The Value of Randomized Solutions in Mixed-Integer Distributionally Robust Optimization Problems

Erick Delage*, Ahmed Saif†

October 13th, 2020

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 roll 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 improvements 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 continuous relaxation. Furthermore, we identify conditions under which the bound is tight. We then develop algorithmic procedures, based on column generation, for solving both single- and 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.

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 (Parys et al. 2017). In DRO, a decision maker typically minimizes the worst-case risk 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 (see Mohajerin Esfahani and Kuhn (2018) for a discussion). On the other hand, DRO avoids the inherent over-conservatism of RO that usually leads to poor expected performances, and allows for 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 often obtain a statistical guarantee on the *out-of-sample* performance of the DRO problem's solution (*e.g.,* see Delage and Ye (2010) and Mohajerin Esfahani and Kuhn (2018)).

Recently, Delage et al. (2019) introduced the idea of exploiting randomized strategies in DRO problems that arise when using worst-case 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

*GERAD & Department of Decision Sciences, HEC Montréal, Montréal, Québec, H3T 2A7, Canada (erick.delage@hec.ca)
†Department of Industrial Engineering, Dalhousie University, Halifax, Nova Scotia, B3J 1B6, Canada (ahmed.saif@dal.ca).

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. (2019, 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 a worst-case risk measure than what can be achieved by any 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 that the value of randomization in mixed-integer DRO problems with convex cost functions and convex risk measures is bounded by the difference between the optimal values of the nominal DRO problem and that of its continuous 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 finitely convergent column generation algorithms for solving single- and two-stage mixed-integer linear DRO problems with randomization and expected value as the risk measure. The two-stage problem algorithm iterates between solving a restricted primal problem to generate new candidates to be added to the list of possible worst-case scenarios, and solving a restricted dual problem to generate new candidates to be added to the randomized strategy. 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 a 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 in Section 3 an example of distributionally robust uncapacitated facility location problem to illustrate how worst-case risk in mixed-integer DRO problems can be reduced through randomization. A second example involving a distributionally robust newsvendor problem can be found in Appendix A. 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 column generation algorithms for solving single- and two-stage distributionally robust integer linear programming problems and explain how they can be modified to solve mixed-integer problems. These algorithms are then used to solve three classical discrete problems in Section 6: the assignment problem, the uncapacitated and capacitated facility location 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. Finally, conclusions are drawn and directions for future research are proposed in Section 8. We note that all the proofs of our theorems are deferred to Appendix B while Appendix G presents a discussion on the difficulty of adoption of randomized strategies in practice.

Remark 1. *It is worth clarifying for the reader the distinction between the two related concepts of "risk aversion" and "ambiguity aversion". Following the work of Ellsberg (1961) and Epstein (1999), the notion of risk refers "to situations where the perceived likelihood of events of interests can be represented by probabilities", whereas ambiguity "refers to situations where the information available to the decision maker is too imprecise to be summarized by a probability measure". Similarly, the notions of risk aversion* *and ambiguity aversion refer to how a decision maker behaves when being exposed to random variables with known, or with unknown distributions, respectively. Both types of random variables emerge in a decision problem with distributional ambiguity, especially in contexts when randomization can be used (see Section 3 for two examples). Throughout the paper, risk aversion will be modeled using a law-invariant convex risk measure, e.g., the expected value (referred to as the "risk neutral" attitude), conditional value-at-risk, etc. On the other hand, following an axiomatic motivation proposed in Delage et al. (2019), ambiguity aversion will be modeled by measuring the risk associated with a random variable when the distribution is fixed to a worst-case realization from the distributional ambiguity set. This gives rise to the notion of "ambiguity averse risk measure" (i.e., worst-case risk measure) minimization, which is the class of DRO problems that we focus on.*

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. Unless specifically indicated, all vectors are column vectors and 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 each has marginal distribution 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 $\mathcal{X}$. We denote by ρ an arbitrary law-invariant convex risk measure, while $\sup_{F_\xi \in \mathcal{D}} \rho_{\xi \sim F_\xi}(X)$ refers to the worst-case risk measured using ρ when letting the distribution of ξ take on any distribution in the set $\mathcal{D}$. Finally, when LP duality is used, the dual/primal variables are included, between parentheses, right after their corresponding primal/dual constraints in the mathematical formulations.

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 distribution 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 parameter space (*e.g.,* the Wasserstein ambiguity sets proposed by Mohajerin Esfahani and Kuhn (2018)). Comparatively, zero-sum games usually treat risk aversion using the 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). Herein, the authors study RO in combinatorial optimization problems against an adaptive online adversary, which acts after the decision maker but only exploits information about the decision maker's randomized strategy (see Ben-David et al. (1990) for an application of this concept in online optimization). Similarly as shown in Delage et al. (2019) for this kind of problems, the authors show that an ambiguity averse risk neutral decision maker can strictly benefit from using randomized instead of a deterministic 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 more 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 players 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 an ambiguity averse risk measure (instead of regret) as the objective. 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 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 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 distribution 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 algorithm that we present in Section 5 address more general and complicated problems than the two-person zero-sum matrix games found in the literature. 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

In this section, we present an example that illustrates how randomization can reduce risks in mixed-integer DRO problems. Note that a second example involving the popular newsvendor problem can be found in Appendix A, where randomization can encourage the decision maker to act less conservatively, and where we hint at how randomized strategies could be implemented in practice by designing contracts that outsource the act of randomization.

The uncapacitated facility location problem seeks to select a subset of a discrete set of potential locations to open facilities and to assign the demands originating from a discrete set of nodes to open facilities in order to minimize the setup and shipping costs. We focus on the case when demands are uncertain and the entire demand of each node must be assigned to a single facility. Consider the following 2-node Distributionally Robust Uncapacitated Facility Location Problem (DRUFLP):

$$
\begin{aligned}
\underset{x,y}{\text{minimize}} \quad & \max_{F_\xi \in \mathcal{D}} \mathbb{E}_{(\xi_1,\xi_2)\sim F_\xi}[f \cdot (x_1 + x_2) + c \cdot (\xi_1 y_{12} + \xi_2 y_{21})] \\
\text{subject to} \quad & \sum_{j \in \{1,2\}} y_{ij} = 1 && \forall i \in \{1,2\} \\
& y_{ij} \le x_j && \forall i \in \{1,2\}, j \in \{1,2\} \\
& x_j, y_{ij} \in \{0,1\} && \forall i \in \{1,2\}, j \in \{1,2\},
\end{aligned}
$$

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. Note that there is no transportation cost if the demand is served by the facility at the same location. We also let the distributional 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

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

which simply captures the fact that the only information available about the random vector $[\xi_1 \ \ \xi_2]^\intercal$ is that the sum of any subset of its terms cannot be strictly greater 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 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) pairs:

$$\begin{aligned} \max_{F_\xi \in \mathcal{D}} \ \mathbb{E}_{(\xi_1,\xi_2)\sim F_\xi}[f \cdot (x_1 + x_2) + c \cdot (\xi_1 y_{12} + \xi_2 y_{21})] \ &= \ \max_{\xi \in \mathcal{U}} f \cdot (x_1 + x_2) + c \cdot (\xi_1 y_{12} + \xi_2 y_{21}) \\ &= \max_{(\xi_1,\xi_2)\in\{(\bar{d},0),\,(0,\,\bar{d}),(0,\,0)\}} f \cdot (x_1 + x_2) + c\cdot(\xi_1 y_{12} + \xi_2 y_{21}) \\ &= \ f\cdot(x_1 + x_2) + \max\{c\bar{d}y_{12}\, ,\, c\bar{d}y_{21}\} \\ &\geq \ f + c\bar{d} = \max_{F_\xi \in \mathcal{D}} \mathbb{E}_{(\xi_1,\xi_2)\sim F_\xi}[f \cdot (1 + 0) + c \cdot (\xi_1 \cdot 0 + \xi_2 \cdot 1)] \, . \end{aligned}$$

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 a location that will serve all the demand, we have that

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

With this randomized strategy, the maximum reduction in worst-case expected cost is realized when $f = c\bar{d}$, in which case the reduction amounts to 25%.

4 The Value of Randomized Solutions

The example presented in the previous section illustrated how randomization can immunize against distributional ambiguity. While the reduction achieved in worst-case expected cost might already make a randomized strategy appear attractive, questions remain on the magnitude of improvements that decision makers can expect through randomization in other instances and problems, and how to efficiently find the optimal randomized strategies in larger-scale problems. This section tries to answer the first question by providing useful bounds on the value of randomized solutions, whereas the second question is addressed in Section 5.

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 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, $h(x, \xi) : \mathcal{X} \times \Xi \mapsto \mathbb{R}$ 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 $\rho_{\xi\sim F_\xi}\left(h(x, \xi)\right) =$

$\mathbb{E}_{\xi\sim F_\xi}[h(x,\xi)]$, as discussed in Remark 1, 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. (2019), 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 is 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. (2019) 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.

Proposition 1. *Given that ρ is a law-invariant 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 in x 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`****.*

Proposition 1 provides a mean of bounding the value of randomization using any convex relaxation $\mathcal{X}'$ of $\mathcal{X}$. In particular, in many applications of DRO, the ambiguity averse 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. Proposition 1 also states that when the risk measure $\rho(\cdot)$ 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 2. *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 $n+1$ points, that achieves optimality in problems* (2) *and* (4).

In Section 6, we show how Proposition 2 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 2 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 3. *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*

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

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

Remark 2. *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 $\mathcal{D} := \{F_\xi \,|\, \mathbb{P}_F(\xi \in \Xi) = 1\}$. Under these conditions, they establish that the bound presented in Proposition 1 is tight and can be solved 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 2), they prove that the ratio v_d/v_r is always bounded by $n+1$, which can be tightened when $\Xi \subset \mathbb{R}_+^m$ is convex, compact 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 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 Linear 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:

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

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 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

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

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

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 optimal value of the recourse problem is bounded for all $x \in \mathcal{X}$ and $\xi \in \Xi$.

Remark 3. *In describing problem (5), we make a list of assumptions that focuses our algorithmic efforts on problems that fall in the class of two-stage linear optimization models. While the structure of these models limits applicability, this class of problems does encompass a large number of interesting decision problems encountered in practice: e.g., in power systems (Dehghan et al. 2017), lot-sizing (Bertsimas and de Ruiter 2016), and supply chain management (Buhayenko and den Hertog 2017). Our assumption about Ξ also prevents us from modeling situations in which one would like to include distributions with unbounded support such as a multi-variate normal or exponential distributions. However, this is a common assumption that is made about DRO problems to help with either the theoretical or numerical analysis (e.g., see Assumption 1 in Delage and Ye (2010)). Finally, in sharp contrast with Section 4, our algorithmic contribution will focus on the ambiguity averse risk neutral case where $\rho(\cdot) = \mathbb{E}[\cdot]$. Extending our proposed algorithms to accommodate more general frameworks constitutes a natural direction for future research.*

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 us 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 strategy selects the feasible action $\bar{x}^k$. Mathematically, the randomized strategy problem can be rewritten as

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

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

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 facts that X and ξ are independent, that X is discrete, and that the expectation operator is linear:

$$\begin{aligned}
&\underset{F_x\in\Delta(\mathcal{X})}{\text{minimize}} \sup_{F_\xi\in\mathcal{D}} \mathbb{E}_{(X,\xi)\sim F_x\times F_\xi}\left[c_1^\intercal X + h(X,\xi)\right] \\
&\equiv \underset{F_x\in\Delta(\mathcal{X})}{\text{minimize}} \sup_{F_\xi\in\mathcal{D}} \mathbb{E}_{X\sim F_x}\left[c_1^\intercal X + \mathbb{E}_{\xi\sim F_\xi}[h(X,\xi)]\right] \\
&\equiv \underset{p:p_k\geq 0,\, \sum_{k\in\mathcal{K}} p_k=1}{\text{minimize}} \sum_{k\in\mathcal{K}} c_1^\intercal \bar{x}^k p_k + \sup_{F_\xi\in\mathcal{D}} \sum_{k\in\mathcal{K}} \mathbb{E}_{\xi\sim F_\xi}\left[h(\bar{x}^k,\xi)\right] p_k \\
&\equiv \underset{p:p_k\geq 0,\, \sum_{k\in\mathcal{K}} p_k=1}{\text{minimize}} \sum_{k\in\mathcal{K}} c_1^\intercal \bar{x}^k p_k + \sup_{F_\xi\in\mathcal{D}} \mathbb{E}_{\xi\sim F_\xi}\left[\sum_{k\in\mathcal{K}} h(\bar{x}^k,\xi)p_k\right] \\
&\equiv \underset{p:p_k\geq 0,\, \sum_{k\in\mathcal{K}} p_k=1}{\text{minimize}} \sum_{k\in\mathcal{K}} c_1^\intercal \bar{x}^k p_k + \sup_{F_\xi\in\mathcal{D}} \mathbb{E}_{\xi\sim F_\xi}\left[g(p,\xi)\right].
\end{aligned}$$

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

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

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 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 be briefly discussed in Remark 4). A notable example of a piecewise linear 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 4.

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

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

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

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

where $\zeta^{\max} := \sup_{l\in\mathcal{L},\xi\in\Xi}\pi_l(\xi)$ and where $\zeta_l\leq\zeta^{\max}$ 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:

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

which can then be reintegrated in the main optimization problem as

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

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

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

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 (2018). 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,

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

where $\mathcal{M}(\Xi)$ is the space of all distributions F supported on Ξ with $\mathbb{E}_{\xi\sim F}\left[\|\xi\|\right] = \int_\Xi \|\xi\| F(\mathrm{d}\xi) < \infty$ 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) and Fournier and Guillin (2015, Theorem 2) 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 probabilistic guarantee of containing the true underlying distribution from which the observations were drawn.

Using similar steps as used in Mohajerin Esfahani and Kuhn (2018), problem (7) can be reformulated as

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

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

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

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

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

where $\zeta^{\max} := \sup_{\xi\in\Xi}\|\xi-\widehat{\xi}_\omega\|_1$ is, again, chosen such that $\zeta\leq\zeta^{\max}$ 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:

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

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

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

5.3 A column generation algorithm for single-stage problems

We just established that under fairly weak assumptions, *i.e.,* piecewise-linear functions $\pi_l(\cdot)$ in (8) and l_1-norm in (10), one can reformulate problem (7) 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. However, before addressing the general two-stage case, this section presents a simple column generation algorithm that can be used to identify an optimal randomized strategy when problem (7) reduces to a single-stage problem, *i.e.,* $h(x,\xi) := \xi^\intercal C_2 x$ with $C_2\in\mathbb{R}^{m\times n}$. Also, for simplicity of exposure, our discussion will focus on the case of a Wasserstein ambiguity set, yet can easily be modified to accommodate problem (9). Hence, the robust constraint (12b), indexed by $\omega\in\Omega$, can be written as $\sup_{(\xi,\zeta)\in\Xi'_\omega}\sum_{k\in\mathcal{K}}\xi^\intercal C_2 x^k p_k - \lambda\zeta \leq t_\omega$, or equivalently

$$\underset{\xi_\omega,\zeta_\omega\geq 0,\delta_\omega\geq 0}{\text{maximize}} \quad \sum_{k\in\mathcal{K}} \xi_\omega^\intercal C_2 x^k p_k - \lambda\zeta_\omega \leq t_\omega \tag{13a}$$

$$\text{subject to} \quad C_\xi\xi_\omega \leq d_\xi \qquad (\alpha_\omega) \tag{13b}$$

$$\zeta_\omega \leq \zeta^{\max} \qquad (\beta_\omega) \tag{13c}$$

$$e^\intercal\delta_\omega \leq \zeta_\omega \qquad (\gamma_\omega) \tag{13d}$$

$$\xi_\omega - \widehat{\xi}_\omega \leq \delta_\omega \qquad (\psi_\omega^+) \tag{13e}$$

$$-\xi_\omega + \widehat{\xi}_\omega \leq \delta_\omega, \qquad (\psi_\omega^-) \tag{13f}$$

where $\alpha_\omega \in \mathbb{R}_+^{s_\xi}$, $\beta_\omega, \gamma_\omega \in \mathbb{R}_+$ and $\delta_\omega, \psi_\omega^+, \psi_\omega^- \in \mathbb{R}_+^m$. Note that the tuple (ξ, ζ) is indexed by ω to emphasize that each constraint (12b) is individually reformulated. By applying LP duality on (13) for each $\omega \in \Omega$ and reintegrating the resulting dual formulations into (12), we get the following large-scale LP:

$$\underset{\substack{p_k \geq 0, \lambda \geq 0 \\ \alpha_\omega \geq 0, \beta_\omega \geq 0, \gamma_\omega \geq 0 \\ \psi_\omega^+ \geq 0, \psi_\omega^- \geq 0}}{\text{minimize}} \quad \sum_{k \in \mathcal{K}} c_1^\intercal x^k p_k + \lambda \epsilon + \frac{1}{|\Omega|} \left(\sum_{\omega \in \Omega} d_\xi^\intercal \alpha_\omega + \zeta^{\max} \beta_\omega + \widehat{\xi}_\omega^\intercal (\psi_\omega^+ - \psi_\omega^-) \right) \tag{14a}$$

$$\text{subject to} \quad C_\xi^\intercal \alpha_\omega + \psi_\omega^+ - \psi_\omega^- = \sum_{k \in \mathcal{K}} C_2 x^k p_k \qquad \forall \omega \in \Omega \qquad (\xi_\omega) \tag{14b}$$

$$\beta_\omega + \lambda \geq \gamma_\omega \qquad \forall \omega \in \Omega \qquad (\zeta_\omega) \tag{14c}$$

$$\psi_\omega^+ + \psi_\omega^- \leq e \gamma_\omega \qquad \forall \omega \in \Omega \qquad (\delta_\omega) \tag{14d}$$

$$\sum_{k \in \mathcal{K}} p_k = 1. \qquad (w) \tag{14e}$$

Given the exponential size of $\mathcal{K}$, it is usually not possible to enumerate all of its elements at the outset. Instead, we will solve problem (14) using a smaller set $\mathcal{K}'$ and progressively add new candidates to this set until an ε-optimal randomized strategy is found. This so-called column generation algorithm can be seen as performing constraint generation on the following dual problem:

$$\underset{\xi_\omega, \zeta_\omega \geq 0, \delta_\omega \geq 0, w}{\text{maximize}} \quad w \tag{15a}$$

$$\text{subject to} \quad w \leq c_1^\intercal x^k + \sum_{\omega \in \Omega} \xi_\omega^\intercal C_2 x^k \qquad \forall k \in \mathcal{K} \qquad (p_k) \tag{15b}$$

$$\sum_{\omega \in \Omega} \zeta_\omega \leq \epsilon \qquad (\lambda) \tag{15c}$$

$$C_\xi \xi_\omega \leq \frac{1}{|\Omega|} d_\xi \qquad \forall \omega \in \Omega \qquad (\alpha_\omega) \tag{15d}$$

$$\left\| \xi_\omega - \frac{\widehat{\xi}_\omega}{|\Omega|} \right\|_1 \leq \zeta_\omega \leq \frac{\zeta^{\max}}{|\Omega|} \qquad \forall \omega \in \Omega. \tag{15e}$$

Hence, the column generation algorithm can be described as follows.

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 the deterministic strategy problem (5)). Set the upper bound $UB := \infty$ and the lower bound $LB := -\infty$.
2. Solve the restricted master problem, *i.e.,* problem (14), with the subset $\mathcal{K}'$ instead of $\mathcal{K}$ to obtain a new upper bound and the dual variables $\bar{\xi}_\omega$ of (14b).
3. Solve the subproblem

$$\underline{w} = \min_{x \in \mathcal{X}} \left(c_1^\intercal + \sum_{\omega \in \Omega} \bar{\xi}_\omega^\intercal C_2 \right) x$$

   to obtain a new feasible solution $\bar{x}^k$ and update the lower bound as $LB := \max(LB, \underline{w})$. If $UB - LB < \varepsilon$, terminate the algorithm and declare the current solution p_k^* of the restricted master problem to be optimal for the randomized strategy problem. Otherwise, add the new index k to $\mathcal{K}'$ and repeat steps 2 and 3.

5.4 A two-layer column generation algorithm

In this section, we propose a two-layer column generation algorithm that can identify an optimal randomized strategy for problem (7) together with its optimal value v_r. For simplicity of exposure, our discussion will focus on the case of a Wasserstein ambiguity set yet can easily be modified to accommodate problem (9). First, we note that since $h(\bar{x}^k, \xi)$ is convex in ξ (see chapter 3.2.5 of Boyd and Vandenberghe (2004) for a proof of convexity) and 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

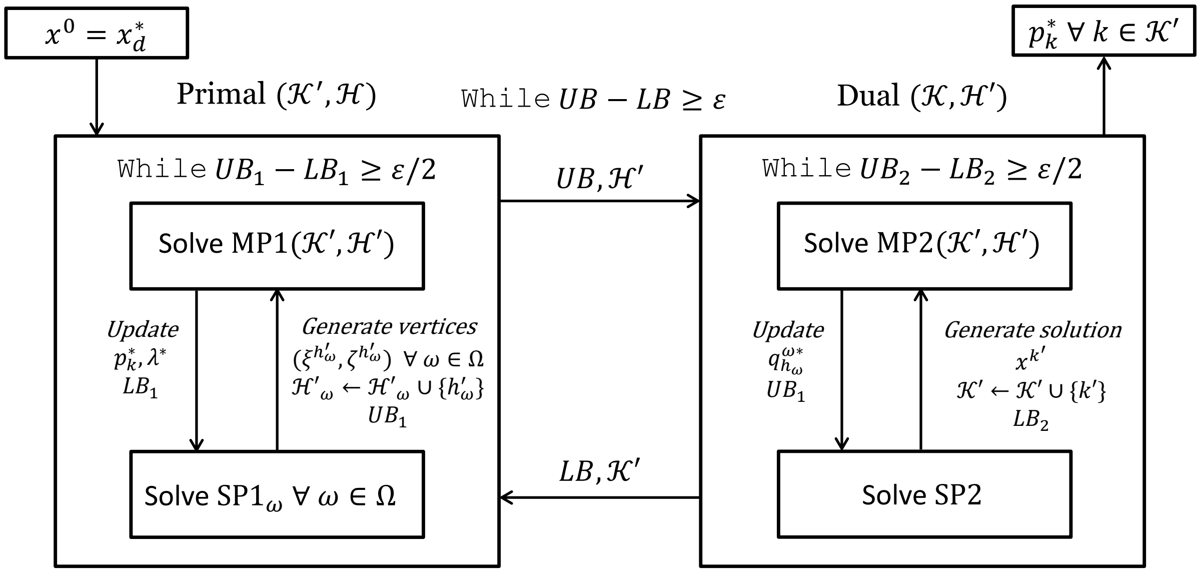

Figure 1: An outline of the two-layers column generation algorithm.

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:

$$\begin{aligned}
\underset{p \geq 0, \lambda \geq 0, \{t_\omega\}_{\omega \in \Omega}}{\text{minimize}} \quad & \sum_{k \in \mathcal{K}} c_1^\intercal \bar{x}^k p_k + \lambda \epsilon + \frac{1}{|\Omega|} \sum_{\omega \in \Omega} t_\omega && \text{(16a)} \\
\text{subject to} \quad & \sum_{k \in \mathcal{K}} h(\bar{x}^k, \xi^{h_\omega}) p_k - \zeta^{h_\omega} \lambda \leq t_\omega \quad \forall h_\omega \in \mathcal{H}_\omega, \; \omega \in \Omega \quad (q_{h_\omega}^\omega) && \text{(16b)} \\
& \sum_{k \in \mathcal{K}} p_k = 1. \quad (w) && \text{(16c)}
\end{aligned}$$

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 (an outline of the algorithm is provided in Fig. 1 and a *pseudocode* description is also presented in Appendix C):

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 $\omega \in \Omega$,

$$\begin{aligned}
[\texttt{Primal}(\mathcal{K}', \mathcal{H})] : \quad \underset{p \geq 0, \lambda \geq 0, \{t_\omega\}_{\omega \in \Omega}}{\text{minimize}} \quad & \sum_{k \in \mathcal{K}'} c_1^\intercal \bar{x}^k p_k + \lambda \epsilon + \frac{1}{|\Omega|} \sum_{\omega \in \Omega} t_\omega \\
\text{subject to} \quad & \sum_{k \in \mathcal{K}'} h(\bar{x}^k, \xi^{h_\omega}) p_k - \zeta^{h_\omega} \lambda \leq t_\omega \quad \forall h_\omega \in \mathcal{H}_\omega, \; \omega \in \Omega \\
& \sum_{k \in \mathcal{K}'} p_k = 1,
\end{aligned}$$

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 (16) 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 satisfies $p^* \geq 0$, $\lambda^* \geq 0$, 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 4.
(d) Let t^*_ω be the optimal value of $\texttt{SP1}_\omega$ for all $\omega \in \Omega$. Update the upper bound as

$$UB_1 := \min\left(UB_1, \sum_{k\in\mathcal{K}} c_1^\intercal \bar{x}^k p^*_k + \lambda^*\epsilon + \frac{1}{|\Omega|}\sum_{\omega\in\Omega} t^*_\omega\right).$$

(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 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 need to be computed once when a new vertex is added to $\mathcal{H}'_\omega$.
(f) If $UB_1 - LB_1 < \varepsilon/2$ 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 $\omega \in \Omega$. This problem is derived by taking the dual of the large-scale LP (16) with multipliers $q^\omega_{h_\omega}$ and w for constraints (16b) and (16c), respectively, and considering, for all $\omega \in \Omega$, the subset $\mathcal{H}'_\omega \subset \mathcal{H}_\omega$ instead of the complete index sets $\mathcal{H}_\omega$. The resulting dual master problem is

$$[\texttt{Dual}(\mathcal{K},\mathcal{H}')] : \quad \underset{w,\{q^\omega\}_{\omega\in\Omega}}{\text{maximize}} \quad w \tag{17a}$$

$$\text{subject to} \quad w \leq c_1^\intercal \bar{x}^k + \sum_{\omega\in\Omega}\sum_{h_\omega\in\mathcal{H}'_\omega} h(\bar{x}^k,\xi^{h_\omega}) q^\omega_{h_\omega} \quad \forall k \in \mathcal{K} \quad (p_k) \tag{17b}$$

$$\sum_{\omega\in\Omega}\sum_{h_\omega\in\mathcal{H}'_\omega} \zeta^{h_\omega} q^\omega_{h_\omega} \leq \epsilon \quad (\lambda) \tag{17c}$$

$$\sum_{h_\omega\in\mathcal{H}'_\omega} q^\omega_{h_\omega} = \frac{1}{|\Omega|} \quad \forall \omega \in \Omega \quad (t_\omega) \tag{17d}$$

$$q^\omega \geq 0 \quad \forall \omega \in \Omega. \tag{17e}$$

The optimal value of $\texttt{Dual}(\mathcal{K},\mathcal{H}')$ provides a lower bound for problem (16) 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 set of reasonable size using the following sub-procedure:

(a) Set $LB_2 := -\infty$ and $UB_2 := UB$
(b) Initialize each $q^{\omega*}_{h_\omega}$ to an arbitrary solution that satisfy constraints (17c), (17d) and (17e). 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 $\underset{x\in\mathcal{X}}{\text{minimize}} \quad c_1^\intercal x + \sum_{\omega\in\Omega}\sum_{h_\omega\in\mathcal{H}'_\omega} h(x,\xi^{h_\omega}) q^{\omega*}_{h_\omega}$, which reduces to the mixed-integer linear program

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

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

Consider the optimal x^* to be a new support point for the optimal randomized strategy.
(d) Let w^* be the optimal value obtained for 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 (`MP2`), defined as the restricted version of `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 once only 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 `SP1`$_\omega$ can be reformulated as a mixed-integer linear program.

Proposition 4. *Problem* `SP1`$_\omega$ *is equivalent to the following mixed-integer linear program:*

$$
\begin{aligned}
\underset{\substack{\xi,\zeta,\delta,\alpha,\beta,\psi,\phi\geq 0\\ Bin^1,Bin^2,Bin^3,Bin^4,Bin^5,Bin^6,Bin^7}}{\text{maximize}} \quad & \sum_{k\in\mathcal{K}'} (W_0\bar{x}^k + b)^\intercal \phi_k + d_\xi^\intercal \alpha + \zeta^{max}\gamma + \widehat{\xi}_\omega^\intercal(\psi^+ - \psi^-) \\
\text{subject to} \quad & A^\intercal \phi_k = c_2 p_k^* \qquad \forall k \in \mathcal{K}' && \text{(18a)}\\
& \sum_{i=1}^m \left(\sum_{k\in\mathcal{K}'} \phi_k^T W_i x_k \right) e_i = C_\xi^T \alpha + \psi^+ - \psi^- && \text{(18b)}\\
& 0 \leq d - C\xi \leq M(1 - Bin^1) && \text{(18c)}\\
& 0 \leq \alpha \leq M Bin^1 && \text{(18d)}\\
& 0 \leq \zeta - e^\intercal \delta \leq M(1 - Bin^2) && \text{(18e)}\\
& 0 \leq \beta \leq M Bin^2 && \text{(18f)}\\
& 0 \leq \zeta^{max} - \zeta \leq M(1 - Bin^3) && \text{(18g)}\\
& 0 \leq \gamma \leq M Bin^3 && \text{(18h)}\\
& 0 \leq \delta - \xi + \widehat{\xi}_\omega \leq M(1 - Bin^4) && \text{(18i)}\\
& 0 \leq \psi^+ \leq M Bin^4 && \text{(18j)}\\
& 0 \leq \delta + \xi - \widehat{\xi}_\omega \leq M(1 - Bin^5) && \text{(18k)}\\
& 0 \leq \psi^- \leq M Bin^5 && \text{(18l)}\\
& 0 \leq \zeta \leq M(1 - Bin^6) && \text{(18m)}\\
& 0 \leq \lambda^* + \gamma - \beta \leq M Bin^6 && \text{(18n)}\\
& 0 \leq \delta \leq M(1 - Bin^7) && \text{(18o)}\\
& 0 \leq \beta - \psi^+ - \psi^- \leq M Bin^7 && \text{(18p)}\\
& Bin^1 \in \{0, 1\}^{s_\xi},\ Bin^2 \in \{0, 1\},\ Bin^3 \in \{0, 1\},\ Bin^4 \in \{0, 1\}^m, \\
& Bin^5 \in \{0, 1\}^m,\ Bin^6 \in \{0, 1\},\ Bin^7 \in \{0, 1\}^m, && \text{(18q)}
\end{aligned}
$$

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* $\psi^- \in \mathbb{R}_+^m$*, while* M *is a large enough constant.*

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 `SP1`$_\omega$.

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

Remark 4. *In Theorem 1, the convergence in a 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 section how the algorithm can be modified to handle the case where* $\mathcal{X}$ *is mixed-integer, namely* $n_2 > 0$*. The case of general ambiguity sets can also be accommodated but requires one to design an alternative scheme for solving* `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 4 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.*

5.5 The case of mixed-integer feasible region

We briefly outline how the algorithm presented in Section 5.4 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 3, 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.,* $\{\bar{x}_1^k\}_{k\in\mathcal{K}} := \{x_1 \in \mathbb{Z}^{n_1} \,|\, \exists x_2 \in \mathbb{R}^{n_2},\, [x_1^\intercal\; x_2^\intercal]^\intercal \in \mathcal{X}\}$.

Proposition 5. *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*

$$\begin{aligned}
\underset{p\geq 0,\{z^k\}_{k\in\mathcal{K}}}{\text{minimize}} \quad & \sum_{k\in\mathcal{K}} c_1^\intercal z^k + \sup_{F_\xi\in\mathcal{D}} \mathbb{E}_{\xi\sim F_\xi}\left[\sum_{k\in\mathcal{K}} h'(p_k, z^k, \xi)\right] && \text{(19a)}\\
\text{subject to} \quad & C_x z^k \leq d_x p_k \quad \forall k\in\mathcal{K} && \text{(19b)}\\
& P_{\mathbb{Z}} z^k = \bar{x}_1^k p_k \quad \forall k\in\mathcal{K} && \text{(19c)}\\
& \sum_{k\in\mathcal{K}} p_k = 1, && \text{(19d)}
\end{aligned}$$

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.,

$$\begin{aligned}
h'(p_k, z^k, \xi) := \min_{y'} \;\; & c_2^\intercal y' \\
\text{s.t.} \;\; & Ay' \geq W(\xi) z^k + b p_k.
\end{aligned}$$

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 5.4, one can obtain the following large-scale convex optimization problem when employing the Wasserstein ambiguity set $\mathcal{D}(\widehat{F}_\xi, \epsilon)$:

$$\begin{aligned}
\underset{p\geq 0,\{z^k\}_{k\in\mathcal{K}},\lambda\geq 0,\{t_\omega\}_{\omega\in\Omega}}{\text{minimize}} \quad & \sum_{k\in\mathcal{K}} c_1^\intercal z^k + \lambda\epsilon + \frac{1}{|\Omega|}\sum_{\omega\in\Omega} t_\omega \\
\text{subject to} \quad & \sum_{k\in\mathcal{K}} h'(p_k, z^k, \xi^{h_\omega}) - \zeta^{h_\omega}\lambda \leq t_\omega \quad \forall h_\omega\in\mathcal{H}_\omega,\ \omega\in\Omega \\
& C_x z^k \leq d_x p_k \quad \forall k\in\mathcal{K} \\
& P_{\mathbb{Z}} z^k = \bar{x}_1^k p_k \quad \forall k\in\mathcal{K} \\
& \sum_{k\in\mathcal{K}} p_k = 1.
\end{aligned}$$

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

$$\begin{aligned}
[\texttt{Primal}'(\mathcal{K}',\mathcal{H})]: \quad \underset{p\geq 0,\{z^k\}_{k\in\mathcal{K}'},\lambda\geq 0,\{t_\omega\}_{\omega\in\Omega}}{\text{minimize}} \quad & \sum_{k\in\mathcal{K}'} c_1^\intercal z^k + \lambda\epsilon + \frac{1}{|\Omega|}\sum_{\omega\in\Omega} t_\omega \\
\text{subject to} \quad & \sum_{k\in\mathcal{K}'} h'(p_k, z^k, \xi^{h_\omega}) - \zeta^{h_\omega}\lambda \leq t_\omega \quad \forall h_\omega\in\mathcal{H}_\omega,\ \omega\in\Omega \\
& C_x z^k \leq d_x p_k \quad \forall k\in\mathcal{K}' \\
& P_{\mathbb{Z}} z^k = \bar{x}_1^k p_k \quad \forall k\in\mathcal{K}' \\
& \sum_{k\in\mathcal{K}'} p_k = 1,
\end{aligned}$$

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

$$\begin{array}{lll}
\underset{\substack{p\geq 0,\lambda\geq 0,\{t_\omega\}_{\omega\in\Omega},\{z^k\}_{k\in\mathcal{K}'}\\ \{y_{k,h_\omega}\}_{k\in\mathcal{K}',h_\omega\in\mathcal{H}'_\omega,\omega\in\Omega}}}{\text{minimize}} & \sum_{k\in\mathcal{K}'} c_1^\intercal z^k + \lambda\epsilon + \frac{1}{|\Omega|}\sum_{\omega\in\Omega} t_\omega & \\
\text{subject to} & \sum_{k\in\mathcal{K}'} c_2^\intercal y_{k,h_\omega} - \zeta^{h_\omega}\lambda \leq t_\omega & \forall h_\omega \in \mathcal{H}'_\omega,\ \omega\in\Omega \\
& Ay_{k,h_\omega} \geq W(\xi_{h_\omega})z^k + bp_k & \forall k\in\mathcal{K}',\ h_\omega\in\mathcal{H}'_\omega,\ \omega\in\Omega \\
& C_x z^k \leq d_x p_k & \forall k\in\mathcal{K}' \\
& P_\mathbb{Z} z^k = \bar{x}_1^k p_k & \forall k\in\mathcal{K}' \\
& \sum_{k\in\mathcal{K}'} p_k = 1, &
\end{array}$$

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 (18).

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 D) that the dual problem takes the form:

$$\begin{aligned}
[\texttt{Dual}'(\mathcal{K},\mathcal{H}')]:\quad & \underset{w,q\geq 0}{\text{maximize}}\quad w && && \text{(20a)}\\
& \text{subject to}\quad w \leq \min_{x^k\in\mathcal{X}_k} c_1^\intercal x^k + \sum_{\omega\in\Omega}\sum_{h_\omega\in\mathcal{H}'_\omega} h(x^k,\xi^{h_\omega})q_{h_\omega} && \forall k\in\mathcal{K} \quad (p_k) && \text{(20b)}\\
& \sum_{\omega\in\Omega}\sum_{h_\omega\in\mathcal{H}'_\omega} \zeta^{h_\omega} q_{h_\omega} \leq \epsilon && (\lambda) && \text{(20c)}\\
& \sum_{h_\omega\in\mathcal{H}'_\omega} q_{h_\omega} = \frac{1}{|\Omega|} && \forall\omega\in\Omega,\quad (t_\omega) && \text{(20d)}
\end{aligned}$$

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 (20b) by the dual problem associated to:

$$\begin{array}{lll}
\underset{x\in\mathcal{X}_k,\{y^\omega_{h_\omega}\}_{h_\omega\in\mathcal{H}_\omega,\omega\in\Omega}}{\text{minimize}} & c_1^\intercal x + \sum_{\omega\in\Omega}\sum_{h_\omega\in\mathcal{H}'_\omega} q_{h_\omega} c_2^\intercal y^\omega_{h_\omega} & \\
\text{subject to} & Ay^\omega_{h_\omega} \geq W(\xi^{h_\omega})x + b & \forall h_\omega\in\mathcal{H}'_\omega,\ \omega\in\Omega\,.
\end{array}$$

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 applications 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 Wasserstein ball and a polyhedral support set $\Xi := \{\xi \mid C_\xi\xi \leq d_\xi\}$.

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:

$$\underset{x\in\mathcal{X}_{\text{AP}}}{\text{minimize}} \sup_{F_\xi\in\mathcal{D}} \mathbb{E}_{\xi\sim F_\xi}\left[\sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}}\xi_{ij}x_{ij}\right] \tag{21}$$

where

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

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

Corollary 1. *For the DRAP presented in equation* (21)*, 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

$$\mathcal{X}'_{AP} := \left\{ x\in[0,\,1]^{|\mathcal{I}|\times|\mathcal{J}|} \;\middle|\; \begin{array}{l} \sum\limits_{j\in\mathcal{J}} x_{ij} = 1 \;\; \forall\, i\in\mathcal{I} \\ \sum\limits_{i\in\mathcal{I}} x_{ij} = 1 \;\; \forall\, j\in\mathcal{J} \end{array} \right\}; \tag{22}$$

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

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

Corollary 1 follows straightforwardly from Proposition 2, which 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 (2018), one can reformulate DRAP as the mixed-integer program

$$\begin{aligned}
&\underset{x\in\mathcal{X}_{\text{AP}},\lambda,\{t_\omega\}_{\omega\in\Omega},\{\nu_\omega\}_{\omega\in\Omega}}{\text{minimize}} && \lambda\epsilon + \frac{1}{|\Omega|}\sum_{\omega\in\Omega} t_\omega && && \text{(23a)}\\
&\text{subject to} && \sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}} \widehat{\xi}_{ij\omega}x_{ij} + \left(d - C\widehat{\xi}_\omega\right)^\intercal \nu_\omega \le t_\omega && \forall\omega\in\Omega && \text{(23b)}\\
& && \|C^\intercal\nu_\omega - x\|_\infty \le \lambda && \forall\omega\in\Omega && \text{(23c)}\\
& && \nu_\omega \ge 0 && \forall\omega\in\Omega\,, && \text{(23d)}
\end{aligned}$$

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

We can also use Corollary 1 to find the optimal value of the randomized strategy problem. Let $x^* = \arg\min\limits_{x\in\mathcal{X}'_{\text{AP}}} \delta^*\left(x|\mathcal{U}\right)$ be the optimal solution of problem (23) with $\mathcal{X}_{\text{AP}}$ replaced with $\mathcal{X}'_{\text{AP}}$. As discussed in Section 4, the optimal randomized strategy can then be found by solving the problem:

$$\begin{aligned}
&\underset{p\ge 0}{\text{minimize}} && \|x^* - \sum_{k\in\mathcal{K}} p_k\bar{x}^k\|_1\\
&\text{subject to} && \sum_{k\in\mathcal{K}} p_k = 1\,.
\end{aligned}$$

This problem can be rewritten as

$$
\begin{aligned}
\underset{p\geq 0,\theta}{\text{minimize}} \quad & \sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}}\theta_{ij} && && \text{(24a)}\\
\text{subject to} \quad & \theta_{ij}+\sum_{k\in\mathcal{K}}p_k\bar{x}^k_{ij}\geq x^*_{ij} && \forall i\in\mathcal{I}, j\in\mathcal{J} \quad && (\psi^+_{ij}) \quad \text{(24b)}\\
& \theta_{ij}-\sum_{k\in\mathcal{K}}p_k\bar{x}^k_{ij}\geq -x^*_{ij} && \forall i\in\mathcal{I}, j\in\mathcal{J} \quad && (\psi^-_{ij}) \quad \text{(24c)}\\
& \sum_{k\in\mathcal{K}}p_k=1, && && (\phi) \quad \text{(24d)}
\end{aligned}
$$

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

$$
\begin{aligned}
\underset{\psi^+\geq 0,\psi^-\geq 0,\phi}{\text{maximize}} \quad & \sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}}(\psi^+_{ij}-\psi^-_{ij})x^*_{ij}-\phi\\
\text{subject to} \quad & \sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}}\left(\psi^+_{ij}-\psi^-_{ij}\right)\bar{x}^k_{ij}\leq\phi && \forall k\in\mathcal{K} && (p_k)\\
& \psi^+_{ij}+\psi^-_{ij}\leq 1 && \forall i\in\mathcal{I},\, \forall j\in\mathcal{J}, && (\theta_{ij})
\end{aligned}
$$

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 (24) to obtain an upper bound. Then, the dual variables ψ^+ and ψ^- are used in the subproblem $\max\limits_{x\in\mathcal{X}_{AP}}\left(\psi^+-\psi^-\right)^\intercal x$ to generate a new extreme point $\bar{x}^k$. 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_j denotes the decision to open a facility at location j, while each y_{ij} denotes the decision to serve the demand at node i using the facility opened at j, in particular

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

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 DRUFLP, in which both the locations 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:

$$
\underset{p_k\geq 0,\sum\limits_{k\in\mathcal{K}}p_k=1}{\text{minimize}}\ \sup_{F_\xi\in\mathcal{D}}\mathbb{E}_{\xi\sim F_\xi}\left[\sum_{k\in\mathcal{K}}\left(\sum_{j\in\mathcal{J}}f_jx^k_j+\sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}}\xi_ic_{ij}y^k_{ij}\right)p_k\right], \tag{25}
$$

where $\mathcal{K}$ is the index set of feasible (x,y) pairs in $\mathcal{X}_{\text{UFLP}}$. One can implement the column generation algorithm described in Section 5.3 to solve this problem. In every iteration, we solve the master problem (14) with a partial set $\mathcal{K}'\subseteq\mathcal{K}$, and with the terms $c_1^\intercal x^k$ and $\xi^\intercal C_2x^k$ replaced, respectively, with $\sum\limits_{j\in\mathcal{J}}f_jx^k_j$ and $\sum\limits_{j\in\mathcal{J}}\xi_ic_{ij}y_{ij}$.

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:

$$\min_{p:p\geq 0, \sum\limits_{k\in\mathcal{K}} p_k=1} \sum_{k\in\mathcal{K}}\sum_{j\in\mathcal{J}} f_j x_j^k p_k + \sup_{F_\xi\in\mathcal{D}} \mathbb{E}_{\xi\sim F_\xi}\left[\sum_{k\in\mathcal{K}} h(x^k,\xi)p_k\right], \tag{26}$$

where

$$h(x,\xi) := \min_{z\geq 0} \quad \sum_{i\in\mathcal{I}}\sum_{j\in\mathcal{J}} c_{ij}z_{ij} \tag{27a}$$

$$\text{s.t.} \quad \sum_{j\in\mathcal{J}} z_{ij} = \xi_i \qquad \forall i\in\mathcal{I} \tag{27b}$$

$$\sum_{i\in\mathcal{I}} z_{ij} \leq v_j x_j \qquad \forall j\in\mathcal{J}, \tag{27c}$$

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 Capacitated Facility Location Problem 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 (26) and (27) 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 5.4 can be used to solve it. The implementation details for solving a DRCFLP with a Wasserstein ambiguity set are provided in Appendix E.

7 Numerical Results

We conducted a series of numerical experiments to assess the quality of the bounds proposed in Section 4 and the numerical performance of the solution algorithms presented in Section 5 on the three applications presented in Section 6. 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. In all algorithms, we set the optimality tolerance to $\varepsilon=0.02$. We further refer the reader to Appendix F for additional experiments that investigate the effect of increasing the sample size (*i.e.,* the support size of the empirical distribution $\widehat{F}_\xi^\Omega$) on the computational times of the algorithms and the solution support sizes. The Matlab codes used to generate the results can be found on GitHub[1] .

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 test instance,

[1]GitHub repository: `https://github.com/AhmedSaifDal/ValueRandomizedSolutions`.

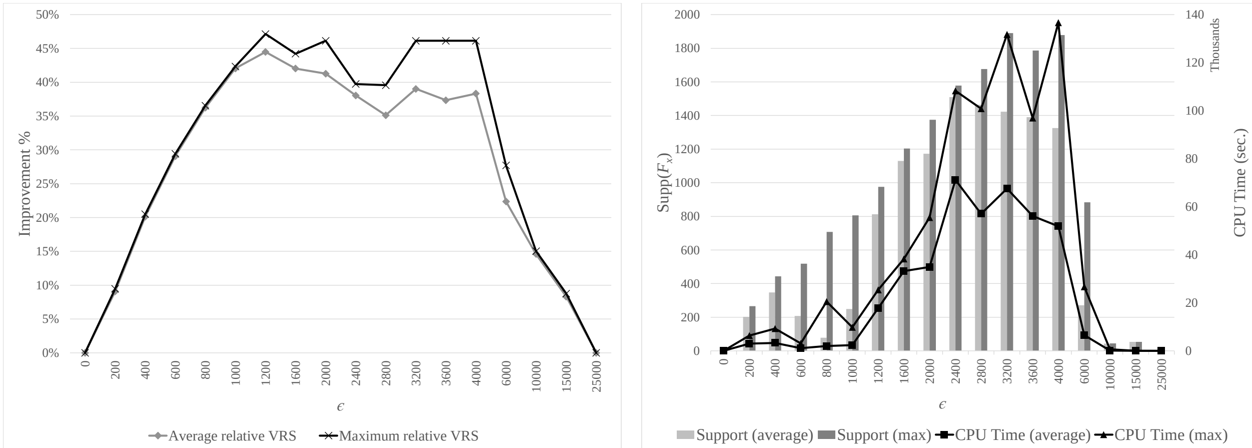

(a) (b)

Figure 2: (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.

the DRAP with a deterministic strategy was first reformulated into a mixed-integer linear program using the approach proposed in Mohajerin Esfahani and Kuhn (2018, Corollary 5.1) and 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 2-(a) presents statistics about the relative improvement achieved by a randomized strategy compared to a deterministic strategy on the 10 test 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 2-(a), one should notice that when the ambiguity set contains only the empirical distribution (*i.e.,* $\epsilon = 0$), there is no value in randomization. This is due to the fact that the problem reduces to a simple expected value minimization problem (with a known distribution) which is known to be "randomization-proof" (see Definition 8 in Delage et al. (2019)). 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 problem reduces to a deterministic assignment problem with $\xi_{ij} = \xi_{ij}^{nom}(1+\Delta_{ij})$, and randomization becomes ineffective for the same reason. Between these two extreme cases, we can confirm that employing a randomized strategy can lead to a significant reduction in the 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 2-(b) presents statistics about the support size 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 size seems quite large, it is still well below the theoretical bound of 10,001 plans obtained from Proposition 2. The average computational time over all instances was 22,682 seconds and, on average, produced 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 a 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

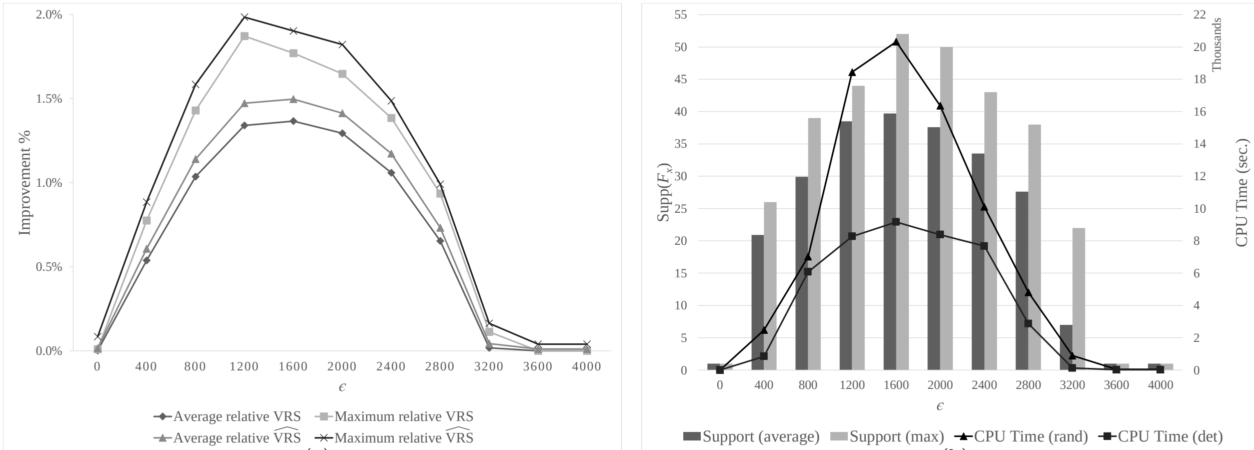

(a) (b)

Figure 3: (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.

(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.

7.2 Experiments with the DRUFLP

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

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

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 wide range of ϵ values. 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 Proposition 1, to compute the bound $\widehat{\text{VRS}}$. To solve the deterministic strategy problem, we first reformulated it into a mixed-integer linear program using the approach proposed in Mohajerin Esfahani and Kuhn (2018, Corollary 5.1).

Figure 3-(a) presents statistics about the relative $\widehat{\text{VRS}}$ bound and about the actual relative improvements achieved by the optimal randomized strategy in the 10 test instances and under different levels 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 it becomes too large. Unlike 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 Proposition 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 test instances to recognize that at $\epsilon = 0$, there was no possible improvement for randomized strategies. Even in the other 5 test instances, the identified potential of improvement was nearly inexistent (always below 0.05%). This evidence supports an observation made by Morris (1978) that solving a linear relaxation of the uncapacitated facility location problem typically leads to identifying an integer solution.

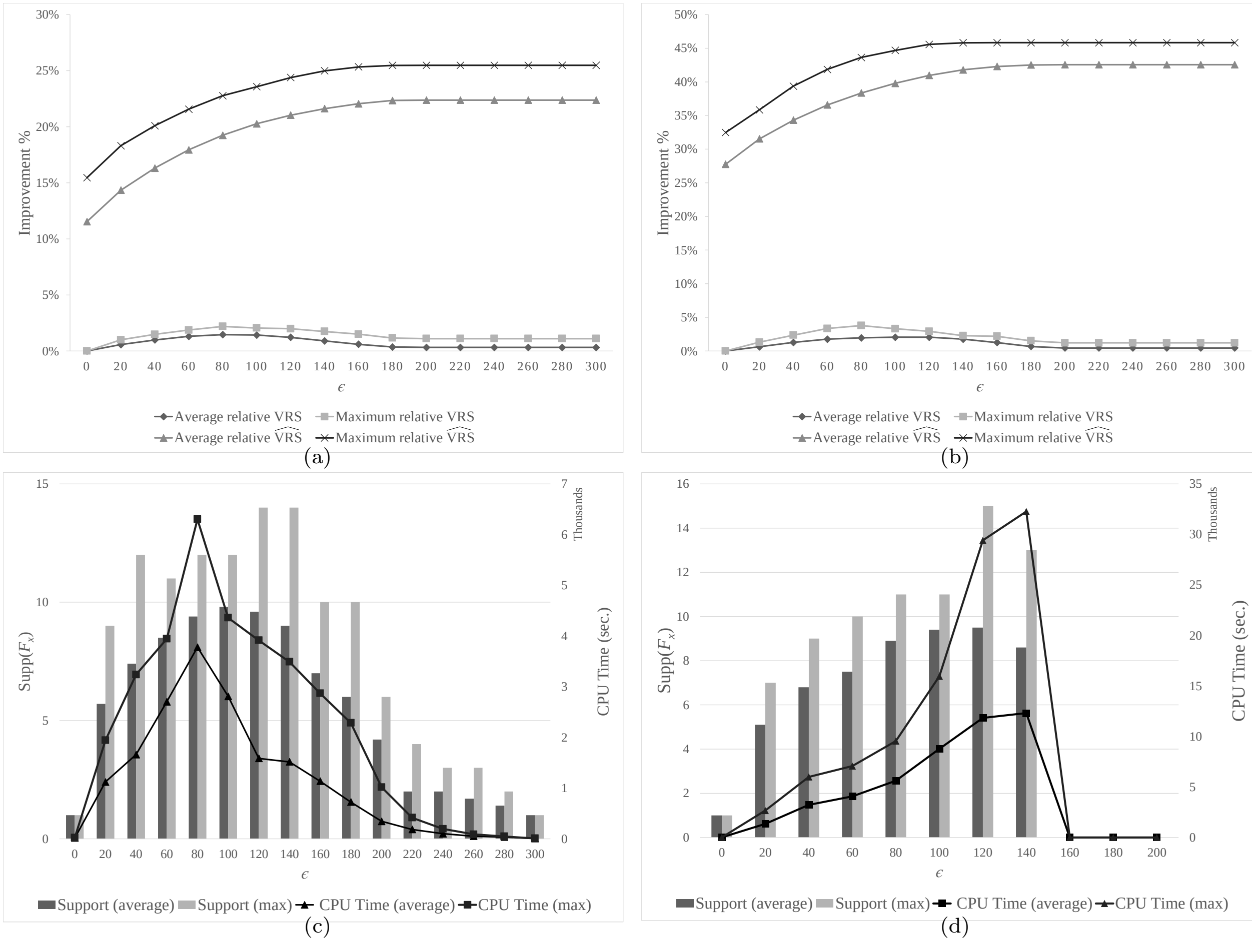

Figure 4: (a,b) Improvements due to randomization and relaxation with $r = 3$ and 5, respectively. (c,d) Optimal solution support size and computational time with $r = 3$ and 5, respectively. All based on 10 DRCFLP test instances.

In Figure 3-(b), we report on the average and the maximum support size of the optimal randomized strategies obtained for the test instances with different values of ϵ, as well as the computational times for the column generation algorithm. One can first notice that the average support size of the optimal randomized strategies is somewhat proportional to the extent of relative improvement. In fact, the largest average support size of 39.7 plans was reached at $\epsilon = 1600$, which nearly coincides with the peak performance improvement reached at $\epsilon = 1200$. One can also confirm that in all cases, the support size of the optimal randomized strategy is always far below the theoretical limit of 301 (see Proposition 2). The average computational time for the DRUFLP with randomization was 7321 seconds whereas the largest computational time among all instances was 44,000 seconds. Most of the computational time was spend on solving the binary integer programming subproblems. The tailing-off effect was also observed in this problem. 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). To solve the deterministic strategy problem, we used the algorithm proposed in Saif and Delage (2020), which is based on the column-and-constraint generation algorithm developed in Zeng and Zhao (2013) for solving two-stage RO problems.

Figures 4-(a) and (b) present statistics about the relative $\widehat{\text{VRS}}$ bound and the actual relative improvement achieved by the optimal randomized strategy in the 10 test 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 $\widehat{\text{VRS}}$ bound is, in general, a poor indicator of the maximum improvement that can be achieved by randomization for this class of problems. 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 test instances with $\epsilon \in [10,\ 300]$, and peaked at 2.05% for problems with $\epsilon = 120$ and $r = 5$.

Figures 4-(c) and (d) show the average and maximum support size of the optimal randomized strategies and the computational time for the column generation algorithm obtained for the different test instances and values of ϵ and r. Among test instances with $\epsilon \in \{20, 40, \ldots, 300\}$, the average optimal support sizes were 5.65 and 5.16 for $r = 3$ and 5, respectively, while the maximum optimal support size was 14 in both cases. This seems to indicate that the structure of optimal randomized strategies becomes slightly simpler as the capacity scarcity is reduced and that, perhaps, in practice the optimal support size remains comparable to n, although Proposition 2 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 were 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 on Gurobi were 3317 and 1870 seconds, respectively. These results clearly show that the proposed algorithm can handle reasonably sized problems quite effectively. We observed also that the bottleneck of the algorithm in terms of efficiency was in solving `SP2`.

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 continuous 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 uncapacitated and capacitated 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, traveling salesman problems, *etc.*

Although the theoretical results presented in Section 4 are general and can be used for DRO problems with any law-invariant convex risk measure and support set, the algorithms described in Section 5 are suitable only for cases when the decision maker has an AARN attitude, *i.e.,* $\rho(\cdot) = \mathbb{E}[\cdot]$, and when the support set Ξ is bounded. As a future extension, it would be interesting to see how the column generation algorithm can be modified to handle unbounded support sets. In such a case, it might be possible to reduce the two-layer column generation algorithm to a single-layer one, given that Hanasusanto and Kuhn (2018) have shown that some two-stage Wasserstein-distance-based DRO problems with no support constraints and an l_1-norm can be tractably reformulated as linear programs.

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.

Acknowledgements

The authors would like to thank the associate editor and the two reviewers for their valuable comments and suggestions. This research was partially supported by the Natural Sciences and Engineering Research Council of Canada [Grant RGPIN-2016-05208], the Canada Research Chair program [Grant 950-230057], and the Groupe d'études et de recherche en analyse des décisions.

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 locationtransportation 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, de Ruiter FJ (2016) Duality in two-stage adaptive linear optimization: Faster computation and stronger bounds. *INFORMS J. Comput.* 28(3):500–511.

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.

Boyd S, Vandenberghe L (2004) *Convex Optimization* (USA: Cambridge University Press).

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

Buhayenko V, den Hertog D (2017) Adjustable robust optimisation approach to optimise discounts for multi-period supply chain coordination under demand uncertainty. *International Journal of Production Research* 55(22):6801–6823.

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.

Dehghan S, Amjady N, Conejo AJ (2017) Adaptive robust transmission expansion planning using linear decision rules. *IEEE Trans. Power Syst.* 32(5):4024–4034.

Delage E, Kuhn D, Wiesemann W (2019) "Dice"-sion making under uncertainty: When can a random decision reduce risk? *Management Science* 65(7):3282–3301.

Delage E, Ye Y (2010) Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Operations Research* 58(3):595–612.

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

Ellsberg D (1961) Risk, ambiguity, and the savage axioms. *The Quarterly Journal of Economics* 75(4):643–669.

Epstein LG (1999) A definition of uncertainty aversion. *The Review of Economic Studies* 66(3):579–608.

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.

Fournier N, Guillin A (2015) On the rate of convergence in wasserstein distance of the empirical measure. *Probability Theory and Related Fields* 162(3):707–738.

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

Hanasusanto GA, Kuhn D (2018) Conic programming reformulations of two-stage distributionally robust linear programs over wasserstein balls. *Operations Research* 66(3):849–869.

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 (2018) Data-driven distributionally robust optimization using the Wasserstein metric: performance guarantees and tractable reformulations. *Mathematical Programming* 171:115166.

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.

Parys BPGV, Esfahani PM, Kuhn D (2017) From data to decisions: Distributionally robust optimization is optimal. *ArXiv* abs/1704.04118.

Saif A, Delage E (2020) Data-driven distributionally robust capacitated facility location problem, preprint.

Scarf H (1958) A min-max solution of an inventory problem. *Studies in The Mathematical Theory of Inventory and Production* 201–209.

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.

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.

A A Distributionally Robust Newsvendor Problem Example

In the classical newsvendor problem, a vendor needs to decide how many newspapers he should order for his stand without knowing in advance the number of customers that will be buying the product. In his seminal work, Scarf (1958) makes a hypothesis that has been heavily popularized, namely that only the mean and variance of the distribution describing the likelihood that any number of customers show up to purchase the item is known. This leads to the following so-called "min-max newsvendor" problem.

$$\max_{x \geq 0} \quad \inf_{F \in \mathcal{D}(\mathbb{R}_+, \mu, \sigma)} \mathbb{E}_{D \sim F}[r \min(x, D) - cx] \,,$$

where $x \in \mathbb{R}$ represents the number of newspapers that are ordered, D is the random number, drawn from the distribution F, of customers that will visit the stand to purchase a newspaper, r is the unit selling price, and c is the unit ordering cost. Note also that $\mathcal{D}(\mathbb{R}_+, \mu, \sigma^2)$ represents the set of all possible distributions for a non-negative random variable with mean μ and variance σ^2.

In this example, we revisit this problem with the notion of a fixed delivery cost f and a penalty cost in case of excessive loss sales. In particular, we consider the following problem:

$$\min_{z \in \{0,1\},\, 0 \leq x \leq Mz} \quad \sup_{F \in \mathcal{D}(\mathbb{R}_+, \mu, \sigma)} \mathbb{E}[cx + fz - r \min(x, D)] + p \max(D - x - \Delta,\, 0) \,,$$

where z captures the decision to make an order of newspapers, f is the fixed delivery cost, and p is the cost for every loss sales that exceeds some threshold $\Delta > 0$. The latter cost can for instance capture the fact that if more than Δ customers are unable to obtain a copy of their favorite newspaper, the reputation of the newsstand will be negatively affected. Letting the problem parameters take on the following values, $r = \$1$, $c = \$0.5$, $f = \$300$, $p = \$1$, $\mu = 1000$, $\sigma = 500$, and $\Delta = 4000$, the newsvendor needs to decide whether he will ask for a delivery of newspapers, and if so in what quantity. In this context, one can actually show numerically that if the decision is not to make an order, then the worst-case expected profit is achieved by the distribution that puts 99.3% probability on a demand of 959 units and 0.007% on a demand of 7041 units and achieves a worst-case expected loss of \$20.7. Alternatively, if he does make an order then it should optimally be of 1000 units to reach a worst-case expected loss of \$50, where the worst-case distribution puts 50% on a demand of 500 and 50% on 1500. Necessarily, this analysis should motivate the newsvendor to make no order.

Actually, there is still hope to convince this ambiguity averse newsvendor to make an order, albeit a randomized one. Let us for instance consider the randomized strategy $0 \oplus_{85\%} 1000$, *i.e.,* a random number U between 1 and 100 is uniformly drawn to select between 0, if $U \leq 85$, and 1000 units otherwise. For a decision maker that minimizes the worst-case expected cost, the risk of this randomized strategy should be computed as

$$\begin{aligned} \inf_{F \in \mathcal{D}(\mathbb{R}_+, \mu, \sigma)} \mathbb{E}[&(c1000 + f) \cdot \mathbf{1}\{U > 85\} - r \min(1000 \cdot \mathbf{1}\{U > 85\}, D) \\ &+ p \max(D - 1000 \cdot \mathbf{1}\{U > 85\} - \Delta,\, 0)]. \end{aligned} \tag{28}$$

Indeed, mathematically one can demonstrate that this strategy has a risk of \$8.2, where the worst-case expected profit is achieved by the distribution that puts 51.6% probability on 589, 48.2% on 1411, and 0.2% on 6889 units. It is therefore a strategy that outperforms the two deterministic ones described earlier. This is particularly interesting considering that with this strategy there is a 15% chance that an order is made.

Note that equation (28) captures the risk at the moment that the newsvendor commits to the randomized strategy and requires the newsvendor to follow through with the strategy once U is observed, rather than reassessing the risks at that point. This might be difficult to accept from an operational perspective. Alternatively, let us consider that the newsvendor is willing to sign a randomized contract with the supplier of the form $0 \oplus_{85\%} 1000$, *i.e.,* that the supplier will take on the responsibility of drawing U and if $U > 85$ to call the newsvendor at a convenient time and announce his visit. In this context, once the $0 \oplus_{85\%} 1000$ contract is signed, the newsvendor becomes committed to following through with it. Nevertheless, based on his ambiguity aversion the newsvendor should be convinced that such a contract constitutes the best option in terms of worst-case risk.

B Proofs

B.1 Proof of Proposition 1

Our proof exploits the extension of Theorem 2 in Delage et al. (2019) 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

$$\begin{aligned}
\min_{x\in\mathcal{X}'} \sup_{F_\xi\in\mathcal{D}} \rho_{\xi\sim F_\xi}(h(x,\xi)) &\leq \min_{x\in\mathcal{C}(\mathcal{X})} \sup_{F_\xi\in\mathcal{D}} \rho_{\xi\sim F_\xi}(h(x,\xi)) \\
&= \min_{g(\cdot)\in\mathcal{G}} \sup_{F_\xi\in\mathcal{D}} \rho_{\xi\sim F_\xi}(g(\xi)) \\
&= \min_{F_g\in\Delta(\mathcal{G})} \sup_{F_\xi\in\mathcal{D}} \rho_{(G,\xi)\sim F_g\times F_\xi}(G(\xi)) \\
&= \min_{F_g\in\Delta(\bar{\mathcal{G}})} \sup_{F_\xi\in\mathcal{D}} \rho_{(G,\xi)\sim F_g\times F_\xi}(G(\xi)) \\
&= \min_{F_x\in\Delta(\mathcal{C}(\mathcal{X}))} \sup_{F_\xi\in\mathcal{D}} \rho_{(X,\xi)\sim F_x\times F_\xi}(h(X,\xi)) \\
&\leq \min_{F_x\in\Delta(\mathcal{X})} \sup_{F_\xi\in\mathcal{D}} \rho_{(X,\xi)\sim F_x\times F_\xi}(h(X,\xi)) = v_r\,,
\end{aligned}$$

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. (2019) 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:

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

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

$$\begin{aligned}
\sup_{F_\xi\in\mathcal{D}} \mathbb{E}_{(X,\xi)\sim F_x^*\times F_\xi}[h(X,\xi)] &= \sup_{F_\xi\in\mathcal{D}} \mathbb{E}_{\xi\sim F_\xi}[h(\mathbb{E}_{X\sim F_x^*}[X],\xi)] \\
&= \min_{x\in\mathcal{C}(\mathcal{X})} \sup_{F_\xi\in\mathcal{D}} \mathbb{E}_{\xi\sim F_\xi}[h(x,\xi)] \\
&= \min_{F_x\in\Delta(\mathcal{X})} \sup_{F_\xi\in\mathcal{D}} \mathbb{E}_{(X,\xi)\sim F_x\times F_\xi}[h(X,\xi)] = v_r\,.
\end{aligned}$$

This completes our proof.

B.2 Proof of Proposition 2

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.

B.3 Proof of Proposition 3

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

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

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

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

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 3 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`.

B.4 Proof of Proposition 4

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 $\texttt{SP1}_\omega$ and linearize the norm constraint to obtain the following bilinear optimization problem:

$$\begin{aligned}
\underset{\xi,\zeta\geq 0,\delta\geq 0,\{\phi_k\}_{k\in\mathcal{K}'}}{\text{maximize}} \quad & \sum_{k\in\mathcal{K}'}(W(\xi)\bar{x}^k+b)^\intercal\phi_k-\lambda^*\zeta && (29a)\\
\text{subject to} \quad & A^\intercal\phi_k=c_2p_k^* \qquad \forall k\in\mathcal{K}' && (29b)\\
& C_\xi\xi\leq d_\xi && (29c)\\
& e^\intercal\delta\leq\zeta && (29d)\\
& \zeta\leq\zeta^{\max} && (29e)\\
& \xi-\widehat{\xi}_\omega\leq\delta && (29f)\\
& \widehat{\xi}_\omega-\xi\leq\delta, && (29g)\\
& \phi_k\geq 0 \qquad \forall k\in\mathcal{K}', && (29h)
\end{aligned}$$

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:

$$\begin{aligned}
\rho(\{\phi_k\}_{k\in\mathcal{K}'}) := \max_{\xi,\zeta\geq 0,\delta\geq 0} \quad & \sum_{k\in\mathcal{K}'}(W(\xi)\bar{x}^k+b)^\intercal\phi_k-\lambda^*\zeta && (30a)\\
\text{s.t.} \quad & C_\xi\xi\leq d_\xi \qquad (\alpha) && (30b)\\
& e^\intercal\delta\leq\zeta \qquad (\beta) && (30c)\\
& \zeta\leq\zeta^{\max} \qquad (\gamma) && (30d)\\
& \xi-\widehat{\xi}_\omega\leq\delta \qquad (\psi^+) && (30e)\\
& \widehat{\xi}_\omega-\xi\leq\delta. \qquad (\psi^-), && (30f)
\end{aligned}$$

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

$$\begin{aligned}
\rho(\{\phi_k\}_{k\in\mathcal{K}'}) = \min_{\substack{\alpha\geq 0,\beta\geq 0,\gamma\geq 0\\ \psi^+\geq 0,\psi^-\geq 0}} \quad & \sum_{k\in\mathcal{K}'}(W_0\bar{x}^k+b)^\intercal\phi_k+d_\xi^\intercal\alpha+\zeta^{\max}\gamma+\widehat{\xi}_\omega^\intercal(\psi^+-\psi^-) && (31a)\\
\text{s.t.} \quad & \sum_{i=1}^m\left(\sum_{k\in\mathcal{K}'}\phi_k^TW_i\bar{x}^k\right)e_i=C_\xi^T\alpha+\psi^+-\psi^- && (31b)\\
& \beta\leq\lambda^*+\gamma && (31c)\\
& \psi^++\psi^-\leq\beta && (31d)
\end{aligned}$$

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 (30). 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 (31). 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 $\texttt{SP1}_\omega$. In some preliminary experiments, we established that our chosen approach had a significant impact on reducing the solution time for $\texttt{SP1}_\omega$.

B.5 Proof of Theorem 1

The proof revolves around the fact that in each iteration, the algorithm either terminates with $UB-LB\leq\varepsilon$ 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.

B.6 Proof of Proposition 5

Since one can verify that $h(x,\xi) := c_1^\intercal x + h(x,\xi)$ is convex in x, based on Proposition 3, we have that the RSP reduces to

$$\begin{aligned}
\underset{p\in\mathbb{R}^{|\mathcal{K}|},\{x^k\}_{k\in\mathcal{K}}}{\text{minimize}} \quad & \sum_{k\in\mathcal{K}} c_1^\intercal p_k x^k + \sup_{F_\xi\in\mathcal{D}} \mathbb{E}_{\xi\sim F_\xi}\left[\sum_{K\in\mathcal{K}} p_k h'(x^k,\xi)\right] \\
\text{subject to} \quad & C_x x^k \leq d_x \,,\ \forall k\in\mathcal{K} \\
& P_{\mathbb{Z}} x^k = \bar{x}_1^k \,,\ \forall k\in\mathcal{K} \\
& p_k \geq 0 \ \ \forall\, k\in\mathcal{K}, \quad \sum_{k\in\mathcal{K}} p_k = 1\,.
\end{aligned}$$

Using a simple change of variable $z_k := x_k p_k$ and $y' := y p_k$, we obtain the reduction presented in problem (19) by exploiting the fact that $\bar{\mathcal{X}} := \{x\in\mathbb{R}^n \,|\, C_x x \leq d_x\}$ 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 $\xi\in\Xi$.

B.7 Proof of Corollary 1

Based on Theorem 1, we can conclude that 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 (22) 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.

C 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**
  `// Solve Primal`$(\mathcal{K}',\mathcal{H})$:
  $LB_1\longleftarrow LB$
  $UB_1\longleftarrow\infty$
  while $UB_1\geq LB_1$ **do**
    $\forall\omega\in\Omega$, solve $\texttt{SP1}_\omega$ to get a new vertex $(\xi^{\bar{h}_\omega},\zeta^{\bar{h}_\omega})$ and t_ω^*
    $\mathcal{H}'_\omega\longleftarrow\mathcal{H}'_\omega\cup\{\bar{h}_\omega\}$for all $\omega\in\Omega$
    $UB_1=\min(UB_1,\sum_{k\in\mathcal{K}}c_1^\intercal\bar{x}^kp_k^*+\lambda^*\epsilon+\frac{1}{|\Omega|}\sum_{\omega\in\Omega}t_\omega^*)$
    Solve `MP1` to obtain new $\{p_k^*\}_{k\in\mathcal{K}'}$ and λ^* and update LB_1
  end
  $UB\longleftarrow UB_1$
  `// Solve Dual`$(\mathcal{K},\mathcal{H}')$:
  $UB_2\longleftarrow UB$
  $LB_2\longleftarrow-\infty$
  while $UB_2\geq LB_2$ **do**
    Solve `SP2` to get a new solution $x^{\bar{k}}$ and the optimal value w^*
    $\mathcal{K}'\longleftarrow\mathcal{K}'\cup\{\bar{k}\}$
    $LB_2=\max(LB_2,w^*)$
    Solve `MP2` to obtain new $q_{h_\omega}^*$ and update UB_2
  end
  $LB\longleftarrow LB_2$
end

Algorithm 1: The Two-layer Column-Generation Algorithm

Figure 5: The Two-Layer column generation Algorithm.

D Derivation of Dual Problem (20)

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

$$
\begin{aligned}
&\min_{p\geq 0,\lambda\geq 0}\ \min_{z^1\in\mathcal{Z}_1(p_1),\ldots,z^{|K|}\in\mathcal{Z}_{|K|}(p_{|K|})} \lambda\epsilon+\sum_{k\in\mathcal{K}}c_1^\intercal z^k+\max_{w\geq 0} w\left(1-\sum_{k\in\mathcal{K}}p_k\right)\\
&\qquad\qquad +\frac{1}{|\Omega|}\sum_{\omega\in\Omega}\max_{h_\omega\in\mathcal{H}_\omega}\sum_{k\in\mathcal{K}}h'(p_k,z^k,\xi^{h_\omega})-\zeta^{h_\omega}\lambda\\
&=\min_{p\geq 0,\lambda\geq 0}\ \min_{x^1\in\mathcal{X}_1,\ldots,x^{|\mathcal{K}|}\in\mathcal{X}_{|\mathcal{K}|}} \lambda\epsilon+\sum_{k\in\mathcal{K}}c_1^\intercal x^k p_k+\max_{w\geq 0} w\left(1-\sum_{k\in\mathcal{K}}p_k\right)\\
&\qquad\qquad +\frac{1}{|\Omega|}\sum_{\omega\in\Omega}\max_{q^\omega\in\mathcal{Q}_\omega}\sum_{h_\omega\in\mathcal{H}_m}q^\omega_{h_\omega}\sum_{k\in\mathcal{K}}p_k h(x^k,\xi^{h_\omega})-\zeta^{h_\omega}\lambda\\
&=\min_{p\geq 0,\lambda\geq 0}\ \max_{w\geq 0,q^1\in\mathcal{Q}_1,\ldots,q^{|\Omega|}\in\mathcal{Q}_{|\Omega|}} \lambda\epsilon-\frac{1}{|\Omega|}\sum_{\omega\in\Omega}\sum_{h_\omega\in\mathcal{H}_\omega}q^\omega_{h_\omega}\zeta^{h_\omega}\lambda+w\left(1-\sum_{k\in\mathcal{K}}p_k\right)\\
&\qquad\qquad +\sum_{k\in\mathcal{K}}p_k\left(\min_{x^k\in\mathcal{X}_k}c_1^\intercal x^k+\frac{1}{|\Omega|}\sum_{\omega\in\Omega}\sum_{h_\omega\in\mathcal{H}_\omega}q^\omega_{h_\omega}h(x^k,\xi^{h_\omega})\right)\\
&\geq \max_{w\geq 0,q^1\in\mathcal{Q}_1,\ldots,q^{|\Omega|}\in\mathcal{Q}_{|\Omega|}}\ \min_{p\geq 0,\lambda\geq 0} \lambda\epsilon-\frac{1}{|\Omega|}\sum_{\omega\in\Omega}\sum_{h_\omega\in\mathcal{H}_\omega}q^\omega_{h_\omega}\zeta^{h_\omega}\lambda+w\left(1-\sum_{k\in\mathcal{K}}p_k\right)\\
&\qquad\qquad +\sum_{k\in\mathcal{K}}p_k\left(\min_{x^k\in\mathcal{X}_k}c_1^\intercal x^k+\frac{1}{|\Omega|}\sum_{\omega\in\Omega}\sum_{h_\omega\in\mathcal{H}_\omega}q^\omega_{h_\omega}h(x^k,\xi^{h_\omega})\right),
\end{aligned}
$$

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 $\texttt{Dual}'(\mathcal{K},\mathcal{H})$ by replacing ${q^\omega_{h_\omega}}':=(1/|\Omega|)q^\omega_{h_\omega}$. 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 $\texttt{Dual}'(\mathcal{K},\mathcal{H})$ 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 $\mathcal{Q}:=\{\{q^\omega\}_{\omega\in\Omega}\,|\,\sum_{\omega\in\Omega}\sum_{h_\omega\in\mathcal{H}_\omega}\zeta^{h_\omega}q^\omega_{h_\omega}\leq\epsilon,\ \sum_{h_\omega\in\mathcal{H}'_\omega}q^\omega_{h_\omega}=\frac{1}{|\Omega|},\ q^\omega\geq 0,\ \forall\omega\in\Omega\}$ has a strict interior point.*

Proof. Proof. We instead demonstrate that

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

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 $\{q^\omega_{h_\omega}\}_{h_\omega\in\mathcal{H}_\omega}$ such that $\xi'_\omega=\sum_{h_\omega\in\mathcal{H}_\omega}\xi^{h_\omega}q^\omega_{h_\omega}$ and $\zeta'_\omega=\sum_{h_\omega\in\mathcal{H}_\omega}\zeta^{h_\omega}q^\omega_{h_\omega}$. 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 $\mu^\xi_\omega:=(1/|\mathcal{H}_\omega|)\sum_{h_\omega\in\mathcal{H}_\omega}\xi^{h_\omega}$ and $\mu^\zeta_\omega:=(1/|\mathcal{H}_\omega|)\sum_{h_\omega\in\mathcal{H}_\omega}\zeta^{h_\omega}$ and identify some $(\nu^\xi_\omega,\nu^\zeta_\omega)\in\Xi'_\omega$ such that $(\xi'_\omega,\zeta'_\omega)$ is the convex combination of $(\mu^\xi_\omega,\mu^\zeta_\omega)$ and $(\nu^\xi_\omega,\nu^\zeta_\omega)$. This is always possible since $(\xi'_\omega,\zeta'_\omega)$ is in the relative interior $\mathcal{H}_\omega$. The convex combination of the representations for $(\mu^{\xi}_{\omega}, \mu^{\zeta}_{\omega})$ and $(\nu^{\xi}_{\omega}, \nu^{\zeta}_{\omega})$ 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 $\{q^{\omega}_{h_\omega}\}_{h_\omega \in \mathcal{H}_\omega}$ is such that

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

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

E Solving the DRCFLP with Randomization

As described in Section 5.4, 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, `MP1` and `MP2`, and two subproblems, $\texttt{SP1}_{\omega}$ and `SP2`. For completeness, we briefly describe the details of these problems below.

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

$$\begin{aligned}
[\texttt{MP1}]: \quad & \underset{p \geq 0, \lambda \geq 0, \{t_\omega\}_{\omega \in \Omega}}{\text{minimize}} && \sum_{k \in \mathcal{K}'} \sum_{j \in \mathcal{J}} f_j x^k_j p_k + \lambda \epsilon + \frac{1}{|\Omega|} \sum_{\omega \in \Omega} t_\omega \\
& \text{subject to} && \sum_{k \in \mathcal{K}'} \sum_{j \in \mathcal{J}} \sum_{i \in \mathcal{I}} c_{ij} z_{ijk} p_k - \zeta^{h_\omega} \lambda \leq t_\omega \quad \forall \omega \in \Omega,\ h_\omega \in \mathcal{H}_\omega \\
& && \sum_{k \in \mathcal{K}'} p_k = 1.
\end{aligned}$$

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

$$\begin{aligned}
[\texttt{SP1}_{\omega}]: \quad \underset{(\xi, \zeta, \delta) \in \Upsilon}{\text{maximize}} \quad \min_{z \geq 0} \quad & \sum_{k \in \mathcal{K}'} \sum_{j \in \mathcal{J}} \sum_{i \in \mathcal{I}} p^*_k c_{ij} z_{ijk} - \lambda^* \zeta \\
\text{s.t.} \ & \sum_{j \in \mathcal{J}} z_{ijk} = \xi_i && \forall i \in \mathcal{I}, k \in \mathcal{K}' \quad (\nu_{ik}) \\
& v_j x^k_j - \sum_{i \in \mathcal{I}} z_{ijk} \geq 0 && \forall j \in \mathcal{J}, k \in \mathcal{K}' \quad (\mu_{jk} \geq 0)\,,
\end{aligned}$$

where $\Upsilon := \{(\xi, \zeta, \delta) \in \mathbb{R}^{|\mathcal{I}|} \times \mathbb{R} \times \mathbb{R}^{|\mathcal{I}|} \,|\, \zeta \geq 0, (29c) - (29g)\}$. Following a similar procedure as used in the proof of Proposition 4, we obtain the following equivalent mixed-integer linear program:

$$\begin{aligned}
& \underset{\substack{\xi, \zeta, \delta, \alpha, \beta, \psi, \nu, \mu \geq 0 \\ Bin^1, Bin^2, Bin^3, Bin^4, Bin^5, Bin^6, Bin^7}}{\text{maximize}} && d_\xi^\intercal \alpha + \bar{\zeta}_{max} \gamma + \sum_{i \in \mathcal{I}} \widehat{\xi}^{\omega}_i (\psi^+_i - \psi^-_i) - \sum_{k \in \mathcal{K}'} \sum_{j \in \mathcal{J}} v_j x^k_j \mu_{jk} \\
& \text{subject to} && \nu_{ik} - \mu_{jk} \leq c_{ij} p^*_k \quad \forall i \in \mathcal{I}, j \in \mathcal{J}, k \in \mathcal{K}' \\
& && (18b) - (18q).
\end{aligned}$$

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

$$\begin{aligned}
[\texttt{MP2}]: \quad & \underset{w, q \geq 0}{\text{maximize}} && w \\
& \text{subject to} && w \leq \sum_{k \in \mathcal{K}} \sum_{j \in \mathcal{J}} \bar{x}^k + \sum_{\omega \in \Omega} \sum_{h_\omega \in \mathcal{H}'_\omega} h(\bar{x}^k, \xi^{h_\omega}) q_{h_\omega} \quad \forall k \in \mathcal{K} \quad (p_k) \\
& && \sum_{\omega \in \Omega} \sum_{h_\omega \in \mathcal{H}'_\omega} \zeta^{h_\omega} q_{h_\omega} \leq \epsilon \quad (\lambda) \\
& && \sum_{h_\omega \in \mathcal{H}'_\omega} q_{h_\omega} = \frac{1}{|\Omega|} \quad \forall \omega \in \Omega \quad (t_\omega),
\end{aligned}$$

while its associated subproblem reduces to

$$\begin{aligned}
[\texttt{SP2}]: \quad & \underset{x \in \mathcal{X}, z \geq 0}{\text{minimize}} && \sum_{j \in \mathcal{J}} f_j x_j + \frac{1}{|\Omega|} \sum_{\omega \in \Omega} \sum_{j \in \mathcal{J}} \sum_{i \in \mathcal{I}} c_{ij} z_{ijh_\omega} \\
& \text{subject to} && \sum_{j \in \mathcal{J}} z_{ijh_\omega} \geq \xi^{h_\omega}_i \quad \forall i \in \mathcal{I},\ h_\omega \in \mathcal{H}'_\omega \\
& && \sum_{i \in \mathcal{I}} z_{ijh_\omega} \leq v_j x_j \quad \forall j \in \mathcal{J},\ h_\omega \in \mathcal{H}'_\omega \,.
\end{aligned}$$

F Effect of sample size on numerical efficiency and solution's support

In this appendix, we investigate the effect of increasing the sample size (*i.e.,* the support size of the empirical distribution $\widehat{F}_{\xi}^{\Omega}$) on the computational times of the algorithms and the solution support sizes. For this round of experiments, we tested on DRAP instances of size $|\mathcal{I}| = |\mathcal{J}| = 50$, DRUFLP instances of size $|\mathcal{I}| = |\mathcal{J}| = 150$ and DRCFLP instances of size $|\mathcal{I}| = |\mathcal{J}| = 10$ with $r = 5$, while keeping all other parameters unchanged. Smaller instances, compared to previous experiments, were used so they could be solved in reasonable times even when much larger sample sizes were used. Specifically, we tested with sample sizes of $|\Omega| = 10, 20, 40, 80, 160$ and 320. Figure 6 shows the results for the three problems with different values of ϵ. The line charts on the left-hand side show the average computational time in seconds, whereas the bar charts on the right-hand side show the average support size of the optimal randomized strategies.

It is clear that for the DRAP instances, the sample size has, virtually, no effect on the computational performance of the algorithm. This observation is easily explainable by the fact that neither the master problem (24) nor the subproblem depends on $|\Omega|$. Conversely, the computational times for the DRUFLP and the DRCFLP significantly increased as larger sample sizes were used, with a dramatically larger effect on the latter. In the case of the DRUFLP, increasing the sample size from 10 to 320 increased the computational time 10-fold on average. In comparison, the computational time increase in the case of the DRCFLP was about 190-fold, on average, for the same increase in the sample size. A possible explanation is that in the DRCFLP, only one subproblem is solved in every iteration and the size of this subproblem does not depend on the sample size, whereas in the DRCFLP we solve $|\Omega|$ primal subproblems, *i.e.,* $\texttt{SP1}_{\omega}$, and a dual subproblem $\texttt{SP2}$ whose size depends on $|\Omega|$. On the other hand, the effect of sample size on the optimal support size in all problems seems insignificant.

One can also easily notice the significant effect of problem size, *i.e.,* $|\mathcal{I}|$ and $|\mathcal{J}|$, on the CPU time. This is to be expected in integer programming problems which are known to be intractable even in the deterministic case. Surprisingly, in many cases (especially for DRAP) and despite the high solution time experienced, solving the randomized strategy problems using the proposed algorithms was easier than solving their corresponding deterministic strategy problems directly on Gurobi. Indeed, one way to improve the computational performance of the proposed algorithms could be to use a more efficient programming language than Matlab like C++ or Julia. We performed some preliminary tests on the DRAP instances using Julia and the results were promising as the computational times were cut by more than half.

G On the difficulty of adoption of randomized strategies

Although randomized/mixed strategies are not entirely new and have been proposed in the game theory literature for decades, practitioners might find the idea of utilizing them with a passive adversary (*i.e.,* nature) as a way to protect against distributional ambiguity unappealing. This skeptical view towards randomization might be caused by the classical rationalization which suggests that randomization can be used to "trick" or "bluff" a malevolent opponent. Obviously, in DRO this rationalization does not apply anymore. Moreover, decision makers might be more reluctant to apply randomized strategies in "one-off" decisions, as opposed to repeated ones where the distribution of actions can be interpreted as frequencies. This feeling might be especially strong when the amounts at stake are significant or when the decisions have long-term implications. Another strong obstacle to the adoption of randomized strategies could come from the operational burden of having to prepare for the different actions that are part of the randomized strategy until the realized action is revealed. In this regard, one might favor revealing the random action early at the price of a risk guarantee that is more difficult to interpret: *i.e.,* as the risk that was faced prior to revealing the action. In fact, this interpretation of risk is exactly the same as what is offered for the statistical estimation of confidence intervals.

While we consider beyond the purpose of this work to provide irrefutable arguments that would convince such reluctant adopters, we list below some situations in which decision makers might find randomized strategies rather appealing.

1. Situations where the value of randomization is large compared to the worst-case expected cost for pure strategies (see Section 7.1 where approximately 50% reduction in risk could be obtained for the DRAP);

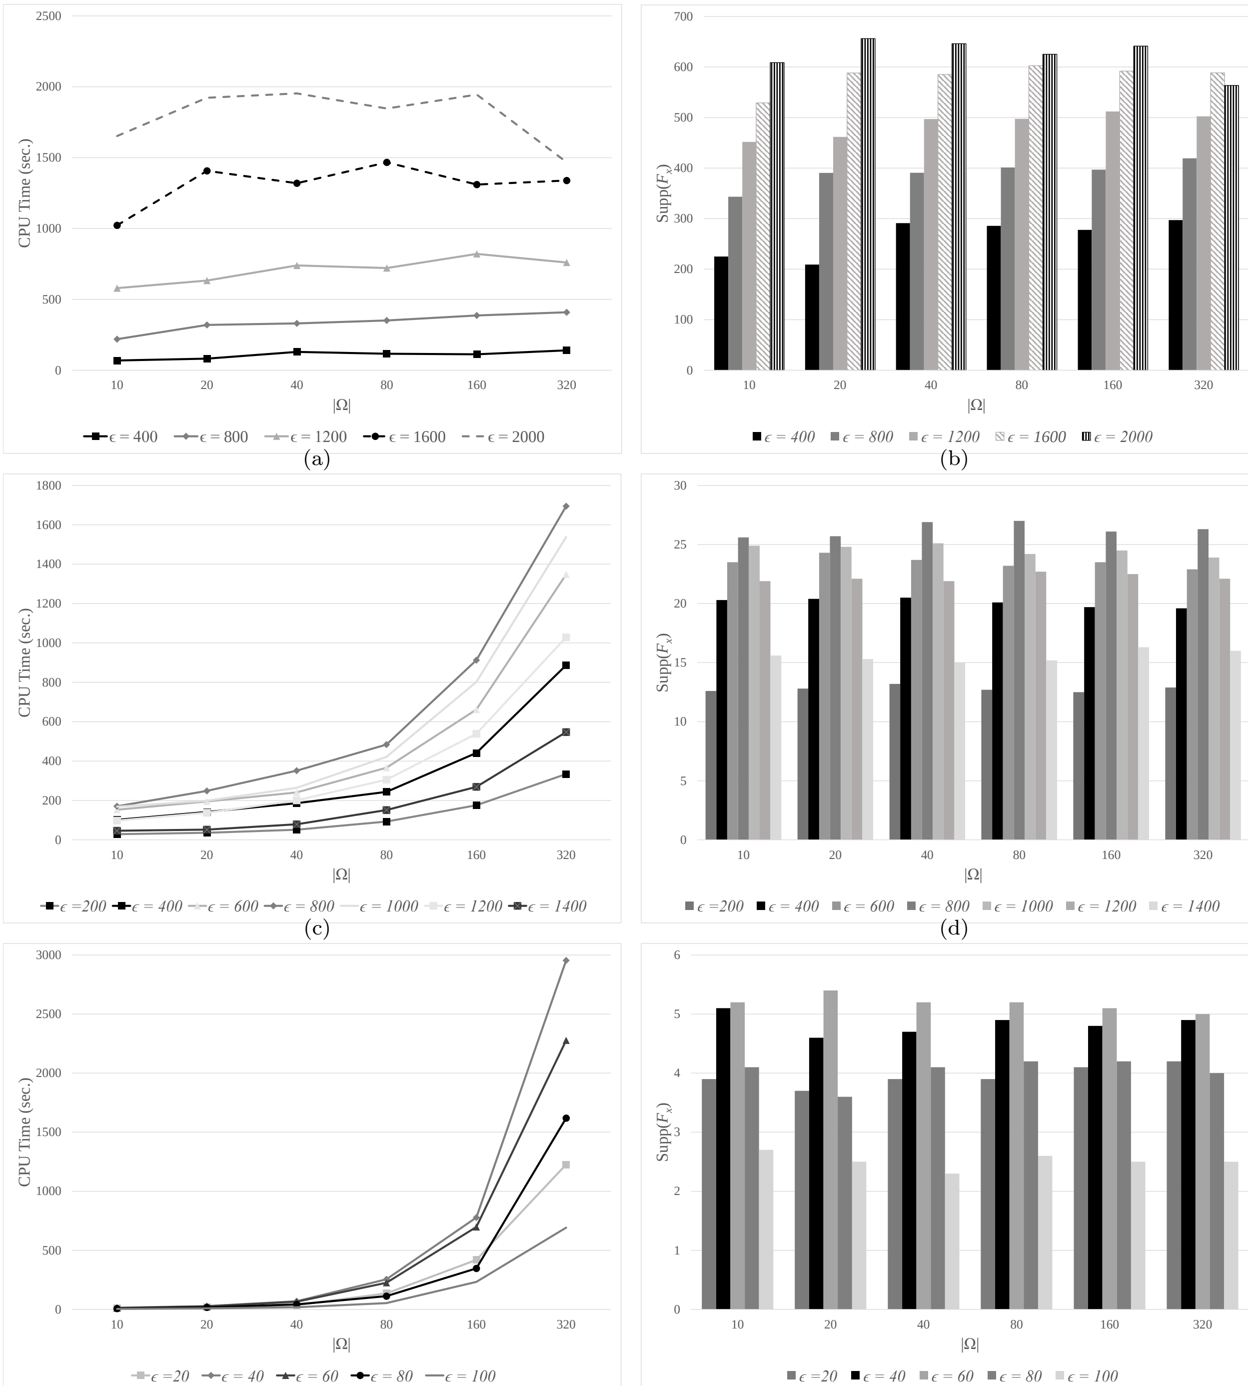

Figure 6: Effect of sample size on the computational time and the size of the solution's support, based on the average of 10 random tests instances of DRAP (a,b) DRUFLP (c,d) and DRCFLP with $r = 5$ (e,f), respectively.

2. Situations where the value of randomization is strictly positive and the individual risks of the different actions that compose the optimal randomized strategy are similar to each other (see the 2-node DRUFLP problem example in Section 3);
3. Situations where distributional ambiguity is high and where randomization might be able to eliminate (or significantly reduce) it (see the newsvendor problem example in Section A);
4. Situations in which the decision problem can be considered as belonging to a larger set of decision problems that will be solved either concurrently or sequentially.

We note that the list above is non-exhaustive and the relevance of each of its elements would still need to be carefully validated in an empirical study.