Risk-Averse Optimization in Multicriteria and Multistage Decision Making

Merve Merakli

A dissertation
submitted in partial fulfillment of the
requirements for the degree of

Doctor of Philosophy

University of Washington

2018

Reading Committee:

Simge Küçükyavuz, Chair

Archis Ghate

Nilay Noyan

Zelda Zabinsky

Program Authorized to Offer Degree:
Industrial & Systems Engineering

University of Washington

Abstract

Risk-Averse Optimization in Multicriteria and Multistage Decision Making

Merve Merakli

Chair of the Supervisory Committee:
Associate Professor Simge Küçükyavuz
Department of Industrial and Systems Engineering

Risk-averse stochastic programming provides means to incorporate a wide range of risk attitudes into decision making. Pioneered by the advances in financial optimization, several risk measures such as Value-at-Risk (VaR) and Conditional-Value-at-Risk (CVaR) are employed in risk-averse stochastic programming for a variety of application areas. In this work, we consider risk-averse modeling approaches for stochastic multicriteria and stochastic sequential decision-making problems. First, we propose a new multivariate definition for CVaR as a set of vectors. We analyze its properties and establish that the new definition remedies some potential drawbacks of the existing definitions for discrete random variables. Motivated by the computational challenges in the optimization of vector-valued multivariate definitions of CVaR, next, we study two-stage stochastic programming problems with multivariate risk constraints utilizing a scalarization scheme. We formulate this problem as a mixed-integer program (MIP) and devise two delayed cut generation algorithms. The effectiveness of the proposed modeling approach and solution methods are demonstrated on a pre-disaster relief network design problem. Finally, we study the Markov Decision Processes (MDPs) under cost and transition probability uncertainty with the objective of optimizing the VaR associated with the expected performance of an MDP model. Based on a sampling approach, we provide an MIP formulation and a branch-and-cut algorithm, and demonstrate our proposed methods on an inventory management problem for long-term humanitarian relief operations.

TABLE OF CONTENTS

Page

LIST OF FIGURES

LIST OF TABLES

ACKNOWLEDGMENTS

First and foremost, I wish to express my deepest gratitude to my advisor, Dr. Simge Küçükyavuz, for giving me this great opportunity to dive into the realm of risk-averse optimization, for many stimulating and inspiring conversations, and for her support and guidance throughout my journey as a Ph.D. student. She has been a great mentor to me with her everlasting passion and enthusiasm for research, immense knowledge, and patience to point me in the right direction whenever I needed. I genuinely feel lucky and privileged to have her as my advisor.

Second, I would like to thank my reading committee members Dr. Archis Ghate, Dr. Nilay Noyan, and Dr. Zelda Zabinsky for their invaluable time and comments that helped me improve this dissertation. I especially thank Dr. Nilay Noyan for our collaborative work presented in Chapter 3. I can only hope that I have learned from such an excellent researcher.

I am also very grateful to the faculty and staff in the Department of Industrial and Systems Engineering for all their help. In addition, I would like to acknowledge that my research was funded by National Science Foundation grant #1733001.

Knowing that "it is the journey that matters in the end," I want to thank all who led me to this point from the very beginning. I am forever indebted to my primary school teacher, Birol Koşak, for nurturing my love for mathematics. He completely changed the life of a little girl from a small village—I can never forget how much he has done for me. I am grateful to Dr. Emre Alper Yıldırım and Dr. Mustafa Çelebi Pınar, who made me fascinated by optimization and integer programming with their excellent teaching during my undergraduate years. No words could be enough to thank my dear M.S. studies advisor, Dr. Hande Yaman. I remember myself in our first meeting as a naive student with no experience in conducting original research. She introduced me into academic research, and provided with her support and encouragement ever since. I am also grateful to Dr. Guzin Bayraksan and Dr. Marc Posner for all their help during my time at the Ohio State University.

During the last three years, I have lived in two cities and had the chance to meet many wonderful people who made me enjoy and treasure every moment. I first thank my fellow traveler Hao-Hsiang Wu for his invaluable companionship and amazing fortunetelling skills. I owe a huge "thanks" to my friends in Seattle, Safiye Çelik, Havva Gürbüz, Burak Boyacıoğlu, Hasan Manzour, Zahra Ghatrani, Mutita Siriruchatanon, Larissa Prates, Aven Samareh, Michelle Song, Tolu Abe and Ting-Yu Ho, for their warm welcome and friendship. I am also extremely thankful to my friends in Columbus, Zeynep Aydoğdu, Seçkin Şahin, Anıl Türkkan, Könül Bayramoğlu Kavlak, Hüseyin Bahtiyar, İsmail Nas, Deniz Sargun, Xiao Liu, Ben Chaiken, Jangho Park, and Leo Cheng-Lung Chen, for not allowing me to feel homesick for a second. This list would not be complete without my dear friends from Bilkent, Gizem Özbaygın, Nihal Berktaş, Ece Demirci, Pınar Aksoy Özpınar, Nil Karacaoğlu, Burcu Tekin, Hüseyin Gürkan, and Sinan Bayraktar, who are always just a call away. Last but not the least, I have no words to describe how blessed I feel to have Esra Ünsal, Cihan Bilge Kayasandık, and Büşra Ünsal as my family by choice, my fellow dreamers, and my partners in crime. We have been thorough ups and downs, joys and sorrows, half of the Europe and U.S., tons of junk foods, and who-knows-how-many LOTR marathons for over a decade, and certainly there are many more to come.

Finally, I would like to thank my family, my mother Halime Meraklı, my father Hasan Meraklı and my dearest sister Feyza Meraklı, whose effect on me is profound and inexpressible, without whom this thesis would truly not be possible. I dedicate this thesis to them for their endless love and support.

Chapter 1

INTRODUCTION

Decision making (optimization) under uncertainty is an intensely explored research area with numerous applications including manufacturing, logistics, robotics, communication systems, and healthcare. In contrast with the deterministic setting, where all parameters of the system are assumed to be known, this framework acknowledges the fact that the decision makers may have incomplete information about the system dynamics. However, incorporating uncertainty into decision making is often a challenging task due to high computational requirements of dealing with probabilistic functions and potential problem-specific characteristics that need to be captured by the uncertainty model.

Classical stochastic optimization problems consider the expected behaviour of a system. For a decision vector $\mathbf{x}$ with a feasible region $\mathcal{X}$ and random input parameters $\boldsymbol{\xi}$, the problem can be stated as

$$\min_{\mathbf{x}\in\mathcal{X}} \quad \mathbb{E}_{\boldsymbol{\xi}}[f(\mathbf{x},\boldsymbol{\xi})], \tag{1.1}$$

assuming for simplicity that the uncertainty is only in the objective function. Problem (1.1) optimizes the expected value of the function $f(\mathbf{x},\boldsymbol{\xi})$. This approach is motivated by the applications in which the decision makers are interested in the long-term performance of the solution over a large sample, as in that case, the law of large numbers ensures reliability of the solution obtained by solving (1.1). On the other hand, this approach ignores the variability in the parameters and therefore is not suitable for certain cases, where the decision makers are interested in obtaining a solution that performs well even under the worst-case realization of the random parameters $\boldsymbol{\xi}$.

Alternatively, robust optimization problems aim to find solutions that perform optimally under the worst-case realization of parameters over an uncertainty set representing all possible values of the random parameters without making any assumptions on their prior information. Assuming an uncertainty set $\mathcal{U}$, the robust counterpart of (1.1) can be formulated as follows

$$\min_{\mathbf{x}\in\mathcal{X}} \max_{\boldsymbol{\xi}\in\mathcal{U}} \quad f(\mathbf{x}, \boldsymbol{\xi}). \tag{1.2}$$

A common criticism for the robust formulations is that the worst-case approach results in conservative solutions that, for most cases, focus solely on scenarios rarely encountered in practice at the expense of underperforming for more likely realizations of random parameters.

The previously mentioned expected value and robust approaches correspond to two extreme attitudes towards risk. The expected value approach as in problem (1.1) assumes that the decision maker is risk-neutral and ignores the variability in the objective function value. The robust model (1.2) on the other hand, mimics a risk-averse decision maker whose only concern is the worst-case performance. More recently, risk-averse stochastic optimization problems are introduced to provide the flexibility to capture a wider range of risk attitudes [98, 100]. In this framework, depending on the problem context, the risk can be incorporated into the decisions in the form of an objective function or constraints. For a risk functional $\rho(\cdot)$, the problem

$$\min_{\mathbf{x}\in\mathcal{X}} \quad \rho\left(f(\mathbf{x}, \boldsymbol{\xi})\right)$$

minimizes the risk associated with the value of the function $f(\mathbf{x}, \boldsymbol{\xi})$ for the selected solution $\mathbf{x}$. In more general terms, the objective function can be defined as a weighted sum of different risk measures. Alternatively, for a random benchmark $\mathbf{Y}$, the risk-constrained stochastic optimization problem

$$\min_{\mathbf{x}\in\mathcal{X}} \quad \mathbb{E}_{\boldsymbol{\xi}}[f(\mathbf{x},\boldsymbol{\xi})]$$
$$\text{s.t.} \quad \rho\left(f(\mathbf{x},\boldsymbol{\xi})\right) \leq \rho\left(\mathbf{Y}\right)$$

aims to achieve an acceptable level of risk limited by $\rho\left(\mathbf{Y}\right)$ while minimizing the expected value of the objective function. Note that the functions considered in the objective and the risk constraints do not need to be the same, and multiple risk constraints or mean-risk objective functions can be included.

In this thesis, we study risk-averse stochastic programming problems using both modeling choices. We are particularly interested in the risk measures Value-at-Risk (VaR) and Conditional Value-at-Risk (CVaR), which are commonly employed in the financial optimization literature. In the next section, we provide the definitions and general properties of the VaR and the CVaR.

1.1 The Value-at-Risk and The Conditional Value-at-Risk

Risk measures are mappings from a set of random variables to the real numbers that enable a direct comparison between random variables with respect to their risk. An ideal risk measure satisfies the properties of coherency and law-invariance. The seminal work of Artzner et al. [6] states that a mapping $\rho : \mathcal{V} \to \mathbb{R}$ defined on a set of random variables $\mathcal{V}$ is a coherent risk measure if the following properties hold for random variables $V, V_1, V_2 \in \mathcal{V}$.

- *Normalized:* $\rho(0) = 0$, implying that the risk associated with holding no sources of risk is zero.

- *Positively homogeneous:* $\rho(kV) = k\rho(V)$ for $k \geq 0$. That is, the risk is proportional to the size of its source.

- *Translation equivariant:* $\rho(V + k) = \rho(V) + k$ for $k \in \mathbb{R}$, which means that increasing the size of a risk source by a sure amount increases its risk by the same amount.

- *Subadditive:* $\rho(V_1 + V_2) \leq \rho(V_1) + \rho(V_2)$. This property is related to diversification. It implies that merging different sources of risk does not create additional risk.

- *Monotone:* $V_1 \leq V_2$ almost surely $\Rightarrow \rho(V_1) \leq \rho(V_2)$, meaning that if an option has better outcomes than another option for almost all possible realizations, then it is also less risky compared to the other option.

Note that in the definition above and hereafter, we use the convention that smaller values of random variables as well as risk measures are preferable (as in random costs). Kusuoka [62] further characterizes the class of coherent law-invariant risk measures. Basically, a risk measure is law-invariant if it depends only on the distribution of the associated random variable.

The initial studies in the portfolio optimization literature mainly focus on the VaR as a risk measure. At a confidence level $\alpha \in (0, 1]$, the VaR, or α-quantile in mathematical terms, corresponds to a threshold for the value of the random variable such that the probability of falling below the VaR value is at least α. For a random variable V at confidence level α, the $\mathrm{VaR}_\alpha(V)$ can be formulated as

$$\begin{aligned} \min \quad & \eta \\ \text{s.t.} \quad & \mathrm{P}(V \leq \eta) \geq \alpha \end{aligned}$$

in the form of a chance-constrained program. Despite its popularity in many applications of finance, VaR is not a coherent risk measure since it is not a convex function in general.

More recently, the CVaR introduced by Rockafellar and Uryasev [93] has been widely used in many applications of risk-averse stochastic optimization problems due to its desirable properties. The CVaR is a coherent [86] and law-invariant risk measure that can be used as a building block for many other risk measures. For a random variable V, the CVaR at confidence level $\alpha \in (0,1]$, denoted as $\mathrm{CVaR}_\alpha(V)$, corresponds to the expected value of random variable V under the condition that it is no better than the $\mathrm{VaR}_\alpha(V)$, and it is given by [93]

$$\mathrm{CVaR}_\alpha(V) = \min_{\eta \in \mathbb{R}} \left\{ \eta + \frac{1}{1-\alpha} \mathbb{E}(V-\eta)_+ \right\}, \tag{1.3}$$

where $(x)_+ = \max(0,x)$, using the convention that smaller values of random variables as well as risk measures are preferable. It is well known that there exists an optimal solution to problem (1.3) with $\eta = \mathrm{VaR}_\alpha(V)$. Assume that V follows a finite discrete distribution. In other words, there is a finite number of scenarios, n, where v_i represents the realization of V under scenario $i \in \{1,\dots,n\}$ with probability q_i. Then an equivalent representation of the CVaR is given by the following linear program

$$\mathrm{CVaR}_\alpha(V) = \min \left\{ \eta + \frac{1}{1-\alpha} \sum_{i \in [n]} q_i w_i \;:\; w_i \geq v_i - \eta,\ \forall\, i \in [n], \quad \mathbf{w} \in \mathbb{R}^n_+,\ \eta \in \mathbb{R} \right\}.$$

Furthermore, the definition of CVaR can be used as a way to build preference relations between univariate random variables. We say that a random variable X is preferable to random variable Y in terms of CVaR at confidence level α if

$$\mathrm{CVaR}_\alpha(X) \leq \mathrm{CVaR}_\alpha(Y). \tag{1.4}$$

In an optimization setting, this relation can be enforced as a constraint to ensure that our optimal decisions are less risky compared to a random benchmark outcome.

In a parallel line of work, the preference relations between random variables are established using the second-order stochastic dominance (SSD) relation, which is successfully applied in many areas due to its consistency with risk-averse preferences [31, 47, 77, 79]. SDD relation is proposed for the cases in which larger values of random variables are preferable. Its counterpart in the reverse convention where smaller values of random variables are preferable is referred to as increasing convex order (ICO). We say that a random variable X

is preferable to random variable Y with respect to ICO if

$$\mathbb{E}(X-\eta)_+ \leq \mathbb{E}(Y-\eta)_+ \quad \forall \eta \in \mathbb{R}.$$

Note that the ICO relation given above is equivalent to a continuum of CVaR relations such that [31]

$$\mathrm{CVaR}_\alpha(X) \leq \mathrm{CVaR}_\alpha(Y) \quad \forall \alpha \in (0,1].$$

Hence, ICO relation is stronger than the CVaR relation, and it may lead to conservative decisions in some cases.

The CVaR and ICO relations described above provide the means to compare scalar random variables. However, in many real-life problems, our decisions are subject to multiple sources of risk. For example, in humanitarian relief management context, one may be interested in comparing alternative relief network designs in a pre-disaster environment under high uncertainty based on their capability to ensure an efficient, fast and equal distribution of supply materials in the aftermath of a disaster. As in this example, multiple stochastic performance measures of interest may be conflicting and/or correlated. Hence, enforcing the preference relations for each performance measure separately likely produces sub-optimal decisions.

In the seminal work [32], Dentcheva and Ruszczyński extend the univariate stochastic order relations to the multivariate case by utilizing a family of scalarization vectors. In this idea, two random vectors are compared based on their scalarized values with respect to all scalarization vectors in a given set $\mathcal{C}$. For two d-dimensional random vectors $\mathbf{X}$ and $\mathbf{Y}$, we say that $\mathbf{X}$ is preferable to $\mathbf{Y}$ based on the ICO relation and the scalarization set $\mathcal{C} \subseteq \mathbb{R}^d_+$ if

$$\mathbb{E}(\mathbf{c}^\top \mathbf{X}-\eta)_+ \leq \mathbb{E}(\mathbf{c}^\top \mathbf{Y}-\eta)_+ \quad \forall \eta \in \mathbb{R},\ \mathbf{c} \in \mathcal{C}.$$

The linear scalarization scheme implies that each scalarization vector corresponds to the relative weights of performance criteria, which may be difficult to determine as a single value in practice. However, using a set of scalarization vectors instead of a single vector allows to incorporate possible ambiguities in the relative importance of criteria.

Based on this idea, Noyan and Rudolf [81] extend the univariate CVaR relation (1.4) for the multivariate random vectors as in the following definition.

Definition 1 ([81])**.** *Let* $\mathbf{X}$ *and* $\mathbf{Y}$ *be two d-dimensional random vectors,* $\mathcal{C} \subset \mathbb{R}^d_+$ *a set of scalarization vectors, and* $\alpha \in (0,1]$ *a specified confidence level. We say that* $\mathbf{X}$ *is preferable to* $\mathbf{Y}$ *in terms of* CVaR *at confidence level* α *with respect to* $\mathcal{C}$*, if*

$$\mathrm{CVaR}_\alpha(\mathbf{c}^\top \mathbf{X}) \leq \mathrm{CVaR}_\alpha(\mathbf{c}^\top \mathbf{Y}) \quad \forall \mathbf{c} \in \mathcal{C}.$$

We remind that the definitions in this section are adjusted according to the convention that smaller values of random variables and risk measures are preferable. Assuming a polyhedral set of scalarization vectors $\mathcal{C}$, the relation given in Definition 1 can be referred to as the multivariate polyhedral CVaR relation.

1.2 Risk-Averse Multistage Decision Making

Many real-world problems require determination of a sequence of decisions over a planning horizon in stochastic environments. In this setting, the decision taken at any period depends on the outcomes of the previous periods.

Two-stage stochastic programming problems are one of the most studied problems in the domain of stochastic multistage optimization problems (see [13, 107] and references therein). The main assumption of two-stage stochastic programming is that the decisions made in the first-stage under the presence of incomplete information can be improved by the recourse actions taken in the second-stage after the uncertainty is revealed. Considering a finite representation of uncertainty, i.e., a probability space $(\Omega, 2^\Omega, \Pi)$ with $\Omega = \{\omega_1, \ldots, \omega_m\}$ and $\Pi(\omega_s) = p_s$, where $S = \{1, \ldots, m\}$ corresponds to the index set of the elementary events (also referred to as scenarios), a traditional two-stage stochastic program can be stated as

$$\min \quad f(\mathbf{x}) + \mathbb{E}(Q(\mathbf{x}, \boldsymbol{\xi}(\omega))) \tag{1.5a}$$

$$\text{s.t.} \quad \mathbf{x} \in \mathcal{X}. \tag{1.5b}$$

Here, $f(\mathbf{x})$ refers to a deterministic objective function for the first-stage solution $\mathbf{x}$, which can take any value in set $\mathcal{X}$, and $Q(\mathbf{x}, \boldsymbol{\xi}(\omega_s))$ is the optimal second-stage objective function value under scenario $s \in S$ considering the realization of random input parameters $\boldsymbol{\xi}(\omega_s) = (\mathbf{q}(\omega_s), T(\omega_s), W(\omega_s))$. The second-stage problem for any scenario $s \in S$ is

$$Q(\mathbf{x}, \boldsymbol{\xi}(\omega_s)) = \min_{\mathbf{y} \in \mathcal{Y}(\mathbf{x}, \boldsymbol{\xi}(\omega_s))} \quad \mathbf{q}_s^\top \mathbf{y}, \tag{1.6}$$

where $\mathcal{Y}(\mathbf{x}, \boldsymbol{\xi}(\omega_s)) = \{\mathbf{y} \in \mathbb{R}_+^{n_2} : T_s\mathbf{x} + W_s\mathbf{y} \geq \mathbf{h}_s\}$ defines the set of feasible second-stage solutions $\mathbf{y}$ (recourse actions) under scenario $s \in S$. For solution algorithms on two-stage stochastic programs, we refer the reader to [10, 12, 112].

In classical two-stage stochastic programs, decision makers are assumed to be risk-neutral; the objective is to minimize the sum of the first-stage cost and the expected second-stage cost as in (1.5a). A natural way to incorporate risk-aversion into the two-stage framework is to replace the expectation term in (1.5a) with a risk measure of interest such as VaR and CVaR (see, e.g., [4, 40, 76, 80, 104, 107]). Alternatively, it is possible to enforce stochastic preference relations as described in the previous section [29, 30, 55] or chance constraints in (1.5b) (see, e.g., [120]). Note that most of the studies in this area considers univariate risk measures and stochastic preference relations. A detailed literature review on two-stage stochastic optimization with multivariate CVaR relation and multivariate SSD relation is provided in Chapter 3.

Risk-averse optimization for stochastic programming problems with more than two stages is not straightforward even considering the univariate risk. In addition to the computational and theoretical challenges of solving multistage optimization problems, one also needs to determine how to model the uncertainty and the risk. In the sequential decision making framework, it is possible to measure the risk in each period separately, or the risk with respect to the overall planning horizon. Additionally, the uncertainty can be assumed to be stationary (uncertain parameters follow the same distribution for the whole planning horizon) or time-dependent. Most of the literature focuses on extending the modeling approaches and solution algorithms for the risk-neutral problems to the risk-averse case using coherent measures of risk (see, for example, [21, 97, 105]).

The multistage stochastic optimization problem of interest in this thesis is the Markov decision processes (MDPs). Different from the general multistage stochastic programs, in MDPs, decisions at any stage of the planning horizon are not affected by the past, but only the current state of system, referred to as the Markov property. Based on this property, an MDP is described by its state space $\mathcal{H}$, action space $\mathcal{A}$, cost vector $\tilde{\mathbf{c}}$, and transition probability matrix $\tilde{\mathbf{P}}$. Note that here we consider an infinite horizon MDP with stationary cost and transition probabilities, hence the time indices are ignored. At any stage, the decision maker takes an action $a \in \mathcal{A}$ considering the current state of the system $i \in \mathcal{H}$ with the aim of optimizing the overall performance of the MDP model, e.g., the total expected cost, the total expected discounted cost or the expected average cost per stage. Based on the selected action, an immediate cost of taking action a in state i, denoted as $\tilde{c}_i(a)$, is incurred and the system transitions to another state $j \in \mathcal{H}$ with probability $\tilde{P}_{ij}(a)$. A policy corresponds to a decision strategy determined for each state (and stage in the case of unstationary parameters), and for a given policy $\boldsymbol{\pi}$, the behavior of an MDP model is characterized by the Bellman equations [9]

$$v_i = \tilde{c}_i(\pi_i) + \gamma \sum_{j \in \mathcal{H}} \tilde{P}_{ij}(\pi_i) v_j, \quad i \in \mathcal{H}, \tag{1.7}$$

where v_i represents the expected total discounted cost given that the system was initially in state $i \in \mathcal{H}$. We assume that the distribution of the initial state is given as $\mathbf{q} \in \mathbb{R}^{|\mathcal{H}|}$, and the future costs are discounted by $\gamma \in (0, 1)$. Bellman equations lead to a natural LP representation of the MDP model minimizing the expected total discounted cost as follows

$$\max \quad \sum_{i\in\mathcal{H}} q_i v_i \tag{1.8a}$$

$$s.t. \quad v_i \le \tilde{c}_i(a) + \gamma \sum_{j\in\mathcal{H}} \tilde{P}_{ij}(a) v_j, \quad i \in \mathcal{H},\ a \in \mathcal{A}. \tag{1.8b}$$

Constraints (1.8b) ensure that Bellman equations (1.7) are satisfied for the policy minimizing the expected total discounted cost. In addition, MDPs can be efficiently solved using the well-known value iteration and policy iteration algorithms [90].

Note that in classical MDPs, the parameters describing the costs ($\tilde{\mathbf{c}}$) and transition probabilities ($\tilde{\mathbf{P}}$) are assumed to be known. They are usually estimated based on the past data and/or expert opinions. However, in many applications, it may be challenging to obtain a single estimation of those parameters that correctly characterizes behavior of the MDP. Motivated by this, in this thesis, we consider MDPs with parameter uncertainty.

1.3 Research Scope and Outline

In this thesis, we consider risk-averse stochastic decision-making problems in two domains. The studies in the first two chapters aim to incorporate risk-aversion into multicriteria optimization problems. In multivariate setting, the challenges in measuring the risk associated with the decisions are coupled by the difficulty of comparing vector-valued, stochastic objective function values of different solutions. We address these structural and computational difficulties in risk-averse multicriteria decision making, and seek new definitions of multivariate risk measures and computationally efficient solution methods. The risk measures used in these chapters are built upon the univariate definition of CVaR.

In Chapter 2, we propose a new definition of multivariate conditional value-at-risk (MCVaR) as a set of vectors for arbitrary probability spaces. We explore the properties of the proposed vector-valued MCVaR (VMCVaR) and show the advantages of VMCVaR over the existing definitions particularly for discrete random variables. In addition, we discuss optimization of problems involving VMCVaR and propose multiobjective optimization formulations.

Motivated by the computational challenges associated with optimization problems involving VMCVaR, in Chapter 3, we consider two classes of multicriteria two-stage stochastic programs in finite probability spaces with multivariate risk constraints utilizing scalarization schemes. The first-stage problem features a multivariate stochastic benchmarking constraint based on a vector-valued random variable representing multiple and possibly conflicting stochastic performance measures associated with the second-stage decisions. In particular, the aim is to ensure that the associated random outcome vector of interest is preferable to a specified benchmark with respect to the multivariate polyhedral CVaR or a multivariate stochastic order relation. In this case, the classical decomposition methods cannot be used directly due to the complicating multivariate stochastic benchmarking constraints. We propose an exact unified decomposition framework for solving these two classes of optimization problems and show its finite convergence. We apply the proposed approach to a stochastic network design problem in a pre-disaster humanitarian logistics context and conduct a computational study concerning the threat of hurricanes in the Southeastern part of the United States. Our numerical results on these large-scale problems show that our proposed algorithm is computationally scalable.

In Chapter 4, we focus on risk-averse sequential decision making problems arising in dynamic stochastic environments. Mainly, we consider Markov Decision Processes (MDPs) under cost and transition probability uncertainty with the purpose of obtaining risk-averse solutions optimizing the expected performance of the MDP with prespecified probabilistic guarantees. To this end, we utilize the VaR to measure the risk associated with the expected total discounted cost of an infinite-horizon MDP model with finite state and action spaces. Considering a sampling approach, we propose a mixed-integer programming formulation that seeks an optimal stationary deterministic policy minimizing the value of VaR, and devise a branch-and-cut algorithm. The proposed modeling approach and solution algorithms are demonstrated on an inventory management problem for humanitarian relief operations during a slow onset disaster such as war, political insurgence, extreme poverty, famine, and drought.

The remainder of this thesis is organized as follows. In Chapter 2, we propose a new definition for multivariate CVaR and explore its properties. In Chapter 3, we consider multi-criteria two-stage stochastic optimization problems under multivariate risk constraints with an application in pre-disaster relief network design. Finally, in Chapter 4, we study MDPs under cost and transition probability uncertainty in the context of long-term humanitarian relief operations.

Chapter 2

VECTOR-VALUED MULTIVARIATE CONDITIONAL VALUE-AT-RISK

2.1 Introduction

This chapter is based on [75]. *Conditional value-at-risk* (CVaR) is a widely used tool, especially in financial optimization, for assessing the risk associated with a certain decision. The value of CVaR at a confidence level p is the expected outcome given that it is unfavorable compared to at least $100p\%$ of the possible realizations. The threshold value corresponds to the p-quantile, which is also known as the *value-at-risk* (VaR). VaR is a suitable risk measure for the cases where the aim is to avoid unfavorable outcomes with high probability. However, it does not measure the magnitude of the unfavorable outcomes. To address this, CVaR has been introduced to quantify the risk as the expected value of the undesired outcomes [93].

CVaR is first introduced as a risk measure for univariate random variables. Due to its desirable properties such as coherence and law invariance, CVaR is widely incorporated into optimization problems to minimize or limit the risk of the corresponding decisions (see e.g., [93], [94], [41]). However, in many real life problems, decision makers are interested in measuring the risk arising from multiple factors rather than a single outcome. One way of extending the univariate definitions of VaR and CVaR to the multivariate case is to use a scalarization vector for turning the random outcome vector into a scalar. However, the relative importance of criteria is usually ambiguous. The weights of criteria with respect to each other are not always quantifiable as a unique vector, as they may be subject to conflicting opinions of a group of decision makers. To address this, recent work considers a robust approach, where the set of possible scalarization vectors is assumed to be known and often assumed to be polyhedral or convex, and a worst-case scalarization vector in this set is used to scalarize the multivariate random vector [61, 68, 81, 82, 84]. These papers not only define the concept of polyhedral multivariate CVaR for finite discrete distributions, but also give optimization models and solution methods when the multivariate random outcome vector is a function of the decisions. For computability of these risk measures, the solution methods rely on sampling from the true distribution. In this context, it is natural to consider finite discrete distributions, described with a finite number of scenarios. In a related line of work, Prékopa and Lee [89] provide an alternative scalar-valued definition, where risk is measured along directions through a given reference point in the multidimensional space, and then the minimum or a combination is calculated with respect to the directions in a cone. In this paper, instead of a (scalarized) real-valued representation of the multivariate CVaR, we introduce a vector-valued risk measure without the need for the specification of an ambiguity set describing the possible weights or a reference point.

The first challenge in defining a multivariate CVaR is the determination of the p-quantile (multivariate VaR). While the p-quantile corresponds to a single real value in the univariate setting, p-quantile of a random vector may point to several vectors in the multivariate context. There are several studies addressing this issue. Prékopa [88] proposes a multivariate definition of VaR as a set of vectors, called p-Level Efficient Points (pLEPs), rather than a single value for arbitrary multivariate distributions. Cousin and Di Bernardino [22] propose an alternative approach for the continuous distribution case that computes a single vector as the expectation of the boundary surface of the multivariate quantile function. Torres et al. [111], on the other hand, introduce a direction parameter and compute the multivariate VaR vector in that direction. Di Bernardino et al. [36] and Adrian and Brunnermeier [2] also provide single, vector-valued definitions of multivariate VaR by conditioning on the information on certain criteria. As in the univariate case, these multivariate VaR definitions are only concerned with the probability of having favorable outcomes, and not the magnitude of the unfavorable outcomes.

The second challenge in defining a multivariate CVaR is related with the ambiguity in the characterization of *undesirable* outcomes in the multivariate case. Prékopa [88] provides a single and real-valued definition of multivariate CVaR that classifies a vector as undesirable if it is not better than or equivalent to any pLEP as defined in [87]. Cousin and Di Bernardino [23] propose a single, vector-valued CVaR definition such that realizations worse than at least one pLEP with respect to every criteria are assumed to be undesirable. The details of these studies providing a definition for multivariate CVaR will be discussed in Section 2.3.

This study is dedicated to defining a vector-valued multivariate CVaR for arbitrary multivariate distributions without characterizing a set that describes the relative importance of criteria in advance. Throughout the paper, smaller values of random variables as well as smaller values of risk measures are assumed to be preferable. We propose a new definition for multivariate CVaR and explore its properties in Section 2.2. In Section 2.3, we review the existing definitions, and demonstrate their shortcomings in the case of finite discrete distributions. We show that our new definition overcomes these problems, which leads to a unified context for multivariate CVaR.

2.2 Multivariate CVaR

We use the multivariate VaR definition in [88] as the p-quantile function for our multivariate CVaR definition. Prékopa [88] defines multivariate VaR (MVaR) using pLEPs introduced for the case of discrete distributions in [87] and extended for the general distribution functions in [88].

Definition 2 (Prékopa [88])**.** *Let $\mathbf{X} \in \mathbb{R}^d$ be a random vector and F its cumulative distribution function. The vector $\mathbf{s} \in \mathbb{R}^d$ is a pLEP (or a p-efficient point) of the distribution of $\mathbf{X}$, if $F(\mathbf{s}) \geq p$ and there is no $\mathbf{y} \leq \mathbf{s},\ \mathbf{y} \neq \mathbf{s}$ such that $F(\mathbf{y}) \geq p$.*

In the multivariate context, there may be infinitely many vectors (pLEPs) satisfying the above conditions. Hence $\text{MVaR}_p(\mathbf{X})$ corresponds to the set of such vectors. Dentcheva et al. [35] show that the set $\text{MVaR}_p(\mathbf{X})$ is nonempty and finite in the case that $\mathbf{X}$ has a finite discrete distribution with integer-valued realizations. The elements in $\text{MVaR}_p(\mathbf{X})$ correspond to the Pareto efficient points of the multi-objective chance-constrained optimization problem:

$$\min \quad \mathbf{v} \tag{2.1a}$$

$$\text{s.t.} \quad \mathrm{P}(\mathbf{X} \leq \mathbf{v}) \geq p, \tag{2.1b}$$

where $\mathbf{X} \leq \mathbf{v}$ denotes a component-wise relation, i.e., $X_i \leq v_i$ for all $i \in [d]$. For an overview of the algorithms to generate the MVaR for the discrete distribution case, we refer the reader to Prékopa [88] and the references therein. As mentioned earlier, $\mathrm{MVaR}_p(\mathbf{X})$ concerns the p-quantile of the multivariate risk only. To measure the magnitude of the risk associated with the $(1-p)100\%$ worst outcomes of multiple risk factors, a multivariate analogue of CVaR is needed.

Next, we propose a new definition for multivariate CVaR as a set of vectors, each associated with a pLEP in MVaR. In our definition, outcomes exceeding a pLEP in at least one criterion are considered as undesirable with respect to that pLEP.

Definition 3 (Vector-Valued Multivariate Conditional Value-at-Risk)**.** *Assume that* $\mathbf{X} \in \mathbb{R}^d$ *is a random vector with a set of pLEPs* $\mathrm{MVaR}_p(\mathbf{X})$ *at confidence level p. The vector-valued multivariate CVaR of* $\mathbf{X}$ *at level p, denoted as* $\mathrm{VMCVaR}_p(\mathbf{X})$ *is defined as,*

$$\mathrm{VMCVaR}_p(\mathbf{X}) = \mathrm{Min}\{\mathrm{MCVaR}_p(\mathbf{X}, \boldsymbol{\eta}) : \ \boldsymbol{\eta} \in \mathrm{MVaR}_p(\mathbf{X})\}, \tag{2.2}$$

where

$$\mathrm{MCVaR}_p(\mathbf{X}, \boldsymbol{\eta}) = \boldsymbol{\eta} + \frac{1}{1-p}\mathbb{E}[(\mathbf{X} - \boldsymbol{\eta})_+]$$

and $[(\mathbf{X} - \boldsymbol{\eta})_+]_i = \max(0, X_i - \eta_i)$ *for all* $i \in [d] := \{1, \ldots, d\}$.

The Min operator in Definition 3 ensures that an element of $\mathrm{VMCVaR}_p(\mathbf{X})$ is non-dominated by another element in this set, i.e., the elements of $\mathrm{VMCVaR}_p(\mathbf{X})$ form an efficient frontier. (In the context that smaller values are preferable, a vector $\mathbf{x} \in \mathbb{R}^d$ is *non-dominated* by a vector $\mathbf{y} \in \mathbb{R}^d$ if $y_i > x_i$ for some $i \in [d]$. In addition, $\mathbf{x}$ dominates $\mathbf{y}$ if $y_i \geq x_i$ for all $i \in [d]$.) Note that if X is a univariate random variable, then there is a unique element $\mathrm{MCVaR}_p(X, \eta)$ for $\eta = \mathrm{VaR}_p(X)$, hence the Min operator is not needed. To see why this operator is needed in the multivariate case, we provide an example next.

Example 1. *Let* $\mathbf{X}$ *be a bivariate random vector with support* $\mathbf{X} \in \{(4, 1.5), (1, 3), (2, 5), (2, 3), (3, 1)\}$ *such that each realization is equally likely. At confidence level* $p = 0.6$*,* $\mathrm{MVaR}_p(\mathbf{X}) = \{(3, 3), (2, 5)\}$*. Using this information, we compute* $\mathrm{MCVaR}_p(\mathbf{X}, (3, 3)) = (3.5, 4)$*, and* $\mathrm{MCVaR}_p(\mathbf{X}, (2, 5)) = (3.5, 5)$*. Clearly,* $\mathrm{MCVaR}_p(\mathbf{X}, (3, 3)) \leq \mathrm{MCVaR}_p(\mathbf{X}, (2, 5))$*, hence we let* $\mathrm{VMCVaR}(\mathbf{X}) = \{(3.5, 4)\}$*.*

Recall that for a univariate random variable V, the CVaR at confidence level p is given as [93]

$$\mathrm{CVaR}_p(V) = \min_{\eta \in \mathbb{R}} \{\eta + \frac{1}{1-p} \mathbb{E}(V - \eta)_+\}. \tag{2.3}$$

Suppose that V follows a finite discrete distribution. In other words, there is a finite number of scenarios, n, where v_i represents the realization of V under scenario $i \in [n]$ with probability q_i. Then an equivalent representation of univariate CVaR is given by a linear program (LP)

$$\mathrm{CVaR}_p(V) = \min\{\eta + \frac{1}{1-p} \sum_{i \in [n]} q_i w_i \ : \ w_i \geq v_i - \eta, \ \forall \ i \in [n], \quad \mathbf{w} \in \mathbb{R}^n_+, \ \eta \in \mathbb{R}\}. \tag{2.4}$$

It is well known that there exists an optimal solution to problems (2.3) and (2.4) with $\eta = \mathrm{VaR}_p(V)$.

The following remark states that for univariate random variables, VMCVaR coincides with the classical definition of CVaR.

Remark 1. *Let* $X \in \mathbb{R}$ *be a univariate random variable such that* $\mathrm{VaR}_p(X) = \eta$ *at confidence level* p*. Then,*

$$\mathrm{MVaR}_p(X) = \{\eta\}, \tag{2.5}$$

and $\mathrm{VMCVaR}_p(X)$ *is a singleton with element*

$$\mathrm{MCVaR}_p(X, \eta) = \mathrm{CVaR}_p(X). \tag{2.6}$$

The statements (2.5) and (2.6) simply follow from Definition (2) and Definition (3), respectively.

Next we consider an alternative representation of VMCVaR as a multiobjective optimization problem.

Proposition 1. *For a random vector* $\mathbf{X} \in \mathbb{R}^d$ *at confidence level* p*,* $\mathrm{VMCVaR}_p(\mathbf{X})$ *is given by the Pareto-efficient optimal solutions to the multiobjective optimization problem*

$$\min_{\boldsymbol{\eta} \in \mathbb{R}^d} \quad \boldsymbol{\eta} + \frac{1}{1-p}\mathbb{E}[(\mathbf{X} - \boldsymbol{\eta})_+] \tag{2.7a}$$

$$s.t. \quad \mathrm{P}(\mathbf{X} \leq \boldsymbol{\eta}) \geq p. \tag{2.7b}$$

Proof. First, note that for given $\boldsymbol{\eta} \in \mathbb{R}^n$, the objective function in (2.7a) corresponds to $\mathrm{MCVaR}_p(\mathbf{X}, \boldsymbol{\eta})$ as stated in Definition 3. We need to show that any optimal objective vector of the above problem is equal to the value of $\mathrm{MCVaR}_p(\mathbf{X}, \boldsymbol{\eta})$ for some $\boldsymbol{\eta} \in \mathrm{MVaR}_p(\mathbf{X})$. Considering the definition of $\mathrm{MVaR}_p(\mathbf{X})$ in (2.1) along with constraint (2.7b), for any feasible solution $\boldsymbol{\eta}^*$ of problem (2.7), there exists $\boldsymbol{\eta} \in \mathrm{MVaR}_p(\mathbf{X})$ such that $\boldsymbol{\eta}^* \geq \boldsymbol{\eta}$. Note that $\boldsymbol{\eta}$ is feasible for problem (2.7) by definition. Let $\eta_i^* = \eta_i + \epsilon_i$ for $i \in [d]$, where $\epsilon_i \geq 0$. (Note that if $\boldsymbol{\eta}^* \neq \boldsymbol{\eta}$, then $\epsilon_j > 0$ for some $j \in [d]$.) Let $\boldsymbol{\vartheta}(\boldsymbol{\eta})$ and $\boldsymbol{\vartheta}(\boldsymbol{\eta}^*)$ be the optimal objective vectors associated with solutions $\boldsymbol{\eta}$ and $\boldsymbol{\eta}^*$, respectively. Then,

$$\boldsymbol{\vartheta}_i^*(\boldsymbol{\eta}^*) = \eta_i^* + \frac{1}{1-p}\mathbb{E}[(X_i - \eta_i^*)_+] \tag{2.8}$$

$$= \eta_i^* + \frac{\mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta}^*)}{1-p}\mathbb{E}[(X_i - \eta_i^*)_+ | \mathbf{X} \nleq \boldsymbol{\eta}^*] \tag{2.9}$$

$$= (\eta_i + \epsilon_i) + \frac{\mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta}^*)}{1-p}\mathbb{E}[(X_i - (\eta_i + \epsilon_i))_+ | \mathbf{X} \nleq \boldsymbol{\eta}^*] \tag{2.10}$$

$$\geq \eta_i + \epsilon_i + \frac{\mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta}^*)}{1-p}\mathbb{E}[(X_i - \eta_i)_+|\mathbf{X} \nleq \boldsymbol{\eta}^*] - \frac{\mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta}^*)}{1-p}\epsilon_i \tag{2.11}$$

$$\geq \eta_i + \frac{1}{1-p}\mathbb{E}[(X_i - \eta_i)_+] - \frac{\mathrm{P}(\mathbf{X} \leq \boldsymbol{\eta}^*, \mathbf{X} \nleq \boldsymbol{\eta})}{1-p}\epsilon_i + \left(1 - \frac{\mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta}^*)}{1-p}\right)\epsilon_i \tag{2.12}$$

$$= \eta_i + \frac{1}{1-p}\mathbb{E}[(X_i - \eta_i)_+] + \left(1 - \frac{\mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta})}{1-p}\right)\epsilon_i \tag{2.13}$$

$$\geq \boldsymbol{\vartheta}_i(\boldsymbol{\eta}) = (\mathrm{MCVaR}_p(\mathbf{X}, \boldsymbol{\eta}))_i. \tag{2.14}$$

Equations (2.8) and (2.9) directly follow from (2.7a) and the definition of conditional probability, respectively. Equation (2.10) is obtained after substituting $\eta_i^* = \eta_i + \epsilon_i$, and inequality (2.11) is based on the fact that $(\mathbf{a} - \mathbf{b})_+ \geq (\mathbf{a})_+ - (\mathbf{b})_+$ (here $\mathbf{a} = X_i - \eta_i$ and $\mathbf{b} = \epsilon_i$). Equation (2.12) is a result of $\mathbb{E}[(X_i - \eta_i)_+] = \mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta}^*)\mathbb{E}[(X_i - \eta_i)_+|\mathbf{X} \nleq \boldsymbol{\eta}^*] + \mathrm{P}(\mathbf{X} \leq \boldsymbol{\eta}^*, \mathbf{X} \nleq \boldsymbol{\eta})\mathbb{E}[(X_i - \eta_i)_+|\mathbf{X} \leq \boldsymbol{\eta}^*, \mathbf{X} \nleq \boldsymbol{\eta}]$ and the assumption that $\boldsymbol{\eta}_i^* - \boldsymbol{\eta}_i = \epsilon_i$. Equation (2.13) follows from $\boldsymbol{\eta}^* \geq \boldsymbol{\eta}$ and the law of total probability. Finally, inequality (2.14) is due to the representation of $\mathrm{MCVaR}_p(\mathbf{X}, \boldsymbol{\eta})$ in Definition (3) and the fact that $1 - \frac{\mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta})}{1-p} \geq 0$ by constraint (2.7b).

The arguments above show that for every feasible solution $\boldsymbol{\eta}^*$ of problem (2.7) with objective function value $\boldsymbol{\vartheta}(\boldsymbol{\eta}^*)$, there exists another feasible solution $\boldsymbol{\eta} \in \mathrm{MVaR}_p(\mathbf{X})$ such that $\boldsymbol{\vartheta}(\boldsymbol{\eta}^*) \geq \boldsymbol{\vartheta}(\boldsymbol{\eta}) = \mathrm{MCVaR}_p(\mathbf{X}, \boldsymbol{\eta})$. Next, we need to show that any element in $\mathrm{VMCVaR}_p(\mathbf{X})$ is a Pareto optimal solution to problem (2.7). By definition, all elements of $\mathrm{VMCVaR}_p(\mathbf{X})$ are feasible for (2.7). Furthermore, the Min operator in Definition 3 ensures that the elements of $\mathrm{VMCVaR}_p(\mathbf{X})$ correspond to the Pareto efficient points of problem (2.7), completing the proof. □

For random vectors with discrete distribution functions, the formulation of VMCVaR presented in Proposition (1) can be alternatively stated as a multiobjective mixed-integer programming problem. Let $\mathbf{x}^s = (x_1^s, \ldots, x_d^s)$ be the realization of $\mathbf{X}$ under scenario s with probability q_s for $s \in [n]$. Then, analogous to the formulation of CVaR for univariate random variables with discrete distributions in (2.4), we obtain the following formulation for the multivariate case,

$$\min \quad \boldsymbol{\eta} + \frac{1}{1-p}\sum_{s\in[n]} q_s \mathbf{w}^s \tag{2.15a}$$

$$\text{s.t.} \quad w_i^s \geq \mathbf{x}_i^s - \eta_i, \quad i \in [d],\ s \in [n], \tag{2.15b}$$

$$\mathbf{w}^s \geq \mathbf{0}, \quad s \in [n], \tag{2.15c}$$

$$\sum_{s\in[n]} q^s \beta_s \leq 1-p, \tag{2.15d}$$

$$\mathbf{x}_i^s \leq \eta_i + M_{is}\beta_s, \quad i \in [d],\ s \in [n], \tag{2.15e}$$

$$\beta_s \in \{0,1\}, \quad s \in [n], \tag{2.15f}$$

whose Pareto-efficient optimal solutions yield the set $\text{VMCVaR}_p(\mathbf{X})$. Here, $\beta_s = 0$ if under scenario s we have $\mathbf{x}^s \leq \boldsymbol{\eta}$ and $M_{is}, i \in [d],\ s \in [n]$ is a large enough number so that constraint (2.15e) is trivially satisfied if $\beta_s = 1$. By Proposition 1, for any optimal objective function value, this problem always has an optimal solution $\boldsymbol{\eta}$ in the set $\text{MVaR}_p(\mathbf{X})$. The complexity of the multivariate CVaR definition is evident from comparing the formulations for the univariate and multivariate case; the former is a convex optimization problem (2.3) or an LP (2.4) for the cases of general and discrete distributions, respectively, whereas the latter is a chance-constrained program (2.7) or a multiobjective mixed-integer program (2.15). However, the main difficulty lies at identifying the MVaR values. Once these are identified, say using the methods reviewed in [88], calculating $\text{MCVaR}_p(\mathbf{X}, \eta)$ for each $\eta \in \text{MVaR}_p(\mathbf{X})$ reduces to a convex optimization problem in general and an LP, which is polynomial in the number of scenarios, n, for the discrete case.

In the next proposition, we provide a more intuitive definition of $\text{MCVaR}_p(\mathbf{X}, \boldsymbol{\eta})$ for the case that $\text{P}(\mathbf{X} \leq \boldsymbol{\eta}) = p$.

Proposition 2. *For a random variable* $\mathbf{X} \in \mathbb{R}^d$ *and* $\boldsymbol{\eta} \in \text{MVaR}_p(\mathbf{X})$ *such that* $\text{P}(\mathbf{X} \leq \boldsymbol{\eta}) =$

p, the following equality holds,

$$\mathrm{MCVaR}_p(\mathbf{X},\boldsymbol{\eta}) = \boldsymbol{\eta} + \frac{1}{1-p}\mathbb{E}[(\mathbf{X}-\boldsymbol{\eta})_+] = \mathbb{E}[\max(\mathbf{X},\boldsymbol{\eta})|\mathbf{X} \nleq \boldsymbol{\eta}] \tag{2.16}$$

where $\mathbf{X} \nleq \boldsymbol{\eta}$ *denotes a component-wise relation such that* $X_i > \eta_i$ *for at least one* $i \in [d]$, *and* $\max(\mathbf{X},\boldsymbol{\eta})$ *represents a vector whose* i^{th} *component is equal to* $\max(X_i, \eta_i)$.

Proof. For a criterion $i \in [d]$, the i^{th} component of $\mathrm{MCVaR}_p(\mathbf{X},\boldsymbol{\eta})$ is equal to,

$$\begin{aligned}
(\mathrm{MCVaR}_p(\mathbf{X},\boldsymbol{\eta}))_i &= \eta_i + \frac{1}{1-p}\mathbb{E}[(X_i-\eta_i)_+] \\
&= \eta_i + \frac{1}{\mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta})}\mathbb{E}[(X_i-\eta_i)_+] \\
&= \eta_i + \mathbb{E}[(X_i-\eta_i)_+|\mathbf{X} \nleq \boldsymbol{\eta}] \\
&= \mathbb{E}[\eta_i|\mathbf{X} \nleq \boldsymbol{\eta}] + \mathbb{E}[(\max(X_i,\eta_i)-\eta_i)|\mathbf{X} \nleq \boldsymbol{\eta}] \\
&= \mathbb{E}[(\eta_i + \max(X_i,\eta_i)-\eta_i)|\mathbf{X} \nleq \boldsymbol{\eta}] \\
&= \mathbb{E}[\max(X_i,\eta_i)|\mathbf{X} \nleq \boldsymbol{\eta}].
\end{aligned}$$

□

This is also analogous to the univariate definition of CVaR since for a univariate random variable V such that $\mathrm{P}(V \leq \mathrm{VaR}_p(V)) = p$,

$$\begin{aligned}
\mathrm{CVaR}_p(V) &= \mathbb{E}(V|V > \mathrm{VaR}_p(V)) = \mathbb{E}(\max(V,\mathrm{VaR}_p(V))|V > \mathrm{VaR}_p(V)) \\
&= \mathbb{E}(\max(V,\mathrm{VaR}_p(V))|V \nleq \mathrm{VaR}_p(V)).
\end{aligned} \tag{2.17}$$

The proposed multivariate risk measure should satisfy some desired properties to be a useful tool for evaluating the risk of random vectors. Artzner et al. [6] axiomatize the desired properties of univariate risk measures under the *coherence* concept. A coherent univariate risk measure is normalized, positively homogeneous, translation equivariant, monotone and subadditive. The risk measure of interest in this study, $\mathrm{VMCVaR}_p(\mathbf{X})$, is a mapping from a d-dimensional random vector to a set of d-dimensional vectors. Hence, we aim to prove an analogous form of these properties for set-valued multivariate risk measures. Recall that for two random vectors, $\mathbf{X}$ and $\mathbf{Y}$, first-order stochastic dominance implies $\mathbf{X} \preceq_{(1)} \mathbf{Y}$ if and only if $\mathbb{E}[f(\mathbf{X})] \leq \mathbb{E}[f(\mathbf{Y})]$ for any bounded monotonic function $f : \mathbb{R}^d \to \mathbb{R}$.

Proposition 3. *For a d-dimensional random vector $\mathbf{X}$, $\mathrm{VMCVaR}_p(\mathbf{X})$ satisfies the following properties,*

(i) Normalized: $\mathrm{VMCVaR}_p(\mathbf{0}) = \{\mathbf{0}\}$.

(ii) Positively homogeneous: $\forall \boldsymbol{\vartheta} \in \mathrm{VMCVaR}_p(\mathbf{X}), k\boldsymbol{\vartheta} \in \mathrm{VMCVaR}_p(k\mathbf{X})$ *for* $k \in \mathbb{R}_+$.

(iii) Translation equivariant: $\forall \boldsymbol{\vartheta} \in \mathrm{VMCVaR}_p(\mathbf{X}), \boldsymbol{\vartheta} + \mathbf{k} \in \mathrm{VMCVaR}_p(\mathbf{X} + \mathbf{k})$ *for* $\mathbf{k} \in \mathbb{R}^d$.

(iv) Monotone with respect to first-order stochastic dominance:
$\mathbf{X} \preceq_{(1)} \mathbf{Y} \Rightarrow \forall \boldsymbol{\vartheta}_1 \in \mathrm{VMCVaR}_p(\mathbf{X}), \exists \boldsymbol{\vartheta}_2 \in \mathrm{VMCVaR}_p(\mathbf{Y})$ *s.t.* $\boldsymbol{\vartheta}_2 \geq \boldsymbol{\vartheta}_1$, *and*
$\mathbf{X} \preceq_{(1)} \mathbf{Y} \Rightarrow \forall \boldsymbol{\vartheta}_2 \in \mathrm{VMCVaR}_p(\mathbf{Y}), \exists \boldsymbol{\vartheta}_1 \in \mathrm{VMCVaR}_p(\mathbf{X})$ *s.t.* $\boldsymbol{\vartheta}_2 \geq \boldsymbol{\vartheta}_1$.

Proof. It can be easily seen that $\mathrm{VMCVaR}_p(\mathbf{X})$ is normalized. We will prove the remaining properties using the fact that for all $\boldsymbol{\vartheta} \in \mathrm{VMCVaR}_p(\mathbf{X})$, $\boldsymbol{\vartheta} = \mathrm{MCVaR}_p(\mathbf{X}, \boldsymbol{\eta})$ for some $\boldsymbol{\eta} \in \mathrm{MVaR}_p(\mathbf{X})$.

(ii) From [64], we know if $\boldsymbol{\eta} \in \mathrm{MVaR}_p(\mathbf{X})$, then $k\boldsymbol{\eta} \in \mathrm{MVaR}_p(k\mathbf{X})$ for $k \in \mathbb{R}_+$. Hence,

$$\begin{aligned} k\,\mathrm{MCVaR}_p(\mathbf{X}, \boldsymbol{\eta}) &= k\left(\boldsymbol{\eta} + \frac{1}{1-p}\mathbb{E}[(\mathbf{X} - \boldsymbol{\eta})_+]\right) \\ &= k\boldsymbol{\eta} + \frac{1}{1-p}\mathbb{E}[(k\mathbf{X} - k\boldsymbol{\eta})_+] = \mathrm{MCVaR}_p(k\mathbf{X}, k\boldsymbol{\eta}). \end{aligned}$$

(iii) Lee and Prékopa [64] show that $\boldsymbol{\eta} + \mathbf{k} \in \mathrm{MVaR}_p(\mathbf{X} + \mathbf{k})$ for $\mathbf{k} \in \mathbb{R}^d$ when $\boldsymbol{\eta} \in \mathrm{MVaR}_p(\mathbf{X})$. Using this,

$$\mathrm{MCVaR}_p(\mathbf{X}, \boldsymbol{\eta}) + \mathbf{k} = \boldsymbol{\eta} + \frac{1}{1-p}\mathbb{E}[(\mathbf{X} - \boldsymbol{\eta})_+] + \mathbf{k}$$

$$= \boldsymbol{\eta} + \mathbf{k} + \frac{1}{1-p}\mathbb{E}[(\mathbf{X} + \mathbf{k} - \boldsymbol{\eta} - \mathbf{k})_+] = \mathrm{MCVaR}_p(\mathbf{X} + \mathbf{k}, \boldsymbol{\eta} + \mathbf{k}).$$

(iv) We first provide the proof of the first statement. By the monotonicity property of MVaR given in [64], $\mathbf{X} \preceq_{(1)} \mathbf{Y}$ implies that for all $\boldsymbol{\eta}_1 \in \mathrm{MVaR}_p(\mathbf{X})$ there exists $\boldsymbol{\eta}_2 \in \mathrm{MVaR}_p(\mathbf{Y})$ such that $\boldsymbol{\eta}_2 \geq \boldsymbol{\eta}_1$ and there is no $\boldsymbol{\eta}_2 \in \mathrm{MVaR}_p(\mathbf{Y})$ such that $\boldsymbol{\eta}_2 < \boldsymbol{\eta}_1$ for some $\boldsymbol{\eta}_1 \in \mathrm{MVaR}_p(\mathbf{X})$. Then,

$$\begin{aligned}
\boldsymbol{\vartheta}_1 &= \mathrm{MCVaR}_p(\mathbf{X}, \boldsymbol{\eta}_1) = \boldsymbol{\eta}_1 + \frac{1}{1-p}\mathbb{E}[(\mathbf{X} - \boldsymbol{\eta}_1)_+] && (2.18)\\
&= \boldsymbol{\eta}_1 + \frac{1}{1-p}\mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta}_1)\mathbb{E}[(\mathbf{X} - \boldsymbol{\eta}_1)_+ | \mathbf{X} \nleq \boldsymbol{\eta}_1] && (2.19)\\
&= \boldsymbol{\eta}_1 + \frac{1}{1-p}\mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta}_1, \mathbf{X} \leq \boldsymbol{\eta}_2)\mathbb{E}[(\mathbf{X} - \boldsymbol{\eta}_1)_+ | \mathbf{X} \nleq \boldsymbol{\eta}_1, \mathbf{X} \leq \boldsymbol{\eta}_2] && (2.20)\\
&\quad + \frac{1}{1-p}\mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta}_1, \mathbf{X} \nleq \boldsymbol{\eta}_2)\mathbb{E}[(\mathbf{X} - \boldsymbol{\eta}_1)_+ | \mathbf{X} \nleq \boldsymbol{\eta}_1, \mathbf{X} \nleq \boldsymbol{\eta}_2]\\
&\leq \boldsymbol{\eta}_1 + \frac{1}{1-p}\mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta}_1, \mathbf{X} \leq \boldsymbol{\eta}_2)\mathbb{E}[(\boldsymbol{\eta}_2 - \boldsymbol{\eta}_1)_+ | \mathbf{X} \nleq \boldsymbol{\eta}_1, \mathbf{X} \leq \boldsymbol{\eta}_2] && (2.21)\\
&\quad + \frac{1}{1-p}\mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta}_2)\mathbb{E}[(\boldsymbol{\eta}_2 - \boldsymbol{\eta}_1) + (\mathbf{X} - \boldsymbol{\eta}_2)_+ | \mathbf{X} \nleq \boldsymbol{\eta}_2]\\
&= \boldsymbol{\eta}_1 + \frac{1}{1-p}\mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta}_1, \mathbf{X} \leq \boldsymbol{\eta}_2)\mathbb{E}[(\boldsymbol{\eta}_2 - \boldsymbol{\eta}_1) | \mathbf{X} \nleq \boldsymbol{\eta}_1, \mathbf{X} \leq \boldsymbol{\eta}_2] && (2.22)\\
&\quad + \frac{1}{1-p}\mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta}_2)\mathbb{E}[(\boldsymbol{\eta}_2 - \boldsymbol{\eta}_1) | \mathbf{X} \nleq \boldsymbol{\eta}_2]\\
&\quad + \frac{1}{1-p}\mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta}_2)\mathbb{E}[(\mathbf{X} - \boldsymbol{\eta}_2)_+ | \mathbf{X} \nleq \boldsymbol{\eta}_2]\\
&= \boldsymbol{\eta}_1 + \frac{1}{1-p}\mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta}_1)(\boldsymbol{\eta}_2 - \boldsymbol{\eta}_1) && (2.23)\\
&\quad + \frac{1}{1-p}\mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta}_2)\mathbb{E}[(\mathbf{X} - \boldsymbol{\eta}_2)_+ | \mathbf{X} \nleq \boldsymbol{\eta}_2]\\
&\leq \boldsymbol{\eta}_2 + \frac{1}{1-p}\mathrm{P}(\mathbf{Y} \nleq \boldsymbol{\eta}_2)\mathbb{E}[(\mathbf{Y} - \boldsymbol{\eta}_2)_+ | \mathbf{Y} \nleq \boldsymbol{\eta}_2] && (2.24)\\
&= \boldsymbol{\eta}_2 + \frac{1}{1-p}\mathbb{E}[(\mathbf{Y} - \boldsymbol{\eta}_2)_+] = \mathrm{MCVaR}_p(\mathbf{Y}, \boldsymbol{\eta}_2) = \boldsymbol{\vartheta}_2, && (2.25)
\end{aligned}$$

where equality (2.18) follows from Definition 3; equalities (2.19) and (2.20) are due to the law of total expectation; inequality (2.21) follows from the condition $\mathbf{X} \leq \boldsymbol{\eta}_2$ in the first expectation term and the fact that $(\mathbf{X}-\boldsymbol{\eta}_1)_+ \leq (\boldsymbol{\eta}_2-\boldsymbol{\eta}_1)+(\mathbf{X}-\boldsymbol{\eta}_2)_+$ for all $\mathbf{X} \nleq \boldsymbol{\eta}_2$ as $\boldsymbol{\eta}_2 \geq \boldsymbol{\eta}_1$; equality (2.22) is obtained using the linearity of expectation; equality (2.23) follows from the law of total expectation, since the events $(\mathbf{X} \nleq \boldsymbol{\eta}_1, \mathbf{X} \leq \boldsymbol{\eta}_2)$ and $(\mathbf{X} \nleq \boldsymbol{\eta}_2)$ are disjoint with union $(\mathbf{X} \nleq \boldsymbol{\eta}_1)$. Inequality (2.24) follows from the fact that $\boldsymbol{\eta}_2 \geq \boldsymbol{\eta}_1 + \frac{\mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta}_1)}{1-p}(\boldsymbol{\eta}_2 - \boldsymbol{\eta}_1)$ as $\mathrm{P}(\mathbf{X} \nleq \boldsymbol{\eta}_1) \leq 1-p$, and $\mathbf{X} \preceq_{(1)} \mathbf{Y}$. Lastly, equality (2.25) is simply due to Definition 3.

The second statement directly results from the formulation of VMCVaR as an optimization problem in (2.7): For any $\boldsymbol{\vartheta}_2 \in \mathrm{VMCVaR}_p(\mathbf{Y})$, there exists some $\boldsymbol{\eta}$ such that $\mathrm{P}(\mathbf{Y} \leq \boldsymbol{\eta}) \geq p$ and $\boldsymbol{\vartheta}_2 = \mathrm{MCVaR}_p(\mathbf{Y}, \boldsymbol{\eta})$. By assumption, we have $\mathrm{P}(\mathbf{X} \leq \boldsymbol{\eta}) \geq p$ and $\mathrm{MCVaR}_p(\mathbf{X}, \boldsymbol{\eta}) \leq \mathrm{MCVaR}_p(\mathbf{Y}, \boldsymbol{\eta})$ since the function $\mathrm{MCVaR}_p(\cdot, \boldsymbol{\eta})$ is nondecreasing. Finally, by definition, there exists some $\boldsymbol{\vartheta}_1 \in \mathrm{VMCVaR}_p(\mathbf{X})$ such that $\boldsymbol{\vartheta}_1 \leq \mathrm{MCVaR}_p(\mathbf{X}, \boldsymbol{\eta})$.

□

Note that Lee and Prékopa [64] show that the properties in Proposition 3 hold for their definition of MCVaR (we will give this definition in detail in Section 2.3). While we use the same definition of MVaR as [64], our definition of MCVaR is different, hence we need to prove these properties for our definition. Lee and Prékopa [64] also show that a multivariate analogue of the subadditivity property, which is included in the coherence definition of univariate risk measures, is not satisfied by their MCVaR definition. We consider this property next.

Remark 2. *Risk measure* MCVaR *does not satisfy the following analogue of the subadditivity property:*

$$\mathrm{MCVaR}_p(\mathbf{X}, \boldsymbol{\eta}_x) + \mathrm{MCVaR}_p(\mathbf{Y}, \boldsymbol{\eta}_y) \geq \mathrm{MCVaR}_p(\mathbf{X}+\mathbf{Y}, \boldsymbol{\eta}_{x+y}),$$
$$\forall \boldsymbol{\eta}_x \in \mathrm{MVaR}_p(\mathbf{X}), \boldsymbol{\eta}_y \in \mathrm{MVaR}_p(\mathbf{Y}), \boldsymbol{\eta}_{x+y} \in \mathrm{MVaR}_p(\mathbf{X}+\mathbf{Y}).$$

Proof. Consider the following counterexample. Let $\mathbf{X}$ and $\mathbf{Y}$ be dependent bivariate random vectors with joint support

$$(\mathbf{X},\mathbf{Y}) \in \{((1,5),(4,1.5)),((3,2),(1,3)),((2,1),(2,5)),((1,4),(2,3)),((5,5),(3,1))\},$$

such that each realization is equally likely. At confidence level $p = 0.6$, $\mathrm{MVaR}_p(\mathbf{X}) = \{(3,4),(2,5)\}$, $\mathrm{MVaR}_p(\mathbf{Y}) = \{(3,3),(2,5)\}$ and $\mathrm{MVaR}_p(\mathbf{X}+\mathbf{Y}) = \{(4,7),(8,6)\}$. Using this information, we compute $\mathrm{VMCVaR}_p(\mathbf{X}) = \{(4,5)\}$, $\mathrm{VMCVaR}_p(\mathbf{Y}) = \{(3.5,4)\}$ and $\mathrm{VMCVaR}_p(\mathbf{X}+\mathbf{Y}) = \{(6.5,7),(8,6.75)\}$. It can be seen that $(4,5)+(3.5,4) \ngeq (8,6.75)$, hence subadditivity property is violated. □

In addition to coherence, from definition (2.4), it is clear that in the univariate case, CVaR is a conservative approximation of VaR, as $\mathrm{CVaR}_p(V) \geq \mathrm{VaR}_p(V)$ for a univariate random variable V. Analogously, it is expected that the multivariate CVaR value should not dominate its corresponding pLEP, i.e., we should not have $\mathrm{MCVaR}_p(\mathbf{X},\boldsymbol{\eta}) < \boldsymbol{\eta}$ for any $\boldsymbol{\eta} \in \mathrm{MVaR}_p(\mathbf{X})$ for a random vector $\mathbf{X} \in \mathbb{R}^d$. By definition (2.2), $\mathrm{VMCVaR}_p(\mathbf{X})$ clearly satisfies this property. In the next section, we show that this may not be the case for some existing definitions of multivariate CVaR.

2.3 Comparison to Existing Definitions

In this section, we briefly review other multivariate CVaR definitions from the literature and compare them with the proposed definition of MCVaR.

Prékopa [88] utilizes the definition of MVaR given in Definition 2 as a basis for their multivariate CVaR definition. Using the definition of $\mathrm{MVaR}_p(\mathbf{X})$, they assume an outcome $\mathbf{X}$ to be desirable if

$$\mathbf{X} \in \bigcup_{\mathbf{s}\in \mathrm{MVaR}_p(\mathbf{X})} (\mathbf{s} + \mathbb{R}^d_-) \tag{2.26}$$

and undesirable if

$$\mathbf{X} \in \bigcap_{\mathbf{s} \in \mathrm{MVaR}_p(\mathbf{X})} (\mathbf{s} + \mathbb{R}^d_-)^c,$$

where A^c is the complement of set A. These definitions imply that an event is undesirable if it is not better than or equivalent to any element in $\mathrm{MVaR}_p(\mathbf{X})$. Here, we assume that set $\mathbb{R}^d_-$ includes the zero vector along the same lines with the univariate definition of VaR. Assume that the set of desirable outcomes, defined in the right-hand-side of (2.26), is denoted as D_p. The set $\overline{D_p^c}$ represents the closure of its complement, which is the set of undesirable outcomes.

Definition 4 (Prékopa [88])**.** *The Multivariate Conditional Value-at-Risk of the random vector* $\mathbf{X}$ *at confidence level* p*, denoted by* $\overline{\mathrm{MCVaR}}_p(\mathbf{X})$*, is defined as*

$$\overline{\mathrm{MCVaR}}_p(\mathbf{X}) = \mathbb{E}[\psi(\mathbf{X})|\mathbf{X} \in \overline{D_p^c}],$$

where ψ *is a d-variate function such that* $\mathbb{E}[\psi(\mathbf{X})]$ *exists.*

In particular, the function $\psi(\mathbf{z})$ in [88] is defined as $\psi(\mathbf{z}) = \sum_{i=1}^{d} c_i z_i$, where $\sum_{i=1}^d c_i = 1$ and $\mathbf{c} \geq \mathbf{0}$. This scalarization function is consistent with the weighted sum approach commonly applied in multicriteria optimization, where c_i is interpreted as the relative weight of criterion $i \in [d]$. Using the expansion on conditional probability

$$\mathbb{E}[\psi(\mathbf{X})] = \mathbb{E}[\psi(\mathbf{X})|\mathbf{X} \notin D_p]\mathrm{P}(\mathbf{X} \notin D_p) + \mathbb{E}[\psi(\mathbf{X})|\mathbf{X} \in D_p]\mathrm{P}(\mathbf{X} \in D_p),$$

the multivariate CVaR value is computed as

$$\overline{\mathrm{MCVaR}}_p(\mathbf{X}) = \mathbb{E}[\psi(\mathbf{X})|\mathbf{X} \notin D_p] = \frac{1}{\mathrm{P}(\mathbf{X} \notin D_p)}(\mathbb{E}[\psi(\mathbf{X})] - \mathbb{E}[\psi(\mathbf{X})|\mathbf{X} \in D_p]\mathrm{P}(\mathbf{X} \in D_p)). \tag{2.27}$$

assuming that the probability of $\mathrm{P}(\mathbf{X} \in \mathrm{MVaR}_p(\mathbf{X}))$ is negligible. This assumption is quite natural in case of continuous distributions. For discrete probability distributions, it provides

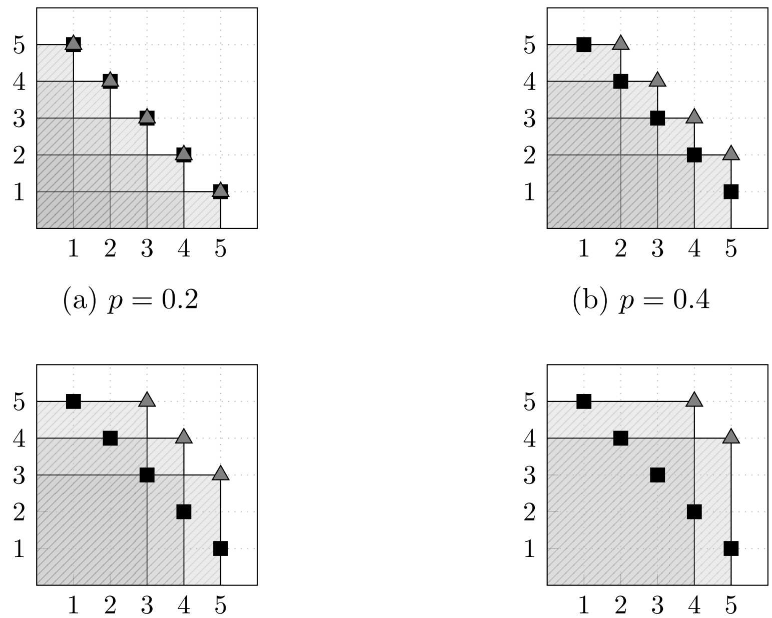

(a) $p = 0.2$

(b) $p = 0.4$

(c) $p = 0.6$

(d) $p = 0.8$

Figure 2.1: Desirable outcomes for different values of p in Example 2

an approximation.

A shortcoming of this definition in (2.27) is that $\overline{\text{MCVaR}}_p(\mathbf{X})$ fails to determine the risk associated with an outcome vector in some discrete cases, which we describe next.

Example 2. *Consider bivariate random variable* $\mathbf{X} \in \mathbb{R}^2$ *with support*

$$\{(1,5),(2,4),(3,3),(4,2),(5,1)\}.$$

All realizations are assumed to have equal probability, 0.2. For different values of p, the set of desirable outcomes is depicted in Figure 2.1 as the gray shaded areas. The square- and triangle-shaped points represent the realizations and pLEPs of $\mathbf{X}$*, respectively. As it can be seen, all possible realizations of* $\mathbf{X}$ *are desirable for each possible value* $0.4 \leq p \leq 0.8$*. Consequently,* $\overline{\text{MCVaR}}_p(\mathbf{X})$ *is undefined at any level* $0.4 \leq p \leq 0.8$*.*

Another problem is the possibility of obtaining an $\overline{\text{MCVaR}}_p(\mathbf{X})$ value that is preferable to $\psi(\boldsymbol{\eta})$ for some $\boldsymbol{\eta} \in \text{MVaR}_p(\mathbf{X})$. Consider the following example.

Example 3. *Let $\mathbf{X} \in \mathbb{R}^2$ be a bivariate random variable with support*

$$\{(1,5),(2,4),(3,3),(4,2),(5,1)\},$$

such that $\text{P}(\mathbf{X} = (1,5)) = \text{P}(\mathbf{X} = (5,1)) = 0.05$ and all the other realizations are equally likely. At confidence level $p = 0.9$, we obtain $\text{MVaR}_p(\mathbf{X}) = \{(4,4)\}$. For pLEP (4,4), the realizations (1,5) and (5,1) are undesirable. The $\overline{\text{MCVaR}}$ value corresponding to this pLEP is $\psi((3,3))$, which is smaller than $\psi((4,4))$, for any increasing linear function $\psi(\cdot)$ including the one proposed in [88].

Hamel et al. [50] propose an alternative multivariate CVaR definition as a convex set (in fact, a polyhedral cone pointed at a vector v) obtained by adding a polyhedral cone, C, containing the non-negative orthant to the vector v each component of which corresponds to the univariate CVaR value of the respective criterion, i.e., $v = (\text{CVaR}_{p_1}(X_1), \dots, \text{CVaR}_{p_d}(X_d))$ for given confidence levels $p_1, \dots, p_d$. For instance, the minimum multivariate CVaR vector according to this definition for the random vector $\mathbf{X}$ in Example 2 at $p_1 = p_2 = 0.8$, where $C = \mathbb{R}^d_+$ is given by $(\text{CVaR}_{0.8}(X_1), \text{CVaR}_{0.8}(X_2)) = (5,5)$. Note that using the individual CVaR values for each criterion does not consider the dependencies between the individual criteria, and may result in a conservative estimate.

Another vector-valued alternative to the $\overline{\text{MCVaR}}$ of Prékopa [88] for defining multivariate CVaR is proposed by Cousin and Di Bernardino [23] under the name lower-orthant conditional tail expectation ($\underline{\text{CTE}}$).

Definition 5 (Cousin and Di Bernardino [23])**.** *For a given random vector $\mathbf{X} \in \mathbb{R}^d$ with distribution function F at level p, the lower-orthant conditional tail expectation is defined as*

$$\underline{CTE}_p(\mathbf{X}) = \mathbb{E}[\mathbf{X} | F(\mathbf{X}) \geq p]. \tag{2.28}$$

This definition includes only the outcomes dominated by or equal to a pLEP in the computation of multivariate CVaR, whereas the $\overline{\text{MCVaR}}$ of Prékopa [88] considers the outcomes worse than all pLEPs in at least one criterion (see also [96] for a related definition). As in the previous definition, $\underline{\text{CTE}}_p(\mathbf{X})$ fails to measure the risk associated with some discrete random vectors (consider Example 2 at any confidence level $p > 0.2$).

Next we compare these risk measures with the proposed definition of MCVaR for the simple case of a random variable with a single pLEP. Assume that we are given a random vector $\mathbf{X} \in \mathbb{R}^d$ such that $\boldsymbol{\eta}$ is the only pLEP of $\mathbf{X}$ at confidence level p, i.e. $\text{MVaR}_p(\mathbf{X}) = \{\boldsymbol{\eta}\}$. By doing so, we ensure that all multivariate CVaR definitions correspond to a single vector and hence they can be compared. Furthermore, we consider the vector-valued adaptation of $\overline{\text{MCVaR}}_p(\mathbf{X})$, denoted by $\overline{\text{VMCVaR}}_p(\mathbf{X})$, where we replace the scalarization function $\psi(\mathbf{X})$ in Definition 4 with the vector $\mathbf{X}$. In particular, we define

$$\overline{\text{VMCVaR}}_p(\mathbf{X}) = \mathbb{E}[\mathbf{X} | X_i \geq \eta_i \text{ for some } i \in [d]]. \tag{2.29}$$

as the expected vector value of undesirable outcomes as defined in [88]. As explained previously, this may result in a multivariate CVaR value that dominates $\boldsymbol{\eta}$. The following results would also hold for the case that all measures are scalarized using the same vector.

We represent definition (2.28) equivalently as,

$$\underline{\text{CTE}}_p(\mathbf{X}) = \mathbb{E}[\mathbf{X} | X_i \geq \eta_i \ \forall i \in [d]].$$

Note that in the following discussion, when the set $\text{VMCVaR}_p(\mathbf{X})$ has a single element in it, we refer to that element as $\text{VMCVaR}_p(\mathbf{X})$, as well.

Proposition 4. *For a given random vector* $\mathbf{X}$ *such that* $\text{MVaR}_p(\mathbf{X}) = \{\boldsymbol{\eta}\}$, $\text{P}(\mathbf{X} \leq \boldsymbol{\eta}) = p$, *and* $\text{P}(\mathbf{X} = \boldsymbol{\eta})$ *is negligible, the following relations hold,*

$$\overline{\text{VMCVaR}}_p(\mathbf{X}) \leq \text{VMCVaR}_p(\mathbf{X}) \leq \underline{CTE}_p(\mathbf{X}),$$

when $\overline{\text{VMCVaR}}_p(\mathbf{X})$ *and* $\underline{CTE}_p(\mathbf{X})$ *are well-defined for* $\mathbf{X}$.

Proof. Note that the conditions of the proposition ensure that the risk measures are comparable. We consider the case that all risk measures are well-defined for $\mathbf{X}$ (recall that $\underline{\text{CTE}}_p(\mathbf{X})$ and $\overline{\text{VMCVaR}}_p(\mathbf{X})$ may not be defined in some cases). Since $\overline{\text{VMCVaR}}_p(\mathbf{X})$ defined in (2.29) takes the expectation of $\mathbf{X}$ while $\text{VMCVaR}_p(\mathbf{X})$ takes the expectation of $\max(\mathbf{X}, \boldsymbol{\eta})$ (see Proposition 2) and the condition in $\text{VMCVaR}_p(\mathbf{X})$ given in Proposition 2 is more restrictive, we have $\overline{\text{VMCVaR}}_p(\mathbf{X}) \leq \text{VMCVaR}_p(\mathbf{X})$. To show that $\text{VMCVaR}_p(\mathbf{X}) \leq \underline{\text{CTE}}_p(\mathbf{X})$, we represent $\underline{\text{CTE}}_p(\mathbf{X})$ equivalently as,

$$\underline{\text{CTE}}_p(\mathbf{X}) = \mathbb{E}[\max(\mathbf{X}, \boldsymbol{\eta}) | X_i \geq \eta_i \ \forall i \in [d]].$$

The first terms in both definitions are the same except for the conditions, $\underline{\text{CTE}}_p(\mathbf{X})$ conditions on the realizations with $\mathbf{X} \geq \boldsymbol{\eta}$, while $\text{VMCVaR}_p(\mathbf{X})$ conditions on the realizations with $\mathbf{X} \not\leq \boldsymbol{\eta}$). Hence the claim follows.

□

In summary, we propose a new definition of a vector-valued multivariate CVaR that is consistent with its univariate counterpart. We compare this definition with alternative vector-valued definitions, and show its advantages. However, we recognize the computational difficulties involved with obtaining this vector, especially when used in an optimization setting. In this case, we believe that the scalarization approaches in [61, 68, 81, 82, 84] provide a practical and reasonable estimation of the multivariate risk.

Chapter 3

TWO-STAGE STOCHASTIC PROGRAMMING UNDER MULTIVARIATE RISK CONSTRAINTS WITH AN APPLICATION TO HUMANITARIAN RELIEF NETWORK DESIGN

3.1 Introduction

This chapter is based on [84]. Two-stage stochastic programming is a vibrant research area, which provides a natural and widely applicable modeling framework for decision making problems under uncertainty in a large variety of fields. Such models are well-suited for the situations where decisions are made in two stages; the first-stage decisions are made before the uncertainty is revealed, and the second-stage (recourse) decisions are made given the predetermined first-stage decisions and the observed realization of the random parameters. For many practical decision making problems, it is essential to evaluate the decisions according to multiple and possibly conflicting stochastic criteria. Along these lines, we focus on two-stage stochastic programming models involving first-stage decisions leading to uncertain outcomes that can be evaluated according to multiple stochastic performance measures associated with the corresponding second-stage decisions.

In multicriteria stochastic optimization, one popular approach is to aggregate the specified multiple stochastic objective criteria and obtain a single objective function. Following this so-called weighted-sum approach, we formulate the second-stage problem as a multiobjective optimization problem and focus on its expected optimal objective value in the first stage. In addition, one may also prefer to impose constraints on some performance measures associated with the second-stage decisions. To this end, we employ stochastic benchmarking preference relations between the vector-valued random outcomes of the second-stage decisions. In particular, in the first class of problems we study, the first-stage problem enforces the multivariate polyhedral CVaR constraint introduced by Noyan and Rudolf [81]. This approach considers a family of linear scalarization functions (a polyhedral set of scalarization weight vectors) and requires all scalarized versions of the decision-based random vector of outcomes to be preferable to the corresponding scalarizations of an existing reference (benchmark) outcome according to the univariate CVaR relation. Furthermore, allowing arbitrary polyhedra as scalarization sets also provides a good balance between flexibility and computational tractability. To the best of our knowledge, our study is a first in developing such a risk-averse two-stage model featuring multivariate risk measure-based constraints. We also note that our proposed modeling framework is not limited to the risk-neutral objective function; it can be easily extended to the case where the objective function features also a risk measure representing the decision makers' risk preferences regarding the random recourse cost. For acceptable risk-averse objectives, which lead to computationally tractable two-stage stochastic programming models, we refer the reader to [4] and [38].

Optimization problems with the multivariate stochastic benchmarking constraints have been receiving increasing attention in the literature. Using multivariate stochastic relations is essential to capture the correlation between the multiple random outcomes. The majority of the existing studies extend a univariate stochastic preference relation, based on second-order stochastic dominance (SSD) or a coherent risk measure such as CVaR, to the multivariate case by considering a family of scalarization functions. While multivariate SSD-constrained problems are more frequently studied (see, e.g., [32, 33, 34, 52, 55]), there is an increasing interest in risk measure-constrained variants (see, e.g., [68, 81]), which provide natural relaxations to overly demanding and conservative SSD-based models. More recently, studies focusing on both multivariate SSD- and risk measure-constrained models also appear in literature [61, 82]. In these existing studies, the scalarization functions are almost exclusively of the linear type with the exception of Noyan and Rudolf [82] who extend the multivariate risk-constrained models by incorporating a general class of scalarization functions (including a variety of nonlinear scalarization functions). Moreover, we point out that the scalarization-based line of research has been dedicated almost entirely to single-stage (static) decision-making problems. The only exception is the work of Dentcheva and Wolfhagen [34], which introduces a two-stage stochastic optimization problem involving a multivariate SSD (more precisely, referred to as the increasing convex ordering in the context of cost minimization) constraint on a random outcome vector of the second-stage decisions with respect to the unit simplex of scalarization vectors. The authors present two approximate decomposition-based solution algorithms which rely on the Lagrangian relaxation of the multivariate stochastic ordering constraint and show their finite convergence to an ε-feasible ε-optimal solution, even if the probability space is not finite; their results and algorithms are adaptations of those presented in their earlier study [33] for the single-stage case. We contribute to this line of research by introducing an alternative two-stage model with multivariate CVaR constraints. In addition, we consider the SSD-based counterpart (as in [33]) and provide a new computationally tractable and exact solution algorithm for this problem class. We defer a detailed discussion on the advantages of our solution method compared to the existing approaches for the multivariate SSD-constrained two-stage models to Section 3.3.

Although, it is not directly related to the multicriteria stochastic optimization problems of our interest, we note that there are a few studies on risk-averse two-stage models with univariate (first- or second-order) stochastic dominance and CVaR-based constraints (see, e.g., [28, 40, 48]). However, they employ the risk constraints to compare scalar-based random variables. Our study extends such risk-averse modeling approaches to the multicriteria case, allowing us to consider additional stochastic criteria associated with the second-stage decisions other than the random recourse cost.

The proposed modeling approach with the multivariate risk constraints is a fairly recent research area, and it has promise to be applied in a wide variety of fields. This approach is particularly well-suited for the field of humanitarian logistics, since incorporating risk is crucial for rarely occurring disaster events (e.g., [80]), and considering multiple conflicting performance criteria (such as efficiency and equity) is often essential for the effectiveness of the relief response systems (e.g., [49, 56, 113]). Motivated by the significance of the long-term pre-disaster planning, we apply the proposed framework to a stochastic pre-disaster relief network design problem. Hence, our study also contributes to the humanitarian relief literature by introducing a new risk-averse two-stage optimization model, which provides a flexible and tractable way of considering decision makers' risk preferences based on multiple stochastic criteria.

Next we summarize our contributions and give an outline of our paper. In Section 3.2, we describe a novel two-stage stochastic program with multivariate CVaR constraints. It is well-known that risk-neutral two-stage stochastic programs are generally hard to solve due to a potentially large number of scenario-dependent recourse decisions. Thus, introducing a multivariate stochastic benchmarking relation, which enforces a collection of risk constraints associated with a scalarization set, further complicates the solution of these optimization models. A common approach in solving such large-scale stochastic optimization models is to employ a Benders-type scenario decomposition approach (e.g., [99]). The classical decomposition methods cannot be directly applied to our model due to the complicating risk constraints. Utilizing successive relaxations of the multivariate polyhedral CVaR relation, we develop two types of exact and finitely convergent delayed cut generation solution algorithms; the iteratively generated cuts – common in both algorithms – are associated with the scalarization vectors for which the risk constraints are relaxed. In addition, the second algorithm adapts a scenario decomposition approach that exploits the decomposable structure of CVaR and the second-stage problems; this further decomposition proves to be useful in solving larger problem instances. We also develop optimality conditions and strong duality results under certain linearity assumptions. While we propose natural primal algorithms to solve the problem of interest, these duality results provide a foundation for the potential future development of alternative dual-based solution methods. Despite our focus on the CVaR-based models, in Section 3.3, we show that our proposed solution methods also apply to the multivariate SSD-constrained two-stage models. In Section 3.4, we apply the proposed modeling approach to a stochastic pre-disaster relief network design problem and present the corresponding mathematical programming formulations. Section 3.5 is dedicated to the computational study.

3.2 Two-Stage Optimization with Multivariate CVaR Constraints

In this section, we introduce the multicriteria stochastic decision-making framework of interest, and present the proposed multivariate CVaR-constrained two-stage stochastic programming model and the corresponding delayed cut generation solution algorithms. We also develop optimality conditions and strong duality results under certain linearity assumptions.

We consider a finite probability space $(\Omega, 2^{\Omega}, \Pi)$ with $\Omega = \{\omega_1, \ldots, \omega_m\}$ and $\Pi(\omega_s) = p_s$. Let $S = \{1, \ldots, m\}$ be the index set of the elementary events (also referred to as scenarios). First, we give the general form of a traditional risk-neutral two-stage stochastic program

$$\min \quad f(\mathbf{x}) + \mathbb{E}(Q(\mathbf{x}, \boldsymbol{\xi}(\omega))) \tag{3.1a}$$

$$\text{s.t.} \quad \mathbf{x} \in \mathcal{X}. \tag{3.1b}$$

Here, $f(\mathbf{x})$ is a convex objective function, and $\mathcal{X} \subset \mathbb{R}_+^{\bar{n}_1} \times \mathbb{Z}_+^{n_1 - \bar{n}_1}$ is a non-empty convex set with additional integrality restrictions (such as the set of integer points in a polyhedron). The set $\mathcal{X}$ is defined by the deterministic constraints on the first-stage decision variables. $\boldsymbol{\xi}(\omega_s) \subset \mathbb{R}^r$ and $Q(\mathbf{x}, \boldsymbol{\xi}(\omega_s))$ designate the vector of the random input parameters of the second-stage problem and the optimal second-stage objective value under scenario s, respectively. More specifically, introducing the notation $\boldsymbol{\xi}(\omega_s) = (\mathbf{q}(\omega_s), T(\omega_s), W(\omega_s),$
$\mathbf{h}(\omega_s)) = (\mathbf{q}_s, T_s, W_s, \mathbf{h}_s)$,

$$Q(\mathbf{x}, \boldsymbol{\xi}(\omega_s)) = \min_{\mathbf{y} \in \mathcal{Y}(\mathbf{x}, \boldsymbol{\xi}(\omega_s))} \quad \mathbf{q}_s^\top \mathbf{y}, \tag{3.2}$$

where $\mathcal{Y}(\mathbf{x}, \boldsymbol{\xi}(\omega_s)) = \{\mathbf{y} \in \mathbb{R}_+^{n_2} : T_s\mathbf{x} + W_s\mathbf{y} \geq \mathbf{h}_s\}$. The second-stage objective function $\mathbf{q}_s^\top \mathbf{y}$ could be a weighted sum of multiple objectives as in [34], or a single objective, such as cost, that is compatible with the first-stage objective.

Next, we describe how we develop a risk-averse variant of the risk-neutral two-stage model (3.1) by incorporating a multivariate risk constraint. For the outcome mapping $\hat{\mathbf{G}}$ corresponding to the performance measures of interest, we introduce the decision-based random vector $\hat{\mathbf{G}}(\mathbf{x},\mathbf{y})$ given by $\left[\hat{\mathbf{G}}(\mathbf{x},\mathbf{y})\right](\omega) = \hat{\mathbf{G}}(\mathbf{x},\mathbf{y}(\omega),\boldsymbol{\xi}(\omega))$. In particular, $\hat{\mathbf{G}}(\mathbf{x},\mathbf{y}(\omega),\boldsymbol{\xi}(\omega))$ represents the multiple random performance measures of interest associated with the second-stage decisions $\mathbf{y}(\omega)$ for given $\mathbf{x}$ and $\boldsymbol{\xi}(\omega)$. Finally, we let $\hat{\mathbf{G}}(\mathbf{x},\mathbf{y}(\omega_s),\boldsymbol{\xi}(\omega_s)) = \hat{\mathbf{g}}_s(\mathbf{x},\mathbf{y}_s) = (\mathrm{g}_s^1(\mathbf{x},\mathbf{y}_s), \mathrm{g}_s^2(\mathbf{x},\mathbf{y}_s), \ldots, \mathrm{g}_s^d(\mathbf{x},\mathbf{y}_s))^\top,\ s \in S$, and assume that $\mathrm{g}_s^i(\mathbf{x},\mathbf{y}_s)$ is an affine function of $\mathbf{x}$ and $\mathbf{y}_s$ such that $\mathrm{g}_s^i(\mathbf{x},\mathbf{y}_s) = \bar{\mathbf{g}}_s^i\mathbf{x} + \tilde{\mathbf{g}}_s^i\mathbf{y}_s$ for all $i \in \{1,\ldots,d\}$ and $s \in S$. We also assume that a corresponding benchmark random outcome vector and a polyhedron of scalarization vectors, each component of which represents the relative importance of each decision criterion, are given (specified by decision makers). Let $\mathbf{Z}$ be the benchmark vector with realizations $\mathbf{z}_1,\ldots,\mathbf{z}_{|\tilde{S}|}$ and associated probabilities $\tilde{p}_i,\ i \in \tilde{S}$; it is often constructed from some benchmark decision in which case we have $\tilde{S} = S$ and $\tilde{p}_i = p_i,\ i \in S$. The scalarization set $\mathcal{C}$ is allowed to be an *arbitrary polyhedron*. Without loss of generality (see, e.g., [82]), considering the interpretation of the scalarization vectors, $\mathcal{C}$ is naturally assumed to be a subset of the unit simplex. In our risk-averse modeling approach, we introduce a multivariate risk constraint to ensure that the decision-based random outcome vector $\hat{\mathbf{G}}(\mathbf{x},\mathbf{y})$ is preferable to the specified benchmark $\mathbf{Z}$ according to the multivariate CVaR relation introduced by Noyan and Rudolf [81]. According to this relation, all scalarized versions of $\hat{\mathbf{G}}(\mathbf{x},\mathbf{y})$ are preferable to the corresponding scalarizations of the benchmark outcome according to the univariate CVaR-based preference relation. In our setting, smaller values of random variables and risk measures are considered to be preferable.

Using the above notations and conventions, we now introduce the class of two-stage stochastic programming problems with multivariate CVaR constraints:

$$\min \quad f(\mathbf{x}) + \sum_{s\in S} p_s Q(\mathbf{x},\boldsymbol{\xi}(\omega_s)) \tag{3.3a}$$

$$\text{s.t.} \quad \mathrm{CVaR}_\alpha(\mathbf{c}^\top\hat{\mathbf{G}}(\mathbf{x},\mathbf{y})) \le \mathrm{CVaR}_\alpha(\mathbf{c}^\top\mathbf{Z}), \quad \forall\ \mathbf{c}\in\mathcal{C}, \tag{3.3b}$$

$$\mathbf{x} \in \mathcal{X}, \tag{3.3c}$$

$$Q(\mathbf{x}, \boldsymbol{\xi}(\omega_s)) = \mathbf{q}_s^\top \mathbf{y}(\omega_s), \quad \forall s \in S, \tag{3.3d}$$

$$\mathbf{y}(\omega_s) \in \mathcal{Y}(\mathbf{x}, \boldsymbol{\xi}(\omega_s)), \quad \forall s \in S, \tag{3.3e}$$

where for a given $\hat{\mathbf{c}}_{(l)} \in \mathcal{C}$, $\mathrm{CVaR}_\alpha(\hat{\mathbf{c}}_{(l)}^\top \hat{\mathbf{G}}(\mathbf{x}, \mathbf{y}))$ is calculated as (see [93])

$$\min \quad \Bigg\{ \eta_l + \frac{1}{1-\alpha} \sum_{s \in S} p_s w_{sl} \ : \ w_{sl} \geq \hat{\mathbf{c}}_{(l)}^\top \hat{\mathbf{G}}(\mathbf{x}, \mathbf{y}(\omega_s), \boldsymbol{\xi}(\omega_s)) - \eta_l \quad \forall\ s \in S,\ w_{sl} \geq 0 \quad \forall\ s \in S,\ \eta_l \in \mathbb{R} \Bigg\}. \tag{3.4}$$

In this model, for a first-stage decision vector $\mathbf{x}$, the corresponding second-stage decisions are determined among those for which the risk constraint (3.3b) is satisfied and each scenario-based recourse cost is minimized. Throughout the paper, we assume that problem (3.3), if feasible, has a finite objective value, i.e., $f(\mathbf{x}), \mathbb{E}(Q(\mathbf{x}, \boldsymbol{\xi}(\omega))) > -\infty$ for all $\mathbf{x} \in \mathcal{X}$.

We note that using a scalarization-based risk constraint such as (3.3b) is useful to address ambiguities and inconsistencies in the weight vectors; for detailed discussions we refer to [54] and [68]. Moreover, CVaR is a widely applied popular risk measure with several desirable properties; it is a law invariant coherent risk measure [6], serves as a fundamental building block for other coherent risk measures, and can be used to capture a wide range of risk preferences, including risk-neutral and worst-case approaches. In addition, [81] highlights that CVaR-based relations provide flexible, meaningful, and computationally tractable relaxations for overly conservative SSD relations. In a broad sense, CVaR can be viewed as a weighted sum of the least favorable outcomes (those exceeding the α-quantile – also known as value-at-risk at confidence level α).

The next proposition presents the deterministic equivalent formulation (DEF) of the proposed risk-averse model (3.3). The proof leads to another equivalent model in the form of a risk-neutral two-stage stochastic program, which proves to be useful in developing a scenario decomposition-based solution algorithm.

Proposition 5. *The following deterministic formulation is equivalent to the two-stage model given by* (3.3)*:*

$$\min \quad f(\mathbf{x}) + \sum_{s \in S} p_s \mathbf{q}_s^\top \mathbf{y}_s \tag{3.5a}$$

$$s.t. \quad \eta_l + \frac{1}{1-\alpha} \sum_{s \in S} p_s w_{sl} \leq \text{CVaR}_\alpha(\hat{\mathbf{c}}_{(l)}^\top \mathbf{Z}), \quad \forall\, l = 1, \ldots, \bar{L}, \tag{3.5b}$$

$$w_{sl} \geq \hat{\mathbf{c}}_{(l)}^\top \hat{\mathbf{g}}_s(\mathbf{x}, \mathbf{y}_s) - \eta_l, \quad \forall\, s \in S,\ l = 1, \ldots, \bar{L}, \tag{3.5c}$$

$$w_{sl} \geq 0, \quad \forall\, s \in S,\ l = 1, \ldots, \bar{L}, \tag{3.5d}$$

$$\mathbf{x} \in \mathcal{X}, \quad \boldsymbol{\eta} \in \mathbb{R}^{\bar{L}}, \tag{3.5e}$$

$$T_s \mathbf{x} + W_s \mathbf{y}_s \geq \mathbf{h}_s, \quad \forall\, s \in S, \tag{3.5f}$$

$$\mathbf{y}_s \in \mathbb{R}_+^{n_2}, \quad \forall\, s \in S, \tag{3.5g}$$

where $\bar{L}$ is a finite integer and $\hat{\mathbf{c}}_{(l)}$, $l = 1, \ldots, \bar{L}$, are the projections of the vertices of the polyhedron $P = \{(\mathbf{c}, \eta, \mathbf{w}) \in \mathcal{C} \times \mathbb{R} \times \mathbb{R}_+^{|\tilde{S}|} : w_i \geq \mathbf{c}^\top \mathbf{z}_i - \eta,\ i \in \tilde{S}\}$.

Proof. First, observe that it is sufficient to consider a finite number of scalarization vectors from set $\mathcal{C}$, specifically the projections of the extreme points (given by $\hat{\mathbf{c}}_{(1)}, \ldots, \hat{\mathbf{c}}_{(\bar{L})}$) of the above defined polyhedron P [81], since P is only characterized by the given fixed (decision-independent) benchmark outcome vector. Hence, constraints (3.3b) can be replaced by

$$\text{CVaR}_\alpha(\hat{\mathbf{c}}_{(l)}^\top \hat{\mathbf{G}}(\mathbf{x}, \mathbf{y})) \leq \text{CVaR}_\alpha(\hat{\mathbf{c}}_{(l)}^\top \mathbf{Z}), \quad \forall\, l = 1, \ldots, \bar{L}. \tag{3.6}$$

Second, the univariate CVaR-relations in (3.6) can be represented by linear inequalities of type (3.5b)-(3.5d). This observation is a simple consequence of the LP representation of CVaR in (3.4), and the fact that $\text{CVaR}_\alpha(\mathbf{c}^\top \mathbf{Z})$ is a known constant given a vector $\mathbf{c}$ and the benchmark $\mathbf{Z}$. It is also easy to see that CVaR constraints (3.5c) can be decomposed over scenarios. This allows us to reformulate the proposed risk-averse model (3.3) as a risk-neutral two-stage program. In particular, w_{sl} and η_l, for $s \in S, l = 1, \ldots, \bar{L}$ are considered as first-stage decision variables and (3.3b) is replaced by (3.5b) and (3.5d), whereas the inequalities associated with $s \in S$ in (3.5c) are added to the constraints of the second-stage problem. As a result, we obtain the following risk-neutral two-stage model which is equivalent to (3.3):

$$\min \quad f(\mathbf{x}) + \sum_{s \in S} p_s \tilde{Q}(\mathbf{x}, \boldsymbol{\eta}, \mathbf{w}, \boldsymbol{\xi}(\omega_s)) \tag{3.7a}$$

$$\text{s.t.} \quad (3.5\text{b}), (3.5\text{d}) - (3.5\text{e}), \tag{3.7b}$$

where

$$\tilde{Q}(\mathbf{x}, \boldsymbol{\eta}, \mathbf{w}, \boldsymbol{\xi}(\omega_s)) = \quad \min \quad \mathbf{q}_s^\top \mathbf{y} \tag{3.8a}$$

$$\text{s.t.} \quad w_{sl} \geq \hat{\mathbf{c}}_{(l)}^\top \hat{\mathbf{g}}_s(\mathbf{x}, \mathbf{y}_s) - \eta_l, \quad \forall\, l = 1, \ldots, \bar{L}, \tag{3.8b}$$

$$T_s \mathbf{x} + W_s \mathbf{y} \geq \mathbf{h}_s, \quad \mathbf{y} \in \mathbb{R}_+^{n_2}. \tag{3.8c}$$

Then, by the well-known interchangeability principle of risk-neutral two-stage models (see, e.g., [99]), the assertion directly follows. □

Remark 3. *[**Mean-Risk Objective Function**] The proposed model can be easily extended to the case where the random objective outcomes are compared according to a mean-risk functional representing the decision makers' risk preferences. In particular,* $\mathbb{E}(Q(\mathbf{x}, \boldsymbol{\xi}(\omega))$ *can be replaced by* $\mathbb{E}(Q(\mathbf{x}, \boldsymbol{\xi}(\omega)) + \lambda \rho(Q(\mathbf{x}, \boldsymbol{\xi}(\omega))$ *in the objective function, where* ρ *is a risk functional such as* CVaR_α *and* λ *is a non-negative trade-off coefficient representing the exchange rate of mean cost for risk. Note that one can alternatively minimize* $\mathbb{E}(Q(\mathbf{x}, \boldsymbol{\xi}(\omega))$ *under the constraint of the form* $\rho(Q(\mathbf{x}, \boldsymbol{\xi}(\omega)) \leq \hat{\rho}$*, where* $\hat{\rho}$ *is a suitably specified threshold parameter. When* ρ *is a non-decreasing function and* $\lambda \geq 0$*, the mean-risk objective function preserves the convexity, and consequently, slight variants of our cutting plane solution algorithms, which we will describe in the next section, are applicable (for similar developments see, e.g., [4, 80]). Thus, our proposed modeling and solution framework is not limited to the risk-neutral objective function.*

We next present two types of algorithms to solve the computationally challenging stochastic programming model (3.3). Then, we provide a theoretical background to our formulations by developing strong duality results and optimality conditions for an important special case. These duality results generalize the previous results for single-stage multivariate CVaR-constrained problems [81]. Related results for two-stage multivariate SSD-constrained optimization are given in [34].

3.2.1 Delayed Cut Generation for Deterministic Equivalent Formulation

The deterministic equivalent formulation (3.5) features finite but potentially an exponential number of scalarization vectors, obtained as projections of the vertices of the polyhedron P. Hence, it is natural to develop a delayed cut generation algorithm, which avoids the impractical vertex enumeration approach required for explicitly constructing the risk constraints in advance. The proposed algorithms rely on successive relaxations of the multivariate polyhedral CVaR relation, and they iteratively generate cuts associated with the scalarization vectors for which the risk constraints are relaxed.

In our first cutting plane approach, at an intermediate iteration, the relaxed master problem (RMP^L) includes the constraint (3.3b) for a subset of $\mathcal{C}$, say $\{\tilde{\mathbf{c}}_{(1)}, \dots, \tilde{\mathbf{c}}_{(L)}\}$. Accordingly, RMP^L takes the form of (3.5), where $\bar{L}$ is replaced by L. Note that under the assumptions that $f(\mathbf{x})$ and $\hat{\mathbf{g}}_s(\mathbf{x}, \mathbf{y}_s)$, $s \in S$ are affine functions and $\mathcal{X}$ is a polyhedron (with additional integrality restrictions), the DEF (3.5) and RMP^L are large-scale linear (mixed-integer) programs. Given an optimal solution to the RMP^L $(\bar{\mathbf{x}}, \bar{\boldsymbol{\eta}}, \bar{\mathbf{w}}, \bar{\mathbf{y}})$, the following separation problem (SP) is solved to identify if there is any violated inequality (3.3b):

$$\max_{\mathbf{c} \in \mathcal{C}} \quad \mathrm{CVaR}_\alpha(\mathbf{c}^\top \hat{\mathbf{G}}(\bar{\mathbf{x}}, \bar{\mathbf{y}})) - \mathrm{CVaR}_\alpha(\mathbf{c}^\top \mathbf{Z}). \tag{3.9}$$

Noyan and Rudolf [81], Küçükyavuz and Noyan [61], and Liu et al. [68] provide alternative mixed-integer programming (MIP) formulations for the cut generation problem (3.9). Because we can treat the DEF as a large-scale single-stage multivariate CVaR-constrained

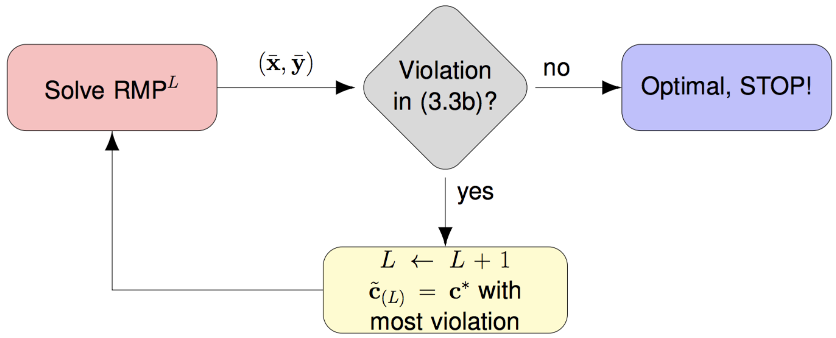

Figure 3.1: The delayed cut generation algorithm for DEF

problem, the convergence of the resulting delayed cut generation algorithm follows from [81]. We illustrate basic steps of the algorithm in Figure 3.1 assuming that the DEF is feasible.

Remark 4. *The cut generation formulations in [81], [61] and [68] are based on the opposite convention that larger values of random variables, as well as larger values of risk measures, are considered to be preferable. Those existing MIP formulations should be altered to reflect this difference in convention.*

Remark 5 (**General Applicability**)**.** *Observe that* $\hat{\mathbf{G}}(\bar{\mathbf{x}}, \bar{\mathbf{y}})$ *is just an input parameter for the separation problem. On a related note, we emphasize that the finite representation results for the multivariate CVaR and SSD relations, and the structure of the corresponding separation problems are independent of the outcome mapping* $\hat{\mathbf{G}}$*. In particular, the finite representations depend only on the benchmark* $\mathbf{Z}$ *and the scalarization set* $\mathcal{C}$ *(and not on the decision-based outcomes); see the polyhedron* P *in Proposition 5. The key implication is that the proposed delayed cut generation algorithm is well-defined for more general cases with nonlinear outcome mappings. In the general case, we can solve the RMP using nonlinear programming techniques. Under the convexity assumption on* $\hat{\mathbf{G}}$*, DEF* (3.5) *and RMP become large-scale (mixed-integer) convex programs, and our algorithm is also guaranteed to provide an optimal solution (if any exists).*

3.2.2 Delayed Cut Generation with Scenario Decomposition

The deterministic equivalent formulation (3.5) contains scenario-dependent variables and constraints, and consequently, its size grows as the number of scenarios increases. Because the number of scenarios is typically large, it is generally impractical to solve the DEF directly, even without the multivariate stochastic benchmarking constraints. The L-shaped method [112], which is a Benders-type scenario decomposition algorithm, is arguably the most commonly used computationally viable method to solve the classical (risk-neutral) two-stage stochastic programs. It relies on LP duality requiring the assumption that the second-stage problem (3.2) is a linear program. We propose a non-trivial extension of this approach to solve our multivariate CVaR-constrained two-stage model. Similar to its risk-neutral counterpart without the risk constraint, our linearity assumptions on the outcome mappings $\mathbf{g}_s(\mathbf{x}, \mathbf{y}_s),\ s \in S$, allow us to use LP duality to obtain finitely many linear Benders cuts.

In the DEF (3.5), constraints (3.5b)–(3.5c), which capture the CVaR relation along with the associated auxiliary variables $\boldsymbol{\eta}$ and $\mathbf{w}$, create further coupling of the scenarios in addition to the original coupling constraint (3.5f). Considering this structure, we handle the $\mathbf{x}$, $\boldsymbol{\eta}$, and $\mathbf{w}$ variables in the first stage, decompose the second-stage problems over scenarios once the first-stage variables are fixed, and solve iteratively the RMP involving only the first-stage decision variables and auxiliary decision variables ($\theta_s,\ s \in S$) for approximating the optimal second-stage objective function values. We obtain the following RMP at an intermediate iteration, where a subset of the scalarization vectors of cardinality L is generated so far:

$$
\begin{aligned}
(\mathrm{MP}^L)\ \min \quad & f(\mathbf{x}) + \sum_{s \in S} p_s \theta_s && (3.10a)\\
\text{s.t.} \quad & (\mathbf{x}, \boldsymbol{\eta}, \mathbf{w}, \boldsymbol{\theta}) \in \mathcal{O}, && (3.10b)\\
& (\mathbf{x}, \boldsymbol{\eta}, \mathbf{w}) \in \mathcal{F}, && (3.10c)\\
& \eta_l + \frac{1}{1-\alpha} \sum_{s \in S} p_s w_{sl} \leq \mathrm{CVaR}_\alpha(\tilde{\mathbf{c}}_{(l)}^\top \mathbf{Z}), \quad \forall\ l = 1, \ldots, L, && (3.10d)
\end{aligned}
$$

$$(3.5d) - (3.5e) \text{ (with } \bar{L} \text{ is replaced by } L). \tag{3.10e}$$

Here $\mathcal{O}$ is a polyhedral set defined by the constraints (referred to as the *optimality cuts*) that give valid lower bounding approximations of $\boldsymbol{\theta}$, and $\mathcal{F}$ is a polyhedral set defined by the constraints (referred to as the *feasibility cuts*) that represent the conditions for the first-stage decision vectors to yield feasible second-stage problems. In what follows, we give the explicit forms of the inequalities in the sets $\mathcal{O}$ and $\mathcal{F}$.

We start the algorithm with a small, potentially empty, subset of the scalarization set $\mathcal{C}$. At each iteration, given the subset $\{\tilde{\mathbf{c}}_{(1)}, \ldots, \tilde{\mathbf{c}}_{(L)}\}$, we first solve the relaxed master problem (MP^L) to obtain an optimal solution $\bar{\mathbf{v}} := (\bar{\mathbf{x}}, \bar{\boldsymbol{\eta}}, \bar{\mathbf{w}}, \bar{\boldsymbol{\theta}})$. Using this solution of the RMP, we solve the second-stage subproblem (PS$_s^L$) for each scenario $s \in S$ given by

$$(\text{PS}_s^L) \quad \bar{Q}_s^L := \min \quad \mathbf{q}_s^\top \mathbf{y} \tag{3.11a}$$

$$\text{s.t.} \quad -\tilde{\mathbf{c}}_{(l)}^\top \tilde{\mathbf{g}}_s \mathbf{y} \geq \tilde{\mathbf{c}}_{(l)}^\top \bar{\mathbf{g}}_s \bar{\mathbf{x}} - \bar{\eta}_l - \bar{w}_{sl}, \quad \forall\, l = 1, \ldots, L, \tag{3.11b}$$

$$W_s \mathbf{y} \geq \mathbf{h}_s - T_s \bar{\mathbf{x}}, \tag{3.11c}$$

$$\mathbf{y} \in \mathbb{R}_+^{n_2}, \tag{3.11d}$$

where constraint (3.11b) is equivalent to constraint (3.5c) after partial substitution of the values of the first-stage variables with $\bar{\mathbf{x}}, \bar{\boldsymbol{\eta}}$ and $\bar{\mathbf{w}}$. Let $\beta_{sl},\ l = 1, \ldots, L$, and $\boldsymbol{\gamma}_s$ for $s \in S$, be the dual variables associated with constraints (3.11b) and (3.11c), respectively. If the primal subproblem (PS$_s^L$) is infeasible for some $s \in S$, then let $(\boldsymbol{\beta}_s, \boldsymbol{\gamma}_s)$ provide an extreme ray (direction of unboundedness) of the corresponding dual feasible region and add the following Benders feasibility cut to the set $\mathcal{F}$:

$$0 \geq \boldsymbol{\gamma}_s^\top (\mathbf{h}_s - T_s \mathbf{x}) + \sum_{l=1}^{L} \beta_{sl}^\top \left(\tilde{\mathbf{c}}_{(l)}^\top \bar{\mathbf{g}}_s \mathbf{x} - \eta_l - w_{sl} \right). \tag{3.12}$$

If the primal subproblem is feasible and the current estimate of the optimal second-stage objective value ($\bar{\theta}_s$) is less than the actual optimal second-stage objective value ($\bar{Q}_s^L$), then let $(\boldsymbol{\beta}_s, \boldsymbol{\gamma}_s)$ provide an optimal dual extreme point and add the following Benders optimality cut to the set $\mathcal{O}$:

$$\theta_s \geq \boldsymbol{\gamma}_s^\top (\mathbf{h}_s - T_s \mathbf{x}) + \sum_{l=1}^{L} \beta_{sl}^\top \left(\tilde{\mathbf{c}}_{(l)}^\top \bar{\mathbf{g}}_s \mathbf{x} - \eta_l - w_{sl} \right). \tag{3.13}$$

Now suppose that all primal subproblems are feasible with the finite objective values $\bar{Q}_s^L, s \in S$ (recall our assumption that the original problem is not unbounded). In this case, given $(\bar{\mathbf{x}}, \bar{\mathbf{y}})$, we solve the separation problem (3.9) to obtain an optimal solution $\mathbf{c}^*$. If the corresponding optimal objective value is non-positive, then there is no violation in the CVaR constraint (3.3b). Furthermore, recall that in this case, all second-stage problems are feasible as well. In other words, the current first-stage solution is feasible for the original problem, so we update the upper bound on the optimal objective value (denoted by u). On the other hand, if the optimal objective value of the SP is positive, then $\mathbf{c}^* \in \mathcal{C}$ creating the most violation in constraint (3.11b) is added as the $L+1$th scalarization vector both to the RMP and the second-stage subproblems with $L \leftarrow L+1$. We repeat these iterations until either an infeasible RMP is detected, or the RMP objective function value is within the given tolerance ϵ of the upper bound u (see Figure 3.2 for an overview of Algorithm 1 assuming feasibility of the DEF). The pseudo-code of the proposed decomposition algorithm is provided in Algorithm 1. Note that the feasible regions of the subproblems change depending on the scalarization vectors included in the formulations at that iteration. In addition, the RMP also grows both in terms of the number of variables and constraints as we add more scalarization vectors to its formulation due to the addition of the η_l and $\mathbf{w}_l$ variables for each scalarization vector $\tilde{\mathbf{c}}_{(l)}$, as illustrated in Figure 3.3. As a result, our proposed algorithm can be seen as a *delayed column and cut generation algorithm*.

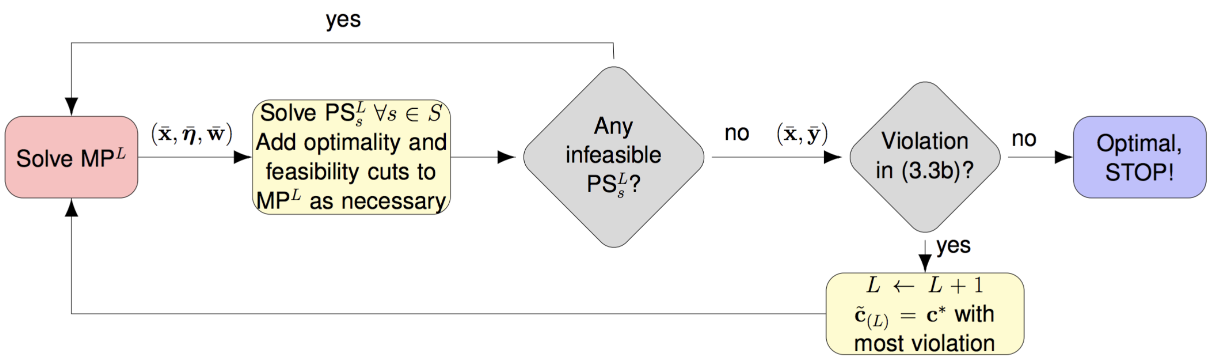

Figure 3.2: The delayed cut generation algorithm with scenario decomposition

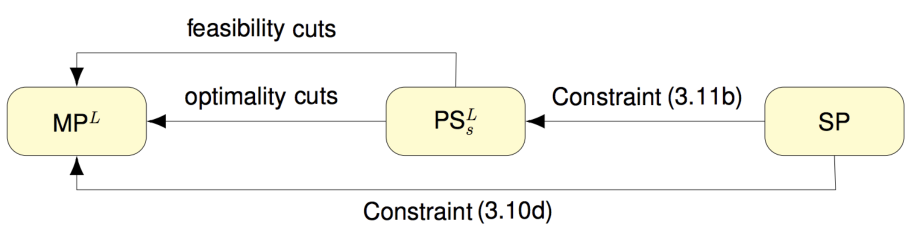

Figure 3.3: Flow chart for the delayed cut generation algorithm with scenario decomposition

Algorithm 1: Decomposition algorithm for the multivariate risk-constrained two-stage optimization

Given L initial scalarization vectors $(\{\tilde{\mathbf{c}}_{(i)}\}_1^L \subset \mathcal{C})$ and a small tolerance parameter $\epsilon > 0$, set $converge = false$ and $u \leftarrow \infty$;

while $converge = false$ **do**

- Solve (MP^L) and get the optimal solution $\bar{\mathbf{v}}$ and objective value $F(\mathrm{MP}^L)$;
- **if** (MP^L) *is infeasible* **then**
  - $converge = true$. The original risk-averse problem is infeasible;
- **else**
  - **if** $(1-\epsilon)u \leq F(\mathrm{MP}^L) \leq (1+\epsilon)u$ **then**
    - Set $converge = true$. The ϵ-optimal solution is $\bar{\mathbf{v}}$;
  - **else**
    - **for** *each scenario* $s \in S$ **do**
      - Solve the primal subproblem (PS_s^L) and get the optimal objective value $\bar{Q}_s^L$;
      - **if** (PS_s^L) *is infeasible* **then**
        - Obtain an extreme ray of the dual feasible region associated with (PS_s^L), given by $(\boldsymbol{\beta}_s, \boldsymbol{\gamma}_s)$, and add the corresponding feasibility cut to $\mathcal{F}$;
      - **else**
        - Obtain an optimal extreme point of the dual of (PS_s^L), given by $(\boldsymbol{\beta}_s, \boldsymbol{\gamma}_s)$, and if $\bar{\theta}_s < Q_s^L$, then add the corresponding optimality cut to $\mathcal{O}$;
    - **if** (PS_s^L) *is feasible for all* $s \in S$ **then**
      - Solve the separation problem (SP);
      - Obtain the optimal solution $\mathbf{c}^*$ (a vertex solution as defined in Proposition 5)
    - and optimal objective value $v(\mathbf{c}^*)$;
    - **if** $v(\mathbf{c}^*) \leq 0$ **then**
      - Current solution is feasible, set $u \leftarrow f(\bar{\mathbf{x}}) + \sum_{s \in S} p_s \bar{Q}_s^L$;
    - **else**
      - Set $L \leftarrow L + 1$ and $\tilde{\mathbf{c}}_{(L)} \leftarrow \mathbf{c}^*$;

Proposition 6. *Suppose that the relaxed master problem* (MP^L) *is given by* (3.10)*, the second-stage subproblems* (PS_s^L) *are given by* (3.11)*, the separation problem is given by* (3.9)*, the feasibility and optimality cuts, respectively, are given by* (3.12) *and* (3.13)*, and the first-stage solution vector is given by* $\bar{\mathbf{v}} = (\bar{\mathbf{x}}, \bar{\boldsymbol{\eta}}, \bar{\mathbf{w}}, \bar{\boldsymbol{\theta}})$*. Then, for a sufficiently small value of the tolerance parameter* ϵ*, Algorithm 1 provides either an optimal solution to problem* (3.1) *or a proof of infeasibility in finitely many iterations.*

Proof. First, observe that the optimality cuts (3.13) generated from subproblems (PS_s^L) for $s \in S$ and the subset $\{\tilde{\mathbf{c}}_{(1)}, \ldots, \tilde{\mathbf{c}}_{(L)}\}$ of $\mathcal{C}$ are valid, because subproblems (PS_s^L) are relaxations of the original subproblems (3.8) given by $(\mathrm{PS}_s^{\bar{L}})$, and hence the duals of the relaxed subproblems provide valid lower bounds on the optimal second-stage objective values. Similarly, the feasibility cuts (3.12) obtained from the relaxed subproblems are valid for the original subproblems (3.8). At an intermediate iteration, if there is a violation in the CVaR constraint (3.3b), then the exact solution of the SP guarantees that a corresponding violating scalarization vector is identified and the associated inequalities and variables are added to the RMP. Furthermore, [81] provides a procedure to guarantee that the $\mathbf{c}^*$ vector generated from the SP is a vertex solution as defined in Proposition 5. Next, recall that the number of scalarization vectors of interest, $\bar{L}$, is finite. Hence, in the worst case, in a finite number of iterations, all $\bar{L}$ scalarization vectors will be added, and the second-stage problems become exact. At this point, the convergence of the algorithm follows directly from the convergence of the classical L-shaped method [112]. □

Remark 6 (General Applicability). *The proposed Benders decomposition-based algorithm can handle a particular class of nonlinear performance measures that can be represented using a finite set of linear inequalities. For instance, performance measures that are expressed as the maximum of finitely many linear terms can be linearized by utilizing an additional variable. In Section 3.4, we follow such a linearization approach in the context of humanitarian relief network design to represent the maximum proportion of unsatisfied demand (see* (3.35j) *and* (3.36)*).*

3.2.3 Linear Programming Formulation and Duality Results

In this section, we develop duality formulations and optimality conditions for the important special case of the multivariate CVaR-constrained two-stage model (3.3), where the mapping f is linear, the set $\mathcal{X}$ is a polyhedral set, and the first-stage decisions are continuous. Let $\mathcal{X} = \{\mathbf{x} \in \mathbb{R}^{n_1}_+ \ : \ A\mathbf{x} \leq \mathbf{b}\}$ and $f(\mathbf{x}) = \mathbf{f}^\top \mathbf{x}$ for some matrix $A \in \mathbb{R}^{m_1 \times n_1}$, and the vectors $\mathbf{f} \in \mathbb{R}^{n_1}$ and $\mathbf{b} \in \mathbb{R}^{m_1}$. Then, the model of interest becomes

$$\text{(LinearP)} \quad \min \quad \mathbf{f}^\top \mathbf{x} + \sum_{s \in S} p_s \mathbf{q}_s^\top \mathbf{y}_s \tag{3.14a}$$

$$\text{s.t.} \quad A\mathbf{x} \leq \mathbf{b}, \quad \mathbf{x} \in \mathbb{R}^{n_1}_+, \tag{3.14b}$$

$$(3.3\text{b}),\ (3.5\text{f}) - (3.5\text{g}). \tag{3.14c}$$

For a finite set $\tilde{\mathcal{C}} = \{\tilde{\mathbf{c}}_{(1)}, \ldots, \tilde{\mathbf{c}}_{(L)}\}$ we consider the following LP, referred to as (FiniteP($\tilde{\mathcal{C}}$)):

$$\min\{\mathbf{f}^\top \mathbf{x} + \sum_{s \in S} p_s \mathbf{q}_s^\top \mathbf{y}_s \ : \ (3.5\text{b})-(3.5\text{d}) \text{ (with } \bar{L} \text{ replaced by } L), (3.14\text{b}),\ \boldsymbol{\eta} \in \mathbb{R}^L,\ (3.5\text{f})-(3.5\text{g})\}.$$

The next lemma shows that (FiniteP($\tilde{\mathcal{C}}$)) is equivalent to (LinearP) for suitable choices of $\tilde{\mathcal{C}}$; it is a simple consequence of Proposition 5.

Lemma 1. *Let $\hat{\mathcal{C}} = \{\hat{\mathbf{c}}_{(1)}, \ldots, \hat{\mathbf{c}}_{(\bar{L})}\}$ as defined in Proposition 5, and assume that the finite set $\tilde{\mathcal{C}}$ satisfies $\hat{\mathcal{C}} \subset \tilde{\mathcal{C}} \subset \mathcal{C}$. Then a vector $(\mathbf{x}, \mathbf{y})$ is a feasible (optimal) solution of* (LinearP) *if and only if $(\mathbf{x}, \mathbf{y}, \boldsymbol{\eta}^{(\mathbf{x},\mathbf{y})}, \mathbf{w}^{(\mathbf{x},\mathbf{y})})$ is a feasible (optimal) solution of* (FiniteP($\tilde{\mathcal{C}}$))*, where $\eta_\ell^{(\mathbf{x},\mathbf{y})} = \mathrm{VaR}_\alpha(\tilde{\mathbf{c}}_{(\ell)}^\top \hat{\mathbf{G}}(\mathbf{x}, \mathbf{y}))$ and $w_{i\ell}^{(\mathbf{x},\mathbf{y})} = [\eta_\ell^{(\mathbf{x},\mathbf{y})} - \tilde{\mathbf{c}}_{(\ell)}^\top \hat{\mathbf{g}}_s(\mathbf{x}, \mathbf{y}_s)]_+$.*

The LP formulation (FiniteP($\tilde{\mathcal{C}}$)) does not immediately offer a tractable solution approach, however it leads to a natural delayed cut generation algorithm as discussed in Sections 3.2.1 and 3.2.2. Moreover, (FiniteP($\tilde{\mathcal{C}}$)) provides an important direct way to derive strong duality results and optimality conditions, presented below. Note that these results provide a theoretical foundation for the potential future development of dual-based (column generation-based or Lagrangian-based) solution methods. First, recall the following notation: $\hat{\mathbf{g}}_s(\mathbf{x},\mathbf{y}_s) = (\mathrm{g}_s^1(\mathbf{x},\mathbf{y}_s), \mathrm{g}_s^2(\mathbf{x},\mathbf{y}_s), \ldots, \mathrm{g}_s^d(\mathbf{x},\mathbf{y}_s))^\top,\ s \in S$, and $\mathrm{g}_s^i(\mathbf{x},\mathbf{y}_s) = \bar{\mathbf{g}}_s^i\mathbf{x} + \tilde{\mathbf{g}}_s^i\mathbf{y}_s,\ i \in \{1,\ldots,d\},\ s \in S$. By abuse of notation, we introduce the random matrices $\bar{\mathbf{g}}\ :\ \Omega \to \mathbb{R}^{d\times n_1}$ and $\tilde{\mathbf{g}}\ :\ \Omega \to \mathbb{R}^{d\times n_2}$, where $\bar{\mathbf{g}}^\top(\omega_s) = \bar{\mathbf{g}}_s^\top$ and $\tilde{\mathbf{g}}^\top(\omega_s) = \tilde{\mathbf{g}}_s^\top$ with the ith columns defined by $\bar{\mathbf{g}}_s^i$ and $\tilde{\mathbf{g}}_s^i$, respectively. Denoting the set of all finitely supported finite non-negative measures on the scalarization polyhedron $\mathcal{C}$ by $\mathcal{M}_+^F(\mathcal{C})$, we obtain the following dual problem to (LinearP):

$$\text{(LinearD)} \quad \max \quad -\int_{\mathcal{C}} \mathrm{CVaR}_\alpha(\mathbf{c}^\top\mathbf{Z})\,\mu(\mathrm{d}\mathbf{c}) - \boldsymbol{\lambda}^\top\mathbf{b} + \mathbb{E}[\boldsymbol{\pi}^\top\mathbf{h}(\omega)] \tag{3.15a}$$

$$\text{s.t.} \quad \mathbb{E}(\nu) = \mu, \tag{3.15b}$$

$$\nu(\omega_s) \le \frac{\mu}{1-\alpha}, \quad \forall\ s \in S, \tag{3.15c}$$

$$\mathbb{E}[\boldsymbol{\pi}^\top T(\omega)] - \mathbb{E}\left(\int_{\mathcal{C}} \mathbf{c}^\top\bar{\mathbf{g}}\,\nu(\mathrm{d}\mathbf{c})\right) \le \mathbf{f}^\top + \boldsymbol{\lambda}^\top A, \tag{3.15d}$$

$$\boldsymbol{\pi}_s^\top W_s - \int_{\mathcal{C}} \mathbf{c}^\top\tilde{\mathbf{g}}_s[\nu(\omega_s)](\mathrm{d}\mathbf{c}) \le \mathbf{q}_s^\top, \quad \forall\ s \in S, \tag{3.15e}$$

$$\boldsymbol{\lambda} \in \mathbb{R}_+^{m_1}, \quad \mu \in \mathcal{M}_+^F(\mathcal{C}), \quad \nu : \Omega \to \mathcal{M}_+^F(\mathcal{C}), \quad \boldsymbol{\pi}_s \in \mathbb{R}_+^{m_2},\ s \in S. \tag{3.15f}$$

The proof of the next duality theorem, which is similar to that of Theorem 3 in Noyan and Rudolf [81], is omitted here for the sake of brevity. It mainly relies on Lemma 1 and the following facts for a finitely supported measure $\mu \in \mathcal{M}_+^F(\mathcal{C})$: $\mathrm{support}(\mu) = \{\mathbf{c} \in \mathcal{C}\ :\ \mu(\{\mathbf{c}\}) > 0\}$ and $\int_{\mathcal{C}} u(\mathbf{c})\,\mu(\mathrm{d}\mathbf{c}) = \sum_{\mathbf{c}\in\mathrm{support}(\mu)} u(\mathbf{c})\mu(\{\mathbf{c}\})$ for a function $u : \mathcal{C} \to \mathbb{R}$.

Theorem 1. *The problem* (LinearP) *has a finite optimum value if and only if* (LinearD) *does, in which case the two optimum values coincide. In addition, a feasible solution* $(\mathbf{x},\mathbf{y})$ *of* (LinearP) *and a feasible solution* $(\boldsymbol{\lambda},\mu,\nu,\boldsymbol{\pi})$ *of* (LinearD) *are both optimal for their respective problems if and only if the following complementary slackness conditions hold:*

$$
\begin{array}{lll}
(i) & \text{support}(\mu) \subset \left\{\mathbf{c} \;:\; \text{CVaR}_\alpha(\mathbf{c}^\top\hat{\mathbf{G}}(\mathbf{x},\mathbf{y})) = \text{CVaR}_\alpha(\mathbf{c}^\top\mathbf{Z})\right\}, & \\
(ii) & \text{support}(\nu(\omega_s)) \subset \left\{\mathbf{c} \;:\; \text{VaR}_\alpha(\mathbf{c}^\top\hat{\mathbf{G}}(\mathbf{x},\mathbf{y})) \le \mathbf{c}^\top\hat{\mathbf{g}}_s(\mathbf{x},\mathbf{y}_s)\right\}, & s \in S, \\
(iii) & \text{support}(\frac{\mu}{1-\alpha} - \nu(\omega_s)) \subset \left\{\mathbf{c} \;:\; \text{VaR}_\alpha(\mathbf{c}^\top\hat{\mathbf{G}}(\mathbf{x},\mathbf{y})) \ge \mathbf{c}^\top\hat{\mathbf{g}}_s(\mathbf{x},\mathbf{y}_s)\right\}, & s \in S, \\
(iv) & \boldsymbol{\lambda}^\top(\mathbf{b} - A\mathbf{x}) = 0, & \\
(v) & \boldsymbol{\pi}_s^\top(T_s\mathbf{x} + W_s\mathbf{y}_s - \mathbf{h}_s) = 0, & s \in S, \\
(vi) & (\mathbf{f}^\top + \boldsymbol{\lambda}^\top A - \mathbb{E}[\boldsymbol{\pi}^\top T(\omega)] + \mathbb{E}\left(\int_{\mathcal{C}} \mathbf{c}^\top\bar{\mathbf{g}}\,\nu(\mathrm{d}\mathbf{c})\right))\mathbf{x} = 0, & \\
(vii) & (\mathbf{q}_s^\top - \boldsymbol{\pi}_s^\top W_s + \int_{\mathcal{C}} \mathbf{c}^\top\tilde{\mathbf{g}}_s[\nu(\omega_s)](\mathrm{d}\mathbf{c}))\mathbf{y}_s = 0, & s \in S.
\end{array}
$$

We remark that our dual formulation is analogous to Haar's dual for semi-infinite linear programs (e.g., [15]). The finite representation of the multivariate CVaR relation (appearing in Proposition 5) proves to be useful to obtain strong duality results and optimality conditions directly from linear programming duality without the need for constraint qualifications.

Lagrangian Duality. The duality results established in Theorem 1 have a natural Lagrangian interpretation. In fact, measures on $\mathcal{C}$ are a natural choice to use as Lagrange multipliers, since the CVaR constraints in (3.3b) are indexed by the scalarization set $\mathcal{C}$. Accordingly, we define the Lagrangian function $\mathcal{L} : (\mathcal{X},\mathcal{Y}) \times \mathcal{M}_+^F(\mathcal{C}) \to \mathbb{R}$ as follows:

$$
\mathcal{L}(\mathbf{x},\mathbf{y},\mu) = \mathbf{f}^\top\mathbf{x} + \sum_{s\in S} p_s\mathbf{q}_s^\top\mathbf{y}_s + \int_{\mathcal{C}} \text{CVaR}_\alpha(\mathbf{c}^\top\hat{\mathbf{G}}(\mathbf{x},\mathbf{y}))\,\mu(\mathrm{d}\mathbf{c}) - \int_{\mathcal{C}} \text{CVaR}_\alpha(\mathbf{c}^\top\mathbf{Z}))\,\mu(\mathrm{d}\mathbf{c}). \tag{3.16}
$$

Here, by abuse of notation, $\mathbf{y} \in \mathbb{R}_+^{|S|n_2}$ represents the collection of decision vectors $\mathbf{y}_s \in \mathbb{R}_+^{n_2}$, $s \in S$, and $\mathbf{y} \in \mathcal{Y}$ refers to $\mathbf{y}_s \in \mathcal{Y}(\mathbf{x},\boldsymbol{\xi}(\omega_s))$, $s \in S$. For the ease of exposition, we assume that the feasible sets $\mathcal{X}$ and $\mathcal{Y}$ are compact; this is a very mild assumption due to the finiteness of the first-stage and second-stage objective function values. Then, (LinearP) is equivalent to $\min\limits_{(\mathbf{x},\mathbf{y})\in(\mathcal{X},\mathcal{Y})} \max\limits_{\mu\in\mathcal{M}_+^F(\mathcal{C})} \mathcal{L}(\mathbf{x},\mathbf{y},\mu)$, while the corresponding Lagrangian dual problem is given by

$$
\text{(LagrangianD)} \max_{\mu\in\mathcal{M}_+^F(\mathcal{C})} \min_{(\mathbf{x},\mathbf{y})\in(\mathcal{X},\mathcal{Y})} \mathcal{L}(\mathbf{x},\mathbf{y},\mu). \tag{3.17}
$$

For any feasible solution of (LinearP) and any $\mu \in \mathcal{M}_+^F(\mathcal{C})$, $\mathcal{L}(\mathbf{x}, \mathbf{y}, \mu) \leq \mathbf{f}^\top \mathbf{x} + \sum_{s\in S} p_s \mathbf{q}_s^\top \mathbf{y}_s$. The weak duality immediately follows: $\mathrm{OBF}_P \geq \mathrm{OBF}_D$, where OBF_P and OBF_D, respectively, denote the optimum objective values of (LinearP) and (LinearD). The next theorem, which is a consequence of Theorem 1, provides the strong duality result and optimality conditions.

Theorem 2. *If the primal problem* (LinearP) *has an optimal solution, then the dual problem* (LinearD) *also has an optimal solution, and the optimal objective values coincide. A primal feasible solution* $(\mathbf{x}^*, \mathbf{y}^*)$ *and a dual feasible solution* μ^* *are simultaneously optimal if and only if they satisfy the equations*

$$\int_{\mathcal{C}} \mathrm{CVaR}_\alpha(\mathbf{c}^\top \hat{\mathbf{G}}(\mathbf{x}^*, \mathbf{y}^*))\, \mu^*(\mathrm{d}\mathbf{c}) = \int_{\mathcal{C}} \mathrm{CVaR}_\alpha(\mathbf{c}^\top \mathbf{Z})\, \mu^*(\mathrm{d}\mathbf{c}), \tag{3.18}$$

$$\mathcal{L}(\mathbf{x}^*, \mathbf{y}^*, \mu^*) = \min_{(\mathbf{x},\mathbf{y})\in(\mathcal{X},\mathcal{Y})} \mathcal{L}(\mathbf{x}, \mathbf{y}, \mu^*). \tag{3.19}$$

Proof. Let us first show the following claim: *If the equations* (3.18)-(3.19) *hold for some primal feasible* $(\mathbf{x}^*, \mathbf{y}^*)$ *and dual feasible* μ^**, then these solutions are simultaneously optimal with coinciding objective values.* According to (3.18) and (3.19), we have

$$\begin{aligned} \mathrm{OBF}_P \leq \mathbf{f}^\top \mathbf{x}^* + \sum_{s\in S} p_s \mathbf{q}_s^\top \mathbf{y}_s^* = \mathcal{L}(\mathbf{x}^*, \mathbf{y}^*, \mu^*) &= \min_{(\mathbf{x},\mathbf{y})\in(\mathcal{X},\mathcal{Y})} \mathcal{L}(\mathbf{x}, \mathbf{y}, \boldsymbol{\mu}^*) \\ &\leq \max_{\mu\in\mathcal{M}_+^F(\mathcal{C})} \min_{(\mathbf{x},\mathbf{y})\in(\mathcal{X},\mathcal{Y})} \mathcal{L}(\mathbf{x}, \mathbf{y}, \mu) = \mathrm{OBF}_D\,. \end{aligned}$$

On the other hand, the weak duality implies that $\mathrm{OBF}_P \geq \mathrm{OBF}_D$, which proves the claim.

We next prove the claim that if $(\mathbf{x}^*, \mathbf{y}^*)$ *is primal optimal and* μ^* *is dual optimal, then they satisfy the equations* (3.18)-(3.19)*. To this end, we first show that for any given primal optimal solution* $(\mathbf{x}^*, \mathbf{y}^*)$ *there exists a corresponding dual feasible solution* $\bar{\mu}$ *such that* (3.18) *and* (3.19) *are satisfied for the choice* $\mu^* = \bar{\mu}$*.* Let us consider an optimal solution $(\mathbf{x}^*, \mathbf{y}^*)$ of (LinearP). By Theorem 1 there exists an optimal solution $(\bar{\boldsymbol{\lambda}}, \bar{\mu}, \bar{\nu}, \bar{\boldsymbol{\pi}})$ of (LinearD) with the same optimal objective value. According to the complementary slackness condition (i), the equality $\mathrm{CVaR}_\alpha(\mathbf{c}^\top\hat{\mathbf{G}}(\mathbf{x}^*,\mathbf{y}^*)) = \mathrm{CVaR}_\alpha(\mathbf{c}^\top\mathbf{Z})$ holds on the support of the measure $\bar{\mu}$, which implies (3.18). *To show that* (3.19) *also holds, we next prove that, substituting* $\mu^* = \bar{\mu}$*,* $\mathcal{L}(\mathbf{x},\mathbf{y},\bar{\mu}) \geq \mathcal{L}(\mathbf{x}^*,\mathbf{y}^*,\bar{\mu})$ *is valid for any primal feasible solution* $(\mathbf{x},\mathbf{y})$ *(and equality holds for* $(\mathbf{x}^*,\mathbf{y}^*)$*).* For ease of exposition, let $U(\mathbf{x},\mathbf{y},\bar{\mu}) = \int_{\mathcal{C}} \mathrm{CVaR}_\alpha(\mathbf{c}^\top\hat{\mathbf{G}}(\mathbf{x},\mathbf{y}))\,\bar{\mu}(\mathrm{d}\mathbf{c})$ and $U(\mathbf{Z},\bar{\mu}) = \int_{\mathcal{C}} \mathrm{CVaR}_\alpha(\mathbf{c}^\top\mathbf{Z}))\,\bar{\mu}(\mathrm{d}\mathbf{c})$. Then, using the definition of the Lagrangian function (3.16) and the dual feasibility conditions (3.15d)-(3.15e) we obtain the following relations:

$$\begin{aligned}
&\mathcal{L}(\mathbf{x},\mathbf{y},\bar{\mu}) + U(\mathbf{Z},\bar{\mu}) = \mathbf{f}^\top\mathbf{x} + \sum_{s\in S} p_s\mathbf{q}_s^\top\mathbf{y}_s + U(\mathbf{x},\mathbf{y},\bar{\mu}) \geq -\bar{\boldsymbol{\lambda}}^\top A\mathbf{x} + \mathbb{E}[\bar{\boldsymbol{\pi}}^\top T(\omega)\mathbf{x}]\\
&- \mathbb{E}\left(\int_{\mathcal{C}} \mathbf{c}^\top\bar{\mathbf{g}}\mathbf{x}\,\bar{\nu}(\mathrm{d}\mathbf{c})\right) + \sum_{s\in S} p_s\left(\bar{\boldsymbol{\pi}}_s^\top W_s\mathbf{y}_s - \int_{\mathcal{C}} \mathbf{c}^\top\tilde{\mathbf{g}}_s\mathbf{y}_s[\bar{\nu}(\omega_s)](\mathrm{d}\mathbf{c})\right) + U(\mathbf{x},\mathbf{y},\bar{\mu})\\
&= -\bar{\boldsymbol{\lambda}}^\top A\mathbf{x} + \sum_{s\in S} p_s\bar{\boldsymbol{\pi}}_s^\top(T_s\mathbf{x} + W_s\mathbf{y}_s) - \mathbb{E}\left(\int_{\mathcal{C}} \mathbf{c}^\top\bar{\mathbf{g}}\mathbf{x}\,\bar{\nu}(\mathrm{d}\mathbf{c})\right)\\
&- \sum_{s\in S} p_s \int_{\mathcal{C}} \mathbf{c}^\top\tilde{\mathbf{g}}_s\mathbf{y}_s[\bar{\nu}(\omega_s)](\mathrm{d}\mathbf{c}) + U(\mathbf{x},\mathbf{y},\bar{\mu}).
\end{aligned} \tag{3.20}$$

Now observe that the complementary slackness condition (iv) implies $-\bar{\boldsymbol{\lambda}}^\top A\mathbf{x} \geq -\bar{\boldsymbol{\lambda}}^\top\mathbf{b} = -\bar{\boldsymbol{\lambda}}^\top A\mathbf{x}^*$, as $A\mathbf{x} \leq \mathbf{b}$ holds for all $\mathbf{x} \in X$, and the Lagrange multiplier $\bar{\boldsymbol{\lambda}}$ is nonnegative. Similarly, as $T_s\mathbf{x} + W_s\mathbf{y}_s \geq \mathbf{h}_s$ holds for all $\mathbf{y}_s \in \mathcal{Y}(\mathbf{x},\boldsymbol{\xi}(\omega_s))$, $s \in S$, and the Lagrange multiplier $\bar{\boldsymbol{\pi}}_s$ is nonnegative for all $s \in S$, the complementary slackness condition (v) implies that $\bar{\boldsymbol{\pi}}_s^\top(T_s\mathbf{x} + W_s\mathbf{y}_s) \geq \bar{\boldsymbol{\pi}}_s^\top\mathbf{h}_s = \bar{\boldsymbol{\pi}}^\top(T_s\mathbf{x}^* + W_s\mathbf{y}_s^*)$. Therefore, the following relation is valid for any primal feasible solution $(\mathbf{x},\mathbf{y})$:

$$-\bar{\boldsymbol{\lambda}}^\top A\mathbf{x} + \sum_{s\in S} p_s\bar{\boldsymbol{\pi}}_s^\top(T_s\mathbf{x} + W_s\mathbf{y}_s) \geq -\bar{\boldsymbol{\lambda}}^\top A\mathbf{x}^* + \sum_{s\in S} p_s\bar{\boldsymbol{\pi}}_s^\top(T_s\mathbf{x}^* + W_s\mathbf{y}_s^*). \tag{3.21}$$

Observing that $U(\mathbf{Z},\bar{\mu})$ is independent of the decision vectors $\mathbf{x}$ and $\mathbf{y}$, by (3.20) and (3.21), it is sufficient to show that the following inequality is valid for any primal feasible solution $(\mathbf{x},\mathbf{y})$ (and equality holds for $(\mathbf{x}^*,\mathbf{y}^*)$):

$$-\mathbb{E}\left(\int_{\mathcal{C}} \mathbf{c}^\top\bar{\mathbf{g}}\mathbf{x}\,\bar{\nu}(\mathrm{d}\mathbf{c})\right) - \sum_{s\in S} p_s \int_{\mathcal{C}} \mathbf{c}^\top\tilde{\mathbf{g}}_s\mathbf{y}_s[\bar{\nu}(\omega_s)](\mathrm{d}\mathbf{c}) + U(\mathbf{x},\mathbf{y},\bar{\mu})$$

$$\geq -\mathbb{E}\left(\int_{\mathcal{C}} \mathbf{c}^\top \bar{\mathbf{g}} \mathbf{x}^* \, \bar{\nu}(\mathrm{d}\mathbf{c})\right) - \sum_{s\in S} p_s \int_{\mathcal{C}} \mathbf{c}^\top \tilde{\mathbf{g}}_s \mathbf{y}_s^* [\bar{\nu}(\omega_s)](\mathrm{d}\mathbf{c}) + U(\mathbf{x}^*, \mathbf{y}^*, \bar{\mu}).$$

Accordingly, we prove that the following chain of inequalities holds for any $(\mathbf{x}, \mathbf{y}) \in (\mathcal{X}, \mathcal{Y})$, and that all inequalities hold with equality for $(\mathbf{x}, \mathbf{y}) = (\mathbf{x}^*, \mathbf{y}^*)$.

$$-\mathbb{E}\left(\int_{\mathcal{C}} \mathbf{c}^\top \bar{\mathbf{g}} \mathbf{x} \, \bar{\nu}(\mathrm{d}\mathbf{c})\right) - \sum_{s\in S} p_s \int_{\mathcal{C}} \mathbf{c}^\top \tilde{\mathbf{g}}_s \mathbf{y}_s [\bar{\nu}(\omega_s)](\mathrm{d}\mathbf{c}) + U(\mathbf{x}, \mathbf{y}, \bar{\mu}) \tag{3.22}$$

$$= -\sum_{s\in S} p_s \int_{\mathcal{C}} \mathbf{c}^\top (\bar{\mathbf{g}}_s \mathbf{x} + \tilde{\mathbf{g}}_s \mathbf{y}_s)[\bar{\nu}(\omega_s)](\mathrm{d}\mathbf{c})$$
$$+ \int_{\mathcal{C}} \left(\mathrm{VaR}_\alpha(\mathbf{c}^\top \hat{\mathbf{G}}(\mathbf{x}, \mathbf{y})) + \frac{1}{1-\alpha} \sum_{s\in S} p_s [\mathbf{c}^\top \hat{\mathbf{g}}_s(\mathbf{x}, \mathbf{y}_s) - \mathrm{VaR}_\alpha(\mathbf{c}^\top \hat{\mathbf{G}}(\mathbf{x}, \mathbf{y}))]_+\right) \bar{\mu}(\mathrm{d}(\mathbf{c})) \tag{3.23}$$

$$= -\sum_{s\in S} p_s \int_{\mathcal{C}} \left(\mathbf{c}^\top \hat{\mathbf{g}}_s(\mathbf{x}, \mathbf{y}_s) - \mathrm{VaR}_\alpha(\mathbf{c}^\top \hat{\mathbf{G}}(\mathbf{x}, \mathbf{y}))\right) [\bar{\nu}(\omega_s)](\mathrm{d}\mathbf{c})$$
$$+ \frac{1}{1-\alpha} \sum_{s\in S} p_s \int_{\mathcal{C}} [\mathbf{c}^\top \hat{\mathbf{g}}_s(\mathbf{x}, \mathbf{y}_s) - \mathrm{VaR}_\alpha(\mathbf{c}^\top \hat{\mathbf{G}}(\mathbf{x}, \mathbf{y}))]_+ \, \bar{\mu}(\mathrm{d}(\mathbf{c})) \tag{3.24}$$

$$\geq \sum_{s\in S} p_s \int_{\mathcal{C}} \Big([\mathbf{c}^\top \hat{\mathbf{g}}_s(\mathbf{x}, \mathbf{y}_s) - \mathrm{VaR}_\alpha(\mathbf{c}^\top \hat{\mathbf{G}}(\mathbf{x}, \mathbf{y}))]_+$$
$$-(\mathbf{c}^\top \hat{\mathbf{g}}_s(\mathbf{x}, \mathbf{y}_s) - \mathrm{VaR}_\alpha(\mathbf{c}^\top \hat{\mathbf{G}}(\mathbf{x}, \mathbf{y})))\Big) [\bar{\nu}(\omega_s)](\mathrm{d}\mathbf{c}) \tag{3.25}$$

$$\geq 0. \tag{3.26}$$

We note that an alternative definition of CVaR for a random variable V is given by $\mathrm{CVaR}_\alpha(V) = \mathrm{VaR}_\alpha(V) + \frac{1}{1-\alpha}\mathbb{E}([V - \mathrm{VaR}_\alpha(V)]_+)$, where $[z]_+ = \max(z, 0)$ [93]. Substituting this definition into the third term and expanding the expected value in the first term (3.22) becomes (3.23). Using the dual feasibility condition (3.15b) to replace $\bar{\mu}$ in the second term of (3.23) provides (3.24). Then, (3.25) follows from (3.15c); equality for $(\mathbf{x}, \mathbf{y}) = (\mathbf{x}^*, \mathbf{y}^*)$ is ensured by the complementary slackness condition (iii). Finally, (3.26) is a consequence of the trivial inequality $[a]_+ \geq a$ and the non-negativity of the random measure ν. Equality is again ensured for $(\mathbf{x}, \mathbf{y}) = (\mathbf{x}^*, \mathbf{y}^*)$, since the complementary slackness condition (ii) implies that $\mathbf{c}^\top \hat{\mathbf{g}}_s(\mathbf{x}, \mathbf{y}_s) - \mathrm{VaR}_\alpha(\mathbf{c}^\top \hat{\mathbf{G}}(\mathbf{x}, \mathbf{y})) \geq 0$ holds on the support of $\bar{\nu}(\omega_s)$ for all $s \in S$.

Finally, let us consider a primal optimal solution $(\mathbf{x}^*, \mathbf{y}^*)$ and a dual optimal solution μ^*.

We just showed above that there exists some dual feasible $\bar{\mu}$ such that $U(\mathbf{x}^*, \mathbf{y}^*, \bar{\mu}) = U(\mathbf{Z}, \bar{\mu})$ and $\mathcal{L}(\mathbf{x}^*, \mathbf{y}^*, \bar{\mu}) = \min_{(\mathbf{x},\mathbf{y})\in(\mathcal{X},\mathcal{Y})} \mathcal{L}(\mathbf{x}, \mathbf{y}, \bar{\mu}) = \mathbf{f}^\top \mathbf{x}^* + \sum_{s\in S} p_s \mathbf{q}_s^\top \mathbf{y}_s^*$. As the set of feasible primal solutions is compact and the Lagrangian function $\mathcal{L}$ is continuous, there also exists some $(\hat{\mathbf{x}}, \hat{\mathbf{y}})$ such that $\mathcal{L}(\hat{\mathbf{x}}, \hat{\mathbf{y}}, \mu^*) = \min_{(\mathbf{x},\mathbf{y})\in(\mathcal{X},\mathcal{Y})} \mathcal{L}(\mathbf{x}, \mathbf{y}, \mu^*)$. Then, by the optimality of μ^*, we have $\mathbf{f}^\top \mathbf{x}^* + \sum_{s\in S} p_s \mathbf{q}_s^\top \mathbf{y}_s^* = \min_{(\mathbf{x},\mathbf{y})\in(\mathcal{X},\mathcal{Y})} \mathcal{L}(\mathbf{x}, \mathbf{y}, \bar{\mu}) \le \min_{(\mathbf{x},\mathbf{y})\in(\mathcal{X},\mathcal{Y})} \mathcal{L}(\mathbf{x}, \mathbf{y}, \mu^*) = \mathcal{L}(\hat{\mathbf{x}}, \hat{\mathbf{y}}, \mu^*)$, and consequently, the following relation holds:

$$
\begin{aligned}
U(\mathbf{x}^*, \mathbf{y}^*, \mu^*) - U(\mathbf{Z}, \mu^*) &= \mathcal{L}(\mathbf{x}^*, \mathbf{y}^*, \mu^*) - \mathbf{f}^\top \mathbf{x}^* - \sum_{s\in S} p_s \mathbf{q}_s^\top \mathbf{y}_s^* \\
&\ge \mathcal{L}(\mathbf{x}^*, \mathbf{y}^*, \mu^*) - \mathcal{L}(\hat{\mathbf{x}}, \hat{\mathbf{y}}, \mu^*) = \mathcal{L}(\mathbf{x}^*, \mathbf{y}^*, \mu^*) - \min_{(\mathbf{x},\mathbf{y})\in(\mathcal{X},\mathcal{Y})} \mathcal{L}(\mathbf{x}, \mathbf{y}, \mu^*) \ge 0.
\end{aligned}
$$

On the other hand, by the primal feasibility of $(\mathbf{x}^*, \mathbf{y}^*)$ and the non-negativity of μ^*, $U(\mathbf{x}^*, \mathbf{y}^*, \mu^*) - U(\mathbf{Z}, \mu^*) \le 0$ holds, which immediately implies (3.18) and (3.19).

□

3.3 Two-Stage Optimization with a Multivariate Stochastic Ordering

The solution algorithms proposed in Section 3.2 can also be applied to the two-stage stochastic programs with the multivariate preference constraints based on increasing convex order (ICO). The counterpart of ICO in the opposite convention, where larger values of random variables are preferred, is referred to as the SSD. Similar to the CVaR case, the univariate ICO relation can be extended to the multivariate case by considering a family of linear scalarization functions; more specifically, a random vector $\mathbf{X}$ is said to dominate another random vector $\mathbf{Z}$ with respect to the ICO and the scalarization set $\mathcal{C}$ if

$$
\mathbb{E}\left(\left[\mathbf{c}^\top \mathbf{X} - \eta\right]_+\right) \le \mathbb{E}\left(\left[\mathbf{c}^\top \mathbf{Z} - \eta\right]_+\right), \quad \forall\, \eta \in \mathbb{R},\ \mathbf{c} \in \mathcal{C}. \tag{3.27}
$$

The risk-averse two-stage optimization model of Dentcheva and Wolfhagen [34] utilizes this multivariate stochastic preference relation to compare the decision-based random outcome vector $\hat{\mathbf{G}}(\mathbf{x}, \mathbf{y})$ with the benchmark outcome vector $\mathbf{Z}$. In particular, the two-stage stochastic programming model with the multivariate ICO constraints takes the form of (3.3), where the multivariate CVaR relation (3.3b) is replaced by (3.27) with $\mathbf{X} = \hat{\mathbf{G}}(\mathbf{x}, \mathbf{y})$. Assuming again that $\mathbf{Z}$ has finitely many realizations $\mathbf{z}_1, \ldots, \mathbf{z}_{|\tilde{S}|}$ with associated probabilities $\tilde{p}_i,\ i \in \tilde{S}$, the inequality (3.27) can be equivalently stated as

$$\mathbb{E}\left(\left[\mathbf{c}^\top \mathbf{X} - \mathbf{c}^\top \mathbf{z}_t\right]_+\right) \leq \mathbb{E}\left(\left[\mathbf{c}^\top \mathbf{Z} - \mathbf{c}^\top \mathbf{z}_t\right]_+\right), \quad \forall\, t \in \tilde{S},\ \mathbf{c} \in \mathcal{C}. \tag{3.28}$$

Relation (3.28) directly follows from (3.27) due to the well-known result [29] that for finite probability spaces the univariate SSD relation remains valid if the corresponding expected shortfall constraints (in our convention, constraints (3.27) featuring the expected excess values for a particular $\mathbf{c}$) only required for the realizations of the benchmark random variable (instead of all $\eta \in \mathbb{R}$).

For finite probability spaces, Homem-de-Mello and Mehrotra [52] show that it is sufficient to consider a finite number of scalarization vectors $\{\hat{\mathbf{c}}_{(1)}, \ldots, \hat{\mathbf{c}}_{(\bar{L})}\}$ in relation (3.28), specifically the projections of the extreme points of polyhedra $P_t,\ t \in \tilde{S}$, defined for each $t \in \tilde{S}$ as $P_t = \{w_i \geq \mathbf{c}^\top \mathbf{z}_i - \mathbf{c}^\top \mathbf{z}_t,\ i \in \tilde{S},\ \mathbf{c} \in \mathcal{C},\ \mathbf{w} \in \mathbb{R}_+^{|\tilde{S}|}\}$. For an alternative finite representation of (3.28), which is based on the vertices of the polyhedron P defined in Proposition 5, we refer to [81]. The fact that these finite representations of (3.28) are only characterized by the given benchmark vector allows us to develop computationally tractable DEFs of the optimization models featuring the multivariate ICO relation as a benchmarking constraint. In particular, for the two-stage multivariate ICO-constrained problem we can obtain the following DEF with finitely many constraints:

$$\min \quad f(\mathbf{x}) + \sum_{s \in S} p_s \mathbf{q}_s^\top \mathbf{y}_s \tag{3.29a}$$

$$\text{s.t.} \quad \sum_{s \in S} p_s w_{slt} \leq \sum_{i \in \tilde{S}} \tilde{p}_i \left[\hat{\mathbf{c}}_{(l)}^\top \mathbf{z}_i - \hat{\mathbf{c}}_{(l)}^\top \mathbf{z}_t\right]_+, \quad \forall\, t \in \tilde{S},\ l = 1, \ldots, \bar{L}, \tag{3.29b}$$

$$w_{slt} \geq \hat{\mathbf{c}}_{(l)}^\top \hat{\mathbf{g}}_s(\mathbf{x}, \mathbf{y}_s) - \hat{\mathbf{c}}_{(l)}^\top \mathbf{z}_t, \quad \forall\, s \in S,\ t \in \tilde{S},\ l = 1, \ldots, \bar{L}, \tag{3.29c}$$

$$w_{slt} \geq 0, \quad \forall\, s \in S,\ t \in \tilde{S},\ l = 1, \ldots, \bar{L}, \tag{3.29d}$$

$$\mathbf{x} \in \mathcal{X}, \quad (3.5f) - (3.5g). \tag{3.29e}$$

Although the number of inequalities (3.29b)–(3.29c) is finite, it is possibly exponential. Hence, as in the CVaR-constrained optimization, the ICO constraints can be added as needed using a delayed cut generation approach. As in Section 3.2.1, we start with solving the DEF with a subset of constraints (3.29b)–(3.29c) for $l = 1, \ldots, L$. For a given solution $\bar{\mathbf{v}} := (\bar{\mathbf{x}}, \bar{\mathbf{w}}, \bar{\mathbf{y}})$ to this relaxed problem, we check if there is a violation in (3.28) (with $\mathbf{X} = \hat{\mathbf{G}}(\bar{\mathbf{x}}, \bar{\mathbf{y}})$). Note that constraints (3.28) are defined for each scalarization vector and for each realization of $\mathbf{Z}$, as such possible violations in constraints (3.28) should be checked for all $\mathbf{z}_t,\ t \in \tilde{S}$. In line with these discussions, given $(\bar{\mathbf{x}}, \bar{\mathbf{w}}, \bar{\mathbf{y}})$, the separation problem corresponding to the t^{th} realization of $\mathbf{Z}$ becomes

$$\min_{\mathbf{c} \in \mathcal{C}} \quad \mathbb{E}\left(\left[\mathbf{c}^\top \mathbf{Z} - \mathbf{c}^\top \mathbf{z}_t \right]_+ \right) - \mathbb{E}\left(\left[\mathbf{c}^\top \hat{\mathbf{G}}(\bar{\mathbf{x}}, \bar{\mathbf{y}}) - \mathbf{c}^\top \mathbf{z}_t \right]_+ \right). \tag{3.30}$$

This cut generation problem can be solved using the adaptations of the MIP formulations provided in [52], [61], and [82] for the opposite convention that larger outcomes are preferred.

As in the CVaR case, the structure of the ICO constraints can be exploited to further decompose the formulation (3.29) over scenarios. Observe that constraints (3.29b) and (3.5f) couple the scenarios together. However, if we handle the original first-stage variables $\mathbf{x}$ and the auxiliary $\mathbf{w}$ variables in the first stage, then the second-stage problems decompose over scenarios once the first-stage variables are fixed. Therefore, we can apply Algorithm 1 with the updated definitions of the problems (MP^L), (PS$_s^L$), and (SP), the feasibility and optimality cuts, and the first-stage solution vector $\mathbf{v}$, which we describe next.

The RMP at an intermediate iteration, where a subset of the scalarization vectors of cardinality L is generated so far, is formulated as

$$(\mathrm{MP}^L) \ \min \quad f(\mathbf{x}) + \sum_{s \in S} p_s \theta_s \tag{3.31a}$$

$$\text{s.t.} \quad (\mathbf{x}, \mathbf{w}, \boldsymbol{\theta}) \in \mathcal{O}, \tag{3.31b}$$

$$(\mathbf{x}, \mathbf{w}) \in \mathcal{F}, \tag{3.31c}$$

$$\mathbf{x} \in \mathcal{X}, \quad (3.29b), (3.29d) \text{ (with } \bar{L} \text{ is replaced by } L). \tag{3.31d}$$

Given a first-stage solution vector $\bar{\mathbf{v}} = (\bar{\mathbf{x}}, \bar{\mathbf{w}}, \bar{\boldsymbol{\theta}})$ of RMP and the L scalarization vectors, the second-stage subproblem under $s \in S$ takes the form of

$$(\text{PS}_s^L) \quad \bar{Q}_s^L := \min \quad \mathbf{q}_s^\top \mathbf{y} \tag{3.32a}$$

$$\text{s.t.} \quad -\tilde{\mathbf{c}}_{(l)}^\top \tilde{\mathbf{g}}_s \mathbf{y} \geq \tilde{\mathbf{c}}_{(l)}^\top \bar{\mathbf{g}}_s \bar{\mathbf{x}} - \tilde{\mathbf{c}}_{(l)}^\top \mathbf{z}_t - \bar{w}_{slt}, \quad \forall\, t \in \tilde{S},\ l = 1, \ldots, L, \tag{3.32b}$$

$$(3.11c) - (3.11d). \tag{3.32c}$$

Here, constraint (3.32b) is equivalent to constraint (3.29c) with partial substitution of the values of the first-stage variables $\bar{\mathbf{x}}$ and $\bar{\mathbf{w}}$. Denoting the dual vectors associated with constraints (3.32b) and (3.11c) by $\beta_{stl}, t \in \tilde{S}, l = 1, \ldots, L$, and $\boldsymbol{\gamma}_s,\ s \in S$, the feasibility and optimality cuts are given by

$$0 \geq \boldsymbol{\gamma}_s(\mathbf{h}_s - T_s \mathbf{x}) + \sum_{t \in \tilde{S}} \sum_{l=1}^{L} \beta_{stl} \left(\tilde{\mathbf{c}}_{(l)}^\top \bar{\mathbf{g}}_s \mathbf{x} - \tilde{\mathbf{c}}_{(l)}^\top \mathbf{z}_t - w_{slt} \right) \text{ and} \tag{3.33}$$

$$\theta_s \geq \boldsymbol{\gamma}_s(\mathbf{h}_s - T_s \mathbf{x}) + \sum_{t \in \tilde{S}} \sum_{l=1}^{L} \beta_{stl} \left(\tilde{\mathbf{c}}_{(l)}^\top \bar{\mathbf{g}}_s \mathbf{x} - \tilde{\mathbf{c}}_{(l)}^\top \mathbf{z}_t - w_{slt} \right). \tag{3.34}$$

Proposition 7. *Suppose that the relaxed master problem* (MP^L) *is given by* (3.31)*, the second-stage subproblems (PS$_s^L$) are given by* (3.32)*, the separation problem is given by* (3.30)*, the feasibility and optimality cuts, respectively, are given by* (3.33) *and* (3.34)*, and the first-stage solution vector is given by* $\bar{\mathbf{v}} = (\bar{\mathbf{x}}, \bar{\mathbf{w}}, \bar{\boldsymbol{\theta}})$*. Then, for a sufficiently small value of the tolerance parameter* ϵ*, Algorithm 1 provides either an optimal solution to problem* (3.29) *or a proof of infeasibility in finitely many iterations.*

Proof. The proof is similar to that of Proposition 6. However, in the SSD case, we refer to [82] for a procedure to guarantee that the $\mathbf{c}^*$ vector generated from the SP is a vertex solution as defined in Proposition 5.

□

Finally, we highlight four major advantages of our proposed solution framework over the existing methods [34]. First, our method reformulates the risk-averse model of interest as a risk-neutral two-stage stochastic program without using the Lagrangian relaxation of the multivariate ICO constraints, and consequently, manages to solve the problem within a single Benders decomposition framework. It maintains a single master problem to solve successively redefined (based on the iteratively enlarged subset of the scalarization vectors) risk-neutral relaxations of the problem. Recall that using a subset of $\mathcal{C}$ provides a relaxation of the problem. These risk-neutral formulations are constructed by using the decomposable structure of the ICO constraints over scenarios. Dentcheva and Wolfhagen [34] also construct successive relaxations by considering a subset of $\mathcal{C}$; however, they obtain approximate risk-neutral relaxations by using the Lagrangian relaxation of the multivariate ICO constraints. In particular, at each iteration of their algorithms, a risk-neutral approximation with an updated Lagrangian objective function is solved (to calculate the dual function given a set of Lagrangian multipliers) using decomposition methods such as Benders decomposition. The need for solving a separate risk-neutral two-stage model at each iteration, in addition to the well-known computational challenges in solving the non-differentiable Lagrangian dual problem, could impose significant computational difficulties. Second, the Lagrangian-based existing algorithms are not exact even for the finite probability spaces; they are shown to finitely converge to an ε-feasible (with respect to the ICO constraint) ε-optimal solution both for continuous and finite probability spaces. In contrast, our finitely-convergent algorithm provides an exact (optimal and feasible) solution for finite probability spaces. Third, the existing algorithms do not provide valid upper bounds on the optimal objective value at intermediate stages, because the Lagrangian dual at an iteration is constructed using the information on a subset of $\mathcal{C}$. In contrast, Algorithm 1 provides valid upper bounds (in line

20) at intermediate stages. Fourth, our method applies to problems that contain discrete variables in the first stage. In contrast, the Lagrangian-based methods cannot guarantee the integer feasibility (see the primal solution recovery step of the algorithms in [34]). In summary, we contribute to the literature by providing a new computationally tractable and exact solution algorithm for the multivariate ICO-constrained two-stage models with integer variables in the first stage under finite probability spaces.

3.4 A Stochastic Optimization Model for Pre-disaster Relief Network Design

An effective and sound pre-disaster relief network design calls for modeling the risk associated with the high level of uncertainty inherent in rarely occurring disaster events (see, e.g., [80]) and considering multiple and possibly conflicting performance criteria such as responsiveness and equity (see, e.g., [49, 56, 83, 113]). Moreover, a two-stage decision making framework is beneficial in this context, since the pre-positioning design decisions must be made before a disaster strikes, whereas the relief distribution decisions should be made in the post-disaster stage. Motivated by the significance of developing such effective pre-disaster plans (see, e.g., [8, 102]), we apply our proposed approach to a stochastic pre-disaster relief network design problem. In this section, we first briefly review the relevant literature, and then describe the problem setting and present the corresponding mathematical programming formulations.

There is a rich and growing literature on developing stochastic programming models for humanitarian logistics (e.g., [18, 65]). The majority of the existing studies focus on pre-disaster relief network design and develop risk-neutral two-stage stochastic programs (e.g., [8, 37, 91, 102]). However, this extensive literature includes only a few studies (e.g., [39, 53, 80, 92]) that provide risk-averse stochastic models, and lacks models with multivariate risk constraints. Here, we mainly discuss the studies that are most closely related to ours. Rawls and Turnquist [91] develop a risk-neutral two-stage stochastic programming model; in the first-stage it determines the cost-wise optimal decisions on locations and capacities of the response facilities, and the inventory levels of relief supplies under uncertainty in demand, damage levels of pre-stocked supplies, and transportation network conditions. Their second-stage problem is formulated as a classical network flow model, which involves detailed distribution decisions representing the flow of relief supplies on each arc. Noyan [80] obtains a risk-averse version of this model by incorporating CVaR as the risk measure on the total cost in addition to its expectation. There also exist chance-constrained variants [53, 92]. A recent study [39] proposes a more elaborate risk-averse extension, which features a mean-risk objective function on the random total cost with CVaR being the risk measure (as in [80]), and enforces a joint chance constraint on the feasibility of the second-stage problem (as in [53]). In contrast to the other variants, the model of Elçi and Noyan [39] relies on an alternative formulation of the second-stage problem, which focuses on assigning the demand points to the facilities instead of determining the detailed arc-flow decisions. In our study, we follow this practically meaningful assignment-based modeling approach for the second-stage problem, and develop a new risk-averse variant of the widely-cited model of Rawls and Turnquist [91].

Problem Description and Mathematical Formulations. We extend the model of Rawls and Turnquist [92] by incorporating the multivariate CVaR relation into the first-stage to evaluate the decisions based on additional multiple stochastic criteria. In particular, we consider the following additional stochastic measures: the maximum proportion of unsatisfied demand and a responsiveness measure based on the total delivery amount-based average travel time. With the term responsiveness we refer to the expeditiousness of the response. A large variety of criteria have been employed in humanitarian logistics and they can be grouped in three main categories (see, e.g., [49, 56]): *efficiency* (related to cost), *efficacy or effectiveness* (related to providing quick and sufficient distribution) and *equity* (related to providing a fair service). According to this classification, our model addresses the issues about efficiency and efficacy of the relief operations using a weighted-sum based objective, which aggregates the costs of opening facilities, demand shortages, and purchasing and shipping the relief supplies. In addition, it addresses the issues about responsiveness and equity (in terms of supply allocation) via the multivariate CVaR constraints. The CVaR-based relation is preferred, since it is a natural relaxation of its SSD-counterpart. Multivariate SSD constraints are typically very demanding and they can be excessively hard to satisfy in practice.

We model the relief distribution system using a network, where I and J denote the sets of nodes representing the demand and candidate facility locations, respectively; we assume without loss of generality that $J \subseteq I$. The set H denotes the multiple facility types, a facility of type h has a given capacity level K_h. Setting up a facility of type h at node j has a fixed cost F_{jh} and a unit variable acquisition cost a_j. We consider a single commodity (as in, e.g., [53]), since the relief items can be supplied as bundles. The demand values, travel times, demand shortage penalty costs, and damage levels of supplies are assumed to be random and their realizations are represented by a finite set of scenarios S with probabilities $p_s,\ s \in S$. We next introduce notation for the realizations of these random parameters under scenario s: demand at node i is d_i^s, travel time from node j to node i is ν_{ji}^s, unit shortage penalty cost is π^s, unit shipment cost from node j to node i is c_{ji}^s, proportion of undamaged supplies at node j is Γ_j^s. Accordingly, the vector of the random parameter realizations under scenario $s \in S$ is given by $\boldsymbol{\xi}(\omega_s) = (\mathbf{d}^s, \boldsymbol{\nu}^s, \pi^s, \mathbf{c}^s, \boldsymbol{\Gamma}^s)$. In contrast to [91], each demand location can be served only by a facility node, which satisfies a responsiveness requirement. In particular, we consider a common upper bound τ on travel times to construct the scenario-dependent coverage sets: $N_i^s = \{j \in J \mid \nu_{ji}^s \leq \tau\}$ and $M_j^s = \{i \in I \mid \nu_{ji}^s \leq \tau\}$, respectively, denote the set of facility nodes that can cover demand node i and the set of demand nodes that can be served by facility node j under scenario s. Enforcing a common threshold serves the objective of providing equitable service in terms of response times.

In our two-stage decision making framework, the first-stage decisions are represented using the following notation: $x_{jh} = 1$ if a facility of size category $h \in H$ is located at node $j \in J$, and $x_{jh} = 0$ otherwise; R_j denotes the amount of supplies pre-located at node $j \in J$. The notation for the second-stage decisions under scenario $s \in S$ is as follows: y_{ji}^s denotes the amount of supplies delivered to demand node $i \in I$ from node $j \in J$; u_i^s designates the amount of the unsatisfied demand at node $i \in I$; $u_{\max}^s$ is the maximum proportion of unsatisfied demand (maximum is calculated over the demand nodes), i.e., $u^s_{\max} = \max_{i\in I} u^s_i/d^s_i$. We next present our risk-averse two-stage stochastic programming model:

$$\min \quad \sum_{j\in J}\sum_{h\in H} F_{jl}x_{jl} + \sum_{j\in J} a_j R_j + \sum_{s\in S} p_s(\sum_{j\in J}\sum_{i\in M^s_j\setminus\{j\}} c^s_{ji}y^s_{ji} + \sum_{i\in I}\pi^s u^s_i) \tag{3.35a}$$

$$\text{s.t.} \quad \text{CVaR}_\alpha(\mathbf{c}^\top \hat{\mathbf{G}}(\mathbf{x},\mathbf{R},\mathbf{y},\mathbf{u},u_{\max}) \le \text{CVaR}_\alpha(\mathbf{c}^\top \mathbf{Z}), \quad \forall\, \mathbf{c}\in\mathcal{C}, \tag{3.35b}$$

$$\sum_{h\in H} x_{jh} \le 1, \quad \forall\, j\in J, \tag{3.35c}$$

$$R_j \le \sum_{h\in H} K_l x_{jh}, \quad \forall\, j\in J, \tag{3.35d}$$

$$x_{jh}\in\{0,1\}, \quad \forall\, j\in J,\ h\in H, \quad R_j\ge 0, \quad \forall\, j\in J, \tag{3.35e}$$

$$u^s_i \ge d^s_i - \sum_{j\in N^s_i} y^s_{ji}, \quad \forall\, i\in I\setminus J, \tag{3.35f}$$

$$u^s_j \ge d^s_j + \sum_{i\in M^s_j\setminus\{j\}} y^s_{ji} - (\Gamma^s_j R_j + \sum_{i\in N^s_j\setminus\{j\}} y^s_{ij}), \quad \forall\, j\in J, \tag{3.35g}$$

$$\sum_{i\in M^s_j\setminus\{j\}} y^s_{ji} \le \Gamma^s_j R_j, \quad \forall\, j\in J, \tag{3.35h}$$

$$\sum_{i\in N^s_j\setminus\{j\}} y^s_{ij} \le (1-\sum_{h\in H} x_{jh})d^s_j, \quad \forall\, j\in J, \tag{3.35i}$$

$$u^s_{\max} \ge \frac{u^s_i}{d^s_i}, \quad \forall\, i\in I, \tag{3.35j}$$

$$y^s_{ji}\ge 0 \quad \forall\, j\in J,\ i\in M^s_j\setminus\{j\}, \tag{3.35k}$$

$$u^s_i \ge 0 \quad \forall\, i\in I, \quad u^s_{\max}\ge 0. \tag{3.35l}$$

The corresponding *risk-neutral first-stage problem* (3.1) is defined by the objective function (3.35a) and constraints (3.35c)-(3.35e), where the underlying second-stage problem (3.2) is defined by the constraints (3.35f)-(3.35l). We refer to this second-stage problem as the *original* one. The objective, defined by (3.35a), is to minimize the expected total cost of opening facilities, purchasing and shipping the relief supplies, and demand shortages. As discussed in Section 3.2, the constraint (3.35b) ensures that the decision-based random outcome vector $\hat{\mathbf{G}}(\mathbf{x}, \mathbf{R}, \mathbf{y}, \mathbf{u}, u_{\max})$ is preferable to the benchmark outcome $\mathbf{Z}$ according to the multivariate polyhedral CVaR relation. Here $\hat{\mathbf{G}}(\mathbf{x}, \mathbf{R}, \mathbf{y}(\omega_s), \mathbf{u}(\omega_s), u_{\max}(\omega_s), \boldsymbol{\xi}(\omega_s)) = \hat{\mathbf{g}}_s(\mathbf{x}, \mathbf{R}, \mathbf{y}^s, \mathbf{u}^s, u^s_{\max})$ is a two-dimensional vector, the components of which are given by

$$\mathrm{g}^1_s(\mathbf{x}, \mathbf{R}, \mathbf{y}^s, \mathbf{u}^s, u^s_{\max}) = u^s_{\max}, \tag{3.36}$$

$$\mathrm{g}^2_s(\mathbf{x}, \mathbf{R}, \mathbf{y}^s, \mathbf{u}^s, u^s_{\max}) = \frac{\sum\limits_{i\in I} \mathrm{ATS}^s_i}{\sum\limits_{i\in I} \max\limits_{j\in N^s_i\setminus\{i\}} \nu^s_{ji}}, \text{where } \mathrm{ATS}^s_i = \sum_{j\in N^s_i\setminus\{i\}} \frac{y^s_{ji}\nu^s_{ji}}{d^s_i}. \tag{3.37}$$

The maximum proportion of unsatisfied demand (3.36) is crucial for equity in terms of supply allocation, whereas the second measure (3.37), the total delivery amount-based average travel time score (ATS), is related to responsiveness. The delivery-based unit average travel time to demand node i could alternatively be calculated by replacing d^s_i with the corresponding exact delivery amount (i.e., y^s_{ji}) in (3.37). However, using this alternative measure would result in nonlinear and nonconvex constraints in our formulations, and consequently, it would not be possible to guarantee an optimal solution. On the other hand, our modeling approach aims to keep the exact delivery amounts closer to the total demand by penalizing the unsatisfied demand in the objective function, and enforcing constraints on the maximum proportion of the unsatisfied demand. In fact, enforcing stricter benchmarking requirements in terms of demand satisfaction generally results in smaller differences between the total delivery amounts and the total demands. Moreover, (3.37) relies on a scaling with an upper bound $(\max\limits_{j\in N^s_i\setminus\{i\}} \nu^s_{ji})$ of $\mathrm{ATS}_i,\ i \in I$, as shown below:

$$\mathrm{ATS}_i = \sum_{j\in N^s_i\setminus\{i\}} \frac{y^s_{ji}\nu^s_{ji}}{d_i} \leq \frac{\sum\limits_{j\in N^s_i\setminus\{i\}} y^s_{ji}\nu^s_{ji}}{\sum\limits_{j\in N^s_i\setminus\{i\}} y^s_{ji}} \leq \frac{\sum\limits_{j\in N^s_i\setminus\{i\}} y^s_{ji} \max\limits_{k\in N^s_i\setminus\{i\}} \nu^s_{ki}}{\sum\limits_{j\in N^s_i\setminus\{i\}} y^s_{ji}} = \max_{j\in N^s_i\setminus\{i\}} \nu^s_{ji}.$$

The first inequality follows from (3.35i) and the second one is obtained by replacing the ν^s_{ji} parameters with the largest possible travel time in the corresponding coverage set N^s_i. Observe that $\max\limits_{j\in N^s_i\setminus\{i\}} \nu^s_{ji}$ is a stronger upper bound on ATS_i than $\sum_{j\in N^s_i\setminus\{i\}} \nu^s_{ij}$ for all $i \in I$. This scaling ensures that both measures take values between 0 and 1 and prevents biased solutions due to the differences in the magnitude of the outcomes. We also note that considering a set of scalarization vectors in (3.35b) is essential to deal with the ambiguities and inconsistencies in the weight vectors. Eliciting the relative weights of even a single decision maker/expert can be challenging (for a related review, see, e.g., [68]), and is further exacerbated if multiple experts are involved as in the case of humanitarian logistics applications.

We next elaborate on the remaining constraints of the proposed model. Constraints (3.35c) ensure that at most one facility is opened at any node $j \in J$. Constraints (3.35d) guarantee that the inventory levels at open facilities do not exceed the corresponding capacity limits and there is a facility located at node j if its inventory level is positive. Constraints (3.35f)-(3.35i) correspond to a restricted version of the classical network flow formulation, which provides a more structured flow considering the coverage issues. More specifically, these constraints enforce that the nodes without any pre-stocked inventory receive service directly from the open facilities that are sufficiently close (according to the upper bound on travel times). Constraints (3.35h) ensure that the amount delivered from a facility node does not exceed its available inventory level and outgoing flow is not possible if there is no facility located. Constraints (3.35i) assure that there is no delivery to a node if a facility is located at that node, and the amount of delivery is limited by its demand, otherwise. Constraints (3.35j) calculate the maximum proportion of demand shortage.

We develop computationally effective solution methods for the model (3.35) by applying our cut generation-based algorithms presented in Section 3.2. Observing that the underlying original second-stage problem is always feasible, it is interesting to note that our scenario decomposition-based algorithm may need to generate the Benders feasibility cuts (3.12). This is the case since the algorithm solves iteratively modified versions of the original second-stage problem with the constraints of type (3.11b), which may lead to infeasibility in a second-stage subproblem.

3.5 Computational Study

In the computational analysis, we apply the two-stage model (3.35) and the proposed solution algorithms to a case study concerning the threat of hurricanes. In the first part, we highlight the benefits of the proposed risk-averse model in the context of humanitarian relief network design, and demonstrate the impact of the essential modeling features - the benchmark vector and scalarization set - on the solutions. In the second part, we illustrate the computational effectiveness of the two types of algorithms proposed in Section 3.2.

As highlighted in earlier studies [61, 82], which consider both multivariate CVaR- and SSD-constrained problems, CVaR-constrained models prove to be significantly more challenging computationally than their SSD-constrained counterparts. The additional challenge in CVaR case stems from the fact that the CVaR calculation in the separation problem requires identifying the corresponding α-quantile. Due to this combinatorial structure of CVaR, additional binary variables are introduced to the separation problem (in the order of scenarios), which significantly increases computational complexity (see, e.g., [61]). Additionally, SSD constraints can be represented by a continuum of CVaR constraints, and hence CVaR-constrained problems provide natural relaxations to their SSD-constrained counterparts. In line with these discussions, we restrict our attention to the CVaR-constrained problem (3.35).

Before proceeding to the numerical results, we describe how the test problem instances are generated. We used the data sets of Elçi and Noyan [39], which are based on a case study concerning the threat of hurricanes in the Southeastern part of the United States [91]; in this case study, the region is represented by a network with 30 nodes and 55 arcs, and only a single set of hurricane scenarios with a cardinality of 51 is considered. In order to generate instances with a larger number of scenarios, Hong et al. [53] propose a scenario generation method which takes into account the dependency structures inherent in disaster relief networks. Their approach randomly identifies the link degradation and node damage levels depending on the proximity to the center of the disaster, and generates the realizations of the random input parameters accordingly. Elçi and Noyan [39] follow this elaborate scenario generation approach, but make the necessary modifications according to their assignment-based formulation of the second-stage problem. More specifically, the authors ignore 55 links of the original network and introduce the links according to the coverage-based responsiveness requirement. In our computational study, we consider the same network structure to distribute water (in units of 1000 gallons), where each node can only be covered by the facilities within the four hours of travel time, i.e., the parameter τ is specified as 4 hours. We refer the reader to [53] and [39] for further details on the values of the relevant parameters, such as cost, facility capacity limits, base values (estimates without the effect of random variations/deviations) of the travel times and water demand, and the sampling probability distribution information, etc. Finally, we note that all scenarios are assumed to be equally likely.

We next discuss the additional issues that arise due to the different features of our new risk-averse model. The stochastic benchmarking constraints of interest additionally require us to specify a confidence level α associated with CVaR, a benchmark random outcome vector, and a scalarization set. It is more straightforward to specify the α parameter, which naturally takes a large probability value such as 0.90 or 0.95 according to the risk-aversion level of the decision makers; we set $\alpha = 0.9$ unless otherwise stated. In general, the benchmark outcome vector can be obtained using two alternative approaches, either in a constructive or a direct manner. The constructive approach requires identifying a reasonable and feasible benchmark first-stage solution and calculating the performance measures of interest associated with the corresponding optimal second-stage decisions. Particularly, a benchmark first-stage solution can be determined based on the existing practices of a relief organization, e.g. Federal Emergency Management Agency (FEMA). We present an application of this approach in Section 3.5.2, which focuses on the computational performance of the proposed solution algorithms. In contrast, in the direct approach, the benchmark random outcome vector does not necessarily correspond to a feasible first-stage solution. The probability distribution of such a benchmark vector can be directly constructed by specifying target performance values at particular probability levels. Due to its flexibility to specify a wide range of benchmark vectors, we use this direct approach in Section 3.5.1 while analyzing the impact of using different benchmarks and scalarization sets on the optimal solutions. Further details on the construction of the benchmark outcome vector are provided in the corresponding sections. In addition, the scalarization set is assumed to have an ordered preference structure as follows: $\mathcal{C} = \mathcal{C}_\gamma = \{\mathbf{c} \in \mathbb{R}^2_+ : c_1 + c_2 = 1, c_2 \geq \gamma c_1\}$ for some $\gamma \geq 0$, where c_1 and c_2 correspond to the relative importance of the demand satisfaction criterion (3.36) and responsiveness criterion (3.37), respectively. This representation is consistent with the well-known analytical hierarchy process (AHP) method, where decision makers are expected to rate the relative importance of each pair of criteria [101]. In our case, the pairwise relative importance rate of interest is represented by the parameter γ. By changing the value of this parameter, we can easily modify the scalarization set in line with the preferences of the decision makers. Furthermore, the parameter γ can be used to adjust the size of the scalarization set since $\mathcal{C}_{\gamma_1} \subseteq \mathcal{C}_{\gamma_2}$ for any $\gamma_1 \geq \gamma_2$.

3.5.1 Model Analysis

We present a detailed computational study to demonstrate the value of incorporating the multivariate risk constraint (3.35b) into the pre-disaster relief network design problem. In Section 5.1.1, our analysis focuses on benchmarking the solutions produced by the proposed risk-averse model against those provided by its risk-neutral counterpart without the multivariate risk constraint (3.35b). The results demonstrate the flexibility of the proposed modeling approach to provide a wide range of solutions that are inclusively aligned with multiple decision makers having different opinions on the relative importance of each criterion. In addition, we illustrate the impact of the choice of a benchmark on the solutions (Section 5.1.2) and the impact of the choice of a scalarization set on the solutions (Section 5.1.3).

In our risk-averse framework, the scalarization set $\mathcal{C}$ and the benchmark outcome vector $\mathbf{Z}$ are among the key model inputs affecting the optimal solutions under the multivariate CVaR constraint. The set $\mathcal{C}$ controls the inclusiveness of different views on the relative importance of multiple criteria. We consider the scalarization sets of the form $\mathcal{C} = \mathcal{C}_\gamma$ for varying values of γ. Note that $\mathcal{C}_\gamma$ enlarges as γ decreases; $\mathcal{C}_{\gamma\to\infty}$ converges to a set with single element $(0,1)$ and $\mathcal{C}_0$ corresponds to the unit simplex. On the other hand, the vector $\mathbf{Z}$ is used to impose benchmarking (threshold) requirements in terms of equity and responsiveness. For ease of analysis, we construct the benchmark vector using the direct approach mentioned at the beginning of Section 3.5. More specifically, we set $\mathbf{Z} = \mathbf{b} \circ \mathbf{Z}^N = (\mathrm{b}_1 Z_1^N, \mathrm{b}_2 Z_2^N)$, where $\mathbf{Z}^N$ is the random outcome vector associated with the optimal solution obtained by solving the risk-neutral model without the multivariate risk constraint, and $\mathbf{b} \in \mathbb{R}_+^2$ is a vector of (benchmark) control parameters. By varying values of the parameter $\mathbf{b}$, we can adjust the level of restriction enforced on the risk-averse solutions in terms of equity and responsiveness measures compared to the risk-neutral solutions. Note that when $\mathbf{b} = (1,1)$, the proposed model becomes equivalent to the risk-neutral counterpart, and using smaller values of $\mathbf{b}$ corresponds to enforcing a more demanding benchmarking constraint to ensure a better performance than that of the risk-neutral solution. In this part of the computational study, we present results on a base test case with $|S| = 100$ and $\alpha = 0.9$ under varying values of the parameters γ and $\mathbf{b}$. We analyze the solutions obtained under different parameter settings mainly in terms of the expected total cost, the CVaR values of the equity ($\mathbf{g}^1$) and responsiveness ($\mathbf{g}^2$) measures, and the CVaR values of the some scalarized versions of the random vector $\hat{\mathbf{G}}(\mathbf{x}, \mathbf{R}, \mathbf{y}, \mathbf{u}, u_{\max}) = (\mathbf{g}^1, \mathbf{g}^2)^\top$.

3.5.1.1 Value of the risk-averse model

We first show that our proposed modeling approach provides flexibility to balance the trade-off between the total cost, equity (in terms of supply allocation) and responsiveness while taking into account the ambiguity and inconsistency in the relative importance of such conflicting criteria. Through Tables 3.1-3.2, we provide a detailed analysis of the solutions obtained under varying benchmarks for the scalarizations sets $\mathcal{C}_1$ and $\mathcal{C}_0$, respectively. For each solution, we report the expected total cost (TC) and its components: the total facility setup cost (FC), the total acquisition cost (AC), and the expectations of the random total distribution cost (DC) and demand shortage cost (SC). In addition, we report the expectation and CVaR values of the equity and responsiveness measures, the number of opened facilities of each size category, and the total pre-stocked inventory levels. We remind that the columns with header $\mathbf{b} = (1, 1)$ correspond to the risk-neutral case.

We observe that the risk-averse model provides better solutions in terms of equity and/or responsiveness measures according to the univariate CVaR-preferability (i.e., random outcomes with smaller CVaR values are preferred) while compromising from the expected total cost. Our numerical results indicate that the risk-neutral solution performs poorly in terms of the equity while achieving a better performance in terms of the responsiveness and expected total cost. For the risk-neutral solution, the expected value of the maximum proportion of unsatisfied demand ($\mathbb{E}(\mathbf{g}^1)$) is calculated as 0.75 and $\mathbf{g}^1$ takes the worst possible value of 1 with a probability of at least 0.1. From a practitioner's standpoint, such a policy is highly undesirable considering the highly likely large values of $\mathbf{g}^1$. On the other hand, the risk-averse model under $\mathcal{C} = \mathcal{C}_1$ and $\mathbf{b} = (0.75, 1)$ provides a solution that achieves 45% and 10% reductions in the expectation and CVaR of $\mathbf{g}^1$, respectively, at the expense of less than 10% increase in the total expected cost. Furthermore, the CVaR of the responsiveness measure $\mathbf{g}^2$ is reduced by 14% while its expectation is increased by less than 9%.

In addition to its flexibility to adjust the performance in terms of equity and responsiveness, the risk-averse modeling approach also enables decision makers to restrain the compromise in the expected total cost by relaxing the benchmarking requirements for some selected criteria compared to the risk neutral solution, i.e., by setting $\mathrm{b}_i > 1$ for some $i \in [d]$. For example, according to Table 3.1, under $\mathcal{C} = \mathcal{C}_1$ and $\mathbf{b} = (0.5, 2)$ (resp., $\mathbf{b} = (0.5, 1)$), we obtain 21% and 15% (resp., 37% and 45%) improvements in $\mathrm{CVaR}_{0.9}(\mathbf{g}^1)$ and $\mathrm{CVaR}_{0.9}(\mathbf{g}^2)$ compared to the risk-neutral solution, while increasing the expected total cost by 27.23% (resp., 75.20%). Note that such a relaxation of the multivariate benchmarking constraint may also lead to a worse performance compared to the risk neutral solution with respect to the univariate CVaR-preferability of some criteria. For example, the solution under $\mathbf{b} = (0.75, 2)$

and $\mathcal{C} = \mathcal{C}_1$ performs slightly worse than the risk-neutral solution in terms of the $\text{CVaR}_{0.9}(\mathbf{g}^2)$ of the responsiveness measure $\mathbf{g}^2$. Depending on the level of willingness of the decision makers to pay for further improving equity and/or responsiveness, the benchmark requirements can be adjusted.

Another interesting observation is that the risk-averse policies tend to stock more inventory as the multivariate benchmarking constraint becomes more demanding by means of a stricter benchmark (with more likely smaller realizations) or a larger scalarizarion set. As expected, specifying more demanding/strict benchmarks in terms of equity in general increases the pre-stocked inventory levels since smaller values of the equity measure $\mathbf{g}^1$ can mainly be attained by delivering more supplies. Interestingly, enforcing a stricter benchmark in terms of responsiveness also has an increasing effect on the inventory levels. Considering the definition of responsiveness measure given in (3.37), this result seems counter-intuitive at first glance. However, this can be attributed to the scalarization scheme used in the multivariate risk constraint. By enforcing a risk constraint on the weighted sum of the outcome vectors, the combined effect of criteria in the benchmark enables the risk-averse model to avoid extreme values in both criteria, and hence achieves more moderate shortage levels and average travel times. This effect becomes more prominent as the set of scalarization vectors gets larger. It can be seen that the inventory level under $\mathbf{b} = (0.75, 0.75)$ is 21375 and 35752 units for the scalarization sets $\mathcal{C}_1$ and $\mathcal{C}_0$, respectively.

Incorporating a multivariate stochastic benchmarking constraint into the two-stage stochastic optimization problem of interest raises concerns regarding the feasibility. Enforcing an overly demanding risk-averse preference relation may lead to infeasibility, even though the risk-neutral counterpart without a risk constraint always has a feasible solution. In Table 3.1, it can be seen that our model has no feasible solution for $\mathbf{b} = (0.25, \leq 2)$ and $\mathcal{C} = \mathcal{C}_1$. Enlarging the scalarization set also magnifies the effects of the benchmark on the problem feasibility. As the scalarization set gets larger, it becomes more challenging to obtain a solution with a multivariate outcome vector that is preferable to the benchmark with respect to the multivariate CVaR relation based on all the possible opinions represented by the scalarization set. For $\mathcal{C} = \mathcal{C}_0$, the problem becomes infeasible for three additional settings (Table 3.2). Hence, to avoid potential infeasibility issues associated with the multivariate risk constraint, it is essential to set reasonable threshold requirements in terms of criteria of interest and the diversity of opinions. In the following subsections, we further elaborate on the impact of the choice of the benchmark and the scalarization set on the solutions.

Table 3.1: Detailed analysis of optimal solutions under different benchmarks and $\mathcal{C} = \mathcal{C}_1$

b	**(1,1)**	**(1,0.75)**	**(1,0.5)**	**(1,0.25)**	**(0.75,2)**	**(0.75,1)**	**(0.75,0.75)**	**(0.5,2)**	**(0.5,1)**	**(0.5,0.75)**	**(0.25,≤2)**
$\mathbb{E}[\text{TC}]$	33.00	33.00	33.04	33.29	33.02	36.26	38.42	41.98	57.81	68.37	INF
FC	0.49	0.51	0.51	0.49	0.51	1.98	2.39	2.82	4.34	5.30	INF
AC	3.55	3.65	3.64	3.06	3.55	10.01	13.84	19.74	39.46	51.05	INF
$\mathbb{E}[\text{DC}]$	0.69	0.68	0.67	0.51	0.69	1.36	1.60	2.01	2.09	1.83	INF
$\mathbb{E}[\text{SC}]$	28.26	28.15	28.22	29.23	28.27	22.90	20.58	17.41	11.92	10.19	INF
$\mathbb{E}[\mathbf{g}^1]$	0.7454	0.7410	0.7611	0.8500	0.7313	0.4070	0.3643	0.2730	0.1393	0.1257	INF
$\mathbb{E}[\mathbf{g}^2]$	0.0523	0.0447	0.0385	0.0239	0.0525	0.0569	0.0535	0.0644	0.0433	0.0298	INF
$\text{CVaR}_{0.9}(\mathbf{g}^1)$	1.0000	1.0000	1.0000	1.0000	1.0000	0.8965	0.8549	0.7934	0.6316	0.6029	INF
$\text{CVaR}_{0.9}(\mathbf{g}^2)$	0.1676	0.1257	0.0842	0.0419	0.1691	0.1442	0.1249	0.1423	0.0919	0.0751	INF
Small facilities	6	7	7	6	7	5	5	2	1	0	INF
Medium facilities	2	2	2	2	2	10	9	10	7	9	INF
Large facilities	0	0	0	0	0	0	2	3	10	12	INF
Inventory	5480	5632	5615	4721	5477	15450	21375	30484	60919	78820	INF

Cost terms $\mathbb{E}[\text{TC}]$, FC, AC, $\mathbb{E}[\text{DC}]$ and $\mathbb{E}[\text{SC}]$ are given in millions, inventory levels are in units of 1000 gallons. "INF" stands for "Infeasible".

Table 3.2: Detailed analysis of optimal solutions under different benchmarks and $\mathcal{C} = \mathcal{C}_0$

b	**(1,1)**	**(1,0.75)**	**(1,0.5)**	**(1,0.25)**	**(0.75,2)**	**(0.75,1)**	**(0.75,0.75)**	**(0.5,2)**	**(0.5,1)**	**(0.5,0.75)**	**(0.25,≤2)**
$\mathbb{E}[TC]$	33.00	33.00	33.04	33.29	43.85	43.85	43.90	INF	INF	INF	INF
FC	0.49	0.51	0.51	0.49	2.90	2.90	2.92	INF	INF	INF	INF
AC	3.55	3.65	3.64	3.06	23.12	23.12	23.16	INF	INF	INF	INF
$\mathbb{E}[DC]$	0.69	0.68	0.67	0.51	2.30	2.30	2.27	INF	INF	INF	INF
$\mathbb{E}[SC]$	28.26	28.15	28.22	29.23	15.54	15.54	15.56	INF	INF	INF	INF
$\mathbb{E}[\mathbf{g}^1]$	0.7454	0.7410	0.7611	0.8500	0.3628	0.2852	0.3384	INF	INF	INF	INF
$\mathbb{E}[\mathbf{g}^2]$	0.0523	0.0447	0.0385	0.0239	0.0833	0.0832	0.0718	INF	INF	INF	INF
$\text{CVaR}_{0.9}(\mathbf{g}^1)$	1.0000	1.0000	1.0000	1.0000	0.7500	0.7500	0.7500	INF	INF	INF	INF
$\text{CVaR}_{0.9}(\mathbf{g}^2)$	0.1676	0.1257	0.0842	0.0419	0.1683	0.1676	0.1257	INF	INF	INF	INF
Small facilities	6	7	7	6	0	0	1	INF	INF	INF	INF
Medium facilities	2	2	2	2	9	9	9	INF	INF	INF	INF
Large facilities	0	0	0	0	4	4	4	INF	INF	INF	INF
Inventory	5480	5632	5615	4721	35690	35690	35752	INF	INF	INF	INF

Cost terms $\mathbb{E}[\text{TC}]$, FC, AC, $\mathbb{E}[\text{DC}]$ and $\mathbb{E}[\text{SC}]$ are given in millions, inventory levels are in units of 1000 gallons. "INF" stands for "Infeasible".

3.5.1.2 Impact of the benchmark

Here we investigate the implications of the choice of the benchmark outcome vector regarding the solution quality in terms of the total cost, equity and responsiveness. In this spirit, we obtain the risk-neutral solution and the risk-averse solutions under varying values of the benchmark control parameter $\mathbf{b}$ and $\mathcal{C} = \mathcal{C}_1$. To provide intuitive insights, Figure 3.4 presents comparative results on $\text{CVaR}_{0.9}$ of the scalarized equity and responsiveness measures under several scalarization vectors from the unit simplex. More specifically, we use the scalarization vector $\mathbf{c}^* = (c_1^*, c_2^*)^\top$, where $c_2^* = 1 - c_1^*$, for varying values of $c_1^* \in [0, 1]$ to compare the outcomes obtained under different benchmarks according to possible different opinions on the relative importance of criteria. Figure 3.4 provides results for the $\mathbf{c}^*$ vectors that are not in $\mathcal{C}_1$ so that the performance of the risk-averse solutions can also be tested against opinions not considered in the decision making process. Conveniently, the extreme cases $c_1^* = 1$ and $c_1^* = 0$ in this setup convey the univariate CVaR of the equity measure $\mathbf{g}^1$ and the responsiveness measure $\mathbf{g}^2$, respectively. In addition, for the same instances, Figure 3.5 presents the break-down of the corresponding cost components.

For each scalarization vector considered in Figure 3.4, the risk-averse solutions outperform the risk-neutral solution in terms of $\text{CVaR}_{0.9}(c_1^*\mathbf{g}^1 + c_2^*\mathbf{g}^2)$ for every benchmark vector used in the analysis. This implies that the risk-averse solutions perform better than the risk-neutral solution with respect to both equity and responsiveness measures. Even the opinions not considered in the decision making process, i.e., $\mathbf{c}^* \notin \mathcal{C}_1$, favor the risk-averse solutions in Figure 3.4 over the risk-neutral solution in terms of $\text{CVaR}_{0.9}(c_1^*\mathbf{g}^1 + c_2^*\mathbf{g}^2)$. As it can be seen from Figure 3.4, specifying more demanding benchmarking requirements in terms of equity (see the lines where $b_1 < 1$) has evidently a more prominent effect on the CVaR values of interest. This is mainly due to the fact that the risk-neutral solution already performs well in terms of the responsiveness criterion, and this leaves more limited room for further improvements.

As in the CVaR case, the restrictions on the responsiveness measure $\mathbf{g}^2$ has a limited effect on the cost objective. Figure 3.5 shows that the expected total cost values associated with the solutions obtained under more demanding responsiveness requirements, where $\mathbf{b} \in \{(1,1),(1,0.75),(1,0.25)\}$, are similar with the risk-neutral case, while $\text{CVaR}_{0.9}(\mathbf{g}^2)$ is improved by up to 75%. On the other hand, specifying a stricter benchmark in terms of equity by setting $\mathbf{b} = (0.75, 1)$ provides approximately 10% and 14% improvements in the $\text{CVaR}_{0.9}$ of the equity and responsiveness measures, respectively, over the risk neutral solution without any significant increase in the expected total cost (see Table 3.1). However, when $\mathbf{b} = (0.50, 1)$, the corresponding 37% (resp., 45%) improvement in the $\text{CVaR}_{0.9}$ of the equity (resp., responsiveness) measure almost doubles the expected total cost. We observe that the increase in the expected total cost is mainly due to the increase in the total acquisition cost caused by highly elevated inventory levels. It appears that the level of reduction in the expected total shortage cost (due to the higher level of supplies) is not sufficient to compensate for the increase in the total acquisition cost. These results demonstrate the flexibility of the proposed modeling approach to provide a wide range of solutions via varying the benchmark, and hence, highlight the significance of performing a sensitivity analysis regarding the choice of the benchmark outcome vector.

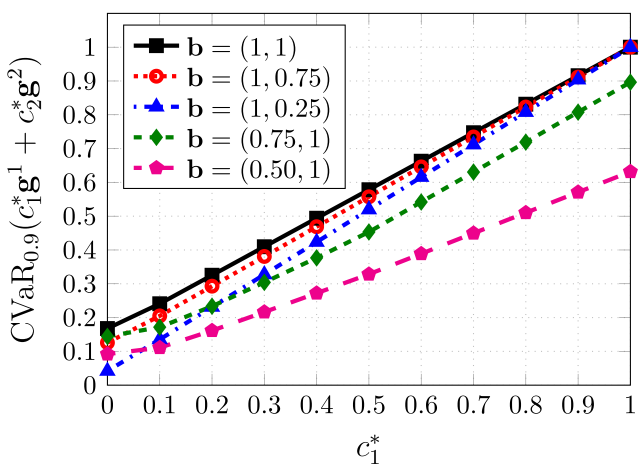

Figure 3.4: The effect of the benchmark vector on the CVaR values under the varying scalarization vector $\mathbf{c}^*$

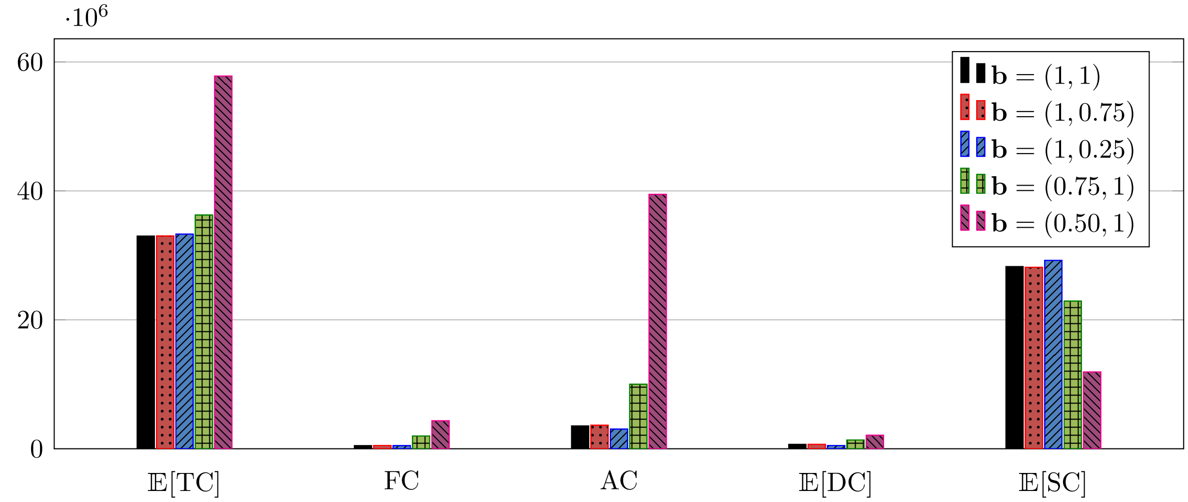

Figure 3.5: The effect of the benchmark vector on the expected total cost and its components

3.5.1.3 Impact of the scalarization set

Next we analyze the effects of the scalarization set on the solution quality. In Table 3.3, we consider several values for the scalarization vector $\mathbf{c}$, which may correspond to a sample representing the ambiguity in the opinion of a single decision maker or the inconsistency in the opinions of different stakeholders such as government, community, non-governmental organizations, engineers, sponsors. With the second interpretation in mind, we compare the solutions obtained under different scalarization sets in terms of the $\mathrm{CVaR}_{0.9}$ of the equity and responsiveness measures scalarized based on the opinion of each stakeholder. The notation $\mathbf{g}_{\mathcal{C}}$ stands for the random outcome vector $(\mathbf{g}^1, \mathbf{g}^2)^\top$ obtained under the scalarization set $\mathcal{C}$. Recall that in the case of $\mathcal{C} = \emptyset$, the risk-averse model becomes equivalent to its risk neutral counterpart. In addition, our construction of the scalarization set $\mathcal{C}_\gamma$ has certain implications on the quality and the sensitivity of the optimal solutions. It can be seen that independent from the value of γ, the vector $(0,1)$ is always included in the set $\mathcal{C}_\gamma$. Thus, the risk-averse solutions obtained by using $\mathcal{C}_\gamma$ for any value of γ are ensured to perform better than the benchmark in terms of $\mathrm{CVaR}_{0.9}$ of the responsiveness measure. Consequently, for the instances under a stricter benchmark in terms of the responsiveness measure, e.g., $\mathrm{b}_2 < 1$, the size of the scalarization set has little effect on the optimal solutions. Accordingly, we restrict our analysis to a benchmark that is more demanding only in terms of equity criterion, i.e., $\mathbf{b} = (0.75, 1)$. The values highlighted in boldface stand for the solutions undesirable for the corresponding stakeholder compared to the benchmark outcome vector. Note that such a situation occurs only when the corresponding $\mathbf{c}$ vector of the stakeholder does not belong to the scalarization set $\mathcal{C}$ featured in the model.

As expected, the number of stakeholders pleased with the performance of the optimal solution in terms of equity and responsiveness increases as a wider set of opinions are represented via a larger set of scalarization set in the model. We observe that all stakeholders are satisfied with the solution obtained under $\mathcal{C} = \mathcal{C}_0$, whereas all but one prefer the benchmark over the risk-neutral solution (when $\mathcal{C} = \emptyset$). The results in Table 3.3 also show that for any stakeholder considered in our experiment, the risk-averse solutions perform at least as good as the risk neutral solution regarding the $\text{CVaR}_{0.9}$ of the scalarized equity and responsiveness measures, although they may fail to surpass the performance of the benchmark for some stakeholders. For example, for $\mathbf{c} = (0.75, 0.25)$, $\text{CVaR}_{0.9}(\mathbf{c}^\top \mathbf{g}_{\mathcal{C}})$ associated with the solution under $\mathcal{C} = \mathcal{C}_1$ is calculated as 0.6748, which is worse than the value of the benchmark ($\text{CVaR}_{0.9}(\mathbf{c}^\top \mathbf{g}_{\mathcal{C}}) = 0.6016$), but still better than the value of the risk-neutral solution ($\text{CVaR}_{0.9}(\mathbf{c}^\top \mathbf{g}_{\mathcal{C}}) = 0.7891$). In addition, using a larger scalarization set does not guarantee that the stakeholders, who are already satisfied with the solution obtained under a smaller scalarization set, will be better off under the new solution, but it ensures that they will remain satisfied compared to the benchmark. For $\mathbf{c} = (0.25, 0.75)$, the corresponding value of $\text{CVaR}_{0.9}(\mathbf{c}^\top \mathbf{g}_{\mathcal{C}})$ increases from 0.2593 to 0.2926 as the scalarization set enlarges from $\mathcal{C}_{0.5}$ to $\mathcal{C}_{0.33}$. This results mainly from the problem setup, since, among two solutions each performing better than the benchmark in terms of the $\text{CVaR}_{0.9}$ of the scalarized equity and responsiveness measures, the model chooses the one with the smaller expected total cost.

Finally, we investigate the variation in the expected total cost with respect to the size of the scalarization set. Figure 3.6 depicts the results on the expected total cost and its components associated with the optimal policies obtained under varing scalarization sets in the increasing order of size. These results indicate that using a larger set of opinions on the relative importance of criteria leads to higher expected total costs. In order to incorporate more diverse opinions into the decision making process, our model compromises from the cost objective; the corresponding difference can be interpreted as the marginal cost of disagreement among stakeholders. The number of stakeholders satisfied with the current decision in Table 3.3 can be increased from one to four for a 9.87% increase in the expected total cost by using the risk-averse model under $\mathcal{C} = \mathcal{C}_1$ instead of the risk-neutral counterpart. In addition, we observe a 32.89% increase in the expected total cost when the unit simplex is specified as the scalarization set in the risk-averse model ($\mathcal{C} = \mathcal{C}_0$). These observations illustrate the effect of the scalarization set on the optimal solutions, and highlight the significance of performing a sensitivity analysis regarding the choice of the scalarization set. On the other hand, the solutions appear to be more robust to the changes in the scalarization set compared to the changes in the benchmark; as mentioned earlier, the expected total cost can be almost doubled depending on the benchmark outcome vector.

Table 3.3: Comparison of the solutions under varying scalarization sets - in terms of the CVaR of interest based on a selected set of scalarization vectors

		$\text{CVaR}_{0.9}(\mathbf{c}^\top \mathbf{g}_{\mathcal{C}})$					
$\mathbf{c}$	$\text{CVaR}_{0.9}(\mathbf{c}^\top Z)$	$\mathcal{C} = \emptyset$	$\mathcal{C} = \mathcal{C}_{10^7}$	$\mathcal{C} = \mathcal{C}_1$	$\mathcal{C} = \mathcal{C}_{0.5}$	$\mathcal{C} = \mathcal{C}_{0.33}$	$\mathcal{C} = \mathcal{C}_0$
(0,1)	0.1676	0.1676	0.1676	0.1442	0.1612	0.1676	0.1676
(0.25,0.75)	0.3049	**0.3674**	**0.3674**	0.2682	0.2593	0.2926	0.2662
(0.33,0.66)	0.3544	**0.4377**	**0.4377**	0.3282	0.3144	0.3424	0.3165
(0.5,0.5)	0.4533	**0.5783**	**0.5783**	0.4533	0.4290	0.4420	0.4179
(0.66,0.33)	0.5522	**0.7189**	**0.7189**	**0.6010**	0.5522	0.5460	0.5285
(0.75,0.25)	0.6016	**0.7891**	**0.7891**	**0.6748**	**0.6191**	0.6016	0.5838
(1,0)	0.7500	**1.0000**	**1.0000**	**0.8965**	**0.8198**	**0.7953**	0.7500

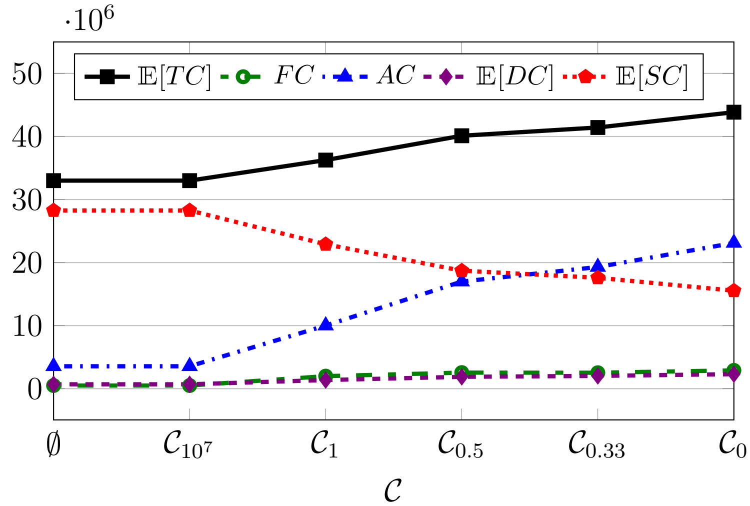

Figure 3.6: The effect of the scalarization set on the expected total cost and its components

3.5.2 Computational Performance of the Solution Methods

We evaluate the computational performance of the proposed solution algorithms in terms of the run times and the number of operations performed at each iteration. We solve the two-stage model (3.35) under varying values of α and $|S|$. On the other hand, we set the scalarization set $\mathcal{C} = \mathcal{C}_1$, and consider a particular type of benchmark random vector obtained by using the constructive approach discussed at the beginning of Section 3.5. In particular, our benchmark construction is inspired by the current practice of FEMA, where the disaster recovery centers in the Southeastern part of the United States are mainly located around Houston and Orlando [42]. In line with this limited information, we specify the benchmark as the random outcome vector associated with the first-stage solution, which locates single large-sized facilities at Houston, Orlando, Baton Rouge, Biloxi, and Atlanta with the maximum possible inventory level (5394 units). The locations in addition to Houston and Orlando are selected based on their proximity to the other demand points with the purpose of guaranteeing a better performance in terms of both criteria of interest. While we focus our numerical study on the effects of α and $|S|$ on the solution time, we note that the computational performance of the algorithms may also be affected by the choice of the scalarization set and the benchmark. For a detailed empirical analysis of the difficulty of the separation problems with respect to different scalarization sets and benchmarks we refer the reader to [61, 68, 81].

All experiments are performed on a single thread of a Windows server with Intel® Xeon® CPU E5-2630 processor at 2.40 GHz and 32 GB of RAM using Java and Cplex 12.6.0. CPLEX is invoked with a time limit of 3600 seconds. MIP gap and feasibility tolerances are set to 10^{-5} and 10^{-9}, respectively; otherwise, we use the default CPLEX settings. The results are obtained for the instances with $\alpha = 0.99, 0.95, 0.90$ and up to 1500 scenarios. For each setting, we report the average over three instances. We implement the cutting plane algorithms using the *lazy constraint callback* function of CPLEX. Considering the fact that using callbacks prevents CPLEX from utilizing dynamic search, we specify the MIP search method for both algorithms as *traditional branch-and-cut.* Following an acceleration technique suggested in [14], we initialize the RMP of the scenario decomposition-based algorithm with optimality and feasibility cuts associated with an initial first-stage solution to benefit from the possible improvement obtained from the automatically created CPLEX cuts. In particular, our initial solution locates two facilities at Mobile and Orlando with the smallest size. Additionally, we employ the *user cut callback* function of CPLEX for adding the Benders cuts at the root node for the fractional solutions. The *usercut callback* is invoked at each iteration at the root node until the improvement in the relative MIP gap is less than 10^{-6} for the last five consecutive iterations.

In Table 3.4, we report several statistics on the performance of the delayed cut generation algorithm for DEF (DCG-DEF) and its scenario decomposition-based counterpart (DCG-SD). Under the "Time" columns, the solution times for only the instances solved within the time limit are reported. For both solution algorithms, we report the total run time and the time spent solving the separation problems under columns "Total/[%gap]", and "Sep.", respectively. If the algorithm terminated at the time limit with a feasible solution for some replications of the instance, then the average of the relative optimality gaps are reported under the column "Total/[%gap]", inside square brackets. The relative optimality gap reported by CPLEX corresponds to the relative gap between the objective function value for the best available integer solution (z^{IP}) and the best LP relaxation bound at the time of termination (z^{LP}), i.e. $100|z^{IP} - z^{LP}|/z^{IP}$. For both algorithms, we report the average number of separation problems solved and the average number of scalarization vectors added under columns "Sep." and "L", respectively. For DCG-SD, we report the average total number of cuts ((3.10d), (3.12), and (3.13)) added to the master problem under column "Cuts". We do not report the number of cuts for DCG-DEF, because in this case, constraints of type (3.5b) and (3.5c) are added when a new scalarization vector is included in the RMP. Hence, the total number of cuts added to the RMP for DCG-DEF is $(|S|+1)L$, where $|S|$ is the number of scenarios and L is the number of scalarization vectors generated.

The results show that DCG-SD outperforms DCG-DEF in terms of the computation times and number of instances that can be solved to optimality. Out of 63 instances, DCG-DEF is able to find an optimal solution for only 35 instances and it terminates without a feasible solution after one hour for seven instances. On the other hand, DCG-SD provides at least a feasible solution for all instances even if it fails to prove optimality for five of them. Considering the instances that can be solved to optimality by both methods, it can be seen that DCG-SD proves optimality in a shorter amount of time than DCG-DEF.

We observe that unlike the single-stage multivariate CVaR-constrained problems in [81], the separation problem is no longer the bottleneck taking over 95% of the solution time. In fact, for DCG-DEF, the time spent on separation is negligible when compared to the overall solution time, with a small number of separation problems solved. On the other hand, a large number of separation problems are solved in DCG-SD, hence it is important to use the strongest formulations available for the separation problem (i.e., those provided in [68] for the equiprobable scenario case) to reduce the time spent on generating the scalarization vectors. Nevertheless, a few instances cannot be solved with DCG-SD within the time limit mainly due to the long solution times of the separation MIPs. However, for these instances, our algorithm is able to find a feasible solution during its course, and hence it provides optimality gaps at termination. To avoid long solution times, one can implement heuristic separation methods. In addition, if ε-feasible solutions are satisfactory, then the violation thresholds can be updated accordingly.

Table 3.4: Performance of the proposed delayed cut generation algorithms ($|I| = |J| = 30$)

		DCG-DEF				DCG-SD				
		Time (s)		Number		Time (s)		Number		
α	# Sce.	Total / [%gap]	Sep.	Sep.	L	Total / [%gap]	Sep.	Cuts	Sep.	L
0.99	200	704.9	0.2	4.3	2.3	46.8	3.1	6101.3	62.0	3.0
	300	1483.5	0.7	5.3	3.0	108.7	5.1	9697.3	78.0	3.3
	400	2622.4	1.2	5.3	3.7	351.5	13.9	18569.7	122.0	5.3
	500	3071.3	0.8	4.0	3.0	363.3	17.7	20805.0	105.3	4.3
	600	[0.2]$^{\dagger **}$	-	-	-	662.1	42.4	27068.7	157.0	5.0
	800	***	-	-	-	987.2	51.7	31994.7	129.7	4.7
	1000	***	-	-	-	1594.2	110.3	43934.0	129.3	5.3
	1500	***	-	-	-	3450.6 [0.4]††	506.0	62436.0	172.0	4.0
0.95	200	408.9	0.7	2.7	0.7	148.6	17.4	9524.7	137.0	4.0
	300	1476.4	1.6	3.3	1.0	277.6	26.1	13280.0	128.7	4.0
	400	3392.7	2.1	3.7	1.7	536.8	45.2	20633.3	162.0	4.0
	500	2753.2 [0.2]†	1.4	2.5	1.5	671.9	150.6	23045.0	161.3	3.3
	600	[0.1]$^{\dagger\dagger *}$	-	-	-	878.0	156.2	26151.3	141.0	3.3
	800	***	-	-	-	1616.4	126.4	33607.7	115.7	3.0
	1000	***	-	-	-	1857.7	274.6	42229.3	140.3	3.3
	1500	***	-	-	-	3390.1 [1.4]††	867.2	52294.0	120.0	2.0
0.90	200	397.4	1.0	2.7	0.7	303.5	36.5	11781.0	180.0	4.3
	300	1128.5	0.6	3.0	0.7	352.5	48.7	13152.7	111.3	3.0
	400	2263.0	1.6	4.0	1.3	623.9	89.3	19482.3	130.3	3.3
	500	1660.2 [0.3]†	2.8	2.0	1.0	1392.1	228.6	25358.0	150.3	3.3
	600	2264.3 [0.1]††	2.2	2.0	1.0	857.5 [7.9]†	229.6	30051.0	153.5	3.0

$\dagger$: Each dagger sign indicates one instance hitting the time limit with an integer feasible solution.
$*$: Each asterisk sign indicates one instance hitting the time limit with no integer feasible solution.

We close this section by noting that the problem sizes we consider are significantly larger (in terms of first- and second-stage variables and constraints and/or the number of scenarios) when compared to the data sets in [81] and [33] for related optimization problems with multivariate stochastic benchmarking constraints. In addition, our first-stage problem involves discrete decisions, which results in the most difficult data sets considered for this problem class to date. In conclusion, our computational experiments showcase the scalability of the proposed scenario decomposition-based method with respect to the number of scenarios.

Chapter 4

MARKOV DECISION PROCESSES UNDER PARAMETER UNCERTAINTY: A CHANCE-CONSTRAINED OPTIMIZATION APPROACH

4.1 Introduction

Markov Decision Processes (MDPs) are effectively used in many applications of sequential decision making in uncertain environments including inventory management, manufacturing, robotics, communication systems, and healthcare [5, 16, 90]. In an MDP model, the decision makers take an action at specified points in time considering the current state of the system with the aim of minimizing their expected loss (resp., maximizing their expected utility), and depending on the action taken, the system transitions to another state. The evolution of the underlying process is mainly characterized by the action costs (resp., rewards) and transition probabilities between the system states, inducing two types of uncertainty. The *internal uncertainty* stems from the probabilistic behaviour of transitions, whereas the *parameter uncertainty* is due to the ambiguities in the parameters representing the costs and transition probabilities. In classical MDPs, these parameters are assumed to be known; they are usually estimated from historical data or learned from previous experiences. However, in practice, it is usually not possible to obtain a single estimate that fully captures the nature of the uncertainties. The actual performance of the system may significantly differ from the anticipated performance of the MDP model due to the inherent variation in the parameters [72].

We consider a humanitarian inventory management problem for relief items required to sustain basic needs of a population affected by a slow-onset disaster, e.g., war, political insurgence, extreme poverty, famine, or drought. Since the progress and impact of a slow-onset disaster generally depend on unpredictable political and/or natural events, the demand for the relief items is highly variable. The supply amounts are also exposed to uncertainty as they mainly rely on voluntary donations. Another critical issue in humanitarian inventory management is the perishability of many relief items such as food and medication. At the beginning of each time period, based on the current inventory level, the decision makers need to determine an additional order quantity to minimize the total inventory holding, stock-out and disposal costs considering the expiration dates and the uncertainty in supply and demand. This problem can be modeled as an MDP, where the current inventory level represents the state of the system, and the uncertainty in supply, demand, and inventory is captured by the transition probabilities between different states [43]. Here, internal uncertainty refers to the probabilistic behaviour of transitions from the current inventory level. Parameter uncertainty, on the other hand, stems from possible fluctuations in the cost terms, demand and supply rates used in the estimation of transition probabilities, and shelf life of perishable relief items. Our focus in this study is on the parameter uncertainty in MDPs motivated by the applications in which the objective function is aggregated over a number of problem instances, e.g., total inventory cost for multiple items and multiple storage facilities with similar characteristics. In such cases, aggregation across multiple instances mitigates the variation due to internal uncertainty, and parameter uncertainty becomes the main source of variation.

Robust optimization is widely used in the literature to incorporate parameter uncertainty into MDPs. In the robust modeling framework, the objective is to optimize the worst-case performance over all possible realizations in a given uncertainty set. This approach is appealing in the sense that it requires no prior information on the distribution of costs or transitions and it gives rise to computationally efficient solution algorithms. However, it often leads to conservative results because the focus is on the worst-case system performance, which may be rarely encountered in practice. In the initial studies on robust MDPs, uncertainty is usually described using a polyhedral set, because it leads to tractable solution algorithms [7, 46, 103, 110, 114]. The uncertainty sets are later extended for more general definitions with the aim of balancing the conservatism of the solutions and the tractability of the solution algorithms. Nilim and El Ghaoui [78] model the uncertainty in transition probabilities using a set of stochastic matrices satisfying rectangularity property, i.e., when there are no correlations between transition probabilities for different states and actions, and devise an efficient dynamic programming algorithm. Similarly, Iyengar [57] studies robust MDPs under transition probability uncertainty with rectangularity assumption and provide robust value and policy iteration algorithms. Sinha and Ghate [108] propose a policy iteration algorithm for robust nonstationary MDPs following the rectangularity assumption. Wiesemann et al. [115] relax the rectangularity assumption in [78] and consider a more general class of uncertainty sets in which the assumption of no correlation between transition probabilities is only made for states, not for actions. Mannor et al. [73] define a tractable subclass of non-rectangular uncertainty sets, namely k-rectangular uncertainty sets, such that the number of possible conditional projections of the uncertainty set is at most k. For relevant studies on continuous-time robust MDPs, we refer the reader to [58, 85, 116].

An alternative modeling approach to balance the conservatism of the solutions is to consider the distributional robustness, where the uncertain parameters are assumed to follow the worst-case distribution from a set of possible distributions described by some general properties such as expectations or moments. Unlike robust MDPs, distributionally robust MDPs incorporate the available incomplete information on the a priori distribution of the uncertain parameters and decrease the degree of conservatism in the optimal solutions. Shapiro and Kleywegt [106] show the relationship between distributionally robust and Bayesian stochastic models by proving that the former is equivalent to the expected value problem with respect to a certain mixture distribution involving a finite number of distributions from the uncertainty set. However, the equivalence requires perfect information on a priori distribution of the transition matrix, which may be not possible in practice. Xu and Mannor [118] address parameter uncertainty in MDPs using a sequence of nested sets of distributions (that is, any set of distributions contains the previous set). The parameters associated with any state are enforced to follow a distribution that belongs to each uncertainty set with at least a specified probability increasing in the size of the set. This improves accuracy of the uncertainty set by adding more flexibility to take several confidence levels into consideration. The authors present a solution framework, which is polynomial time under some mild conditions. Yu and Xu [119] extend this study for a more general uncertainty set, which can be nested or disjoint, and possibly include linear inequalities. Their definition allows for inclusion of means and variance measures in the uncertainty set.

Bayesian approaches to address parameter uncertainty attract more attention in the recent literature. In this approach, unknown parameters are treated as random variables with corresponding probability distributions. Steimle et al. [109] consider a multi-model MDP, where the aim is to find a policy maximizing the weighted sum of expected total rewards over a finite horizon associated with different sets of parameters obtained by different estimation methods. This modeling approach is analogous to an expected value problem considering a finite number of scenarios in stochastic optimization, because each set of parameters can be treated as a scenario in which the corresponding scenario probability is set as the normalized weight value. They show that the problem is NP-hard and provide a mixed-integer programming formulation and a heuristic algorithm. A potential drawback of this modeling framework in the stochastic optimization context is that the expected value model ignores the risk arising from the parameter uncertainty. Considering this issue, Xu and Mannor [117] address reward uncertainty in MDPs with respect to parametric regret. The MDP of interest is a finite horizon, discounted model with finite state and action spaces. They consider two different objectives: minimax regret based on the robust approach and mean-variance trade-off of the regret based on the Bayesian approach. They propose a nonconvex quadratic program for the minimax regret problem and a convex quadratic program for the mean-variance trade-off problem. Delage and Mannor [27] consider reward and transition probability uncertainty separately and propose a chance-constrained model in the form of percentile optimization, which corresponds to the risk measure, value-at-risk. The authors give a formulation for infinite horizon MDPs with finite state and action spaces, and stationary policies. They show that the problem is intractable in the general case but can be efficiently solved when the rewards follow a Gaussian distribution or transition probabilities are modelled using independent Dirichlet priors. Chen and Bowling [19] also take a Bayesian approach to deal with reward and transition probability uncertainty. Instead of making distributional assumptions, they focus on determining objective functions that lead to optimal solutions that are easy to approximate for any distribution. They investigate a family of percentile-based objectives that can represent many classical objective functions such as expectation, value-at-risk, conditional value-at-risk, etc. Then they propose a tractable subclass, namely k-of-N measures, and approximate an optimal policy using a no-regret algorithm used in game theory. Similarly, Adulyasak et al. [3] focus on finding objective functions that are separable over realizations of uncertain parameters under a sampling framework. They propose two objectives: average value maximization and confidence probability maximization. They represent two mixed-integer nonlinear programming formulations to find stationary policies on finite horizon MDPs with finite state and action spaces. The proposed formulations are solved using a Lagrangian dual decomposition algorithm.

In this study, we adopt a Bayesian approach to address MDPs under cost and transition probability uncertainty. Our aim is to obtain a stationary policy that optimize the quantile function value at a certain confidence level α, also known as risk measure value-at-risk (VaR_α), with respect to parameter uncertainty. Quantile-based performance measures are used in many applications in the service industry because of their clear interpretation and correspondence with the service-level requirements [11, 26]. In the humanitarian inventory management problem, the VaR_α objective has a natural interpretation: to find a procurement strategy that minimizes the budget required to cover the expected total discounted inventory costs with at least α probability. Note that it is also possible to employ different risk measures based on the problem context and preferences of the decision makers. In the recent literature, the Conditional Value-at-Risk (CVaR) is widely used as an alternative to the VaR due to its desirable computational and theoretical properties. Different than the VaR, the CVaR focuses on the magnitude of the extreme realizations; it approximately corresponds to the expectation of the worst $100(1-\alpha)\%$ outcomes. The main criticism over the VaR is that it violates the subadditivity axiom of coherency, whereas the CVaR is a coherent risk measure [6]. Subadditivity is mainly associated with diversification, that is, merging two sources of risk does not create additional risk. However, it may be irrelevant in certain settings such as humanitarian inventory management problem, where the aim is to obtain an optimal ordering policy. Some studies in the recent literature argue necessity of the subadditivity axiom and suggest a reasonable relaxation, comonotonic subadditivity, which is a property satisfied by the VaR [51, 59]. Comonotonic subadditivity enforces subadditivity for only comonotone random variables, i.e., random variables with perfectly aligned realizations. In addition, the CVaR is computationally attractive due to its convexity, while the optimization of VaR is often challenging because it is not a convex function in general. Nevertheless, the main challenge in our case is the combinatorial nature of policies independent from the choice of the risk measure, which makes the problem NP-hard even for the expected value objective [109]. Moreover, a reliable estimation of the VaR requires smaller sample sizes compared to the CVaR as it is statistically more robust to imperfect data.

Unlike [27], which also considers the VaR objective under the name percentile criterion, we assume that the distributions of uncertain parameters can be discretized using a sampling method. Along the same lines, Zhang et al. [120] utilize a sampling approach on a finite horizon stochastic MIP problem with dynamic decisions, and enforce joint chance constraints on the service level in each period. They represent the uncertainty using a scenario tree as a time-dependent model and show that branch-and-cut algorithms can be effectively used for the multistage stochastic optimization problems with finite horizon. Different than their study, we consider infinite horizon problems with the VaR_α objective assuming the uncertain parameters are stationary over time. This way, we provide a mathematical model, propose computational enhancements and devise a branch-and-cut algorithm that can be used for infinite horizon MDPs with general distributions on parameters. The proposed modeling approach and solution algorithms are illustrated on an inventory management problem in long-term humanitarian relief operations.

4.2 Problem Formulation and Structural Properties

Consider a discrete time infinite horizon MDP model with finite state space $\mathcal{H}$ and finite action space $\mathcal{A}$. We define $\tilde{c}_i(a)$ as the immediate expected cost of taking action $a \in \mathcal{A}$ in state $i \in \mathcal{H}$ and $\tilde{P}_{ij}(a)$ as the probability of transitioning from state $i \in \mathcal{H}$ to state $j \in \mathcal{H}$ under action $a \in \mathcal{A}$. The future costs are discounted by $\gamma \in [0, 1)$ and the distribution of the initial state is given as $|\mathcal{H}|$-dimensional vector $\mathbf{q}$. A stationary policy $\boldsymbol{\pi} = (\pi_1, \ldots, \pi_{|\mathcal{H}|})$ refers to a sequence of decision rules π_i describing the action strategy for each state $i \in \mathcal{H}$. Assuming that the cost and transition probability parameters are nonnegative, stationary and bounded, the expected total discounted cost of the underlying Markov chain for a given policy $\boldsymbol{\pi}$ and known system parameters $(\tilde{\mathbf{c}}, \tilde{\mathbf{P}})$ can be stated as

$$C(\boldsymbol{\pi}, \tilde{\mathbf{c}}, \tilde{\mathbf{P}}) = \mathbb{E}_{\mathbf{x} \in \mathcal{H}} \left(\sum_{t=0}^{\infty} \gamma^t \tilde{c}_{x_t}(\pi_{x_t}) | x_0 \propto \mathbf{q}, \pi \right),$$

where x_0 and x_t denote the initial state and the state of the system at decision epoch $t > 0$, respectively. Let Π be the set of stationary Markov policies. In an MDP model, the aim is to find a policy $\boldsymbol{\pi} \in \Pi$ minimizing the expected total discounted cost, i.e.,

$$\min_{\boldsymbol{\pi} \in \Pi} \quad C(\boldsymbol{\pi}, \tilde{\mathbf{c}}, \tilde{\mathbf{P}}). \tag{4.1}$$

This problem is known to have a stationary and deterministic optimal solution over the set of all policies, and it can be solved easily using several well-known methods such as the value iteration algorithm, the policy iteration algorithm, and linear programming [90]. Using the Bellman equation [9], problem (4.1) can be alternatively stated as

$$\min_{\boldsymbol{\pi} \in \Pi,\ \mathbf{v} \in \mathbb{R}^{|\mathcal{H}|}} \quad \sum_{i \in \mathcal{H}} q_i v_i \tag{4.2a}$$

$$\text{s.t. } v_i = \tilde{c}_i(\pi_i) + \gamma \sum_{j \in \mathcal{H}} \tilde{P}_{ij}(\pi_i) v_j, \quad i \in \mathcal{H}, \tag{4.2b}$$

where v_i is the expected sum of discounted costs under the selected policy $\boldsymbol{\pi}$ when starting from state $i \in \mathcal{H}$.

In this chapter, we consider a more general setting, where the elements of the cost vector $\tilde{\mathbf{c}}$ and transition probability matrix $\tilde{\mathbf{P}}$ are assumed to be random variables instead of known parameters. We seek a policy $\boldsymbol{\pi} \in \Pi$ minimizing the α-quantile of the expected total discounted cost with respect to parameter uncertainty, which corresponds to the optimal policy of the following problem

$$\min_{y \in \mathbb{R},\ \boldsymbol{\pi} \in \Pi} \quad y \tag{4.3a}$$

$$s.t. \quad \mathbb{P}_{\tilde{\mathbf{c}},\tilde{\mathbf{P}}}\left(C(\boldsymbol{\pi}, \tilde{\mathbf{c}}, \tilde{\mathbf{P}}) \leq y\right) \geq \alpha. \tag{4.3b}$$

The formulation above ensures that the expected total discounted cost for the optimal policy $\boldsymbol{\pi}^*$ is less than or equal to the optimal objective function value y^* with probability at least α under the distributions of $\tilde{\mathbf{c}}$ and $\tilde{\mathbf{P}}$. Note that the optimal value y^* corresponds to the value-at-risk (VaR) of the expected total discounted cost at confidence level α. Such formulations optimizing VaR (α-quantile) are also referred as quantile or percentile optimization in the literature. Delage and Mannor [27] consider the quantile optimization problem (4.3) under cost and transition probability uncertainty separately. The authors assume that the cost parameters follow a Gaussian distribution in the former case, and the transition probabilities in the latter are modeled using Dirichlet priors. Different than their model, we consider the cost and transition probability uncertainty simultaneously without any assumptions on the distributions of uncertain parameters other than the applicability of a sampling method.

Using the nominal MDP model in (4.2), we obtain an alternative chance-constrained formulation for problem (4.3),

$$\min_{y \in \mathbb{R},\ \boldsymbol{\pi} \in \Pi,\ \mathbf{v} \in \mathbb{R}^{|\mathcal{H}|}} \quad y \tag{4.4a}$$

$$s.t. \quad \mathbb{P}_{\tilde{\mathbf{c}},\tilde{\mathbf{P}}}\left(\sum_{i \in \mathcal{H}} q_i v_i^{(\tilde{\mathbf{c}},\tilde{\mathbf{P}})} \leq y\right) \geq \alpha. \tag{4.4b}$$

$$v_i^{(\tilde{\mathbf{c}},\tilde{\mathbf{P}})} = \tilde{c}_i(\pi_i) + \gamma \sum_{j\in\mathcal{H}} \tilde{P}_{ij}(\pi_i) v_j^{(\tilde{\mathbf{c}},\tilde{\mathbf{P}})} \qquad i \in \mathcal{H}. \tag{4.4c}$$

The distribution of $\mathbf{v}^{(\tilde{\mathbf{c}},\tilde{\mathbf{P}})}$ depends on the joint distribution of random parameters $\tilde{\mathbf{c}}$ and $\tilde{\mathbf{P}}$ and the selected policy $\boldsymbol{\pi}$. In general, such problems with chance constraints are highly challenging to solve since they require computation of the joint distribution function, which usually involves numerical integration in multidimensional spaces [24, 25, 44, 45]. On the other hand, using a discrete representation of the distribution function obtained by a sampling method significantly reduces the computational complexity and provides reliable approximations to the chance-constrained problem for a sufficiently large sample size [17, 70]. In stochastic optimization, each sample of parameters is referred to as a *scenario*. Note that this approach yields optimal solutions for multi-model MDPs in which parameter uncertainty can be finitely discretized. For example in medical applications for designing optimal treatment and screening protocols, the system state usually represents patient health status and transition probabilities between states can be computed using multiple tools from the clinical literature, which often produce different parameters [109]. In this context, the parameters computed using each tool can be treated as a scenario.

Another challenge in solving problem (4.4) is the possibility that none of the optimal policies is deterministic. As mentioned before, for unconstrained infinite horizon discounted MDPs with discrete state space, finite action space and known parameters, it is ensured that there always exists a deterministic stationary optimal policy. However, this is not necessarily true for MDPs under parameter uncertainty or additional constraints; all optimal policies of problem (4.4) may be randomized as demonstrated in the following example.

Example 4. *Consider a single state infinite horizon MDP with two actions, a and b, under two scenarios with equal probabilities. Under scenario 1,* $\tilde{c}(a) = 0$ *and* $\tilde{c}(b) = 2$*, and under scenario 2,* $\tilde{c}(a) = 2$ *and* $\tilde{c}(b) = 0$*. In case a stationary deterministic policy is applied, i.e., either action a or b is chosen, the optimal objective function value of the problem with* $\alpha = 0.9$ *is* $2/(1-\gamma)$*. However, if a randomized policy of selecting action a or b with equal* *probabilities is applied, then the optimal objective function value is* $1/(1-\gamma)$*, proving that no stationary deterministic policy is optimal.*

Despite the possibility that all optimal policies are randomized, we limit our attention to deterministic policies motivated by applications where implementation of deterministic policies may be preferred over randomized policies. These include the cases in which the decision makers are prone to making errors [20], and the cases that making randomized decisions raises ethical concerns as in the health care systems [109] and humanitarian relief operations. Note that our results and solution algorithm, can be easily adjusted for the finite horizon MDPs with nonstationary/stationary deterministic optimal policies by introducing a time dimension on the decisions. Here we focus on the infinite horizon case, in which the stationarity assumption is desired for practicality and tractability.

We propose a mixed-integer programming formulation for problem (4.4) assuming the existence of a discrete representation for parameter uncertainty. Considering the previous arguments, we make two important assumptions for obtaining an efficient solution method. In particular, we restrict our attention to the set of stationary deterministic policies, denoted as Π_D, and assume that the joint distribution of $\tilde{\mathbf{c}}$ and $\tilde{\mathbf{P}}$ is represented as a finite set of scenarios $\mathcal{S} = \{1, \ldots, n\}$ with corresponding probabilities $p^1, \ldots, p^n$. Under these assumptions, problem (4.4) can be restated as the following deterministic equivalent formulation,

$$\min \; y \tag{4.5a}$$

$$s.t. \; \sum_{a \in \mathcal{A}} w_{ia} = 1, \quad i \in \mathcal{H}, \tag{4.5b}$$

$$\sum_{s \in \mathcal{S}} z^s p^s \geq \alpha, \tag{4.5c}$$

$$\sum_{i \in \mathcal{H}} q_i v_i^s \leq y + (1 - z^s) M, \quad s \in \mathcal{S}, \tag{4.5d}$$

$$v_i^s \geq \sum_{a \in \mathcal{A}} \tilde{c}_i^s(a) w_{ia} + \gamma \sum_{a \in \mathcal{A}} \sum_{j \in \mathcal{H}} \tilde{P}_{ij}^s(a) v_j^s w_{ia}, \quad i \in \mathcal{H}, \; s \in \mathcal{S}, \tag{4.5e}$$

$$z^s \in \{0, 1\}, \quad s \in \mathcal{S}, \tag{4.5f}$$

$$w_{ia} \in \{0, 1\}, \quad i \in \mathcal{H},\ a \in \mathcal{A}, \tag{4.5g}$$

where z^s is a binary variable equal to 1 if scenario s satisfies the chance constraint (4.4b) and $z^s = 0$ otherwise, w_{ia} is a binary variable taking value 1 if action a is taken in state i under the optimal strategy, and M represents a large number such that constraint (4.5e) is redundant if $w_{ia} = 0$. Note that the variables $\mathbf{w}$ are required for enforcing that a single policy is selected across all scenarios. The objective is to find the α-quantile for the expected total discounted cost of the given MDP with respect to the uncertainty in $\tilde{\mathbf{c}}$ and $\tilde{\mathbf{P}}$. Constraint (4.5b) combined with the domain constraint (4.5g) ensures that at each state exactly one action is taken, which is effective for all stages. Constraints (4.5c) and (4.5f) assure that the chance constraint is satisfied with probability at least α. Constraints (4.5d) and (4.5e) enforce the Bellman optimality conditions as in formulation (4.2), and ensure that the optimal value of y is equal to the maximum realization of the expected total discounted cost over the selected scenarios ($z^s = 1$).

This formulation is nonlinear due to constraints (4.5e), which contain the bilinear term $v_j^s w_{ia}$. The nonlinearity can be eliminated by replacing (4.5e) with the linear inequality

$$v_i^s \geq \tilde{c}_i^s(a) + \gamma \sum_{j \in \mathcal{H}} \tilde{P}_{ij}^s(a) v_j^s - (1 - w_{ia}) M_{is}, \quad i \in \mathcal{H},\ a \in \mathcal{A},\ s \in \mathcal{S}. \tag{4.6}$$

Constraint (4.6) assures that the second requirement in the chance constraint (4.4b) is satisfied for the selected scenarios, i.e., $z^s = 1$, and the state-action pairs (i, a) such that $w_{ia} = 1$. Note that when the integrality constraint on $\mathbf{w}$ is relaxed, formulation (4.5) allows for randomized stationary policies. However, in this case, the nonlinear constraint (4.5e) cannot be replaced by constraint (4.6), since the latter requires the optimal policy to be deterministic. Instead, the bilinear term $v_j^s w_{ia}$ can be aproximated by linear inequalities using McCormick envelopes [74]. Additionally, under the deterministic policy assumption, McCormick envelopes provide an exact linearization of the bilinear term, but it requires the addition of $n|\mathcal{A}||\mathcal{H}|^2$ variables and $4n|\mathcal{A}||\mathcal{H}|^2$ constraints, hence linearization through constraint (4.6) is preferred.

4.3 Computational Enhancements

In this section, we examine the deterministic equivalent formulation (4.5) and propose methods that can potentially improve its computational performance.

4.3.1 Preprocessing

Existing solution algorithms for MDPs usually suffer from the size of state and action spaces, often referred to as the curse of dimensionality. In our case, the computational complexity of problem (4.5) is additionally amplified with the number of scenarios. Hence it is worthwhile to search for methods to reduce the size of the problem beforehand. In what follows, we provide preprocessing procedures to obtain tighter bounds on the optimal value of y and hence to narrow down the solution space before solving the original formulation.

Let y^* be the optimal objective function value of the original formulation (4.5). First, we relax the requirement that the same policy should be selected over all scenarios, and use monotonicity property of the VaR function to find bounds on the value of y^*. Let $\bar{b}$ be a random variable representing the maximum expected total discounted cost of the relaxed MDP model. The realization of $\bar{b}$ under scenario $s \in \mathcal{S}$ can be computed by solving the following linear programming problem

$$\bar{b}^s = \min \sum_{i\in\mathcal{H}} q_i v_i \tag{4.7a}$$

$$s.t. \quad v_i \geq \tilde{c}_i^s(a) + \gamma \sum_{j\in\mathcal{H}} \tilde{P}_{ij}^s(a) v_j, \quad i \in \mathcal{H},\ a \in \mathcal{A}. \tag{4.7b}$$

Then we order the its realizations under each scenario as

$$\bar{b}^{s_1} \leq \bar{b}^{s_2} \leq \cdots \leq \bar{b}^{s_{|\mathcal{S}|}},$$

and let $k_1 \in \{1, \dots, n\}$ be the order of the scenario such that $\sum_{t=1}^{k_1} p^{s_t} \geq \alpha$ and $\sum_{t=1}^{k_1-1} p^{s_t} < \alpha$, and $b_u = \bar{b}^{s_{k_1}}$. Note that b_u corresponds to the VaR of random variable $\bar{b}$ at confidence level α. Hence b_u provides an upper bound on y^*, i.e., $b_u \geq y^*$, since $\bar{b}^s \geq C(\boldsymbol{\pi}, \tilde{\mathbf{c}}^s, \tilde{\mathbf{P}}^s)$ for all $s \in \mathcal{S}$ and $\boldsymbol{\pi} \in \Pi_D$.

Similarly, let $\underline{b}$ be a random variable representing the minimum expected total discounted cost of the MDP with relaxed policy selection requirements, whose realizations under each scenario $s \in \mathcal{S}$ can be computed by solving the linear programs

$$\underline{b}^s = \max \sum_{i \in \mathcal{H}} q_i v_i \tag{4.8a}$$

$$s.t. \quad v_i \leq \tilde{c}_i^s(a) + \gamma \sum_{j \in \mathcal{H}} \tilde{P}_{ij}^s(a) v_j, \quad i \in \mathcal{H}, \; a \in \mathcal{A}. \tag{4.8b}$$

Let b_l correspond to the VaR of random variable $\underline{b}$ at confidence level α. The value of b_l provides a lower bound on y^*, i.e., $b_l \leq y^*$ as $\underline{b}^s \leq C(\boldsymbol{\pi}, \tilde{\mathbf{c}}^s, \tilde{\mathbf{P}}^s)$ for all $s \in \mathcal{S}$ and $\boldsymbol{\pi} \in \Pi_D$ by definition.

Problems (4.7) and (4.8) can be efficiently solved using a policy iteration or value iteration algorithm in polynomial time. Using these bounds, we can conclude that any scenario s, whose lower bound is greater than the upper bound on the quantile value, i.e., $\underline{b}^s > b_u$, cannot satisfy the chance constraint in the optimal policy, hence we can set $z^s = 0$. Similarly, any scenario $s \in \mathcal{S}$ with an upper bound value smaller than the lower bound on the quantile function value, i.e., $\bar{b}^s < b_l$ must satisfy the chance constraint in the optimal solution, implying that the corresponding scenario variable can be set as $z^s = 1$. This result can be used for reducing the size of the scenario space. Furthermore, we can add the constraints

$$b_l \leq y \leq b_u$$

into formulation (4.5). As previously mentioned, validity of these bounds follows from monotonicity property of the VaR function. Additionally, the inequalities

$$y \geq \underline{b}^s z^s, \quad s \in \mathcal{S} \tag{4.9}$$

are valid due to constraints (4.5d) and the fact that $\underline{b}^s \leq C(\boldsymbol{\pi}, \tilde{\mathbf{c}}^s, \tilde{\mathbf{P}}^s)$ for all $s \in \mathcal{S}$, $\boldsymbol{\pi} \in \Pi_D$. Note that Küçükyavuz and Noyan [61] propose similar bounding and scenario elimination ideas and demonstrate their effectiveness in the context of multivariate CVaR-constrained stochastic optimization.

The information obtained by solving problems (4.7) and (4.8) can also be used to obtain tighter values of the big-M terms in constraints (4.5d) and (4.6). Clearly, we can set $M = b_u - b_l$ in (4.5d) because its value is bounded by

$$M \geq \max_{\mathbf{v}:(4.5b)-(4.5g)} \sum_{i \in \mathcal{H}} q_i v_i^s - \min_{y:(4.5b)-(4.5g)} y \quad \text{for } s \in \mathcal{S} \text{ s.t. } z^s = 1.$$

Let $\bar{v}_i^s$ and $\underline{v}_i^s$ be the optimal values of variable v_i in problem (4.7) and (4.8), respectively, for state $i \in \mathcal{H}$ and scenario $s \in \mathcal{S}$. Steimle et al. [109] show that the big-M term in constraint (4.6) can be set as $M_{is} = \bar{v}_i^s - \underline{v}_i^s$ for all $i \in \mathcal{H},\ s \in \mathcal{S}$.

4.3.2 Obtaining a Feasible Solution

Our preliminary computational experiments suggested that the upper bound b_u described in the previous section can be further improved by finding a feasible solution to problem (4.5). Here, we propose a polynomial time heuristic algorithm (Algorithm 2), which benefits from the connection between a substructure of our problem with the robust MDPs to attain a policy that performs reasonably well.

Algorithm 2 follows two sequential steps: scenario selection and policy selection. In the first step, we decide on which scenarios will be enforced to satisfy the chance constraint, i.e., $\hat{z}^s = 1$. As in the previous section, we relax the requirement that the selected policy should be the same over all scenarios, and solve problems (4.8) independently to obtain the optimal

Algorithm 2: *initialSolution()*

Given distinct cost and transition matrices $\{\tilde{\mathbf{c}}, \tilde{\mathbf{P}}\}_{s\in\mathcal{S}}$ with corresponding probabilities $\{p\}_{s\in\mathcal{S}}$, set $\hat{\mathbf{z}} \leftarrow \mathbf{0}$;

for *each scenario* $s \in \mathcal{S}$ **do**

  Solve problem (4.8) to obtain its optimal objective function value $\underline{b}^s$;

Compute $\mathrm{VaR}_\alpha(\underline{b})$ and find a subset of scenarios $\bar{\mathcal{S}} \subseteq \mathcal{S}$ such that $\underline{b}^s \leq \mathrm{VaR}_\alpha(\underline{b})$ for each scenario $s \in \bar{\mathcal{S}}$ and $\sum_{s\in\bar{\mathcal{S}}\setminus\{s'\}} p^s < \alpha$ for all $s' \in \bar{\mathcal{S}}$;

for *each scenario* $s \in \bar{\mathcal{S}}$ **do**

  Set $\hat{z}^s \leftarrow 1$;

Return *robustPolicySelection(*$\hat{\mathbf{z}}$*)*;

objective function value $\underline{b}^s$ corresponding to each scenario $s \in \mathcal{S}$. Then, we set $\hat{z}^s = 1$ for each scenario s in a subset $\bar{\mathcal{S}} \subseteq \mathcal{S}$ such that $\underline{b}^s \leq \mathrm{VaR}_\alpha(\underline{b})$ for each scenario $s \in \bar{\mathcal{S}}$ and $\sum_{s\in\bar{\mathcal{S}}\setminus\{s'\}} p^s < \alpha$ for all $s' \in \bar{\mathcal{S}}$.

In the policy selection phase provided in Algorithm 3, we use the scenarios selected in the first step to obtain a feasible policy. In other words, we let $\mathcal{S}(\mathbf{z}) = \bar{\mathcal{S}}$ in Line 1 of Algorithm 3. Based on this set, we obtain a policy by taking advantage of a substructure of the problem for fixed scenario selection decisions. Using the definition of $C(\boldsymbol{\pi}, \tilde{\mathbf{c}}^s, \tilde{\mathbf{P}}^s)$, original formulation (4.5) can be restated as

$$\min_{y\in\mathbb{R},\ \boldsymbol{\pi}\in\Pi_D} \quad y \tag{4.10a}$$

$$\text{s.t.} \quad y \geq C(\boldsymbol{\pi}, \tilde{\mathbf{c}}^s, \tilde{\mathbf{P}}^s) - (1 - z^s)M, \qquad s \in \mathcal{S}, \tag{4.10b}$$

$$\sum_{s\in\mathcal{S}} z^s p^s \geq \alpha, \tag{4.10c}$$

$$z^s \in \{0, 1\}, \quad s \in \mathcal{S}. \tag{4.10d}$$

For any given feasible vector $\mathbf{z}$, problem (4.10) is equivalent to

$$\text{RMDP}(\mathbf{z}): \quad \min_{\boldsymbol{\pi} \in \Pi_D} \max_{s \in \mathcal{S}(\mathbf{z})} C(\boldsymbol{\pi}, \tilde{\mathbf{c}}^s, \tilde{\mathbf{P}}^s), \tag{4.11}$$

where $\mathcal{S}(\mathbf{z}) := \{s \in \mathcal{S} |\ z^s = 1\}$. Note that the substructure (4.11) can be considered as a robust MDP, where the uncertainty in the transition matrix is coupled across the time horizon and the state space, i.e., a single realization of the the cost and transition matrices is selected randomly at the beginning and it holds for all decision epochs and states of the system. Using the scenario vector $\hat{\mathbf{z}}$ obtained in the scenario selection step, in Algorithm 3, we find a policy by solving a relaxed version of RMDP($\hat{\mathbf{z}}$), which we describe next.

Algorithm 3: *robustPolicySelection*($\mathbf{z}$)

---

1 Given a small tolerance parameter $\epsilon > 0$, set some positive $\hat{\mathbf{v}}_1 \in \mathbb{R}^{|\mathcal{H}|}$, $\mathcal{S}(\mathbf{z}) \leftarrow \{s \in \mathcal{S} |\ z^s = 1\}$ and $t \leftarrow 1$;

2 **for** *each state* $i \in \mathcal{H}$ **do**

3   **for** *each action* $a \in \mathcal{A}$ **do**

4   Set $\sigma_{ia} \leftarrow \max\limits_{s \in \mathcal{S}(\mathbf{z})} \left\{ c_i^s(a) + \gamma \sum_{j \in \mathcal{H}} \tilde{P}_{ij}^s(a) \hat{v}_t(j) \right\}$;

5   Compute $\hat{\mathbf{v}}_{t+1}(i) \leftarrow \min\limits_{a \in \mathcal{A}} \ \sigma_{ia}$;

6 **if** $||\hat{\mathbf{v}}_{t+1} - \hat{\mathbf{v}}_t|| < \frac{(1-\gamma)\epsilon}{\gamma}$ **then**

7   Go to line 10;

8 **else**

9   Set $t \leftarrow t + 1$ and go to line 2;

10 Return $\boldsymbol{\pi}$ such that $\pi_i \in \arg\max\limits_{a \in \mathcal{A}} \ \sigma_{ia}$;

Nilim and El Ghaoui [78] consider a relaxed variant of RMDP for the case of transition matrix uncertainty assuming a certain rectangularity property, i.e., the transition probabilities are independent over different states and actions. Additionally, the transition matrix is allowed to be time-varying. The authors prove that the optimal policy in this setting obeys a set of optimality conditions and propose a robust value iteration algorithm. In the case of stationary transition matrix uncertainty represented by a finite number of scenarios as in our problem framework, we show that this algorithm may produce sub-optimal policies on a non-trivial counterexample with 21 scenarios, eight states and five actions. Here, we further generalize their algorithm for the case of uncertainty in both cost and transition matrices to find a policy that performs well in terms of the VaR objective, but is not necessarily optimal. For the vector $\hat{\mathbf{z}}$ computed in the first step, we construct the set $\mathcal{S}(\hat{\mathbf{z}}) := \{s \in \mathcal{S} \mid \hat{z}^s = 1\}$ and solve the following equations through a variant of the value iteration algorithm as presented in Algorithm 3

$$v_i = \min_{a\in\mathcal{A}} \max_{s\in\mathcal{S}(\mathbf{z})} \left\{ c_i^s(a) + \gamma \sum_{j\in\mathcal{H}} \tilde{P}_{ij}^s(a) v_j \right\}, \quad i \in \mathcal{H}.$$

Note that the policy obtained by Algorithm 3 and the corresponding value functions provide upper bounds on the results of the robust problem RMDP stated in (4.11).

4.4 A Branch-and-Cut Algorithm

Because policy decisions are made before the actual realization of the random data, formulation (4.5) can be restated as a two-stage chance-constrained problem proposed by Luedtke et al. [71] and Liu et al. [66] with the following first-stage problem,

$$\begin{aligned}
\text{(MP)} \quad \min \quad & y && \text{(4.12a)}\\
s.t. \quad & \sum_{a\in\mathcal{A}} w_{ia} = 1, \quad i \in \mathcal{H}, && \text{(4.12b)}\\
& \sum_{s\in\mathcal{S}} z^s p^s \geq \alpha, && \text{(4.12c)}\\
& z^s \in \{0,1\}, \quad s \in \mathcal{S}, && \text{(4.12d)}\\
& w_{ia} \in \{0,1\}, \quad i \in \mathcal{H}, a \in \mathcal{A}, && \text{(4.12e)}\\
& y \geq 0 && \text{(4.12f)}\\
& (y, \mathbf{w}, \mathbf{z}) \in \mathcal{C}, && \text{(4.12g)}
\end{aligned}$$

where $\mathcal{C}$ is the set of Benders cuts obtained by solving the second-stage subproblems for all scenarios $s \in \mathcal{S}$. At any iteration of the branch-and-cut algorithm, we solve a relaxation of (MP) including only a finite subset of the Benders cuts, and obtain an incumbent first-stage solution $(\hat{y}, \hat{\mathbf{w}}, \hat{\mathbf{z}})$. The additional constraint (4.12f) ensures that $\hat{y} \geq 0$ at any iteration, and provides a natural bound on the objective function value since all cost terms are nonnegative. Note that we only need to ensure feasibility of the subproblems corresponding to the selected scenarios, i.e., $\hat{z}^s = 1$. If an incumbent first-stage solution is feasible for all selected subproblems, then it is also feasible for the original problem. Otherwise a feasibility cut must be added to eliminate infeasibilities. In addition, the policy corresponding to the first-stage solution can be stated as $\hat{\boldsymbol{\pi}} = (\bar{a}_1, \bar{a}_2, \dots, \bar{a}_{|\mathcal{H}|})$, where $\bar{a}_i$ is the action taken in state i in the incumbent first-stage solution, i.e., $\hat{w}_{i\bar{a}_i} = 1$.

Using the first-stage solution $(\hat{y}, \hat{\mathbf{w}}, \hat{\mathbf{z}})$, the second-stage subproblem for scenario $s \in S$ with $\hat{z}^s = 1$ can be stated as

$$(\text{PS}^s) \quad \min \; 0 \tag{4.13a}$$

$$s.t. \quad \sum_{i \in \mathcal{H}} q_i v_i \leq \hat{y} + (1 - \hat{z}^s) M, \tag{4.13b}$$

$$v_i \geq \tilde{c}_i^s(a) + \gamma \sum_{j \in \mathcal{H}} \tilde{P}_{ij}^s(a) v_j - (1 - \hat{w}_{ia}) M_{is}, \quad i \in \mathcal{H}, \; a \in \mathcal{A}. \tag{4.13c}$$

Scenario indices on state value variables $\mathbf{v}$ are omitted for brevity. Note that constraints (4.13c) corresponding to actions such that $\hat{w}_{ia} = 0$ could be omitted since they are redundant. Let $-\rho$ and $\boldsymbol{\eta}$ be the dual variables associated with constraints (4.13b) and (4.13c), respectively. Then, dual of the second-stage problem for scenario s is

$$(\text{DS}^s) \quad \max \quad \sum_{i \in \mathcal{H}} \sum_{a \in \mathcal{A}} (\tilde{c}_i^s(a) - (1 - \hat{w}_{ia}) M_{is}) \eta_{ia} - \rho \hat{y} - \rho (1 - \hat{z}^s) M \tag{4.14a}$$

$$s.t. \quad \sum_{a \in \mathcal{A}} \eta_{ia} - \gamma \sum_{j \in \mathcal{H}} \sum_{a \in \mathcal{A}} \tilde{P}_{ji}^s(a) \eta_{ja} - q_i \rho = 0, \quad i \in \mathcal{H}, \tag{4.14b}$$

$$\rho \geq 0, \quad \eta_{ia} \geq 0, \; i \in \mathcal{H}, \; a \in \mathcal{A}, \tag{4.14c}$$

and the corresponding Benders cut for scenario s becomes

$$\hat{\rho} y + \hat{\rho}(1 - z^s)M \geq \sum_{i\in\mathcal{H}} \sum_{a\in\mathcal{A}} (\tilde{c}_i^s(a) - (1 - w_{ia})M_{is})\hat{\eta}_{ia}. \tag{4.15}$$

These cuts improve the lower bound on y only for the specific policy for which the dual variables $(\hat{\rho}, \hat{\boldsymbol{\eta}})$ are computed, and the scenarios with $z^s = 1$. Instead, we can employ a stronger class of valid inequalities proposed by Luedtke et al. [71] (see also [1, 60, 67, 121] for further strengthenings) that requires an extreme ray $(\hat{\rho}, \hat{\boldsymbol{\eta}})$ of the dual second-stage problem problem (DS^s).

The following lemma provides a characterization for an extreme ray that makes problem (DS^s) unbounded when the primal second-stage problem (PS^s) is infeasible for some $s \in \mathcal{S}$.

Lemma 2. *Given a first-stage solution* $(\hat{y}, \hat{\mathbf{w}}, \hat{\mathbf{z}})$ *and its corresponding policy* $\hat{\boldsymbol{\pi}} = (\bar{a}_1, \bar{a}_2, \ldots, \bar{a}_{|\mathcal{H}|})$*, if the second-stage subproblem* (PS^s) *is infeasible for some scenario* $s \in \mathcal{S}$ *with* $\hat{z}^s = 1$*, then there exists a direction of unboundedness for problem* (DS^s)*, where the dual variable* $\boldsymbol{\eta}$ *satisfies* $\eta_{ia} = 0$ *for* $i \in \mathcal{H},\ a \in \mathcal{A}$ *unless* $a = \bar{a}_i$*. Consequently,* $\sum_{a\in\mathcal{A}} \eta_{ia} = \eta_{i\bar{a}_i}$*.*

Proof. Consider some scenario $s \in \mathcal{S}$ with $\hat{z}^s = 1$ such that the second-stage subproblem (PS^s) is infeasible. Note that the last term in (4.14a) cancels out, because $(1 - \hat{z}^s) = 0$.

Constraint (4.14b) and the objective (4.14a) can be scaled by $\rho > 0$. Hence, without loss of generality, we consider two cases: $\rho = 1$ and $\rho = 0$. In the first case, the problem is equivalent to the dual of a classical MDP maximizing the expected total discounted cost with updated cost values $\check{c}_i^s(a) = \tilde{c}_i^s(a) - (1 - \hat{w}_{ia})M_{is}$ for $i \in \mathcal{H},\ a \in \mathcal{A}$. Therefore, assuming sufficiently large M_{is} values for all $i \in \mathcal{H}$, we can obtain a direction of unboundedness for (DS^s) with $\eta_{ia} = 0$ for $i \in \mathcal{H},\ a \in \mathcal{A}$ unless $a = \bar{a}_i$.

When $\rho = 0$, the dual second-stage problem reduces to

$$\max \quad \sum_{i\in\mathcal{H}} \sum_{a\in\mathcal{A}} (\tilde{c}_i^s(a) - (1 - \hat{w}_{ia})M_{is})\eta_{ia} \tag{4.16a}$$

$$s.t. \quad \sum_{a\in\mathcal{A}} \eta_{ia} - \gamma \sum_{j\in\mathcal{H}} \sum_{a\in\mathcal{A}} \tilde{P}_{ji}^s(a)\eta_{ja} = 0, \quad i \in \mathcal{H}, \tag{4.16b}$$

$$\eta_{ia} \geq 0, \ i \in \mathcal{H}, \ a \in \mathcal{A}, \tag{4.16c}$$

which is the dual of

$$\min \ 0 \tag{4.17a}$$

$$s.t. \ \ v_i \geq (\tilde{c}_i^s(a) - (1 - \hat{w}_{ia})M_{is}) + \gamma \sum_{j \in \mathcal{H}} \tilde{P}_{ij}^s(a) v_j, \quad i \in \mathcal{H}, \ a \in \mathcal{A}. \tag{4.17b}$$

Formulation (4.17) is always feasible, because its feasible region corresponds to the feasible region of the LP formulation of an MDP with updated cost terms, which is known to have a feasible solution. Hence, weak duality ensures that the objective function of (4.16) cannot be positive. Consequently, when $\rho = 0$, there cannot exist a direction of unboundedness for (DS^s), which contradicts with our initial assumption that (PS^s) is infeasible. □

Given an extreme ray $(\hat{\rho}, \hat{\boldsymbol{\eta}})$ of (DS^s) as described in Lemma 2, we need to solve the following problem for each scenario $s \in \mathcal{S}$

$$\begin{aligned}
\min \quad & \hat{\rho} y + \sum_{i \in \mathcal{H}} \sum_{a \in \mathcal{A}} (1 - w_{ia}) M_{is} \hat{\eta}_{ia} \\
s.t. \quad & \sum_{i \in \mathcal{H}} q_i v_i^s \leq y, \\
& v_i^s \geq \tilde{c}_i^s(a) + \gamma \sum_{j \in \mathcal{H}} \tilde{P}_{ij}^s(a) v_j^s - (1 - w_{ia}) M_{is}, \quad i \in \mathcal{H}, \ a \in \mathcal{A}, \\
& w_{ia} \in \{0, 1\}, \quad i \in \mathcal{H}, \ a \in \mathcal{A}.
\end{aligned}$$

Observe that the optimal objective function value of this problem is equal to the expected total cost under scenario s and policy $\hat{\boldsymbol{\pi}} = (\bar{a}_1, \bar{a}_2, \ldots, \bar{a}_{|\mathcal{H}|})$ multiplied by the dual variable $\hat{\rho} \geq 0$, since otherwise the big-M terms and Lemma 2 ensure that the objective function takes a larger value than the objective function value corresponding to policy $\hat{\boldsymbol{\pi}}$. Hence we can solve the following analogous problem

$$\min \quad \sum_{i\in\mathcal{H}} q_i v_i^s \tag{4.18a}$$

$$s.t. \quad v_i^s \geq \tilde{c}_i^s(\bar{a}_i) + \gamma \sum_{j\in\mathcal{H}} \tilde{P}_{ij}^s(\bar{a}_i) v_j^s, \quad i \in \mathcal{H}, \tag{4.18b}$$

which is equal to $C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^s, \tilde{\mathbf{P}}^s)$, the expected total discounted cost of the underlying Markov chain for policy $\hat{\boldsymbol{\pi}} \in \Pi$ under scenario $s \in \mathcal{S}$. This step can be regarded as the evaluation of the policy selected at the first-stage for n distinct MDPs considering different scenarios. The optimal value $C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^s, \tilde{\mathbf{P}}^s)$ for each scenario s can simply be computed by solving a system of $|\mathcal{H}|$ linear equalities as in the policy evaluation step of the policy iteration algorithm. Let σ be a permutation of scenarios such that $C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{\sigma_1}, \tilde{\mathbf{P}}^{\sigma_1}) \geq C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{\sigma_2}, \tilde{\mathbf{P}}^{\sigma_2}) \geq \cdots \geq C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{\sigma_{|\mathcal{S}|}}, \tilde{\mathbf{P}}^{\sigma_{|\mathcal{S}|}})$ and let $k := \max\{i : \sum_{j=1}^{i} p^{\sigma_j} \leq 1 - \alpha\}$. In the following proposition, we present an alternative set of feasibility cuts based on the new subproblems (4.18) and constraint (4.12c), which is equivalent to the knapsack inequality $\sum_{s\in\mathcal{S}}(1 - z^s)p^s \leq 1 - \alpha$.

Proposition 8. *Let $T = \{t_1, \ldots, t_l\} \subseteq \{\sigma_1, \ldots, \sigma_k\}$ be such that $C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{t_i}, \tilde{\mathbf{P}}^{t_i}) \geq C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{t_{i+1}}, \tilde{\mathbf{P}}^{t_{i+1}})$ for $i = 1, \ldots, l$, where $t_1 = \sigma_1$ and $t_{l+1} = \sigma_{k+1}$, and $\bar{M}$ represent a large number such that inequality* (4.19) *becomes redundant when $w_{j\bar{a}_j} = 0$ for some $j \in \mathcal{H}$. Then the inequality*

$$y + \sum_{i=1}^{l} \left(C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{t_i}, \tilde{\mathbf{P}}^{t_i}) - C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{t_{i+1}}, \tilde{\mathbf{P}}^{t_{i+1}}) \right) (1 - z^{t_i}) \geq C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{t_1}, \tilde{\mathbf{P}}^{t_1}) - \bar{M} \sum_{j\in\mathcal{H}} (1 - w_{j\bar{a}_j}) \tag{4.19}$$

is valid for problem (4.5).

Proof. Using Theorem 1 in [69] by substituting $\alpha x = \hat{\rho} y + \sum_{j\in\mathcal{H}} \sum_{a\in\mathcal{A}} (1 - w_{ja}) \hat{\eta}_{ja} M_{is}$ and the knapsack inequality $\sum_{s\in\mathcal{S}}(1 - z^s)p^s \leq 1 - \alpha$, we obtain the following valid inequality

$$\hat{\rho} y + \sum_{j\in\mathcal{H}} \sum_{a\in\mathcal{A}} (1 - w_{ja}) \hat{\eta}_{ja} M_{is} + \sum_{i=1}^{l} \left(\hat{\rho} C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{t_i}, \tilde{\mathbf{P}}^{t_i}) - \hat{\rho} C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{t_{i+1}}, \tilde{\mathbf{P}}^{t_{i+1}}) \right) (1 - z^{t_i}) \geq \hat{\rho} C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{t_1}, \tilde{\mathbf{P}}^{t_1}).$$

Due to Lemma 2, we know that $\hat{\eta}_{ia} = 0$ unless $a = \bar{a}_i$, which implies that the inequality above can be equivalently written as

$$\hat{\rho}y + \sum_{j\in\mathcal{H}}(1-w_{j\bar{a}_j})\hat{\eta}_{j\bar{a}_j}M_{is} + \sum_{i=1}^{l}\left(\hat{\rho}C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{t_i}, \tilde{\mathbf{P}}^{t_i}) - \hat{\rho}C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{t_{i+1}}, \tilde{\mathbf{P}}^{t_{i+1}})\right)(1-z^{t_i}) \geq \hat{\rho}C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{t_1}, \tilde{\mathbf{P}}^{t_1}).$$

Here the big-M terms ensure that this inequality is only effective for policy $\hat{\boldsymbol{\pi}} = (\bar{a}_1, \bar{a}_2, \ldots, \bar{a}_{|\mathcal{S}|})$ and otherwise redundant. Therefore, it can be restated as

$$\hat{\rho}y + \sum_{i=1}^{l}\left(\hat{\rho}C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{t_i}, \tilde{\mathbf{P}}^{t_i}) - \hat{\rho}C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{t_{i+1}}, \tilde{\mathbf{P}}^{t_{i+1}})\right)(1-z^{t_i}) \geq \hat{\rho}C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{t_1}, \tilde{\mathbf{P}}^{t_1}) - \hat{\rho}\bar{M}\sum_{j\in\mathcal{H}}(1-w_{j\bar{a}_j}).$$

so that the dependency on the value of the extreme ray $(\hat{\rho}, \hat{\boldsymbol{\eta}})$ can be eliminated. After scaling the valid inequality above by $\hat{\rho} \geq 0$, we obtain inequality (4.19). $\square$

Using valid inequalities (4.19), we no longer need to solve the subproblems PS^s. Instead, any policy $\hat{\boldsymbol{\pi}}$ selected in the first-stage (given by the incumbent solution $\hat{\mathbf{w}}$) is evaluated by computing $C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^s, \tilde{\mathbf{P}}^s)$ for all scenarios $s \in \mathcal{S}$. Then, inequalities (4.19) cut off any infeasible first-stage solution $(\hat{y}, \hat{\mathbf{w}}, \hat{\mathbf{z}})$ and provide a lower bound on y using feasible first-stage solutions.

The valid inequalities (4.19) still contain the $\bar{M}$ term that may significantly effect computational efficiency of the proposed method. Ideally, it should correspond to the largest possible decrease in the value of y by policy changes. In the following corollary, we provide a valid choice of $\bar{M}$.

Corollary 1. *By setting* $\bar{M} = C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{t_1}, \tilde{\mathbf{P}}^{t_1}) - b_l$ *in* (4.19)*, we obtain the following valid inequalities*

$$\begin{aligned} y + \sum_{i=1}^{l}\left(C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{t_i}, \tilde{\mathbf{P}}^{t_i}) - C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{t_{i+1}}, \tilde{\mathbf{P}}^{t_{i+1}})\right)(1-z^{t_i}) & \\ \geq C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{t_1}, \tilde{\mathbf{P}}^{t_1}) - \left(C(\hat{\boldsymbol{\pi}}, \tilde{\mathbf{c}}^{t_1}, \tilde{\mathbf{P}}^{t_1}) - b_l\right)\sum_{j\in\mathcal{H}}(1-w_{j\bar{a}_j}). & \end{aligned} \tag{4.20}$$

Proof. Arranging the terms in inequality (4.19), we obtain

$$\bar{M}\sum_{j\in\mathcal{H}}(1-w_{j\bar{a}_j}) \geq C(\hat{\boldsymbol{\pi}},\tilde{\mathbf{c}}^{t_1},\tilde{\mathbf{P}}^{t_1}) - \sum_{i=1}^{l}\Big(C(\hat{\boldsymbol{\pi}},\tilde{\mathbf{c}}^{t_i},\tilde{\mathbf{P}}^{t_i}) - C(\hat{\boldsymbol{\pi}},\tilde{\mathbf{c}}^{t_{i+1}},\tilde{\mathbf{P}}^{t_{i+1}})\Big)(1-z^{t_i}) - y,$$

and the following lower bound on $\bar{M}$

$$\bar{M} \geq \max_{\mathbf{z}\in\{0,1\}^{|\mathcal{S}|}:(4.5\text{c})}\left\{C(\hat{\boldsymbol{\pi}},\tilde{\mathbf{c}}^{t_1},\tilde{\mathbf{P}}^{t_1}) - \sum_{i=1}^{l}\Big(C(\hat{\boldsymbol{\pi}},\tilde{\mathbf{c}}^{t_i},\tilde{\mathbf{P}}^{t_i}) - C(\hat{\boldsymbol{\pi}},\tilde{\mathbf{c}}^{t_{i+1}},\tilde{\mathbf{P}}^{t_{i+1}})\Big)(1-z^{t_i})\right\} - \min y.$$

Hence, we can set $\bar{M} = C(\hat{\boldsymbol{\pi}},\tilde{\mathbf{c}}^{t_1},\tilde{\mathbf{P}}^{t_1}) - b_l$ in (4.19) and obtain inequality (4.20) which is valid due to the lower bound $y \geq b_l$, the nonnegativity of the second term on the left-hand-side and the fact that $C(\hat{\boldsymbol{\pi}},\tilde{\mathbf{c}}^{t_1},\tilde{\mathbf{P}}^{t_1}) - b_l \geq 0$. When $w_{j\bar{a}_j} = 1$ for all $j \in \mathcal{H}$, the validity of inequality (4.20) directly follows from Proposition 8. Otherwise, $\sum_{j\in\mathcal{H}}(1-w_{j\bar{a}_j}) \geq 1$ and consequently the right-hand-side of inequality (4.20) is at most b_l, which is a lower bound on y.

□

Note that although MDPs with known parameters can be modeled as an LP with implicit policy decisions, our formulation (4.5) cannot benefit from this property and requires integer policy variables $\mathbf{w}$, which significantly increases computational complexity and hinders efficiency of the decomposition algorithms. This is mainly because of the requirement that the selected policy should be the same over all scenarios. We further elaborate on modeling challenges and alternative decomposition approaches in Appendix A.

4.5 An Application to Inventory Management in Long-Term Humanitarian Relief Operations

Long-term humanitarian relief operations, alternatively referred to as continuous aid operations, are vital to sustain daily basic needs of a population affected by a slow onset disaster including war (Syria, Afghanistan, Iraq), political insurrection (Syria), famines (Yemen, South Sudan, Somalia), droughts (Ethiopia) and extreme poverty (Niger, Liberia). Unlike the sudden onset disasters (e.g., earthquakes, hurricanes, terrorist attacks), they require delivery of materials such as food, water and medical supplies to satisfy a chronic need over a long period of time. Since the progress of a slow onset disaster usually presents irregularities in terms of its scale and location, the demand is highly uncertain. In addition, the supply levels are also exposed to uncertainty as they mainly depend on donations from multiple resources. More than 90% of the people affected by slow onset disasters lives in developing countries, hence the required relief items are usually outsourced from resources in various locations around the world [95]. Another important consideration in long-term humanitarian operations is the perishability of the relief items. Many items needed during and after a disaster, e.g., food and medication, have a limited shelf life. In addition to the possible uncertainty in the initial shelf life, the remaining shelf life upon arrival is also affected by the uncertainty in the lead times due to unknown location of the donations. Hence, high level of uncertainty inherent in the supply chain makes the system prone to unwanted shortages and disposals. To prevent interruptions, the policy makers may interfere through different actions such as campaigns and advertisements that provide additional relief items.

In this section, we consider an inventory management model for a single perishable item in long-term humanitarian relief operations proposed by Ferreira et al. [43]. They formulate the problem as an infinite horizon MDP with finite state and action spaces, where the states represent possible levels of inventory. The aim is to minimize the long-term average expected cost. Their model assumes that both demand and supply is uncertain following Poisson distributions with known demand and supply rates, respectively, and the expiration time of the supply items is deterministic. These parameters are generally obtained using various forecasting methods. However, due to multiple sources of uncertainty in the context of long-term humanitarian operations, the forecasted values may be erroneous, which may consequently affect the performance of the resulting policy. To handle the parameter uncertainty, different from [43], we assume that demand and supply rates and the expiration time probabilistically take value from a finite set of scenarios $\mathcal{S}$ and the objective in our model is to minimize the α-quantile of the total discounted expected total cost with respect to the uncertainty in these parameters.

In what follows, we elaborate on the components and assumptions of the MDP model based on [43].

<u>*States:*</u> The state of the system is represented by the number of available items in the inventory at the beginning of each decision epoch that will not expire before delivery. It is assumed that the inventory has a capacity K, e.g., the state space $\mathcal{H} = \{0, 1, \ldots, K\}$.

<u>*Actions:*</u> At the beginning of each decision epoch, the decision maker takes an action a from a finite set $\mathcal{A}$ that provides an additional n_a number of relief items. The set of actions taken for each possible state of the system that minimizes the objective function is referred to as the optimal policy.

<u>*Transition Probabilities:*</u> Under each scenario $s \in \mathcal{S}$, the demand (D) and supply (U) for the item follow Poisson distribution with probability mass functions f_d^s and f_u^s of rates μ_d^s and μ_u^s, respectively, and the expiration time takes value t_e^s. We assume that the demand rate is at least as large as the supply rate. Let $\Delta_{\min}$ and $\Delta_{\max}$ be the minimum and maximum possible difference between the supply and demand at any decision epoch considering a $100(1-\epsilon)\%$ confidence level for a small $\epsilon > 0$. Note that the minimum possible demand and supply amounts are assumed as zero so that $\Delta_{\min}$ and $\Delta_{\max}$ represent the negative of maximum possible demand and the maximum possible supply, respectively. Then, ignoring the perishability of the item and the actions taken by the decision maker, the probability of transitioning from inventory level $i \in \mathcal{H}$ to $j \in \mathcal{H}$ under scenario $s \in \mathcal{S}$ becomes

$$\hat{P}_{ij}^s = \begin{cases} \sum_{\Delta=\Delta_{\min}}^{j-i} \sum_{D=0}^{-\Delta_{\min}} f_d^s(D) f_u^s(D+\Delta), & \text{if } j = 0, \\ \sum_{\Delta=j-i}^{\Delta_{\max}} \sum_{D=0}^{-\Delta_{\min}} f_d^s(D) f_u^s(D+\Delta), & \text{if } j = K, \\ \sum_{D=0}^{-\Delta_{\min}} f_d^s(D) f_u^s(D+j-i), & \text{otherwise.} \end{cases}$$

Instead of keeping track of the number of items that expire in each period, the model is simplified by assuming that the items procured in each decision epoch have the same expiration time. As a result, at the beginning of each decision epoch, the expiration probability for the whole batch of newly arriving items is computed as the probability of not being able to consume all available items before the expiration time. Note that since the demand follows Poisson distribution with rate μ_d^s, the probability of consuming k items before time t_e follows Erlang distribution with parameters μ_d^s and k, which can be stated as $(1 - f_e^s(t_e^s, k))$. Incorporating perishability of the items based on this simplification, the probability of a transition from inventory level $i \in \mathcal{H}$ to inventory level $j \in \mathcal{H}$ under scenario $s \in \mathcal{S}$ when action $a \in \mathcal{A}$ is taken can be computed as follows.

$$\tilde{P}_{ij}^s(a) = \begin{cases} (1 - f_e^s(t_e^s, j))\ \hat{P}_{i+n_a,j}^s, & \text{if } j > i + n_a, \\ \displaystyle\sum_{j'=i+n_a+1}^{\min(K, i+n_a+\Delta_{\max})} f_e^s(t_e^s, j')\ \hat{P}_{i+n_a,j'}^s + \hat{P}_{i+n_a,j}^s, & \text{if } j = i + n_a, \\ \hat{P}_{i+n_a,j}^s, & \text{otherwise.} \end{cases} \tag{4.21}$$

Equation (4.21) follows from the fact that if an arriving batch of items is known to expire, the new arrivals are immediately used to fulfill the demand with priority over the existing inventory (the existing inventory can be used later since it is guaranteed not to expire by the previous assumption), and the remaining items in the batch are disposed. Hence, the case $j > i + n_a$ can occur only if the new items do not expire. Similarly, $j = i + n_a$ implies either that the incoming supply is at least as much as the demand in the current period and the supply surplus is disposed, or demand is equal to the sum of supply and additional items acquired by the action taken.

Cost function: The expected cost of taking action $a \in \mathcal{A}$ at inventory level $i \in \mathcal{H}$ under scenario $s \in \mathcal{S}$, stated as $\tilde{c}_i^s(a)$, consists of three main components: the inventory holding cost, the expected shortage cost and the expected disposal cost. Assuming a unit disposal cost of $\tilde{d}$, the total expected disposal cost when action $a \in \mathcal{A}$ is taken at inventory level $i \in \mathcal{H}$ under scenario $s \in \mathcal{S}$ becomes

$$DC(i,a,s) = \sum_{\Delta=1}^{\min(K-i-n_a,\Delta_{\max})} f_e^s(t_e^s, i+n_a+\Delta)\hat{P}^s_{i+n_a,i+n_a+\Delta}\ \Delta\ \tilde{d}$$
$$+ \sum_{j=K}^{K+n_a+\Delta_{\max}} \sum_{D=0}^{-\Delta_{\min}} f_d^s(D) f_u^s(D+j-i-n_a)\ (j-K)\ \tilde{d}.$$

The first summation term is due to the expired items, while the second term is because of the inventory surplus. Similarly, the shortage costs are stated as

$$SC(i,a,s) = \sum_{j=i+n_a+\Delta_{\min}}^{-1} \sum_{D=0}^{-\Delta_{\min}} f_d^s(D) f_u^s(D+j-i-n_a)\ (-j)\ \tilde{u}, \quad i \in \mathcal{H},\ a \in \mathcal{A},\ s \in \mathcal{S},$$

where $\tilde{u}$ is the unit shortage cost. Assuming that the unit inventory holding cost is h, the total expected cost of taking action $a \in \mathcal{A}$ at inventory level $i \in \mathcal{H}$ under scenario $s \in \mathcal{S}$ can be computed using the following equation

$$\tilde{c}_i^s(a) = hi + DC(i,a,s) + SC(i,a,s), \quad i \in \mathcal{H},\ a \in \mathcal{A},\ s \in \mathcal{S}.$$

Note that the cost structure may be different depending on the characteristics of the environment in consideration. The MDP formulation provides decision makers the flexibility to use even non-convex cost functions. Our methodology only requires the cost terms to be bounded.

4.6 Computational Study

In this section, we conduct computational experiments on the long-term humanitarian relief operations inventory management problem described in Section 4.5 to examine the effects of incorporating parameter uncertainty into MDP models and to compare the efficiency of different solution approaches. The problem instances are generated based on the experiments provided in [43] considering the inventory management for a blood center collecting and distributing blood packs to support humanitarian relief operations. We assume that the inventory replenishment decisions are made on a weekly basis, and the weekly demand and supply rates for the blood packs take values in the intervals $[80, 100]$ and $[50, 70]$, respectively. After donation, each blood pack has a shelf life of six weeks, however unknown lead times may affect the shelf life remained at the time of arrival to the blood center. Hence, the shelf life is assumed to be random on the interval $[1, 6]$. In case of need, the center may procure additional supply of blood packs by sending up to $|\mathcal{A}| \in \{2, 3, 4\}$ blood collection vehicles to distant areas. Each vehicle collects additional 20 blood packs at the expense of incurring a certain cost. The cost parameters for additional procurement, inventory holding costs, disposal costs and shortage costs are given in Table 4.1. We additionally assume that the blood center has a capacity of $K \in \{50, 100, 150\}$ units and the blood packs are in batches of 10 units so that $|\mathcal{H}| \in \{5, 10, 15\}$. Different than [43], we generate random instances with $|\mathcal{S}| \in \{5, 10, 20\}$ equiprobable scenarios, where the parameters of demand and supply rates, shelf life, and disposal and shortage costs for each scenario randomly take value on their respective intervals stated above. The distribution on the initial state is assumed to be uniform as well.

All experiments are performed using single thread of a Windows server with Intel(R) Xeon(R) CPU E5-2630 processor at 2.40 GHz and 32 GB of RAM using Python 3.6.2 and Gurobi Optimizer 8.0. The time limit for each instance on Gurobi is set to 3600 seconds and we use the default settings for the MIP gap and feasibility tolerances. The results are obtained for the instances with $\alpha = 0.80$ and $\gamma = 0.99$.

4.6.1 Model Analysis

We first examine the impact of incorporating parameter uncertainty into the MDP model in terms of the value gained using the stochastic information on parameters and the structure of the optimal policies minimizing the VaR objective.

In Table 4.2, we compare the objective function value of the quantile optimization problem (OPT) with two benchmark cases. The first one assumes that the decision maker waits

Table 4.1: Cost parameters

Source	Cost (in dollars)
No additional procurement	60.00
One vehicle	120.00
Two vehicles	170.00
Three vehicles	225.00
Four vehicles	280.00
Unit inventory holding cost (h)	1.00
Unit disposal cost ($\tilde{d}$)	*Uniform*(900.00,1100.00)
Unit shortage cost ($\tilde{u}$)	*Uniform*(900.00,1100.00)

until observing the actual parameter realizations and makes a decision for each scenario independently. This approach does not provide a feasible solution as it may produce distinct policies for all scenarios. Clearly the quantile value in this case corresponds to a lower bound on the OPT since it can be stated as LB = $\mathrm{VaR}_\alpha(\underline{b})$, where the realization under scenario $s \in \mathcal{S}$ is $\underline{b}^s = \min_{\boldsymbol{\pi} \in \Pi_D} C(\boldsymbol{\pi}, \tilde{\mathbf{c}}^s, \tilde{\mathbf{P}}^s)$. Using this value, we compute the value of perfect information on MDP parameters as VPI = OPT − LB. The second case considers the quantile function value corresponding to a policy obtained by solving a single MDP with expected parameter values, referred as the mean value problem in stochastic programming context. The quantile function value in this case, denoted as MV, provides an upper bound on OPT. Based on MV, the value of incorporating stochasticity of parameters into the MDP model can be measured as VSS = MV − OPT. The VPI and VSS values for different number of scenarios $|\mathcal{S}|$, states $|\mathcal{H}|$ and actions $|\mathcal{A}|$ are presented in Table 4.2. Each row corresponds to the average of two replications and only the instances which can be solved to optimality within the time limit for both replications are included. The results show that by incorporating uncertainty in the parameters, the α-quantile value can be improved by \$9124.39 on average with a maximum improvement of \$31501.97. Furthermore, the average and maximum losses in the quantile function value due to not knowing the true parameter realizations is \$20168.71 and \$42384.58, respectively.

Table 4.2: Value of perfect information and the quantile optimization solution

$\|\mathcal{S}\|$	$\|\mathcal{H}\|$	$\|\mathcal{A}\|$	OPT	LB	MV	VPI	VSS
5	6	3	1323266.89	1320642.13	1325802.94	2624.76	2536.05
5	6	4	1275652.00	1233267.43	1285268.76	42384.57	9616.76
5	6	5	1275652.01	1233267.43	1285268.76	42384.58	9616.76
5	11	3	630088.05	629180.00	630088.05	908.06	0.00
5	11	4	496600.75	489783.17	510460.50	6817.58	13859.74
5	16	5	237153.12	237153.12	243500.16	0.00	6347.04
10	6	3	1435969.99	1405018.51	1448186.73	30951.48	12216.75
10	6	4	1312221.11	1288666.06	1312221.11	23555.04	0.00
10	6	5	1312221.11	1288666.06	1312221.11	23555.04	0.00
10	11	3	779098.65	778993.64	800260.81	105.01	21162.16
10	11	4	507803.55	484532.90	530678.26	23270.65	22874.71
10	16	3	504524.22	504524.22	514766.17	0.00	10241.95
20	6	3	1480946.44	1454615.00	1489090.57	26331.44	8144.14
20	6	4	1349030.91	1310024.28	1349030.91	39006.63	0.00
20	6	5	1349030.91	1307089.02	1349030.91	41941.88	0.00
20	11	3	828917.02	827008.07	860418.99	1908.95	31501.97
20	11	4	529109.88	491987.43	536106.44	37122.45	6996.56
Average						20168.71	9124.39
Maximum						42384.58	31501.97

Figure 4.6.1 illustrates behaviour of the optimal policies under different settings for a particular instance with five scenarios, an inventory capacity of 150 units and at most two vehicles, e.g., $|\mathcal{S}| = 5$, $|\mathcal{H}| = 16$, $|\mathcal{A}| = 3$. The red line marked with circles corresponds to the quantile-optimizing policy, where two vehicles are dispatched in each period with at most 70 units in the inventory, one vehicle if the inventory level is in the interval $(70, 90]$ and no vehicles otherwise. Similarly, the line marked with squares depicts the mean value policy and the dashed lines represent the optimal policies for each scenario independently. It can be seen that the optimal policy for the quantile optimization problem is more aligned with the extreme scenarios, while the mean value policy represents an average behaviour as expected. Our experiments also show that the optimal policies for the quantile optimization problem are in general characterized as threshold policies, where the number of vehicles to

be dispatched decreases as the inventory level increase, as illustrated in Figure 4.6.1. Note that in some cases a certain structure of the policies may be desired by the decision makers even though not optimal. Using such information on the characteristics of the desired policy, it is possible to add the following constraint into our model to enforce a threshold structure

$$w_{ia} \leq \sum_{a' \in \mathcal{A}:\ a' \leq a} w_{i'a'}, \quad a \in \mathcal{A},\ i, i' \in \mathcal{H}:\ i' > i. \tag{4.22}$$

Constraint (4.22) ensures that if n_a vehicles are dispatched at the inventory level i, then for any inventory level $i' > i$, the number of vehicles dispatched can be at most n_a (assuming $n_a > n'_a$ for $a, a' \in \mathcal{A}$ such that $a > a'$, as given in the problem statement). Incorporating more information on the characteristics of the desired policy may also provide computational advantages as it reduces the solution space.

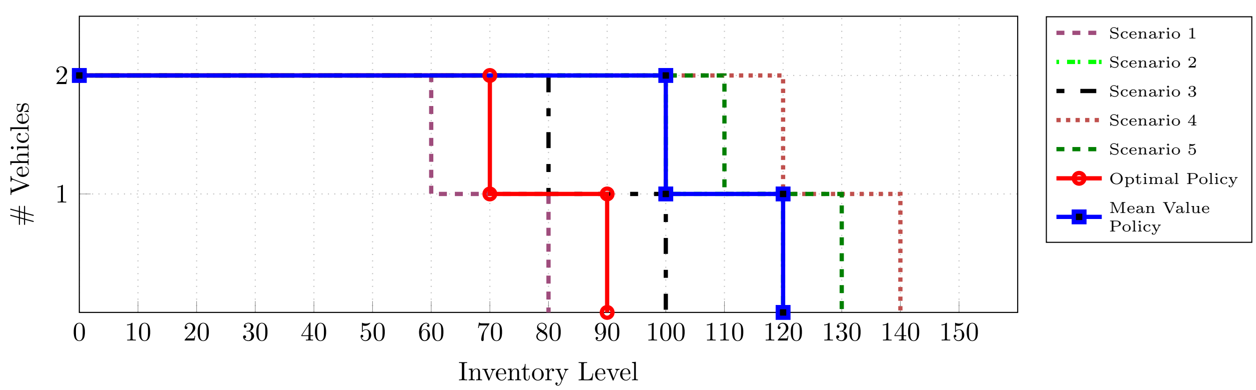

Figure 4.1: Comparison of policies computed under different settings

4.6.2 Comparison of Solution Approaches

In this section, we evaluate computational performance of the proposed methods in terms of solution times and optimality gaps achieved within the time limit of one hour. Since the policy obtained by solving the quantile-optimization problem incorporates the anticipated uncertainties in the parameters, it is safe to assume there is no need to recompute the optimal policy frequently unless a major change in the environment. Hence, the decision makers more likely prefer an optimal policy to a heuristic policy at the expense of higher computation times.

We first examine accuracy of the heuristic solutions generated by Algorithm 2 and the mean value (MV) problem. Table 4.3 reports the computation time "hTime" (resp., mvTime) and the optimality gap "hGap" (resp., mvGap) for policy of Algorithm 2 (resp., the MV policy), and the difference between (mvGap-hGap). The optimality gap for a policy with objective function value obj is computed as $100 \times \frac{(obj-obj^*)}{obj^*}$, where obj^* refers to the optimal objective function value of the quantile minimization problem. Each row corresponds to an average of two replications. The results show that both Algorithm 2 and the MV problem produce fairly good heuristic solutions. Algorithm 2 provides a better accuracy as it has a lower average optimality gap (1.33% vs 2.50%) and a lower maximum optimality gap (4.15% vs 7.91%) than the MV policy. On the other hand, in general, it takes shorter time to solve the MV problem using a value iteration algorithm compared to Algorithm 2. The average solution time of Algorithm 2 is 9.58 seconds, whereas the MV problem can be solved in 1.26 seconds in the average.

Next, we evaluate computational performance of the proposed exact solution methods under different settings as listed below.

- *DEF-base*: corresponds to formulation (4.5) with constraints (4.6) replacing (4.5e), and the big-M terms are set to $M = b_u$ and $M_{is} = \bar{v}_i^s,\ i \in \mathcal{H}, s \in \mathcal{S}$, where $\bar{v}_i^s$ the optimal value of variable v_i in problem (4.7) for scenario s.

- *DEF-P*: Different than *DEF-base*, it utilizes the lower and upper bound values obtained by preprocessing. The additional constraints (4.9) and scenario elimination given in Section 4.3.1 in which the upper bound is set to the best value obtained by Algorithm 2 and the MV problem are applied.

- *DEF-PMI*: In addition to *DEF-P*, the big-M terms are set to the tighter values $M =$

Table 4.3: Evaluation of the policy provided by Algorithm 2 and the MV policy

$\lvert\mathcal{S}\rvert$	$\lvert\mathcal{H}\rvert$	$\lvert\mathcal{A}\rvert$	hTime	hGap	mvTime	mvGap	mvGap-hGap
5	6	3	1.78	0.21	0.37	0.21	0.00
5	6	4	2.27	0.74	0.46	0.74	0.00
5	6	5	2.78	0.74	0.55	0.74	0.00
5	11	3	4.45	0.00	0.92	0.00	0.00
5	11	4	5.94	0.35	1.14	2.90	2.56
5	11	5	7.29	0.10	1.37	3.09	2.99
5	16	3	8.49	2.45	1.67	4.29	1.84
5	16	4	11.15	3.64	2.11	3.64	0.00
5	16	5	14.14	2.64	2.57	2.64	0.00
10	6	3	3.24	0.85	0.36	0.85	0.00
10	6	4	4.29	0.00	0.46	0.00	0.00
10	6	5	5.36	0.00	0.56	0.00	0.00
10	11	3	8.74	2.65	0.93	2.65	0.00
10	11	4	11.53	0.88	1.16	4.49	3.61
10	11	5	14.67	4.15	1.43	7.91	3.76
10	16	3	16.64	0.74	1.74	2.00	1.26
10	16	4	22.16	0.00	2.21	5.12	5.11
10	16	5	27.50	3.77	2.64	3.77	0.00
Average			9.58	1.33	1.26	2.50	1.17
Maximum			27.50	4.15	2.64	7.91	5.11

$b_u - b_l$ and $M_{is} = \bar{v}_i^s - \underline{v}_i^s,\ i \in \mathcal{H}, s \in \mathcal{S}$ as described in Section 4.3.1, and Gurobi solver is provided with a initial policy, which corresponds to the best policy obtained using Algorithm 2 and the MV problem.

- *BnC-mixing*: A branch-and-cut algorithm is applied to solve the relaxed master problem (4.12) with additional constraints (4.9) by iteratively adding feasibility cuts (4.20) violated by any incumbent solution encountered in the branch-and-cut tree.

In addition, we compare the results of the formulations above with the results of formulation *DEF-TS* obtained by adding the threshold structure constraints (4.22) to *DEF-base*.

In Tables 4.4-4.5, we report the best objective function value ("Obj.") and optimality gap ("Gap") achieved within the time limit, and the total solution time in seconds ("Time") for each setting. The optimality gap values are computed as $100 \times \frac{ub-lb}{lb}$, where ub and lb correspond to the best achieved upper and lower bounds on the optimal objective function value, respectively. Each row corresponds to the average of two replications and each blue dagger (†) in column "Gap" indicates a replication with positive optimality gap obtained within the time limit. For the decomposition algorithm *BnC-mixing*, we additionally report the average time in seconds spent solving a single subproblem (4.18) and a separation problem checking for violated optimality cuts under the column "AvSubTim" and the number of times the subproblems are invoked by an incumbent solution under the column "# Sub".

We first focus on different implementations of the DEF. Table 4.4 compares the base case *DEF-base*, *DEF-P* with preprocessing, *DEF-PMI* with preprocessing, stronger big-M terms and initial solutions, and *DEF-TS* with additional threshold structure constraints in terms of computational times and optimality gaps. Note that *DEF-TS* is not guaranteed to provide an optimal solution because it is a restriction of the original problem. Hence, for *DEF-TS*, we also report the gap between the objective function value of *DEF-TS* ($objT$) and the best objective function value obtained by the other three settings ($objB$) as $100 \times \frac{(objT - objB)}{objB}$ under the column "dGap".

Table 4.4: Comparison of different implementations of the DEF

			DEF-base			*DEF-P*			*DEF-PMI*			*DEF-TS*			
$\lvert\mathcal{S}\rvert$	$\lvert\mathcal{H}\rvert$	$\lvert\mathcal{A}\rvert$	Obj.	Gap	Time	Obj.	Gap	Time	Obj.	Gap	Time	Obj.	dGap	Gap	Time
5	6	3	1323266.89	0.00	0.34	1323266.89	0.00	0.27	1323266.89	0.00	0.16	1323266.89	0.00	0.00	0.14
5	6	4	1275652.01	0.00	2.02	1275652.00	0.00	2.92	1275652.01	0.00	1.95	1275652.01	0.00	0.00	0.13
5	6	5	1275652.01	0.00	5.98	1275652.01	0.00	3.85	1275652.01	0.00	4.18	1275652.01	0.00	0.00	0.20
5	11	3	630088.05	0.00	117.17	630088.05	0.00	53.64	630088.05	0.00	32.48	630088.05	0.00	0.00	0.16
5	11	4	496600.75	35.54†	3592.18	496600.75	0.00	638.26	496600.75	0.00	867.14	496600.75	0.00	0.00	0.34
5	11	5	475023.83	93.37††	3600.06	475023.83	0.03†	1966.05	475176.51	0.06†	1800.03	475023.83	0.00	0.00	1.15
5	16	3	415590.55	88.85††	3600.08	415590.55	1.36††	3600.06	415590.55	1.36††	3600.03	415590.55	0.00	0.00	0.38
5	16	4	256166.41	94.99††	3600.05	256108.06	0.55††	3600.04	265309.78	4.04††	3600.04	256108.06	0.00	0.00	1.92
5	16	5	237153.12	97.91††	3600.08	237153.12	0.00	74.90	243500.16	2.51†	1800.06	237153.12	0.00	0.00	4.68
10	6	3	1435969.99	0.00	0.95	1435969.99	0.00	0.50	1435969.99	0.00	0.58	1435969.99	0.00	0.00	0.13
10	6	4	1312221.11	0.00	7.48	1312221.11	0.00	3.63	1312221.10	0.00	3.06	1312221.11	0.00	0.00	0.22
10	6	5	1312221.11	0.00	47.84	1312221.11	0.00	8.49	1312221.11	0.00	7.36	1312221.11	0.00	0.00	0.27
10	11	3	779098.65	0.00	396.45	779098.65	0.00	44.20	779098.65	0.00	68.10	779098.65	0.00	0.00	0.29
10	11	4	507803.55	86.35††	3600.05	507803.55	2.64†	2742.55	507803.55	0.00	2866.96	507803.55	0.00	0.00	0.94
10	11	5	478713.67	96.16††	3600.06	478713.67	4.16††	3600.04	478713.67	4.16††	3600.03	478713.67	0.00	0.00	2.47
10	16	3	504524.22	95.36††	3600.06	521417.24	3.32†	3254.54	508058.45	0.74†	1800.04	504524.22	0.00	0.00	0.55
10	16	4	323605.12	98.24††	3600.06	306013.43	6.80††	3600.06	304779.96	6.43†	1800.05	304779.96	0.00	0.00	1.76
10	16	5	257883.51	98.09††	3600.09	257883.51	5.16†	1918.50	265974.01	8.30†	1800.07	257883.51	0.00	0.00	14.55

The results in Table 4.4 show that incorporating the bounds obtained by preprocessing significantly improves solution times and optimality gaps. *DEF-P* succeeds to solve 24 out of 36 instances to optimality within the time limit while *DEF-base* can only solve 17 instances. For the instances that cannot be solved to optimality within the time limit, *DEF-P* provides reasonable optimality gaps ($< 7\%$), whereas *DEF-base* outputs large optimality gap values ($> 35\%$). As it can be seen under column "*DEF-PMI*", setting stronger big-M terms and providing an initial solution to the solver decreases the number of instances that cannot be solved to optimality from 12 to 11, however there is no strict decrease in the solution times or optimality gaps. On the other hand, addition of threshold structure constraint provides significant improvements in the solution times and optimality gaps. *DEF-TS* can solve all instances in Table 4.4 within 15 seconds. It also achieves the best objective function value obtained by the three other implementations of DEF for all instances. These results clearly demonstrate the importance of incorporating additional information on the desired structure of the optimal policy into the problem formulation.

Next, we provide preliminary results on the performance of the proposed decomposition algorithm *BnC-mixing*. The results in Table 4.5 clearly shows that the DEF formulation *DEF-P* outperforms the decomposition algorithm *BnC-mixing* in terms of computational efficiency. The *BnC-mixing* terminates with a positive optimality gap for 11 out of 18 instances, while *DEF-P* can solve all but five instances to optimality. The *BnC-mixing* spends a considerable amount of time checking for violations in the optimality cuts due to the large number of times the subproblem is invoked by an incumbent solution. This drawback may be resulting from the big-M term in the feasibility cuts (4.20), which makes the inequality effective only for the policy in the corresponding incumbent solution.

4.7 Conclusions

In this chapter, we investigate the risk associated with parameter uncertainty in MDPs. We formulate the problem with the objective of minimizing the VaR and explore the structure of the optimal policies. Assuming a discrete representation of uncertainty and determin-

Table 4.5: Comparison of *DEF-base* and the decomposition algorithm *BnC-mixing*

			DEF-P			*BnC-mixing*				
$\lvert\mathcal{S}\rvert$	$\lvert\mathcal{H}\rvert$	$\lvert\mathcal{A}\rvert$	Obj.	Gap	Time	Obj.	Gap	Time	AvSubTim	# Sub
5	6	3	1323266.89	0.00	0.27	1323266.89	0.00	6.03	0.0153	388.50
5	6	4	1275652.00	0.00	2.92	1275652.01	0.00	65.63	0.0153	4112.00
5	6	5	1275652.01	0.00	3.85	1275652.01	0.00	317.59	0.0176	15635.50
5	11	3	630088.05	0.00	53.64	630088.05	0.16†	3362.96	0.0517	49026.00
5	11	4	496600.75	0.00	638.26	590510.48	16.74††	3601.94	0.0520	48382.00
5	11	5	475023.83	0.03†	1966.05	510692.51	6.91††	3601.97	0.0512	40829.00
5	16	3	415590.55	1.36††	3600.06	465373.31	13.60††	3602.23	0.1033	28575.00
5	16	4	256108.06	0.55††	3600.04	399552.20	35.53††	3601.17	0.1065	29741.00
5	16	5	237153.12	0.00	74.90	472340.65	48.68††	3600.83	0.1059	30034.00

istic policies, we provide an MIP formulation, propose reprocessing methods and heuristic algorithms, and devise a branch-and-cut algorithm. The proposed modeling approach and solution algorithms are tested on an inventory management problem in the long term humanitarian relief operations context.

The contributions in this chapter can be improved in several directions. First, as also indicated by our computational studies, the mathematical models and algorithms suffer from the size of the action space, often referred to as the curse of dimensionality. Hence it is worthwhile to search for action elimination methods that could improve the computational efficiency of our solution methods. Furthermore, the strength of the feasibility cut (4.19) and other possible alternatives can be explored.

Chapter 5

CONTRIBUTIONS AND FUTURE DIRECTIONS

This chapter summarizes the contributions of this thesis and elaborates on relevant research questions that can be further explored. In this thesis, we work on risk-averse stochastic programming problems encountered in the areas of multicriteria and sequential decision-making based on the risk measures VaR and CVaR. We propose mathematical formulations and devise state-of-art solution algorithms using various mixed-integer programming methods including valid inequalities, decomposition algorithms and stochastic dynamic programming.

Chapters 2 and 3 aim to incorporate risk-aversion into multicriteria decision making. In many real-life applications, decision makers need to consider the risk arising from multiple sources. For example, in the context of humanitarian relief operations, one may be interested in minimizing the risk associated with the operational costs, supply shortages, response speed, and equal distribution of supplies simultaneously. These applications bring out the need for appropriate multivariate definitions of risk measures and comparison relations between random vectors. A majority of the studies on multivariate risk extend the existing univariate definitions and relations to the multivariate case by using a scalarization scheme that requires identification of the relative importance of each criterion. However, in certain cases, it may not be possible to obtain a unique representation of the relative weights of criteria with respect to each other, as they may be subject to a group of decision makers with conflicting opinions.

Motivated by the possible ambiguities in the relative importance of criteria, in Chapter 2, we propose a multivariate definition of CVaR, referred as VMCVaR, as a set of vectors. The new definition addresses certain drawbacks of the existing vector-valued definitions of multivariate CVaR, which mostly consider continuous probability distributions, and provides an alternative that can be used for both continuous and discrete distribution cases. We show that the VMCVaR satisfies the properties of an analogous definition of coherence for random vectors except subadditivity, and provide multiobjective optimization formulations.

Remark 2 demonstrates that the VMCVaR does not satisfy a certain analogue of the subadditivity property enforced for all corresponding pLEPs. This result implies that there may be some element in VMCVaR, where merging two sources of risk actually increases the risk with respect to at least one criterion. As a result, it would be interesting to consider relaxed variants of this subadditivity property and investigate if there always exists an element in VMCVaR that makes merging two sources of risk preferable in terms of all criteria, or if merging two sources of risk is preferable in terms of at least one criterion for all elements in VMCVaR.

Considering the challenges in the computation of VMCVaR, in Chapter 3, we study two-stage stochastic programs with multivariate risk constraints utilizing a scalarization scheme. The aim is to ensure that the random outcome vector of interest associated with the optimal solution is prefereble to a specified benchmark random outcome vector with respect to the multivariate polyhedral CVaR or a multivariate stochastic order relation. We provide a mixed-integer programming formulation and two delayed cut generation algorithms that can be used for both variants. The proposed modeling approach and solution methods are demonstrated on a pre-disaster relief network design problem.

The delayed cut generation algorithm with scenario decomposition proposed in Section 3.2.2 is only applicable for the cases in which all criteria of interest can be linearly represented. This assumption is mainly due to the fact that the Benders cuts are only effective when the subproblems satisfy the property of strong duality. To this end, it may be worthwhile to explore ways of extending our scenario decomposition algorithm for a more general set of performance criteria.

Chapter 4 focuses on an application of risk-averse optimization in the stochastic sequential decision making context. In this chapter, we study the risk, particularly VaR, associated with the expected performance of an MDP arising from the uncertainty in the cost and transition probability parameters. We formulate the problem as a mixed-integer program that aims to obtain a policy minimizing the VaR of the expected total discounted cost of an MDP at a specified confidence level with respect to parameter uncertainty, and propose techniques that can potentially improve computational efficiency of the formulation. In addition, we devise a branch-and-cut algorithm based on a scenario decomposition approach. The proposed methodology is applied to an inventory management problem for long-term humanitarian relief operations.

Our experiments indicate that the proposed formulation and solution algorithms suffer from the size of the solution space, which is constituted of the state, action and scenario spaces. Hence, applying scenario and policy elimination methods may be crucial for a better computational performance. In Section 4.3.1, we propose a scenario elimination technique that employs a relaxation of the original problem, which allows for different policies in each scenario. Our computational results demonstrate the effectiveness of the proposed scenario elimination technique and imply that our solution methods may benefit from further reductions in the solution space. Hence, a promising future research direction is to explore for scenario and action elimination methods that can be applied to MDPs in general. Furthermore, we observed that the optimal policies in our instances of the humanitarian inventory management problem has a threshold structure, i.e., the number of vehicles to be dispatched decreases with an increase in the inventory. It would be interesting to investigate if there always exists an optimal policy with threshold structure for the humanitarian inventory management problem.

Appendix A
ALTERNATIVE SOLUTION ALGORITHMS

A classical MDP minimizing the total expected discounted cost with known cost ($\tilde{\mathbf{c}}$) and transition probabilities ($\tilde{\mathbf{P}}$) can be stated as the LP

$$\max \quad \sum_{i \in \mathcal{H}} q_i v_i \tag{A.1a}$$

$$s.t. \quad v_i \leq \tilde{c}_i(a) + \gamma \sum_{j \in \mathcal{H}} \tilde{P}_{ij}(a) v_j, \quad i \in \mathcal{H},\ a \in \mathcal{A}, \tag{A.1b}$$

where q_i is the probability of starting in state $i \in \mathcal{H}$ and the future costs are discounted by $\gamma \in (0,1)$. In this formulation, the policy decisions are implicit. Based on the optimal solution of the LP, the optimal policy can be determined by selecting an action $a \in \mathcal{A}$ in each state $i \in \mathcal{H}$ that satisfies the associated constraints (A.1b) as equality. This property follows from the fact that the value function $\mathbf{v}$ for any deterministic optimal policy satisfies

$$v_i = \min_{a \in \mathcal{A}} \left\{ \tilde{c}_i(a) + \gamma \sum_{j \in \mathcal{H}} \tilde{P}_{ij}(a) v_j \right\}, \quad i \in \mathcal{H}, \tag{A.2}$$

which is concave because the minimum of piecewise-linear functions is concave.

However, the quantile optimization problem (4.5) cannot utilize such implicit representation of the policies due to two main reasons. First, the objective function in (4.5) requires minimization of a nonnegative weighted sum of $\mathbf{v}$, which is concave as stated in (A.2), and minimization of concave functions is hard, in general. Second, problems with an objective function maximizing a nonnegative weighted sum of $\mathbf{v}$ still require an explicit representation of the policies as binary variables $\mathbf{w}$, because the same policy should be imposed across all scenarios. Note that, in this case, the condition (A.2) does not necessarily hold for any independent scenario.

The deterministic equivalent form, denoted as DEF, (4.5) contains two sets of binary decisions: scenario selection variables $\mathbf{z}$ and policy selection variables $\mathbf{w}$. In Section 4.4, we propose a branch-and-cut algorithm (*BnC-mixing*) with scenario decomposition by assuming that both variables $\mathbf{z}$ and $\mathbf{w}$ are fixed in the first stage. The second-stage problems become LPs that can be solved in polynomial time. Nevertheless, our computational experiments show that *BnC-mixing* is outperformed by the DEF in terms of computational efficiency. This result is mainly due to the fact that fixing both $\mathbf{z}$ and $\mathbf{w}$ in the first stage oversimplifies the second-stage problems, whereas the first-stage problem stays difficult to solve. Considering the trade-off between the computational effort required to solve the first- and second-stage problems, here we propose alternative solution algorithms and evaluate their potential challenges.

Note that when either the scenario or policy selection variables are considered in the second stage, the second-stage problems do not decompose over scenarios anymore. Policy decisions ($\mathbf{w}$) are required to be the same across all scenarios, and the scenario variables ($\mathbf{z}$) are linked by constraint (4.5c). Additionally, we need to deal with MIP problems in the second stage.

First assume that only the policy decisions are made in the first stage. Then the first-stage problem becomes

$$\begin{aligned}
\min \quad & \theta \\
s.t. \quad & \sum_{a \in \mathcal{A}} w_{ia} = 1, \quad i \in \mathcal{H}, \\
& w_{ia} \in \{0,1\}, \quad i \in \mathcal{H}, a \in \mathcal{A}, \\
& (\theta, \mathbf{w}) \in \mathcal{C}_w,
\end{aligned}$$

where θ is an auxiliary variable representing the second-stage objective function value, and

$\mathcal{C}_w$ is the set of optimality cuts obtained by solving the second-stage problem

$$\min \quad y \tag{A.4a}$$

$$s.t. \quad \sum_{s\in\mathcal{S}} z^s p^s \geq \alpha, \tag{A.4b}$$

$$\sum_{i\in\mathcal{H}} q_i v_i^s \leq y + (1 - z^s)M, \tag{A.4c}$$

$$v_i^s \geq \tilde{c}_i^s(a) + \gamma \sum_{j\in\mathcal{H}} \tilde{P}_{ij}^s(a) v_j^s - (1 - \hat{w}_{ia})M_{is}, \quad i \in \mathcal{H},\ a \in \mathcal{A},\ s \in \mathcal{S}, \tag{A.4d}$$

$$z^s \in \{0, 1\}, \quad s \in \mathcal{S}, \tag{A.4e}$$

for a given first-stage solution $\hat{\mathbf{w}}$. Note that subproblem (A.4) is always feasible, and it can be efficiently solved even though it contains binary scenario selection variables. The optimal value of y for the incumbent solution $\hat{\mathbf{w}}$, denoted by $y_{\hat{\mathbf{w}}}$, can be obtained by computing the expected total discounted cost corresponding to each scenario $s \in \mathcal{S}$ for the policy given by $\hat{\mathbf{w}}$, and then calculating the VaR_α associated with this random variable. Note that computation of the VaR_α for a random variable with known finite distribution takes polynomial time.

Although subproblem (A.4) can be solved efficiently, we still need to incorporate the information obtained in the second stage into the first-stage problem. Here, the classical Benders cuts cannot be utilized because the strong duality property does not hold for general MIPs. Instead, we employ an alternative set of optimality cuts proposed by Laporte and Louveaux [63]

$$\theta \geq (y_{\hat{\mathbf{w}}} - L)\left(\sum_{(i,a):\hat{w}_{ia}=1} w_{ia} - \sum_{(i,a):\hat{w}_{ia}=0} w_{ia}\right) - (y_{\hat{\mathbf{w}}} - L)(|\mathcal{H}| - 1) + L, \tag{A.5}$$

where L is a lower bound on the second-stage objective function value. Using the preprocessing results presented in Section 4.3.1, we can set $L = b_l$. In this framework, for each incumbent solution $\hat{\mathbf{w}}$ obtained in the branch-and-cut tree, we solve subproblem (A.4) to check for violations in the associated optimality cut (A.5). Our preliminary computational results demonstrate that this solution algorithm performs worse than the *BnC-mixing* in terms of computational efficiency. Hence it is not included in the main text.

Another alternative modeling approach is to fix only the scenario selection variables $\mathbf{z}$ in the first stage. In this case, the first-stage problem becomes

$$\begin{aligned} \min \ & \theta \\ s.t. \ & \sum_{s \in \mathcal{S}} z^s p^s \geq \alpha, \\ & z^s \in \{0,1\}, \quad s \in \mathcal{S}, \\ & (\theta, \mathbf{z}) \in \mathcal{C}_z, \end{aligned}$$

where θ is an auxiliary variable representing the second-stage objective function value, and $\mathcal{C}_z$ is the set of optimality cuts obtained by solving the second-stage problem

$$\min \ y \tag{A.7a}$$

$$s.t. \ \sum_{a \in \mathcal{A}} w_{ia} = 1, \quad i \in \mathcal{H}, \tag{A.7b}$$

$$\sum_{i \in \mathcal{H}} q_i v_i^s \leq y + (1 - \hat{z}^s) M, \quad s \in \mathcal{S}, \tag{A.7c}$$

$$v_i^s \geq \tilde{c}_i^s(a) + \gamma \sum_{j \in \mathcal{H}} \tilde{P}_{ij}^s(a) v_j^s - (1 - w_{ia}) M_{is}, \quad i \in \mathcal{H}, \ a \in \mathcal{A}, \ s \in \mathcal{S}, \tag{A.7d}$$

$$w_{ia} \in \{0,1\}, \quad i \in \mathcal{H}, \ a \in \mathcal{A}, \tag{A.7e}$$

for an incumbent scenario selection variable $\hat{\mathbf{z}}$. As previously stated in Section 4.3.2, subproblem (A.7) for incumbent solution $\hat{\mathbf{z}}$ can be equivalently stated as

$$\text{RMDP}(\hat{\mathbf{z}}): \quad \min_{\boldsymbol{\pi} \in \Pi_D} \max_{s \in \mathcal{S}(\hat{\mathbf{z}})} C(\boldsymbol{\pi}, \tilde{\mathbf{c}}^s, \tilde{\mathbf{P}}^s),$$

where $\mathcal{S}(\hat{\mathbf{z}}) := \{s \in \mathcal{S} | \ \hat{z}^s = 1\}$, and $C(\boldsymbol{\pi}, \tilde{\mathbf{c}}^s, \tilde{\mathbf{P}}^s)$ is the expected total discounted cost of policy $\boldsymbol{\pi}$ under scenario s. The RMDP($\hat{\mathbf{z}}$) can be considered as a robust MDP with stationary cost and transition probability uncertainty, i.e., a single realization of parameters apply for all states and time periods. There are studies in the literature that propose value and policy iteration algorithms for robust stationary MDPs under transition probability matrix uncertainty making certain independency assumptions on states and actions (see [57, 78]). However, for our uncertainty model, we have a randomly generated non-trivial counterexample with 21 scenarios, eight states and five actions at $\gamma = 0.90$ that proves the suboptimality of the policies obtained by these algorithms even for the case of uncertainty only in the transition matrix. Hence, subproblem (A.7) needs to be solved as an MIP problem for each incumbent solution $\hat{\mathbf{z}}$, and consequently fixing only the scenario decisions in the first-stage problem does not lead to an efficient exact solution algorithm.

BIBLIOGRAPHY

[1] A. Abdi and R. Fukasawa. On the mixing set with a knapsack constraint. *Mathematical Programming*, 157(1):191–217, 2016.

[2] T. Adrian and M.K. Brunnermeier. CoVaR. *The American Economic Review*, 106(7): 1705–1741, 2016.

[3] Y. Adulyasak, P. Varakantham, A. Ahmed, and P. Jaillet. Solving uncertain MDPs with objectives that are separable over instantiations of model uncertainty. In *AAAI*, pages 3454–3460, 2015.

[4] S. Ahmed. Convexity and decomposition of mean-risk stochastic programs. *Mathematical Programming*, 106(3):433–446, 2006. ISSN 0025-5610.

[5] E. Altman. *Constrained Markov decision processes*, volume 7. CRC Press, 1999.

[6] P. Artzner, F. Delbaen, J.M. Eber, and D. Heath. Coherent measures of risk. *Mathematical Finance*, 9(3):203–228, 1999.

[7] J. Bagnell, A.Y. Ng, and J. Schneider. Solving uncertain Markov decision problems. *Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, Tech. Rep. CMU-RI-TR-01-25*, 2001.

[8] B. Balcik and B.M. Beamon. Facility location in humanitarian relief. *International Journal of Logistics: Research and Applications*, 11(2):101–121, 2008.

[9] R. Bellman. *Dynamic Programming.* Courier Corporation, 2013.

[10] J.F. Benders. Partitioning procedures for solving mixed-variables programming problems. *Numerische Mathematik*, 4(1):238–252, 1962.

[11] D.F. Benoit and D. Van den Poel. Benefits of quantile regression for the analysis of customer lifetime value in a contractual setting: An application in financial services. *Expert Systems with Applications*, 36(7):10475–10484, 2009.

[12] J.R. Birge and F.V. Louveaux. A multicut algorithm for two-stage stochastic linear programs. *European Journal of Operational Research*, 34(3):384–392, 1988.

[13] J.R. Birge and F.V. Louveaux. *Introduction to Stochastic Programming*. Springer, New York, 1997.

[14] M. Bodur and J. Luedtke. Mixed-integer rounding enhanced Benders decomposition for multiclass service-system staffing and scheduling with arrival rate uncertainty. *Management Science*, 63(7):2073–2091, 2016.

[15] J.F. Bonnans and A. Shapiro. *Perturbation Analysis of Optimization Problems*. Springer Series in Operations Research. Springer-Verlag, New York, 2000.

[16] R.J. Boucherie and N.M. van Dijk. *Markov Decision Processes in Practice*. Springer, 2017.

[17] G.C. Calafiore and M.C. Campi. The scenario approach to robust control design. *IEEE Transactions on Automatic Control*, 51(5):742–753, 2006.

[18] M. Çelik, O. Ergun, B. Johnson, P. Keskinocak, A. Lorca, P. Pekgün, and J. Swann. Humanitarian logistics. In *INFORMS Tutorials in Operations Research*, pages 18–49. INFORMS, Maryland, 2012.

[19] K. Chen and M. Bowling. Tractable objectives for robust policy optimization. In *Advances in Neural Information Processing Systems*, pages 2069–2077, 2012.

[20] R.C. Chen and G.L. Blankenship. Dynamic programming equations for constrained stochastic control. In *Proceedings of the 2002 American Control Conference*, volume 3, pages 2014–2022. IEEE, 2002.

[21] R.A. Collado, D. Papp, and A. Ruszczyński. Scenario decomposition of risk-averse multistage stochastic programming problems. *Annals of Operations Research*, 200(1): 147–170, 2012.

[22] A. Cousin and E. Di Bernardino. On multivariate extensions of value-at-risk. *Journal of Multivariate Analysis*, 119:32–46, 2013.

[23] A. Cousin and E. Di Bernardino. On multivariate extensions of conditional-tail-expectation. *Insurance: Mathematics and Economics*, 55:272–282, 2014.

[24] I. Deák. Multidimensional integration and stochastic programming. *Numerical techniques for stochastic optimization*, pages 187–200, 1988.

[25] I. Deák. Subroutines for computing normal probabilities of setscomputer experiences. *Annals of Operations Research*, 100(1-4):103–122, 2000.

[26] G. DeCandia, D. Hastorun, M. Jampani, G. Kakulapati, A. Lakshman, A. Pilchin, S. Sivasubramanian, P. Vosshall, and W. Vogels. Dynamo: Amazon's highly available key-value store. *ACM SIGOPS Operating Systems Review*, 41(6):205–220, 2007.

[27] E. Delage and S. Mannor. Percentile optimization for Markov decision processes with parameter uncertainty. *Operations Research*, 58(1):203–213, 2010.

[28] D. Dentcheva and G. Martinez. Two-stage stochastic optimization problems with stochastic ordering constraints on the recourse. *European Journal of Operational Research*, 219(1):1 – 8, 2012.

[29] D. Dentcheva and A. Ruszczyński. Optimization with stochastic dominance constraints. *SIAM Journal on Optimization*, 14(2):548–566, 2003.

[30] D. Dentcheva and A. Ruszczyński. Optimality and duality theory for stochastic optimization problems with nonlinear dominance constraints. *Mathematical Programming*, 99(2):329–350, 2004.

[31] D. Dentcheva and A. Ruszczyński. Portfolio optimization with stochastic dominance constraints. *Journal of Banking & Finance*, 30(2):433–451, 2006.

[32] D. Dentcheva and A. Ruszczyński. Optimization with multivariate stochastic dominance constraints. *Mathematical Programming*, 117(1):111–127, 2009.

[33] D. Dentcheva and E. Wolfhagen. Optimization with multivariate stochastic dominance constraints. *SIAM Journal on Optimization*, 25(1):564–588, 2015.

[34] D. Dentcheva and E. Wolfhagen. Two-stage optimization problems with multivariate stochastic order constraints. *Mathematics of Operations Research*, 41(1):1–22, 2016.

[35] D. Dentcheva, A. Prékopa, and A. Ruszczyński. Concavity and efficient points of discrete distributions in probabilistic programming. *Mathematical Programming*, 89 (1):55–77, 2000.

[36] E. Di Bernardino, J.M. Fernández-Ponce, F. Palacios-Rodríguez, and M.R. Rodríguez-Griñolo. On multivariate extensions of the conditional value-at-risk measure. *Insurance: Mathematics and Economics*, 61:1–16, 2015.

[37] A. Döyen, N. Aras, and G. Barbarosoğlu. A two-echelon stochastic facility location model for humanitarian relief logistics. *Optimization Letters*, 6(6):1123–1145, 2012. ISSN 1862-4472.

[38] J. Dupačová. Risk objectives in two-stage stochastic programming models. *Kybernetika*, 44(2):227–242, 2008.

[39] Ö. Elçi and N. Noyan. A chance-constrained two-stage stochastic programming model for humanitarian relief network design. *Transportation Research Part B: Methodological*, 108:55–83, 2018.

[40] C.I. Fábián. Handling CVaR objectives and constraints in two-stage stochastic models. *European Journal of Operational Research*, 191(3):888 – 911, 2008.

[41] C.I. Fábián and A. Veszprémi. Algorithms for handling CVaR-constraints in dynamic stochastic programming models with applications to finance. *Journal of Risk*, 10(3): 111–131, 2008.

[42] Federal Emergency Management Agency (FEMA). `https://egateway.fema.gov/ESF6/DRCLocator#`, Accessed: 2017-12-08.

[43] G.O. Ferreira, E.F. Arruda, and L.G. Marujo. Inventory management of perishable items in long-term humanitarian operations using Markov decision processes. *International Journal of Disaster Risk Reduction*, 31:460–469, 2018.

[44] A. Genz. Numerical computation of multivariate normal probabilities. *Journal of Computational and Graphical Statistics*, 1(2):141–149, 1992.

[45] A. Genz and F. Bretz. *Computation of Multivariate Normal and t Probabilities*, volume 195. Springer Science & Business Media, 2009.

[46] R. Givan, S. Leach, and T. Dean. Bounded-parameter Markov decision processes. *Artificial Intelligence*, 122:71109, 2000.

[47] R. Gollmer, F. Neise, and R. Schultz. Stochastic programs with first-order dominance constraints induced by mixed-integer linear recourse. *SIAM Journal on Optimization*, 19(2):552–571, 2008.

[48] R. Gollmer, U. Gotzes, and R. Schultz. A note on second-order stochastic dominance constraints induced by mixed-integer linear recourse. *Mathematical Programming*, 216 (1):179–190, 2011.

[49] W.J. Gutjahr and P.C. Nolz. Multicriteria optimization in humanitarian aid. *European Journal of Operational Research*, 252(2):351 – 366, 2016.

[50] A.H. Hamel, B. Rudloff, and M. Yankova. Set-valued average value at risk and its computation. *Mathematics and Financial Economics*, 7(2):229–246, 2013.

[51] C.C. Heyde, S.G. Kou, and X.H. Peng. What is a good risk measure: Bridging the gaps between data, coherent risk measures, and insurance risk measures. *Preprint, Columbia University*, 2006.

[52] T. Homem-de-Mello and S. Mehrotra. A cutting surface method for uncertain linear programs with linear stochastic dominance constraints. *SIAM Journal on Optimization*, 20(3):1250–1273, 2009.

[53] X. Hong, M.A. Lejeune, and N. Noyan. Stochastic network design for disaster preparedness. *IIE Transactions*, 47(4):329–357, 2015.

[54] J. Hu and S. Mehrotra. Robust and stochastically weighted multiobjective optimization models and reformulations. *Operations Research*, 60(4):936–953, 2012.

[55] J. Hu, T. Homem-de Mello, and S. Mehrotra. Sample average approximation of stochastic dominance constrained programs. *Mathematical Programming*, 133(1-2):171–201, 2012.

[56] M. Huang, K. Smilowitz, and B. Balcik. Models for relief routing: Equity, efficiency and efficacy. *Transportation Research Part E*, 48(1):2–18, 2012.

[57] G.N. Iyengar. Robust dynamic programming. *Mathematics of Operations Research*, 30(2):257–280, 2005.

[58] S. Kalyanasundaram, E.K.P. Chong, and N.B. Shroff. Markov decision processes with uncertain transition rates: Sensitivity and robust control. In *Proceedings of the 41st IEEE Conference on Decision and Control*, volume 4, pages 3799–3804. IEEE, 2002.

[59] S.G. Kou, X.H. Peng, and C.C. Heyde. External risk measures and Basel accords. *Mathematics of Operations Research*, 38(3):393–417, 2013.

[60] S. Küçükyavuz. On mixing sets arising in chance-constrained programming. *Mathematical Programming*, 132(1-2):31–56, 2012.

[61] S. Küçükyavuz and N. Noyan. Cut generation for optimization problems with multivariate risk constraints. *Mathematical Programming*, 159(1-2):165–199, 2016.

[62] S. Kusuoka. On law invariant coherent risk measures. *Advances in Mathematical Economics*, 3:83–95, 2001.

[63] G. Laporte and F.V. Louveaux. The integer L-shaped method for stochastic integer programs with complete recourse. *Operations Research Letters*, 13(3):133–142, 1993.

[64] J. Lee and A. Prékopa. Properties and calculation of multivariate risk measures: MVaR and MCVaR. *Annals of Operations Research*, 211(1):225–254, 2013.

[65] F. Liberatore, C. Pizarro, C. S. de Blas, M.T. Ortuño, and B. Vitoriano. Uncertainty in humanitarian logistics for disaster management. A review. In *Decision Aid Models for Disaster Management and Emergencies*, pages 45–74. Springer, 2013.

[66] X. Liu, S. Küçükyavuz, and J. Luedtke. Decomposition algorithms for two-stage chance-constrained programs. *Mathematical Programming*, 157(1):219–243, 2016.

[67] X. Liu, F. Kılınç-Karzan, and S. Küçükyavuz. On intersection of two mixing sets with applications to joint chance-constrained programs. *Mathematical Programming*, pages 1–40, 2017.

[68] X. Liu, S. Küçükyavuz, and N. Noyan. Robust multicriteria risk-averse stochastic programming models. *Annals of Operations Research*, 259(1):259–294, 2017.

[69] J. Luedtke. A branch-and-cut decomposition algorithm for solving chance-constrained mathematical programs with finite support. *Mathematical Programming*, 146(1-2): 219–244, 2014.

[70] J. Luedtke and S. Ahmed. A sample approximation approach for optimization with probabilistic constraints. *SIAM Journal on Optimization*, 19(2):674–699, 2008.

[71] J. Luedtke, S. Ahmed, and G.L. Nemhauser. An integer programming approach for linear programs with probabilistic constraints. *Mathematical Programming*, 122(2): 247–272, 2010.

[72] S. Mannor, D. Simester, P. Sun, and J.N. Tsitsiklis. Bias and variance approximation in value function estimates. *Management Science*, 53(2):308–322, 2007.

[73] S. Mannor, O. Mebel, and H. Xu. Robust MDPs with k-rectangular uncertainty. *Mathematics of Operations Research*, 41(4):1484–1509, 2016.

[74] G.P. McCormick. Computability of global solutions to factorable nonconvex programs: Part I - Convex underestimating problems. *Mathematical Programming*, 10(1):147–175, 1976.

[75] M. Meraklı and S. Küçükyavuz. Vector-valued multivariate conditional value-at-risk. *Operations Research Letters*, 46(3):300–305, 2018.

[76] N. Miller and A. Ruszczyński. Risk-averse two-stage stochastic linear programming: Modeling and decomposition. *Operations Research*, 59(1):125–132, 2011.

[77] Y. Nie, X. Wu, and T. Homem-de-Mello. Optimal path problems with second-order stochastic dominance constraints. *Networks and Spatial Economics*, 12(4):561–587, 2012.

[78] A. Nilim and L. El Ghaoui. Robust control of Markov decision processes with uncertain transition matrices. *Operations Research*, 53(5):780–798, 2005.

[79] N. Noyan. Alternate risk measures for emergency medical service system design. *Annals of Operations Research*, 181(1):559–589, 2010.

[80] N. Noyan. Risk-averse two-stage stochastic programming with an application to disaster management. *Computers & Operations Research*, 39(3):541–559, 2012.

[81] N. Noyan and G. Rudolf. Optimization with multivariate conditional value-at-risk constraints. *Operations Research*, 61(4):990–1013, 2013.

[82] N. Noyan and G. Rudolf. Optimization with stochastic preferences based on a general class of scalarization functions. *Operations Research*, 66(2):463–486, 2018.

[83] N. Noyan, B. Balcik, and S. Atakan. A stochastic optimization model for designing last mile relief networks. *Transportation Science*, 50(3):1092–1113, 2016.

[84] N. Noyan, M. Meraklı, and S. Küçükyavuz. Two-stage stochastic programming under multivariate risk constraints with an application to humanitarian relief network design. *Submitted for publication*, Optimization Online `http://www.optimization-online.org/DB_FILE/2017/01/5804.pdf`, 2017.

[85] T. Osogami. Robustness and risk-sensitivity in Markov decision processes. In *Advances in Neural Information Processing Systems*, pages 233–241, 2012.

[86] G.C. Pflug. Some remarks on the value-at-risk and the conditional value-at-risk. In *Probabilistic Constrained Optimization: Methodology and Applications*. Kluwer Academic Publishers, Dordrecht, 2000.

[87] A. Prékopa. Dual method for the solution of a one-stage stochastic programming problem with random RHS obeying a discrete probability distribution. *Zeitschrift für Operations Research*, 34(6):441–461, 1990.

[88] A. Prékopa. Multivariate value at risk and related topics. *Annals of Operations Research*, 193(1):49–69, 2012.

[89] A. Prékopa and J. Lee. Risk tomography. *European Journal of Operational Research*, 265(1):149–168, 2018.

[90] M.L. Puterman. *Markov decision processes: Discrete stochastic dynamic programming.* John Wiley & Sons, 2014.

[91] C.G. Rawls and M.A. Turnquist. Pre-positioning of emergency supplies for disaster response. *Transportation Research Part B: Methodological*, 44(4):521–534, 2010.

[92] C.G. Rawls and M.A. Turnquist. Pre-positioning planning for emergency response with service quality constraints. *OR Spectrum*, 33(3):481–198, 2011.

[93] R.T. Rockafellar and S. Uryasev. Optimization of conditional value-at-risk. *Journal of Risk*, 2:21–42, 2000.

[94] R.T. Rockafellar and S. Uryasev. Conditional value-at-risk for general loss distributions. *Journal of Banking & Finance*, 26(7):1443–1471, 2002.

[95] B. Rottkemper, K. Fischer, and A. Blecken. A transshipment model for distribution and inventory relocation under uncertainty in humanitarian operations. *Socio-Economic Planning Sciences*, 46(1):98–109, 2012.

[96] L. Rüschendorf. Risk measures for portfolio vectors. In *Mathematical Risk Analysis*, pages 167–188. Springer Series in Operations Research and Financial Engineering, 2013.

[97] A. Ruszczyński. Risk-averse dynamic programming for Markov decision processes. *Mathematical Programming*, 125(2):235–261, 2010.

[98] A. Ruszczyński. Advances in risk-averse optimization. *INFORMS Tutorials in Operations Research*, pages 168–190, 2013.

[99] A. Ruszczyński and A. Shapiro. *Stochastic Programming, Handbooks in Operations Research and Management Science Volume 10*. Elsevier, Amsterdam, 2003.

[100] A. Ruszczyński and A. Shapiro. Risk averse optimization. In *Lectures on Stochastic Programming: Modeling and Theory*, pages 253–332. SIAM, 2009.

[101] T.L. Saaty. *Decision Making for Leaders the Analytic Hierarchy Process for Decisions in a Complex World*. RWS publications, 1990.

[102] J. Salmerón and A. Apte. Stochastic optimization for natural disaster asset prepositioning. *Production and Operations Management*, 19(5):561–574, 2010.

[103] J.K. Satia and R.E. Lave Jr. Markovian decision processes with uncertain transition probabilities. *Operations Research*, 21(3):728–740, 1973.

[104] R. Schultz and S. Tiedemann. Conditional value-at-risk in stochastic programs with mixed-integer recourse. *Mathematical Programming*, 105(2-3):365–386, 2006.

[105] A. Shapiro. Minimax and risk averse multistage stochastic programming. *European Journal of Operational Research*, 219(3):719–726, 2012.

[106] A. Shapiro and A. Kleywegt. Minimax analysis of stochastic problems. *Optimization Methods and Software*, 17(3):523–542, 2002.

[107] A. Shapiro, D. Dentcheva, and A. Ruszczyński. *Lectures on stochastic programming: modeling and theory*. SIAM, 2009.

[108] S. Sinha and A. Ghate. Policy iteration for robust nonstationary Markov decision processes. *Optimization Letters*, 10(8):1613–1628, 2016.

[109] L.N. Steimle, D.L. Kaufman, and B.T. Denton. Multi-model Markov decision processes: A new method for mitigating parameter ambiguity. Optimization Online `http://www.optimization-online.org/DB_FILE/2018/01/6434.pdf`, 2018.

[110] A. Tewari and P.L. Bartlett. Bounded parameter Markov decision processes with average reward criterion. In *International Conference on Computational Learning Theory*, pages 263–277. Springer, 2007.

[111] R. Torres, R.E. Lillo, and H. Laniado. A directional multivariate value at risk. *Insurance: Mathematics and Economics*, 65:111–123, 2015.

[112] R.M. Van Slyke and R. Wets. L-shaped linear programs with applications to optimal control and stochastic programming. *SIAM Journal of Applied Mathematics*, 17(4): 638–663, 1969.

[113] B. Vitoriano, M.T. Ortuño, G. Tirado, and J. Montero. A multi-criteria optimization model for humanitarian aid distribution. *Journal of Global Optimization*, 51(2):189–208, 2011.

[114] C.C. White III and H.K. Eldeib. Markov decision processes with imprecise transition probabilities. *Operations Research*, 42(4):739–749, 1994.

[115] W. Wiesemann, D. Kuhn, and B. Rustem. Robust Markov decision processes. *Mathematics of Operations Research*, 38(1):153–183, 2013.

[116] H. Xu and S. Mannor. The robustness-performance tradeoff in Markov decision processes. In *Advances in Neural Information Processing Systems*, pages 1537–1544, 2007.

[117] H. Xu and S. Mannor. Parametric regret in uncertain Markov decision processes. In *Proceedings of the 48th IEEE Conference on Decision and Control*, pages 3606–3613. IEEE, 2009.

[118] H. Xu and S. Mannor. Distributionally robust Markov decision processes. *Mathematics of Operations Research*, 37(2):288–300, 2012.

[119] P. Yu and H. Xu. Distributionally robust counterpart in Markov decision processes. *IEEE Transactions on Automatic Control*, 61(9):2538–2543, 2016.

[120] M. Zhang, S. Küçükyavuz, and S. Goel. A branch-and-cut method for dynamic decision making under joint chance constraints. *Management Science*, 60(5):1317–1333, 2014.

[121] M. Zhao, K. Huang, and B. Zeng. A polyhedral study on chance constrained program with random right-hand side. *Mathematical Programming*, 166(1-2):19–64, 2017.