals and the integers, respectively. $F_1 \times F_2$ refers to a distribution on the product space such that for $(\xi_1, \xi_2) \sim F_1 \times F_2$, ξ_1 is independent of ξ_2 and both have marginal distributions F_1 and F_2, respectively. We use $\mathcal{C}(\mathcal{X})$ to denote the convex hull of a set $\mathcal{X}$, and $\Delta(\mathcal{X})$ as the set of all probability measures on the measurable space $(\mathcal{X}, \mathcal{B}_\mathcal{X})$, with $\mathcal{B}_\mathcal{X}$ as the Borel σ-algebra over $\mathcal{X}$.

2   Related work

The idea of using randomization in DRO problems is related to the concept of a *mixed strategy* in two-person zero-sum games (von Neumann 1928), where players choose and communicate probability distributions over their respective set of actions. In both fields, a decision maker is considered to solve a minimax problem over a set of distributions. It is important, though, to note two differences between the use of randomization in DRO compared to game theory. First, the notion of an "adversary" is not explicit in DRO problems but rather follows from axiomatic assumptions that are made about how the decision maker perceives risk in an ambiguous environment. Second, randomization in DRO raises significant computational challenges given that such models can employ risk measures that are non-linear with respect to the distribution function and optimize over highly structured distribution sets defined on a continuous parameters space (*e.g.,* the Wasserstein-based ambiguity sets proposed by Mohajerin Esfahani and Kuhn (2017)). Comparatively, zero-sum games usually treat risk aversion using expected utility, which is linear with respect to the distribution functions, and consider discrete action spaces and a simple probability simplex for the distribution sets.

Our study of randomized strategies in DRO problems is also related to some recent unpublished and independent work of Bertsimas et al. (2018), who studied RO in combinatorial optimization problems against an *adaptive online adversary*, which acts after the decision maker but can only exploit information about the decision maker's randomized strategy as opposed to the exact action that is implemented (see Ben-David et al. (1990) for an application of this concept in online optimization). Similarly as shown in Delage et al. (2018) for this kind of problems, an *ambiguity averse risk neutral*

decision maker can strictly benefit from using randomized instead of a *pure* strategy. They also show that the value of randomization can be computed in polynomial time if the cost function is linear and the nominal problem is tractable. They, however, leave open the question of identifying an optimal randomized strategy. This work significantly extends these results to the case where a general risk measure, cost function, and ambiguity set are used, and propose numerical schemes for determining optimal randomized strategies.

Another closely related work is that of Mastin et al. (2015). These authors studied a randomized version of a regret minimization problem, where the optimizing player selects a probability distribution (corresponding to a mixed strategy) over solutions and the adversary selects a cost function with knowledge of the player's distribution, but not of its realization. They studied two special cases of uncertainty, namely uncertainty representable through discrete scenarios and interval (*i.e.,* box) uncertainty. For these two cases, they showed that if the nominal problem is polynomially solvable, then the randomized regret minimization problem can also be obtained in polynomial time. However, they do not address more general convex uncertainty sets arguing that the problem becomes NP-hard for these cases. They also provide uniform bounds for the value of randomization for the two cases of interest. Our work, in contrast, addresses more general uncertainty models, *i.e.,* moment-based and "Wasserstein" distributional ambiguity sets, with convex risk measures (instead of regret) as the objective. This implies that both the decision maker and the adversary can employ randomized strategies. We devise exact solution algorithms applicable for single-stage and two-stage decision problems with a mixed-integer (instead of purely combinatorial) action space. Finally, the numerical bounds that we described can be computed for any ambiguity-averse risk measure and convex support set.

In this work, we extensively use column-generation algorithms to solve problems with large discrete feasible sets efficiently. Column- and/or constraint-generation algorithms have been utilized frequently for solving robust and distributionally-robust optimization problems. Atamtürk and Zhang (2007) used a cutting-plane algorithm for solving a two-stage network flow and design problem, in which separation problems are solved iteratively to eliminate infeasibility and to tighten the bound. A Benders decomposition (*i.e.,* delayed constraint-generation) algorithm was proposed by Thiele et al. (2010) to solve robust linear optimization problems with recourse. Similar Benders-type constraint generation algorithms were used, for example in Brown et al. (2009), Agra et al. (2018), and in Ardestani-Jaafari and Delage (2018). Zhao and Guan (2018) utilized this Benders-based approach to solve a two-stage DRO problem with a Wasserstein ambiguity set, similar to the *deterministic strategy problem* presented in Section 4. Recently, Luo and Mehrotra (2017) proposed a decomposition approach to solve DRO problems with a Wasserstein ambiguity set. They proposed an exchange method to solve the formulated problem for the general nonlinear model to ε-optimality and a central cutting-surface algorithm to solve the special case when the function is convex with respect to the decision variables. Another approach for solving two-stage robust optimization problems is the column-and-constraint generation method proposed by Zeng and Zhao (2013). They showed that it computationally outperforms Benders-based approaches. Chan et al. (2018) used a similar row-and-column generation approach to solve a robust defibrillator deployment problem.

Both the Benders-based constraint generation and the column-and-constraint generation algorithms are well-suited for deterministic strategy problems, for which the objective is to find a pure strategy. Since we are dealing with a randomized strategy problem that aims to find a probability distributions over multiple solutions, we devise a new two-layer column-generation algorithm that iterates between a primal perspective to generate feasible adversary actions and a dual perspective to generate feasible actions for the decision maker. We note that our algorithm is similar in spirit to the *double oracle* method proposed by McMahan et al. (2003) for large-scale zero-sum matrix games, which has found applications particularly in security games (Jain et al. 2011, Yang et al. 2018) and Natural Language Processing (Wang et al. 2017). However, as mentioned earlier, the algorithms that we present in Section 5 address more general and complicated problems than the two-person zero-sum matrix games found in the literate. Namely, in our model both "players"' action spaces can be continuous and the set of feasible mixture strategies is richly configured thus making the application of this type of algorithm far from trivial.

3   Illustrative example

Consider the following 2-node Distributionally Robust Uncapacitated Facility Location Problem (DRU-FLP):

$$\underset{x,y}{\text{minimize}} \quad \max_{F_\xi \in \mathcal{D}} \mathbb{E}_{(\xi_1,\xi_2)\sim F_\xi}[f(x_1 + x_2) + c(\xi_1 y_{12} + \xi_2 y_{21})]$$

$$\text{subject to} \quad \sum_{j\in\{1,2\}} y_{ij} = 1 \qquad\qquad\qquad \forall i \in \{1,2\}$$

$$y_{ij} \leq x_j \qquad\qquad\qquad \forall i \in \{1,2\}, j \in \{1,2\}$$

$$x_j, y_{ij} \in \{0,1\} \qquad\qquad\qquad \forall i \in \{1,2\}, j \in \{1,2\},$$

where each x_j denotes the decision to open a facility at location j, each y_{ij} denotes the decision to assign all the demand at location i to the facility at location j, and where $\xi_i \in \mathbb{R}_+$ is the random demand realized at location i and (ξ_1, ξ_2) are jointly distributed according to F_ξ. Moreover, the coefficients $f \in \mathbb{R}_+$ and $c \in \mathbb{R}_+$, respectively, denote the facility setup cost and unit transportation cost. We also let the ambiguity set take the form:

$$\mathcal{D} := \left\{ F_\xi : \mathbb{P}_{(\xi_1,\xi_2)\sim F_\xi}((\xi_1,\xi_2) \in \mathcal{U}) = 1 \right\},$$

with

$$\mathcal{U} := \left\{ (\xi_1,\xi_2) \in \mathbb{R} \times \mathbb{R} \mid \xi_1 \in [0,\bar{d}], \ \xi_2 \in [0,\bar{d}], \ \xi_1 + \xi_2 \leq \bar{d} \right\},$$

which simply captures the fact that the only information available about the random vector $[\xi_1 \ \xi_2]^\mathsf{T}$ is that the sum of any subset of its terms cannot be strictly larger than $\bar{d}$.

When $f > c\bar{d}$, *i.e.*, the setup costs are larger than the worst-case transportation costs, one can easily demonstrate that opening a single facility at either location 1 or 2, in order to serve the entire demand, is optimal and reaches a worst-case expected total cost of $f + c\bar{d}$. In particular, for all feasible (x, y) pair:

$$\max_{F_\xi \in \mathcal{D}} \ \mathbb{E}_{(\xi_1,\xi_2)\sim F_\xi}[f(x_1 + x_2) + c(\xi_1 y_{12} + \xi_2 y_{21})] \ = \ \max_{\xi\in\mathcal{U}} f(x_1 + x_2) + c(\xi_1 y_{12} + \xi_2 y_{21})$$

$$\geq \max_{(\xi_1,\xi_2)\in\{(\bar{d},0),\,(0,\bar{d}),\,(0,0)\}} f(x_1 + x_2) + c(\xi_1 y_{12} + \xi_2 y_{21})$$

$$= \ f(x_1 + x_2) + \max\{c\bar{d}y_{12} \,,\, c\bar{d}y_{21}\}$$

$$\geq \ f + c\bar{d} = \max_{F_\xi \in \mathcal{D}} \mathbb{E}_{(\xi_1,\xi_2)\sim F_\xi}[f(1 + 0) + c(\xi_1 \cdot 0 + \xi_2 \cdot 1)].$$

One can, however, verify that the following randomized strategy reduces the worst-case expected cost to $f + c\bar{d}/2$:

$$(X_1, X_2, Y_{11}, Y_{12}, Y_{21}, Y_{22}) \sim F_{x,y} := \begin{cases} (1,0,1,0,1,0) & \text{with probability } 50\% \\ (0,1,0,1,0,1) & \text{with probability } 50\% \end{cases}.$$

Indeed, for this strategy, which randomly chooses between the two locations which one will serve all the demand, we have that

$$\underset{F_\xi \in \mathcal{D}}{\text{maximize}} \ \ \mathbb{E}_{(X,Y)\sim F_{x,y},(\xi_1,\xi_2)\sim F_\xi}[f(X_1 + X_2) + c(\xi_1 Y_{12} + \xi_2 Y_{21})] = \max_{\xi\in\mathcal{U}} f + \frac{1}{2}c\xi_1 + \frac{1}{2}c\xi_2$$

$$= \max_{d\in\mathcal{U}} f + \frac{1}{2}c(\xi_1 + \xi_2) = f + c\bar{d}/2.$$

With this randomized strategy, the maximum reduction in worst-case expected cost is realized when $f = cd$, in which case the reduction amounts to 25%. While such a reduction in worst-case expected cost might already make a randomized strategy appear attractive, the question remains of whether a larger reduction can be achieved for this problem. More generally, one might ask how large is this reduction for any given instance of DRUFLP and whether an optimal randomized policies can be identified for large scale versions of this problem. These will be the questions addressed in the rest of this paper.

4   The value of randomized solutions

Consider the mixed-integer distributionally robust optimization problem, which we also refer to as the *deterministic strategy problem*,

$$[\texttt{DSP}]:\quad \underset{x\in\mathcal{X}}{\text{minimize}}\ \sup_{F_\xi\in\mathcal{D}}\ \rho_{\xi\sim F_\xi}\left[h(x,\xi)\right], \tag{1}$$

where $\mathcal{X}\subset\mathbb{Z}^{n_1}\times\mathbb{R}^{n_2}$ is a compact set and $\xi\in\mathbb{R}^m$ is a random vector on the measure having a multivariate distribution function F_ξ that belongs to the distributional set $\mathcal{D}$ containing distributions supported on some $\Xi\subseteq\mathbb{R}^m$. Finally, $h(x,\xi)$ is a cost function and $\rho_{\xi\sim F_\xi}[h(x,\xi)]$ refers to a law-invariant convex risk measure on the probability space $(\Xi,\mathcal{B}_\Xi,F_\xi)$, with $\mathcal{B}_\Xi$ as the Borel σ-algebra over Ξ.

In DSP, the decision maker selects a single action (*i.e.,* a deterministic strategy) $x^*\in\mathcal{X}$ aiming to minimize the worst-case risk associated to a random cost $h(x,\xi)$. For example, when $\rho_{\xi\sim F_\xi}[h(x,\xi)]=\mathbb{E}_{\xi\sim F_\xi}[h(x,\xi)]$, we will say that the decision maker has an *ambiguity averse risk neutral* (AARN) attitude. Intuitively, such an attitude can be interpreted as the attitude of a player trying to achieve the lowest expected cost when playing a game against nature (the adversary) who chooses the distribution F_ξ from $\mathcal{D}$. More generally and reasonably speaking, as shown by Delage et al. (2018), the decision model referred as DSP emerges in any context where the decision maker is considered ambiguity averse (satisfies the axioms of ambiguity aversion and ambiguity monotonicity) and considered to agree with the monotonicity, convexity, and translation invariance axioms of convex risk measure (Föllmer and Schied 2002).

An important result in Delage et al. (2018) consists in establishing that whenever the risk measure $\rho(\cdot)$ satisfies the Lebesgue property, an ambiguity averse decision maker might benefit from employing a randomized strategy instead of a deterministic action. Namely, the decision maker's overall risk might be reduced by solving the *randomized strategy problem*:

$$[\texttt{RSP}]\quad \underset{F_x\in\Delta(\mathcal{X})}{\text{minimize}}\ \sup_{F_\xi\in\mathcal{D}}\ \rho_{(X,\xi)\sim F_x\times F_\xi}\left[h(X,\xi)\right], \tag{2}$$

where $\Delta(\mathcal{X})$ is the set of all probability measures on the measurable space $(\mathcal{X},\mathcal{B}_\mathcal{X})$, with $\mathcal{B}_\mathcal{X}$ as the Borel σ-algebra over $\mathcal{X}$. Moreover, (X,ξ) should be considered as a pair of independent random vectors with marginal probability measures characterized by F_x and F_ξ, respectively.

Definition 1 *Let v_d and v_r refer to the optimal value of problems* (1) *and* (2), *respectively. We define the* value of randomized solutions *as the difference between v_d and v_r,*

$$VRS := v_d - v_r\,.$$

Conceptually, the VRS (and bounds on this value) serves a similar purpose to what is known as the *value of stochastic solutions* for a stochastic program. Namely, It allows one to judge whether it is worth investing a significant amount of additional computational efforts in the resolution of problem (2). Yet, VRS might additionally be used to quantify whether the additional implementation difficulties (both operational and psychological) associated to randomized strategies are worth the investment.

While we will later provide an algorithmic procedure to solve the RSP, or at least bound the VRS, we start here with a tractably more attractive way of bounding this quantity.

Theorem 1 *Given that ρ is a convex risk measure and $h(x,\xi)$ a convex function with respect to x for all $\xi\in\Xi$. Let $\mathcal{X}'$ be any closed set known to contain the convex hull of $\mathcal{X}$, then*

$$VRS \le \widehat{VRS} := v_d - \min_{x\in\mathcal{X}'}\ \sup_{F_\xi\in\mathcal{D}}\ \rho_{\xi\sim F_\xi}[h(x,\xi)]\,. \tag{3}$$

Moreover, if the following conditions are satisfied:

1. *the decision maker has an AARN attitude, i.e., $\rho(\cdot) = \mathbb{E}[\cdot]$,*
2. *the function $h(\cdot, \xi)$ is affine for all $\xi \in \Xi$,*
3. *the set $\mathcal{X}'$ is the convex hull of $\mathcal{X}$,*

then this bound is tight and v_r is achieved by any strategy $F_x^* \in \Delta(\mathcal{X})$ such that $\mathbb{E}_{F_x^*}[X] \in \arg\min_{x \in \mathcal{X}'} \sup_{F_\xi \in \mathcal{D}} \mathbb{E}_{\xi \sim F_\xi}[h(x, \xi)]$, hence F_x^* is optimal in RSP.

Theorem 1 provides a mean of bounding the value of randomization using any convex relaxation $\mathcal{X}'$ of $\mathcal{X}$. In particular, in many application of DRO, the worst-case risk measure $\sup_{F_\xi \in \mathcal{D}} \rho_{\xi \sim F_\xi}[h(x, \xi)]$ is known to be conic representable (see Bertsimas et al. (2017)); Hence, when it is also the case for $\mathcal{X}'$, evaluating VRS can be numerically as difficult as solving the deterministic strategy problem. Theorem 1 also states that when the risk measure is an expected value and the cost function is affine in x, determining the optimal value of the randomized strategy problem reduces to solving the deterministic strategy problem over the convex hull of $\mathcal{X}$, and that the unresolved optimal solution x^* of this new DSP provides the expected decision vector under some optimal randomized strategy F_x^* for the RSP. Using this result, an optimal randomized strategy should therefore be found by solving:

$$\underset{F_x \in \Delta(\mathcal{X})}{\text{minimize}} \|x^* - \mathbb{E}_{X \sim F_x}[X]\|_1. \tag{4}$$

Since, by definition, we have that $x^* \in \mathcal{C}(\mathcal{X})$, problem (4) is necessarily feasible and has an optimal value of 0. Another interesting property of problem (4) and the RSP is presented in the following proposition.

Proposition 1 *If the decision maker's attitude is AARN, i.e., $\rho(\cdot) = \mathbb{E}[\cdot]$, and the function $h(\cdot, \xi)$ is affine for all $\xi \in \Xi$, then there necessarily exists a discrete distribution F_x^*, supported on at most $n + 1$ point, that achieves optimality in problems (2) and (4).*

In Section 6, we show how Proposition 1 can be applied to find the optimal randomized strategy for a stochastic assignment problem under distributional ambiguity. Note that, in general, the solution of (4), and more generally problem (2), is not unique. In other words, multiple randomized strategies might achieve the same optimal value v_r. These randomized strategies are supported on different subsets of $\mathcal{X}$ of arbitrarily large sizes, potentially even infinite. From an algorithmic point of view, the hope is to quickly identify the $n + 1$ support points that are needed to characterize an optimal F_x^*.

In the context of applications that involve cost functions $h(x, \xi)$ that are convex in x, while it is unclear whether a result similar to Proposition 1 still holds, we can nevertheless guarantee that there always exists an optimal randomized strategy that takes the form of a discrete distribution with finite support. This result will be used in Section 5 to design exact solution schemes.

Proposition 2 *If the decision maker's attitude is AARN, i.e., $\rho(\cdot) = \mathbb{E}[\cdot]$, and the function $h(\cdot, \xi)$ is convex for all $\xi \in \Xi$, then there necessarily exists a discrete distribution F_x^*, supported on a finite number of points, that achieves optimality in problem (2). Moreover, F_x^* can be parameterized using $\{(\bar{x}_1^k, x_2^k, p_k)\}_{k \in \mathcal{K}} \subset \mathbb{Z}^{n_1} \times \mathbb{R}^{n_2} \times \mathbb{R}$, such that $\mathbb{P}_{F_x^*}(x = [\bar{x}_1^{k\intercal} \ x_2^{k\intercal}]^\intercal) = p_k$ where*

$$\{\bar{x}_1^k\}_{k \in \mathcal{K}} = \{x_1 \in \mathbb{Z}^{n_1} \mid \exists x_2 \in \mathbb{R}^{n_2}, [x_1^\intercal \ x_2^\intercal]^\intercal \in \mathcal{X}\}$$

is the set of feasible joint assignment for the integer variables and $\mathcal{K}$ is its set of indexes.

Remark 1 *In a private communication we received (Bertsimas et al. 2018), the authors study the VRS for the case where the DSP reduces to a robust linear programming problem, i.e., $\rho(\cdot) = \mathbb{E}[\cdot]$, $h(x, \xi) := \xi^T x$, and $\mathcal{D} := \{F_\xi \mid \mathbb{P}_F(\xi \in \Xi) = 1\}$. Under these conditions, they establish that the bound presented in Theorem 1 is tight and can be obtained in polynomial time if the DSP can be solved in polynomial time. Furthermore, using an argument that is based on Carathéodory's theorem (similar to our proof of Proposition 1), they prove that the ratio v_d/v_r is always bounded by $n + 1$, which can*

be tightened when $\Xi \subset \mathbb{R}_+^m$ and "nearly symmetric". They however do not present any method for identifying optimal randomized strategies. In comparison, this work studies a DSP *model that is more general with respect to the risk attitude, the structure of the cost function and of the ambiguity set. Given that in this case Equation* (3) *does not always provide a tight bound, we will focus next on developing numerical procedures that tighten this gap and as a side product identify optimal (or nearly optimal) randomized strategies.*

5   Exact algorithms for two-stage AARN problems

In this section, we propose an algorithmic procedure based on column generation to find the optimal randomized strategy in a class of discrete two-stage linear DRO problems described as follows:

$$\underset{x \in \mathcal{X}}{\text{minimize}} \ \sup_{F_\xi \in \mathcal{D}} \ c_1^\mathsf{T} x + \mathbb{E}_{\xi \sim F_\xi}[h(x, \xi)], \tag{5}$$

where $\mathcal{X} := \bar{\mathcal{X}} \cap \mathbb{Z}^{n_1} \times \mathbb{R}^{n_2}$ for some bounded polyhedron $\bar{\mathcal{X}} := \{x \in \mathbb{R}^n \,|\, C_x x \le d_x\}$ with $C_x \in \mathbb{R}^{s_x \times n}$ and $d_x \in \mathbb{R}^{s_x}$, $c_1 \in \mathbb{R}^n$, and $\xi \in \mathbb{R}^m$ is a random vector with a distribution known to be supported on a subset of the bounded polyhedron $\Xi := \{\xi \in \mathbb{R}^m \,|\, C_\xi \xi \le d_\xi\}$ with $C_\xi \in \mathbb{R}^{s_\xi \times m}$ and $d \in \mathbb{R}^{s_\xi}$. The expectation, which expresses a risk neutral attitude for a random cost with known distribution, is taken with respect to the probability distribution F_ξ that belongs to an ambiguity set $\mathcal{D}$, and is applied to the objective value of the second-stage problem:

$$h(x, \xi) := \underset{y}{\text{minimize}} \quad c_2^\mathsf{T} y \tag{6a}$$

$$\text{subject to} \quad Ay \ge W(\xi)x + b. \tag{6b}$$

where $W : \mathbb{R}^m \to \mathbb{R}^{s \times m}$ is an affine mapping defined as $W(\xi) := \sum_{i=1}^m W_i \xi_i + W_0$ for some $W_i \in \mathbb{R}^{s \times n}$ for each $i = 1, \ldots, m$. We assume that problem (5) has relatively complete recourse, *i.e.,* for all $x \in \mathcal{X}$ and $\xi \in \Xi$, the recourse problem (6) has a feasible solution, and assume that the recourse problem is bounded for all $x \in \mathcal{X}$ and $\xi \in \Xi$.

To simplify the exposition, we start by making the following assumption which will be relaxed in Section 5.5.

Assumption 1 *The feasible set is a discrete set, i.e., $\mathcal{X} := \bar{\mathcal{X}} \cap \mathbb{Z}^n$.*

Now, instead of choosing a single action/solution x^*, let's consider the case as in RSP where the decision maker can randomize between multiple actions/solutions. Following Assumption 1, since $\mathcal{X}$ is a discrete set, let $\mathcal{K} := \{1, 2, \ldots, |\mathcal{X}|\}$ be the index set of all members of $\mathcal{X}$, *i.e.,* $\mathcal{X} = \{\bar{x}^k\}_{k \in \mathcal{K}}$. The randomized strategy problem then reduces to determining an optimal distribution function F_x parametrized by $p \in \mathbb{R}_+^{|\mathcal{K}|}$ such that $\mathbb{P}_{F_x}(X = \bar{x}^k) = p_k$, *i.e.,* each p_k is the probability that the randomized policy selects the feasible action $\bar{x}^k$. Mathematically, the randomized strategy problem can be rewritten as

$$\underset{p \in \mathbb{R}^{|\mathcal{K}|}}{\text{minimize}} \qquad \sum_{k \in \mathcal{K}} c_1^\mathsf{T} \bar{x}^k p_k + \sup_{F_\xi \in \mathcal{D}} \mathbb{E}_{\xi \sim F_\xi}[g(p, \xi)] \tag{7a}$$

$$\text{subject to} \qquad p_k \ge 0 \,, \, \forall \, k \in \mathcal{K}, \quad \sum_{k \in \mathcal{K}} p_k = 1 \,, \tag{7b}$$

where $g(p, \xi) := \sum_{k \in \mathcal{K}} h(\bar{x}^k, \xi) p_k$ is used for ease of exposition. Note that to obtain this reformulation, one should start from problem (2) and exploit the fact that X and ξ are independent, that X is discrete, and that the expectation operator is linear:

$$\underset{F_x \in \Delta(\mathcal{X})}{\text{minimize}} \ \underset{F_\xi \in \mathcal{D}}{\sup} \ \mathbb{E}_{(X,\xi) \sim F_x \times F_\xi} \left[c_1^\intercal X + h(X, \xi) \right]$$

$$\equiv \ \underset{F_x \in \Delta(\mathcal{X})}{\text{minimize}} \ \underset{F_\xi \in \mathcal{D}}{\sup} \ \mathbb{E}_{X \sim F_x} \left[c_1^\intercal X + \mathbb{E}_{\xi \sim F_\xi}[h(X, \xi)] \right]$$

$$\equiv \ \underset{p : p_k \geq 0, \ \sum_{k \in \mathcal{K}} p_k = 1}{\text{minimize}} \ \sum_{k \in \mathcal{K}} c_1^\intercal \bar{x}^k p_k + \underset{F_\xi \in \mathcal{D}}{\sup} \ \sum_{k \in \mathcal{K}} \mathbb{E}_{\xi \sim F_\xi} \left[h(\bar{x}^k, \xi) \right] p_k$$

$$\equiv \ \underset{p : p_k \geq 0, \ \sum_{k \in \mathcal{K}} p_k = 1}{\text{minimize}} \ \sum_{k \in \mathcal{K}} c_1^\intercal \bar{x}^k p_k + \underset{F_\xi \in \mathcal{D}}{\sup} \ \mathbb{E}_{\xi \sim F_\xi} \left[\sum_{k \in \mathcal{K}} h(\bar{x}^k, \xi) p_k \right]$$

$$\equiv \ \underset{p : p_k \geq 0, \ \sum_{k \in \mathcal{K}} p_k = 1}{\text{minimize}} \ \sum_{k \in \mathcal{K}} c_1^\intercal \bar{x}^k p_k + \underset{F_\xi \in \mathcal{D}}{\sup} \ \mathbb{E}_{\xi \sim F_\xi} \left[g(p, \xi) \right].$$

As usual, the first step in dealing with DRO problems is to try to reformulate them as finite dimensional robust optimization problems. Indeed, whether such a reformulation exists depends on the definition of the ambiguity set $\mathcal{D}$. In what follows, we show how to reformulate problem (7) for two important classes of ambiguity sets: moment-based and Wasserstein ambiguity sets.

5.1 A reformulation for moment-based ambiguity sets

We first consider a moment-based ambiguity set defined as

$$\mathcal{D}(\Xi, \mu, \gamma) := \left\{ F_\xi \mid \mathbb{P}_{\xi \sim F_\xi}(\xi \in \Xi) = 1, \ \mathbb{E}_{\xi \sim F_\xi}[\xi] = \mu, \ \mathbb{E}_{\xi \sim F_\xi}[\pi_l(\xi)] \leq \gamma_l, \ \forall l \in \mathcal{L} \right\}, \quad (8)$$

where $\mu \in \mathbb{R}^m$ is the known mean of ξ, and where for each $l \in \mathcal{L}$ with $|\mathcal{L}|$ finite, the moment function $\pi_l : \mathbb{R}^m \to \mathbb{R}$ is piecewise-linear convex and its expectation bounded by $\gamma_l \in \mathbb{R}$. This ambiguity set can be considered a special case of the set presented in Bertsimas et al. (2017), where both Ξ and $\pi_l(\cdot)$ are considered second-order cone representable (implications of our results to this more general case will bebriefly discussed in Remark 2). A notable example of piecewise linear moment function is when $\pi_l(\xi) := |a_l^\intercal(\xi - \mu)|$, which places an upper bound on absolute deviation along the direction of a_l. On the other hand, a function that places an upper bound on variance, *i.e.*, $\pi_l(\xi) = (\xi - \mu)^2$ would need to be treated as discussed in Remark 2.

Following the work of Wiesemann et al. (2014), we can redefine $\mathcal{D}(\Xi, \mu, \gamma)$ using a *lifting* to the space of $[\xi^\intercal \ \zeta^\intercal]^\intercal$ with $\zeta \in \mathbb{R}^{|\mathcal{L}|}$ capturing a vector of random bounds on each $\pi_l(\xi)$ so that problem (7) with $\mathcal{D}(\Xi, \mu, \gamma)$ is equivalent to

$$\underset{p : p_k \geq 0, \ \sum_{k \in \mathcal{K}} p_k = 1}{\text{minimize}} \ \sum_{k \in \mathcal{K}} c_1^\intercal \bar{x}^k p_k + \underset{F_{(\xi,\zeta)} \in \mathcal{D}(\Xi', [\mu^\intercal \ \gamma^\intercal]^\intercal)}{\sup} \ \mathbb{E}_{(\xi,\zeta) \sim F_{(\xi,\zeta)}} \left[g(p, \xi) \right]$$

where

$$\mathcal{D}(\Xi', [\mu^\intercal \ \gamma^\intercal]^\intercal) := \left\{ F_{(\xi,\zeta)} \ \middle| \ \begin{array}{l} \mathbb{P}_{(\xi,\zeta) \sim F_{(\xi,\zeta)}}((\xi,\zeta) \in \Xi') = 1 \\ \mathbb{E}_{(\xi,\zeta) \sim F_{(\xi,\zeta)}}[\xi] = \mu \\ \mathbb{E}_{(\xi,\zeta) \sim F_{(\xi,\zeta)}}[\zeta] = \gamma \end{array} \right\},$$

with

$$\Xi' := \left\{ (\xi, \zeta) \ \middle| \ \begin{array}{c} \xi \in \Xi \\ \pi_l(\xi) \leq \zeta_l \leq \bar{\zeta}_{\max}, \ \forall l \in \mathcal{L} \end{array} \right\},$$

where $\bar{\zeta}_{\max} := \sup_{l \in \mathcal{L}, \xi \in \Xi} \pi_l(\xi)$ and where $\zeta_l \leq \bar{\zeta}_{\max}$ is added to make Ξ' bounded without affecting the quality of the reformulation. One can readily verify that Ξ' is polyhedral under the piecewise-linear convexity assumption of each $\pi_l(\xi)$. Using the reformulation proposed by Wiesemann et al. (2014), which is based on strong duality of semi-infinite conic programs (Shapiro 2001, Theorem 3.4), one can simplify the worst-case expectation expression as follows:

$$\sup_{F_{(\xi,\zeta)} \in \mathcal{D}(\Xi', [\mu^\intercal \ \gamma^\intercal]^\intercal)} \mathbb{E}_{(\xi,\zeta) \sim F_{(\xi,\zeta)}} [g(p,\xi)] = \max_{[\xi^\intercal \ \zeta^\intercal]^\intercal \in \Xi'} \inf_{q,\lambda} g(p,\xi) + (\mu - \xi)^\intercal q + (\gamma - \zeta)^\intercal \lambda$$

$$= \inf_{q,\lambda} \mu^\intercal q + \gamma^\intercal \lambda + \max_{[\xi^\intercal \ \zeta^\intercal]^\intercal \in \Xi'} g(F_x, \xi) - \xi^\intercal q - \zeta^\intercal \lambda$$

which can then be reintegrated in the main optimization problem as

$$\underset{p,q,\lambda,t}{\text{minimize}} \qquad \sum_{k \in \mathcal{K}} c_1^\intercal \bar{x}^k p_k + \mu^\intercal q + \lambda^\intercal \gamma + t \tag{9a}$$

$$\text{subject to} \qquad \max_{(\xi,\zeta) \in \Xi'} g(p,\xi) - \xi^\intercal q - \zeta^\intercal \lambda \le t \tag{9b}$$

$$p \ge 0, \quad \sum_{k \in \mathcal{K}} p_k = 1. \tag{9c}$$

We are left with a finite dimensional robust two-stage linear optimization problem which could in theory be solved either approximately using linear decision rules (see Ben-Tal et al. (2004)) or exactly using, for example, the column-and-constraint generation method in Zeng and Zhao (2013). Unfortunately, in both cases the problem is highly intractable since it potentially involves an exponential number of decision variables due to $|\mathcal{K}|$. The numerical difficulty associated with the exact resolution of this problem will be addressed shortly using a two-layer column generation method.

5.2 A reformulation for Wasserstein ambiguity sets

The second class of ambiguity sets that we consider consists of an ambiguity set defined by a Wasserstein ball centered at some empirical distribution $\widehat{F}_\xi$ as introduced in Mohajerin Esfahani and Kuhn (2017). Specifically, we let $\mathcal{D}(\widehat{F}_\xi, \epsilon)$ be a ball of radius $\epsilon > 0$ centered at the empirical distribution $\widehat{F}_\xi^\Omega$ constructed based on a set $\{\widehat{\xi}_\omega\}_{\omega \in \Omega} \subset \Xi$ of i.i.d. observations. More specifically,

$$\mathcal{D}(\widehat{F}_\xi^\Omega, \epsilon) := \left\{ F_\xi \in \mathcal{M}(\Xi) \mid d_\mathrm{W}(F_\xi, \widehat{F}_\xi) \le \epsilon \right\}, \tag{10}$$

where $\mathcal{M}(\Xi)$ is the space of all distributions supported on Ξ and $d_\mathrm{W} : \mathcal{M}(\Xi) \times \mathcal{M}(\Xi) \to \mathbb{R}$ is the Wasserstein metric defined as

$$d_\mathrm{W}(F_1, F_2) := \inf \left\{ \int_{\Xi^2} \|\xi_1 - \xi_2\| \Pi(\mathrm{d}\xi_1, \mathrm{d}\xi_2) \ \middle| \ \begin{array}{l} \Pi \text{ is a joint distribution of } \xi_1 \text{ and } \xi_2 \\ \text{with marginals } F_1 \text{ and } F_2 \text{ respectively} \end{array} \right\},$$

where $\| \cdot \|$ represents an arbitrary norm on $\mathbb{R}^m$. This ambiguity set has become very popular in the recent years given that it can directly incorporate the information obtained from past observations of ξ while letting the decision maker control, through his selection of ϵ, the optimism of the model regarding how close the future realization will be from any of the observed ones. We refer the reader to Kantorovich and Rubinshtein (1958) for more technical details about $\mathcal{D}(\widehat{F}_\xi^\Omega, \epsilon)$ and for statistical methods that can be used to calibrate ϵ so that $\mathcal{D}(\widehat{F}_\xi^\Omega, \epsilon)$ has a certain guarantee of containing the true underlying distribution from which the observations were drawn.

Using similar steps as used in Mohajerin Esfahani and Kuhn (2017), problem (7) can be reformulated as

$$\underset{p,\lambda \ge 0, \{t_\omega\}_{\omega \in \Omega}}{\text{minimize}} \qquad \sum_{k \in \mathcal{K}} c_1^\intercal \bar{x}^k p_k + \lambda \epsilon + \frac{1}{|\Omega|} \sum_{\omega \in \Omega} t_\omega \tag{11a}$$

$$\text{subject to} \qquad \max_{\xi \in \Xi} \left(g(p,\xi) - \lambda \|\xi - \widehat{\xi}_\omega\| \right) \le t_\omega \quad \forall \omega \in \Omega \tag{11b}$$

$$p \ge 0, \sum_{k \in \mathcal{K}} p_k = 1, \tag{11c}$$

where each $t_\omega \in \mathbb{R}$.

In order to make problem (11) take the form of a finite dimensional robust two-stage linear optimization problem as was done for the moment-based ambiguity set in (9) for each of the constraints indexed by $\omega \in \Omega$, we assume that the l_1-norm is used in the Wasserstein metric and use the lifted bounded polyhedral uncertainty set

$$\Xi'_\omega := \left\{ (\xi, \zeta) \in \mathbb{R}^m \times \mathbb{R} \;\middle|\; \begin{array}{c} \xi \in \Xi \\ \|\xi - \widehat{\xi}_\omega\|_1 \leq \zeta \leq \bar{\zeta}_{\max} \end{array} \right\},$$

where $\bar{\zeta}_{\max} := \sup_{\xi \in \Xi} \|\xi - \widehat{\xi}_\omega\|_1$ is again chosen such that $\zeta \leq \bar{\zeta}_{\max}$ makes Ξ'_ω bounded while preserving the exactness of the reformulation. With that, our two-stage DRO problem with randomization and a Wasserstein ambiguity set can be reformulated as the robust two-stage linear optimization problem:

$$\underset{p \geq 0, \lambda \geq 0, \{t_\omega\}_{\omega \in \Omega}}{\text{minimize}} \quad \sum_{k \in \mathcal{K}} c_1^\mathsf{T} \bar{x}^k p_k + \lambda \epsilon + \frac{1}{|\Omega|} \sum_{\omega \in \Omega} t_\omega \tag{12a}$$

$$\text{subject to} \quad \sup_{(\xi, \zeta) \in \Xi'_\omega} \left(\sum_{k \in \mathcal{K}} h(\bar{x}^k, \xi) p_k - \lambda \zeta \right) \leq t_\omega \quad \forall \omega \in \Omega \tag{12b}$$

$$\sum_{k \in \mathcal{K}} p_k = 1. \tag{12c}$$

5.3 A two-layer column generation algorithm

We just established that under fairly weak assumptions, *i.e.,* piecewise-linear moment functions in (8) and l_1-norm in (10), one can reformulate problem (5) under both the moment-based and the Wasserstein ambiguity sets as robust two-stage linear optimization problems yet with an excessively large number of decision variables. In this section, we propose a two-layer column generation algorithm that can identify an optimal randomized strategy for problem (5) together with its optimal value v_r. For simplicity of exposition, our discussion will focus on the case of a Wasserstein ambiguity set yet can easily be modifed to accommodate problem (9).

First, we note that since the inner optimization in Equation (12b) is a convex maximization over a bounded polyhedral set, its maxima is attained at one of the vertices of Ξ'_ω. Hence, we replace the maximization over Ξ'_ω with a maximization over the set of vertices $\{\xi^{h_\omega}, \zeta^{h_\omega}\}_{h_\omega \in \mathcal{H}_\omega}$ for each $\omega \in \Omega$, where $\mathcal{H}_\omega$ is the index set for the vertices of Ξ'_ω. Hence, problem (12) can be rewritten as the large-scale LP:

$$\underset{p \geq 0, \lambda \geq 0, \{t_\omega\}_{\omega \in \Omega}}{\text{minimize}} \quad \sum_{k \in \mathcal{K}} c_1^\mathsf{T} \bar{x}^k p_k + \lambda \epsilon + \frac{1}{|\Omega|} \sum_{\omega \in \Omega} t_\omega \tag{13a}$$

$$\text{subject to} \quad \sum_{k \in \mathcal{K}} h(\bar{x}^k, \xi^{h_\omega}) p_k - \zeta^{h_\omega} \lambda \leq t_\omega \quad \forall h_\omega \in \mathcal{H}_\omega, \; \omega \in \Omega \tag{13b}$$

$$\sum_{k \in \mathcal{K}} p_k = 1. \tag{13c}$$

Except for very small instances, it is impossible to enumerate and include the entire sets of all $|\mathcal{K}|$ decision variables and $\sum_{\omega \in \Omega} |\mathcal{H}_\omega|$ vertices in this problem. Instead, we begin with subsets $\mathcal{K}' \subset \mathcal{K}$ and $\mathcal{H}'_\omega \subset \mathcal{H}_\omega$ for each $\omega \in \Omega$, and employ a two-layer column generation algorithm to generate and add new elements iteratively, as needed. The algorithm operates as follows (a *pseudocode* description is also presented in Appendix B):

1. Initialize the subset $\mathcal{K}' \subset \mathcal{K}$ to any singleton (*e.g.,* so that $\{x_k\}_{k \in \mathcal{K}'}$ contains the solution x_d^* to DSP). Initialize the sets $\mathcal{H}'_\omega = \emptyset$ for all $\omega \in \Omega$. Finally, set the upper bound $UB = \infty$ and the lower bound $LB = -\infty$.

2. Solve the primal master problem which seeks an optimal randomized strategy supported on $\{x_k\}_{k \in \mathcal{K}'}$ while considering the full set of vertices $\mathcal{H}_\omega$ for all $\omega \in \Omega$,

$$[\texttt{Primal}(\mathcal{K}', \mathcal{H})]: \quad \underset{p \geq 0, \lambda \geq 0, \{t_\omega\}_{\omega \in \Omega}}{\text{minimize}} \quad \sum_{k \in \mathcal{K}'} c_1^{\mathsf{T}} \bar{x}^k p_k + \lambda \epsilon + \frac{1}{|\Omega|} \sum_{\omega \in \Omega} t_\omega$$

$$\text{subject to} \quad \sum_{k \in \mathcal{K}'} h(\bar{x}^k, \xi^{h_\omega}) p_k - \zeta^{h_\omega} \lambda \leq t_\omega \quad \forall h_\omega \in \mathcal{H}_\omega, \ \omega \in \Omega \quad (q_{h_\omega}^\omega)$$

$$\sum_{k \in \mathcal{K}'} p_k = 1, \qquad\qquad\qquad\qquad\qquad (w),$$

where $p \in \mathbb{R}^{|\mathcal{K}'|}$ and where $\mathcal{H} := \{\mathcal{H}_\omega\}_{\omega \in \Omega}$. The solution of $\texttt{Primal}(\mathcal{K}', \mathcal{H})$ provides a feasible solution to problem (13) hence an upper bound which will be used to update UB. Unfortunately, one cannot handle all the constraints of this problem indexed with $h_\omega \in \mathcal{H}_\omega$. Therefore, we generate the ones that are needed to confirm optimality through the following sub-procedure:

   (a) Set $LB_1 := LB$ and $UB_1 := \infty$.

   (b) Initialize p^* and λ^* to any arbitrary solution that satisfy $p^* \geq 0$, $\lambda^* \geq 0$, and $\sum_{k \in \mathcal{K}'} p_k^* = 1$.

   (c) For each $\omega \in \Omega$, solve the subproblem

$$[\texttt{SP1}_\omega] \quad \underset{(\xi, \zeta) \in \Xi_\omega'}{\text{maximize}} \sum_{k \in \mathcal{K}'} p_k^* h(\bar{x}^k, \xi) - \lambda^* \zeta,$$

   to generate a new worst-case vertex $(\xi^{h_\omega}, \zeta^{h_\omega})$ in each support set Ξ_ω'. This can be done by solving a mixed-integer linear program as will be described in Proposition 3.

   (d) Let t_ω^* be the optimal value of $\texttt{SP1}_\omega$ for all $\omega \in \Omega$. Update the upper bound as

$$UB_1 := \min \left(UB_1, \sum_{k \in \mathcal{K}} c_1^{\mathsf{T}} \bar{x}^k p_k^* + \lambda^* \epsilon + \frac{1}{|\Omega|} \sum_{\omega \in \Omega} t_\omega^* \right).$$

   (e) Add the index of each new vertex generated $(\xi^{h_\omega}, \zeta^{h_\omega})$ to its respective index subset $\mathcal{H}_\omega'$, $\omega \in \Omega$ and solve the primal master problem (MP1), defined as the restricted version of $\texttt{Primal}(\mathcal{K}', \mathcal{H}')$ where each set $\mathcal{H}_\omega$ is replaced with its subset $\mathcal{H}_\omega'$, to update (p^*, λ^*) and obtain a lower bound LB_1. Note that for each $k \in \mathcal{K}'$ the values of $h(\bar{x}^k, \xi^{h_\omega})$ for all $h_\omega \in \mathcal{H}_\omega'$ only needs to be computed once when a new vertex is added to $\mathcal{H}_\omega'$.

   (f) If $UB_1 - LB_1 < \varepsilon/2$, where ε is a sufficiently small tolerance, terminate the sub-algorithm and set $UB = UB_1$. Otherwise, return to Step 2c.

3. Solve the dual master problem which seeks an optimal randomized strategy supported on the whole $\{x_k\}_{k \in \mathcal{K}}$ set while considering the set of vertices $\mathcal{H}_\omega'$ for all $\omega \in \Omega$,

$$[\texttt{Dual}(\mathcal{K}, \mathcal{H}')]: \quad \underset{w, \{q^\omega\}_{\omega \in \Omega}}{\text{maximize}} \quad w \qquad\qquad\qquad (14a)$$

$$\text{subject to} \quad w \leq c_1^{\mathsf{T}} \bar{x}^k + \sum_{\omega \in \Omega} \sum_{h_\omega \in \mathcal{H}_\omega'} h(\bar{x}^k, \xi^{h_\omega}) q_{h_\omega}^\omega \quad \forall k \in \mathcal{K} \quad (p_k) \qquad (14b)$$

$$\sum_{\omega \in \Omega} \sum_{h_\omega \in \mathcal{H}_\omega'} \zeta^{h_\omega} q_{h_\omega}^\omega \leq \epsilon \qquad\qquad\qquad (\lambda) \qquad (14c)$$

$$\sum_{h_\omega \in \mathcal{H}_\omega'} q_{h_\omega}^\omega = \frac{1}{|\Omega|} \qquad\qquad\qquad \forall \omega \in \Omega \quad (t_\omega) \qquad (14d)$$

$$q^\omega \geq 0 \qquad\qquad\qquad\qquad \forall \omega \in \Omega. \qquad (14e)$$

The optimal value of $\texttt{Dual}(\mathcal{K}, \mathcal{H}')$ provides a lower bound for problem (13) which should be used to update LB. Note that the optimal dual variables can also be used to initialize p^* and λ^* in Step 2b. Since we cannot handle the full set of actions $\mathcal{X}$ implemented by the randomized strategy, we progressively construct an optimal support of reasonable size using the following sub-procedure:

   (a) Set $LB_2 := -\infty$ and $UB_2 := UB$

(b) Initialize each $q_{h_\omega}^{\omega*}$ to an arbitrary solution that satisfy constraints (14c), (14d) and (14e). In practice, one can obtain such a valid assignment based on the optimal assignment for the dual variables of $\texttt{Primal}(\mathcal{K}', \mathcal{H})$.

(c) Solve the subproblem $\displaystyle\minimize_{x\in\mathcal{X}}\ c_1^\intercal x + \sum_{\omega\in\Omega}\sum_{h_\omega\in\mathcal{H}'_\omega} h(x, \xi^{h_\omega}) q_{h_\omega}^{\omega*}$, which reduces to the mixed-integer linear program

$$[\texttt{SP2}]: \quad \minimize_{x\in\mathcal{X}, y}\quad c_1^\intercal x + \sum_{\omega\in\Omega}\sum_{h_\omega\in\mathcal{H}'_\omega} q_{h_\omega}^{\omega*} c_2^\intercal y_{h_\omega}$$

$$\text{subject to}\quad Ay_{h_\omega} - W(\xi^{h_\omega})x \geq b \qquad \forall h_\omega \in \mathcal{H}'_\omega.$$

Consider the optimal x^* to be a new support point for the optimal randomized strategy.

(d) Let w^* be the optimal value obtained for $\texttt{SP2}$. The lower bound is updated as $LB_2 := \max(LB_2, w^*)$.

(e) Add the index of the new support point x^* to $\mathcal{K}'$ and solve the dual master problem ($\texttt{MP2}$), defined as the restricted version of $\texttt{Dual}(\mathcal{K}', \mathcal{H}')$ where $\mathcal{K}$ is replaced with $\mathcal{K}'$, to update $q_{h_\omega}^{\omega*}$ and obtain an upper bound UB_2. Note that for each $h_\omega \in \mathcal{H}'_\omega$ the values of $h(\bar{x}^k, \xi^{h_\omega})$ for all $k \in \mathcal{K}'$ needs to be computed only once when a new support point is added to $\mathcal{K}'$.

(f) If $UB_2 - LB_2 < \varepsilon/2$, terminate the sub-algorithm and set $LB = LB_2$. Otherwise, return to Step 3c.

4. Iterate between the steps (2) and (3) until $UB - LB < \varepsilon$.

To complete the presentation of the two-layer column generation algorithm, we present how problem $\texttt{SP1}_\omega$ can be reformulated as a mixed-integer linear program.

Proposition 3 *Problem $\texttt{SP1}_\omega$ is equivalent to the following mixed-integer linear program:*

$$\maximize_{\substack{\xi,\zeta,\delta,\alpha,\beta,\psi,\phi\geq 0 \\ Bin^1, Bin^2, Bin^3, Bin^4, Bin^5, Bin^6, Bin^7}} \quad \sum_{k\in\mathcal{K}'} (W_0\bar{x}^k + b)^\intercal\phi_k + d_\xi^\intercal\alpha + \bar{\zeta}_{max}\gamma + \widehat{\xi_\omega^\intercal}(\psi^+ - \psi^-)$$

$$\text{subject to}\quad A^\intercal\phi_k = c_2 p_k^* \qquad\qquad \forall k\in\mathcal{K}' \tag{15a}$$

$$\sum_{i=1}^m \left(\sum_{k\in\mathcal{K}'}\phi_k^T W_i x_k\right) e_i = C_\xi^T\alpha + \psi^+ - \psi^- \tag{15b}$$

$$0 \leq d - C\xi \leq M(1 - Bin^1) \tag{15c}$$

$$0 \leq \alpha \leq MBin^1 \tag{15d}$$

$$0 \leq \zeta - e^\intercal\delta \leq M(1 - Bin^2) \tag{15e}$$

$$0 \leq \beta \leq MBin^2 \tag{15f}$$

$$0 \leq \bar{\zeta}_{max} - \zeta \leq M(1 - Bin^3) \tag{15g}$$

$$0 \leq \gamma \leq MBin^3 \tag{15h}$$

$$0 \leq \delta - \xi + \widehat{\xi}_\omega \leq M(1 - Bin^4) \tag{15i}$$

$$0 \leq \psi^+ \leq MBin^4 \tag{15j}$$

$$0 \leq \delta + \xi - \widehat{\xi}_\omega \leq M(1 - Bin^5) \tag{15k}$$

$$0 \leq \psi^- \leq MBin^5 \tag{15l}$$

$$0 \leq \zeta \leq M(1 - Bin^6) \tag{15m}$$

$$0 \leq \lambda^* + \gamma - \beta \leq MBin^6 \tag{15n}$$

$$0 \leq \delta \leq M(1 - Bin^7) \tag{15o}$$

$$0 \leq \beta - \psi^+ - \psi^- \leq MBin^7 \tag{15p}$$

$$Bin^1 \in \{0, 1\}^{s_\xi},\ Bin^2 \in \{0, 1\},\ Bin^3 \in \{0, 1\},\ Bin^4 \in \{0, 1\}^m,$$

$$Bin^5 \in \{0, 1\}^m,\ Bin^6 \in \{0, 1\},\ Bin^7 \in \{0, 1\}^m, \tag{15q}$$

where $\phi_k \in \mathbb{R}_+^s,\ \alpha \in \mathbb{R}_+^{s_\xi},\ \beta \in \mathbb{R}_+,\ \gamma \in \mathbb{R}_+,\ \psi^+ \in \mathbb{R}_+^m$ and $\psi^- \in \mathbb{R}_+^m$.

It is worth emphasizing that, in contrast to the reformulation that would be obtained by applying the scheme of Zeng and Zhao (2013) directly, our MILP reformulation ensures that the number of binary variables does not increase with the size of the support $\mathcal{K}'$ of the randomized strategy. This alternative approach leads to a significant reduction in the solution time for $\mathtt{SP1}_\omega$.

Theorem 2 *The two-layer column generation algorithm presented in Algorithm 1 converges in a finite number of iterations.*

Remark 2 *In Theorem 2, the convergence in finite number of iteration follows from our assumption that $\mathcal{X}$ is a bounded discrete set, Ξ is a bounded polyhedron and that the Wasserstein ambiguity set employs a metric that is based on the l_1-norm (or alternatively that $\pi(\xi)$ is piecewise-linear in the case of a moment-based ambiguity set). However, Algorithm 1 is generic and can be used for more general forms of decision spaces and ambiguity sets. In particular, we will discuss in the next two sections how the algorithm can be modified to handle applications where there is no recourse, and where $\mathcal{X}$ is mixed-integer, namely $n_2 > 0$. The case of general ambiguity set can also be accommodated but requires one to design an alternative scheme for solving $\mathtt{SP1}_\omega$. In particular, one might suspect that, following the work of Zeng and Zhao (2013), if Ξ'_ω (for Wasserstein ambiguity set) or Ξ' (for the moment-based ambiguity set) are second-order cone representable, then similar arguments as those used in the proof of Proposition 3 could be used to design an equivalent mixed-integer second-order cone programming problem. The question of whether the algorithm would still be guaranteed to converge in a finite number of iterations for such ambiguity sets remain open for future research.*

5.4 The case of single-stage problems

The proposed algorithm can be used for single-stage DRO problems where $h(x,\xi) := \xi^\intercal C_2 x$ with $C_2 \in \mathbb{R}^{m \times n}$. The fact that there is no recourse problem makes problem (12) reduce to a simpler single-stage problem which is amenable to classical derivation of the robust counterpart model:

$$\underset{\substack{p_k \geq 0, \lambda \geq 0, t_\omega \\ \alpha_\omega, \beta_\omega, \gamma_\omega \\ \psi_\omega^+, \psi_\omega^-}}{\text{minimize}} \quad \sum_{k \in \mathcal{K}} c_1^\intercal \bar{x}^k p_k + \lambda \epsilon + \frac{1}{|\Omega|} \sum_{\omega \in \Omega} t_\omega \tag{16a}$$

$$\text{subject to} \quad d_\xi^\intercal \alpha_\omega + \bar{\zeta}_{\max} \gamma_\omega + \widehat{\xi}_\omega^\intercal (\psi_\omega^+ - \psi_\omega^-) \leq t_\omega \,, \forall \omega \in \Omega \tag{16b}$$

$$C_\xi^\intercal \alpha_\omega + \psi_\omega^+ - \psi_\omega^- = \sum_{k \in \mathcal{K}} p_k C_2 \bar{x}^k \,, \forall \omega \in \Omega \qquad (\xi_\omega) \tag{16c}$$

$$e\beta_\omega \leq \lambda + \gamma_\omega \,, \forall \omega \in \Omega \qquad\qquad (\zeta_\omega) \tag{16d}$$

$$\psi_\omega^+ + \psi_\omega^- \leq \beta_\omega \,, \forall \omega \in \Omega \tag{16e}$$

$$\sum_{k \in \mathcal{K}} p_k = 1 \,. \tag{16f}$$

where $\alpha_\omega \in \mathbb{R}_+^{s_\xi}$, $\beta_\omega \in \mathbb{R}_+$, $\gamma_\omega \in \mathbb{R}_+$, $\phi_\omega^+ \in \mathbb{R}_+^m$ and $\phi_\omega^- \in \mathbb{R}_+^m$. One still needs to employ an algorithm such as Algorithm 1 to get a solution because of the size of $\mathcal{K}$. $\mathtt{Primal}(\mathcal{K}', \mathcal{H})$ is however unnecessary since vertices do not need to be enumerated anymore. Hence, the two-layer column generation reduces to a single-layer version that focuses on Step 3. In particular, for this step, $\mathtt{Dual}(\mathcal{K}, \mathcal{H})$ reduces to:

$$\underset{w, q' \geq 0}{\text{maximize}} \quad w$$

$$\text{subject to} \quad w \leq c_1^\intercal \bar{x}^k + \frac{1}{|\Omega|} \sum_{\omega \in \Omega} \bar{x}^k C_2^\intercal \left(\sum_{h_\omega \in \mathcal{H}_\omega} \xi^{h_\omega} {q_{h_\omega}^\omega}' \right) \quad \forall k \in \mathcal{K} \qquad (p_k)$$

$$\frac{1}{|\Omega|} \sum_{\omega \in \Omega} \sum_{h_\omega \in \mathcal{H}_\omega} \zeta^{h_\omega} {q_{h_\omega}^\omega}' \leq \epsilon \qquad\qquad (\lambda)$$

$$\sum_{h_\omega \in \mathcal{H}_\omega} {q_{h_\omega}^\omega}' = 1 \qquad\qquad \forall \omega \in \Omega,$$

where we replaced $q' := |\Omega|q$. In this reformulation, one recognizes that $\{q_{h_\omega}^{\omega}{}'\}_{h_\omega \in \mathcal{H}_\omega, \omega \in \Omega}$ serves as a way of encoding convex combinations of all the vertices of Ξ'_ω. We therefore can simply replace these to simplify the model and obtain:

$$[\texttt{Dual}'(\mathcal{K}, \mathcal{H})]: \quad \underset{w, \{\xi_\omega, \zeta_\omega\}_{\omega \in \Omega}}{\text{maximize}} \quad w$$

$$\text{subject to} \quad w \leq c_1^{\mathsf{T}} \bar{x}^k + \bar{x}^k C_2^{\mathsf{T}} \left(\frac{1}{|\Omega|} \sum_{\omega \in \Omega} \xi_\omega \right) \quad \forall k \in \mathcal{K} \quad (p_k)$$

$$\frac{1}{|\Omega|} \sum_{\omega \in \Omega} \zeta_\omega \leq \epsilon \quad\quad\quad\quad\quad\quad (\lambda)$$

$$C_\xi \xi_\omega \leq d_\xi \quad\quad\quad\quad \forall \omega \in \Omega$$

$$\|\hat{\xi}_\omega - \xi_\omega\|_1 \leq \zeta_\omega \quad\quad\quad \forall \omega \in \Omega.$$

Furthermore, the subproblem straighforwardly reduces to

$$[\texttt{SP2}']: \quad \underset{x \in \mathcal{X}}{\text{minimize}} \quad \left(c_1 + C_2^{\mathsf{T}} \frac{1}{|\Omega|} \sum_{\omega \in \Omega} \xi_\omega^* \right)^{\mathsf{T}} x.$$

Intuitively, the restricted master problem $\texttt{MP} \equiv \texttt{Dual}'(\mathcal{K}', \mathcal{H})$ identifies a worst-case $\mathbb{E}_{F_\xi}[\xi]$ among all $F_\xi \in \mathcal{D}(\widehat{F}_\xi, \epsilon)$ while the sub-problem $\texttt{SP2}'$ identifies the member of $\mathcal{X}$ that is best-suited for this worst-case mean vector. Section 6.2 presents a customized version of this single layer column generation algorithm for the DRUFLP described in Section 3.

5.5 The case of mixed-integer feasible region

We briefly outline how the algorithm presented in Section 5.3 can be modified to handle more general mixed-integer feasible sets $\mathcal{X}$ that do not satisfy Assumption 1, *i.e.*, $\mathcal{X} = \bar{\mathcal{X}} \cap \mathbb{Z}^{n_1} \times \mathbb{R}^{n_2}$ with $n_2 \neq 0$. Similarly as was done in Proposition 2, we let $\{\bar{x}_1^k\}_{k \in \mathcal{K}} \subset \mathbb{Z}^{n_1}$ describe the finite set, indexed using $k \in \mathcal{K}$, of feasible assignments for the integer decision variables, *i.e.*, $\{\bar{x}_1^k\}_{k \in \mathcal{K}} := \{x_1 \in \mathbb{Z}^{n_1} \mid \exists x_2 \in \mathbb{R}^{n_2}, [x_1^{\mathsf{T}} \ x_2^{\mathsf{T}}]^{\mathsf{T}} \in \mathcal{X}\}$.

Proposition 4 *Let $\mathcal{X}$ be mixed-integer, the decision maker's attitude be AARN, and the cost function $h(x, \xi)$ capture a two-stage decision problem as described in Equation* (6). *Then, the* RSP *presented in Equation* (2) *is equivalent to*

$$\underset{p \in \mathbb{R}^{|\mathcal{K}|}, \{z^k\}_{k \in \mathcal{K}}}{\text{minimize}} \quad \sum_{k \in \mathcal{K}} c_1^{\mathsf{T}} z^k + \sup_{F_\xi \in \mathcal{D}} \mathbb{E}_{\xi \sim F_\xi} \left[\sum_{k \in \mathcal{K}} h'(p_k, z^k, \xi) \right] \tag{17a}$$

$$\text{subject to} \quad C_x z^k \leq d_x p_k, \, \forall k \in \mathcal{K} \tag{17b}$$

$$P_{\mathbb{Z}} z^k = \bar{x}_1^k p_k, \, \forall k \in \mathcal{K} \tag{17c}$$

$$p_k \geq 0, \, \forall k \in \mathcal{K}, \quad \sum_{k \in \mathcal{K}} p_k = 1, \tag{17d}$$

where each $z^k \in \mathbb{R}^n$, where $P_{\mathbb{Z}} \in \mathbb{R}^{n_1 \times n}$ is the projection matrix that retrieves the n_1 first elements of a vector in $\mathbb{R}^n$, i.e., $P_{\mathbb{Z}} := [I \ 0]$, and finally where $h'(p_k, z^k, \xi)$ denotes the perspective of the recourse function $h(x, \xi)$, i.e.,

$$h'(p_k, z^k, \xi) := \min_{y'} \ c_2^{\mathsf{T}} y'$$

$$\text{s.t.} \ Ay' \geq W(\xi) z^k + b p_k.$$

In particular, both problems achieve the same optimal value and an optimal randomized strategy F_x^* for (2) *is supported on the collection of points $\{z^{k*}/p_k^*\}_{k \in \mathcal{K} : p_k^* \neq 0}$ with respective probabilities $\{p_k^*\}_{k \in \mathcal{K} : p_k^* \neq 0}$.*

Following similar steps as in Section 5.3, one can obtain the following large-scale convex optimization problem when employing the Wasserstein ambiguity set $\mathcal{D}(\widehat{F}_\xi, \epsilon)$:

$$\underset{p \geq 0, \{z^k\}_{k \in \mathcal{K}}, \lambda \geq 0, \{t_\omega\}_{\omega \in \Omega}}{\text{minimize}} \quad \sum_{k \in \mathcal{K}} c_1^\mathsf{T} z^k + \lambda \epsilon + \frac{1}{|\Omega|} \sum_{\omega \in \Omega} t_\omega$$

$$\text{subject to} \quad \sum_{k \in \mathcal{K}} h'(p_k, z^k, \xi^{h_\omega}) - \zeta^{h_\omega} \lambda \leq t_\omega \quad \forall h_\omega \in \mathcal{H}_\omega, \ \omega \in \Omega$$

$$C_x z^k \leq d_x p_k \qquad\qquad \forall k \in \mathcal{K}$$

$$P_{\mathbb{Z}} z^k = \bar{x}_1^k p_k \qquad\qquad \forall k \in \mathcal{K}$$

$$\sum_{k \in \mathcal{K}} p_k = 1.$$

We can proceed as before meaning that we start with two sets $\mathcal{K}' \subset \mathcal{K}$ and $\mathcal{H}' \subset \mathcal{H}$ and progressively identify which indexes to add to both $\mathcal{K}'$ and $\mathcal{H}'$. The so-called primal problem now takes the shape of

$$[\texttt{Primal}'(\mathcal{K}', \mathcal{H}')] : \quad \underset{p \geq 0, \{z^k\}_{k \in \mathcal{K}'}, \lambda \geq 0, \{t_\omega\}_{\omega \in \Omega}}{\text{minimize}} \quad \sum_{k \in \mathcal{K}'} c_1^\mathsf{T} z^k + \lambda \epsilon + \frac{1}{|\Omega|} \sum_{\omega \in \Omega} t_\omega$$

$$\text{subject to} \quad \sum_{k \in \mathcal{K}'} h'(p_k, z^k, \xi^{h_\omega}) - \zeta^{h_\omega} \lambda \leq t_\omega \quad \forall h_\omega \in \mathcal{H}_\omega, \ \omega \in \Omega$$

$$C_x z^k \leq d_x p_k \qquad\qquad \forall k \in \mathcal{K}'$$

$$P_{\mathbb{Z}} z^k = \bar{x}_1^k p_k \qquad\qquad \forall k \in \mathcal{K}'$$

$$\sum_{k \in \mathcal{K}'} p_k = 1,$$

which in its restricted form can be reformulated as an LP that integrates the recourse variables:

$$\underset{\substack{p \geq 0, \lambda \geq 0, \{t_\omega\}_{\omega \in \Omega}, \{z^k\}_{k \in \mathcal{K}'} \\ \{y_{k,h_\omega}\}_{k \in \mathcal{K}', h_\omega \in \mathcal{H}'_\omega, \omega \in \Omega}}}{\text{minimize}} \quad \sum_{k \in \mathcal{K}'} c_1^\mathsf{T} z^k + \lambda \epsilon + \frac{1}{|\Omega|} \sum_{\omega \in \Omega} t_\omega$$

$$\text{subject to} \quad \sum_{k \in \mathcal{K}'} c_2^\mathsf{T} y_{k,h_\omega} - \zeta^{h_\omega} \lambda \leq t_\omega \quad \forall h_\omega \in \mathcal{H}'_\omega, \ \omega \in \Omega$$

$$A y_{k,h_\omega} \geq W(\xi_{h_\omega}) z^k + b p_k \quad \forall k \in \mathcal{K}', \ h_\omega \in \mathcal{H}'_\omega, \ \omega \in \Omega$$

$$C_x z^k \leq d_x p_k \qquad\qquad \forall k \in \mathcal{K}'$$

$$P_{\mathbb{Z}} z^k = \bar{x}_1^k p_k \qquad\qquad \forall k \in \mathcal{K}'$$

$$\sum_{k \in \mathcal{K}'} p_k = 1,$$

with a subproblem

$$[\texttt{SP1}'_\omega] : \quad \underset{(\xi, \zeta) \in \Xi'_\omega}{\text{maximize}} \sum_{k \in \mathcal{K}'} p_k^* h(z^{k*}/p_k^*, \xi) - \lambda^* \zeta \,,$$

which can again be cast as the mixed-integer linear program presented in (15).

An additional challenge arises when attempting to identify a new support point $k \in \mathcal{K}$ to add to $\mathcal{K}'$. First, one needs to show (see Appendix C) that the dual problem takes the form:

$$[\texttt{Dual}'(\mathcal{K}, \mathcal{H}')] : \quad \underset{w, q \geq 0}{\text{maximize}} \quad w \tag{18a}$$

$$\text{subject to} \quad w \leq \min_{x^k \in \mathcal{X}_k} c_1^\mathsf{T} x^k + \sum_{\omega \in \Omega} \sum_{h_\omega \in \mathcal{H}'_\omega} h(x^k, \xi^{h_\omega}) q_{h_\omega} \quad \forall k \in \mathcal{K} \quad (p_k) \tag{18b}$$

$$\sum_{\omega \in \Omega} \sum_{h_\omega \in \mathcal{H}'_\omega} \zeta^{h_\omega} q_{h_\omega} \leq \epsilon \qquad\qquad (\lambda) \tag{18c}$$

$$\sum_{h_\omega \in \mathcal{H}'_\omega} q_{h_\omega} = \frac{1}{|\Omega|} \qquad\qquad \forall \omega \in \Omega \quad (t_\omega), \tag{18d}$$

where $\mathcal{X}_k := \{x \in \mathbb{R}^n \,|\, C_x x \leq d_x,\ P_{\mathbb{Z}} x = \bar{x}_1^k\}$. Hence, the restricted dual problem with $\mathcal{K}' \subset \mathcal{K}$ can be solved as a LP by replacing constraint (18b) by the dual problem associated to :

$$\underset{x \in \mathcal{X}_k, \{y_{h_\omega}^\omega\}_{h_\omega \in \mathcal{H}_\omega, \omega \in \Omega}}{\text{minimize}} \qquad c_1^\mathsf{T} x + \sum_{\omega \in \Omega} \sum_{h_\omega \in \mathcal{H}'_\omega} q_{h_\omega} c_2^\mathsf{T} y_{h_\omega}^\omega$$

$$\text{subject to} \qquad A y_{h_\omega}^\omega \geq W(\xi^{h_\omega}) x + b \qquad \forall h_\omega \in \mathcal{H}'_\omega,\ \omega \in \Omega \,.$$

On the other hand, the subproblem used in Step 3 to add a new support point in $\mathcal{K}'$ takes the same form as SP2.

6　Application on distributionally robust integer problems

In this section, we apply the column generation algorithms presented in Section 5 to solve the RSP that emerges in three classical application of discrete optimization: the assignment problem (as an example of problems with integer polyhedron feasibility sets), the uncapacitated facility location problem (as an example of single-stage problems), and the capacitated facility location problem (as an example of two-stage problems). To simplify exposition, we again focus on the case when $\mathcal{D}$ is the Wasserstein ambiguity set defined in Section 5.2 with a l_1-norm Wassrstein ball and a polyhedral support set $\Xi := \{\xi \,|\, C_\xi \xi \leq d_\xi\}$.

6.1　Distributionally robust assignment problem

The assignment problem aims to find the minimum weighted matching over a bipartite graph. It belongs to a class referred to as minimum-cost network flow (MCNF) problems. It is well-known that the constraint matrix of this class of problems is *totally unimodular*, meaning that, under mild conditions, the relaxed feasible set is an integer polyhedron, hence $\bar{\mathcal{X}} = \mathcal{C}(\mathcal{X})$. For more details about MCNF problems and total unimodularity, the reader is referred to Ahuja et al. (1993).

The distributionally robust assignment problem (DRAP) can be stated as follows:

$$\underset{x \in \mathcal{X}_{\text{AP}}}{\text{minimize}} \ \underset{F_\xi \in \mathcal{D}}{\sup} \ \mathbb{E}_{\xi \sim F_\xi} \left[\sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \xi_{ij} x_{ij} \right] \tag{19}$$

where

$$\mathcal{X}_{\text{AP}} := \left\{ x \in \{0,\,1\}^{|\mathcal{I}| \times |\mathcal{J}|} \ \middle| \ \begin{array}{l} \sum_{j \in \mathcal{J}} x_{ij} = 1,\, \forall\, i \in \mathcal{I} \\ \sum_{i \in \mathcal{I}} x_{ij} = 1,\, \forall\, j \in \mathcal{J} \end{array} \right\}.$$

In this formulation, $\mathcal{I}$ and $\mathcal{J}$ are sets of demand and supply points, respectively, each x_{ij} is a binary assignment variable and ξ_{ij} is an uncertain assignment cost.

Proposition 5 *For the DRAP presented in Equation* (19)*, the value of randomized solutions is equal to*

$$VRS = \min_{x \in \mathcal{X}_{AP}} \delta^* \left(x | \mathcal{U} \right) - \min_{x' \in \mathcal{X}'_{AP}} \delta^* \left(x | \mathcal{U} \right) ,$$

where

$$\mathcal{X}'_{AP} := \left\{ x \in [0,\,1]^{|\mathcal{I}| \times |\mathcal{J}|} \ \middle| \ \begin{array}{l} \sum_{j \in \mathcal{J}} x_{ij} = 1,\, \forall\, i \in \mathcal{I} \\ \sum_{i \in \mathcal{I}} x_{ij} = 1,\, \forall\, j \in \mathcal{J} \end{array} \right\}; \tag{20}$$

while $\delta^*(v|\mathcal{U}) := \sup_{\mu \in \mathcal{U}} \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} v_{ij} \mu_{ij}$ *is the support function of the set:*

$$\mathcal{U} := \mathcal{C} \left(\{ \mu \in \mathbb{R}^{|\mathcal{I}| \times |\mathcal{J}|} \,\middle|\, \exists\, F_\xi \in \mathcal{D},\ \mu_{ij} = \mathbb{E}_{\xi \sim F_\xi}[\xi_{ij}],\ \forall\, i \in \mathcal{I},\ \forall\, j \in \mathcal{J} \} \right).$$

Theorem 5 enables us to find the value of the randomized solution simply by solving a continuous relaxation of DRAP. Note that while an integer solution can always be obtained by solving the continuous relaxation of the deterministic assignment problem, this is not true for DRAP since the worst-case expected cost function is convex with respect to x (instead of being linear). Let, for instance, the distributional set $\mathcal{D}$ be a Wasserstein set similar to the one presented in Section 5.2. Using Corollary 5.1 in Mohajerin Esfahani and Kuhn (2017), one can reformulate DRAP as the mixed-integer program

$$\underset{x \in \mathcal{X}_{\mathrm{AP}}, \lambda, \{t_\omega\}_{\omega \in \Omega}, \{\nu_\omega\}_{\omega \in \Omega}}{\text{minimize}} \quad \lambda \epsilon + \frac{1}{|\Omega|} \sum_{\omega \in \Omega} t_\omega \tag{21a}$$

$$\text{subject to} \quad \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \widehat{\xi}_{ij\omega} x_{ij} + \left(d - C \widehat{\xi}_\omega \right)^\mathsf{T} \nu_\omega \leq t_\omega, \quad \forall \omega \in \Omega \tag{21b}$$

$$\| C^\mathsf{T} \nu_\omega - x \|_\infty \leq \lambda \qquad\qquad \forall \omega \in \Omega \tag{21c}$$

$$\nu_\omega \geq 0 \qquad\qquad \forall \omega \in \Omega, \tag{21d}$$

where each $\nu_\omega \in \mathbb{R}^{s_\xi}$ and where $\widehat{\xi}_\omega$ is short for $[\widehat{\xi}_{11\omega} \ \widehat{\xi}_{21\omega} \ldots \widehat{\xi}_{|\mathcal{I}|1\omega} \ \ \widehat{\xi}_{12\omega} \ldots \widehat{\xi}_{|\mathcal{I}|2\omega} \ldots \widehat{\xi}_{|\mathcal{I}||\mathcal{J}|\omega}]^\mathsf{T}$.

We can also use Proposition 5 to find the optimal value of the randomized strategy problem. Let $x^* = \arg\min_{x \in \mathcal{X}'_{\mathrm{AP}}} \delta^*(x|\mathcal{U})$ be the optimal solution of problem (21) with $\mathcal{X}_{\mathrm{AP}}$ replaced with $\mathcal{X}'_{\mathrm{AP}}$. As discussed in Section 4, the optimal randomized strategy can then be found by solving the problem:

$$\underset{p \geq 0}{\text{minimize}} \quad \| x^* - \sum_{k \in \mathcal{K}} p_k \bar{x}^k \|_1$$

$$\text{subject to} \quad \sum_{k \in \mathcal{K}} p_k = 1.$$

This problem can be rewritten as

$$\underset{p \geq 0, \theta}{\text{minimize}} \quad \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \theta_{ij} \tag{22a}$$

$$\text{subject to} \quad \theta_{ij} + \sum_{k \in \mathcal{K}} p_k \bar{x}^k_{ij} \geq x^*_{ij} \quad \forall i \in \mathcal{I}, j \in \mathcal{J} \quad (\psi^+_{ij}) \tag{22b}$$

$$\theta_{ij} - \sum_{k \in \mathcal{K}} p_k \bar{x}^k_{ij} \geq -x^*_{ij} \quad \forall i \in \mathcal{I}, j \in \mathcal{J} \quad (\psi^-_{ij}) \tag{22c}$$

$$\sum_{k \in \mathcal{K}} p_k = 1 \qquad\qquad (\phi), \tag{22d}$$

where $\theta \in \mathbb{R}^{|\mathcal{I}| \times |\mathcal{J}|}$. Since the number of extreme points is extremely large, we generate the elements of $\mathcal{K}$ and add them progressively, as needed, using a column generation approach. To do so, we first write the dual problem as

$$\underset{\psi^+ \geq 0, \psi^- \geq 0, \phi}{\text{maximize}} \quad \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} (\psi^+_{ij} - \psi^-_{ij}) x^*_{ij} - \phi$$

$$\text{subject to} \quad \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \left(\psi^+_{ij} - \psi^-_{ij} \right) \bar{x}^k_{ij} \leq \phi, \quad \forall k \in \mathcal{K} \qquad (p_k)$$

$$\psi^+_{ij} + \psi^-_{ij} \leq 1 \qquad\qquad \forall i \in \mathcal{I}, \forall j \in \mathcal{J} \quad (\theta_{ij}),$$

where $\psi^+ \in \mathbb{R}^{|\mathcal{I}| \times |\mathcal{J}|}$, $\psi^- \in \mathbb{R}^{|\mathcal{I}| \times |\mathcal{J}|}$, and $\phi \in \mathbb{R}$. We then choose a subset $\mathcal{K}' \subset \mathcal{K}$ (which can initially include only the index of the deterministic problem's solution) and solve the restricted version of problem (22) to obtain an upper bound. Then, the dual variables ψ^+ and ψ^- are used in the subproblem $\max_{x \in \mathcal{X}_{\mathrm{AP}}} (\psi^+ - \psi^-)^\mathsf{T} x$ to generate a new extreme point $x^{k'}$. The set $\mathcal{K}'$ is updated by adding the index of the new extreme point and the restricted master problem is resolved to obtain a new upper bound and new values for the dual variables. The algorithm iterates between the restricted master problem and the subproblem until the solution of the restricted master problem's objective value becomes smaller than some tolerance $\varepsilon \geq 0$.

6.2 Distributionally robust uncapacitated facility location

We now revisit the DRUFLP problem presented in Section 3 in its more general form:

$$\underset{(x,y)\in\mathcal{X}_{\text{UFLP}}}{\text{minimize}}\ \underset{F_\xi\in\mathcal{D}}{\text{maximize}}\ \mathbb{E}_{\xi\sim F_\xi}\left[\sum_{j\in\mathcal{J}}f_jx_j+\sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}}\xi_ic_{ij}y_{ij}\right],$$

where each x_i denotes the decision to open a facility at location i, while each y_{ij} denotes the decision to serve the demand at j using facility i, in particular

$$\mathcal{X}_{\text{UFLP}}:=\left\{(x,y)\in\{0,1\}^{|\mathcal{I}|}\times\{0,1\}^{|\mathcal{I}|\times|\mathcal{J}|}\ \middle|\ \begin{array}{l}\sum_{j\in\mathcal{J}}y_{ij}=1\,,\,\forall\,i\in\mathcal{I}\\ y_{ij}\le x_j\,,\,\forall\,i\in\mathcal{I},j\in\mathcal{J}\end{array}\right\}.$$

The uncertain parameters in this problem are ξ_i, *i.e.*, the demand at each customer location $i\in\mathcal{I}$ and that must be served by one of the open facilities. We focus on the classical single-stage, single assignment version of the UFLP, in which both the location of facilities and the assignment of demand to them are determined before the demand is revealed and each demand is fully assigned to a single open facility.

With a randomized strategy, the problem can be stated as:

$$\underset{p_k\ge0,\,\sum\limits_{k\in\mathcal{K}}p_k=1}{\text{minimize}}\ \underset{F_\xi\in\mathcal{D}}{\sup}\ \mathbb{E}_{\xi\sim F_\xi}\left[\sum_{k\in\mathcal{K}}\left(\sum_{j\in\mathcal{J}}f_j\bar{x}_j^k+\sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}}\xi_ic_{ij}\bar{y}_{ij}^k\right)p_k\right],\tag{23}$$

where $\mathcal{K}$ is the index set of feasible (x,y) pairs in $\mathcal{X}_{\text{UFLP}}$. One can implement the single-stage variant of the column generation algorithm described in Section 5.4 to solve this problem. In every iteration, we solve the master problem (16) with a partial set $\mathcal{K}'\subseteq\mathcal{K}$, and with constraint (16c) replaced with

$$C_\xi^\mathsf{T}\alpha_\omega+\psi_\omega^+-\psi_\omega^-=\sum_{k\in\mathcal{K}}\sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}}c_{ij}\bar{y}_{ij}^k\,,\,\forall\,\omega\in\Omega\qquad(\xi_\omega).\tag{24}$$

Next, the optimal solution $\{\xi_\omega^*\}_{\omega\in\Omega}$ for the dual variables corresponding to constraint (24), are used in the subproblem SP2′:

$$\underset{(x,y)\in\mathcal{X}_{\text{UFLP}}}{\text{minimize}}\ \sum_{j\in\mathcal{J}}f_jx_j+\frac{1}{|\Omega|}\sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}}\sum_{\omega\in\Omega}\xi_{i\omega}^*c_{ij}y_{ij},$$

to generate a new solution $(\bar{x}^{k'},\bar{y}^{k'})$ and update the lower bound (LB). The set $\mathcal{K}'$ is then updated by including k' and the master problem is resolved to update $\{\xi_\omega^*\}_{\omega\in\Omega}$ and obtain an upper bound (UB). The algorithm terminates when $UB-LB\le\varepsilon$.

6.3 Distributionally robust capacitated facility location problem

We formulate the distributionally-robust capacitated facility location problem (DRCFLP) with randomization as a two-stage stochastic program with distributional ambiguity. Similar to the DRUFLP presented in the previous section, we consider uncertainty in the demand quantity ξ_i. The *here-and-now* decision is a potentially randomized set of facility locations parametrized using a probability vector $p\in\mathbb{R}^{|\mathcal{K}|}$ where $\mathcal{K}$ captures the set of indices for all members of $\{0,1\}^{|\mathcal{J}|}$. With that, the DRCFLP can be stated as follows:

$$\underset{p:p\ge0,\,\sum\limits_{k\in\mathcal{K}}p_k=1}{\min}\ \sum_{k\in\mathcal{K}}\sum_{j\in\mathcal{J}}f_j\bar{x}_j^kp_k+\underset{F_\xi\in\mathcal{D}}{\sup}\ \mathbb{E}_{\xi\sim F_\xi}\left[\sum_{k\in\mathcal{K}}h(\bar{x}^k,\xi)p_k\right],\tag{25}$$

where

$$h(x, \xi) := \min_{z \geq 0} \quad \sum_{j \in \mathcal{J}} \sum_{i \in \mathcal{I}} c_{ij} z_{ij}$$

$$\text{s.t.} \quad \sum_{j \in \mathcal{J}} z_{ij} = \xi_i \qquad \forall i \in \mathcal{I}$$

$$\sum_{i \in \mathcal{I}} z_{ij} \leq v_j x_j \quad \forall j \in \mathcal{J},$$

is the second-stage (*recourse*) problem that is solved to find the assignment of demand to opened facilities once the uncertain demand quantities become known. Unlike the classical formulation of the CFLP that uses the variable $y_{ij} \in [0, 1]$ to denote the *fraction* of customer i's demand served by facility j (see, *e.g.,* Fernández and Landete (2015)), we equivalently use $z_{ij} = \xi_i y_{ij}$ to denote the *quantity* of customer i's demand served by facility j, as the recourse decision variable. This choice of the allocation variables enables us to write the second-stage problem in the form presented in Equation (6).

The DRCFLP formulation provided in Equations (25) and (26) can be seen as a special case of the problem described by Equations (5) and (6), respectively. Therefore, the two-layer column generation algorithm presented in Section 5.3 can be used to solve it. The implementation details for solving a DRCFLP with a Wasserstein ambiguity set are provided in Appendix D.

7 Numerical results

We conducted a series of numerical experiments to assess quality of the bounds proposed in Section 4 and of the numerical performance of the solution algorithms presented in Section 5. All algorithms were implemented using Matlab R2017a, and Gurobi 5.7.1 was called to solve the master and subproblems. All tests were run on a personal computer with an Intel Core i-7 7700 3.6 GHz processor and 16 GB of RAM. For all problems, we used a sample set of 10 observations (*i.e.,* $|\Omega| = 10$) selected uniformly at random from the set Ξ to construct the empirical distribution $\widehat{F}_\xi$. The ambiguity set was then defined as a Wasserstein ball of radius ϵ around the empirical distribution.

7.1 Experiments with the DRAP

We experimented with 10 random instances of size $|\mathcal{I}| = |\mathcal{J}| = 100$. The support set for ξ was a hypercube defined as

$$\Xi := \left\{ \xi \in \mathbb{R}_+^{|\mathcal{I}| \times |\mathcal{J}|} : \xi_{ij}^{nom}(1 - \Delta_{ij}) \leq \xi_{ij} \leq \xi_{ij}^{nom}(1 + \Delta_{ij}) \right\},$$

where ξ^{nom} was a nominal cost vector drawn uniformly at random in $[10, 20]^{|\mathcal{I}| \times |\mathcal{J}|}$ while, for each $i \in \mathcal{I}$ and $j \in \mathcal{J}$, the relative maximum deviation Δ_{ij} was drawn uniformly and independently at random from $[0.5, 1]$. We studied for each instance a range of different values of $\epsilon \in [0, 25000]$. For each instance, the DRAP with a deterministic strategy was first solved to obtain v_d. We then solved a continuous relaxation of the deterministic strategy problem and used its solution to find an optimal randomized strategy using the algorithm outlined in Section 6.1. Figure 1a presents statistics about the relative improvement achieved by a randomized strategy compared to a deterministic strategy on 10 problem instances and as a function of ϵ. Both the average and the maximum improvements observed in the test instances are reported. Note that the bound obtained from Theorem 1 is exact for this application thus $\text{VRS} = \widehat{\text{VRS}}$

Looking at Figure 1a, one should notice that when the ambiguity set contains only the empirical distribution (*i.e.,* $\epsilon = 0$), there is no value for randomization. This is due to the fact that the problem reduces to a simple expected value minimization problem (with known distribution) which is known to be randomization-proof. On the other hand, when ϵ becomes very large, the adversary can place all the probability mass at any vertex of Ξ. Given that a box support set is used, the

(a)   (b)
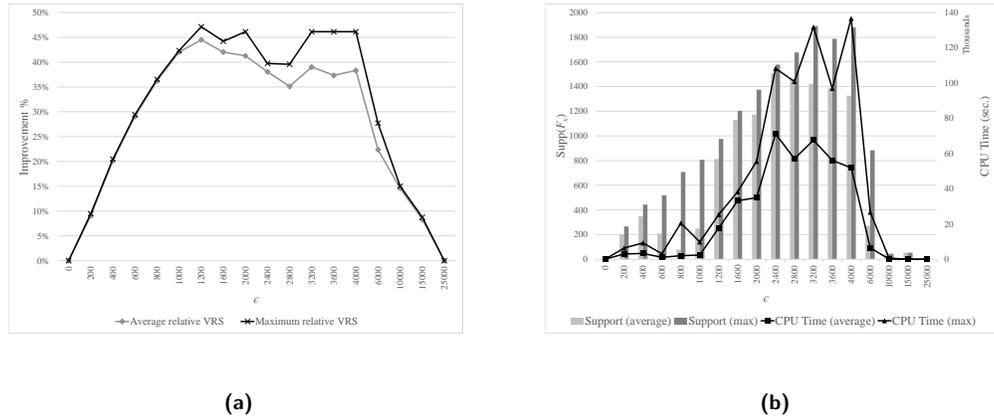
Figure 1: (a) Average and maximum improvement achieved by randomization. (b) Average and maximum support size of the optimal randomized strategy and computational time. Both based on 10 DRAP instances.

problem reduces to a deterministic assignment problem with $\xi = \xi^{nom}(1 + \Delta_{ij})$, and randomization becomes ineffective for the same reason. Between these two extreme cases, we can confirm that employing a randomized strategy can lead to significant reduction in worst-case expected assignment cost. For example, at $\epsilon = 1200$ an average improvement of 44.48% was achieved, whereas the maximum improvement observed in the 10 test instances was 47.14%. Intuitively, this might be explained by the fact that randomization allows the decision maker to mitigate his ambiguity aversion by diversifying the types of cost c_{ij} to which his expected cost is sensitive. In particular, while a deterministic strategy's expected cost only depends on the quality of n different cost values, a randomized one has the potential of making the expected cost depend on the quality of all n^2 terms in the cost matrix c, effectively distributing the risks accordingly.

Figure 1b presents statistics about the size of the support set of the optimal randomized strategy in the DRAP instances tested and the computational time of the solution algorithm. One can notice that in order to reap the full benefit of randomization, the optimal strategy randomizes among a number of feasible solutions (*i.e.,* assignment plans) that ranges between 5 and 1890 plans (excluding the case of $\epsilon = 0$, where randomization has no value). Although the support set size seems quite large, it is still well below the theoretical bound of 10,001 plans obtained from Proposition 1. The average computational time for all instances was 22,682 seconds and on average produces optimal randomized strategies supported on 649 assignment plans. The longest computational time observed was 136,651 seconds and in this case produced a strategy supported on 1879 assignment plans. In comparison, solving the deterministic strategy DRAP took an average of 33,268 seconds, with some instances taking more than 48 hours to solve on Gurobi. We note that while a near-optimal solution is approached considerably fast, only little progress per iteration is made close to the optimum, a well-known phenomenon in column generation algorithms known as the *tailing-off effect* (see Gilmore and Gomory (1961), for instance). Therefore, despite the apparently long computational time and large support size at optimality, most (typically > 90%) of the improvement was achieved in the first few iterations. Hence, the decision maker can terminate the algorithm prematurely and still obtain a feasible randomized strategy that considerably outperforms the deterministic strategy.

We experimented with 10 random instances of size $|\mathcal{I}| = |\mathcal{J}| = 300$. The coordinates of demand points (which are also the potential facility locations) were selected uniformly at random on a unit square, and we used the Euclidean distance between any two points as the unit shipping cost c_{ij}. The set-up cost was $f_j = 10$ for all potential locations while the uncertain demand was supported on a hypercube defined as

$$\Xi := \left\{ \xi \in \mathbb{R}_+^{|I|} : \xi_i^{nom}(1 - \Delta_i) \leq \xi_i \leq \xi_i^{nom}(1 + \Delta_i) \right\},$$

where each nominal demand $\xi_i^{nom} \in [10, 20]$ and each maximum relative deviation $\Delta_i \in [0.5, 1]$, were uniformly drawn at random. We again studied the performances under a range of ϵ values between 0 and 4000. For each problem instance, we solved the DRUFLP without and with randomization to obtain v_d and v_r, respectively, and a relaxed version of the deterministic problem, as prescribed by Theorem 1, to compute the bound $\widehat{\text{VRS}}$.

Figure 2a presents statistics about the relative $\widehat{\text{VRS}}$ bound and about the actual relative improvements achieved by the best randomized strategy in a set of 10 problem instances and under different amount of distributional ambiguity, parameterized by $\epsilon \in [0, 4000]$. Similar to the case of the DRAP, we observe again that randomization does not lead to a reduction in worst-case expected cost when $\epsilon = 0$ or when becomes too large. Unlike for the case of DRAP, however, we observe here that the improvement obtained from randomization remains relatively small for this set of 10 instances. In particular, it never exceeds 1.87% and peaks with an average improvement of 1.37% at $\epsilon = 1600$. Interestingly, it appears that for this set of problem instances, one does not actually need to solve the DRUFLP with randomization to draw this conclusion. Indeed, the $\widehat{\text{VRS}}$ bound proposed in Theorem 1 can be computed much more efficiently and already confirms that the improvement is below 2% for all problem instances and values of ϵ. The $\widehat{\text{VRS}}$ bound was even able in 5 problem instances to recognize that at $\epsilon = 0$, there was no possible improvement for randomized strategies. Even in the case of the 5 other problem instances, the identified potential of improvement was nearly inexistant (always below 0.05%). This evidence supports an observation made by Morris (1978) that solving a linear relaxation of the UFLP typically leads to identifying an integer solution.
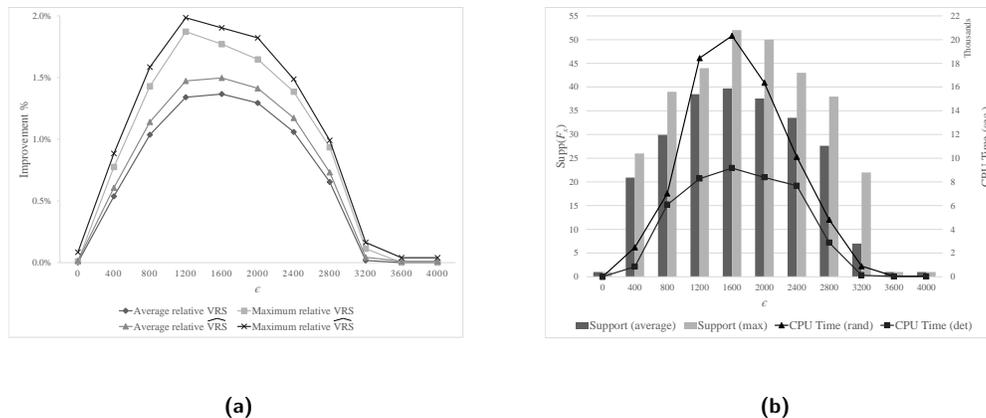
7.2   Experiments with the DRUFLP

(a)   (b)

Figure 2: (a) Average and maximum $\widehat{\text{VRS}}$ bound and actual improvement achieved by randomization. (b) Average and maximum support size of the optimal randomized strategy and computational time. Both based on 10 DRUFLP instances.

In Figure 2b, we report on the average and maximum size of the support of the optimal randomized strategies obtained for the different problem instances and different values of ϵ, as well as computational times for the column generation algorithm. One can first notice that the size of the optimal support set is somewhat proportional to the extent of relative improvement. In fact, the largest average support set size, of 39.7 plans, was reached under $\epsilon = 1600$, which nearly coincides with the peak performance improvement reached at $\epsilon = 1200$. One can also confirm that in all cases, the size of the optimal support for the randomized strategy is always far below the theoretical limit of 301 (see Proposition 1). The average computational time for the DRUFLP with randomization was 7321 seconds whereas the largest computational time among all instances was 44,000 seconds. The *tailing-off effect* was also observed in this case. In comparison, the deterministic strategy problem took on average 3951 seconds to solve.

7.3   Experiments with the DRCFLP

We studied the value of randomization in 10 randomly generated instances of the DRCFLP of size $|\mathcal{I}| = |\mathcal{J}| = 20$. For each of these instances, the setup costs f, transportation costs c, and the support set Ξ were constructed exactly as in Section 7.2. We assumed that all facilities have the same capacity of $v = \frac{r \sum_{i \in \mathcal{I}} \xi_i^{nom}}{|\mathcal{J}|}$, where r controls how scarce the capacity is (*i.e.*, larger r implies less scarcity). Figures 3a and 3b present statistics about the relative $\widehat{\mathrm{VRS}}$ bound and the actual relative improvements achieved by the optimal randomized strategy in a set of 10 problem instances under different levels of distributional ambiguity, parameterized by $\epsilon \in [0, 300]$, and with $r \in \{3, 5\}$, respectively. Looking at these figures, one immediately notices that the $\widehat{\mathrm{VRS}}$ bound is, in general, a poor indicator of the maximum improvement that can be achieved by randomization for this class of problem. The quality of this bound also seems to degrade as the capacity becomes less scarce. On the other hand, the actual improvement achieved by randomized strategies in this class of problems appears to be more significant than in the DRUFLP. It reaches a maximum of 3.77% in a problem instance with $\epsilon = 80$ and $r = 5$. Otherwise, the average relative improvement was at 0.67% and 1.00% for problems with $r = 3$ and 5 respectively, when computed over all problems instances with $\epsilon \in [10, 300]$, and peaked at 2.05% for problems with $\epsilon = 120$ and $r = 5$.

(a)

(b)

(c)

(d)
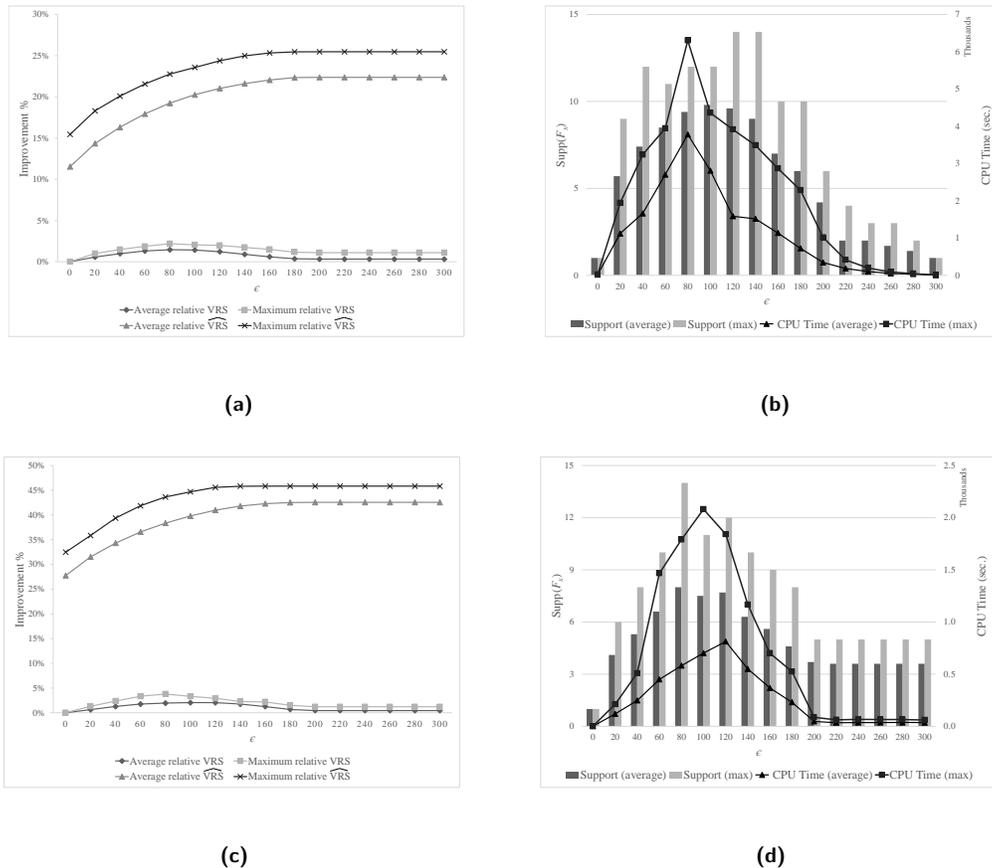
Figure 3: (a,b) Improvements due to randomization and relaxation with $r = 3$ and 5, respectively. (c,d) Optimal solution support size and computational time with $r = 3$ and 5, respectively. All based on 10 DRCFLP test instances.

Figures 3c and 3d show the average and maximum size of the support of the optimal randomized strategies and the computational time for the column generation algorithm obtained for the different problem instances and values of ϵ and r. Among problem instances with $\epsilon \in \{20, 40, \ldots, 300\}$, the average optimal support set sizes were 5.65 and 5.16 for $r = 3$ and 5, respectively, while the maximum

optimal support set size was 14 in both cases. This seems to indicate that the structure of optimal randomized strategies become slightly simpler as the capacity scarcity is reduced and that perhaps in practice the optimal support size still remains comparable to n although Proposition 1 does not apply for this class of problems. The effect of ϵ on the value of randomization is similar to what was observed in the experiments with the DRAP and the DRUFLP, namely that the value of randomization peaks at mid-range values for ϵ while it degrades to zero as ϵ gets closer to zero or grows to infinity. The average computational times needed to solve the proposed two-layer column generation algorithm for problem instances with $r = 3$ and 5 was 1111 and 268 seconds, respectively, whereas the largest computational time in all tested instances was less than two hours. In comparison, the average computational times for the deterministic strategy problems using Gurobi were 3317 and 1870 seconds, respectively. These results clearly show that the proposed algorithm can handle reasonably sized problems quite effectively.

8   Conclusions and future directions

In this paper, we investigated the value of randomization in a general class of distributionally robust two-stage linear program with mixed-integer first stage decisions. We established for the first time how the value of randomization in problems where the cost function and risk measures are both convex can be bounded by the difference between the optimal value of the nominal DRO problem and of its convex relaxation. We further demonstrated that if the decision maker is AARN, then a finitely supported optimal randomized strategy always exists. This allowed us to design an efficient two-layer column generation algorithm for identifying this compact optimal support and its associated probability weights in two-stage problems where uncertainty appears in the right-hand side of the constraint sets. Our numerical experiments provided empirical evidence that 1) the proposed algorithm can address reasonably sized version of assignment problems, and both capacitated and uncapacitated facility location problems; 2) randomization is especially effective in applications, such as the assignment problem, where the vector of binary variables is constrained to be very sparse compared to the number of potential perturbations, *e.g.*, $O(\sqrt{n})$ non-zeroes compared to $O(n)$ perturbations in the assignment problem. The latter should imply that randomization can be especially beneficial in other problem classes that have this property such as shortest path, travelling salesman problems, etc.

From an algorithmic point of view, it should be possible to extend relatively easily the algorithm to two-stage problems with uncertainty in the objective function, or problems with worst-case expected utility objectives where the utility function is piecewise linear. On the other hand, significant additional development would need to be achieved in order to handle more general two-stage decision problems or risk functions such as value-at-risk, conditional value-at-risk, expectiles, etc. For instance, it remains open to establish whether a finitely supported optimal randomized strategy necessarily exists under more general conditions than the AARN attitude.

A   Proofs

A.1   Proof of Theorem 1

Our proof exploits the extension of Theorem 2 in Delage et al. (2018) which states that if the objective function is an ambiguity averse convex risk measure, which canonical form is as presented in Equation (1), and the set of all feasible random costs $\{h(x,\xi) : x \in \mathcal{X}\}$ is a convex set, then there is no benefit in adopting a randomized strategy. Indeed, from this we can conclude that

$$\min_{x \in \mathcal{X}'} \sup_{F_\xi \in \mathcal{D}} \rho_{\xi \sim F_\xi}[h(x,\xi)] \leq \min_{x \in \mathcal{C}(\mathcal{X})} \sup_{F_\xi \in \mathcal{D}} \rho_{\xi \sim F_\xi}[h(x,\xi)]$$

$$= \min_{g(\cdot) \in \mathcal{G}} \sup_{F_\xi \in \mathcal{D}} \rho_{\xi \sim F_\xi}[g(\xi)]$$

$$= \min_{F_g \in \Delta(\mathcal{G})} \sup_{F_\xi \in \mathcal{D}} \rho_{(G,\xi) \sim F_g \times F_\xi}[G(\xi)]$$

$$= \min_{F_g \in \Delta(\bar{\mathcal{G}})} \sup_{F_\xi \in \mathcal{D}} \rho_{(G,\xi) \sim F_g \times F_\xi}[G(\xi)]$$

$$= \min_{F_x \in \Delta(\mathcal{C}(\mathcal{X}))} \sup_{F_\xi \in \mathcal{D}} \rho_{(X,\xi) \sim F_x \times F_\xi}[h(X,\xi)]$$

$$\leq \min_{F_x \in \Delta(\mathcal{X})} \sup_{F_\xi \in \mathcal{D}} \rho_{(X,\xi) \sim F_x \times F_\xi}[h(X,\xi)] = v_r \,,$$

where

$$\mathcal{G} := \{g : \mathbb{R}^m \to \mathbb{R} \,|\, \exists\, x \in \mathcal{C}(\mathcal{X}),\, g(\xi) \geq h(x,\xi)\,\forall\xi\}$$

is a convex set of random variables, and where

$$\bar{\mathcal{G}} := \{g : \mathbb{R}^m \to \mathbb{R} \,|\, \exists\, x \in \mathcal{C}(\mathcal{X}),\, g(\xi) = h(x,\xi)\,\forall\xi\}\,.$$

The first inequality follows from the fact that $\mathcal{X}' \supseteq \mathcal{C}(\mathcal{X})$, *i.e.,* the convex hull of $\mathcal{X}$. The next four steps follow, respectively, from the fact that $\rho(\cdot)$ is monotone, the fact that $\mathcal{G}$ is a convex set hence the extension of Theorem 2 in Delage et al. (2018) holds, again the fact that $\rho(\cdot)$ is monotone, and finally based on the definition of $\bar{\mathcal{G}}$. The last inequality follows from the fact that $\mathcal{C}(\mathcal{X}) \supseteq \mathcal{X}$.

To identify the special case where the bound is tight, we can proceed as follows:

$$\min_{x \in \mathcal{C}(\mathcal{X})} \sup_{F_\xi \in \mathcal{D}} \mathbb{E}_{\xi \sim F_\xi}[h(x,\xi)] \;=\; \min_{F_x \in \Delta(\mathcal{X})} \sup_{F_\xi \in \mathcal{D}} \mathbb{E}_{\xi \sim F_\xi}[h(\mathbb{E}_{X \sim F_x}[X],\xi)]$$

$$= \min_{F_x \in \Delta(\mathcal{X})} \sup_{F_\xi \in \mathcal{D}} \mathbb{E}_{(X,\xi) \sim F_x \times F_\xi}[h(X,\xi)]\,,$$

where the first step follows from the fact that $\mathcal{C}(\mathcal{X})$ is the convex hull of $\mathcal{X}$ and the second step from the linearity of the expectation operator. Hence, if F_x^* is such that $\mathbb{E}_{F_x^*}[X] \in \arg\min_{x \in \mathcal{C}(\mathcal{X})} \sup_{F_\xi \in \mathcal{D}} \mathbb{E}_{\xi \sim F_\xi}[h(x,\xi)]$, one can confirm that

$$\sup_{F_\xi \in \mathcal{D}} \mathbb{E}_{(X,\xi) \sim F_x^* \times F_\xi}[h(X,\xi)] = \sup_{F_\xi \in \mathcal{D}} \mathbb{E}_{\xi \sim F_\xi}[h(\mathbb{E}_{X \sim F_x^*}[X],\xi)]$$

$$= \min_{x \in \mathcal{C}(\mathcal{X})} \sup_{F_\xi \in \mathcal{D}} \mathbb{E}_{\xi \sim F_\xi}[h(x,\xi)]$$

$$= \min_{F_x \in \Delta(\mathcal{X})} \sup_{F_\xi \in \mathcal{D}} \mathbb{E}_{(X,\xi) \sim F_x \times F_\xi}[h(X,\xi)] = v_r\,.$$

This completes our proof.

A.2   Proof of Proposition 1

The result follows from Carathéodory's theorem (see, *e.g.,* Eckhoff (1993)), which states that any vector $x \in \mathcal{C}(\mathcal{X})$ can be represented as a convex combination, parameterized by $\{\theta_k\}_{k=1}^{n+1}$, of at most $n+1$ affinely independent vectors $\{\bar{x}^k\}_{k=1}^{n+1}$, with each $\bar{x}_k \in \mathcal{X}$. We can therefore establish that for any $x^* \in \mathcal{C}(\mathcal{X})$

$$x^* = \sum_{k=1}^{n+1} \theta_k x_k = \mathbb{E}_{X \sim \bar{F}_x^\theta}[X]$$

where $\bar{F}_x^\theta$ is defined as the discrete distribution that puts probabilities of $\theta_1, \theta_2, \ldots, \theta_{n+1}$ on the points $\bar{x}_1, \bar{x}_2, \ldots, \bar{x}_{n+1}$. Note that in problem (4), x^* is given and assumed to be a member of $\mathcal{C}(\mathcal{X})$, while in problem (2), $x^* \in \arg\min_{x \in \mathcal{C}(\mathcal{X})} \sup_{F_\xi \in \mathcal{D}} \mathbb{E}_{\xi \sim F_\xi}[h(x,\xi)]$ and the result follows from Theorem 1.

A.3 Proof of Proposition 2

We start with some definitions. Consider the set of feasible integer vectors

$$\mathcal{X}_{\mathbb{Z}} := \{x_1 \in \mathbb{Z}^{n_1} \,|\, \exists x_2 \in \mathbb{R}^{n_2}, \, [x_1^{\intercal} \; x_2^{\intercal}]^{\intercal} \in \mathcal{X}\} \,.$$

and for each $x_1 \in \mathcal{X}_{\mathbb{Z}}$, consider the "slice" of $\mathcal{X}$ defined as

$$\mathcal{X}_{\mathbb{R}}(x_1) := \{x_2 \in \mathbb{R}^{n_2} \,|\, [x_1^{\intercal} \; x_2^{\intercal}]^{\intercal} \in \mathcal{X}\} \,.$$

Since $\mathcal{X}$ is bounded, it is clear that $|\mathcal{X}_{\mathbb{Z}}|$ is finite hence $\mathcal{X}_{\mathbb{Z}} = \{\bar{x}_1^k\}_{k \in \mathcal{K}}$ where $\mathcal{K} = \{1, \ldots, |\mathcal{X}_{\mathbb{Z}}|\}$ is an index set for all members of $\mathcal{X}_{\mathbb{Z}}$. Furthermore, we have that, for all $x_1 \in \mathcal{X}_{\mathbb{Z}}$, the set $\mathcal{X}_{\mathbb{R}}(x_1)$ is convex.

The proof of Proposition 2 consists in showing that there exists an optimal discrete randomized strategy parametrized as $\{(p_k, x^k)\}_{k \in \mathcal{K}}$, where each p_k is the probability of drawing action x^k and where each $x^k = [\bar{x}_1^k \; x_2^k]$ for some $x_2^k \in \mathcal{X}_{\mathbb{R}}(\bar{x}_1^k)$. To do so, we consider an arbitrary optimal randomized strategy F_x^* for the RSP. Next, we can argue that

$$
\begin{aligned}
\min_{F_x \in \Delta(\mathcal{X})} \sup_{F_\xi \in \mathcal{D}} \mathbb{E}_{(X,\xi) \sim F_x \times F_\xi}[h(x,\xi)] \;&=\; \sup_{F_\xi \in \mathcal{D}} \mathbb{E}_{(X,\xi) \sim F_x^* \times F_\xi}[h(X,\xi)] \\
&= \sup_{F_\xi \in \mathcal{D}} \sum_{k \in \mathcal{K}} \mathbb{P}_{X \sim F_x^*}(P_{\mathbb{Z}}X = \bar{x}_1^k) \mathbb{E}_{(X,\xi) \sim F_x^* \times F_\xi}[h(X,\xi)|P_{\mathbb{Z}}X = \bar{x}_1^k] \\
&= \sup_{F_\xi \in \mathcal{D}} \sum_{k \in \mathcal{K}} \mathbb{P}_{X \sim F_x^*}(P_{\mathbb{Z}}X = \bar{x}_1^k) \mathbb{E}_{(X_2,\xi) \sim F_{x_2|\bar{x}_1^k}^* \times F_\xi}[h([\bar{x}_1^{k\intercal} \; X_2^{\intercal}]^{\intercal}, \xi)] \\
&\geq \sup_{F_\xi \in \mathcal{D}} \sum_{k \in \mathcal{K}} \mathbb{P}_{X \sim F_x^*}(P_{\mathbb{Z}}X = \bar{x}_1^k) \mathbb{E}_{\xi \sim F_\xi}[h([\bar{x}_1^{k\intercal} \; \mathbb{E}_{X_2 \sim F_{x_2|\bar{x}_1^k}^*}[X_2^{\intercal}]]^{\intercal}, \xi)] \\
&= \sup_{F_\xi \in \mathcal{D}} \sum_{k \in \mathcal{K}} p_k^* \mathbb{E}_{\xi \sim F_\xi}[h([\bar{x}_1^{k\intercal} \; \mu_2^{k*\intercal}]^{\intercal}, \xi)] \\
&= \sup_{F_\xi \in \mathcal{D}} \mathbb{E}_{(X,\xi) \sim \bar{F}_x^* \times F_\xi}[h(X,\xi)] \\
&\geq \min_{F_x \in \Delta(\mathcal{X})} \sup_{F_\xi \in \mathcal{D}} \mathbb{E}_{(X,\xi) \sim F_x \times F_\xi}[h(X,\xi)]
\end{aligned}
$$

where $F_{x_2|\bar{x}_1^k}^*$ denotes the conditional distribution of X_2 given that $X_1 = \bar{x}_1^k$, where $P_{\mathbb{Z}} \in \mathbb{R}^{n_1 \times n}$ is the projection matrix that retrieves the n_1 first elements of a vector in $\mathbb{R}^n$, *i.e.*, $P_{\mathbb{Z}} := [I \; 0]$, and where $p_k^* := \mathbb{P}_{X \sim F_x^*}(P_{\mathbb{Z}}X = \bar{x}_1^k)$ and $\mu_2^{k*} := \mathbb{E}_{X_2 \sim F_{x_2|\bar{x}_1^k}^*}[X_2]$ are the parametrization of a discrete distribution $\bar{F}_x^*$. We note that the first inequality in this derivation follows from Jensen's inequality. The second inequality follows from the fact that $\bar{F}_x^* \in \Delta(\mathcal{X})$ since each $[\bar{x}_1^{k\intercal} \; \mu_2^{k*\intercal}]^{\intercal} \in \mathcal{X}$ given that

$$\mu_2^{k*} = \mathbb{E}_{X_2 \sim F_{x_2|\bar{x}_1^k}^*}[X_2] \in \mathcal{X}_{\mathbb{R}}(\bar{x}_1^k) \,,$$

where we exploited the fact that $F_{x_2|\bar{x}_1^k}^*$ is supported on $\mathcal{X}_{\mathbb{R}}(\bar{x}_1^k)$ which is a convex set. This confirms that there always exists a discrete randomized strategy of the form proposed by the proposition that achieves the optimal value of the RSP.

A.4 Proof of Proposition 3

To obtain the proposed reformulation, one can follow similar steps as are proposed in Zeng and Zhao (2013) yet before doing so one must employ the so-called "dualized reformulation" trick proposed in de Ruiter et al. (2014) in order to prevent the introduction of a set of binary variables with size proportional to $|\mathcal{K}'|$. Specifically, we begin by dualizing the minimization problem that defines $h(\bar{x}^k, \xi)$ for each $k \in \mathcal{K}'$. We then reintegrate the equivalent maximization problem in SP1$_\omega$ and linearize the norm constraint to obtain the following bilinear optimization problem:

$$\underset{\xi, \zeta \geq 0, \delta \geq 0, \{\phi_k\}_{k \in \mathcal{K}'}}{\text{maximize}} \quad \sum_{k \in \mathcal{K}'} (W(\xi)\bar{x}^k + b)^{\intercal}\phi_k - \lambda^* \zeta \tag{27a}$$

$$\text{subject to} \qquad A^\intercal \phi_k = c_2 p_k^* \qquad\qquad \forall k \in \mathcal{K}' \tag{27b}$$

$$C_\xi \xi \le d_\xi \tag{27c}$$

$$e^\intercal \delta \le \zeta \tag{27d}$$

$$\zeta \le \bar{\zeta}_{\max} \tag{27e}$$

$$\xi - \widehat{\xi}_\omega \le \delta \tag{27f}$$

$$\widehat{\xi}_\omega - \xi \le \delta, \tag{27g}$$

$$\phi_k \ge 0 \qquad\qquad \forall k \in \mathcal{K}', \tag{27h}$$

where each $\phi_k \in \mathbb{R}^s$ is the dual vector associated to constraint (6b). Next, we employ the dualized reformulation method presented in de Ruiter et al. (2014). This is done by replacing the maximization problem:

$$\rho(\{\phi_k\}_{k \in \mathcal{K}'}) := \max_{\xi, \zeta \ge 0, \delta \ge 0} \ \sum_{k \in \mathcal{K}'} (W(\xi)\bar{x}^k + b)^\intercal \phi_k - \lambda^* \zeta \tag{28a}$$

$$\text{s.t.} \quad C_\xi \xi \le d_\xi \qquad\qquad (\alpha) \tag{28b}$$

$$e^\intercal \delta \le \zeta \qquad\qquad (\beta) \tag{28c}$$

$$\zeta \le \bar{\zeta}_{\max} \qquad\qquad (\gamma) \tag{28d}$$

$$\xi - \widehat{\xi}_\omega \le \delta \qquad\qquad (\psi^+) \tag{28e}$$

$$\widehat{\xi}_\omega - \xi \le \delta. \qquad\qquad (\psi^-), \tag{28f}$$

which is feasible when each $\widehat{\xi}_\omega \in \Xi$, with its equivalent dual problem:

$$\rho(\{\phi_k\}_{k \in \mathcal{K}'}) = \min_{\substack{\alpha \ge 0, \beta \ge 0, \gamma \ge 0 \\ \psi^+ \ge 0, \psi^- \ge 0}} \ \sum_{k \in \mathcal{K}'} (W_0 \bar{x}^k + b)^\intercal \phi_k + d_\xi^\intercal \alpha + \bar{\zeta}_{\max} \gamma + \widehat{\xi}_\omega^\intercal (\psi^+ - \psi^-) \tag{29a}$$

$$\text{s.t.} \quad \sum_{i=1}^m \left(\sum_{k \in \mathcal{K}'} \phi_k^T W_i \bar{x}^k \right) e_i = C_\xi^T \alpha + \psi^+ - \psi^- \tag{29b}$$

$$\beta \le \lambda^* + \gamma \tag{29c}$$

$$\psi^+ + \psi^- \le \beta \tag{29d}$$

where $\alpha \in \mathbb{R}_+^{s_\xi}$, $\beta \in \mathbb{R}_+$, $\gamma \in \mathbb{R}_+, \psi^+ \in \mathbb{R}_+^m$ and $\psi^- \in \mathbb{R}_+^m$ are the dual variables of the constraints in (28). We can now apply the linearization scheme employed in Zeng and Zhao (2013) on the worst-case linear recourse problem:

$$\max_{\{\phi_k\}_{k \in \mathcal{K}'}} \rho(\{\phi_k\}_{k \in \mathcal{K}'}),$$

with $\rho(\{\phi_k\}_{k \in \mathcal{K}'})$ as defined in (29). This gives rise to the mixed-integer linear program that appears in the proposition.

Note that this MILP reformulation is such that the number of binary variables does not increase with the size of the support $\mathcal{K}'$ of the randomized strategy. This would not be the case if one would apply the linearization scheme of Zeng and Zhao (2013) directly on $\mathtt{SP1}_\omega$. In some preliminary experiments, we established that our chosen approach had a significant impact on reducing the solution time for $\mathtt{SP1}_\omega$.

A.5   Proof of Theorem 2

The proof revolves around the fact that in each iteration, the algorithm either terminates with $UB - LB \le \varepsilon$ or adds at least one new member for either the set $\mathcal{K}' \subseteq \mathcal{K}$ or $\mathcal{H}'_\omega \subseteq \mathcal{H}_\omega$ among all $\omega \in \Omega$. Since the size of $\mathcal{K}$ is finite and the number of vertices of each Ξ'_ω is finite, the algorithm is guaranteed to converge in a finite number of steps.

A.6   Proof of Proposition 4

Since one can verify that $h(x,\xi) := c_1^\mathsf{T} x + h(x,\xi)$ is convex in x, based on Proposition 2, we have that the RSP reduces to

$$
\begin{aligned}
\underset{p\in\mathbb{R}^{|\mathcal{K}|},\{x^k\}_{k\in\mathcal{K}}}{\text{minimize}} \quad & \sum_{k\in\mathcal{K}} c_1^\mathsf{T} p_k x^k + \sup_{F_\xi\in\mathcal{D}} \mathbb{E}_{\xi\sim F_\xi}\left[\sum_{K\in\mathcal{K}} p_k h'(x^k,\xi)\right] \\
\text{subject to} \quad & C_x x^k \le d_x \,, \ \forall k \in \mathcal{K} \\
& P_\mathbb{Z} x^k = \bar{x}_1^k \,, \ \forall k \in \mathcal{K} \\
& p_k \ge 0 \,, \ \forall\, k \in \mathcal{K}, \quad \sum_{k\in\mathcal{K}} p_k = 1 \,.
\end{aligned}
$$

Using a simple change of variable $z_k := x_k p_k$ and $y' := y p_k$, we obtain the reduction presented in problem (17) by exploiting the fact that $\bar{\mathcal{X}} := \{x \in \mathbb{R}^n \,|\, C_x x \le d_x\}$ is assumed to describe a bounded set, and the fact that the recourse problem was assumed to be bounded and feasible for all $x \in \mathcal{X}$ and all $\xi \in \Xi$.

A.7   Proof of Proposition 5

Based on Theorem 1, we can conclude that $\text{VRS} = v_d - \psi$, where

$$
v_d := \min_{x\in\mathcal{X}_{\text{AP}}} \sup_{F_\xi\in\mathcal{D}} \mathbb{E}_{\xi\sim F_\xi}\left[\sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}} \xi_{ij} x_{ij}\right] \,.
$$

and

$$
\psi := \min_{x\in\mathcal{C}(\mathcal{X}_{\text{AP}})} \sup_{F_\xi\in\mathcal{D}} \mathbb{E}_{\xi\sim F_\xi}\left[\sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}} \xi_{ij} x_{ij}\right] \,.
$$

Given that the constraint matrix of the assignment problem embodies the total unimodularity property, the convex hull of $\mathcal{X}_{\text{AP}}$ is directly captured by its continuous relaxation $\mathcal{X}'_{\text{AP}}$. Specifically, the polyhedron defined in Equation (20) has only integer vertices.

Furthermore, the adversarial problems involved in computing v_d and ψ both take the form:

$$
\begin{aligned}
\max_{F_\xi\in\mathcal{D}} \mathbb{E}_{\xi\sim F_\xi}\left[\sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}} \xi_{ij} x_{ij}\right] &= \max_{F_\xi\in\mathcal{D}} \sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}} \mathbb{E}_{\xi\sim F_\xi}[\xi_{ij}] x_{ij} \\
&= \max_{\mu\in\{\mu\in\mathbb{R}^{|\mathcal{I}|\times|\mathcal{J}|}\,|\,\exists\, F_\xi\in\mathcal{D},\, \mu_{ij}=\mathbb{E}_{\xi\sim F_\xi}[\xi_{ij}]\}} \sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}} \mu_{ij} x_{ij} \\
&= \max_{\mu\in\mathcal{U}} \sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}} \mu_{ij} x_{ij} \\
&= \delta^*\left(x|\mathcal{U}\right) \,,
\end{aligned}
$$

where the first equality follows from linearity of expectation and the third equality follows from the fact that the maximum of a linear function over a convex set is always achieved at an extreme point of the set. Finally, the last equality follows from the definition of $\delta^*(\cdot|\mathcal{U})$. This completes our proof.

B Pseudocode description of two-layer column generation algorithm

Input: $\left\{\widehat{\xi}_\omega,\ \omega \in \Omega\right\} \subset \Xi,\ \epsilon \geq 0,\ x_d^*,\ \varepsilon \geq 0$
Output: ε-optimal randomized strategy F_x^* parametrized with $\{(x_k, p_k^*)\}_{k \in \mathcal{K}'}$
$\mathcal{K}' \longleftarrow \{k : \bar{x}^k = x_d^*\}$
$\mathcal{H}_n' \longleftarrow \emptyset\,, \forall \omega \in \Omega$
$LB \longleftarrow -\infty$
$UB \longleftarrow \infty$
Initialize $\{p_k^*\}_{k \in \mathcal{K}'}$ and λ^*
Initialize $q_{h_\omega}^{\omega *}$
while $UB - LB < \varepsilon$ **do**
$\quad$ // Solve `Primal`$(\mathcal{K}', \mathcal{H})$:
$\quad LB_1 \longleftarrow LB$
$\quad UB_1 \longleftarrow \infty$
$\quad$ **while** $UB_1 \geq LB_1$ **do**
$\qquad \forall \omega \in \Omega$, solve `SP1`$_\omega$ to get a new vertex $(\xi^{\bar{h}_\omega}, \zeta^{\bar{h}_\omega})$ and t_ω^*
$\qquad \mathcal{H}_\omega' \longleftarrow \mathcal{H}_\omega' \cup \{\bar{h}_\omega\}$for all $\omega \in \Omega$
$\qquad UB_1 = \min(UB_1, \sum\limits_{k \in \mathcal{K}} c_1^\mathsf{T} \bar{x}^k p_k^* + \lambda^* \epsilon + \frac{1}{|\Omega|} \sum\limits_{\omega \in \Omega} t_\omega^*)$
$\qquad$ Solve `MP1` to obtain new $\{p_k^*\}_{k \in \mathcal{K}'}$ and λ^* and update LB_1
$\quad$ **end**
$\quad UB \longleftarrow UB_1$
$\quad$ // Solve `Dual`$(\mathcal{K}, \mathcal{H}')$:
$\quad UB_2 \longleftarrow UB$
$\quad LB_2 \longleftarrow -\infty$
$\quad$ **while** $UB_2 \geq LB_2$ **do**
$\qquad$ Solve `SP2` to get a new solution $x^{\bar{k}}$ and the optimal value w^*
$\qquad \mathcal{K}' \longleftarrow \mathcal{K}' \cup \{\bar{k}\}$
$\qquad LB_2 = \max(LB_2, w^*)$
$\qquad$ Solve `MP2` to obtain new $q_{h_\omega}^*$ and update UB_2
$\quad$ **end**
$\quad LB \longleftarrow LB_2$
end

Algorithm 1: The Two-layer column-generation algorithm.

C Derivation of dual problem (18)

Using a Lagrangean duality approach, the optimal value of problem `Primal`$'(\mathcal{K}, \mathcal{H})$ can be reformulated as follows:

$$
\begin{aligned}
\min_{p \geq 0, \lambda \geq 0} \ & \min_{z^1 \in \mathcal{Z}_1(p_1), \ldots, z^{|K|} \in \mathcal{Z}_{|K|}(p_{|K|})} \lambda\epsilon + \sum_{k \in \mathcal{K}} c_1^\mathsf{T} z^k + \max_{w \geq 0} w \left(1 - \sum_{k \in \mathcal{K}} p_k\right) \\
& + \frac{1}{|\Omega|} \sum_{\omega \in \Omega} \max_{h_\omega \in \mathcal{H}_\omega} \sum_{k \in \mathcal{K}} h'(p_k, z^k, \xi^{h_\omega}) - \zeta^{h_\omega} \lambda \\
= \ & \min_{p \geq 0, \lambda \geq 0} \ \min_{x^1 \in \mathcal{X}_1, \ldots, x^{|\mathcal{K}|} \in \mathcal{X}_{|\mathcal{K}|}} \lambda\epsilon + \sum_{k \in \mathcal{K}} c_1^\mathsf{T} x^k p_k + \max_{w \geq 0} w \left(1 - \sum_{k \in \mathcal{K}} p_k\right) \\
& + \frac{1}{|\Omega|} \sum_{\omega \in \Omega} \max_{q^\omega \in \mathcal{Q}_\omega} \sum_{h_\omega \in \mathcal{H}_m} q_{h_\omega}^\omega \sum_{k \in \mathcal{K}} p_k h(x^k, \xi^{h_\omega}) - \zeta^{h_\omega} \lambda \\
= \ & \min_{p \geq 0, \lambda \geq 0} \ \max_{w \geq 0, q^1 \in \mathcal{Q}_1, \ldots, q^{|\Omega|} \in \mathcal{Q}_{|\Omega|}} \lambda\epsilon - \frac{1}{|\Omega|} \sum_{\omega \in \Omega} \sum_{h_\omega \in \mathcal{H}_\omega} q_{h_\omega}^\omega \zeta^{h_\omega} \lambda + w \left(1 - \sum_{k \in \mathcal{K}} p_k\right) \\
& + \sum_{k \in \mathcal{K}} p_k \left(\min_{x^k \in \mathcal{X}_k} c_1^\mathsf{T} x^k + \frac{1}{|\Omega|} \sum_{\omega \in \Omega} \sum_{h_\omega \in \mathcal{H}_\omega} q_{h_\omega}^\omega h(x^k, \xi^{h_\omega})\right)
\end{aligned}
$$

$$\geq \max_{w \geq 0, q^1 \in \mathcal{Q}_1, \ldots, q^{|\Omega|} \in \mathcal{Q}_{|\Omega|}} \min_{p \geq 0, \lambda \geq 0} \lambda\epsilon - \frac{1}{|\Omega|} \sum_{\omega \in \Omega} \sum_{h_\omega \in \mathcal{H}_\omega} q_{h_\omega}^\omega \zeta^{h_\omega} \lambda + w \left(1 - \sum_{k \in \mathcal{K}} p_k \right)$$

$$+ \sum_{k \in \mathcal{K}} p_k \left(\min_{x^k \in \mathcal{X}_k} c_1^\mathsf{T} x^k + \frac{1}{|\Omega|} \sum_{\omega \in \Omega} \sum_{h_\omega \in \mathcal{H}_\omega} q_{h_\omega}^\omega h(x^k, \xi^{h_\omega}) \right),$$

where $\mathcal{Z}_k(p_k) := \{z \in \mathbb{R}^n \,|\, C_x z^k \leq d_x p_k, \; P_\mathbb{Z} z^k = \bar{x}_1^k p_k\}$, and where $\mathcal{Q}_\omega := \{q \in \mathbb{R}^{|\mathcal{H}_\omega|} \,|\, q \geq 0, \sum_{h_\omega \in \mathcal{H}_\omega} q_{h_\omega} = 1\}$. It is then straightforward to show that the final maximization operation reduces to problem $\mathtt{Dual}'(\mathcal{K}, \mathcal{H})$ by replacing ${q_{h_\omega}^\omega}' := (1/|\Omega|) q_{h_\omega}^\omega$. We are left with explaining each step in order, and demonstrating that the last inequality is actually tight. In order, the first step follows from replacing $z^k := p_k x^k$ and replacing $\max_{h_\omega \in \mathcal{H}_\omega} a_{h_\omega}$ with $\max_{q \in \mathbb{R}_+^{|\mathcal{H}_\omega|}: \sum_{h_\omega \in \mathcal{H}_\omega} q_{h_\omega} = 1} \sum_{\omega \in \mathcal{H}_\omega} q_{h_\omega} a_{h_\omega}$. The second step, follows from applying the minimax theorem on $\min_{x^1 \in \mathcal{X}_1, \ldots, x^{|K|} \in \mathcal{X}_{|K|}} \max_{w \geq 0, q^1 \in \mathcal{Q}_1, \ldots, q^{|\Omega|} \in \mathcal{Q}_{|\Omega|}}$ which applies since each $\mathcal{X}_k$ is bounded and the function that is optimized over these two sets of variables is convex in x^k's and affine in w and each q^ω. The last step follows from weak minimax theory. One can also confirm that duality is strong here by finding a strictly feasible point for problem $\mathtt{Dual}'(\mathcal{K}, \mathcal{H})$ which implies that Slater's condition is satisfied. The following lemma completes this proof.

Lemma 1 *Given that $\epsilon > 0$ and that the relative interior of Ξ is non-empty, the polyhedron defined by $\mathcal{Q} := \{\{q^\omega\}_{\omega \in \Omega} \,|\, \sum_{\omega \in \Omega} \sum_{h_\omega \in \mathcal{H}_\omega} \zeta^{h_\omega} q_{h_\omega}^\omega \leq \epsilon, \sum_{h_\omega \in \mathcal{H}_\omega'} q_{h_\omega}^\omega = \frac{1}{|\Omega|}, q^\omega \geq 0, \; \forall \omega \in \Omega\}$ has a strict interior point.*

Proof. We instead demonstrate that

$$\mathcal{Q}' := \left\{ \{q^\omega\}_{\omega \in \Omega} \,|\, \frac{1}{|\Omega|} \sum_{\omega \in \Omega} \sum_{h_\omega \in \mathcal{H}_\omega} \zeta^{h_\omega} q_{h_\omega}^\omega \leq \epsilon, \; \sum_{h_\omega \in \mathcal{H}_\omega} q_{h_\omega}^\omega = 1, \; q^\omega \geq 0, \; \forall \omega \in \Omega \right\}$$

has a non-empty strict interior. The claim of the Lemma then follows straightforwardly since $\mathcal{Q}$ is a scaled version of $\mathcal{Q}'$. We construct a strict interior point as follows. First, we perturb each $\hat{\xi}_\omega$ to get $\xi_\omega' \in \Xi$ such that $\|\hat{\xi}_\omega - \xi_\omega'\|_1 < \min(\epsilon, \zeta_{\max})$ for all $\omega \in \Omega$ and such that each ξ_ω' is in the relative interior of Ξ. We then let $\zeta_\omega' := \|\hat{\xi}_\omega - \xi_\omega'\|_1$. Given that each pair $(\xi_\omega', \zeta_\omega') \in \Xi_\omega'$, by Carathéodory's theorem (see, *e.g.*, Eckhoff (1993)), there must therefore exist, for each $\omega \in \Omega$ a convex combination parameterized by $\{q_{h_\omega}^\omega\}_{h_\omega \in \mathcal{H}_\omega}$ such that $\xi_\omega' = \sum_{h_\omega \in \mathcal{H}_\omega} \xi^{h_\omega} q_{h_\omega}^\omega$ and $\zeta_\omega' = \sum_{h_\omega \in \mathcal{H}_\omega} \zeta^{h_\omega} q_{h_\omega}^\omega$. Furthermore, since $(\xi_\omega', \zeta_\omega')$ is in the relative interior of Ξ_ω', there must actually be an assignment for which each $q^\omega > 0$, for all $\omega \in \Omega$. In particular, one can first construct $\mu_\omega^\xi := (1/|\mathcal{H}_\omega|) \sum_{h_\omega \in \mathcal{H}_\omega} \xi^{h_\omega}$ and $\mu_\omega^\zeta := (1/|\mathcal{H}_\omega|) \sum_{h_\omega \in \mathcal{H}_\omega} \zeta^{h_\omega}$ and identify some $(\nu_\omega^\xi, \nu_\omega^\zeta) \in \Xi_\omega'$ such that $(\xi_\omega', \zeta_\omega')$ is the convex combination of $(\mu_\omega^\xi, \mu_\omega^\zeta)$ and $(\nu_\omega^\xi, \nu_\omega^\zeta)$. This is always possible since $(\xi_\omega', \zeta_\omega')$ is in the relative interior $\mathcal{H}_\omega$. The convex combination of the representations for $(\mu_\omega^\xi, \mu_\omega^\zeta)$ and $(\nu_\omega^\xi, \nu_\omega^\zeta)$ provides us with a representation for $(\xi_\omega', \zeta_\omega')$ that has $q^\omega > 0$ for all $\omega \in \Omega$. The constructed assignment for each $\{q_{h_\omega}^\omega\}_{h_\omega \in \mathcal{H}_\omega}$ is such that

$$\frac{1}{|\Omega|} \sum_{\omega \in \Omega} \sum_{h_\omega \in \mathcal{H}_\omega} \zeta^{h_\omega} q_{h_\omega}^\omega = \frac{1}{|\Omega|} \sum_{\omega \in \Omega} \zeta_\omega' = \frac{1}{|\Omega|} \sum_{\omega \in \Omega} \|\hat{\xi}_\omega - \xi_\omega'\|_1 < \epsilon.$$

We can therefore conclude that the identified assignment for $\{q^\omega\}_{\omega \in \Omega}$ must lie in the strict interior of $\mathcal{Q}'$. $\qquad\square$

D   Solving the DRCFLP with randomization

As described in Section 5.3, the two-layer column generation algorithm that is proposed to solve the DRCFLP with randomization iteratively solves four sets of optimization problems, two master problems, $\mathtt{MP1}$ and $\mathtt{MP2}$, and two subproblems, $\mathtt{SP1}_\omega$ and $\mathtt{SP2}$. For completeness, we briefly describe the details of these problems below.

The primal master problem takes the form of the following linear program:

$$[\text{MP1}]: \quad \underset{p\geq 0,\lambda\geq 0,\{t_\omega\}_{\omega\in\Omega}}{\text{minimize}} \quad \sum_{k\in\mathcal{K}'}\sum_{j\in\mathcal{J}} f_j x_j^k p_k + \lambda\epsilon + \frac{1}{|\Omega|}\sum_{\omega\in\Omega} t_\omega$$

$$\text{subject to} \quad \sum_{k\in\mathcal{K}'}\sum_{j\in\mathcal{J}}\sum_{i\in\mathcal{I}} c_{ij} z_{ijk} p_k - \zeta^{h_\omega}\lambda \leq t_\omega \quad \forall\omega\in\Omega,\ h_\omega\in\mathcal{H}_\omega$$

$$\sum_{k\in\mathcal{K}'} p_k = 1.$$

Each of the primal subproblems, indexed by $\omega\in\Omega$, takes the form of the following max-min problem:

$$[\text{SP1}_\omega]: \quad \underset{(\xi,\zeta,\delta)\in\Upsilon}{\text{maximize}} \quad \underset{z\geq 0}{\min} \quad \sum_{k\in\mathcal{K}'}\sum_{j\in\mathcal{J}}\sum_{i\in\mathcal{I}} p_k^* c_{ij} z_{ijk} - \lambda^*\zeta$$

$$\text{s.t.} \sum_{j\in\mathcal{J}} z_{ijk} = \xi_i \qquad \forall i\in\mathcal{I}, k\in\mathcal{K}' \qquad (\nu_{ik})$$

$$v_j x_j^k - \sum_{i\in\mathcal{I}} z_{ijk} \geq 0 \qquad \forall j\in\mathcal{J}, k\in\mathcal{K}' \qquad (\mu_{jk}\geq 0),$$

where $\Upsilon := \{(\xi,\zeta,\delta)\in\mathbb{R}^{|\mathcal{I}|}\times\mathbb{R}\times\mathbb{R}^{|\mathcal{I}|}\,|\,\zeta\geq 0, (27c)-(27g)\}$. Following a similar procedure as used in the proof of Proposition 3, we obtain the following equivalent mixed-integer linear program:

$$\underset{\substack{\xi,\zeta,\delta,\alpha,\beta,\psi,\nu,\mu\geq 0 \\ Bin^1,Bin^2,Bin^3,Bin^4,Bin^5,Bin^6,Bin^7}}{\text{maximize}} \quad d_\xi^\intercal\alpha + \bar{\zeta}_{max}\gamma + \sum_{i\in\mathcal{I}} \widehat{\xi}_i^\omega(\psi_i^+ - \psi_i^-) - \sum_{k\in\mathcal{K}'}\sum_{j\in\mathcal{J}} v_j x_j^k \mu_{jk}$$

$$\text{subject to} \quad \nu_{ik} - \mu_{jk} \leq c_{ij} p_k^* \quad \forall i\in\mathcal{I}, j\in\mathcal{J}, k\in\mathcal{K}'$$

$$(15b)-(15q).$$

On the other hand, the dual master problem takes the form of the following linear program:

$$[\text{MP2}]: \quad \underset{w,q\geq 0}{\text{maximize}} \quad w$$

$$\text{subject to} \quad w \leq \sum_{k\in\mathcal{K}}\sum_{j\in\mathcal{J}} \bar{x}^k + \sum_{\omega\in\Omega}\sum_{h_\omega\in\mathcal{H}_\omega'} h(\bar{x}^k,\xi^{h_\omega}) q_{h_\omega} \quad \forall k\in\mathcal{K} \qquad (p_k)$$

$$\sum_{\omega\in\Omega}\sum_{h_\omega\in\mathcal{H}_\omega'} \zeta^{h_\omega} q_{h_\omega} \leq \epsilon \qquad (\lambda)$$

$$\sum_{h_\omega\in\mathcal{H}_\omega'} q_{h_\omega} = \frac{1}{|\Omega|} \qquad \forall\omega\in\Omega \qquad (t_\omega),$$

while its associated subproblem reduces to

$$[\text{SP2}]: \quad \underset{x\in\mathcal{X},z\geq 0}{\text{minimize}} \quad \sum_{j\in\mathcal{J}} f_j x_j + \frac{1}{|\Omega|}\sum_{\omega\in\Omega}\sum_{j\in\mathcal{J}}\sum_{i\in\mathcal{I}} c_{ij} z_{ijh_\omega}$$

$$\text{subject to} \quad \sum_{j\in\mathcal{J}} z_{ijh_\omega} \geq \xi_i^{h_\omega} \qquad \forall i\in\mathcal{I},\ h_\omega\in\mathcal{H}_\omega'$$

$$\sum_{i\in\mathcal{I}} z_{ijh_\omega} \leq v_j x_j \qquad \forall j\in\mathcal{J},\ h_\omega\in\mathcal{H}_\omega'.$$

References

Agra A, Christiansen M, Hvattum LM, Rodrigues F (2018) Robust optimization for a maritime inventory routing problem. Transportation Science 52(3):509–525.

Ahuja RK, Magnanti TL, Orlin JB (1993) Network Flows: Theory, Algorithms, and Applications (Upper Saddle River, NJ, USA: Prentice-Hall, Inc.).

Ardestani-Jaafari A, Delage E (2018) The value of flexibility in robust location–transportation problems. Transportation Science 52(1):189–209.

Atamtürk A, Zhang M (2007) Two-stage robust network flow and design under demand uncertainty. Operations Research 55(4):662–673.

Ben-David S, Borodin A, Karp R, Tardos G, Wigderson A (1990) On the power of randomization in online algorithms. Algorithmica, 379–386.

Ben-Tal A, Goryashko A, Guslitzer E, Nemirovski A (2004) Adjustable robust solutions of uncertain linear programs. Mathematical Programming 99(2):351–376.

Bertsimas D, Nasrabadi E, Orlin J (2018) On the power of nature in robust discrete optimization. Private communication.

Bertsimas D, Sim M, Zhang M (2017) A practically efficient approach for solving adaptive distributionally robust linear optimization problems. Working Paper.

Brown GG, Carlyle WM, Harney RC, Skroch EM, Wood RK (2009) Interdicting a nuclear-weapons project. Operations Research 57(4):866–877.

Chan TCY, Shen ZM, Siddiq A (2018) Robust defibrillator deployment under cardiac arrest location uncertainty via row-and-column generation. Operations Research 66(2):358–379.

de Ruiter F, Ben-Tal A, Brekelmans R, Hertog D (2014) Adjustable robust optimization with decision rules based on inexact revealed data. Technical report, CentER.

Delage E, Kuhn D, Wiesemann W (2018) "Dice"-sion making under uncertainty: When can a random decision reduce risk? Management Science (accepted).

Eckhoff J (1993) Helly, Radon, and Carathéodory type theorems. Gruber P, Wills J, eds., Handbook of Convex Geometry, 389 – 448 (Amsterdam: North-Holland).

Fernández E, Landete M (2015) Fixed-Charge Facility Location Problems, 47–77 (Cham: Springer International Publishing).

Föllmer H, Schied A (2002) Convex measures of risk and trading constraints. Finance and Stochastics 6(4):429–447.

Gilmore PC, Gomory RE (1961) A linear programming approach to the cutting-stock problem. Operations Research 9(6):849–859.

Jain M, Korzhyk D, Vaněk O, Conitzer V, Pěchouček M, Tambe M (2011) A double oracle algorithm for zero-sum security games on graphs. The 10th International Conference on Autonomous Agents and Multiagent Systems - Volume 1, 327–334, AAMAS '11 (Richland, SC).

Kantorovich LV, Rubinshtein GS (1958) On a space of completely additive functions. Vestnik Leningrad. Univ. 13(7):52–59.

Luo F, Mehrotra S (2017) Decomposition Algorithm for Distributionally Robust Optimization using Wasserstein Metric. ArXiv e-prints.

Mastin A, Jaillet P, Chin S (2015) Randomized minmax regret for combinatorial optimization under uncertainty. Elbassioni K, Makino K, eds., Algorithms and Computation, 491–501 (Berlin, Heidelberg: Springer Berlin Heidelberg).

McMahan H, Gordon GJ, Blum A (2003) Planning in the presence of cost functions controlled by an adversary. Proceedings of the Twentieth International Conference on Machine Learning (ICML).

Mohajerin Esfahani P, Kuhn D (2017) Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations. Mathematical Programming.

Morris JG (1978) On the extent to which certain fixed-charge depot location problems can be solved by LP. Journal of the Operational Research Society 29(1):71–76.

Shapiro A (2001) On Duality Theory of Conic Linear Problems, 135–165 (Boston, MA: Springer US).

Smith JE, Winkler RL (2006) The optimizer's curse: Skepticism and postdecision surprise in decision analysis. Management Science 52(3):311–322.

Thiele A, Terry T, Epelman M (2010) Robust linear optimization with recourse. Working draft.

Van Parys BPG, Mohajerin Esfahani P, Kuhn D (2017) From Data to Decisions: Distributionally Robust Optimization is Optimal. Working draft.

von Neumann J (1928) Zur theorie der gesellschaftsspiele. Mathematische Annalen 100(1):295–320.

Wang H, Rezaei A, Ziebart BD (2017) Adversarial Structured Prediction for Multivariate Measures. ArXiv e-prints.

Wiesemann W, Kuhn D, Sim M (2014) Distributionally robust convex optimization. Operations Research 62(6):1358–1376.

Yang Z, Zhu J, Teng L, Xu J, Zhu Z (2018) A double oracle algorithm for allocating resources on nodes in graph-based security games. Multimedia Tools and Applications 77(9):10961–10977.

Zeng B, Zhao L (2013) Solving two-stage robust optimization problems using a column-and-constraint generation method. Operations Research Letters 41(5):457–461.

Zhao C, Guan Y (2018) Data-driven risk-averse stochastic optimization with Wasserstein metric. Operations Research Letters 46(2):262–267.