

Stochastic Programming Methods for Casualty Response
Planning Problems with Different Classes of Patients, Blood
Inventory Management, and Hospital Evacuations

Pedram Farghadani-Chaharsooghi

A Thesis
in
Mechanical, Industrial & Aerospace Engineering (MIAE)

Presented in Partial Fulfillment of the Requirements
For the Degree of
Doctor of Philosophy (Industrial Engineering) at
Concordia University
Montréal, Québec, Canada

September 2025

© Pedram Farghadani-Chaharsooghi, 2025

CONCORDIA UNIVERSITY
School of Graduate Studies

This is to certify that the thesis prepared

By: **Pedram Farghadani-Chaharsooghi**

Entitled: **Stochastic Programming Methods for Casualty Response Planning Problems with Different Classes of Patients, Blood Inventory Management, and Hospital Evacuations**

and submitted in partial fulfillment of the requirements for the degree of

Doctor of Philosophy (Industrial Engineering)

complies with the regulations of this University and meets the accepted standards with respect to originality and quality.

Signed by the final examining committee:

_____	Chair
<i>Dr. Xia Li</i>	
_____	External Examiner
<i>Dr. Borzou Rostami</i>	
_____	External to Program
<i>Dr. Arash Mohammadi</i>	
_____	Examiner
<i>Dr. Onur Kuzgunkaya</i>	
_____	Examiner
<i>Dr. Tommaso Schettini</i>	
_____	Thesis Supervisor
<i>Dr. Hossein Hashemi Doulabi</i>	
_____	Co-supervisors
<i>Dr. Michel Gendreau & Dr. Walter Rei</i>	

Approved by _____
Dr. Ramin Sedaghati, Graduate Program Director

September, 2025 _____
Dr. Mourad Debbabi, Dean
Gina Cody School of Engineering and Computer Science

Abstract

Stochastic Programming Methods for Casualty Response Planning Problems with Different Classes of Patients, Blood Inventory Management, and Hospital Evacuations

Pedram Farghadani-Chaharsooghi, Ph.D.
Concordia University, 2025

Casualty Response Planning (CRP) is a critical component of disaster management, where timely and coordinated medical intervention can greatly reduce mortality and improve outcomes for affected populations. Large-scale disasters often cause widespread damage to infrastructure, disrupt hospital operations, and generate uncertain and fluctuating demands across diverse patient groups. In these uncertain environments, effective planning requires advanced optimization techniques that can support both immediate and evolving decisions. In this thesis, we present a sequence of stochastic programming models and solution methods that address key operational challenges in casualty response, including the management of different classes of patients, blood inventory logistics, and hospital evacuation, with a focus on optimizing system-wide performance under uncertainty.

First, we present a two-stage stochastic mixed-integer programming model for CRP that jointly considers uncertain patient demand and hospital treatment capacity. In the first stage, we determine the location of Alternative Care Facilities (ACFs) and allocate resources such as rescue vehicles, physicians, and medical equipment. In the second stage, we assign patients with varying injury types to hospitals or ACFs and allow for transfers between facilities if needed. We also introduce an equivalent but more compact model that improves computational performance. To solve these models efficiently, we develop enhanced

decomposition algorithms, including an accelerated branch-and-Benders-cut method enriched with Lagrangian decomposition, multi-cut strategies, and Pareto-optimal cuts. We demonstrate the effectiveness of our approach through extensive computational experiments and a case study based on the 2011 Van earthquake in Turkey.

Second, we extend the CRP framework to a multi-stage stochastic setting that reflects the dynamic and evolving nature of post-disaster operations. We develop a model that incorporates both apheresis and whole blood collection methods, deploys portable apheresis machines, and enables platelet transshipments among medical facilities. We treat patient demands, blood donor availability, and treatment capacity as uncertain parameters that unfold over time. Our objective is to optimize timely surgeries and compatible platelet transfusions, improving patient outcomes. To solve this problem, we implement a stochastic dual dynamic programming algorithm and enhance it using randomized quasi-Monte Carlo sampling, K-means++ clustering for scenario reduction, and strong valid inequalities. We evaluate our method through computational experiments and demonstrate its real-world potential using the Van earthquake case study.

Third, we develop a two-stage stochastic model for hospital evacuation, an often overlooked aspect of disaster response. In the first stage, we make strategic decisions regarding the location of ACFs, allocation of land and aerial rescue vehicles, and designation of backup hospitals. In the second stage, after uncertainty is revealed, we schedule patient evacuations and transfers over time, aiming to minimize evacuation risk, transportation risk, and operational cost. To solve this complex and large-scale problem, we propose a deep-reinforcement-learning-based metaheuristic that integrates deep reinforcement learning with adaptive large neighborhood search. We preserve scenario diversity using decision-based clustering and improve solution quality using quasi-Monte Carlo sampling. Our algorithm delivers near-optimal solutions for small instances and scales efficiently to larger ones, outperforming benchmark methods such as Progressive Hedging.

Acknowledgments

In a world that moves at an ever-accelerating pace, where time often feels like a luxury, pursuing a PhD is both a challenge and a profound privilege. It is the rare opportunity to slow down, to think deeply, to search with intention, and sometimes, to discover. It is the privilege of learning, and with learning, gaining a little more freedom at each step. I started this PhD journey to find a kind of freedom that comes from learning and understanding. However, going through such a big experience and finishing it would not have been possible without the help of others. I want to thank these people with all my heart.

I must start with my research supervisors for accepting me and then supervising and supporting me throughout my doctorate. I am deeply grateful to Profs. Hossein Hashemi Doulabi, Walter Rei, and Michel Gendreau for their constant guidance, patience, and encouragement. Each of you played a unique and important role in shaping my work and helping me grow as a researcher. Your trust in me, your time, and your valuable feedback have meant a great deal, and I feel truly lucky to have had the chance to work with you. In addition, I would like to thank Profs. Borzou Rostami, Arash Mohammadi, Onur Kuzgunkaya, and Tommaso Schettini for accepting to join the jury.

I am also eternally grateful to my caring and supportive family: my encouraging mother, Azita; my motivating brother, Peiman; my affectionate sister, Saharnaz; and my lifelong inspiration, my father, Majid. I cannot express strongly enough how truly blessed I am. Last but by no means least, I owe so much to my lovely partner, Bahar, for standing by me through all the ups and downs. Her cheerfulness and energy are the driving force behind

everything I do.

A heartfelt thank you also goes to my dear friends who have been with me since my early years in Montreal — Kayhan, Negar, Arian, and Yasi — named here in no particular order. Your friendship has been a source of comfort, laughter, and strength throughout this journey. I’m also grateful to my officemates at Concordia (EV 14.101), my friends at CIRRELT, and fellow members of the CIRRELT Student Committee. Thank you for creating such a vibrant and supportive community, whether through “Happy Wednesdays,” coffee and cookie gatherings, or simply working side by side. You all helped make this experience so much more meaningful.

Furthermore, I gratefully acknowledge the Natural Sciences and Engineering Research Council of Canada (NSERC), the Canada Research Chair in Stochastic Optimization of Transport and Logistics Systems, and the Concordia University Doctoral Graduate Fellowship for providing full financial support during my PhD. I also express my sincere gratitude to CIRRELT for supporting my research through several scholarships.

Finally, thank you to everyone who has contributed to my academic journey, whether mentioned by name or not. Your encouragement, guidance, and presence are vital to this achievement.

Contribution of Authors

This dissertation is presented in a manuscript-based format and includes three articles that have either been published or are currently under review in academic journals. These manuscripts were submitted in the following chronological order:

The first article, titled “Stochastic Casualty Response Planning with Multiple Classes of Patients,” was published online in July 2025 in the journal *Computers & Operations Research*. It was also selected as a finalist in the Healthcare SIG Student Paper Competition by the Canadian Operational Research Society (CORS) and was presented at CORS 2025 in Edmonton. The second manuscript, “Stochastic Dual Dynamic Programming for Multi-Stage Casualty Response Planning with Platelet Inventory Management,” was submitted to the *INFORMS Journal on Computing* in July 2025. Finally, the third manuscript, “A Deep-Reinforcement Learning-Based Matheuristic for Stochastic Casualty Response Planning with Hospital Evacuation,” was submitted for publication to the *INFORMS Journal on Computing* in August 2025.

All three manuscripts were co-authored with Dr. Hossein Hashemi Doulabi, Dr. Walter Rei, and Dr. Michel Gendreau. The author of this thesis served as the principal researcher, responsible for tasks such as developing the models and algorithms, implementing the solution methods, analyzing computational results, and writing the initial drafts.

Contents

List of Figures	xii
List of Tables	xiii
1 Introduction	1
2 Stochastic Casualty Response Planning with Multiple Classes of Patients	6
2.1 Introduction	7
2.2 Problem Description	14
2.3 Primary Two-Stage Stochastic Mixed-Integer Programming Model (P1) . .	18
2.3.1 First-Stage Model	18
2.3.2 Second-Stage Model	20
2.4 Improved Two-Stage Stochastic Mixed-Integer Programming Model (P2) .	25
2.5 L-shaped Algorithm and Enhancements	28
2.5.1 Multiple Cuts and Warm-up Strategies	30
2.5.2 Lower Bounding Functional Valid Inequalities (LBF)	30
2.5.3 Strong (Pareto-optimal) Cuts (PO)	31
2.5.4 Benders Dual Approach	32
2.5.5 Branch-and-Benders-Cut Scheme	35
2.5.6 Heuristic Strategies for Strengthened Cut Generation	35
2.6 Computational Results	37
2.6.1 Instance Generation	37

2.6.2	Performance of Mathematical Models	39
2.6.3	Performance of Algorithms	39
2.6.4	Marginal Impacts of Accelerators	43
2.6.5	Impact of Casualty Shortage Penalty Parameters on Model Behavior	45
2.7	Case Study	47
2.7.1	Evaluating Out-of-Sample Performance and the Impact of Hospital Capacity Uncertainty	50
2.8	Conclusion	52

3 Stochastic Dual Dynamic Programming for Multi-Stage Casualty Response

	Planning with Platelet Inventory Management	54
3.1	Introduction	55
3.2	Problem Description	59
3.3	Multi-stage Stochastic Integer Programming Model	62
3.3.1	First-Stage Model	62
3.3.2	Subproblem of Stages $t \in \{2, \dots, T\}$	65
3.3.3	Last Stage Subproblem ($T + 1$)	68
3.4	Stochastic Dual Dynamic Programming (SDDP) and Enhancements	69
3.4.1	Advanced Scenario Sampling Method	72
3.4.2	Scenario Clustering Method using K-means++	72
3.4.3	Warm-up and Multiple Cuts Strategies	73
3.4.4	Strong (Pareto-optimal) Cuts	74
3.4.5	Lower Bounding Functional Valid Inequalities (LBFVIs)	74
3.5	Computational Results	75
3.5.1	Implementation Details	75
3.5.2	Performance on Out-of-Sample Scenarios	76
3.5.3	Performance of the Algorithms on In-Sample Scenarios	77
3.6	Case Study	80

3.7	Conclusion	81
4	A Deep-Reinforcement-Learning-Based Matheuristic for Stochastic Casualty Response Planning with Hospital Evacuation	83
4.1	Introduction	84
4.2	Problem Description	88
4.3	Two-Stage Stochastic Integer Programming Model	91
4.3.1	First-Stage Model	92
4.3.2	Second-Stage Model	93
4.4	Deep-Reinforcement-Learning-Based Matheuristic	98
4.4.1	Advanced Scenario Sampling	100
4.4.2	Decision-Based Scenario Clustering	101
4.4.3	Basic Adaptive Large Neighborhood Search (ALNS)	102
4.4.4	Deep Reinforcement Learning-Based ALNS (DRL-ALNS)	103
4.5	Computational Results	106
4.5.1	Implementation Details	108
4.5.2	Performance on Out-of-Sample Scenarios	109
4.5.3	Performance of the Algorithms on In-Sample Scenarios	111
4.6	Case Study	113
4.7	Conclusion	116
5	Conclusion	118
	Bibliography	122
	Appendix A Appendices for Paper 1	131
A.1	Proof of Theorem 1	131
A.2	Proof of Proposition 1	132
A.3	Flowchart of the BBC_I Algorithm	135

A.4	Progressive Hedging Algorithm (PHA)	135
A.5	Performance Analysis of Benders Decomposition Enhancements	138
A.6	L-Shaped Algorithm Performance	140
Appendix B Appendices for Paper 2		142
B.1	Hyperplane Formulation in the Backward Pass	142
B.2	Illustration of Scenario Trees	143
B.3	Illustration of Single and Multiple Cuts Strategies	143
B.4	Formulations for LBFVIs	144
B.5	Comparison of Lower Bounds for SDDP Variants	146
B.6	Case Study Details	147

List of Figures

1	Network flow of model P1	27
2	Network flow of model P2	27
3	Percentages of Instances Solved to Optimality Within Given Computing Times.	45
4	Impact of Shortage Penalty Parameter on Model Behavior.	46
5	Comparison of ACF Establishments in Van District for Deterministic and Stochastic Models.	48
6	Van Network: (a) Full layout, (b) StoExpRisk layout, (c) StoConsRisk Hos., (d) StoExpRisk Hos.. . . .	115
A.1	Flowchart of the BBC_I algorithm.	136
B.2	Symmetric Scenario Tree	143
B.3	Recombining Tree	143
B.4	Single Cut Strategy	144
B.5	Multiple Cut Strategy	144
B.6	(Color online) LB progression across different SDDP variants.	146
B.7	Van Network: (a) Full layout, (b) Det, (c) StoNoT, (d) StoT.	147

List of Tables

1	An Overview of Latest Related Works.	12
2	Parameters Used for Instance Generation.	38
3	Comparison of Mathematical Models P1 and P2.	40
4	Computational Results of the B&BC Algorithm for the CRP problem with Stochastic Demand and Hospital Treatment Capacity.	42
5	Reverse Marginal Impact.	44
6	In-Sample and Out-of-Sample Performance of BBC_I with Stochastic and Fixed Hospital Caps.	51
7	Parameters Used for Instance Generation	76
8	Out-of-Sample SDDP Performance with Various Scenario Generation and Clustering Methods.	78
9	In-Sample SDDP Performance with Various Enhancements.	79
10	Risk Parameter Settings for Threat and Transportation (taken from Bish et al. (2014))	109
11	Parameters Used for Instance Generation	109
12	Out-of-Sample ALNS Performance with Various Scenario Generation and Clustering Methods.	112
13	In-Sample Algorithm Performance.	113
A.5.1	Performance Comparison of Different Enhancement Combinations in the Algorithm.	139
A.2	Computational Results of the L-Shaped Algorithm.	141

Chapter 1

Introduction

Catastrophic health events (CHEs), whether natural disasters like earthquakes and tsunamis or man-made incidents such as terrorist attacks and chemical spills, pose significant challenges to healthcare systems. These crises result in large numbers of injuries within a short time frame, overwhelming hospitals and disrupting infrastructure. The urgency of medical attention in such scenarios cannot be overstated. For instance, the 2023 Turkey-Syria earthquakes caused over 56,000 deaths and left more than 100,000 people injured; the 2008 Sichuan earthquake led to 87,587 fatalities. In these situations, the window of opportunity for saving lives narrows dramatically. This highlights the critical role of casualty response planning (CRP), which is a coordinated effort across organizations to manage limited healthcare resources, triage and treat patients, and transport them to functioning facilities. Effective CRP can significantly increase survival rates during the crucial hours following a disaster.

One widely adopted approach in CRP is the deployment of Alternative Care Facilities, or ACFs, which are temporary medical stations established in open areas such as stadiums, schoolyards, or parks to supplement overwhelmed hospitals. ACFs play a crucial role in extending emergency care capacity. However, the success of ACFs depends heavily on appropriate location-allocation decisions. A poorly located facility can drastically reduce patient access and delay treatment. A notable example is the 1995 Oklahoma City bombing,

where a triage center was placed too far from the attack site, preventing medical personnel from assessing nearly half of the victims (Larson et al., 2006). Such cases underscore the importance of strategic ACF deployment to ensure that care reaches those in need as quickly as possible (Caunhye and Nie, 2018).

Another often-overlooked complexity in CRP is that disaster victims frequently sustain multiple injuries, each requiring specialized care. For example, after the Wenchuan earthquake, 45.8% of patients had two injuries, and nearly 10% had three or more (Lu-Ping et al., 2012). Similar patterns were observed in the 2009 L'Aquila and 2013 Lushan earthquakes (Del Papa et al., 2019; Kang et al., 2015). Despite the prevalence of multiple injuries, most CRP models treat patients as a homogeneous group or classify them into only two priority levels. This simplification can lead to suboptimal outcomes, such as assigning patients to ACFs that lack the required specialists. Chapter 2 addresses this gap by proposing a CRP model that classifies patients by injury type and matches them with ACFs accordingly, ensuring more accurate and efficient care allocation.

This classification becomes even more critical when considering resource-intensive treatments such as surgeries, which often require reliable access to blood products, particularly platelets. Platelets are among the most essential yet logistically challenging resources in post-disaster medical care due to their short five-day shelf life. After the 2008 Wenchuan earthquake, over 30% of collected blood products were not distributed effectively (Ma et al., 2019), underscoring the need for better blood inventory management. Two primary methods exist for platelet collection: whole blood collection (WBC), which involves time-consuming processing, and apheresis blood donation (ABD), which allows for same-day platelet use. New portable apheresis machines, such as the *Trima Accel 7*, can now be deployed directly at ACFs to address urgent platelet demand (Hosseini et al., 2020; Xu and Szmerekovsky, 2023; Samani and Hosseini-Motlagh, 2023). Chapter 3 focuses on optimizing platelet collection and distribution during disasters, including the allocation of portable apheresis machines and planning for transshipment between medical

facilities.

The coordination of such complex resource and treatment logistics becomes even more challenging when hospitals themselves must be evacuated, which is an aspect often overlooked in traditional CRP frameworks. Emergency plans frequently assume that hospitals will remain operational, yet many healthcare facilities are located in hazard-prone areas. For example, it is estimated that 50–70% of healthcare facilities in Latin America and the Caribbean are situated in high-risk zones (Bagaria et al., 2009). When hospitals are compromised, integrated evacuation planning becomes crucial. However, guidelines on evacuation prioritization vary widely. Some recommend evacuating the most critical patients first, while others suggest starting with those who can be moved quickly (Childers and Taaffe, 2010). In 1997, following a chemical explosion near Helena Regional Medical Center, the lack of a plan for low-priority patients resulted in makeshift care at a community college cafeteria (Squillace, 2010). These challenges are addressed in Chapter 4, which extends the CRP framework to incorporate hospital evacuation and proposes practical strategies to guide prioritization and resource reallocation under such circumstances.

Finally, effective CRP must account for numerous operational constraints, such as matching physicians with patient needs and allocating renewable (staff, vehicles) and non-renewable (medical supplies) resources. For instance, orthopedic patients cannot be treated in ACFs lacking the appropriate specialists or equipment like Kirschner wires. Meanwhile, the role of *rescue vehicles* is paramount in both evacuating injured individuals and delivering medical personnel and supplies. A well-planned allocation strategy ensures that emergency resources reach the right locations on time, enhancing the effectiveness of the entire response effort. These practical concerns are deeply examined across Chapters 2, 3, and 4, each of which tackles different but interconnected aspects of casualty response under uncertainty.

The contributions of this dissertation can be summarized as follows:

- **Problem Modeling:** This dissertation advances the field of CRP by developing novel

stochastic optimization models that address real-world complexities:

- **Chapter 2:** For the first time, we investigate a stochastic CRP model that considers multiple patient classes with varying injury types. The model also includes the assignment of appropriate physician types, non-renewable medical resources, and rescue vehicles to ACFs based on each patient’s specific injuries.
 - **Chapter 3:** We explore, for the first time, a stochastic CRP model that integrates ACF establishment with platelet inventory management using the latest portable apheresis technologies. The model matches platelet blood types and ages to casualties, assigns rescue vehicles to ACFs, and enables transshipment between medical facilities (MFs).
 - **Chapter 4:** We propose the first stochastic CRP model that jointly considers ACF establishment and hospital evacuation planning. The framework also incorporates the allocation of land and aerial rescue vehicles to ACFs under uncertainty.
- **Algorithmic Development:** Each of the proposed models is paired with advanced algorithms designed to efficiently solve the stochastic formulations and improve solution quality:
 - **Chapter 2:** We develop an L-shaped algorithm and a branch-and-Benders-cut (B&BC) algorithm for solving the two-stage stochastic CRP model. These algorithms are significantly enhanced through Benders dual decomposition (BDD), Lagrangian dual decomposition (LDD), a multi-cut reformulation, Pareto-optimal (PO) cuts, and lower bounding functional valid inequalities (LBFVIs), resulting in improved computational performance.
 - **Chapter 3:** We design Stochastic Dual Dynamic Programming (SDDP) and Nested Benders Decomposition (NBD) algorithms to address the multi-stage

CRP model. Key improvements include the use of strong cuts, warm-up strategies, multiple cuts, randomized quasi-Monte Carlo (RQMC) sampling, LBFVIs, and K-means++ scenario clustering to enhance both convergence speed and robustness.

- **Chapter 4:** We propose a deep-reinforcement-learning-based matheuristic that combines reinforcement learning with an adaptive large neighborhood search (ALNS) algorithm. The method is further refined using RQMC sampling and a decision-based scenario clustering approach, making it suitable for large-scale evacuation-integrated CRP problems.
- **Managerial Insights:** In addition to methodological advancements, this dissertation delivers important managerial insights for real-world disaster response planning. Extensive computational experiments were conducted to evaluate the performance of the proposed algorithms across in-sample and out-of-sample scenarios. The results highlight the effectiveness of algorithmic enhancements and provide evidence-based guidance on managing patient triage, resource allocation, and facility planning. Furthermore, the proposed models and solution methods were validated through a real-world case study inspired by a large-scale earthquake in Turkey, demonstrating their practical applicability and strategic value in guiding emergency medical decisions under uncertainty.

Chapter 2

Stochastic Casualty Response Planning with Multiple Classes of Patients

Abstract

In this paper, we study the stochastic casualty response planning problem (CRP) in the context of providing treatments to multiple classes of patients with different types of injuries. In this general setting, both patients' demands and hospitals' treatment capacity are considered uncertain. To the best of our knowledge, this is the first time that this problem is solved. We propose a novel two-stage stochastic mixed-integer programming model which, in the first stage, determines the location of the Alternative Care Facilities (ACFs) and allocates different resources, such as rescue vehicles, medical equipment, and physicians, to them. In the second stage, this model helps decide how to allocate patients with multiple injuries to either ACFs or hospitals, considering their care itineraries and available resources. Moreover, it recommends potential patient transfers between ACFs and hospitals when required. Furthermore, we introduce an alternative two-stage stochastic model that is more compact than the first. This formulation significantly reduces solution times. We also provide an equivalency proof between the two formulations. As the solution method, we develop both the L-shaped algorithm, a pure cutting-plane method tailored to our stochastic mathematical model, and the branch-and-Benders-cut (B&BC) algorithm. To further enhance the efficiency of these algorithms, we develop a wide range of acceleration techniques, including Benders dual decomposition, Lagrangian dual decomposition, a multi-cut reformulation, Pareto-optimal cuts, and the inclusion of lower bounding functional valid inequalities. We carry out extensive computational experiments demonstrating that these algorithmic enhancements dramatically improve the performance of the B&BC

algorithm, reducing the average optimality gap from 7898% in the standard B&BC algorithm to just 0.92% in the enhanced version. Additionally, we benchmark our approach against the progressive hedging algorithm (PHA), a widely used decomposition method in disaster response operations, to further assess its effectiveness. Finally, we present a case study from the 2011 Van earthquake in Turkey, demonstrating the applicability and efficiency of our optimization methods.

2.1 Introduction

Catastrophic health events (CHEs) are natural or man-made disasters that result in large numbers of injured individuals and put healthcare systems under pressure to take emergency actions. After CHEs, injured people require immediate medical attention, and their life often depends on a quick reaction from emergency healthcare providers. As a result, efficient resource management in the initial few hours following a disaster is crucial to enhance the casualty survival rate. Casualty response planning refers to the process of collaboration between organizations to efficiently manage limited resources to help survivors in affected areas and transfer them to hospitals. There have been numerous devastating natural disasters in the past few decades, which have motivated many researchers to work on casualty response planning problems. Notably, Turkey's 2023 earthquake claimed 59,259 lives, Haiti's 2010 earthquake resulted in a staggering 250,000 deaths, and the 2004 South Asia tsunami took 227,898 lives, marking these events as some of history's most devastating natural disasters.

In the context of CRP, using Alternative Care Facilities (ACFs) is one of the most common ways to diagnose an overwhelming number of casualties and treat them after a disaster. ACFs are set up as temporary healthcare facilities with the primary objective of delivering essential care during times of heightened demand, such as in the aftermath of a natural disaster. Stadiums, parks, universities, and schools are some instances of locations where ACFs can be established (Caunhye and Nie, 2018). The advantages of using ACFs will not realize without a suitable location-allocation plan. For example, during the 1995 Oklahoma City bombing, emergency medical personnel set up an ACF in the form of a triage center. However, they could not assess almost half of the victims because the triage center was initially located too far from the assault site (Larson et al., 2006).

The literature on CRP can be categorized into three main categories: deterministic, stochastic, or robust approaches. Deterministic models, while being widely used, have a notable drawback. They may yield infeasible plans or lead to unexpected penalty costs

due to ignoring the uncertainty of demands and hospitals' treatment capacity. Stochastic models address these limitations by incorporating uncertainty into the planning through scenarios, acknowledging that the full scope of disaster needs becomes clear over time.

Such needs are addressed by first-stage and recourse decisions in the stochastic optimization paradigm. These decisions typically evaluate needs using a combination of pre- and post-disaster data, including governmental records, technology outputs like satellite imagery, and on-the-ground reports from key informants, such as community leaders (Grass et al., 2023). Key informants play a crucial role in estimating casualty demand by providing informed assessments based on their expertise, observations, and available data sources. These key informants, who may include medical experts, emergency responders, and disaster assessment teams, do not attempt to predict specific injuries but rather classify casualties into broad patient categories requiring different treatment types. This classification allows for a structured medical resource allocation strategy that aligns with expected treatment needs. Since different disaster types result in varying casualty distributions, practitioners rely on structured scenario-generation processes that incorporate multiple sources of information to construct a realistic representation of potential demand.

In real-world disaster response planning, scenario generation typically relies on historical data, field assessments, and expert evaluations. Historical data provides baseline distributions of casualties from past events, but given the uniqueness of each disaster, updated assessments become crucial. Ground teams and emergency response personnel commonly refine these estimates using various assessment methods, including aerial observations, satellite imagery, and direct communication with community-based informants (Clingerman, 2019). In practice, assessments often involve multiple key informants from the affected communities who independently provide estimates of casualties, injury types, and the severity of impact. For example, in the aftermath of the 2004 tsunami, assessment teams in Indonesia relied on information from multiple informants within communities to verify and update casualty and damage estimates (Clingerman, 2019). Given that no single informant's evaluation is likely to be entirely accurate, decision-makers aggregate these independent assessments to generate a representative range of plausible scenarios. Consequently, the scenarios used in stochastic optimization reflect combined expert perspectives, historical trends, and observed real-world conditions, ensuring robustness in disaster planning decisions.

Practitioners who handle casualty response planning often work with scenarios in this way. In some cases, they simply rely on deterministic models using average estimates from multiple sources, which may lead to overly rigid plans that do not account for the true

variability of demand. While averaging assessments can provide a simplified representation, it inherently discards crucial information about uncertainty, potentially leading to inefficient or infeasible response strategies. Stochastic optimization preserves this uncertainty by explicitly incorporating scenario variations, allowing decision-makers to develop more resilient and adaptive plans. Instead of committing to a single forecast, a stochastic model ensures that resource allocation decisions remain flexible and capable of adapting to different demand realizations.

Another approach to addressing uncertainty is robust optimization, which focuses on finding the best possible decision under the worst-case scenario. Choosing between stochastic optimization and robust optimization depends largely on the confidence level in the available data. If field assessments provide sufficient and reliable information, stochastic optimization offers a balanced approach that accounts for a range of possible outcomes. However, when field data is scarce or unreliable, robust optimization provides a more conservative planning strategy that prepares for the most extreme scenarios. Given our assumption that key informants provide credible assessments backed by multiple sources, we adopt a scenario-based two-stage stochastic approach to optimize decision-making under uncertainty. This methodology ensures that casualty response planning remains realistic, data-driven, and adaptable to the complexities of post-disaster environments. While we briefly review the major works in this area, we refer interested readers to the comprehensive surveys by Farahani et al. (2020) and Dönmez et al. (2021) for a more in-depth literature review.

In the first category, many researchers assumed that all parameters are deterministic and given (e.g., see Apte et al. (2015); Caunhye et al. (2015); Liu et al. (2019); Setiawan et al. (2019)). Apte et al. (2015) developed a mixed-integer programming model to identify the best locations for casualty collection points and to allocate vehicles, resources, and injured people to them, in order to maximize the number of survivors. Caunhye et al. (2015) offered a location-allocation optimization model to minimize the total transportation time, which simultaneously locates ACFs, prioritizes casualties, and considers self-evacuees effects. Setiawan et al. (2019) formulated three mathematical programming models to take into account different levels of coordination for the distribution of relief resources and the allocation of vehicles to minimize the number of victims. Liu et al. (2019) considered ambulances and helicopters to move casualties. Their model determines the optimal locations for temporary healthcare services and recommends the allocation of medical services to injured people.

On the other hand, many researchers in the literature have proposed stochastic optimization (Caunhye and Nie, 2018; Alizadeh et al., 2019; Oksuz and Satoglu, 2020; Chang et al., 2023) and robust optimization models (Hu et al., 2019a; Li et al., 2020; Sun et al., 2021; Yin et al., 2024) to deal with uncertainty in catastrophic health events. Caunhye and Nie (2018) proposed a three-stage stochastic programming model which determines the location of ACFs in the first stage by considering self-evacuees' movement. Then, in the second and third stages, the model addresses the assignment of casualties to triage and treatment, respectively. Alizadeh et al. (2019) proposed a two-stage robust stochastic optimization model with uncertain levels of injuries. They demonstrated that incorporating uncertainty into strategic-level decision-making leads to better design outcomes for the casualty collection points. Oksuz and Satoglu (2020) developed a two-stage stochastic programming model for locating ACFs in the event of an earthquake. Their stochastic model aimed to minimize establishing ACFs and casualties' transportation costs. They incorporated innovative elements into their model by accounting for triage and α -reliability constraints. α -reliability constraints are employed to guarantee a predefined level of reliability (α) within the solutions. The reliability concept ensures that the solution meets all demands in scenarios that encompass at least 100α percent of all possible outcomes. Most recently, Chang et al. (2023) developed a strategy to determine the locations of ACFs and efficiently allocate the limited resources of emergency medical services (EMS) for quick casualty transportation to suitable hospitals, thereby enhancing casualty survival rates. They introduced a novel two-stage simulation-optimization algorithm to tackle the ACF location and EMS resource allocation problem, which combines binary and integer variables. The primary objective was to minimize the expected time for complete casualty delivery from the disaster area to hospitals.

In the context of robust optimization, Hu et al. (2019a) paid attention to diurnal population shifts in urban areas and developed a multi-objective robust optimization model. They proposed an accelerated Benders decomposition (BD) algorithm to solve large-size problems. While all previous works covered only the primary disaster, Li et al. (2020) took secondary shocks into account by offering a scenario-based robust programming model. It is worth mentioning that they assumed all patients have the same medical needs and belong to the same category. However, CRPs are significantly more complex since patients have multiple types of injuries resulting in different categories of patients with needs for different surgical specialties. Sun et al. (2021) introduced a bi-objective robust optimization model for strategic and operational response within a three-level rescue chain comprising casualty clusters, temporary facilities, and general hospitals. Their approach incorporated

uncertainties related to relief supply distribution, demands, and transportation times. To address this model, they used Lagrangian relaxation and ϵ -constraint techniques. Yin et al. (2024) presented a distributionally robust model for complex disaster response, involving facility location, supply inventory, and evacuation planning. Their solution utilized a branch-and-Benders-cut (B&BC) algorithm.

A summary of the CRP literature and the contributions of this study are presented in Table 1. As shown, our work goes beyond integrating multiple well-established elements, such as facility location, resource allocation, casualty allocation, and the use of heterogeneous rescue vehicles. A key novelty of our study is the introduction of care itineraries through the consideration of multiple injuries, i.e., different patient classes. By incorporating care itineraries, we recognize the need for physician specialty assignment — another aspect that, to our knowledge, has not been explicitly addressed in previous CRP models. These contributions enhance the realism of post-disaster planning by accounting for diverse medical needs while also introducing an additional combinatorial dimension to the problem.

Various injuries after CHEs can significantly impact the health and survival of individuals affected by the disaster. For example, earthquakes can cause a variety of injuries, including fractures, lacerations, nerve injuries, and burns (Del Papa et al., 2019). These injuries can be life-threatening and may require immediate medical attention. Additionally, the combination of different injuries can increase the risk of infection, blood loss, and other complications. Lu-Ping et al. (2012) reported that after the Wenchuan earthquake in 2008, 45.8%, 9%, and 1.5% of injured people had two, three, and more than four injuries, respectively. Del Papa et al. (2019) found that 52% of people injured in the 2009 L’Aquila earthquake had multiple injuries. In the 2013 Lushan earthquake in China, this number was 55.7% (Kang et al., 2015). Although patients commonly sustain multiple injuries following disasters, the studies previously reviewed do not consider patients with multiple injuries post-disaster. This is an important point to consider in casualty response planning since it can significantly affect the allocation of medical resources and patients to ACFs. Ignoring the variety of injuries and classifying patients only as a high or low priority can result in sending them to ACFs that are not appropriately equipped for any of their specific injuries.

In addition to multiple injuries, there are several important operational details that have yet to be addressed in the literature. For example, different patients have different needs for renewable and non-renewable resources. Therefore, we have to plan ahead for assigning physicians to ACFs accordingly. For instance, we cannot send a patient with a fracture injury to an ACF that does not have an orthopedic doctor. On the other hand, non-renewable

Table 1: An Overview of Latest Related Works.

Authors	FL	RA	VA	CA	PSA	ML-Inj	MI-PC	HV	Uncertainty type	UA	Objective function	Solution approach
Apte et al. (2015)	✓	✓	✓	✓					Deterministic model	N/A	Minimizing transportation time	MIP
Caunhye et al. (2015)	✓			✓		✓			Deterministic model	N/A	Minimizing transportation time	MIP
Caunhye and Nie (2018)	✓			✓		✓			Number of casualties Proportion of self-evacuees Transportation times	SP	Minimizing transportation time	Benders decomposition
Liu et al. (2019)	✓			✓		✓			Deterministic model	N/A	Maximizing expected survivals Minimizing total operational cost	ϵ -constraint method
Setiawan et al. (2019)	✓	✓	✓	✓				✓	Deterministic model	N/A	Minimizing unmet demand and victim suffering	Heuristic algorithm
Alizadeh et al. (2019)	✓			✓		✓			Number of casualties Transportation capacity	HRS	Minimizing transportation cost and mortality	Sample average approximation
Hu et al. (2019a)	✓			✓					Disaster occurrence time	RO	Minimizing total operational cost and maximum travel distance	Benders decomposition
Li et al. (2020)	✓	✓		✓					Disaster severity Transportation time Number of casualties	HRS	Minimizing setup, operational and evacuation costs	Progressive hedging algorithm
Oksuz and Satoglu (2020)	✓			✓		✓			Earthquake magnitude	SP	Minimizing setup and transportation costs	MIP
Sun et al. (2021)	✓	✓		✓		✓			Number of casualties Number of rescue supplies Transportation time	RO	Minimizing injury severity score and total operational cost	ϵ -constraint method
Chang et al. (2023)	✓		✓						Triage and treatment time Transportation time	SP	Minimizing expected casualty delivery time	Simulation optimization
Yin et al. (2024)	✓	✓		✓					Number of casualties	RO	Minimizing setup, transportation and inventory costs	Benders decomposition
This Study	✓	✓	✓	✓	✓	✓	✓	✓	Number of casualties Hospital treatment capacity	SP	Minimizing setup, transportation, operational, and unmet demand costs	Accelerated B&BC algorithm

FL: Facility location; RA: Resource allocation; VA: Vehicle allocation; CA: Casualty allocation; PSA: Physician specialty assignment; ML-Inj: Multi-level injuries; MI-PC: Multiple injuries (patient classes);

HV: Heterogeneous vehicle; UA: Uncertainty approach; N/A: Not applicable; SP: Stochastic programming; HRS: Hybrid robust and stochastic method; RO: Robust optimization; MIP: Mixed-integer programming model without a specific solution approach.

resources such as medical supplies and drugs also play an important role in the functioning of hospitals and ACFs. For example, Kirschner wires (K-wires) are small, thin, metal pins commonly used in orthopedic surgeries. However, their demand may spike during an emergency, and it is impossible to perform orthopedic surgery without them. This highlights the importance of proper resource allocation in CRPs. Additionally, rescue vehicles are crucial for evacuating people from hazardous locations and transporting emergency personnel and supplies. A suitable rescue vehicle allocation after a disaster helps to ensure that emergency medical services personnel and supplies reach the affected population as quickly and efficiently as possible. It can also improve the effectiveness of the emergency response and help save more lives.

Although there is a rich literature on CRPs, to the best of our knowledge, no paper has addressed patients with multiple types of injuries and the above mentioned features simultaneously. The main contributions of this research are as follows.

- For the first time, we investigate the stochastic casualty response planning problem concerning treating multiple patient classes with varying injury types. Our approach also involves assigning the appropriate type of physicians, non-renewable resources, and rescue vehicles to ACFs based on each patient's specific injuries.
- We propose a primary two-stage stochastic mixed-integer programming model to formulate this problem. Then, we offer an alternative two-stage stochastic formulation, which is more compact than the first. As a result, the second model is significantly more efficient. We also present a proof establishing the equivalence between the two formulations.
- We develop an L-shaped algorithm and a B&BC algorithm to solve our more efficient stochastic programming model. We significantly improve the algorithms by using the Benders dual decomposition (BDD), Lagrangian dual decomposition (LDD), a multi-cut reformulation, Pareto-optimal (PO) cuts, and the inclusion of lower bounding functional valid inequalities (LBF). These enhancements improve the overall performance of the proposed algorithms.
- We present extensive computational results for the CRP when dealing with multiple injuries for each patient. In our computational experiments, we compare the proposed two-stage stochastic programming models and evaluate the performance of the proposed decomposition algorithms and the effect of the developed enhancements.

- To assess the effectiveness of our proposed approach, we benchmark it against the PHA, a widely used decomposition method in disaster response optimization, ensuring a meaningful comparison in terms of solution quality and computational efficiency.
- We also evaluate the efficiency of our best model and algorithm in an earthquake case study from Turkey.

The structure of this paper is organized as follows. In Section 2.2, we describe the problem setting. Then, we present the original and improved two-stage stochastic programming models in Sections 2.3 and 2.4, respectively. Section 2.5 details the proposed L-shaped algorithm and explains the various enhancement techniques we use to improve the algorithm’s performance. Our extensive computational results on the CRP are presented in Section 2.6. In Section 2.7, we present a case study and evaluate out-of-sample performance, as well as the impact of hospital capacity uncertainty. Finally, we draw conclusions and discuss future directions in Section 2.8.

2.2 Problem Description

Generally, disaster management includes four main phases: mitigation, preparedness, response, and recovery. Mitigation involves understanding vulnerabilities to potential hazards and implementing protective measures to reduce risk and increase resilience. Preparedness involves evaluating plans to save lives and coordinate response efforts before a disaster strikes. The primary goal is to achieve an adequate state of preparedness for emergency response through initiatives that strengthen the technical and managerial capabilities of governments, organizations, and communities. The response phase aims to provide immediate assistance to victims, distribute relief supplies, and evacuate affected populations to safe areas. The recovery phase occurs after the emergency and focuses primarily on activities such as debris removal, rebuilding damaged structures, and restoring critical infrastructure. Mitigation and preparedness occur before a tragedy, while response and recovery occur thereafter (Farahani et al., 2020).

The problem that we study in this paper deals with the response phase, which occurs immediately post-disaster. These problems are generally approached through static modeling, which optimizes location and resource allocation decisions while incorporating uncertainties (Caunhye and Nie, 2018; Liu et al., 2023). Static models provide solutions for a specific point in time, making them well-suited for rapid decision-making in post-disaster

scenarios (Dalal and Üster, 2018; Bayram and Yaman, 2018; Wang and Nie, 2022). For example, Dalal and Üster (2018) explored a static humanitarian response network that integrates relief supply and evacuation, solving it using a BD algorithm. Similarly, Bayram and Yaman (2018) developed a two-stage stochastic static evacuation planning model that accounts for uncertainty in both demand and travel time. Their approach optimally determines shelter locations and assigns evacuees while ensuring that travel distances remain within an acceptable limit. In another study, Wang and Nie (2022) examined a static location–allocation problem that incorporates traffic congestion effects into relief supply planning, formulating a two-stage mixed-integer stochastic programming model and solving it with a generalized BD algorithm. While some studies have explored dynamic approaches that account for evolving demand and uncertainty over time, static models remain widely used for their ability to provide effective solutions in time-sensitive disaster response settings.

In our study, we adopt a static approach to model the post-disaster response phase, ensuring effective medical care supply planning within a predetermined time horizon. The time horizon for the decision-making environment begins once the CHEs commence and continues until all affected individuals have been successfully evacuated. The nature of the disaster determines the duration of this time frame. Our study assumes that the critical period, which includes the time during which victims are transported to hospitals or ACFs, is constrained within this predetermined time horizon, which is set at 7 days. Therefore, our primary objective is to design the appropriate medical care supply to meet the overall demand observed during the post-disaster response phase. In this problem, we make three important decisions after a CHE that includes:

1. Determining the location of ACFs,
2. Allocating renewable and non-renewable medical resources such as physicians, rescue vehicles, medications, and medical equipment to ACFs, and
3. Transporting casualties to ACFs and hospitals.

The first two sets of decisions are made immediately after the disaster based on the initial information received about the intensity and geographical dispersion of the disaster. Therefore, we have considered these two sets of decisions in the first stage of our proposed two-stage stochastic programming models. However, the third set of decisions, i.e., the transportation of casualties to ACFs and hospitals, is made after obtaining more information about uncertain parameters, including *(1) the number of casualties with different*

injuries in various areas and (2) the active treatment capacities of hospitals, which may differ from their nominal values due to infrastructure damage and operational disruptions. For instance, after the Great East Japan Earthquake in 2011, Fukushima Prefectural Ono Hospital suffered infrastructure damage, initially losing gas, power, and water supplies. This forced the staff to consolidate all patients into a less affected area while continuing care under emergency conditions (Sawano et al., 2023). Similarly, the 2005 Kashmir earthquake led to the complete destruction of 291 healthcare facilities, with an additional 74 partially damaged, severely impacting hospital capacity in the region (Helminen et al., 2006). The 1999 Taiwan earthquake further exemplifies this issue: Christian Hospital experienced a 10% capacity reduction, Veterans Hospital lost 50% of its treatment capacity, and Shiu-Tuan Hospital was completely closed, leaving the largest medical facility in the area unavailable to earthquake victims (Lee and Loh, 1999). Given such real-world instances of hospital infrastructure failures, incorporating uncertainty in treatment capacities is crucial for effective casualty response planning.

We assume that emergency responders will perform a triage and primary survey in the disaster field. This primary survey follows a well-established protocol, represented by the mnemonic A.B.C.D.E., which stands for Airway, Breathing, Circulation, Disability, and Environment. The survey is designed to swiftly identify and address life-threatening conditions while determining the necessary type and extent of medical interventions. For instance, consider a scenario where a victim is alert, able to speak, oriented, and demonstrating movement in all extremities. In such cases, we can confidently conclude that their airway is clear, oxygen is reaching the brain adequately, and there is no significant central neurological injury (Sever et al., 2012). Please note that while we briefly discuss the triage and primary survey here, a detailed examination of this process is outside the scope of this article. Our primary focus is on outlining the care itinerary for high-priority and low-priority casualties within the healthcare network.

We categorize patients into two distinct groups: high-priority and low-priority casualties (Caunhye and Nie, 2018). Each patient group follows a different care itinerary within the healthcare network. High-priority patients are those with breathing problems, significant bleeding, or mental issues that prevent them from following commands. They are prioritized for immediate transportation to hospitals, where they are considered fully treated upon arrival. We assume that hospitals have limited rescue vehicles and treatment capacities but are fully equipped with the necessary renewable and non-renewable resources to treat patients with diverse injuries.

Low-priority patients, however, have a different itinerary. They may need to go through

more than one ACF before reaching the status of “fully treated.” Although Caunhye and Nie (2018) assumed that low-priority patients are only treatable at ACFs, we recognize the need for a more adaptable and flexible approach to disaster response. As such, we have considered a refinement to their model, allowing for low-priority patients to initially receive care at ACFs while retaining the option to refer them to either another ACF or a hospital if their condition necessitates further treatment that cannot be provided at the initial ACF. This adjustment addresses the critical issue of hospital over-occupation, particularly in the immediate aftermath of a catastrophic event, which becomes even more pronounced when accounting for casualties who independently seek treatment at hospitals without assistance from emergency responders. Their actions can disrupt response planning and strain valuable resources that could have been allocated to more severe cases (Runge and Buddemeier, 2009). To implement this, first, we transfer low-priority patients to ACFs for the appropriate care and stabilization of their conditions. Their status as fully treated is granted only when the necessary physicians and resources have addressed all of their injuries. This assumption introduces complexity, as ACFs may not always have the essential physicians and medical resources available. In such cases, patients may receive partial treatment at one ACF before being transferred to another ACF or hospital for further care.

Effective management of casualty flows throughout the network is a central aspect of our problem. This includes making decisions about the assignment of patients to specific facilities based on their facility care itinerary. This intricate process involves coordinating the movement of patients to ensure timely access to the necessary care and resources, even when ACFs may lack certain required physicians and medical supplies. Our research focuses on optimizing these facility care itineraries, addressing the diverse nature of patient conditions and resource availability.

In addition to the previously defined patient classification, we may have patients with multiple injuries, further underscoring the need for a flexible approach to care. Therefore, we may need more than one type of physician. Regarding the allocation of physicians to ACFs in the second decision set, we mainly need four types of physicians in order to cover all common injuries after disasters: (i) Orthopedic physicians for fractures, crushes, and other similar injuries, (ii) Wound specialists for lacerations, contusions, and open wounds; (iii) Neurologists for nerve injuries, and (iv) Plastic physicians or dermatologists for burns (Tanaka et al., 1999; Doocy et al., 2013). Physicians require a certain amount of resources to be operational, such as their medical kit and other essential equipment. An injured person may have any combination of the aforementioned injuries. Regarding the allocation of non-renewable resources to ACFs in the second decision set, we consider four classes

of them, including drugs, general surgery instruments, orthopedic surgery instruments, and specialized surgery instruments, to treat all mentioned injuries (World Health Organization, 2019). These resources must be supplied by hospitals to ACFs using their limited supply capacities.

The first-stage objective function minimizes the total operational cost of the system, including the establishment cost of ACFs and the transportation cost of non-renewable medical resources from hospitals to ACFs. Moreover, the second-stage objective function minimizes the transportation cost of patients and the penalty corresponding to patients that are not fully recovered due to limited resources in ACFs and hospitals and also the limited number of rescue vehicles.

2.3 Primary Two-Stage Stochastic Mixed-Integer Programming Model (P1)

We propose a two-stage stochastic mixed-integer programming model for the CRP defined in Section 2.2. The first stage includes decisions on:

- Opening ACFs in potential locations.
- Assigning non-renewable medical resources from hospitals to ACFs.
- Allocating rescue vehicles of different types and physicians with various specialties to ACFs.

Then, after the realization of uncertain demands and hospital treatment capacities, in the second-stage model, we determine the transportation of casualties with varying injuries from different demand points to ACFs and hospitals based on their respective itineraries. High-priority patients receive top priority, ensuring immediate transportation to hospitals, where they are promptly deemed “fully treated” upon arrival. On the other hand, low-priority patients follow a distinct itinerary, which may involve visits to more than one ACF before achieving the “fully treated” status. In the following, we present the first- and second-stage models separately.

2.3.1 First-Stage Model

We use the following notation to formulate the first-stage model:

Sets:

- $\mathcal{J}$: The set of potential locations $i \in \mathcal{J}$ for opening ACFs.
- $\mathcal{H}$: The set of existing hospitals $h \in \mathcal{H}$.
- $\mathcal{M}$: The set of rescue vehicle types $m \in \mathcal{M}$.
- $\mathcal{J}$: The set of possible injuries $j \in \mathcal{J}$.
- $\mathcal{J}_i$: The subset of injuries that can be treated at ACF $i \in \mathcal{J}$, where $\mathcal{J}_i \subseteq \mathcal{J}$.
- $\mathcal{K}$: The set of non-renewable resources $k \in \mathcal{K}$.

Parameters:

- f_i : Fixed cost of opening an ACF at location i .
- n_m^{ACF} : The available number of rescue vehicles of type m that can be assigned to ACFs.
- n_j^{spec} : The available number of physicians with specialty j .
- c_{ik}^{Inv} : The inventory capacity for the non-renewable resource k for an ACF located in i .
- c_{hk}^{Supply} : The maximum supply capacity of non-renewable resource k from hospital h .
- ϱ_k : The unit cost associated with the procurement of non-renewable medical resource k for ACFs, covered by the responsible disaster management organization.

First stage variables:

- x_i : 1 if an ACF is opened at location i ; 0 otherwise.
- ϑ_{mi} : The number of rescue vehicles of type m assigned to ACF at location i .
- v_{ij} : The number of physicians with specialty j assigned to ACF at location i (Note that ACF at location $i \in \mathcal{J}$ should be able to treat injury type j , i.e., $j \in \mathcal{J}_i$).
- r_{ihk} : The number of non-renewable resources k assigned to ACF at location i from hospital h .

Using the above notation, we formulate the first-stage model as follows:

$$\min_{\mathbf{x}, \boldsymbol{\vartheta}, \mathbf{v}, \mathbf{r}} \left(\sum_{i \in \mathcal{J}} f_i x_i + \sum_{i \in \mathcal{J}} \sum_{h \in \mathcal{H}} \sum_{k \in \mathcal{K}} \varrho_k r_{ihk} + Q(\mathbf{x}, \boldsymbol{\vartheta}, \mathbf{v}, \mathbf{r}) \right) \quad (1)$$

Subject to:

$$\sum_{\substack{i \in \mathcal{I}: \\ j \in \mathcal{J}_i}} v_{ij} \leq n_j^{spec} \quad j \in \mathcal{J} \quad (2)$$

$$\sum_{i \in \mathcal{I}} \vartheta_{mi} \leq n_m^{ACF} \quad m \in \mathcal{M} \quad (3)$$

$$\sum_{h \in \mathcal{H}} r_{ihk} \leq c_{ik}^{Inv} x_i \quad i \in \mathcal{I}, k \in \mathcal{K} \quad (4)$$

$$\sum_{i \in \mathcal{I}} r_{ihk} \leq c_{hk}^{Supply} \quad h \in \mathcal{H}, k \in \mathcal{K} \quad (5)$$

$$\vartheta_{mi} \leq n_m^{ACF} x_i \quad i \in \mathcal{I}, m \in \mathcal{M} \quad (6)$$

$$v_{ij} \leq n_j^{spec} x_i \quad i \in \mathcal{I}, j \in \mathcal{J}_i \quad (7)$$

$$x_i \in \{0, 1\} \quad i \in \mathcal{I} \quad (8)$$

$$\vartheta_{mi}, r_{ihk}, v_{ij} \in \mathbb{Z}^+ \quad i \in \mathcal{I}, h \in \mathcal{H}, m \in \mathcal{M}, k \in \mathcal{K}, j \in \mathcal{J}_i \quad (9)$$

Objective function (1) includes the total fixed cost of opening ACFs, the procurement cost of non-renewable medical resources covered by the responsible disaster management organization, and the second-stage cost $Q(\mathbf{x}, \boldsymbol{\vartheta}, \mathbf{v}, \mathbf{r})$. Here, $Q(\mathbf{x}, \boldsymbol{\vartheta}, \mathbf{v}, \mathbf{r})$ is the expected second-stage cost that is evaluated in the second-stage problem presented in Section 2.3.2. We have $Q(\mathbf{x}, \boldsymbol{\vartheta}, \mathbf{v}, \mathbf{r}) = E_{s \in S}[Q(\mathbf{x}, \boldsymbol{\vartheta}, \mathbf{v}, \mathbf{r}, \xi(s))]$ where $E_{s \in S}[\cdot]$ calculates the expected value over scenario $s \in S$ and $\xi(s)$ is the vector of uncertain parameters including demands and hospitals' treatment capacity in scenario s . Constraint (2) indicates that the number of physicians with any specialty j assigned to ACFs must be less than or equal to the number of available physicians of the same type. Constraint (3) ensures that the total number of any rescue vehicle of type m assigned to different ACFs cannot exceed the corresponding available number of the same type of rescue vehicles. Constraints (4) and (5) imply the restrictions on the inventory and supply capacity levels of ACFs and hospitals, respectively. Also, according to constraints (6) and (7), we can allocate rescue vehicles and physicians with different specialties only to ACFs that have been established. Constraints (8) and (9) represent integrality constraints for the first-stage variables.

2.3.2 Second-Stage Model

To formulate the second-stage model, we use the following notation:

Sets:

- $\mathcal{L}$: The set of demand locations $l \in \mathcal{L}$.
- $\mathcal{P}$: The set of patient classes. This set is the power set of the injury set $\mathcal{J}$ meaning that each member of this set is a subset of the injury set. Each member of $\mathcal{P}$ represents a class of patients with different injuries. For example, $\{1, 3, 4\} \in \mathcal{P}$ represents a class of patients that have the specific injuries 1, 3, and 4.
- S : The set of all possible scenarios $s \in S$.

Parameters:

- c_i^{ACF} : The treatment capacity of ACF i at location i .
- $c_{hs}^{hospital}$: The treatment capacity of hospital h in scenario s .
- c_m^{rescue} : The capacity of the rescue vehicle of type m , measured in terms of transportation time. It represents the total time that a rescue vehicle can spend transporting patients over the considered horizon, multiplied by the number of patients they can carry each time.
- d_{pls}^{low} : The number of low-priority patients of class p realized in demand location l in scenario s .
- d_{pls}^{high} : The number of high-priority patients of class p realized in demand location l in scenario s .
- $n_{mh}^{hospital}$: The available number of rescue vehicles of type m in hospital h .
- t_j : The duration of treatment for a single injury type j .
- τ_j : The available working time of a physician with specialty j in the considered horizon.
- π_{ij} : Estimated travel cost (distance) between point i and point j ($i, j \in \mathcal{J} \cup \mathcal{H} \cup \mathcal{L}$).
- g_{jk} : The amount of non-renewable medical resource k required for the treatment of a single injury type j .
- ρ_{pl}^{low} : Penalty of the unsatisfied low-priority demand of class p at demand location l .
- ρ_{pl}^{high} : Penalty of the unsatisfied high-priority demand of class p at demand location l , where $\rho_{pl}^{low} \ll \rho_{pl}^{high}$.
- $\xi(s)$: The vector of uncertain parameters including $c_{hs}^{hospital}$, d_{pls}^{high} , and d_{pls}^{low} in scenario s .

Second stage variables:

- q_{plims} : The number of patients of class p moved by a rescue vehicle of type m from demand location l to ACF i in scenario s .
- $u_{pp'ihms}$: The number of patients of class p' , who were previously classified as class p , moved by a rescue vehicle of type m from ACF i to hospital h in scenario s .
- ζ_{plhms} : The number of patients of class p moved by a rescue vehicle of type m from demand location l to hospital h in scenario s .
- $\eta_{pp'ii'ms}$: The number of patients of class p' , who were previously classified as class p , moved by a rescue vehicle of type m from ACF i to ACF i' in scenario s .
- μ_{pls}^{low} : Quantity of unsatisfied low-priority demand of class p at demand location l in scenario s .
- μ_{pls}^{high} : Quantity of unsatisfied high-priority demand of class p at demand location l in scenario s .
- σ_{pis} : Patients of class p who are fully treated at ACF i in scenario s (Note that class- p patient should be fully treatable at ACF i).

Based on the given notation, we formulate the second-stage model for each scenario $s \in S$ as follows:

$$\begin{aligned}
 Q(\mathbf{x}, \boldsymbol{\vartheta}, \mathbf{v}, \mathbf{r}, \xi(s)) = & \min_{\substack{q, u, \eta, \zeta, \\ \mu^{low}, \mu^{high}, \sigma}} \sum_{p \in \mathcal{P}} \sum_{m \in \mathcal{M}} \left(\sum_{l \in \mathcal{L}} \sum_{i \in \mathcal{J}} \pi_{li} q_{plims} + \sum_{\substack{p' \subset p \\ p' \neq \emptyset}} \sum_{i \in \mathcal{J}} \sum_{h \in \mathcal{H}} \pi_{ih} u_{pp'ihms} \right. \\
 & + \sum_{l \in \mathcal{L}} \sum_{h \in \mathcal{H}} \pi_{lh} \zeta_{plhms} + \sum_{\substack{p' \subset p \\ p' \neq \emptyset}} \sum_{i \in \mathcal{J}} \sum_{\substack{i' \in \mathcal{J} \\ i' \neq i}} \pi_{ii'} \eta_{pp'ii'ms} \Big) \\
 & + \sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \left(\rho_{pl}^{low} \mu_{pls}^{low} + \rho_{pl}^{high} \mu_{pls}^{high} \right) \tag{10}
 \end{aligned}$$

Subject to:

$$\sum_{i \in \mathcal{J}} \sum_{m \in \mathcal{M}} q_{plims} + \mu_{pls}^{low} = d_{pls}^{low} \quad p \in \mathcal{P}, l \in \mathcal{L} \tag{11}$$

$$\sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} \zeta_{plhms} + \mu_{pls}^{high} = d_{pls}^{high} \quad p \in \mathcal{P}, l \in \mathcal{L} \tag{12}$$

$$\begin{aligned}
 \sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} q_{plims} + \sum_{\substack{p' \in \mathcal{P}: \\ p \subset p', p' \neq \emptyset}} \sum_{i' \in \mathcal{J} \setminus \{i\}} \sum_{m \in \mathcal{M}} \eta_{p'p'ii'ms} = \sigma_{pis} + \sum_{\substack{p' \subset p: \\ (p \setminus p') \subseteq \mathcal{J}_i, p' \neq \emptyset}} \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} u_{pp'ihms} \\
 + \sum_{\substack{p' \subset p: \\ (p \setminus p') \subseteq \mathcal{J}_i, p' \neq \emptyset}} \sum_{i' \in \mathcal{J} \setminus \{i\}} \sum_{m \in \mathcal{M}} \eta_{pp'ii'ms} \quad i \in \mathcal{J}, p \in \mathcal{P} \tag{13}
 \end{aligned}$$

$$\sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} \zeta_{plhms} + \sum_{p \in \mathcal{P}} \sum_{\substack{p' \in \mathcal{P}: \\ p \subset p'}} \sum_{i \in \mathcal{J}} \sum_{m \in \mathcal{M}} u_{p'p'ihms} \leq c_{hs}^{hospital} \quad h \in \mathcal{H} \tag{14}$$

$$\sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} q_{plims} + \sum_{p \in \mathcal{P}} \sum_{\substack{p' \in \mathcal{P}: \\ p \subset p'}} \sum_{i' \in \mathcal{I} \setminus \{i\}} \sum_{m \in \mathcal{M}} \eta_{p'pi'ims} \leq c_i^{ACF} x_i \quad i \in \mathcal{I} \quad (15)$$

$$t_j \left(\sum_{\substack{p \in \mathcal{P}: \\ j \in p}} \sigma_{pis} + \sum_{\substack{p \in \mathcal{P}: \\ j \in p}} \sum_{\substack{p' \subset p: \\ j \notin p'}} \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} u_{pp'ihms} + \sum_{\substack{p \in \mathcal{P}: \\ j \in p}} \sum_{\substack{p' \subset p: \\ j \notin p'}} \sum_{i' \in \mathcal{I} \setminus \{i\}} \sum_{m \in \mathcal{M}} \eta_{pp'ii'ims} \right) \leq \tau_j v_{ij} \quad i \in \mathcal{I}, j \in \mathcal{J}_i \quad (16)$$

$$\sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \pi_{lh} \zeta_{plhms} \leq c_m^{rescue} n_{mh}^{hospital} \quad m \in \mathcal{M}, h \in \mathcal{H} \quad (17)$$

$$\sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \pi_{li} q_{plims} + \sum_{p \in \mathcal{P}} \sum_{\substack{p' \subset p: \\ p' \neq \emptyset}} \sum_{h \in \mathcal{H}} \pi_{ih} u_{pp'ihms} + \sum_{p \in \mathcal{P}} \sum_{\substack{p' \subset p: \\ p' \neq \emptyset}} \sum_{i' \in \mathcal{I} \setminus \{i\}} \pi_{ii'} \eta_{pp'ii'ims} \leq c_m^{rescue} \vartheta_{mi} \quad m \in \mathcal{M}, i \in \mathcal{I} \quad (18)$$

$$\sum_{p \in \mathcal{P}} \sum_{j \in \{p \cap \mathcal{J}_i\}} \left(g_{jk} \left(\sigma_{pis} + \sum_{\substack{p' \subset p: \\ j \notin p'}} \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} u_{pp'ihms} \right. \right. \\ \left. \left. + \sum_{\substack{p' \subset p: \\ j \notin p'}} \sum_{\substack{i' \in \mathcal{I}: \\ i' \neq i}} \sum_{m \in \mathcal{M}} \eta_{pp'ii'ims} \right) \leq \sum_{h \in \mathcal{H}} r_{ihk} \quad i \in \mathcal{I}, k \in \mathcal{K} \quad (19)$$

$$q_{plims}, u_{pp'ihms}, \zeta_{plhms}, \eta_{pp'ii'ims}, \mu_{pls}^{low}, \mu_{pls}^{high}, \sigma_{pis} \geq 0, \quad i \in \mathcal{I}, j \in \mathcal{J}_i, m \in \mathcal{M}, \\ h \in \mathcal{H}, k \in \mathcal{K}, p \in \mathcal{P}, l \in \mathcal{L} \quad (20)$$

The objective function (10) represents the second-stage cost in scenario s , including 1) the transportation time of casualties between demand locations, ACFs, and hospitals and 2) the penalties corresponding to not serving some patients due to limited resources. Constraints (11) and (12) allocate low- and high-priority patients to ACFs and hospitals, respectively. In these constraints, unmet patient demands may occur, leading to penalties in the objective function. Constraint (13) indicates that the number of patients of class p admitted to ACF i should match the combined count of patients who have either fully recovered within ACF i or partially recovered at ACF i before being transferred, with a different status $p' \subset p$, to hospitals or other ACFs for additional treatment. The left-hand side of this constraint computes the number of class- p patients entering ACF i from demand points:

$$\sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} q_{plims}$$

and from other ACFs:

$$\sum_{\substack{p' \in \mathcal{P}: \\ p \subset p', p' \neq \emptyset}} \sum_{i' \in \mathcal{I} \setminus \{i\}} \sum_{m \in \mathcal{M}} \eta_{p'pi'ims}.$$

On the other hand, the right-hand side computes the total number of patients that leave

ACF i as fully recovered σ_{pis} , or are transferred to hospitals:

$$\sum_{\substack{p' \subset p: \\ (p \setminus p') \subseteq \mathcal{J}_i, p' \neq \emptyset}} \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} u_{pp'ihms},$$

or to other ACFs:

$$\sum_{\substack{p' \subset p: \\ (p \setminus p') \subseteq \mathcal{J}_i, p' \neq \emptyset}} \sum_{i' \in \mathcal{I} \setminus \{i\}} \sum_{m \in \mathcal{M}} \eta_{pp'ii'ms}.$$

These last two terms ensure that patients transferred to hospitals or other ACFs have already been treated for at least one injury in ACF i . Therefore, class- p patients who are not fully treated in ACF i are transferred as new patient classes $p' \subset p$. Constraints (14) and (15) ensure that the treatment capacity of hospitals and ACFs is respected. Constraint (16) guarantees that the required treatment time for injury type j at ACF i does not exceed the available time of the corresponding physician at that ACF. It is important to note that we assume each physician treats only one type of injury. Constraints on the rescue vehicle capacity in hospitals and ACFs are formulated by (17) and (18), respectively. On the left-hand side of these constraints, we have computed the total transportation time of patients imposed on the rescue vehicles of ACFs and hospitals, while on the right-hand side, we have the available capacity. Constraint (19) states that the total non-renewable medical resource type k used at ACF i for treating patients must be less than or equal to the amount of this resource supplied by hospitals. Constraint (20) sets the nonnegativity of second-stage decision variables.

Later in our computational results, we will demonstrate a crucial issue with the previously defined second-stage model (**P1**). Unfortunately, it becomes computationally intractable for realistic-size problems due to its substantial number of variables. Consequently, in the upcoming section, we present a novel and enhanced mathematical model, named model (**P2**), which significantly outperforms model (**P1**). In support of this claim, we have established Theorem 2.1 in online Appendix A.1, showing the equivalence between these two mathematical models.

2.4 Improved Two-Stage Stochastic Mixed-Integer Programming Model (P2)

In this section, we present an enhanced version of the model (P1) achieved by reducing the dimensionality of the original model, particularly by minimizing the number of decision variables in the second stage. This reduction is accomplished through the manipulation of patient indices within the variables $u_{pp'ihms}$ and $\eta_{pp'ii'ms}$. Sets, parameters, first-stage variables, and a part of second-stage variables remain unchanged. We define the new second-stage variables used in the model (P2) as follows:

- u_{pihms} : The number of patients of class p moved by a rescue vehicle of type m from ACF i to hospital h in scenario s . Here, p refers to the state of the patient while being transferred to the hospital, not its primary state at the ACF.
- $\eta_{pii'ms}$: The number of class- p patients moved by a rescue vehicle of type m from ACF i to ACF i' in scenario s . Here, p refers to the state of the patient while being transferred to ACF i' , not its primary state at ACF i .
- σ_{pjis} : The number of class- p patients at ACF i that are first treated for their injury j in scenario s . These patients could be treated for other injuries in the same ACF, too, but injury j is the first injury type that they are treated for. This variable is the main key for formulating the improved model presented in this section.

The first-stage model (1)–(9) remain unchanged. In the following, we only present the improved second-stage model.

$$Q(\mathbf{x}, \boldsymbol{\vartheta}, \mathbf{v}, \mathbf{r}, \xi(s)) = \min_{\substack{q, u, \eta, \zeta, \\ \mu^{low}, \mu^{high}}} \sum_{p \in \mathcal{P}} \sum_{m \in \mathcal{M}} \left(\sum_{l \in \mathcal{L}} \sum_{i \in \mathcal{I}} \pi_{li} q_{plims} + \sum_{i \in \mathcal{I}} \sum_{h \in \mathcal{H}} \pi_{ih} u_{pihms} + \sum_{l \in \mathcal{L}} \sum_{h \in \mathcal{H}} \pi_{lh} \zeta_{plhms} \right. \\ \left. + \sum_{i \in \mathcal{I}} \sum_{\substack{i' \in \mathcal{I} \\ i' \neq i}} \pi_{ii'} \eta_{pii'ms} \right) + \sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \left(\rho_{pl}^{low} \mu_{pls}^{low} + \rho_{pl}^{high} \mu_{pls}^{high} \right) \quad (21)$$

Subject to:

$$\sum_{i \in \mathcal{J}} \sum_{m \in \mathcal{M}} q_{plims} + \mu_{pls}^{low} = d_{pls}^{low} \quad p \in \mathcal{P}, l \in \mathcal{L} \quad (22)$$

$$\sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} \zeta_{plhms} + \mu_{pls}^{high} = d_{pls}^{high} \quad p \in \mathcal{P}, l \in \mathcal{L} \quad (23)$$

$$\begin{aligned} & \sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} q_{plims} + \sum_{i' \in \mathcal{J} \setminus \{i\}} \sum_{m \in \mathcal{M}} \eta_{pi'ims} + \sum_{j \in \mathcal{J}_i} \sum_{\substack{p' \in \mathcal{P}: \\ p' \setminus \{j\} = p \\ p' \setminus \{j\} \neq \emptyset}} \sigma_{p'jis} \\ &= \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} u_{pihms} + \sum_{i' \in \mathcal{J} \setminus \{i\}} \sum_{m \in \mathcal{M}} \eta_{pii'ms} + \sum_{j \in (\mathcal{J}_i \cap p)} \sigma_{pjis} \end{aligned} \quad i \in \mathcal{J}, p \in \mathcal{P} \quad (24)$$

$$\sum_{j \in (\mathcal{J}_i \cap p)} \sigma_{pjis} \geq \sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} q_{plims} + \sum_{i' \in \mathcal{J} \setminus \{i\}} \sum_{m \in \mathcal{M}} \eta_{pi'ims} \quad i \in \mathcal{J}, p \in \mathcal{P} \quad (25)$$

$$\sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} \zeta_{plhms} + \sum_{p \in \mathcal{P}} \sum_{i \in \mathcal{J}} \sum_{m \in \mathcal{M}} u_{pihms} \leq c_{hs}^{hospital} \quad h \in \mathcal{H} \quad (26)$$

$$\sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} q_{plims} + \sum_{p \in \mathcal{P}} \sum_{i' \in \mathcal{J} \setminus \{i\}} \sum_{m \in \mathcal{M}} \eta_{pi'ims} \leq c_i^{ACF} x_i \quad i \in \mathcal{J} \quad (27)$$

$$\sum_{\substack{p \in \mathcal{P}: \\ j \in p}} t_j \sigma_{pjis} \leq \tau_j v_{ij} \quad i \in \mathcal{J}, j \in \mathcal{J}_i \quad (28)$$

$$\sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \pi_{lh} \zeta_{plhms} \leq c_m^{rescue} n_{mh}^{hospital} \quad m \in \mathcal{M}, h \in \mathcal{H} \quad (29)$$

$$\begin{aligned} & \sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \pi_{li} q_{plims} + \sum_{p \in \mathcal{P}} \sum_{h \in \mathcal{H}} \pi_{ih} u_{pihms} \\ &+ \sum_{p \in \mathcal{P}} \sum_{i' \in \mathcal{J} \setminus \{i\}} \pi_{ii'} \eta_{pii'ms} \leq c_m^{rescue} \vartheta_{mi} \end{aligned} \quad m \in \mathcal{M}, i \in \mathcal{J} \quad (30)$$

$$\sum_{p \in \mathcal{P}} \sum_{j \in \{p \cap \mathcal{J}_i\}} g_{jk} \sigma_{pjis} \leq \sum_{h \in \mathcal{H}} r_{ihk} \quad i \in \mathcal{J}, k \in \mathcal{K} \quad (31)$$

$$\begin{aligned} & q_{plims}, u_{pihms}, \zeta_{plhms}, \eta_{pii'ms}, \mu_{pls}^{low}, \mu_{pls}^{high}, \sigma_{pjis} \geq 0, \\ & p \in \mathcal{P}, l \in \mathcal{L}, i \in \mathcal{J}, m \in \mathcal{M}, h \in \mathcal{H}, \\ & k \in \mathcal{K}, j \in \mathcal{J}_i, i' \in \mathcal{J}, i' \neq i \end{aligned} \quad (32)$$

The objective function (21) minimizes transportation times and penalties corresponding to not serving patients. Constraint (22) ensures that low-priority patients are allocated to ACFs, and constraint (23) assigns high-priority patients to the hospitals. Constraint (24), together with variables σ_{pjis} , are the key components of this model that make it possible to formulate the revised model with significantly fewer variables compared to the primary model (**P1**). This constraint works as a balance constraint in the improved model. The left-hand side of constraint (24) computes the number of patients of class p in an ACF that are either received from demand locations (i.e., $\sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} q_{plims}$) and other ACFs (i.e., $\sum_{i' \in \mathcal{J} \setminus \{i\}} \sum_{m \in \mathcal{M}} \eta_{pi'ims}$), or are internally generated in the same ACF from treating the first injuries of other patients of classes $p' \in \mathcal{P}_{|p' \setminus \{j\} = p}$, for $j \in \mathcal{J}_i$. The other side of this constraint either transfers these class- p patients to hospitals (i.e., $\sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} u_{pihms}$) or to other ACFs (i.e., $\sum_{i' \in \mathcal{J} \setminus \{i\}} \sum_{m \in \mathcal{M}} \eta_{pii'ms}$), or treats them for another injury (i.e., $\sum_{j \in (\mathcal{J}_i \cap p)} \sigma_{pjis}$).

Constraint (25) indicates that all class- p patients sent to ACF i must be treated for at

least one injury. The hospital and ACF treatment capacities are respected by constraints (26) and (27), respectively. Physicians of different specialties have limited work time during the planning horizon, which is considered by constraint (28). Constraints (29) and (30) formulate the rescue vehicle capacities for hospitals and ACFs, respectively. Constraint (31) indicates the restriction on the availability of non-renewable medical resources. Constraint (32) declares the nonnegativity of second-stage decision variables.

Figures 1 and 2 illustrate, through a simple example, the equivalency of the subproblems in models **(P1)** and **(P2)**, respectively. This scenario involves one hospital and two types of injuries, which leads to $2^2 - 1$ (excluding the empty set) or 3 patient classes — represented as $p = 1$ (e.g., fractures), $p = 2$ (e.g., nerve injuries), and $p = 3$ (both fractures and nerve injuries), with their corresponding binary representations shown in Figure 2. There is a single type of rescue vehicle for patient transport and one demand location served by two ACFs. ACF #1 is equipped to handle all injury types, while ACF #2 specializes in only injury type #2. Both ACFs are constrained by the limited number of physicians and non-renewable resources available for patient care. As a result, patients may need to be transferred after receiving initial treatment for one injury type to another ACF or hospital for further care, depending on the availability of the required resources and physicians. The yellow cylinders are low-priority patients, and the red circles represent high-priority patients, with p and d indicating their patient class and number of demands, respectively. All patient demands are assumed to be fully satisfied within the depicted scenario.




Figure 1: Network flow of model P1




Figure 2: Network flow of model P2

Theorem 2.1. *Models **(P1)** and **(P2)** are equivalent and, when solved, produce the same first-stage optimal solutions with the same optimal objective values.*

PROOF. We have provided the proof in Appendix A.1. The main idea of this proof is that we show that any feasible solution of model **(P1)** can be projected to a feasible solution of

model **(P2)** with the same first-stage solution and the same objective value, and vice versa. Therefore, both models will result in the same first-stage optimal solution and the same optimal objective value.

2.5 L-shaped Algorithm and Enhancements

In this section, we propose an L-shaped algorithm (Van Slyke and Wets, 1969) with various enhancements to efficiently solve the model **(P2)**. The L-shaped method is an iterative approach that alternates between solving the master problem (MP) and subproblems (SPs) for each scenario. After solving the MP, the obtained first-stage solution is fixed in the SPs, which are then solved for all scenarios $s \in S$. Using the dual information from the optimally solved SPs, optimality cuts are generated and added to refine the MP. This process repeats until the maximum acceptable optimality gap is reached or the time limit is met. In the following, we first present the master problem and subproblem of the L-shaped algorithm. We formulate the MP as follows:

$$(MP) \min_{\mathbf{x}, \boldsymbol{\theta}, \mathbf{v}, \mathbf{r}, \theta} \sum_{i \in \mathcal{I}} f_i x_i + \sum_{i \in \mathcal{I}} \sum_{h \in \mathcal{H}} \sum_{k \in \mathcal{K}} q_k r_{ihk} + \theta \quad (33)$$

Subject to:

$$(2) - (9) \\ \theta \geq \sum_{s \in S} \phi_s \theta_s \quad (34)$$

$$\begin{aligned} \theta_s \geq & \sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \left(d_{pls}^{low} \bar{w}_{spl}^{(1)} + d_{pls}^{high} \bar{w}_{spl}^{(2)} \right) - \sum_{h \in \mathcal{H}} c_{hs}^{hospital} \bar{w}_{sh}^{(5)} \\ & - \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} c_m^{rescue} n_{mh}^{hospital} \bar{w}_{smh}^{(9)} - \sum_{i \in \mathcal{I}} c_i^{ACF} x_i \bar{w}_{si}^{(6)} - \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \tau_j v_{ij} \bar{w}_{sij}^{(7)} \\ & - \sum_{i \in \mathcal{I}} \sum_{m \in \mathcal{M}} c_m^{rescue} \vartheta_{mi} \bar{w}_{sim}^{(10)} - \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}} \left(\sum_{h \in \mathcal{H}} r_{ihk} \right) \bar{w}_{sik}^{(14)} \quad s \in S \quad (35) \end{aligned}$$

$$\theta_s \geq 0 \quad s \in S \quad (36)$$

In MP, θ approximates the second-stage cost and ϕ_s represents the probability of scenario s . Optimality cuts (35) are iteratively added after solving SPs as we will discuss later. In constraint (35), $\bar{w}_{spl}^{(1)}$, $\bar{w}_{spl}^{(2)}$, $\bar{w}_{sh}^{(5)}$, $\bar{w}_{si}^{(6)}$, $\bar{w}_{sij}^{(7)}$, $\bar{w}_{smh}^{(9)}$, $\bar{w}_{sim}^{(10)}$, and $\bar{w}_{sik}^{(14)}$ are the optimal dual variables of constraints (22), (23), (26), (27), (28), (29), (30), and (31), respectively. Observe that Benders feasibility cuts are not necessary because the feasibility of the primal subproblem is ensured by constraints (22) and (23). The subproblem (SP_s) is formulated

for each scenario $s \in S$ as follows:

$$(SP_s) \quad \min_{\substack{q, u, \eta, \zeta, \\ \mu^{low}, \mu^{high}, \sigma}} \sum_{p \in \mathcal{P}} \sum_{m \in \mathcal{M}} \left(\sum_{l \in \mathcal{L}} \sum_{i \in \mathcal{I}} \pi_{li} q_{plims} + \sum_{i \in \mathcal{I}} \sum_{h \in \mathcal{H}} \pi_{ih} u_{pihms} + \sum_{l \in \mathcal{L}} \sum_{h \in \mathcal{H}} \pi_{lh} \zeta_{plhms} \right. \\ \left. + \sum_{i \in \mathcal{I}} \sum_{\substack{i' \in \mathcal{I} \\ i' \neq i}} \pi_{ii'} \eta_{pii'ms} \right) + \sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \left(\rho_{pl}^{low} \mu_{pls}^{low} + \rho_{pl}^{high} \mu_{pls}^{high} \right) \quad (37)$$

Subject to:

$$(22) - (26), (29), (32)$$

$$\sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} q_{plims} + \sum_{p \in \mathcal{P}} \sum_{i' \in \mathcal{I} \setminus \{i\}} \sum_{m \in \mathcal{M}} \eta_{pi'im s} \leq c_i^{ACF} \hat{x}_i \quad i \in \mathcal{I} \quad (38)$$

$$\sum_{\substack{p \in \mathcal{P}: \\ j \in p}} t_j \sigma_{pjis} \leq \tau_j \hat{v}_{ij} \quad i \in \mathcal{I}, j \in \mathcal{J}_i \quad (39)$$

$$\sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \pi_{li} q_{plims} + \sum_{p \in \mathcal{P}} \sum_{h \in \mathcal{H}} \pi_{ih} u_{pihms} + \sum_{p \in \mathcal{P}} \sum_{i' \in \mathcal{I} \setminus \{i\}} \pi_{ii'} \eta_{pii'ms} \leq c_m^{rescue} \hat{\vartheta}_{mi} \quad m \in \mathcal{M}, i \in \mathcal{I} \quad (40)$$

$$\sum_{p \in \mathcal{P}} \sum_{j \in \{p \cap \mathcal{J}_i\}} g_{jk} \sigma_{pjis} \leq \sum_{h \in \mathcal{H}} \hat{r}_{ihk} \quad i \in \mathcal{I}, k \in \mathcal{K} \quad (41)$$

In the SP, $\hat{x}_i$, $\hat{v}_{ij}$, $\hat{\vartheta}_{mi}$, and $\hat{r}_{ihk}$ are constants that are set to the values of their corresponding variables obtained by the MP. In the remainder of this section, we propose the following eight types of enhancements to accelerate the convergence of our L-shaped algorithm:

- Multiple cuts strategy (Birge and Louveaux, 1988)
- Warm-up strategy (McDaniel and Devine, 1977)
- Lower bounding functional valid inequalities (LBF) (Batun et al., 2011)
- Strong (Pareto-optimal) cuts (PO) (Sherali and Lunday, 2013)
- Benders dual decomposition (BDD) (Rahmaniani et al., 2020)
- Lagrangian dual decomposition (LDD) (Zou et al., 2019)
- Branch-and-Benders-cut scheme (B&BC) (Codato and Fischetti, 2006)
- Heuristic strategies for strengthened cut generation

2.5.1 Multiple Cuts and Warm-up Strategies

In this section, we discuss two enhancements used to improve our L-shaped algorithm. The first enhancement involves incorporating multiple Benders cuts into the MP at each iteration. There are two strategies to add optimality cuts. In the first strategy, we add the following aggregated cut after solving all SPs:

$$\begin{aligned}
\theta \geq & \sum_{s \in S} \phi_s \left(\sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \left(d_{pls}^{low} \bar{w}_{spl}^{(1)} + d_{pls}^{high} \bar{w}_{spl}^{(2)} \right) - \sum_{h \in \mathcal{H}} c_{hs}^{hospital} \bar{w}_{sh}^{(5)} \right. \\
& - \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} c_m^{rescue} n_{mh}^{hospital} \bar{w}_{smh}^{(9)} - \sum_{i \in \mathcal{I}} c_i^{ACF} x_i \bar{w}_{si}^{(6)} - \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \tau_j v_{ij} \bar{w}_{sij}^{(7)} \\
& \left. - \sum_{i \in \mathcal{I}} \sum_{m \in \mathcal{M}} c_m^{rescue} \vartheta_{mi} \bar{w}_{sim}^{(10)} - \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}} \left(\sum_{h \in \mathcal{H}} r_{ihk} \right) \bar{w}_{sik}^{(14)} \right) \quad (42)
\end{aligned}$$

In constraint (42), $\bar{w}_{spl}^{(1)}$, $\bar{w}_{spl}^{(2)}$, $\bar{w}_{sh}^{(5)}$, $\bar{w}_{si}^{(6)}$, $\bar{w}_{sij}^{(7)}$, $\bar{w}_{smh}^{(9)}$, $\bar{w}_{sim}^{(10)}$, and $\bar{w}_{sik}^{(14)}$ are the optimal dual variables of constraints (22), (23), (26), (27), (28), (29), (30), and (31), respectively. In the second strategy, after solving each subproblem, a single optimality cut is added per scenario:

$$\begin{aligned}
\theta_s \geq & \sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \left(d_{pls}^{low} \bar{w}_{spl}^{(1)} + d_{pls}^{high} \bar{w}_{spl}^{(2)} \right) - \sum_{h \in \mathcal{H}} c_{hs}^{hospital} \bar{w}_{sh}^{(5)} \\
& - \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} c_m^{rescue} n_{mh}^{hospital} \bar{w}_{smh}^{(9)} - \sum_{i \in \mathcal{I}} c_i^{ACF} x_i \bar{w}_{si}^{(6)} - \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \tau_j v_{ij} \bar{w}_{sij}^{(7)} \\
& - \sum_{i \in \mathcal{I}} \sum_{m \in \mathcal{M}} c_m^{rescue} \vartheta_{mi} \bar{w}_{sim}^{(10)} - \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}} \left(\sum_{h \in \mathcal{H}} r_{ihk} \right) \bar{w}_{sik}^{(14)} \quad s \in S \quad (43)
\end{aligned}$$

In this research, we apply the second strategy since, in the literature, it is computationally proven to be more efficient (Batun et al., 2011).

The second enhancement, the warm-up strategy, involves solving the MP as a linear program (LP) for a predefined number of iterations denoted by α^{warm} , or until the number of iterations that the gap is not improved exceeds β^{warm} (McDaniel and Devine, 1977). Then, the integrality constraints are added to the MP and it is solved as a mixed integer linear programming (MILP).

2.5.2 Lower Bounding Functional Valid Inequalities (LBF)

In the following, we present lower bounding functional valid inequalities that improve the lower bound obtained by the master problem. The main idea of LBF valid inequalities is to give the master problem a rough estimation of the second-stage cost based on the

information of an average scenario. This approach is in line with the principles of partial Benders decomposition, as shown in the work of Crainic et al. (2021). Based on Jensen's Inequality (Jensen, 1906), we add the following constraints to MP to impose the lower bounding functional:

$$\begin{aligned} \theta \geq & \sum_{p \in \mathcal{P}} \sum_{m \in \mathcal{M}} \left(\sum_{l \in \mathcal{L}} \sum_{i \in \mathcal{I}} \pi_{li} q_{plim}(\bar{s}) + \sum_{i \in \mathcal{I}} \sum_{h \in \mathcal{H}} \pi_{ih} u_{pihm}(\bar{s}) + \sum_{l \in \mathcal{L}} \sum_{h \in \mathcal{H}} \pi_{lh} \zeta_{plhm}(\bar{s}) \right. \\ & \left. + \sum_{i \in \mathcal{I}} \sum_{\substack{i' \in \mathcal{I} \\ i' \neq i}} \pi_{ii'} \eta_{pii'm}(\bar{s}) \right) + \sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \left(\rho_{pl}^{low} \mu_{pl}^{low}(\bar{s}) + \rho_{pl}^{high} \mu_{pl}^{high}(\bar{s}) \right) \end{aligned} \quad (44)$$

$$(22) - (32) \quad \text{for the average scenario, i.e., } \bar{s}, \text{ rather than scenario } s \quad (45)$$

According to the above relations, the incorporation of LBF valid inequalities requires adding a copy of all second-stage variables and their corresponding constraints for an average scenario to the MP. This utilization of LBF valid inequalities aligns with the principles of partial Benders decomposition, which allows us to generate valid inequalities based on artificial scenarios. Importantly, even if only one such scenario is produced, the best bound that can be obtained is with the average scenario, as demonstrated in Crainic et al. (2021).

2.5.3 Strong (Pareto-optimal) Cuts (PO)

When applying BD, it is common to find multiple optimal solutions when solving the SPs, especially when dealing with network optimization problems such as facility location problems (Magnanti and Wong, 1981). All optimal solutions to the SPs can be used to produce valid cuts for the master problem. However, if the strongest cuts (i.e., Pareto-optimal cuts) are used, the number of iterations and solution time can be significantly decreased. Sherali and Lunday (2013) suggested solving a pre-emptive priority multiple-objective program. According to this method, a perturbed version of the dual subproblem is created to generate Pareto-optimal cuts, represented as model (46). The modification involves incorporating a perturbation term into the dual objective function using a goal-programming weight ($\mathcal{E}$) and core points of the feasible solution set (i.e., x_i^0 , v_{ij}^0 , v_{mi}^0 , and r_{ihk}^0). When $\mathcal{E}$ takes a small value, for example 10^{-11} in our case, solving the perturbed model (46) yields an optimal

solution that corresponds to a Pareto-optimal solution, as stated in Proposition 1.

$$\begin{aligned}
& \max \sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \left(d_{pls}^{low} w_{spl}^{(1)} + d_{pls}^{high} w_{spl}^{(2)} \right) - \sum_{h \in \mathcal{H}} c_{hs}^{hospital} w_{sh}^{(5)} - \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} c_m^{rescue} n_{mh}^{hospital} w_{smh}^{(9)} \\
& - \sum_{i \in \mathcal{I}} c_i^{ACF} \hat{x}_i w_{si}^{(6)} - \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \tau_j \hat{v}_{ij} w_{sij}^{(7)} - \sum_{i \in \mathcal{I}} \sum_{m \in \mathcal{M}} c_m^{rescue} \hat{\vartheta}_{mi} w_{sim}^{(10)} - \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}} \left(\sum_{h \in \mathcal{H}} \hat{r}_{ihk} \right) w_{sik}^{(14)} \\
& + \mathcal{E} \left(\sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \left(d_{pls}^{low} w_{spl}^{(1)} + d_{pls}^{high} w_{spl}^{(2)} \right) - \sum_{h \in \mathcal{H}} c_{hs}^{hospital} w_{sh}^{(5)} - \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} c_m^{rescue} n_{mh}^{hospital} w_{smh}^{(9)} \right. \\
& \left. - \sum_{i \in \mathcal{I}} c_i^{ACF} x_i^0 w_{si}^{(6)} - \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \tau_j v_{ij}^0 w_{sij}^{(7)} - \sum_{i \in \mathcal{I}} \sum_{m \in \mathcal{M}} c_m^{rescue} \vartheta_{mi}^0 w_{sim}^{(10)} - \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}} \left(\sum_{h \in \mathcal{H}} r_{ihk}^0 \right) w_{sik}^{(14)} \right) \quad (46)
\end{aligned}$$

Subject to:

Constraints for the dual form of the model (21) to (32)

Proposition 1. *If we define $opt(\hat{x}, \hat{\vartheta}, \hat{v}, \hat{r})$ and $sec(\hat{x}, \hat{\vartheta}, \hat{v}, \hat{r})$ as the optimal and second-optimal values of the dual sub-problem corresponding to the solution of $\hat{x}$, $\hat{\vartheta}$, $\hat{v}$, and $\hat{r}$, and $opt(x^0, \vartheta^0, v^0, r^0)$ as the optimal value of the dual sub-problem associated with x^0 , ϑ^0 , v^0 , and r^0 . Then when the goal-programming weight $\mathcal{E}$ meets the condition:*

$$0 < \mathcal{E} < \frac{sec(\hat{x}, \hat{\vartheta}, \hat{v}, \hat{r}) - opt(\hat{x}, \hat{\vartheta}, \hat{v}, \hat{r})}{opt(x^0, \vartheta^0, v^0, r^0)},$$

an optimal solution to the perturbed model (46) will also be a Pareto-optimal solution.

Proof. We have provided the proof in Appendix A.2. □

2.5.4 Benders Dual Approach

We develop the BDD cuts proposed by Rahmaniani et al. (2020) based on Lagrangian duality. In this method, instead of substituting the values of the first-stage solution obtained from MP in the subproblem, we impose the solution to the subproblem as coupling constraints. This process involves creating a new MILP model for the subproblem, where copied variables are introduced to mirror those in the MP with equality constraints (51)–(54). Then, we price out those constraints into the objective function using the associated dual multipliers. By applying this relaxation step, the copied variables no longer equal the first-stage solutions. Consequently, we can enforce integrity constraints on these variables (e.g., constraints (57) and (58)). To develop BDD for our problem, we first replace

constraints (38)–(41) in the subproblem with constraints (47)–(55).

$$\sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} q_{plims} + \sum_{p \in \mathcal{P}} \sum_{i' \in \mathcal{I} \setminus \{i\}} \sum_{m \in \mathcal{M}} \eta_{pi'ims} \leq c_i^{ACF} x_i^{Copy} \quad s \in S, i \in \mathcal{I} \quad (47)$$

$$\sum_{\substack{p \in \mathcal{P}: \\ j \in p}} t_j \sigma_{pjis} \leq \tau_j v_{ij}^{Copy} \quad s \in S, i \in \mathcal{I}, j \in \mathcal{J}_i \quad (48)$$

$$\begin{aligned} \sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \pi_{li} q_{plims} + \sum_{p \in \mathcal{P}} \sum_{h \in \mathcal{H}} \pi_{ih} u_{pihms} \\ + \sum_{p \in \mathcal{P}} \sum_{i' \in \mathcal{I} \setminus \{i\}} \pi_{ii'} \eta_{pii'ms} \leq c_m^{rescue} \vartheta_{mi}^{Copy} \end{aligned} \quad s \in S, m \in \mathcal{M}, i \in \mathcal{I} \quad (49)$$

$$\sum_{p \in \mathcal{P}} \sum_{j \in \{p \cap \mathcal{J}_i\}} g_{jk} \sigma_{pjis} \leq \sum_{h \in \mathcal{H}} r_{ihk}^{Copy} \quad s \in S, i \in \mathcal{I}, k \in \mathcal{K} \quad (50)$$

$$x_{is}^{Copy} = \hat{x}_i \quad s \in S, i \in \mathcal{I} \quad (51)$$

$$v_{ijs}^{Copy} = \hat{v}_{ij} \quad s \in S, i \in \mathcal{I}, j \in \mathcal{J}_i \quad (52)$$

$$\vartheta_{mis}^{Copy} = \hat{\vartheta}_{mi} \quad s \in S, m \in \mathcal{M}, i \in \mathcal{I} \quad (53)$$

$$r_{ihks}^{Copy} = \hat{r}_{ihk} \quad s \in S, i \in \mathcal{I}, h \in \mathcal{H}, k \in \mathcal{K} \quad (54)$$

$$x_{is}^{Copy}, v_{ijs}^{Copy}, \vartheta_{mis}^{Copy}, r_{ihks}^{Copy} \in \mathbb{R}^+ \quad s \in S, i \in \mathcal{I}, j \in \mathcal{J}_i, m \in \mathcal{M}, h \in \mathcal{H}, k \in \mathcal{K} \quad (55)$$

Then we let λ_{is}^{xCopy} , λ_{ijs}^{vCopy} , $\lambda_{mis}^{\varthetaCopy}$, and λ_{ihks}^{rCopy} denote the dual multiplier vector corresponding to the constraints (51)–(54), respectively, and form the corresponding Lagrangian relaxation problem as (56)–(58).

$$\begin{aligned} \max_{\lambda^{xCopy}, \lambda^{vCopy}, \lambda^{\varthetaCopy}, \lambda^{rCopy} \in \mathbb{R}^n} \min_{p \in \mathcal{P}} \sum_{m \in \mathcal{M}} \left(\sum_{l \in \mathcal{L}} \sum_{i \in \mathcal{I}} \pi_{li} q_{plims} + \sum_{i \in \mathcal{I}} \sum_{h \in \mathcal{H}} \pi_{ih} u_{pihms} + \sum_{l \in \mathcal{L}} \sum_{h \in \mathcal{H}} \pi_{lh} \zeta_{plhms} \right. \\ \left. + \sum_{\substack{i \in \mathcal{I} \\ i' \in \mathcal{I} \\ i' \neq i}} \pi_{ii'} \eta_{pii'ms} \right) + \sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \left(\rho_{pl}^{low} \mu_{pls}^{low} + \rho_{pl}^{high} \mu_{pls}^{high} \right) - \left(\lambda_{is}^{xCopy} (x_{is}^{Copy} - \hat{x}_i) \right. \\ \left. + \lambda_{ijs}^{vCopy} (v_{ijs}^{Copy} - \hat{v}_{ij}) + \lambda_{mis}^{\varthetaCopy} (\vartheta_{mis}^{Copy} - \hat{\vartheta}_{mi}) + \lambda_{ihks}^{rCopy} (r_{ihks}^{Copy} - \hat{r}_{ihk}) \right) \end{aligned} \quad (56)$$

Subject to:

$$(2) - (9), (22) - (26), (29), (32), (47) - (50)$$

$$x_{is}^{Copy} \in \{0, 1\} \quad i \in \mathcal{I} \quad (57)$$

$$\vartheta_{mis}^{Copy}, v_{ijs}^{Copy}, r_{ihks}^{Copy} \in \mathbb{Z}^+ \quad i \in \mathcal{I}, m \in \mathcal{M}, h \in \mathcal{H}, k \in \mathcal{K}, j \in \mathcal{J}_i \quad (58)$$

Note that, in the model (56)–(58), we apply the integrality requirements on x_i^{Copy} , v_{ij}^{Copy} , ϑ_{mi}^{Copy} , and r_{ihk}^{Copy} . Then, for any values of λ_i^{xCopy} , λ_{ij}^{vCopy} , $\lambda_{mi}^{\varthetaCopy}$, and λ_{ihk}^{rCopy} , after solving the inner minimization of the Lagrangian relaxation model (56)–(58), we can add the

strengthened cut (59) to the MP. Notably, in BDD, the strengthened cut serves as a valid lower bound, eliminating the need to optimize the Lagrangian multipliers.

$$\begin{aligned}
\bar{\theta}^s \geq & \sum_{p \in \mathcal{P}} \sum_{m \in \mathcal{M}} \left(\sum_{l \in \mathcal{L}} \sum_{i \in \mathcal{I}} \pi_{li} \bar{q}_{plims} + \sum_{i \in \mathcal{I}} \sum_{h \in \mathcal{H}} \pi_{ih} \bar{u}_{pihms} + \sum_{l \in \mathcal{L}} \sum_{h \in \mathcal{H}} \pi_{lh} \bar{\zeta}_{plhms} + \sum_{i \in \mathcal{I}} \sum_{\substack{i' \in \mathcal{I} \\ i' \neq i}} \pi_{ii'} \bar{\eta}_{pii'ms} \right) \\
& + \sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \left(\rho_{pl}^{low} \bar{\mu}_{pls}^{low} + \rho_{pl}^{high} \bar{\mu}_{pls}^{high} \right) - \left(\lambda_i^{xCopy} (x_i - \bar{x}_i^{LR}) + \lambda_{ij}^{vCopy} (v_{ij} - \bar{v}_{ij}^{LR}) \right. \\
& \left. + \lambda_{mi}^{\vartheta Copy} (\vartheta_{mi} - \bar{\vartheta}_{mi}^{LR}) + \lambda_{ihk}^{rCopy} (r_{ihk} - \bar{r}_{ihk}^{LR}) \right) \quad (59)
\end{aligned}$$

In the strengthened cut (59), $\bar{x}_i^{LR}$, $\bar{v}_{ij}^{LR}$, $\bar{\vartheta}_{mi}^{LR}$, and $\bar{r}_{ihk}^{LR}$ are parameters that have already been determined in the Lagrangian relaxation problem. Rahmaniani et al. (2020) showed that when the optimal solution for the MP is an integer solution, the optimality cut (59) is at most as strong as the general optimality cut (43). However, when dealing with fractional solutions in the MP, the strengthened cut (59) can outperform the generalized Benders cut (43).

We also investigated the utilization of the LDD method. In this approach, in dealing with (56)–(58), we first optimize the Lagrangian multiplier values, i.e., λ_i^{xCopy} , $\lambda_{mi}^{\vartheta Copy}$, λ_{ij}^{vCopy} , λ_{ihk}^{rCopy} , through the heuristic model (60)–(61) and then generate the cut (59). We refer to the obtained cuts as LDD cuts.

$$\max_{\eta^{Heuristic} \in \mathbb{R}^1, \lambda \in \mathbb{R}^n} \left(\eta^{Heuristic} - \frac{\delta}{2} \|\lambda^{(t-1)} - \lambda\|_2^2 \right) \quad (60)$$

Subject to:

$$\eta^{Heuristic} \leq SP_{obj}^{\mathcal{V}} + (\bar{y} - \bar{z}^{\mathcal{V}}) \lambda \quad \mathcal{V} = 1, 2, \dots, t-1 \quad (61)$$

The objective function (60) aims to increase the amount of cut lifting minus the distance value. The distance value allows the algorithm to find a new vector λ^t close to the vector obtained from the previous iteration $\lambda^{(t-1)}$. In this model, $\bar{y}$ is the vector of fractional values corresponding to the first-stage variable fixed in the MP. The parameter δ is selected as the step size. In (61), $SP_{obj}^{\mathcal{V}}$ denotes the value of the objective function (56) without the Lagrangian penalties obtained through solving the Lagrangian relaxation model in the $\mathcal{V}^{th}$ iteration. Also, $\bar{z}^{\mathcal{V}}$ represents the value vector of the corresponding first-stage variables. After solving the heuristic model (60)–(61), we have the new Lagrangian multiplier values, build the inner minimization problem (56)–(58), and obtain the relaxed first- and second-stage variables. We then add a new constraint to (61) and solve the revised model (60)–(61) again. This process continues until either the new multiplier values show less improvement than the prior iteration or the maximum iteration limit is reached.

2.5.5 Branch-and-Benders-Cut Scheme

In the classical BD, a single MILP problem is solved to optimally at each iteration. This requires constructing a new branch-and-bound tree each time, leading to potential inefficiencies as previously eliminated candidate solutions are revisited. In this article, whenever we mention “*L-shaped*,” we are actually referring to the classical BD, which is a pure cutting-plane method applied to our stochastic mathematical model. An alternative approach involves constructing a single search tree and generating constraints for both integer and fractional solutions encountered within the tree. This alternative approach leads to the same optimal solution and is commonly known as the branch-and-Benders-cut scheme. In this article, we will refer to this alternative method as the “*Branch-and-Benders-Cut (B&BC)*” algorithm, in which we add cuts for fractional solutions only at the root node, and then only when new incumbent solutions are found. These cuts and valid inequalities are incorporated into the problem using the Callback procedures available in the IBM CPLEX Optimization Solver’s Concert Technology library. We will apply both the “*L-shaped*” and “*B&BC*” algorithms and compare the results in Section 2.6.3 and online Appendix A.6.

2.5.6 Heuristic Strategies for Strengthened Cut Generation

In this section, we introduce two enhanced cut generation strategies aimed at reducing computational effort, particularly when optimizing the Lagrangian dual problem for BDD and LDD cuts, as detailed in Section 2.5.4. These acceleration techniques, including the warm-up, BDD, and LDD, are exclusively applied at the root node of the B&C tree (where at most half of the time limit is dedicated to solving the LP at the root node). However, Modifications are still needed to reduce the computational burden of generating strengthened cuts, as solving a large number of mixed-integer programming (MIP) subproblems at each iteration is time-consuming. This is because the method introduced by Rahmaniani et al. (2020) was meant for a single-scenario problem, while we aim to solve a stochastic problem with at least 100 scenarios. It is important to note that, apart from the root node, we use only generalized Benders cuts (43) at all other nodes in the tree.

The first strategy, which we refer to as BBC_I , is inspired by the three-phase approach proposed by Rahmaniani et al. (2020). In this strategy, we first quickly obtain valid cuts by solving the LP relaxation of the MP. We do this because, during the early stages of the Benders algorithm, the solutions of the MP tend to have low quality. Subsequently, we follow an iterative approach by solving the LP relaxation of the MP and SPs while incorporating generalized Benders cuts in each iteration. We continue this process for a specific

number of iterations ($\alpha^{warm} = 15$) or until the number of iterations with unimproved gaps exceeds a predefined threshold ($\beta^{warm} = 3$). Then, we create strengthened Benders cuts, as explained in Section 2.5.4 considering the modified version. In the modified version for BDD and LDD, when we encounter a fractional solution in the MP at the root node, we do not solve all SPs to produce strengthened cuts for every scenario. Instead, we tackle only a subset of them in each iteration. For example, in the first iteration, we solve the first $\alpha^{cluster} = 10\%$ of scenarios; in the second iteration, we solve the next $\alpha^{cluster} = 10\%$, and so on. This process is repeated for a limited number of iterations ($\alpha^{BDD} = 10$) or when the B&BC gap shows no improvement over three consecutive iterations ($\beta^{BDD} = 3$).

In our second strategy, known as BBC_{II} , we employ a unique approach by applying the warm-up process, BDD, and LDD consecutively but on specific scenario clusters. This means that when we encounter a fractional solution in the master problem at the root node, we handle one cluster at a time. For instance, in the first iteration, we focus on the initial $\alpha^{cluster} = 10\%$ of scenarios, utilizing the warm-up process along with the corresponding generalized Benders cuts. In the subsequent iteration, we move to the next $\alpha^{cluster} = 10\%$ of scenarios and apply strengthened Benders cuts tailored to this cluster, and so on. This process of addressing each cluster one at a time is repeated until we reach a set number of iterations ($\alpha^{BBC_{II}} = 15$) or experience a failure to reduce the gap for three consecutive iterations ($\beta^{BBC_{II}} = 3$). So, in essence, we employ the warm-up, BDD, and LDD techniques individually for each cluster rather than consecutively, making it a more step-by-step or “layered” approach. In BBC_{II} , the heuristic model (60)–(61) is terminated either if we reach the maximum number of iterations, $\alpha^{Heuristic} = 10$, or if we fail to lift the cut by at least $\gamma^{Heuristic} = 10^{-1}$. We have also selected the following values for the remaining parameters in this model: $\delta = 10^5$ and $\mathcal{V} = 10^2$. We choose a large number for δ since we aim to remain close to what we obtained as Lagrangian multiplier values after dualizing of constraints (51)–(54).

We have also developed two improvement methods, namely partially relaxed subproblems and ϵ -optimal cuts, for both strategies. In the former method, when we are applying the BDD or LDD, the integrality constraints are only imposed on a subset of first-stage variables, i.e., x_i and ϑ_{mi} . In that case, we produce a weaker optimality cut compared to the Lagrangian cut due to relaxed integrality constraints on variables v_{ij} and r_{ihk} . However, it still provides a stronger cut than a general Benders cut. In the latter method, we stop solving the MIP subproblems in BDD and LDD once we reach the ϵ -optimal solutions. These solutions have objective values no more than $\epsilon = 10^{-2}$ units higher than the lower bound. The overall implementation of the BBC_I algorithm is illustrated as a flowchart in

online Appendix A.3.

2.6 Computational Results

We carried out the computational tests on a machine with two 2.40 GHz Intel Xeon Gold 6148 processors using C++ and IBM ILOG CPLEX Optimization Studio V22.1 software. One core was used to execute each test instance. In Section 2.6.1, we describe the process by which the set of CRP instances with uncertain demand and hospital treatment capacity was obtained. In Section 2.6.2, we compare the performance of the two proposed mathematical models, and then we assess the computational efficiency of the different enhancements used in our algorithm in Sections 2.6.3. The effectiveness of each accelerator is also evaluated in Section 2.6.4. Finally, in Section 2.6.5, we analyze the impact of casualty shortage penalty parameters on model behavior.

2.6.1 Instance Generation

In this section, we explain the details of generating CRP instances with uncertain demand and hospital treatment capacity. For generating a major part of our instances, we used the data provided by Oksuz and Satoglu (2020) and Sun et al. (2021). However, since we consider some details in CRPs for the very first time, we needed to generate the newly defined parameters as detailed later in this section. *Parameter Settings:* We set the number of ACFs and demand locations to $\{10, 15, 20, 25, 30\}$ and the number of hospitals to $\{10\}$. We assume there are three types of rescue vehicles with capacities of 1, 2, and 3 patients per trip (Rambha et al., 2021).

We uniformly generated the demand locations, ACFs, and hospitals in a square area with a side length of 50 km. For each demand location, we generated the number of patients by a uniform distribution of $[50, 200]$. We use the uniform distribution to generate demand and other uncertain parameters because our model is intended to be general and applicable to various disaster scenarios rather than being tailored to a specific event. Different disasters may lead to different demand distributions, but by using a uniform distribution over a realistic range of values, we ensure that our scenario generation captures a broad spectrum of possibilities. This approach provides a diverse set of instances, allowing us to evaluate the performance of our algorithm across a wide range of possible conditions. Additionally, the uniform distribution is well-suited when there is limited information about the actual demand distribution (Jr and Taskin, 2008). In casualty response planning, demand

estimation immediately after a disaster is often uncertain, making the uniform distribution a reasonable choice. Furthermore, within a given range, the uniform distribution is the maximum entropy distribution, meaning that it makes the fewest assumptions about the likelihood of specific values and treats all possibilities equally (Wu et al., 2024). While specific events may require different distributions, our model remains flexible and can be adjusted accordingly for particular disaster cases.

Regarding the injuries of patients, we considered four common categories of disaster-related injuries, including fractures/crush (FC), laceration/contusions/open wounds (LCO), peripheral nerve injury (PNI), and burns (BU). The occurrence probabilities of these injuries are 63%, 33%, 2%, and 2% of injuries in disasters, respectively (Tanaka et al., 1999; Doocy et al., 2013). To determine which patient suffers from which injuries, we first determine the number of injuries for each patient utilizing an integer uniform distribution $[1, 4]$. Then, we generate the patients' injuries using the above occurrence probabilities. This results in $2^4 - 1$ categories of patients, including all possible injury subsets except the empty set. Moreover, Sun et al. (2021) supposed that 40% to 45% of patients are high-priority. In our study, we assume that patients with 1, 2, 3, and 4 injuries are categorized as high-priority patients, with a uniform distribution range of 20 – 30, 35 – 40, 45 – 50, and 55 – 60, respectively.

For the duration of treatments, we used the means and standard deviations for orthopedics, general surgeries, neurosurgery, and plastic surgeries provided by Costa (2017). The non-renewable resources are also classified into four general groups of drug kits and three other surgical kits for general, orthopedic, and specialized surgeries (World Health Organization, 2019). The cost of these kits is \$7.5 for drugs and \$13.5 for other surgical kits (Sun et al., 2021). All patients require the first group of non-renewable resources, while using the other kits depends on the type of injury. We provide the rest of the parameters in Table 2.

Table 2: Parameters Used for Instance Generation.

<i>Parameters</i>	<i>Value</i>	<i>Reference</i>
Probability of occurrence of scenario s	$\phi^s = \frac{1}{ S }$	
The treatment capacity of ACF i	$c_i^{ACF} = U[400, 1200]$	(Oksuz and Satoglu, 2020)
Fixed cost of opening an ACF at point i	$f_i = c_i^{ACF} \times 1.3 \times 25$	(Oksuz and Satoglu, 2020)
The treatment capacity of hospital h in scenario s	$c_{hs}^{hospital} = U[1000, 2000]$	(Caunhye et al., 2015)
The inventory capacity for non-renewable resource k at ACF i	$c_{ik}^{Inv} = c_i^{ACF} \times U[7, 9]$	(Sun et al., 2021)
The maximum supply capacity of non-renewable resource k from hospital h	$c_{hk}^{Supply} = c_{hs}^{hospital} \times U[14, 16]$	(Sun et al., 2021)
Penalty cost of the unsatisfied high-priority demand of class p at location l	$\rho_{pl}^{high} = \$150,000$	(Li et al., 2020)
Penalty cost of the unsatisfied low-priority demand of class p at location l	$\rho_{pl}^{low} = \frac{\rho_{pl}^{high}}{2}$	

2.6.2 Performance of Mathematical Models

As discussed in Section 2.3.2, we provided an improved model (**P2**) that is an enhanced version of the basic mathematical model (**P1**) for the class of CRPs with multiple injuries. In Table 3, we have compared the performance of these two mathematical models. In the instances of Table 3, we have set the number of scenarios and hospitals to 10. Each row of this table presents the average value for 5 different instances. The first two columns of Table 3 show instances' information, including the number of ACFs and the number of demand points. Under “# Vars.” and “# Const.”, we have provided the number of variables and constraints of each model. Under the “Time (sec)” column, we report the time required to solve the problem optimally using CPLEX directly. We have reported the three recent columns for both models (**P1**) and (**P2**). Finally, the last column reports the improvement in the solution time resulting from using model (**P2**) instead of model (**P1**), i.e., $\Delta^{\text{time}} = 100 \times \left(\frac{P1's \text{ time} - P2's \text{ time}}{P1's \text{ time}} \right)$. Table 3 shows that model (**P2**) significantly outperforms model (**P1**) in terms of solution time and is, on average, 42.56% faster. This saving in solution time is mainly a result of the reduction in the number of variables by approximately 88% in the model (**P2**) compared to (**P1**). Since model (**P2**) shows significantly stronger performance, we present all numerical results in the following sections based on this model.

2.6.3 Performance of Algorithms

The main goal of this subsection is to assess the performance of accelerators and compare the two strengthened cut generation strategies explained in Section 2.5.6. Table 4 presents the computational results for the B&BC algorithm and PHA, which serves as our benchmark (detailed information on PHA is provided in online Appendix A.4). Specifically, BBC_0 is the basic implementation of the B&BC scheme, excluding any additional acceleration techniques as outlined in Section 2.5. In contrast, BBC_I and BBC_{II} are advanced versions of the algorithm. These incorporate all enhancement techniques discussed from Section 2.5.1 through 2.5.5, in addition to two distinct strengthened cut generation strategies detailed in Section 2.5.6.

Each row of the Table 4 represents the average results across five instances. We have considered a time limit of 24 hours for solving instances with these four algorithms. In this table, the first two columns provide the number of potential locations for the ACFs and the number of demand locations, respectively. “Nodes No.” and “Time (sec)” refer to the number of explored nodes in the branch-and-bound tree and the total solution time, respectively. Column “Gap (%)” computes the gap between the lower- and the upper-bound, which are

Table 3: Comparison of Mathematical Models P1 and P2.

Problem Size		P1 Model			P2 Model			Δ^{time} (%)
# ACF	# De.Loc.	# Vars.	# Const.	Time (sec)	# Vars.	# Const.	Time (sec)	
20	10	4,056,924	8,671	34,578	403,444	11,671	18,303	47.79
	15	4,125,926	10,195	19,050	472,622	13,195	9,320	48.61
	20	4,194,924	11,669	24,422	541,428	14,669	16,326	33.33
	25	4,263,928	13,195	36,833	610,622	16,195	18,749	41.14
	30	4,332,922	14,682	21,363	679,525	17,682	10,924	44.01
25	10	5,902,912	10,074	42,414	549,268	13,824	23,610	42.89
	15	5,983,156	11,479	41,463	628,821	15,229	24,902	37.13
	20	6,063,406	12,973	35,645	709,023	16,723	19,174	43.57
	25	6,143,656	14,486	52,620	789,370	18,236	29,359	44.92
	30	6,223,908	15,997	54,769	869,701	19,747	29,810	44.93
30	10	8,086,392	11,318	54,778	716,425	15,818	26,652	51.38
	15	8,177,885	12,752	67,389	807,439	17,252	39,199	41.59
	20	8,269,391	14,327	60,052	899,490	18,827	37,193	40.38
	25	8,360,882	15,754	62,961	990,455	20,254	38,681	38.00
	30	8,452,387	17,292	63,979	1,082,231	21,792	39,074	38.80
Average		6,175,907	12,991	44,821	716,658	16,741	25,418	42.56

Note. To save space, the results for #ACF=10 and #ACF=15 are not presented in this table.

presented under columns “*LB*” and “*UB*” respectively. Additionally, for PHA, we include “*No. Iter.*”, which represents the total number of iterations. Column “*Obj.*” provides the final objective function value obtained using PHA. Lastly, “ $\Delta^{Obj.}$ ” measures the relative objective improvement achieved by BBC_I over PHA, calculated as $100 \times \left(\frac{Obj_{PHA} - UB_{BBC_I}}{Obj_{PHA}} \right)$. This metric highlights the percentage reduction in objective function value when using our proposed approach compared to the benchmark.

Since the BBC_I has resulted in the best performance, we reported more information for this algorithm. “*EC (%)*,” “*PC (%)*,” “*TC (%)*,” and “*SC (%)*” give the contribution percentage of establishment cost, procurement cost, transportation cost, and penalty cost in the total cost for the solutions by solving the stochastic model using the BBC_I algorithm. “*VSS (%)*” also computes the value of the stochastic solution in percentage. This value is computed as follows: $VSS = 100 \times \left(\frac{EEV - UB}{EEV} \right)$, where EEV is the expected objective

value of the expected value solution computed using the scenario set. To obtain EEV, we first solve a deterministic problem using the expected value for the stochastic parameters, also known as the “mean-value” problem. Then, we save the optimal variables obtained in the first stage, fix them in the stochastic model, and proceed to solve the second stage for the different scenarios.

According to Table 4, BBC_0 , BBC_{II} , and BBC_I algorithms have resulted in an average optimality gap of 563.1%, 0.4%, and 0.1%, respectively. These results show that the proposed acceleration techniques significantly contribute to improving the performance of BBC_I and BBC_{II} compared to the BBC_0 algorithm. For the largest instances, the average optimality gap of the BBC_0 algorithm has skyrocketed to 7,900%, while we have obtained 2.4% and 1.0% gaps for BBC_{II} and BBC_I , highlighting the accelerators’ effectiveness. Both BBC_I and BBC_{II} algorithms performed well and found optimal solutions in 21 instance settings. In other instances, BBC_I performed better than BBC_{II} in 3 cases, whereas BBC_{II} outperformed BBC_I on a single instance. The solution time analysis reveals that BBC_I achieved the optimal solutions faster when compared to BBC_{II} in 57% of the cases. Moreover, both BBC_I and BBC_{II} algorithms performed better than the BBC_0 approach in terms of solution time across all instances.

Comparing our best algorithm, BBC_I , with the benchmark PHA, we observe a significant improvement in both computational efficiency and solution quality. PHA failed to converge within the 24-hour time limit for even the smallest instances, with an average solution time of 86,400 seconds, whereas BBC_I required only 19,193 seconds on average. This highlights the efficiency of our proposed approach in reaching high-quality solutions within a practical computational timeframe. Furthermore, BBC_I achieved an average objective function value of 111,165, whereas PHA resulted in an objective value of 9,060,122. This corresponds to a 90.76% improvement in the objective function when transitioning from PHA to BBC_I , demonstrating the superiority of our algorithm in solving the CRP problem effectively.

In addition, the results show that the VSS is approximately 95% for all instances. This highlights the importance of accounting for stochasticity when modeling the CRP problem, as solutions obtained by solving the “mean-value” problem are not good enough for the stochastic problem. To respect the page limit, we have presented the results of the L-shaped algorithm in online Appendix A.6. However, it is worth noting that the B&BC algorithm outperformed the L-shaped algorithm when using the same combination of accelerators.

Table 4: Computational Results of the B&BC Algorithm for the CRP problem with Stochastic Demand and Hospital Treatment Capacity.

Instance		BBC_0			BBC_{II}		BBC_I										PHA			
$ J $	$ \mathcal{L} $	Nodes No.	Time (sec)	Gap (%)	Time (sec)	Gap (%)	Nodes No.	Time (sec)	LB	UB	Gap (%)	EC (%)	PC (%)	TC (%)	SC (%)	VSS (%)	No. Iter.	Time (sec)	Obj.	$\Delta^{Obj.}$ (%)
10	10	20,352	8,742	0.0	2,398	0.0	2,673	2,082	71,294	71,294	0.0	55.8	21.6	22.5	0.0	95.98	317	86,400	17,836,640	99.25
	15	57,222	29,572	0.0	3,102	0.0	4,517	3,206	84,936	84,936	0.0	47.1	26.1	26.8	0.0	95.88	205	86,400	10,205,713	95.34
	20	77,248	65,904	12.7	4,740	0.0	6,703	6,453	112,396	112,396	0.0	44.5	26.8	28.7	0.0	95.25	151	86,400	12,967,317	95.6
	25	64,102	86,400	37.4	4,381	0.0	2,512	3,674	134,560	134,560	0.0	43.5	25.9	30.7	0.0	94.97	119	86,400	2,269,492	83.23
	30	53,991	86,400	72.6	5,911	0.0	6,015	9,149	152,864	152,864	0.0	44.0	27.1	28.9	0.0	94.83	91	86,400	3,950,736	86.07
15	10	44,817	52,929	32.50	3,620	0.0	4,670	3,863	70,144	70,144	0.0	54.2	23.4	22.5	0.0	95.84	128	86,400	16,349,430	99.43
	15	49,260	74,733	49.90	7,013	0.0	4,887	3,492	86,515	86,515	0.0	45.3	24.8	28.2	1.7	95.70	100	86,400	11,947,294	98.74
	20	51,057	86,400	105.4	5,800	0.0	7,152	8,848	106,614	106,614	0.0	45.2	26.4	28.4	0.0	95.69	73	86,400	1,635,050	89.71
	25	40,426	86,400	174.7	11,442	0.0	8,099	11,198	129,213	129,213	0.0	44.7	27.1	28.2	0.0	95.43	55	86,400	2,861,579	86.06
	30	36,343	86,400	176.4	13,265	0.0	6,218	11,585	148,799	148,799	0.0	45.9	27.9	26.3	0.0	95.25	46	86,400	826,801	81.19
20	10	44,755	82,091	82.1	8,209	0.0	8,027	8,210	89,341	89,341	0.0	52.5	24.9	22.5	0.0	96.46	64	86,400	19,528,625	99.50
	15	46,615	86,400	83.5	6,809	0.0	3,719	6,333	81,913	81,913	0.0	44.4	26.4	29.2	0.0	96.24	52	86,400	17,980,166	99.03
	20	33,866	86,400	193.7	14,799	0.0	3,447	11,397	135,098	135,098	0.0	45.1	27.7	27.2	0.0	96.40	43	86,400	1,806,150	87.73
	25	37,311	86,400	219.5	17,773	0.0	8,083	16,728	129,052	129,052	0.0	45.1	27.3	27.6	0.0	95.28	30	86,400	2,391,572	86.32
	30	28,487	86,400	493	39,170	0.0	9,713	19,861	189,248	189,248	0.0	44.1	27.9	28.0	0.0	95.58	24	86,400	1,127,985	81.65
25	10	38,567	86,076	174.4	11,123	0.0	4,639	9,288	62,969	62,969	0.0	49.2	24.6	26.2	0.0	96.47	36	86,400	24,211,720	99.74
	15	31,126	86,400	224.0	10,572	0.0	4,412	12,310	81,884	81,884	0.0	44.0	26.3	29.7	0.0	96.06	26	86,400	13,590,664	94.87
	20	30,653	86,400	190.9	25,843	0.0	2,889	20,392	103,774	103,774	0.0	43.8	27.0	29.3	0.0	95.70	19	86,400	7,763,895	93.27
	25	28,500	86,400	827.9	51,736	0.6	11,691	34,941	127,287	127,287	0.0	43.8	27.7	28.5	0.0	95.19	16	86,400	10,000,720	86.79
	30	22,462	86,400	255.0	52,985	2.9	10,058	51,805	144,366	146,494	1.5	44.4	28.3	27.2	0.1	95.27	13	86,400	3,344,404	84.29
30	10	24,411	86,400	410.6	19,230	0.0	2,924	24,869	65,726	65,726	0.0	52.7	22.9	24.4	0.0	96.22	21	86,400	14,594,000	99.42
	15	28,085	86,400	208.9	32,332	0.0	7,054	34,350	82,406	82,406	0.0	46.7	26.5	26.8	0.0	96.13	14	86,400	6,311,763	87.81
	20	25,575	86,400	256.5	40,186	0.0	8,453	46,085	101,943	102,028	0.1	45.4	28.1	26.5	0.0	95.81	11	86,400	16,734,792	94.20
	25	21,031	86,400	1,897	77,125	4.0	9,719	57,124	124,152	125,422	1.1	43.5	27.7	28.8	0.0	95.37	9	86,400	5,555,426	82.77
	30	17,531	86,400	7,899	83,403	2.4	7,879	62,579	157,632	159,152	1.0	45.0	28.3	26.8	0.0	95.36	8	86,400	711,119	77.01
Average		38,152	78,210	563.1	22,119	0.4	6,246	19,193	110,965	111,165	0.1	46.4	26.4	27.2	0.07	95.69	67	86,400	9,060,122	90.76

Note. To save space, details on LB and UB for the BBC_0 and BBC_{II} approaches are not presented in this table.

2.6.4 Marginal Impacts of Accelerators

To evaluate the impact of individual enhancements, we designate BBC_I as the benchmark and sequentially disable each accelerator, examining their effects on the same instances detailed in Table 4. The summarized results can be found in Table 5. In this tables, under “ $BBC_{(I \setminus \{LBF\})}$ ” and “ $BBC_{(I \setminus \{PO\})}$ ”, the performance of the best algorithm, i.e., BBC_I , without LBF valid inequalities and PO cuts, is given, respectively. Additionally, we report the results for the BBC_I algorithm without warm-up and BDD under columns “ $BBC_{(I \setminus \{Warm-up\})}$ ”, and “ $BBC_{(I \setminus \{BDD\})}$ ”, respectively. Column “ $BBC_{(I + \{LDD\})}$ ” gives information about BBC_I algorithm when LDD is added to it. Comparison of the results in these figures with the gaps reported for the BBC_I algorithm in Table 4 shows that each accelerator used in the BBC_I algorithm positively marginalizes the gap. Furthermore, we have analyzed the combined effects of various enhancements in detail in Online Appendix A.5, where the impact of different strategies on solution time, gap reduction, and algorithmic efficiency is thoroughly evaluated.

Additionally, our analysis revealed that among the enhancements to our algorithm, the LBF valid inequalities have the most significant marginal impact on efficiency, whereas the PO cuts have the least. Excluding LBF valid inequalities leads to longer solution times and prevents the algorithm from finding optimal solutions in 15 instances. Notably, within our algorithm, the inclusion of BDDs uniquely extends the average solution time — from 16,510 seconds without BDDs to 19,193 seconds when they are incorporated. This increase likely stems from the necessity of solving additional MIP subproblems at each iteration. However, despite the longer solution times, the use of BDDs offers substantial benefits, specifically for large-size instances, by contributing to a reduction in both the optimality gap and the number of nodes explored in the search tree by approximately 2-3% and 7,000 nodes, respectively. To easily compare the performance of the algorithms, we visualize their performance in the performance profile of Figure 3.

Table 5: Reverse Marginal Impact.

<i>Problem Size</i>		$BBC_{I \setminus \{LBF\}}$		$BBC_{I \setminus \{PO\}}$		$BBC_{I \setminus \{Warm-up\}}$		$BBC_{I \setminus \{BDD\}}$		$BBC_{I \setminus \{LDD\}}$	
#	#	<i>Time</i>	<i>Gap</i>	<i>Time</i>	<i>Gap</i>	<i>Time</i>	<i>Gap</i>	<i>Time</i>	<i>Gap</i>	<i>Time</i>	<i>Gap</i>
<i>ACF</i>	<i>De.Loc.</i>	(<i>sec</i>)	(%)	(<i>sec</i>)	(%)	(<i>sec</i>)	(%)	(<i>sec</i>)	(%)	(<i>sec</i>)	(%)
10	10	4,855	0.0	1,540	0.0	2,488	0.0	1,844	0.00	5,319	0.00
	15	10,384	0.0	2,764	0.0	2,762	0.0	2,605	0.00	6,300	0.00
	20	13,379	0.0	4,521	0.0	4,509	0.0	3,530	0.00	9,270	0.00
	25	20,115	0.0	5,554	0.0	5,872	0.0	3,381	0.00	15,466	0.00
	30	42,556	0.0	7,104	0.0	5,601	0.0	4,649	0.00	14,015	0.00
15	10	8,614	0.0	4,584	0.0	3,140	0.0	4,427	0.00	13,774	0.00
	15	29,093	0.0	4,931	0.0	6,798	0.0	6,978	0.00	15,937	0.00
	20	51,963	5.1	7,919	0.0	4,738	0.0	5,831	0.00	40,206	0.00
	25	80,137	15.9	7,013	0.0	9,182	0.0	8,818	0.00	44,029	0.00
	30	78,381	24.0	13,473	0.0	11,200	0.0	14,692	0.00	44,741	0.00
20	10	34,132	0.0	18,346	0.0	11,284	0.0	11,795	0.00	43,523	0.00
	15	31,421	0.0	6,654	0.0	5,022	0.0	4,444	0.00	36,722	0.00
	20	86,400	25.2	13,144	0.0	12,985	0.0	11,877	0.00	54,277	0.00
	25	86,400	52.2	28,912	0.3	25,119	0.0	15,468	0.00	52,442	0.00
	30	86,400	56.0	38,802	1.8	35,559	0.0	20,069	0.00	56,968	0.76
25	10	40,512	18.5	11,671	0.0	8,481	0.0	6,725	0.00	49,882	0.00
	15	67,763	10.8	11,544	0.0	11,922	0.0	7,780	0.00	51,754	0.00
	20	86,400	39.3	19,162	0.0	25,453	0.0	12,802	0.00	54,936	0.00
	25	86,400	62.2	35,694	0.0	36,465	0.0	29,719	0.00	61,904	0.00
	30	86,400	58.0	63,452	2.3	34,335	0.0	37,542	0.00	67,092	1.70
30	10	53,806	0.0	28,587	0.0	20,156	0.0	5,296	0.00	52,372	0.00
	15	86,400	41.7	45,210	0.0	27,481	0.0	20,331	0.00	60,992	0.00
	20	86,400	79.4	46,737	0.7	38,126	0.0	25,750	0.00	62,452	0.48
	25	86,400	184.3	84,704	2.8	70,184	3.1	78,520	4.03	81,956	3.54
	30	86,400	498.1	82,585	2.6	86,400	5.1	67,884	2.86	83,595	409.97
<i>Average</i>		57,244	46.8	23,784	0.4	20,210	0.3	16,510	0.28	43,197	16.66

This performance profile reports the computing time in seconds on the horizontal axis and the cumulative percentage of instances being solved to optimality within a given CPU time on the vertical axis. Unlike Table 5, which averages the results of 5 instances per row, Figure 3 aggregates the outcomes for all 125 instances utilized in evaluating the algorithms. It shows that our benchmark algorithm, BBC_I , outperforms the alternatives.




Figure 3: Percentages of Instances Solved to Optimality Within Given Computing Times.


2.6.5 Impact of Casualty Shortage Penalty Parameters on Model Behavior

To evaluate the sensitivity of our model to variations in the penalty cost for unmet casualty demand, we conducted a series of computational experiments. Specifically, we examined how changes in the shortage penalty parameters, ρ_{pl}^{high} and ρ_{pl}^{low} , influence optimal decisions and computational performance.

For this analysis, we selected an instance with $|\mathcal{J}| = 10$, $|\mathcal{L}| = 30$, and $|\mathcal{S}| = 100$. The penalty costs ρ_{pl}^{high} and ρ_{pl}^{low} determine the trade-off between minimizing total costs and prioritizing casualty treatment. We systematically varied the penalty cost multipliers, starting from an extremely small value and gradually increasing it up to 1. Throughout this process, we maintained ρ_{pl}^{high} at twice the value of ρ_{pl}^{low} to consistently prioritize high-priority casualties. The objective of this experiment was to assess:

- The impact of increasing penalty costs on total costs, including transportation, procurement, and shortage costs for both casualty priority levels.
- How the model adapts its resource allocation and casualty transportation decisions as the penalty for untreated casualties increases.
- The computational burden imposed by different penalty levels, particularly in terms of solution time.

Figure 4a illustrates how total costs and their components — transportation, procurement, and shortage costs for both priority levels — evolve as a function of the penalty multiplier. Figure 4b presents the proportion of total solution time spent at different penalty levels.



(a) Impact of Shortage Penalty Parameter on Costs.

(b) Impact of Shortage Penalty Parameter on Solution Time.

Figure 4: Impact of Shortage Penalty Parameter on Model Behavior.

From Figure 4a, we observe that when the shortage penalty is set to an extremely small value, the model has no incentive to allocate resources for casualty transportation. Consequently, 100% of casualties remain untreated, leading to the highest possible shortage rate for both high- and low-priority casualties. This is reflected in Figure 4b, where solution times are minimal, as the optimization process is straightforward under these conditions.

As the penalty multiplier increases, the cost associated with high-priority casualties rises steadily, though at a decreasing rate, since the model prioritizes their treatment. Low-priority casualties, however, continue to be neglected. A crucial turning point occurs at approximately 10^{-3} , where the model is forced to reassess its allocation strategy. At this stage, we observe a sharp decline in low-priority shortage costs in Figure 4a, indicating that the model starts allocating resources to these casualties. Simultaneously, Figure 4b shows a peak in computational effort, suggesting that the model faces its most complex trade-off at this stage.


Beyond this threshold, as the penalty cost continues to increase, the model prioritizes treating all casualties, ensuring no one is left unattended. This results in a minimal shortage rate for both priority levels. Interestingly, solution time slightly decreases beyond this point, as the model transitions from making complex trade-offs to following a more structured and straightforward allocation strategy.

These findings underscore the importance of carefully selecting penalty parameters in casualty response planning. Given the life-and-death nature of post-disaster decisions, an inadequately chosen penalty structure could lead to a large number of untreated casualties, potentially worsening the humanitarian impact of the disaster.

2.7 Case Study

Apart from evaluating the generated instances, we also examined the effectiveness of our proposed algorithm in solving a CRP problem derived from an earthquake case study in Turkey with real data. Our study aimed to identify the location of ACFs in Van Province, Turkey, following the earthquake of 2011. Van district in Turkey is located in a seismically active region and, as a result, has experienced numerous earthquakes throughout its history. The eastern Turkish province of Van experienced a catastrophic earthquake on October 23, 2011, which measured 7.2 on the Richter scale and resulted in numerous fatalities. A second harmful tremor (5.6 on the Richter scale) struck near Van City two weeks after the initial earthquake in October. These two earthquakes had a profound impact, as over 600 individuals lost their lives, and 2,500 suffered injuries (International Federation of Red Cross and Red Crescent Societies, 2012). The 2011 Van earthquake has been the subject of numerous scholarly studies, covering a range of aspects such as last-mile relief network design (Noyan et al., 2016; Noyan and Kahvecioğlu, 2018), humanitarian relief network design (Elçi et al., 2018), shelter location (Kılıcı et al., 2015) and assessment planning (Balcik and İhsan Yanıkoğlu, 2020). We contribute to this stream of research by studying the location-allocation of ACFs in the Van district after the disaster utilizing a dataset from the related natural disaster. We mainly follow the data scheme presented by Oksuz and Satoglu (2020) and Noyan et al. (2016). Below, we outline the key parameters for the case study.

We have 94 demand neighborhoods affected by the earthquake (squares in Figure 5a, with size indicating demand) and 31 potential ACF locations (circles in Figure 5a, with size indicating capacity). Noyan et al. (2016) also grouped these 94 points into 30-point and 60-point clusters using a p -median model, which minimizes the total demand weighted travel time. We also evaluate our model on these two versions. Each neighborhood is classified into one of three damage intensity levels, “destructive,” “damaging,” and “strong or very strong,” based on the earthquake’s intensity and distance from the fault line. Each damage intensity level corresponds to different damage state ratios, including “no damage,” “slight and medium damage,” and “heavy damage and collapse” for buildings in the neighborhood. The base demand for each location was then determined by considering the neighborhood’s population, the damage intensity level, the damage state, and the total number of injuries resulting from the 2011 Van earthquake (see the online Appendix in Noyan et al. (2016), for details). To realize demands for each scenario, we multiply the base demands by a deviation factor. These deviation factors are sampled from uniform distributions, which have



(a) Van's Demands, ACFs, and Hospitals. (b) Van District ACFs (Deterministic). (c) Van District ACFs (Stochastic).

Figure 5: Comparison of ACF Establishments in Van District for Deterministic and Stochastic Models.

intervals of $[0.75, 1.10]$, $[0.75, 1.20]$, and $[0.75, 1.30]$ for intensity levels of “destructive,” “damaging,” and “strong or very strong,” respectively.

In the Van district, five primary hospitals are available to assist patients during a disaster (represented by circles with a plus inside in Figure 5a, with size indicating capacity). However, in the event of an earthquake, hospital treatment capacities become uncertain. Therefore, we assume their capacity fluctuates between 70% and 120% of their nominal capacity (Peleg and Kellermann, 2009; Oksuz and Satoglu, 2020). We also compute the distance from the geographical coordinates.

The average optimality gap of approximately 0.79% signifies that our proposed algorithm is capable of generating high-quality solutions. Another important criterion for evaluating a stochastic mathematical model is the Value of the Stochastic Solution (VSS). VSS is a standard metric that quantifies the expected benefit of using a stochastic model instead of its deterministic counterpart. It is computed by comparing the solution values of the stochastic and expected-value deterministic versions of the problem (Birge and Louveaux, 2011). A high VSS indicates that explicitly accounting for uncertainty leads to significantly better solutions, justifying the additional computational effort required for stochastic optimization. For a detailed explanation of the formulation and computation of VSS in our study, we refer the reader to Section 2.6.3.

In our case study, the substantial VSS values, averaging 97.79%, highlight the critical role of stochasticity in disaster management. This indicates that by explicitly accounting for uncertainty and using the stochastic model instead of its deterministic counterpart, overall costs can be reduced by approximately 98% on average. This underscores the importance of incorporating uncertainty into decision-making to enhance response effectiveness. Let

us now delve deeper into the solutions. In both deterministic and stochastic models, we established ACFs with specific capacities and the ability to treat certain injuries. In both models, we established the following ACFs:

- ACF #1 (with a capacity of 750, capable of treating the following injuries: fractures/crush (Injury #1), laceration/contusions/open wounds (Injury #2), peripheral nerve injury (Injury #3), and burns (Injury #4)),
- ACF #2 (with a capacity of 570, treating Injuries #2 and #3),
- and ACF #11 (with a capacity of 52, treating injury #2).

However, in the deterministic model, we opened an additional ACF, #10, with a capacity of 70 patients capable of treating injuries #1, #2, and #4 concurrently (Figure 5b). In the stochastic model, considering 100 different scenarios with varying patient numbers and hospital treatment capacities, we established two more ACFs: ACF #9 (treating injuries #1 and #2) and ACF #21 (specializing in treating Injuries #1 and #4), with capacities of 130 and 100, respectively (Figure 5c). Notably, the decision to incorporate two additional ACFs, namely #9 and #21, rather than just one (e.g., ACF #10), results in a total capacity increase of 160 patients compared to the deterministic model. This decision not only allows us to provide care for more patients but also enhances our geographical coverage, leading to reduced patient transport times.

Examining the challenging scenarios that can arise after an earthquake, particularly when we base decisions on the average demand in the deterministic model, we come across a critical situation in scenario #22, among 100 scenarios generated for our case study. This scenario has the highest number of low-priority cases associated with Injuries #1 and #2. Moreover, if first-stage decisions are made based on the average scenario in the deterministic model, this scenario also leads to the highest number of unsatisfied low-priority patients. In scenario #22, we encounter a remarkable 1,596 low-priority patients, with 402 patients falling into class #3, the most heavily populated patient class within this scenario. Upon closer examination of class #3, we discover that it exclusively comprises patients with Injuries #1 and #2. This revelation underscores the necessity of an ACF specializing in treating these two specific injuries, much like ACF #9 in the stochastic model. This level of responsiveness to the unique demands of various scenarios is a notable advantage of the stochastic approach, and it addresses an oversight in the deterministic model where such specialized care was not adequately considered.

Effective collaboration among various organizations is crucial for ensuring the availability of necessary equipment, an often overlooked aspect in post-disaster scenarios. The differences between the deterministic and stochastic models extend beyond the number, capacity, and treatment capability of the established ACFs. For example, the stochastic model significantly increases the number of rescue vehicles of various types assigned to ACFs to ensure all low-priority patients can be attended to in every scenario, including worst-case scenarios. Concerning the number of physicians, while the deterministic model accurately estimated the number of required orthopedic physicians, it significantly underestimated the need for wound specialists and plastic physicians or dermatologists and overestimated the number of required neurologists. This highlights the importance of a collaborative effort among companies and organizations to ensure that the right mix of medical professionals and resources is readily available in post-disaster situations, as the requirements vary considerably between deterministic and stochastic models.

2.7.1 Evaluating Out-of-Sample Performance and the Impact of Hospital Capacity Uncertainty

To assess the robustness of our stochastic model and ensure its real-world applicability, we conducted an out-of-sample evaluation. While our model is trained using 100 in-sample scenarios derived from historical data of the 2011 Van earthquake, it is crucial to test how well it generalizes to previously unseen disaster conditions. This ensures that the first-stage decisions remain effective even when faced with variations in patient demand and hospital treatment capacities beyond those explicitly considered in the optimization model used to determine the first-stage plan.

Our evaluation followed these steps: First, we solved the casualty response planning problem for the Van earthquake using 100 in-sample scenarios, generated based on historical data and expert assessments (as discussed in the introduction of Section 2.7). Once the model produced an optimal first-stage solution — determining the locations of ACFs and resource allocations — these decisions were fixed. Next, we tested their robustness on 2,000 new out-of-sample scenarios, generated from real-world data distributions but not included in the in-sample set.

Since our model penalizes unsatisfied casualties heavily, even a single untreated casualty can significantly impact the objective function, making it an unreliable measure for evaluating solution quality. To provide a more practical assessment, we define a new key

performance indicator (KPI): *Percentage of Unsatisfied Demand*, which quantifies the proportion of casualties left without treatment when applying the first-stage solution to out-of-sample scenarios. This metric ensures that our optimized response plan remains effective even under previously unseen disaster conditions. The results are presented in Table 6, under two main columns:

- **Out-of-Sample (Trained with Sto. Hos. Treat. Cap.):** Evaluates performance when the model was trained using stochastic demand and uncertain hospital treatment capacities.
- **Out-of-Sample (Trained with Det. Hos. Treat. Cap.):** Evaluates performance when the model was trained using stochastic demand and fixed hospital treatment capacities.

In this table, the first two columns provide the number of potential locations for the ACFs and the number of demand locations, respectively. “*Time (sec)*” refers to the total solution time. Column “*Gap (%)*” computes the gap between the lower- and the upper-bound for in-sample scenarios. Additionally, for out-of-sample evaluations under both stochastic and deterministic hospital treatment capacities, we include “*Ave. Obj.*”, which represents the average objective function value. Columns “*Unsat. Dem. (#)*” and “*Unsat. Dem. (%)*” show the number and percentage of casualties left untreated in disaster areas, respectively. Lastly, “ $\Delta^{Dem.} (%)$ ” measures the reduction in untreated casualties when accounting for stochastic hospital capacities in the in-sample scenarios.

Table 6: In-Sample and Out-of-Sample Performance of BBC_I with Stochastic and Fixed Hospital Caps.

Instance		In-Sample		Out-of-Sample (Trained with Sto. Hos. Treat. Cap.)			Out-of-Sample (Trained with Det. Hos. Treat. Cap.)			$\Delta^{Dem.}$ (%)
				Ave. Obj.	Unsat. Dem. (#)	Unsat. Dem. (%)	Ave. Obj.	Unsat. Dem. (#)	Unsat. Dem. (%)	
31	30	75,536	0.00	129,935	1	0.01	309,783	3	0.12	66.67
	60	86,400	2.38	595,402	8	0.30	906,262	12	0.48	33.33
	94	80,160	0.00	1,705,770	24	0.93	16,475,700	232	9.27	89.66
Average		80,699	0.79	810,369	11	0.41	5,897,248	82	3	63.22

From Table 6, we observe that the average optimality gap of BBC_I in the case study is 0.79%, demonstrating that our algorithm consistently finds high-quality solutions. More importantly, the out-of-sample evaluation shows that, on average, only 0.41% of total casualties remain unsatisfied when applying the optimized first-stage decisions to scenarios unseen during model training. This exceptionally low percentage suggests that decisions

made while explicitly considering both the assessment information provided by key informants and historical disaster data can remain efficient when evaluated using independently generated scenarios designed to assess the impact of uncertainty.

We also examine the role of hospital capacity uncertainty in decision-making. The results highlight a significant difference: When hospital capacity uncertainty is ignored in the in-sample optimization, out-of-sample casualty shortages increase significantly. Specifically, when comparing solutions obtained with and without hospital capacity uncertainty, the percentage of unsatisfied casualties increases from 0.01%, 0.30%, and 0.93% to 0.12%, 0.48%, and 9.27%, respectively. This corresponds to a 63.22% higher number of unsatisfied casualties when hospital capacity uncertainty is neglected. These findings underscore the importance of incorporating hospital capacity uncertainty into casualty response planning models. Failing to account for it leads to significant misallocations and increased casualty shortages in real-world conditions.

2.8 Conclusion

In this paper, we addressed the challenge of a casualty response planning problem with patients having multiple injuries. To tackle this problem, we developed two two-stage stochastic mixed-integer programming models to formulate the CRP with uncertain demands and hospital treatment capacities. The first model is a basic mathematical model, whereas the second one is a more intricate version that significantly reduces solution times. We proved that the two proposed models are equivalent. As our solution methods, we introduced both the L-shaped algorithm, a pure cutting-plane method tailored to our stochastic mathematical model, and the branch-and-Benders-cut algorithm. In order to significantly improve the efficiency of these algorithms, we have incorporated a diverse set of acceleration techniques, including Benders dual decomposition, Lagrangian dual decomposition, a multi-cut reformulation, the integration of Pareto-optimal cuts, the utilization of lower bounding functional valid inequalities, partially relaxed subproblems, and heuristic cut generation strategies.

Our computational results showed that the improved two-stage stochastic programming model is around 42.5% faster than the initial model. We demonstrated that the proposed acceleration techniques significantly improve the performance of the algorithms and reduce the average optimality gap from 563.1% to 0.1%. Particularly, in some instances, the gap improved from nearly 7,900% to around 1%. We also observed that the average VSS is 95.7%, that demonstrates the importance of addressing uncertainty in casualty response

planning problem and highlights the advantage of our proposed two-stage stochastic programming model compared to its corresponding deterministic model. We also provided extensive results to measure the contribution of enhancements separately. To highlight the applicability of our proposed models and algorithms, we ran our best algorithm on a case study from Turkey and demonstrated that we can find near-optimal solution with an average gap of 0.8%. Additionally, we benchmarked our best-performing algorithm, BBC_I , against PHA, a widely used decomposition approach in disaster response operations. The results showed that BBC_I improved the objective function by 90.76% over PHA while requiring significantly less computation time, demonstrating its superior efficiency and solution quality.

We suggest four important directions for future research. First, given the significant impact of road conditions and travel times in post-disaster scenarios, incorporating debris clearance as an integral component of the problem discussed in this paper would be invaluable. Second, considering secondary disasters within the context of CRP, which includes uncertain demands, stochastic hospital treatment capacities, and multiple injuries, would be interesting. It is important to understand the uncertainty and severity of these secondary disasters and assess their likelihood of affecting hospitals. Third, expanding the model to incorporate uncertainty in both travel times and service times could significantly enhance the realism and effectiveness of post-disaster response planning. While our current work focuses on treatment capacity uncertainty, future research could explore how fluctuating travel conditions and dynamic service times impact resource allocation and patient care itineraries. Such considerations are especially relevant in scenarios where road accessibility changes due to debris accumulation and infrastructure failures. Fourth, exploring dynamic resource management in ACFs, where treatment capabilities adapt to real-time demand and resource availability, could enhance flexibility in medical equipment and personnel deployment. This would improve responsiveness to uncertainty and optimize casualty care in disaster scenarios.

Chapter 3

Stochastic Dual Dynamic Programming for Multi-Stage Casualty Response Planning with Platelet Inventory Management

Abstract

In this paper, we study the stochastic casualty response planning problem and propose a multi-stage stochastic programming model, where initial decisions — such as the location of alternative care facilities (ACFs) and rescue vehicle assignments — are fixed, while patient assignments and allocations of apheresis machines (AM) for blood extraction are updated dynamically as uncertainty unfolds. In this framework, patient demands, blood donor availability, and hospital treatment capacities are treated as uncertain parameters. Our model focuses on deploying newly introduced portable AMs, like the *Trima Accel 7*, and incorporates both apheresis blood donation and whole blood collection methods, along with platelet transshipments among medical facilities to meet varying demands. This approach aims to improve patient outcomes by optimizing three key performance indicators: timely transport from disaster areas to medical facilities, timely surgeries, and compatible platelet transfusions. By matching platelet age and blood type to injury severity, we reduce the risks associated with mismatched transfusions. To address the scalability challenges in multistage stochastic optimization, we employ stochastic dual dynamic programming (SDDP) and benchmark it against nested Benders decomposition (NBD). To further enhance the efficiency of these algorithms, we develop a wide range of acceleration techniques, including strong cuts, multiple cuts, lower bounding functional valid inequalities

(LBFVIs), and warm-up strategies. Effective scenario sampling is achieved through randomized quasi-Monte Carlo (RQMC), while K-means++ clustering is applied for scenario reduction, decreasing the number of scenarios by several orders of magnitude. We carry out extensive computational experiments demonstrating that these enhancements significantly boost the performance of the SDDP algorithm, resulting in a 6,078% reduction in the objective function across 2,500 out-of-sample scenarios and a 40% improvement in the value of the stochastic solution. Finally, a case study from the 2011 Van earthquake in Turkey demonstrates the practical applicability and efficiency of our optimized approach.

3.1 Introduction

The world is facing an increase in natural and man-made disasters, such as earthquakes, floods, and industrial accidents, collectively known as catastrophic health events (CHEs). After a CHE, there are often large numbers of injured people who urgently need medical attention. However, the damaged infrastructure, which can include hospitals that may no longer function properly, puts additional pressure on healthcare systems. This makes it crucial to manage resources effectively in the immediate aftermath of a disaster. Casualty response planning (CRP) involves various organizations working together to make the best use of limited resources to assist survivors and ensure they are transported to hospitals. In recent years, the frequency of severe natural disasters has motivated many researchers to work on CRP problems. For example, the 2023 Turkey-Syria earthquakes caused more than 56,000 deaths and 100,000 injuries, the 2011 Tohoku earthquake and tsunami led to 19,729 deaths and 6,233 injuries, and the 2008 Sichuan earthquake resulted in 87,587 deaths. These events underscore the importance of effective CRP in the wake of such devastating disasters.

In the context of CRP, the use of existing infrastructures, such as schoolyards, parks, and stadiums, as alternative care facilities (ACFs) is a well-established strategy to handle the overwhelming number of casualties. The primary goal of ACFs is to provide additional capacity when hospitals are overwhelmed and unable to handle the large number of injured people. However, the benefits of ACFs can only be fully realized with a well-planned location and allocation strategy. For example, during the 1995 Oklahoma City bombing, emergency medical personnel set up an ACF as a triage center. Unfortunately, they were unable to assess nearly half of the victims because the triage center was initially placed too far from the attack site (Larson et al., 2006).

In addition to the reduction of healthcare treatment capacity after CHEs, these disasters

also lead to a sharp increase in the demand for emergency supplies, particularly blood and its components. Blood transfusion is crucial in providing medical care after large-scale, sudden-onset disasters, especially when surgeries are required, whether in hospitals or ACFs. Despite its importance, significant challenges remain in blood collection and distribution following such events. For example, after the 2008 Wenchuan Earthquake, a total of 107,911 units of blood products were collected, but more than 30% of these units were not distributed (Ma et al., 2019). The challenge is even greater when managing platelet inventory, as platelets have only a 5-day shelf life compared to 42 days for red blood cells (RBCs) and 365 days for plasma.

This challenge motivates us to explore advanced blood collection technologies. There are two common methods for blood collection: whole blood collection (WBC) and apheresis blood donation (ABD). In WBC, blood is collected at centers like hospitals, then tested, processed, and separated into components — a procedure that takes 1–2 days, reducing platelet shelf life, which is critical in disasters. In contrast, ABD directly collects specific components, such as platelets, using an apheresis machine (AM), with remaining blood returned to the donor. This method provides platelets for same-day use, making it highly advantageous in emergencies (Hosseini et al., 2020). While AMs are often associated with permanent facilities (Xu and Szmerekovsky, 2023; Samani and Hosseini-Motlagh, 2023), new portable devices like the *Trima Accel 7* can be deployed at ACFs daily post-disaster. Leveraging such innovations can significantly enhance the availability and effectiveness of life-saving resources in disaster response.

The literature on CRP involving the establishment of ACFs, also called casualty collection points or field hospitals, can be classified into three groups: deterministic, stochastic, and robust optimization approaches. Deterministic models, while common, often yield infeasible plans or high penalties due to their inability to explicitly handle uncertainty. Stochastic models address disaster response challenges by incorporating scenarios that reflect evolving needs, using a combination of pre- and post-disaster data. Advances in technology, such as Synthetic Aperture Radar, UAVs (drones), and Mobile Phone and GPS data, provide real-time insights into disaster impact, complementing traditional sources like government records and key informants such as community leaders (Grass et al., 2023). While robust optimization is well-suited for cases with highly uncertain or limited data, the availability of these advanced technologies allows us to obtain reliable situational awareness. Leveraging these advancements, we adopt a scenario-based multi-stage stochastic approach, where uncertainty is represented by different scenarios derived from satellite

analysis, government reports, and insights from key informants. For a comprehensive review, see Farahani et al. (2020) and Dönmez et al. (2021).

Deterministic approaches have been widely explored (Salman and Gül, 2014; Caunhye et al., 2015; Gu et al., 2018; Liu et al., 2019). Salman and Gül (2014) proposed a multi-period mixed-integer linear programming (MILP) model to optimize ACF locations, minimizing casualty travel time and facility costs, solved via CPLEX. Caunhye et al. (2015) developed a location-allocation model for radiological disaster response, tested using a Los Angeles scenario. Gu et al. (2018) formulated an MILP model for ACF locations and medical supply allocation under budget constraints, solving it with a hybrid meta-heuristic. Liu et al. (2019) introduced a bi-objective MILP model optimizing ACF locations and casualty transportation strategies to maximize survivors while minimizing costs. They solved the model using the ϵ -constraint method.

Alternatively, researchers have developed stochastic (Caunhye and Nie, 2018; Oksuz and Satoglu, 2020; Chang et al., 2023; Oksuz and Satoglu, 2024) and robust optimization (Hu et al., 2019a; Li et al., 2020; Sun et al., 2022; Yin et al., 2024) models to address uncertainties in CHES. Caunhye and Nie (2018) proposed a three-stage stochastic programming (SP) model for locating ACFs and triaging casualties, solved using a Benders decomposition (BD) algorithm. Oksuz and Satoglu (2020) introduced a two-stage SP model to optimize ACF locations and capacities, considering hospital availability, tested on a real case study in Turkey with CPLEX. Chang et al. (2023) focused on ACF location and emergency medical service allocation to optimize casualty transport and survival rates. They developed a hybrid simulation-optimization approach, validated with an earthquake scenario in Tainan City, Taiwan. Most recently, Oksuz and Satoglu (2024) proposed a multi-objective SP model for ACF locations and resource allocation during the critical first 72 hours post-earthquake, solved using an augmented ϵ -constraint method.

In the context of robust optimization, Hu et al. (2019a) proposed a robust multi-objective MILP model for ACF location and casualty allocation, considering urban diurnal population shifts. They developed an accelerated BD algorithm for efficient problem-solving. Li et al. (2020) introduced a scenario-based three-stage hybrid model combining robust and stochastic optimization to design emergency response networks under uncertainties from primary and secondary disasters. They solved this model with a customized progressive hedging algorithm using augmented Lagrangian relaxation. Sun et al. (2022) presented a scenario-based robust bi-objective model integrating ACF location, casualty transport, and relief allocation, accounting for casualty uncertainty and facility disruption risks, solved using the ϵ -constraint method. Recently, Yin et al. (2024) addressed evacuation planning,

ACF location, and supply distribution in a multi-period distributionally robust model, using an enhanced branch-and-Benders-cut algorithm for efficient resolution.

Various blood types can significantly impact the health and survival of individuals affected by disasters. Osorio et al. (2018) highlighted that not only RBCs but also platelet transfusions can lead to severe reactions if the blood types are not appropriately matched, mainly because platelets may contain a small number of RBCs. This can result in adverse outcomes for certain age groups and individuals with specific medical conditions. Additionally, the age of platelets or their freshness plays a crucial role in the success rate of surgeries, which in turn significantly impacts patient survival rates. For instance, platelets that are less than one day old are superior for therapeutic use, especially for patients with weakened immune systems. These fresh platelets are highly preferred for patients with severe injuries, such as high-priority patients with serious bleeding who require immediate surgery. Young platelets, aged 2 to 3 days, are suitable for patients with moderate injuries, such as medium-priority oncologic and hematologic patients and those undergoing cardiac surgery. Patients with less severe injuries, such as low-priority patients requiring general surgery, can be treated with older platelets, up to the maximum shelf life of 5 days (Kamyabniya et al., 2021). While it is expected to have patients with different blood groups and varying injury levels after CHEs, none of the studies reviewed earlier have considered platelet inventory management at ACFs, including the use of newly introduced portable AMs. Ignoring the need for suitable platelet transfusion types and ages can result in patients being sent to ACFs that are not adequately equipped to handle their specific injuries.

In addition to the type and freshness of platelets, several important operational details have yet to be fully addressed in the literature. For example, rescue vehicles play a crucial role in evacuating people from disaster areas and transporting emergency personnel and supplies. Proper rescue vehicle allocation after a disaster ensures that emergency medical services personnel can reach the affected population as quickly and efficiently as possible, thereby improving the effectiveness of the emergency response and potentially saving more lives. On the other hand, the limited availability of newly launched portable AMs means that it may not be possible to equip all ACFs with these devices. As a result, if an ACF is equipped with a few portable AMs and begins collecting fresh platelets, it might be necessary to transship some of these platelets to another ACF that lacks sufficient AMs or blood donors. To address this, we decided to study the impact of transshipment as a tool for enhancing the performance of inventory systems and supply chains within the context of CRP. Although there is extensive literature on CRP, to the best of our knowledge, no paper

has simultaneously addressed the assignment of portable AMs along with ACF establishment, rescue vehicle allocation, platelet inventory management, and casualty allocation. The main contributions of this paper are as follows.

- For the first time, we explore stochastic CRP that incorporates ACF establishment and platelet inventory management, utilizing the latest portable apheresis technologies. Our approach matches platelet blood type and age to casualties while also assigning the appropriate rescue vehicles to ACFs and enabling transshipment between medical facilities (MFs).
- This is the first multi-stage SP model that integrates CRP with platelet inventory management, explicitly accounting for uncertainty in casualty counts, platelet demand, and supply.
- We develop SDDP and NBD algorithms to solve the SP model. We significantly improve the algorithms by incorporating strong cuts, a warm-up strategy, a multiple cuts strategy, RQMC sampling, LBFVIs, and K-means++ for scenario clustering. These enhancements improve the overall performance of the algorithms.
- We present extensive computational results for the CRP, addressing multiple injury levels and a variety of blood types. In our experiments, we evaluate the performance of the decomposition algorithms on 1,000 in-sample and 2,500 out-of-sample scenarios and assess the impact of the developed enhancements.
- We also assess the efficiency of our best algorithm using a real-world case study based on an earthquake scenario in Turkey.

The paper is structured as follows: In Section 3.2, we introduce the problem setting, followed by the presentation of our multi-stage SP model in Section 3.3. Section 3.4 details the SDDP algorithm and the enhancement techniques employed to improve its performance. In Section 3.5, we present extensive computational results for the CRP, followed by a case study in Section 3.6. Finally, Section 3.7 concludes the paper with a summary of findings and suggestions for future research directions.

3.2 Problem Description

Disaster management is commonly divided into four main phases: mitigation, preparedness, response, and recovery (Farahani et al., 2020). The first two phases occur before

a disaster, while the last two take place afterward. During the mitigation and preparedness phases, the focus is on reducing exposure to potential disasters and developing action plans for implementation when a disaster strikes. Once a disaster occurs, the goal of the response phase is to manage available resources efficiently and effectively to save lives and preserve financial and physical assets. Finally, the recovery phase involves restoring both the environment and supply chain networks to their normal conditions.

The problem addressed in this paper focuses on the response phase immediately after a disaster. It involves short-term decisions for dynamically allocating portable AMs and transshipping platelets between MFs over a finite planning horizon. We model this problem as a multi-stage stochastic MILP, with five key decisions following a CHE: (i) determining ACF locations, (ii) allocating rescue vehicles to ACFs, (iii) assigning portable AMs to ACFs, (iv) transshipping platelets between hospitals, ACFs, and hospitals-ACFs, and (v) transporting casualties to MFs and assigning platelet units for treatment. The first two sets of decisions are made immediately after the disaster, based on the initial information received about the disaster's intensity and geographical dispersion. These decisions are determined in the initial period and remain fixed for all subsequent periods. Each period is defined as approximately one day, during which we assume that transshipping platelets, transporting casualties to MFs, and assigning platelet units for treatment can be fully carried out. We assume that each period represents a stage where information is revealed. The corresponding decisions are made dynamically in each period based on updated information about key uncertain parameters, such as the number of casualties by injury level and blood type, available hospital treatment capacities, and the number of platelet donors identified via WBC and ABD methods.

We assume that emergency responders will conduct a triage and initial assessment at the disaster site. This initial assessment follows a widely recognized protocol, often summarized by the mnemonic A.B.C.D.E., which stands for Airway, Breathing, Circulation, Disability, and Exposure/Environment. This assessment aims to quickly identify and manage any life-threatening issues while determining the type and urgency of medical care required. For example, imagine a situation where a victim is conscious, able to communicate, aware of their surroundings, and capable of moving all limbs. In such cases, it can be inferred that their airway is open, they are receiving sufficient oxygen to the brain, and there is no major central neurological damage. While we refer to the triage and initial assessment process to contextualize our problem setting, a comprehensive treatment lies beyond the scope of this paper.

Based on the initial assessment, we can classify patients into three groups according to

the urgency of their medical needs and the amount of platelet units required for their treatment: high-, medium-, and low-priority. High-priority patients, such as those with serious bleeding, breathing difficulties, or severe mental impairment, require a substantial amount of platelet units and must be transported directly to hospitals for immediate and comprehensive care. Medium-priority patients, including those with oncologic and hematologic conditions, require a moderate amount of platelet units. Depending on availability, these patients can be treated at either hospitals or ACFs. Low-priority patients with less critical conditions, such as those needing general surgery, require the fewest platelet units and can also be treated at either hospitals or ACFs. Each group follows a distinct care itinerary within the temporary response healthcare system, with hospitals assumed to have limited rescue vehicles and treatment capacities but equipped to handle the diverse needs of these patient categories.

As previously discussed, we prefer not to send a patient to a hospital or an ACF where suitable platelets are not available for their surgeries. Platelet logistics face two key challenges: ensuring correct blood group substitution and maintaining the freshness of platelets to match them effectively with casualties based on the severity of their injuries. For example, fresh platelets that are just one day old and obtained exclusively through apheresis technology are particularly effective for therapeutic purposes in cases of severe bleeding, such as during bone marrow or organ transplants (Muylle et al., 1992). Concerning the first challenge, ABO-incompatible platelet transfusions can sometimes lead to hemolysis, increasing the risk of fatal outcomes. Therefore, it is crucial to match platelets appropriately based on ABO types for different patient groups. However, in disaster situations with extreme shortages of blood products, the urgent need for life-saving transfusions may outweigh the risk of adverse reactions, making it necessary to substitute one blood group for another (Blumberg et al., 2015). Regarding the second challenge, we categorize platelets into three groups: fresh (one day old), young (2 to 3 days old), and old (4 to 5 days old). High-priority patients should be treated exclusively with fresh platelets, while medium-priority patients can be treated with either fresh or young platelets. Low-priority patients, however, can be treated with platelets of any age up to the maximum shelf life of five days. Beyond five days, platelets are considered outdated and should be discarded.

In a nutshell, the proposed approach begins immediately after a disaster strikes. First, potential locations are selected for establishing ACFs within the constraints of a limited establishment budget. Then, a suitable number of rescue vehicles are assigned to each operational ACF to facilitate patient transfers from demand locations. These initial decisions are fixed for the entire planning horizon. As each new period begins, with no prior

knowledge of that period's uncertain parameters, a central agency must decide how many portable AMs to assign to each operational ACF and determine the necessary transshipments between MFs. Hospitals are equipped with their own AMs and have the capability to collect whole blood, a process that can only be performed within the hospital setting. Once the uncertainties of the period are revealed, casualties are transported to either hospitals or ACFs based on the severity of their injuries and the availability of resources. Due to limited healthcare service capacity, patient transfers may be delayed to subsequent periods, incurring significant costs in the model. At the MFs, decisions must be made about allocating platelets of different ages and blood types to the patients, with a strong preference for matching the patient's blood type. Mismatches result in penalties within the model. If there is insufficient platelet inventory, surgeries may need to be postponed. While this is not ideal, it is better than leaving patients in the disaster areas without stabilizing their conditions. The model also determines the amount of platelet inventory, manages outdated platelets, and helps us decide on the number of accepted donors for each method — whether WBC or ABD — once the uncertainties are known in each period.

3.3 Multi-stage Stochastic Integer Programming Model

Each decision stage corresponds to the arrival of new information on casualties, hospital treatment capacities, and donor availability for the WBC and ABD methods. The first stage represents the initial state, where uncertain parameters are unknown. At this stage, decisions include establishing ACFs, assigning rescue vehicles, allocating portable AMs, and coordinating the transshipment of platelets between MFs for period 1. At stage $t \in \{2, \dots, T\}$, uncertain parameters from period $t - 1$ are known, enabling sequential decisions such as allocating portable AMs to ACFs and determining platelet transshipment between MFs for period t . In the last stage, after all uncertain parameters are revealed, the model computes the allocation of injured patients to MFs, postponements of surgeries, costs of platelet inventory and wastage, and platelet extraction via WBC and ABD methods. This section presents subproblems for the first stage, stages $t \in \{2, \dots, T\}$, and the final stage $T + 1$.

3.3.1 First-Stage Model

The first-stage decisions are made when the uncertain parameters for periods $\{1, \dots, t\}$ are still unknown. These decisions include establishing ACFs (x_i) and assigning rescue

vehicles to them (ϑ_{mi}) across all potential locations i and for all types of rescue vehicles m , where these decisions are fixed across all periods. Additionally, during this stage, we determine the number of portable AMs to be allocated to the opened ACFs (y_{1i}) and the amount of platelet transshipment between MFs (b_{1crhi} , $b'_{1crii'}$, and $b''_{1crhh'}$) for period 1. Since other decisions depend on the realization of the uncertain parameters in period 1, they will be made in subsequent stages. We denote the scenario realization at period t by ω^t , the history of scenario realizations up to period t by $\omega^{[t]} := (\omega^1, \dots, \omega^t)$, and the history of realized random parameters by $\xi(\omega^{[t]})$, using the same scenario notation as introduced by Gul et al. (2015). Below is a detailed description of the model.

Sets:

- $\mathcal{T}$: Set of periods $t \in \mathcal{T}$ in the planning horizon.
- $\mathcal{J}$: Set of potential locations $i \in \mathcal{J}$ for opening ACFs.
- $\mathcal{H}$: Set of existing hospitals $h \in \mathcal{H}$.
- $\mathcal{U}$: Set of MFs $u \in \mathcal{U}$ in the network, i.e., $\mathcal{U} = \mathcal{H} \cup \mathcal{J}$.
- $\mathcal{M}$: Set of rescue vehicle types $m \in \mathcal{M}$.
- $\mathcal{C}$: Set of blood groups $c \in \mathcal{C}$.
- $\mathcal{R}$: Set of platelet ages $r \in \mathcal{R}$.

Deterministic Parameters:

- B : The total available budget for opening ACFs.
- f_i : Fixed cost of opening an ACF at location i .
- E_t : Total number of available portable AMs that can be assigned to ACFs in period t .
- A_{cru} : Initial inventory level of platelets of blood group c and age r at facility u .
- π_{ij} : Estimated travel time between points i and j , where $i, j \in (\mathcal{J} \cup \mathcal{H} \cup \mathcal{L})$.
- n_m^{ACF} : Available number of rescue vehicles of type m that can be assigned to ACFs.
- T : Length of the planning horizon for CRP.

First Stage Variables:

- x_i : 1 if ACF is opened at location i ; 0 otherwise.
- ϑ_{mi} : The number of rescue vehicles of type m assigned to ACF i .

Recourse Variables:

- $y_{ti}(\omega^{[t]})$: The number of portable AMs assigned to ACF i at period t under scenario $\omega^{[t]}$.
- $b_{tcrhi}(\omega^{[t]})$: Number of units of platelets of blood group c and age r transshipped from hospital h to ACF i at period t under scenario $\omega^{[t]}$.
- $b'_{tcrii'}(\omega^{[t]})$: Number of units of platelets of blood group c and age r transshipped from ACF i to ACF i' at period t under scenario $\omega^{[t]}$.

$b''_{tcrhh'}(\omega^{[t]})$: Number of units of platelets of blood group c and age r transshipped from hospital h to hospital h' at period t under scenario $\omega^{[t]}$.

$\mathbf{Y}_t(\omega^{[t]})$: The vector of values for all decision variables defined for and before period t under scenario $\omega^{[t]}$.

The first-stage problem can be formulated as follows:

$$\min \sum_{r \in \mathcal{R}} \left(\sum_{h \in \mathcal{H}} \sum_{i \in \mathcal{J}} \pi_{hi} b_{1crhi} + \sum_{h \in \mathcal{H}} \sum_{\substack{h' \in \mathcal{H}: \\ h' \neq h}} \pi_{hh'} b''_{1crhh'} + \sum_{i \in \mathcal{J}} \sum_{\substack{i' \in \mathcal{J}: \\ i' \neq i}} \pi_{ii'} b'_{1crii'} \right) + \mathcal{F}_2(\mathbf{Y}_1(\omega^{[1]})) \quad (62)$$

Subject to:

$$\sum_{i \in \mathcal{J}} f_i x_i \leq B \quad (63)$$

$$\sum_{i \in \mathcal{J}} \vartheta_{mi} \leq n_m^{ACF} \quad m \in \mathcal{M} \quad (64)$$

$$\vartheta_{mi} \leq n_m^{ACF} x_i \quad i \in \mathcal{J}, m \in \mathcal{M} \quad (65)$$

$$\sum_{i \in \mathcal{J}} y_{1i} \leq E_1 \quad (66)$$

$$y_{1i} \leq E_1 x_i \quad i \in \mathcal{J} \quad (67)$$

$$\sum_{i \in \mathcal{J}} b_{1crhi} + \sum_{\substack{h' \in \mathcal{H}: \\ h' \neq h}} b''_{1crhh'} \leq A_{crh} \quad c \in \mathcal{C}, r \in \mathcal{R} \setminus \{1\}, h \in \mathcal{H} \quad (68)$$

$$\sum_{\substack{i' \in \mathcal{J}: \\ i' \neq i}} b'_{1crii'} \leq A_{cri} \quad c \in \mathcal{C}, r \in \mathcal{R} \setminus \{1\}, i \in \mathcal{J} \quad (69)$$

$$x_i \in \{0, 1\} \quad i \in \mathcal{J} \quad (70)$$

$$\vartheta_{mi} \in \mathbb{Z}^+ \quad i \in \mathcal{J}, m \in \mathcal{M} \quad (71)$$

$$y_{1i}, b_{1crhi}, b'_{1crii'}, b''_{1crhh'} \geq 0 \quad i, i' \in \mathcal{J}, c \in \mathcal{C}, r \in \mathcal{R}, h, h' \in \mathcal{H} \quad (72)$$

The objective function (62) accounts for the transshipment costs between facilities and the future cost $\mathcal{F}_2(\mathbf{Y}_1(\omega^{[1]}))$, which depends on the decisions made in the first stage. Constraints (63) and (64) account for the limited budget for establishing ACFs and the limited number of rescue vehicles that can be assigned to ACFs, respectively. Constraint (65) ensures that rescue vehicles are only assigned to opened ACFs. Constraints (66) and (67) ensure the number of portable AMs assigned to opened ACFs does not exceed the available number. Constraints (68) and (69) state that in the first period, the amount of platelets that can be transshipped among facilities cannot exceed their initial inventories. Finally, constraints (70)–(72) define the domains of all first-stage decision variables.

3.3.2 Subproblem of Stages $t \in \{2, \dots, T\}$

The notation used to formulate the subproblem for stages $t \in \{2, \dots, T\}$ is introduced below.

Sets:

- $\mathcal{L}$: Set of demand locations $l \in \mathcal{L}$.
- $\mathcal{G}_c$: Set of blood groups that can serve as substitutes for blood group $c \in \mathcal{C}$.
- $\mathcal{J}$: Set of patient injury levels, $\mathcal{J} = \{1, 2, 3\}$, where 1 = high, 2 = medium, and 3 = low.
- $\mathcal{R}_j$: Set of platelet ages suitable for injury level $j \in \mathcal{J}$, e.g., $\mathcal{R}_1 = \{1\}$, $\mathcal{R}_2 = \{1, 2, 3\}$, and $\mathcal{R}_3 = \{1, 2, 3, 4, 5\}$.
- $\mathcal{J}_u$: Set of injury levels that can be treated at facility $u \in \mathcal{U}$, e.g., $\mathcal{J}_{h \in \mathcal{H}} = \{1, 2, 3\}$ and $\mathcal{J}_{i \in \mathcal{I}} = \{2, 3\}$.
- $\mathcal{J}_r$: Set of injury levels that can be treated using platelets of age $r \in \mathcal{R}$, e.g., $\mathcal{J}_1 = \{1, 2, 3\}$.

Deterministic Parameters:

- τ : Number of units of platelets extracted from one donor by an AM.
- e : The processing and production time of whole blood collected at hospitals.
- $c^{Aph.}$: The platelet production capacity of an AM in each period.
- ψ_u : Platelet wastage cost at facility u .
- Δ_u : Platelet inventory cost at facility u .
- β_j : The average number of units of platelets needed for a patient with injury level j .
- P_j : Postponing cost of a surgery for patients with injury level j .
- $W_{c'c}$: Penalty cost for substituting platelets of blood group c with those of blood group c' .
- c_i^{ACF} : Treatment capacity of ACF at location i .
- c_m^{rescue} : Patient capacity of a rescue vehicle of type m , measured in terms of transportation time.
- ρ_j : Penalty for the unsatisfied casualties with injury level j .
- $n_h^{Aph.}$: Number of AMs in hospital h .
- $n_{mh}^{Hos.}$: Available number of rescue vehicles of type m at hospital h .

Random Parameters and Sets:

- $d_{tjcl}(\omega^{[t]})$: The random number of patients with injury level j at location l who requires platelets of blood group c at period t under scenario $\omega^{[t]}$.
- $c_{th}^{Hos.}(\omega^{[t]})$: The random treatment capacity of hospital h at period t under scenario $\omega^{[t]}$.
- $\theta_{tch}^{Whole}(\omega^{[t]})$: The random number of available whole-blood donors at hospital h , for platelets of blood group c on period t under scenario $\omega^{[t]}$.
- $\theta_{tcu}^{Aph.}(\omega^{[t]})$: The random number of available apheresis donors at facility u , for platelets of blood group c on period t under scenario $\omega^{[t]}$.

Recourse Variables:

- $q_{tjclum}(\omega^{[t]})$: The number of patients with injury level j who require platelets of blood group c and are transported from demand location l to facility u using a rescue vehicle of type m at period t under scenario $\omega^{[t]}$.
- $\mu_{tjcl}(\omega^{[t]})$: The number of unsatisfied casualties with injury level j at demand location l at period t who require platelets of blood group c under scenario $\omega^{[t]}$.
- $\lambda_{tcu}(\omega^{[t]})$: Number of units of platelets of blood group c extracted by an AM at facility u at period t under scenario $\omega^{[t]}$.
- $\varrho_{tch}(\omega^{[t]})$: Number of units of platelets of blood group c extracted from donated whole blood at hospital h at period t under scenario $\omega^{[t]}$ that becomes available after a processing time of e days.
- $\eta_{tcru}(\omega^{[t]})$: The inventory level of platelets of blood group c and age r at facility u at the end of period t under scenario $\omega^{[t]}$.
- $\varsigma_{tu}(\omega^{[t]})$: Number of units of outdated platelets at facility u at period t under scenario $\omega^{[t]}$.
- $v_{tjc'cru}(\omega^{[t]})$: The number of patients with injury level j who require platelets of blood group c , served by platelets of blood group c' with age r , at facility u at period t under scenario $\omega^{[t]}$. (Here, if $c' = c$, it means the same blood group was used and no substitution occurred).
- $\zeta_{tjcu}(\omega^{[t]})$: The number of patients with injury level j at facility u who have not received platelets of blood group c at period t , resulting in the postponement of their surgeries under scenario $\omega^{[t]}$.

Letting $\mathcal{F}_{t+1}(\mathbf{Y}_t(\omega^{[t]})) = E_{\xi^{[t]}}[Q^{t+1}(\mathbf{Y}_t(\omega^{[t]}), \xi^{[t]}(\omega^{[t]}))]$, we obtain the recursion model (73)–(92). Here, $\hat{X}$ denotes the values of decision variables X determined in previous stages before t .

$$\begin{aligned}
Q^t(\mathbf{Y}_{t-1}(\omega^{[t-1]}), \xi^{[t-1]}(\omega^{[t-1]})) = & \min \sum_{c \in \mathcal{C}} \left(\sum_{j \in \mathcal{J}} \left(\sum_{l \in \mathcal{L}} \sum_{u \in \mathcal{U}} \sum_{m \in \mathcal{M}} \pi_{lu} q_{(t-1)jclum}(\omega^{[t-1]}) + \sum_{u \in \mathcal{U}} P_j \zeta_{(t-1)jcu}(\omega^{[t-1]}) \right. \right. \\
& + \left. \sum_{l \in \mathcal{L}} \rho_j \mu_{(t-1)jcl}(\omega^{[t-1]}) \right) + \sum_{r \in \mathcal{R}} \sum_{u \in \mathcal{U}} \Delta_u \eta_{(t-1)cru}(\omega^{[t-1]}) + \sum_{u \in \mathcal{U}} \left(\psi_u \varsigma_{(t-1)u}(\omega^{[t-1]}) \right. \\
& + \left. \sum_{j \in \mathcal{J}} \sum_{c' \in \mathcal{C}} \sum_{c \in \mathcal{C}} \sum_{r \in \mathcal{R}} W_{c'c} v_{(t-1)jc'cru}(\omega^{[t-1]}) \right) + \sum_{r \in \mathcal{R}} \left(\sum_{h \in \mathcal{H}} \sum_{i \in \mathcal{J}} \pi_{hi} b_{tcrhi}(\omega^{[t-1]}) \right. \\
& + \left. \sum_{h \in \mathcal{H}} \sum_{\substack{h' \in \mathcal{H}: \\ h' \neq h}} \pi_{hh'} b''_{tcrhh'}(\omega^{[t-1]}) + \sum_{i \in \mathcal{J}} \sum_{\substack{i' \in \mathcal{J}: \\ i' \neq i}} \pi_{ii'} b'_{tcrii'}(\omega^{[t-1]}) \right) + \mathcal{F}_{t+1}(\mathbf{Y}_t(\omega^{[t]}))
\end{aligned} \tag{73}$$

Subject to:

$$\sum_{\substack{u \in \mathcal{U}: m \in \mathcal{M} \\ j \in \mathcal{J}_u}} q_{(t-1)jclum}(\omega^{[t-1]}) + \mu_{(t-1)jcl}(\omega^{[t-1]}) = d_{(t-1)jcl}(\omega^{[t-1]}) + \hat{\mu}_{(t-2)jcl} \quad j \in \mathcal{J}, l \in \mathcal{L}, c \in \mathcal{C} \tag{74}$$

$$\begin{aligned}
& \sum_{c' \in \mathcal{G}_c} \sum_{r \in \mathcal{R}_j} v_{(t-1)jc'cru}(\omega^{[t-1]}) + \zeta_{(t-1)jcu}(\omega^{[t-1]}) = \\
& \sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} q_{(t-1)jclum}(\omega^{[t-1]}) + \hat{\zeta}_{(t-2)jcu} \quad j \in \mathcal{J} \setminus \{1\}, c \in \mathcal{C}, u \in \mathcal{U} \tag{75}
\end{aligned}$$

$$\sum_{c' \in \mathcal{G}_c} v_{(t-1)jc'c1h}(\omega^{[t-1]}) + \zeta_{(t-1)jch}(\omega^{[t-1]}) = \sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} q_{(t-1)jclhm}(\omega^{[t-1]}) + \hat{\zeta}_{(t-2)jch} \quad j = 1, c \in \mathcal{C}, h \in \mathcal{H} \quad (76)$$

$$\sum_{j \in \mathcal{J}_i} \sum_{c \in \mathcal{C}} \left(\sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} q_{(t-1)jclim}(\omega^{[t-1]}) + \hat{\zeta}_{(t-2)jci} \right) \leq c_i^{ACF} \hat{x}_i \quad i \in \mathcal{I} \quad (77)$$

$$\sum_{j \in \mathcal{J}_h} \sum_{c \in \mathcal{C}} \left(\sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} q_{(t-1)jclhm}(\omega^{[t-1]}) + \hat{\zeta}_{(t-2)jch} \right) \leq c_h^{Hos.}(\omega^{[t-1]}) \quad h \in \mathcal{H} \quad (78)$$

$$\sum_{j \in \mathcal{J}_h} \sum_{c \in \mathcal{C}} \sum_{l \in \mathcal{L}} \pi_{lh} q_{(t-1)jclhm}(\omega^{[t-1]}) \leq c_m^{rescue} n_{mh}^{Hos.} \quad h \in \mathcal{H}, m \in \mathcal{M} \quad (79)$$

$$\sum_{j \in \mathcal{J}_i} \sum_{c \in \mathcal{C}} \sum_{l \in \mathcal{L}} \pi_{li} q_{(t-1)jclim}(\omega^{[t-1]}) \leq c_m^{rescue} \hat{\vartheta}_{mi} \quad i \in \mathcal{I}, m \in \mathcal{M} \quad (80)$$

$$\sum_{c \in \mathcal{C}} \lambda_{(t-1)ci}(\omega^{[t-1]}) \leq c^{Aph.} \hat{y}_{(t-1)i} \quad i \in \mathcal{I} \quad (81)$$

$$\sum_{c \in \mathcal{C}} \lambda_{(t-1)ch}(\omega^{[t-1]}) \leq c^{Aph.} n_h^{Aph.} \quad h \in \mathcal{H} \quad (82)$$

$$\lambda_{(t-1)cu}(\omega^{[t-1]}) \leq \tau \theta_{(t-1)cu}^{Aph.}(\omega^{[t-1]}) \quad c \in \mathcal{C}, u \in \mathcal{U} \quad (83)$$

$$\varrho_{(t-1)ch}(\omega^{[t-1]}) \leq \theta_{(t-1)ch}^{Whole}(\omega^{[t-1]}) \quad c \in \mathcal{C}, h \in \mathcal{H} \quad (84)$$

$$\varsigma_{(t-1)u}(\omega^{[t-1]}) = \sum_{c \in \mathcal{C}} \eta_{(t-1)c|R|u}(\omega^{[t-1]}) \quad u \in \mathcal{U} \quad (85)$$

$$\sum_{i \in \mathcal{I}} y_{ti}(\omega^{[t-1]}) \leq E_t \quad (86)$$

$$y_{ti}(\omega^{[t-1]}) \leq E_t \hat{x}_i \quad i \in \mathcal{I} \quad (87)$$

$$\sum_{i \in \mathcal{I}} b_{tcrhi}(\omega^{[t-1]}) + \sum_{h' \in \mathcal{H}:} b'_{tcrhh'}(\omega^{[t-1]}) \leq \eta_{(t-1)c(r-1)h}(\omega^{[t-1]}) \quad c \in \mathcal{C}, r \in \mathcal{R}, h \in \mathcal{H} \quad (88)$$

$$\sum_{\substack{i' \in \mathcal{I}: \\ i' \neq i}} b'_{tcrii'}(\omega^{[t-1]}) \leq \hat{\eta}_{(t-1)c(r-1)i}(\omega^{[t-1]}) \quad c \in \mathcal{C}, r \in \mathcal{R} \setminus \{1\}, i \in \mathcal{I} \quad (89)$$

$$\begin{aligned} & \lambda_{(t-1)c'h}(\omega^{[t-1]}) \mathbf{1}_{(r=1)} + \left(\hat{\eta}_{(t-2)c'(r-1)h} + \sum_{\substack{h' \in \mathcal{H}: \\ h' \neq h}} (\hat{b}'_{(t-1)c'rh'h} - \hat{b}'_{(t-1)c'rh'h'}) \right) \\ & - \sum_{i \in \mathcal{I}} \hat{b}_{(t-1)c'rh'i} \mathbf{1}_{(r \neq 1)} + \hat{b}_{(t-e-1)c'h} \mathbf{1}_{((r=e) \wedge (t-e-1) \geq 0)} + A_{c'rh} \mathbf{1}_{(t=2)} = \\ & \sum_{\substack{j \in \mathcal{J}_h \\ j \in \mathcal{J}_r}} \sum_{\substack{c \in \mathcal{C}: \\ c' \in \mathcal{G}_c}} \beta_j v_{(t-1)jc'crh}(\omega^{[t-1]}) + \eta_{(t-1)c'rh}(\omega^{[t-1]}) \quad c' \in \mathcal{C}, r \in \mathcal{R}, h \in \mathcal{H} \quad (90) \end{aligned}$$

$$\begin{aligned} & \lambda_{(t-1)c'i}(\omega^{[t-1]}) \mathbf{1}_{(r=1)} + \left(\hat{\eta}_{(t-2)c'(r-1)i} + \sum_{\substack{i' \in \mathcal{I}: \\ i' \neq i}} (\hat{b}'_{(t-1)c'ri'i} - \hat{b}'_{(t-1)c'ri'i'}) \right) \\ & + \sum_{h \in \mathcal{H}} \hat{b}_{(t-1)c'rh'i} \mathbf{1}_{(r \neq 1)} = \sum_{\substack{j \in \mathcal{J}_i \\ j \in \mathcal{J}_r}} \sum_{\substack{c \in \mathcal{C}: \\ c' \in \mathcal{G}_c}} \beta_j v_{(t-1)jc'cri}(\omega^{[t-1]}) + \eta_{(t-1)c'ri}(\omega^{[t-1]}) \quad c' \in \mathcal{C}, r \in \mathcal{R}, i \in \mathcal{I} \quad (91) \end{aligned}$$

$$\begin{aligned} & y_{ti}(\omega^{[t-1]}), b_{tcrhi}(\omega^{[t-1]}), b'_{tcrii'}(\omega^{[t-1]}), b'_{tcrhh'}(\omega^{[t-1]}) \\ & q_{(t-1)jclum}(\omega^{[t-1]}), \mu_{(t-1)jcl}(\omega^{[t-1]}), \zeta_{(t-1)jcu}(\omega^{[t-1]}) \quad i, i' \in \mathcal{I}, j \in \mathcal{J}, c, c' \in \mathcal{C}, l \in \mathcal{L}, \\ & \lambda_{(t-1)cu}(\omega^{[t-1]}), \varrho_{(t-1)ch}(\omega^{[t-1]}), \eta_{(t-1)cru}(\omega^{[t-1]}) \quad u \in \mathcal{U}, m \in \mathcal{M}, h, h' \in \mathcal{H}, r \in \mathcal{R} \\ & \varsigma_{(t-1)u}(\omega^{[t-1]}), v_{(t-1)jc'cru}(\omega^{[t-1]}) \geq 0 \quad (92) \end{aligned}$$

The objective function (73) minimizes transportation costs, penalties for unserved patients and postponed surgeries, platelet inventory and wastage costs, ABO-compatible substitution penalties, transshipment expenses between MFs, and future costs $\mathcal{F}_{t+1}(\cdot)$. Constraint (74) assigns patients to MFs based on injury severity. On the right-hand side of

this constraint, unserved patients from the previous period are added to the patient demand of the current period. Constraints (75)–(76) allocate suitable platelets to low- and high-priority patients or postpone surgeries. On the right-hand side of these constraints, postponed surgeries from the previous period are reconsidered for platelet allocation. Constraints (77) and (78) respect the treatment capacities of ACFs and hospitals. Constraints (79) and (80) define rescue vehicle capacities for hospitals and ACFs. Constraints (81) and (82) imply that the number of extracted platelet units is limited by the number of available AMs in ACFs and hospitals. Platelet production from ABD and WBC is constrained by available apheresis and whole-blood donors per (83) and (84). Platelets exceeding $|R|$ days of storage are considered expired and must be discarded, as enforced by constraint (85). Constraints (86) and (87) cap portable AM assignments to established ACFs. Transshippable platelets are restricted by prior-period inventory, enforced by (88) and (89) for hospitals and ACFs. Constraints (90) and (91) ensure platelet flow balance at hospitals and ACFs. Their first component calculates fresh platelets ($r = 1$) from ABD, the second handles non-fresh platelets ($r > 1$) based on prior inventory and transshipments, and the third addresses platelets from WBC (exclusive to hospitals, requiring e periods). Initial inventory is included in the first period, as specified in the fourth component of (90). The right-hand side accounts for platelet needs for patient treatment and storage. Finally, constraint (92) enforces the non-negativity of the decision variables.

3.3.3 Last Stage Subproblem ($T + 1$)

On period $T + 1$, there is no need for apheresis assignment or transshipment, as this is the last stage of the planning horizon. Instead, the focus is on transferring patients to ACFs and hospitals, scheduling surgeries based on platelet availability, evaluating end-of-horizon platelet inventory and wastage, and determining whole blood and apheresis production based on available donors. The recourse function for the last period is formulated as follows:

$$\begin{aligned}
Q_{T+1}(\mathbf{Y}_T(\omega^{[T]}), \xi^{[T]}(\omega^{[T]})) = \\
\min_{j \in \mathcal{J}} \sum_{c \in \mathcal{C}} \left(\sum_{l \in \mathcal{L}} \sum_{u \in \mathcal{U}} \sum_{m \in \mathcal{M}} \pi_{lu} q_{Tjclum}(\omega^{[T]}) + \sum_{u \in \mathcal{U}} P_j \zeta_{Tjcu}(\omega^{[T]}) + \sum_{l \in \mathcal{L}} \rho_j \mu_{Tjcl}(\omega^{[T]}) \right. \\
\left. + \sum_{r \in \mathcal{R}} \sum_{u \in \mathcal{U}} \Delta_u \eta_{Tcru}(\omega^{[T]}) \right) + \sum_{u \in \mathcal{U}} \left(\psi_u \varsigma_{Tu}(\omega^{[T]}) + \sum_{\substack{j \in \mathcal{J}: \\ j \in \mathcal{J}_u}} \sum_{c' \in \mathcal{C}} \sum_{c \in \mathcal{C}} \sum_{r \in \mathcal{R}} W_{c'c} v_{Tjc'cru}(\omega^{[T]}) \right) \quad (93)
\end{aligned}$$

Subject to:

Constraints (74)-(85) and (90)-(91), where t is replaced by $(T + 1)$

$$\begin{aligned}
q_{Tjclum}(\omega^{[T]}), \mu_{Tjcl}(\omega^{[T]}), \zeta_{Tjcu}(\omega^{[T]}), \lambda_{Tcu}(\omega^{[T]}), & \quad j \in \mathcal{J}, c, c' \in \mathcal{C}, l \in \mathcal{L}, \\
\varrho_{Tch}(\omega^{[T]}), \eta_{Tcru}(\omega^{[T]}), \varsigma_{Tu}(\omega^{[T]}), v_{Tjc'cru}(\omega^{[T]}) \geq 0 & \quad u \in \mathcal{U}, m \in \mathcal{M}, h \in \mathcal{H}, r \in \mathcal{R} \quad (94)
\end{aligned}$$

Objective (93) is similar to objective function (73), but it does not include the platelet transshipment costs between facilities. Finally, constraint (94) enforces the non-negativity of decision variables.

3.4 Stochastic Dual Dynamic Programming (SDDP) and Enhancements

In this section, we introduce an enhanced SDDP algorithm (Pereira and Pinto, 1991) to efficiently solve the proposed model. SDDP decomposes multi-stage stochastic optimization problems by stage, estimating future costs using the cost-to-go function, which accounts for costs from stage $t + 1$ to the final stage. Since this function is initially unknown, the algorithm iteratively constructs an outer approximation. To manage the exponential growth of states, SDDP assumes stagewise independence of random variables, forming a recombining scenario tree where nodes represent stochastic parameter realizations and paths define scenarios. Illustrations of a stagewise independent scenario tree and its recombining counterpart are provided in Appendix B.2.

Generally, SDDP consists of three main steps. First, the *sampling procedure* selects a subset of scenarios Ξ from the full set Φ using methods such as Monte Carlo sampling. This step is crucial for SDDP, as it helps limit computational complexity while maintaining solution accuracy (line 2 of Algorithm 1). Second, the *forward pass* begins at the root node with the initial state $\mathbf{Y}_0$. The algorithm sequentially solves subproblems across the planning horizon, passing decisions as initial values to subsequent stages. This process is repeated for all sampled scenarios Ξ , ensuring that each realization of uncertainty is accounted for

(lines 3-10 of Algorithm 1). Third, during the *backward pass*, the cost-to-go functions at each stage t are approximated using hyperplanes $\mathcal{F}_{t+1}^k \in \mathcal{V}_t$. These approximations refine the cost estimates iteratively, improving solution accuracy while maintaining computational efficiency (lines 12-19 of Algorithm 1). These steps collectively enable SDDP to efficiently solve large-scale multi-stage stochastic problems while effectively managing computational complexity.

Algorithm 1 Stochastic Dual Dynamic Programming Algorithm for CRP Problem

```

1: while not converged do
2:   Sample a set of scenario paths  $\Xi$  from the entire scenario set  $\Phi$ .
3:   for each sampled scenario path  $\xi \in \Xi$  do (forward pass)
4:     Solve the first-stage model (95)-(97).
5:     Fix the values of  $x_i$  and  $\vartheta_{mi}$  for the entire planning horizon, and obtain the
       values of  $y_{1i}$ ,  $b_{1crhi}$ ,  $b'_{1crii'}$ , and  $b''_{1crhh'}$  to be used in the following stages.
6:     for stage  $t = 2, \dots, T$  do
7:       Solve each subproblem (98)-(100).
8:       Update the cost-to-go functions and pass the state variables to the next
       stage.
9:     end for
10:  end for
11:  Estimate the statistical upper bound and store the deterministic lower bound.
12:  for each sampled scenario path  $\xi \in \Xi$  do (backward pass)
13:    for stage  $t = T + 1, \dots, 2$  do
14:      for each scenario  $\omega^t \in \Omega_t$  do
15:        Modify the model (98)-(100) for stage  $t$  according to the uncertain
        parameters of scenario  $\omega^t$  and the decisions made in the previous stages (from the first
        stage to stage  $t - 1$  in the forward pass), and solve it.
16:      end for
17:      Create a new optimality cut for stage  $t - 1$  to improve the approximation
      of the cost-to-go function (e.g., the cut (159) from stage  $T + 1$  for stage  $T$  in the first
      iteration of the algorithm).
18:    end for
19:  end for
20: end while

```

In the following, we provide further details on each step. The first step, the *sampling procedure*, is explained in detail in Section 3.4.1. For the second step, the forward models for the first stage and stage- t are given in (95)–(97) and (98)–(100), respectively.

$$\min \sum_{r \in \mathcal{R}} \left(\sum_{h \in \mathcal{H}} \sum_{i \in \mathcal{I}} \pi_{hi} b_{1crhi} + \sum_{h \in \mathcal{H}} \sum_{\substack{h' \in \mathcal{H} \\ h' \neq h}} \pi_{hh'} b''_{1crhh'} + \sum_{i \in \mathcal{I}} \sum_{\substack{i' \in \mathcal{I} \\ i' \neq i}} \pi_{ii'} b'_{1crii'} \right) + F_2 \quad (95)$$

Subject to:

$$(63) - (72)$$

$$F_2 \geq \mathcal{F}_2^k(\mathbf{Y}_1(\omega^{[1]})) \quad k \in \mathcal{V}_1 \quad (96)$$

$$F_2 \geq 0 \quad (97)$$

Here, F_2 represents future costs and is calculated using the general constraints (96) and (97) defined for stage 2, or more generally, constraints (99) and (100) for any stage t . Since the hyperplane set $\mathcal{V}_1$ (used for F_2) is initially empty and built iteratively, a lower bound must be assumed, as specified by constraint (97) for the first stage, or by constraint (100) for any stage t . Similarly, for $1 < t < T + 1$, the cost-to-go function is approximated by hyperplanes $\mathcal{V}_t$. In the last stage, however, no cost-to-go function exists, as there are no future stages.

$$\begin{aligned} \min \sum_{j \in \mathcal{J}} \left(\sum_{c \in \mathcal{C}} \left(\sum_{l \in \mathcal{L}} \sum_{u \in \mathcal{U}} \sum_{m \in \mathcal{M}} \pi_{lu} q_{(t-1)jclum}(\omega^{[t-1]}) + \sum_{u \in \mathcal{U}} P_j \zeta_{(t-1)jcu}(\omega^{[t-1]}) + \sum_{l \in \mathcal{L}} \rho_j \mu_{(t-1)jcl}(\omega^{[t-1]}) \right. \right. \\ \left. \left. + \sum_{r \in \mathcal{R}} \sum_{u \in \mathcal{U}} \Delta_u \eta_{(t-1)cru}(\omega^{[t-1]}) \right) \right) + \sum_{u \in \mathcal{U}} \left(\psi_u \varsigma_{(t-1)u}(\omega^{[t-1]}) + \sum_{\substack{j \in \mathcal{J} \\ j \in \mathcal{J}_u}} \sum_{c' \in \mathcal{C}} \sum_{c \in \mathcal{C}} \sum_{r \in \mathcal{R}} W_{c'c} v_{(t-1)jc'ru}(\omega^{[t-1]}) \right) \\ + \sum_{r \in \mathcal{R}} \left(\sum_{h \in \mathcal{H}} \sum_{i \in \mathcal{I}} \pi_{hi} b_{tcrhi}(\omega^{[t-1]}) + \sum_{h \in \mathcal{H}} \sum_{\substack{h' \in \mathcal{H} \\ h' \neq h}} \pi_{hh'} b''_{tcrhh'}(\omega^{[t-1]}) + \sum_{i \in \mathcal{I}} \sum_{\substack{i' \in \mathcal{I} \\ i' \neq i}} \pi_{ii'} b'_{tcrii'}(\omega^{[t-1]}) \right) + F_{t+1} \quad (98) \end{aligned}$$

Subject to:

$$(74) - (92)$$

$$F_{t+1} \geq \mathcal{F}_{t+1}^k(\mathbf{Y}_t(\omega^{[t]})) \quad k \in \mathcal{V}_t \quad (99)$$

$$F_{t+1} \geq 0 \quad (100)$$

The final step, the *backward pass*, starts at the last stage and proceeds backward. In this step, the subproblem for each stage t is solved for all scenarios $\omega^t \in \Omega_t$, where Ω_t represents the set of possible realizations of the uncertain parameters at stage t . Cuts derived from the dual of stage t subproblems are added to the stage $t - 1$ subproblem to estimate future costs. No cuts are added in the last stage, as no future costs remain. Further details on the first hyperplane added in the backward pass and its formulation are provided in Appendix B.1. At the end of each iteration, a *convergence test* determines whether the algorithm meets stopping criteria. Common criteria include reaching a predefined maximum number of iterations (Bruno et al., 2016), a time limit (Thevenin et al., 2022), a stable lower bound (LB) (Yu et al., 2019), or achieving a difference between the LB and upper bound (UB) below a tolerance ϵ (Pereira and Pinto, 1991). Another approach estimates

the optimality gap (Shapiro, 2011) by comparing the deterministic LB (first-stage solution) with the statistical UB (average total cost of sampled scenarios in the forward pass).

To accelerate SDDP convergence and improve out-of-sample performance, we introduce six enhancements: (i) advanced scenario sampling, (ii) K-means++ scenario clustering, (iii) warm-up strategy (McDaniel and Devine, 1977), (iv) multiple cuts (Birge and Louveaux, 1988), (v) strong (Pareto-optimal) cuts (Papadakos, 2008), and (vi) lower-bounding functional valid inequalities (Batun et al., 2011).

3.4.1 Advanced Scenario Sampling Method

The full set of uncertain parameter scenarios is often too large to solve the stochastic problem directly. As a result, the SDDP algorithm typically relies on a subset of scenarios. However, improving model accuracy requires exploring more advanced scenario generation methods that better capture the stochastic nature of the future. For example, the Quasi-Monte Carlo (QMC) method, unlike traditional ones, uses low-discrepancy sequences to generate points that cover the space more uniformly, leading to lower variance in the estimates of integrals or expected values. Nevertheless, QMC alone lacks the randomness needed for unbiased estimates, which is where RQMC method comes into play (Cranley and Patterson, 1976). RQMC introduces randomness into the low-discrepancy sequences while retaining their uniformity, thereby combining the benefits of reduced variance with the ability to produce unbiased estimators and confidence intervals. This makes RQMC particularly effective for generating more accurate approximations with fewer scenarios. In this research, to leverage the advantages of RQMC, we utilize the Halton sequence — a popular low-discrepancy sequence — and randomize it to achieve the desired balance of uniformity and randomness.

3.4.2 Scenario Clustering Method using K-means++

Clustering groups data into clusters to maximize intra-cluster similarity and inter-cluster dissimilarity. K-means (Lloyd, 1982) and K-means++ (Arthur and Vassilvitskii, 2006) are popular methods for partitioning scenarios based on feature similarity. In our largest instance, we reduce scenarios from 10^5 to 10^4 by minimizing the intra-cluster sum of squared distances. This scenario clustering analysis partitions the generated scenarios and retains a single representative scenario from each cluster, effectively reducing the scenario set size while preserving diversity. In K-means, scenarios are treated as n -dimensional points in

Euclidean space, but its random initialization of cluster centers can lead to suboptimal solutions. K-means++ improves this by selecting the first k centers more systematically: the first is random, while subsequent centers are chosen with probabilities proportional to their contribution to the error. This enhances solution quality and convergence. We implemented K-means++ in the SDDP algorithm using *scikit-learn* (Pedregosa et al., 2011). The integration of scenario clustering into the SDDP algorithm is inspired by Emelogu et al. (2016), who applied the K-means algorithm within the classical sample average approximation method. We apply K-means++ to group scenarios into k clusters, with SDDP calculating the expected objective value using cluster centroids instead of the full scenario set. In Section 3.5.2, we compare K-means++ with standard K-means and the Self-Organizing Map (SOM) clustering method, an unsupervised neural network approach introduced by Kohonen (1982). For further details on using K-means, K-means++, and SOM for scenario clustering, see Emelogu et al. (2016) and Rahmati et al. (2023).

3.4.3 Warm-up and Multiple Cuts Strategies

In this section, we present two enhancements designed to improve our SDDP algorithm. The first enhancement involves a warm-up strategy where we initially solve the first-stage MILP model (95)-(97) as a linear program. This linear relaxation is executed for a predefined number of iterations, denoted by α^{warm} , or until the solution gap remains constant for more than β^{warm} consecutive iterations. Once either of these conditions is met, the integrality constraints are reintroduced, and the model is solved as an MILP.

The second enhancement incorporates a multiple cuts strategy during the backward step. There are generally two approaches for adding optimality cuts. The first approach, known as the single cut strategy, aggregates the information from all scenarios at stage $t + 1$ into a single cut, which is then applied to stage t . In contrast, the multiple cuts strategy computes the cost-to-go function $F_{t+1}(\omega^t)$ separately for each scenario ω^t at stage t . Each scenario-specific cut is then added to the previous stage t , and the expected cost-to-go function is calculated as $F_{t+1} = \sum_{\omega^t \in \Omega_t} p(\omega^t) F_{t+1}(\omega^t)$. For a visual explanation of these strategies, refer to Figures (B.4) and (B.5), provided in Appendix B.3. In this research, we adopt the multiple cuts strategy, as it has been shown to be more efficient (Batun et al., 2011).

3.4.4 Strong (Pareto-optimal) Cuts

In BD, it is common to encounter multiple optimal solutions when solving subproblems, particularly in network optimization scenarios like facility location problems. While all these optimal solutions can generate valid cuts for the master problem, using the strongest cuts, known as Pareto-optimal cuts, can significantly reduce the number of iterations and the overall solution time. Although Pareto-optimal cuts are well-known in BD, their application in SDDP remains limited.

In the traditional Magnanti-Wong procedure, solving the dual subproblem precedes addressing the Magnanti-Wong problem in each iteration, which can hinder performance when the dual subproblem is computationally intensive (Magnanti and Wong, 1981). To overcome this, we adopt the independent Magnanti-Wong algorithm from Papadakos (2008), which generates Pareto-optimal cuts directly, bypassing the dual subproblem. This method is implemented in the backward pass of the SDDP algorithm to address a multi-stage stochastic MILP. Specifically, when solving the subproblem of stage t , state variables from $t - 1$ are not fixed to forward pass decisions but to approximate core points, defined as convex combinations of optimal solutions from $t - 1$ in current and previous iterations. For the first iteration, the core points $\mathbf{Y}_t^{Core}$ are set to the forward pass solution. In subsequent iterations α , they are updated as: $\mathbf{Y}_t^{Core} = \phi \mathbf{Y}_t^{\alpha-1} + (1 - \phi) \mathbf{Y}_t^\alpha$, where $\phi < 1$ is set to 0.5 in our implementation.

3.4.5 Lower Bounding Functional Valid Inequalities (LBFVIs)

In this section, we introduce LBFVIs designed to enhance the lower bound achieved by the first-stage model. The core concept behind these inequalities is to provide the first-stage problem with an approximate estimation of the second-stage cost, using the information from the average scenario path. Birge and Louveaux (2011) demonstrated that, for a given stochastic problem SP , the cost associated with solving SP using the scenario path of the expected values of the uncertain parameters is less than or equal to the optimal solution of SP . Consequently, the expected cost-to-go function $\mathcal{F}_{t+1}(\cdot)$ from period $t + 1$ to the final period is always greater than or equal to the cost-to-go function derived from the average scenario path. The mathematical formulation of these constraints and their detailed structure are provided in Appendix B.4.

3.5 Computational Results

We conducted the computational tests on a machine with two 2.40 GHz Intel Xeon Gold 6148 processors, using Python and Gurobi 11.0.1, with each instance running on a single core. Test instances varied in planning horizon, ranging from 72 hours post-disaster (3-stage model) to 5 days (5-stage model), aligning with platelet shelf life. This setup evaluates algorithm performance across different post-disaster timeframes. In the SDDP algorithm’s forward step, we use a randomly selected sample path, though the algorithm can handle any subset of possible scenarios. To provide a detailed evaluation, Section 3.5.1 introduces the test instances, while Sections 3.5.2 and 3.5.3 examine the algorithm’s enhancements in out-of-sample and in-sample scenarios, respectively. After reaching the stopping criterion (a time limit of $5|T|$ hours), we use Monte Carlo simulation to construct a 95% confidence statistical upper bound for in-sample performance. For instances with 1,000 or fewer in-sample scenarios, the upper bound is calculated from the full set; otherwise, 1,000 scenarios are sampled. To assess out-of-sample performance, we simulate 2,500 scenario paths representing potential realizations of uncertain parameters over the planning horizon. These paths, sampled independently via Monte Carlo based on probability distributions, are used consistently across methods and instances.

3.5.1 Implementation Details

In this section, we provide a detailed explanation of how CRP instances with uncertain demand, platelet donors, and hospital treatment capacity were generated. The majority of our instances are based on data from Oksuz and Satoglu (2020), Sun et al. (2021), and Kamyabniya et al. (2021). However, as we introduce certain aspects in CRPs for the first time, we needed to generate new parameters, which are described later in this section.

Instances vary by the number of ACFs ($i \in \{5, 10, 15, 20\}$), demand locations ($l \in \{5, 10, 15, 20\}$), and five hospitals with capacities ranging from 50 to 300 beds. Three rescue vehicle types are considered, with capacities of 1, 2, and 3 patients per trip (Rambha et al., 2021). The model accounts for four ABO blood groups. Demand locations, ACFs, and hospitals are uniformly distributed in a 50 km by 50 km area, with patient numbers generated from $U[0, 250]$ for each demand location. Donors for WBC and ABD methods are also uniformly distributed, with WBC donors drawn from $U[200, 600]$ and ABD donors from $U[100, 500]$ per region and period (Saha, 2022). Patients and donors are allocated to blood groups based on Canadian Blood Services (2023), and casualties are

categorized by priority: 30% high, 30% medium, and 40% low. Platelets have a maximum shelf life of 5 days. At the beginning of each period, transshipment decisions are made based on the previous period’s platelet inventory levels. WBC-derived platelets require 1 day for processing (Samani and Hosseini-Motlagh, 2023), while ABD-derived platelets are available for use on the same day (Hosseini-fard et al., 2020). Each apheresis session takes 90 minutes per donor, with 40 portable AMs available nationwide for post-disaster ACF assignment. Platelet demand varies by priority: high-priority patients require 3 units per surgery, medium-priority 2 units, and low-priority 1 unit. Remaining parameters are listed in Table 7. Benchmark instances and results are also available at <https://github.com/Pedram-Fdi/DynamicResAllocCRP.git>.

Table 7: Parameters Used for Instance Generation

<i>Parameter</i>	<i>Value</i>	<i>Parameter</i>	<i>Value</i>
B	85% of total required budget for all ACFs	τ	8
c_i^{ACF}	$U[100, 400]; \forall i \in \mathcal{I}$	f_i	$c_i^{\text{ACF}} \times 1.3 \times 25; \forall i \in \mathcal{I}$
A_{cru}	$50; \forall c \in \mathcal{C}, r \in \mathcal{R} \text{ if } u \in H, \text{ otherwise } 0$	ρ_j	$\{150000, 112500, 75000\}; \forall j \in \{1, 2, 3\}$
P_j	$\{9000, 4500, 3000\}; \forall j \in \{1, 2, 3\}$	Δ_u	$10; \forall u \in \mathcal{U}$
ψ_u	$2 \times \Delta_u; \forall u \in \mathcal{U}$	n_h^{Aph}	$10; \forall h \in \mathcal{H}$

3.5.2 Performance on Out-of-Sample Scenarios

This section evaluates the performance of scenario sampling and clustering methods for the SDDP algorithm on 2,500 out-of-sample scenarios. Table 8 compares the SDDP algorithm with all enhancements (Sections 3.4.3–3.4.5) using these combinations: *QNSDDP*: QMC with no clustering, *RNSDDP*: RQMC with no clustering, *RSSDDP*: RQMC with SOM clustering, *RKSDDP*: RQMC with K-means clustering, *QKSDDP++*: QMC with K-means++ clustering, *RKSDDP++*: RQMC with K-means++ clustering.

The column “*Prb. Size*” in Table 8 describes each instance, including the number of stages (“ST”), total scenarios (“ $|S|$ ”), ACF locations (“ $|\mathcal{I}|$ ”), and demand locations (“ $|\mathcal{L}|$ ”). Each row in this table presents the average values across five instances. Key performance indicators (KPIs) include: - “DT” (Delayed Transfers): the average number of patients not transferred on time, - “DS” (Delayed Surgeries): the average number of surgeries delayed, - “DB” (Different Blood Type): the average number of patients who received a different platelet type than their own during surgery. The “*Gap*” column shows the performance gap, $Gap = 100 \times \left(\frac{UB(m) - UB^*}{UB^*} \right)$, where $UB(m)$ is the upper bound for each method and UB^* is the best approximated upper bound among all methods. The “VSS (%)” column

reports the value of the stochastic solution as $VSS = 100 \times \left(\frac{EEV - UB}{EEV} \right)$, with EEV as the expected value solution. Bold values indicate the best results in each column.

According to Table 8, $RKSDDP^{++}$ is the most effective method, combining RQMC sampling with K-means++ clustering to achieve the best average objective function and lowest delays. $RKSDDP$ also performed well but was slightly less effective compared to $RKSDDP^{++}$. In contrast, $QNSDDP$ performed the worst, showing the highest delays, mismatched platelet transfusions, and objective gaps. Comparing $QKSDDP^{++}$ to $QNSDDP$ highlights K-means++ clustering’s impact, reducing the objective gap by 96.82% even without RQMC sampling. $RKSDDP^{++}$ consistently achieved the lowest DT and balanced all KPIs effectively, excelling in timely evacuations and minimizing DS and DB. While $RKSDDP$ occasionally surpassed $RKSDDP^{++}$ on specific KPIs, $RKSDDP^{++}$ remained the overall best performer. Notably, DS has a higher weight than DB in the objective function due to its critical importance, though this weighting can be adjusted based on managerial needs. Finally, the average VSS of 40% underscores the importance of stochasticity in modeling multi-stage CRPs. In the next section, we evaluate the impact of other enhancements on the SDDP algorithm.

3.5.3 Performance of the Algorithms on In-Sample Scenarios

This section evaluates the effectiveness of SDDP algorithm enhancements by analyzing the optimality gap across $\max(|S|, 1000)$ in-sample scenarios. As shown in Section 3.5.2, the default $RKSDDP^{++}$ algorithm, combining RQMC sampling, K-means++ clustering, and enhancements from Sections 3.4.3, 3.4.4, and 3.4.5, performed best on out-of-sample scenarios. Here, we assess the individual impact of each enhancement on the algorithm’s optimality gap. Table 9 compares $RKSDDP^{++}$ (Default), nested Benders decomposition (“NBD”), and the default algorithm with each enhancement removed individually: NoSC (without strong cuts), NoMC (without multiple cuts), NoLBFVIs (without LBFVIs), NoWU (without the warm-up strategy), and NoEnh. (without any enhancements). The table reports the average lower bound (LB) and 95%-confidence statistical upper bound (UB) for the default algorithm, along with the optimality gap, $Gap = 100 \times \left(\frac{UB - LB}{UB} \right)$, for all SDDP variants, averaged over five instances. Instances where memory constraints prevented completion are marked as “OM”. Bold values indicate the best results among the different SDDP variants. Figures (B.6a) and (B.6b) in Appendix B.5 show the LB progression at each iteration for various algorithm variants.

Table 8: Out-of-Sample SDDP Performance with Various Scenario Generation and Clustering Methods.

Prb. Size				RNSDDP				RKSDDP ⁺⁺				QNSDDP				RSSDDP				RKSDDP				QKSDDP ⁺⁺				VSS	
ST	S	J	L	DT	DS	DB	Gap (%)	DT	DS	DB	Gap (%)	DT	DS	DB	Gap (%)	DT	DS	DB	Gap (%)	DT	DS	DB	Gap (%)	DT	DS	DB	Gap (%)	(%)	
3	10 ²	5	5	2	3	5	12.96	1	3	4	5.08	462	463	465	19,663	3	4	6	39.82	3	4	6	46.12	15	16	18	348	95.76	
			10	7	127	21	18.44	3	123	16	4.21	951	1,065	992	4,276	5	125	19	11.16	3	123	17	6.83	289	404	304	1,104	70.31	
			15	6	370	82	5.54	4	369	79	2.87	1,366	1,698	1,520	2,037	4	368	80	3.23	1	366	78	0.54	837	1,172	943	1,164	65.76	
			20	48	775	352	0.71	47	774	350	0.22	1,966	2,601	2,304	1,173	49	777	349	1.58	47	775	347	0.51	1,360	2,012	1,684	749	50.51	
		10	5	2	3	4	25.29	1	2	3	1.17	476	478	480	21,603	1	2	4	18.40	2	3	5	43.80	4	5	6	125	94.14	
			10	1	120	14	0.55	1	120	14	0.09	931	1,045	971	4,253	1	120	14	0.02	1	120	14	0.06	76	193	92	281	64.62	
			15	1	367	77	0.02	1	367	76	0.01	1,269	1,602	1,425	1,875	3	369	80	2.91	4	370	80	4.07	274	622	376	361	60.90	
			20	2	740	380	0.67	1	739	338	0.06	1,812	2,450	2,153	1,211	1	740	333	0.06	1	740	338	0.13	1,019	1,682	1,295	645	40.12	
		15	5	1	2	3	0.04	1	2	3	0.02	458	460	461	23,118	1	2	3	0.03	1	2	3	0.01	1	2	3	0.05	85.88	
			10	1	120	13	0.01	1	120	13	0.00	914	1,027	956	4,219	1	120	14	0.32	1	120	13	0.00	30	148	44	105	65.98	
			15	1	367	80	0.04	1	368	77	0.01	1,372	1,705	1,529	2,043	1	367	83	0.10	1	368	82	0.03	560	901	664	763	47.76	
			20	2	743	305	0.57	1	742	310	0.02	1,786	2,425	2,128	1,198	2	742	321	0.28	1	743	319	0.01	921	1,591	1,242	582	36.50	
	20	5	1	2	3	0.02	1	2	3	0.02	470	472	473	22,368	1	2	3	0.03	1	2	3	0.03	2	4	5	51.29	93.62		
		10	2	122	15	4.34	1	120	13	0.01	811	925	853	3,646	1	120	13	0.01	1	120	13	0.00	2	122	18	4.47	39.39		
		15	0	367	76	0.04	0	367	77	0.02	253	609	418	359	9	374	87	10.55	0	366	87	0.05	0	366	143	0.18	37.73		
		20	1	745	313	0.02	1	745	309	0.01	167	890	873	98.54	1	744	321	0.03	1	745	310	0.00	1	742	376	0.18	33.44		
	3-ST Average			5	311	109	4.33	4	310	105	0.86	967	1,245	1,125	7,071	5	311	108	5.53	4	310	107	6.39	337	624	451	393	61.40	
	4	10 ³	5	5	1	3	18	13.19	0	3	18	3.00	583	584	603	31,832	0	3	18	2.24	1	4	19	26.89	2	4	19	48.06	90.28
				10	8	313	91	6.58	3	308	85	1.72	1,192	1,403	1,332	2,233	3	309	85	1.75	4	308	85	2.29	280	553	386	452	49.43
				15	32	828	198	12.82	13	812	176	4.96	1,908	2,472	2,268	1,228	7	807	170	2.67	6	807	167	2.23	861	1,520	1,107	488	28.13
				20	68	1,465	378	0.37	66	1,464	378	0.17	2,494	3,511	2,940	709	67	1,463	377	0.17	67	1,464	381	0.29	2,347	3,378	2,787	660	14.04
			10	5	0	3	18	0.13	0	3	17	0.05	428	430	446	21,366	0	3	18	0.06	1	3	18	14.97	0	3	17	0.15	68.98
				10	0	304	83	0.09	0	304	84	0.02	612	857	735	1,094	1	305	84	0.41	0	304	83	0.07	1	304	94	0.52	25.66
				15	7	799	176	2.55	0	792	180	0.05	1,285	1,895	1,591	768	0	794	169	0.01	0	796	167	0.01	722	1,424	950	455	35.01
20				25	1,423	364	4.73	2	1,402	344	0.06	1,748	2,867	2,241	539	20	1,420	349	3.62	2	1,403	316	0.00	641	1,957	1,196	172	20.56	
15			5	0	3	19	0.12	0	3	19	0.04	396	398	416	20,452	0	3	19	0.04	0	3	19	0.08	0	3	19	0.25	62.47	
			10	0	303	86	0.02	0	303	84	0.01	476	724	607	828	0	303	84	0.00	0	302	85	0.06	1	303	104	0.89	12.29	
			15	26	819	194	10.80	0	794	170	0.01	585	1,290	865	350	1	791	176	0.16	0	792	171	0.07	0	788	239	0.06	25.03	
			20	53	1,448	425	10.60	2	1,400	335	0.16	1,510	2,670	2,032	479	26	1,418	350	4.97	2	1,397	326	0.01	250	1,607	823	67.12	15.66	
20		5	1	3	19	7.07	0	3	19	0.03	320	321	339	16,002	0	3	19	0.03	0	3	19	0.02	0	3	19	0.47	61.77		
		10	0	301	88	0.01	0	301	90	0.01	779	1,005	914	1,389	0	300	90	0.00	0	301	89	0.01	0	300	113	0.09	13.43		
		15	0	786	183	0.04	0	790	171	0.01	0	787	267	0.06	0	789	185	0.06	4	786	184	1.53	0	785	244	0.07	14.75		
		20	3	1,402	330	0.33	2	1,402	329	0.02	32	1,425	651	6.36	19	1,415	350	3.49	2	1,403	346	0.00	2	1,397	620	0.09	14.31		
4-ST Average			14	638	167	4.34	6	630	156	0.64	897	1,415	1,140	6,205	9	633	159	1.23	6	630	155	3.03	319	896	546	146	34.49		
10 ⁴		5	5	1	5	39	29.73	0	4	38	0.05	495	497	544	24,627	1	5	38	14.71	1	5	38	29.07	1	5	39	31.06	81.04	
			10	14	541	198	6.53	7	536	188	2.63	1173	1486	1415	1,152	14	541	197	6.95	6	534	186	1.59	6	531	211	2.06	39.87	
			15	124	1392	388	7.81	91	1363	358	1.65	652	1815	1017	137	110	1381	379	5.08	99	1367	366	3.17	674	1830	1026	139	12.35	
			20	294	2300	641	0.87	290	2296	634	0.47	1595	3365	2164	177	306	2316	653	2.18	289	2296	634	0.38	2246	3854	2739	289	6.91	
		10	5	0	4	35	0.08	0	4	35	0.01	278	280	323	11,897	0	4	36	0.96	0	4	35	0.03	0	4	36	0.23	43.22	
			10	0	527	189	0.08	0	527	186	0.00	904	1271	1141	888	0	522	188	0.06	0	523	190	0.09	7	525	212	4.03	18.48	
			15	26	1281	304	4.54	7	1260	290	0.40	116	1349	489	28.04	60	1308	344	12.61	5	1263	284	0.12	111	1346	411	26.30	16.50	
	20		87	2102	438	1.89	82	2093	441	1.39	208	2215	806	15.73	95	2112	456	2.67	87	2101	449	1.94	823	2768	1383	95.13	10.78		
	15	5	0	4	36	0.15	0	4	35	0.02	373	375	421	17,771	0	4	35	0.05	0	4	36	0.03	0	4	36	0.29	55.64		
		10	0	525	179	0.01	0	178	526	0.00	764	1148	1005	726	7	531	193	4.50	0	529	179	0.01	0	516	199	0.06	12.01		
		15	14	1271	301	2.05	4	1263	288	0.01	102	1349	462	23.28	23	1274	310	3.96	5	1264	291	0.28	5	1260	321	0.19	15.25		
		20	84	2094	450	1.80	70	2081	431	0.33	76	2088	709	1.04	95	2110	481	3.01	66	2078	422	0.00	391	2378	949	39.14	7.48		
20	5	0	5	36	0.10	0	5	35	0.02	450	451	499	20,938	0	5	35	0.22	0	5	35	0.02	0	5	36	0.86	27.38			
	10	0	523	187	0.01	0	456	249	0.00	1029	1363	1269	994	1	521	195	0.48	0	524	183	0.01	0	525	201	0.05	6.99			
	15	5	1239	298	0.01	5	1243	296	0.01	6	1238	364	0.23	5	1236	299	0.02	5	1242	303	0.01	5	1233	333	0.02	7.80			
	20	36	2065	373	0.00	36	2068	372	0.00	37	2066	597	0.05	37	2063	388	0.12	32	2070	368	0.01	209	2230	757	20.08	5.53			
5-ST Average			43	992	256	3.48	37	961	275	0.44	516	1397	827	4,961	47	996	264	3.60	37	988	250	2.30	280	1188	556	40.58	22.95		
Total Average			21	647	177	4.05	16	634	179	0.65	793	1352	1031	6,079	20	647	177	3.45	16										

The default algorithm, $RKSDDP^{++}$, achieved the best performance, with an average in-sample optimality gap of 0.86%. The NBD algorithm struggled with scalability, failing to solve instances with more than three stages due to computational limitations. Among the tested variants, $NoEnh.$ and $NoLBFVIs$ experienced the highest memory issues, with 37.5% and 16.6% of instances encountering out-of-memory errors, highlighting the importance of valid inequalities in improving efficiency. When evaluating individual enhancements, removing strong cuts and the warm-up strategy slightly increased the gap to 0.98% and 1.19%, respectively, suggesting a limited but positive impact on solution quality. However, excluding multi-cuts significantly worsened performance, with the gap increasing to 55.66%, underscoring their critical role in convergence and accuracy. Similarly, $NoLBFVIs$ resulted in an 83.06% gap, the largest among algorithms that solved all instances, compared to 74.43% for $NoEnh.$, though OM errors prevent direct comparison. Overall, these results demonstrate that while some enhancements, like strong cuts and warm-up, contribute incrementally, multi-cuts and LBFVIs are essential for achieving robust and scalable performance.

Table 9: In-Sample SDDP Performance with Various Enhancements.

Prb. Size				RKSDDP ⁺⁺ (Default)			NBD	NoSC	NoMC	NoLBFFVs	NoWU	NoEnh.		
ST	S	I	L	LB	UB	Gap								
3	10 ²	5	5	120,865	120,865	0.00	1.52	0.00	34.25	48.90	0.00	45.64		
			10	2,097,727	2,097,727	0.00	0.73	0.00	86.70	98.43	0.00	98.47		
			15	8,669,450	8,669,450	0.00	0.00	0.00	57.79	80.21	0.00	79.26		
			20	34,267,956	35,015,929	2.43	0.26	3.00	65.55	OM	2.54	OM		
		10	5	119,035	119,094	0.05	3.09	0.12	0.74	0.24	0.09	0.10		
			10	2,094,254	2,094,965	0.03	4.14	0.05	27.31	100.00	0.04	OM		
			15	5,969,306	5,969,306	0.00	4.78	0.00	44.93	100.00	1.01	OM		
			20	13,599,365	13,599,365	0.00	0.00	0.00	32.98	OM	0.00	OM		
		15	5	115,923	115,924	0.00	10.63	0.09	12.34	32.79	0.00	99.94		
			10	2,088,293	2,089,264	0.05	3.97	0.09	43.73	100.00	0.06	100.00		
			15	5,830,955	5,859,627	0.47	4.64	0.88	54.08	100.00	3.54	100.00		
			20	12,285,550	12,285,550	0.00	0.00	0.00	49.71	100.00	0.00	OM		
	20	5	109,097	111,594	2.23	23.45	2.31	31.74	OM	2.28	100.00			
		10	2,056,089	2,058,326	0.11	1.63	0.13	79.26	100.00	0.10	100.00			
		15	5,829,917	5,919,517	1.47	3.19	3.80	81.76	100.00	3.93	100.00			
		20	12,152,260	12,152,260	0.00	0.00	0.00	82.44	OM	0.00	OM			
	3-ST Avg.			6,712,878	6,767,423	0.43	3.88	0.65	49.08	80.05	0.85	82.34		
	4	10 ³	5	5	149,209	149,330	0.08	23.79	0.07	40.19	68.83	2.01	56.99	
				10	4,268,061	4,342,205	1.71	OM	1.71	84.11	99.49	1.70	OM	
				15	13,907,119	14,017,503	0.79	OM	0.79	72.16	92.90	0.79	OM	
				20	34,941,724	35,537,720	1.62	OM	1.69	81.82	OM	1.67	OM	
			10	5	144,142	146,196	1.40	OM	3.42	26.04	4.52	2.51	5.14	
				10	4,022,006	4,100,395	1.91	OM	1.90	34.05	63.13	1.90	3.99	
				15	12,110,120	12,210,922	0.83	OM	0.79	64.18	100.00	0.84	OM	
20				22,992,122	23,123,729	0.57	OM	0.60	51.36	OM	0.85	OM		
15			5	138,781	141,677	2.04	OM	2.00	41.25	45.98	2.07	23.95		
			10	4,129,196	4,203,767	1.78	OM	1.77	76.58	100.00	1.77	100.00		
			15	12,095,596	12,198,062	0.84	OM	0.83	86.71	100.00	0.84	100.00		
			20	26,419,265	26,630,050	0.76	OM	0.60	83.50	100.00	0.61	OM		
20		5	137,207	140,355	2.24	OM	2.20	32.07	100.00	2.24	82.15			
		10	3,991,995	4,063,748	1.76	OM	1.76	81.44	100.00	1.77	100.00			
		15	12,203,201	12,333,763	1.06	OM	1.01	90.68	100.00	1.01	100.00			
		20	22,336,014	22,460,981	0.56	OM	0.61	86.61	100.00	0.61	OM			
4-ST Avg.			10,874,110	10,987,525	1.25	23.79	1.36	64.55	83.92	1.45	63.58			
5		10 ⁴	5	5	170,797	171,128	0.19	OM	0.16	71.86	76.21	0.15	81.06	
				10	6,758,625	6,832,588	1.08	OM	1.10	79.89	99.51	1.09	99.80	
				15	19,298,462	19,428,567	0.67	OM	0.51	77.30	98.16	0.51	98.12	
				20	41,811,533	42,139,418	0.60	OM	0.60	90.79	OM	0.37	OM	
			10	5	165,810	168,102	1.37	OM	1.53	54.58	0.38	1.60	40.95	
				10	6,427,143	6,501,177	1.14	OM	1.12	1.24	80.20	1.13	40.91	
				15	18,533,072	18,636,703	0.55	OM	0.52	22.91	100.00	0.66	OM	
	20			33,087,669	33,265,560	0.53	OM	0.47	46.83	OM	0.46	OM		
	15		5	161,435	164,983	2.15	OM	2.19	56.13	50.29	2.26	18.99		
			10	6,494,232	6,569,726	1.15	OM	1.13	70.33	100.00	1.15	100.00		

Table 9: In-Sample SDDP Performance with Various Enhancements (continued).

ST	Prb. Size			RKSDDP ⁺⁺ (Default)			NBD	NoSC	NoMC	NoLBFVIs	NoWU	NoEnh.
	S	I	L	LB	UB	Gap						
			15	18,199,602	18,298,179	0.54	OM	0.99	18.80	100.00	0.61	OM
			20	33,434,599	33,590,296	0.46	OM	0.39	20.75	100.00	0.45	OM
		20	5	165,363	168,454	1.86	OM	1.81	87.08	100.00	1.88	71.29
			10	6,482,476	6,554,621	1.10	OM	1.10	73.66	88.13	1.12	100.00
			15	19,203,464	19,338,511	0.70	OM	0.68	60.73	100.00	0.73	100.00
			20	31,374,605	31,515,965	0.45	OM	0.64	20.79	100.00	6.03	100.00
		5-ST Avg.		15,110,555	15,208,999	0.91	OM	0.93	53.35	85.21	1.26	77.37
Total Average				10,899,181	10,987,982	0.86	13.83	0.98	55.66	83.06	1.19	74.43

3.6 Case Study

To evaluate the effectiveness of our proposed algorithm, we tested it on a CRP problem based on real data from the 2011 earthquake in Van Province, Turkey. This study focuses on identifying optimal locations for ACFs and dynamic resource allocation within the critical first 72 hours after the disaster. The case study includes 94 demand neighborhoods, 31 potential ACF locations, and four hospitals with uncertain capacities. Patient demand and donor availability are generated using uniform distributions, reflecting realistic disaster scenarios. See Appendix B.6 for a detailed description of the instance setup and intensity levels.

To assess performance, we conducted three analyses: (1) using deterministic parameters (**Det**), (2) using a stochastic approach without transshipment (**StoNoT**), and (3) using a stochastic approach with transshipment (**StoT**). The spatial distribution and allocation of ACFs under each model are illustrated in Appendix B.6. Below, we discuss the results.

Comparison of Patient Delays and Treatment Quality: From the results, we observe a substantial reduction in patient transportation and surgery delays when using stochastic models, particularly in **StoT**. For example, in **Det**, the number of patients experiencing transport delays to medical facilities is approximately 7.6 times higher than in **StoNoT** and 23 times higher than in **StoT**. This significant reduction highlights the advantages of accounting for uncertainty in platelet supply, demand, and treatment capacity, as well as enabling transshipments between facilities in **StoT**. Moreover, the high average VSS value of 24.64% emphasizes the critical importance of stochasticity in disaster management. A similar trend is evident when considering surgery delays. **Det** shows a much higher frequency of delayed surgeries compared to the stochastic models, with **StoT** performing the best in minimizing delays. This pattern also extends to blood transfusions where incorrect blood type infusions occur far more frequently in **Det**. By addressing uncertain demand through a multi-stage stochastic approach and enabling transshipment, **StoT** demonstrates

a higher capacity to meet patient needs promptly and accurately, thereby reducing potential medical complications associated with delays and mismatched platelet types.

Impact on ACF Establishment and Resource Allocation: In a setting with 94 demand points and 31 potential ACF locations, **Det** establishes 16 ACFs, while **StoNoT** and **StoT** establish 14 and 13 ACFs, respectively. This reduction in ACFs when transshipment is enabled indicates that resource sharing between facilities decreases the need for additional ACFs, which would otherwise be required to meet fluctuating demands individually. This result underscores transshipment’s role in CRP, allowing time-sensitive resources, such as platelets, to be redistributed efficiently among medical facilities.

Resource Allocation Efficiency: While **Det** requires a larger number of ACFs to be established, the total number of rescue vehicles allocated to ACFs in **StoNoT** is around 22% higher compared to **Det** (56 versus 46 vehicles). **StoT**, however, significantly improves allocation efficiency, deploying approximately double the number of rescue vehicles as **Det** (93 versus 46) by concentrating them strategically across fewer ACFs. This efficient allocation in **StoT** reflects the model’s ability to adapt to variable demand and maximize rescue vehicle utilization through transshipment. Overall, these results underscore the superiority of **StoT** over **StoNoT**, and **StoNoT** over **Det**. **StoT** stands out for its ability to dynamically respond to demand fluctuations through both effective resource allocation and transshipment. By allowing facilities to share platelet inventories, **StoT** provides a critical advantage in terms of patient coverage and resource flexibility, reducing the need for additional ACFs while ensuring quick and efficient medical support.

3.7 Conclusion

In this paper, we tackled the CRP problem, focusing on managing the deployment of new portable apheresis machine technologies. To address this complex problem, we introduced the SDDP algorithm and benchmarked its performance against the NBD algorithm. To substantially enhance the efficiency of these methods, we incorporated a range of acceleration techniques, including strong cuts, multiple cuts, LBFVIs, and warm-up strategies. Additionally, we conducted a comprehensive computational study on scenario sampling methods by combining QMC and RQMC sampling with four clustering strategies: no clustering, SOM, K-means, and K-means++. This resulted in eight different combinations of sampling and clustering methods. Among these, the combination of RQMC sampling

and K-means++ clustering proved to be the most effective, significantly reducing the number of scenarios while preserving accuracy. Together, these enhancements resulted in our most efficient algorithm, denoted as $RKSDDP^{++}$. While the NBD algorithm encountered memory limitations for instances with more than three stages, $RKSDDP^{++}$ successfully solved all tested instances, achieving an average optimality gap of 0.86%. To demonstrate the practical applicability of our proposed models and algorithms, we tested $RKSDDP^{++}$ on a case study from Turkey, achieving a near-optimal solution with an average optimality gap of 1.96%. Future research could extend this model by accounting for road conditions and travel times in post-disaster settings, incorporating debris clearance as a critical component of CRP. Additionally, exploring the effects of secondary disasters on hospitals within the CRP context could provide valuable insights into their potential impact on response strategies.

Chapter 4

A Deep-Reinforcement-Learning-Based Matheuristic for Stochastic Casualty Response Planning with Hospital Evacuation

Abstract

Hospital evacuation is a critical yet overlooked component of casualty response planning during large-scale disasters. We propose a two-stage stochastic programming model that jointly addresses the location of alternative care facilities, allocation of land and aerial rescue vehicles, and assignment of backup hospitals for disrupted facilities. In the first stage, these strategic decisions are made immediately after a disaster using preliminary information. In the second stage, once uncertainty is realized, such as casualty distribution, hospital disruptions, and patients requiring evacuation, we determine casualty transfers and patient evacuations over time. The objective is to minimize evacuation threat risk, transportation risk, and associated operational costs. As the solution method, we develop a learning-based matheuristic that combines deep reinforcement learning with adaptive large neighborhood search. Scenario diversity is preserved through decision-based clustering, and randomized quasi-Monte Carlo sampling enhances solution quality. Extensive computational experiments demonstrate that our method maintains a 0.52% optimality gap on small instances and outperforms the Progressive Hedging Algorithm on large-scale instances, achieving up to an 83.24% cost improvement across out-of-sample scenarios where exact solvers fail. A real-world case study based on the 2011 Van earthquake in Turkey demonstrates the practical applicability of our approach, showing that it can reduce

the number of non-evacuated patients by up to 73.31% compared to a deterministic model, thereby potentially saving hundreds of additional lives in large-scale disasters.

4.1 Introduction

Catastrophic health events (CHEs), whether caused by natural forces such as floods, hurricanes, or earthquakes, or by human actions such as radiological incidents, pose significant threats to human life, infrastructure, and economic stability. For example, the 2023 Turkey-Syria earthquakes resulted in over 56,000 deaths and 100,000 injuries, while the 2011 earthquake in Japan caused more than 15,000 fatalities, over 6,000 injuries, 2,585 missing persons, and an estimated \$235 billion in economic damage (Amani et al., 2025). In the aftermath of such disasters, a large number of injured individuals require immediate medical attention. However, damaged healthcare infrastructure, including hospitals that may be non-functional or in need of evacuation, places extreme strain on response systems. This makes the efficient use of limited medical resources within the first few hours following a disaster absolutely critical. Casualty response planning (CRP) plays a key role in this process, coordinating multiple agencies to allocate resources, treat the injured, and transport patients to operational medical facilities.

Within CRP, repurposing existing public spaces, such as school grounds, parks, or stadiums, into alternative care facilities (ACFs) is a common and practical strategy for coping with sudden surges in casualty numbers. ACFs are temporary facilities that can be rapidly deployed in disaster-affected areas to provide emergency health care to the impacted population. Their effectiveness, however, depends heavily on strategic decisions regarding their locations and resource allocation. For instance, after the 1995 Sarin attack in Tokyo's subway system, the lack of a coordinated patient distribution plan, along with poor communication and the geographical dispersion of victims, led to significant delays. Some ambulances spent precious time searching for available beds, while others transported patients to nearby hospitals without verifying capacity or mistakenly chose distant facilities (Larson et al., 2006). As highlighted by Larson et al. (2006), cities must develop clear procedures for assigning patients to hospitals based on their conditions, locations, and real-time hospital capacity.

Compounding this challenge, hospitals are not always spared in disaster scenarios — they can themselves become victims. Emergency plans often assume hospitals remain operational, but this is not always the case. In fact, many healthcare facilities are located in hazard-prone regions. For example, it is estimated that between 50% and 70% of healthcare

facilities in Latin America and the Caribbean are situated in high-risk zones, with similar figures reported in the United Kingdom (Bagaria et al., 2009). When hospitals are compromised, casualty response becomes far more complex, and the need for robust evacuation plans becomes critical. Despite the importance of hospital evacuation planning, there is still no universally accepted approach, and existing guidelines vary significantly. For example, the Center for Bioterrorism Preparedness and Planning recommends evacuating the most critical patients first, whereas the Association for periOperative Registered Nurses and the New York Centers for Terrorism Preparedness and Planning suggest starting with the least critical (Childers and Taaffe, 2010). Real-world cases reveal the consequences of inadequate planning. In 1997, after a chemical explosion forced Helena Regional Medical Center to evacuate, high-priority patients were successfully transferred to other hospitals, but with no clear plan for low-priority patients, they were relocated to a local community college, where cafeteria tables were used as makeshift hospital beds (Squillace, 2010). These challenges underscore the urgent need for integrated planning that addresses both ACF deployment and hospital evacuation in the disaster response phase.

Research on CRP that involves setting up ACFs generally falls into three methodological categories: deterministic, stochastic, and robust optimization. Deterministic models are widely used but ignore uncertainty altogether, which can lead to unrealistic or overly costly solutions in practice. In contrast, stochastic models incorporate different disaster scenarios, both before and after an event, to better capture shifting demands and uncertain conditions. The increasing availability of real-time data sources, such as Synthetic Aperture Radar, drones, and mobile phone or GPS signals, has significantly improved situational awareness. These sources complement traditional inputs like official reports and local knowledge from community leaders (Grass et al., 2023). Although robust optimization is valuable when data are highly uncertain or scarce, recent technological advances now enable more reliable and timely information gathering. Building on this, our work employs a scenario-based, two-stage stochastic modeling approach that draws on satellite imagery, government data, and local informant input to represent uncertainty. For further discussion on these modeling frameworks, refer to Farahani et al. (2020) and Dönmez et al. (2021).

Deterministic models have been widely applied in CRP and ACF planning, assuming fixed inputs and known outcomes (Caunhye et al., 2015; Gu et al., 2018; Liu et al., 2019; Mousavi et al., 2024). Caunhye et al. (2015) formulated a location–allocation model aimed at deploying ACFs and managing the evacuation of different categories of victims following severe radiological disasters. Gu et al. (2018) presented a mixed-integer programming

(MIP) approach to determine ACF locations under budget constraints, focusing on maximizing access to healthcare services. Their solution method involved a hybrid metaheuristic algorithm. Liu et al. (2019) proposed a bi-objective mixed-integer linear programming (MIP) framework to simultaneously identify optimal ACF locations and casualty allocation, aiming to increase survival rates while reducing overall costs. The model was solved using the ϵ -constraint method. Mousavi et al. (2024) proposed a MIP model to identify optimal locations for ACFs and developed a hierarchical assignment structure linking injured individuals to ACFs, clinics, and hospitals. They employed a simulation-based optimization to evaluate system performance, showing reductions in congestion, cost, and treatment time.

Stochastic models address uncertainty by incorporating disaster scenarios that reflect varying conditions (Caunhye and Nie, 2018; Ghasemi et al., 2019; Oksuz and Satoglu, 2020; Chang et al., 2023; Oksuz and Satoglu, 2024). Caunhye and Nie (2018) applied a stochastic programming approach to address uncertainties in travel times and casualty demand, incorporating both triage and self-evacuation into the model. They solved the problem using a Benders decomposition algorithm. Ghasemi et al. (2019) developed a multi-objective, multi-commodity, and multi-vehicle MIP model for CRP, aiming to minimize facility location costs and reduce shortages of relief supplies. To solve it, they used a modified multi-objective particle swarm optimization (MMOPSO) algorithm, the Non-Dominated Sorting Genetic Algorithm II (NSGA-II), and the ϵ -constraint method. Oksuz and Satoglu (2020) proposed a two-stage stochastic model to determine the optimal number and locations of ACFs, accounting for existing hospitals, casualty types, and potential road and hospital damages. They validated their model using a real-world case study from Turkey and solved it with CPLEX. Chang et al. (2023) introduced a hybrid simulation-optimization method to locate ACFs and allocate emergency medical services, aiming to minimize the total delivery time for transferring all casualties to hospitals. Recently, Oksuz and Satoglu (2024) presented a multi-objective stochastic programming model that simultaneously optimizes ACF locations, resource allocations, and medical personnel assignments, solving it with an enhanced ϵ -constraint method.

In the context of robust optimization, Li et al. (2020) addressed a joint problem of ACF location and capacity planning, integrating the movement of evacuees and relief supplies across both initial and secondary disaster scenarios. They solved the model using a customized progressive hedging algorithm (PHA). Sun et al. (2022) focused specifically on medical services, studying the evacuation of injured individuals and proposing a three-stage deprivation cost function for severely injured patients. They solved the model using

the ϵ -constraint approach. Yin et al. (2024) developed a distributionally robust optimization model for multi-period humanitarian relief network design, aiming to coordinate ACF locations, supply inventories, and casualty evacuation under uncertainty. To solve it efficiently, they introduced an enhanced branch-and-Benders-cut algorithm. Most recently, Zhang et al. (2025) proposed a new robust optimization model for a four-echelon healthcare logistics network that optimizes ACF locations and resource allocation, while accounting for deprivation costs and casualty triage. Their model was applied to a real-world case study of the 2022 Luding Earthquake.

Beyond the integration of ACF location and hospital evacuation planning, several key operational aspects remain underexplored in the literature. One critical factor is the allocation of rescue vehicles, which play an essential role in evacuating patients from affected areas and transporting emergency personnel and supplies. Effective rescue vehicle assignment ensures that medical teams can reach casualties quickly and that patients are transferred safely and efficiently, significantly improving the overall response effectiveness. In practice, different types of land-based rescue vehicles are available, such as basic life support (BLS) ambulances, advanced life support (ALS) ambulances, and ambulance buses (AMBUS). However, not all vehicle types are suitable for transporting every kind of patient; for instance, severely injured individuals may require more specialized transport. In addition to land vehicles, aerial evacuation options such as helicopters are often deployed to expedite patient transfers, especially over long distances or when road networks are compromised. Importantly, each mode of transportation carries a certain level of risk for the patients being transported. These transportation risks must be carefully weighed against the threat risk posed by delayed evacuation. Although the CRP literature is extensive, to the best of our knowledge, no study has simultaneously considered the assignment of both land and aerial rescue vehicles, alongside ACF establishment, hospital evacuation, and casualty allocation. Addressing this gap, the main contributions of this paper are as follows:

- We propose the first stochastic CRP model that addresses the integrated ACF establishment and hospital evacuation planning. Our framework also incorporates the assignment of appropriate numbers of land and aerial rescue vehicles to each ACF.
- We develop a deep-reinforcement-learning-based metaheuristic that integrates reinforcement learning with an adaptive large neighborhood search (ALNS) algorithm. The algorithm is further enhanced through the use of randomized quasi-Monte Carlo (RQMC) sampling and a novel decision-based scenario clustering approach.
- We conduct extensive computational experiments to evaluate the proposed CRP model,

with a focus on multi-modal vehicle allocation and the associated transportation and evacuation risks.

- To evaluate the performance of our approach, we compare it against the PHA, a widely recognized decomposition method in the disaster response optimization literature, ensuring a valid comparison in terms of solution quality and computational efficiency.
- Finally, we validate the effectiveness of our best-performing algorithm using a real-world earthquake that occurred in Turkey.

The remainder of this paper is organized as follows. In Section 4.2, we describe the problem context. In Section 4.3, we formulate the corresponding two-stage stochastic programming model. In Section 4.4, we present the proposed learning-based matheuristic and the scenario clustering methods designed to enhance its performance on out-of-sample scenarios. In Section 4.5, we provide comprehensive computational results on the CRP considering hospital evacuation. In Section 4.6, we illustrate a real-world application through a detailed case study. Finally, in Section 4.7, we conclude the paper with a summary of key insights and directions for future work.

4.2 Problem Description

Disaster management encompasses four key phases (Tippong et al., 2022). The first phase, *mitigation*, involves anticipating potential hazards and developing action plans to reduce the impact of forthcoming disaster events. The second phase, *preparedness*, focuses on activities aimed at preparing for a disaster, such as training personnel. Both mitigation and preparedness are proactive phases, occurring before any disaster strikes. Once a disaster occurs, the *response* phase begins, which involves immediate actions taken to minimize the losses caused by the catastrophic event. The final phase, *recovery*, focuses on efforts to restore normality in the affected society.

This paper focuses on the response phase. Specifically, it addresses the critical decisions made in the aftermath of a Catastrophic health event (CHE). These decisions include transferring casualties from disaster-affected areas, evacuating patients from damaged hospitals, and assigning rescue vehicles to medical facilities (MFs), which include both ACFs and hospitals, during the early hours after the disaster. To clarify, in this paper, the term *patient* refers to an injured individual who, prior to the catastrophic event, was already

assigned to a specific hospital and, as a result of the event, must be evacuated to another medical facility because the original hospital can no longer provide the necessary care. In contrast, the term *casualty* refers to individuals who require evacuation from the affected region due to injuries sustained during the catastrophic event itself.

We model this response phase as a two-stage stochastic MIP problem, with five key decision-making areas following a disaster: (i) locating ACFs, (ii) assigning appropriate land and aerial rescue vehicles to ACFs, (iii) identifying backup hospitals, (iv) transporting casualties from disaster sites to MFs based on available resources and infrastructure, and (v) evacuating patients from damaged hospitals to functional MFs. The first three decisions are determined in advance of all humanitarian operations, as they establish general guidelines for how the affected population can be treated immediately in the aftermath of the catastrophic event. In particular, they determine how the available healthcare resources can be organized to support needs that are, by nature, uncertain. These advance decisions are based on initial information, such as the disaster's severity and geographical distribution. The remaining two decisions are made after obtaining additional information about uncertain parameters, which include: (1) the number and distribution of casualties across disaster areas, (2) the identification of hospitals that have been compromised or destroyed, (3) the total number of patients in damaged hospitals requiring transfer, and (4) the percentage of patients discharged from functional hospitals during each period after the disaster.

We assume that emergency responders will perform triage and an initial evaluation of casualties at the disaster site. This evaluation follows a standard protocol, often referred to by the acronym A.B.C.D.E., which stands for Airway, Breathing, Circulation, Disability, and Exposure/Environment. The purpose of this initial assessment is to rapidly identify and address any immediate life-threatening concerns while also determining the necessary type and urgency of the medical interventions. For instance, if a victim is conscious, able to communicate, aware of their environment, and can move all limbs, it suggests that their airway is clear, they are receiving adequate oxygen, and no significant central nervous system damage is present. While we provide an overview of the triage and initial assessment process here, an in-depth discussion is outside the scope of this paper.

Regarding the categorization of injured individuals, we follow a widely accepted classification found in the literature (for example, Caunhye and Nie (2018)), which divides casualties and patients into two groups based on the urgency of their medical needs: high-priority and low-priority. High-priority individuals, whether casualties or patients, include those experiencing severe bleeding, respiratory difficulties, or significant mental impairment, and must be transferred directly to hospitals for immediate, specialized care. In

contrast, low-priority individuals, both casualties and patients with less critical conditions, such as those requiring general surgeries, can receive treatment at either hospitals or ACFs.

An essential aspect of casualty transportation is addressing individual needs and minimizing transfer times. To optimize this process, we incorporate a multi-modal transportation system that combines both land and aerial modes. Land transportation includes three types of vehicles: Basic Life Support (BLS) ambulances, Advanced Life Support (ALS) ambulances, and ambulance buses (AMBUS). The primary difference between BLS and ALS ambulances is their capacity, specifically the number of individuals they can carry at a time. High-priority individuals can be transferred using any of these vehicles, including BLS, ALS, or aerial transport vehicles. However, AMBUSes are reserved exclusively for low-priority individuals. In fact, low-priority individuals can be transported using any available land- or aerial-rescue vehicles. Aerial transport is provided by helicopters, which can only land at ACFs that have designated staging areas, provided these ACFs are already established. One distinction between aerial and road modes of transportation is that helicopters are limited to facilities with landing capabilities.

Transferring patients from damaged hospitals or delaying their evacuation always involves some level of risk. These risks primarily stem from two factors: *threat risk*, which refers to the reason for evacuation, and *transportation risk* (Bish et al., 2014). The threat risk depends on the severity of the disaster and its impact on different patient groups, as well as how the threat evolves over time. For instance, an earthquake could immediately harm patients through injuries from falling debris or structural damage to the facility. Additionally, the long-term effects of an earthquake, such as damage to medical equipment and infrastructure, could complicate evacuation efforts and patient care. Similarly, a flood might pose little initial risk, but as water levels rise, the danger to the facility and patients increases. High-priority patients or those reliant on life-support equipment are especially vulnerable compared to patients in stable condition. As such, the threat risk varies depending on the type of disaster and its progression. On the other hand, transportation risk is influenced by the patient's condition, the characteristics of the chosen vehicle, and the time required to transport the patient to a receiving hospital.

In this CRP problem, our objective is to minimize the collaboration costs between hospitals that have agreed to serve as backup facilities in the first stage, as well as the overall evacuation risk for patients being moved, the threat risk for those left behind, the penalty associated with casualties remaining in disaster areas, and the transportation costs in the second stage.

In a nutshell, the proposed approach is initiated immediately after a disaster strikes and

is designed to support decision-making throughout the critical response phase. The first set of decisions focuses on the location of ACFs, selected from a set of candidate locations and subject to a limited establishment budget. Some of these potential locations are equipped with pre-existing staging areas, making them suitable for aerial rescue vehicle operations. Simultaneously, a suitable number of land-based rescue vehicles are allocated to each operational ACF to facilitate the transportation of casualties from disaster zones. In parallel, the model determines compatibility-based backup hospital assignments. A hospital may serve as a backup for another only if their medical service profiles align. For example, if a hospital specializes in burn care or pediatric treatment, a backup must offer equivalent services to be considered compatible. These initial decisions, including ACF location, rescue vehicle allocation, and backup hospital pairing, are fixed for the entire planning horizon.

As time progresses, additional information becomes available, allowing second-stage decisions to be made. Casualties are then transported from affected areas to available medical facilities, with their destinations determined by the severity of their injuries and the availability of resources. High-priority casualties are transferred exclusively to hospitals for specialized care, while low-priority casualties may be directed to either hospitals or ACFs. Due to limited treatment capacity at medical facilities and restricted transportation resources, some casualty transfers may need to be postponed to future periods, incurring delay penalties in the model. Simultaneously, patients located in disrupted hospitals must also be evacuated. Low-priority patients can be moved to ACFs using land-based rescue vehicles, while both high- and low-priority patients may be transferred to designated backup hospitals. This can be done via land or aerial transport, though aerial transfers are only possible when routed through an ACF with a staging area. Capacity constraints, whether in transportation or in backup hospital availability, can cause patient evacuations to be delayed, which in turn leads to elevated risk and increased penalty costs in the model. These penalties capture both the threat risk of patients who remain in compromised facilities and the transportation risk associated with delayed or prolonged transfers.

4.3 Two-Stage Stochastic Integer Programming Model

We develop a two-stage stochastic MIP model to address the CRP described in Section 4.2. In this framework, the first-stage decisions are made immediately after a disaster occurs, before the extent and effects of the catastrophic event are revealed. Uncertainty is then revealed, including updated data on casualty numbers, hospital status, patient discharges, and evacuation needs. Based on this realized information, second-stage decisions are made.

In the following, we present the formulations for the first and second stages separately.

4.3.1 First-Stage Model

We use the following notation to formulate the first-stage model:

Sets:

- $\mathcal{J}$: Set of potential locations i for establishing ACFs.
- $\mathcal{J}^A$: Subset of $\mathcal{J}$ representing ACF locations equipped with a staging area for aerial evacuation.
- $\mathcal{H}$: Set of existing hospitals h .
- $\mathcal{U}$: Set of all medical facilities u , where $\mathcal{U} = \mathcal{H} \cup \mathcal{J}$.
- $\mathcal{M}$: Set of land-based rescue vehicle types m .
- $\mathcal{K}_h$: Subset of $\mathcal{H}$ representing hospitals eligible to serve as backups for hospital $h \in \mathcal{H}$.

Parameters:

- B : Total available budget for establishing ACFs.
- f_i : Fixed cost of opening an ACF at location i .
- n_m^{ACF} : Number of available land-based rescue vehicles of type m that can be assigned to ACFs.
- $\varrho_{hh'}$: Coordination cost incurred when hospital h' serves as a backup for hospital h .

First-Stage Variables:

- x_i : 1 if an ACF is established at location i ; 0 otherwise.
- ϑ_{im} : Number of land-based rescue vehicles of type m assigned to ACF i .
- $w_{hh'}$: 1 if hospital h' is selected as a backup for hospital h ; 0 otherwise.

For clarity, bold notation is used to represent the decision variables in the following problem formulation. We now formulate the first-stage problem as follows:

$$\min_{\mathbf{x}, \boldsymbol{\vartheta}, \mathbf{w}} \left(\sum_{h \in \mathcal{H}} \sum_{\substack{h' \in \mathcal{H} \\ h' \neq h}} \varrho_{hh'} w_{hh'} + Q(\mathbf{x}, \boldsymbol{\vartheta}, \mathbf{w}) \right) \quad (101)$$

Subject to:

$$\sum_{i \in \mathcal{I}} f_i x_i \leq B \quad (102)$$

$$\sum_{i \in \mathcal{I}} \vartheta_{im} \leq n_m^{ACF} \quad m \in \mathcal{M} \quad (103)$$

$$\vartheta_{im} \leq n_m^{ACF} x_i \quad i \in \mathcal{I}, m \in \mathcal{M} \quad (104)$$

$$x_i, w_{hh'} \in \{0, 1\} \quad i \in \mathcal{I}, h, h' \in \mathcal{H} \quad (105)$$

$$\vartheta_{im} \in \mathbb{Z}^+ \quad i \in \mathcal{I}, m \in \mathcal{M} \quad (106)$$

Objective function (101) minimizes the total coordination cost between hospitals when one hospital is designated as a backup for another, in addition to the expected second-stage cost $Q(\mathbf{x}, \boldsymbol{\vartheta}, \mathbf{w})$. The term $Q(\mathbf{x}, \boldsymbol{\vartheta}, \mathbf{w})$ represents the expected value of the second-stage cost, which is evaluated through the second-stage problem presented in Section 4.3.2. It is defined as $Q(\mathbf{x}, \boldsymbol{\vartheta}, \mathbf{w}) = \mathbb{E}_{s \in \mathcal{S}}[Q(\mathbf{x}, \boldsymbol{\vartheta}, \mathbf{w}, \xi(s))]$, where $\mathbb{E}_{s \in \mathcal{S}}[\cdot]$ denotes the expected value over all scenarios $s \in \mathcal{S}$, and $\xi(s)$ is the vector of uncertain parameters associated with scenario s . Constraints (102) and (103) enforce the budget limitation for establishing ACFs and the availability limit on land-based rescue vehicles, respectively. Constraint (104) ensures that vehicles are only assigned to ACFs that are actually opened. Finally, constraints (105) and (106) specify the domains of all first-stage decision variables.

4.3.2 Second-Stage Model

To formulate the second-stage model, we use the following notation:

Sets:

- $\mathcal{T}$: Set of time periods t in the planning horizon.
- $\mathcal{L}$: Set of disaster locations l .
- $\mathcal{J}$: Set of injury severity levels, $\mathcal{J} = \{1, 2\}$, where 1 denotes high-priority and 2 denotes low-priority.
- $\mathcal{J}_u$: Set of injury levels that can be treated at facility $u \in \mathcal{U}$; $\mathcal{J}_h = \{1, 2\}$ for hospitals $h \in \mathcal{H}$, and $\mathcal{J}_i = \{2\}$ for ACFs $i \in \mathcal{I}$.
- $\mathcal{J}'_m$: Set of injury levels that can be transported using land-based rescue vehicle type $m \in \mathcal{M}$.
- $\mathcal{S}$: Set of scenarios s .

Parameters:

- c_i^{ACF} : Treatment capacity of ACF at location i .
- c_h^H : Treatment capacity of hospital h .
- π_{ij}^L : Estimated land travel time between locations i and j , where $i, j \in \mathcal{J} \cup \mathcal{H} \cup \mathcal{L}$.
- π_{ih}^A : Estimated aerial travel time between ACF $i \in \mathcal{J}^A$ and hospital h .
- ρ_j : Penalty for leaving casualties with injury level j in disaster areas.
- ν_j : Risk cost associated with evacuating a patient with injury level j .
- $n_{mh}^{LHos.}$: Number of available land-based rescue vehicles of type m at hospital h .
- ψ_h : Number of aerial rescue vehicles available at hospital h , used to evacuate patients from ACFs with staging areas.
- c_m^{rescue} : Capacity of land-based rescue vehicle type m , measured in terms of transportation time.
- c^{AR} : Capacity of an aerial rescue vehicle, measured in terms of transportation time.
- γ_m^L : Loading/unloading time for land-based rescue vehicle type m .
- γ^A : Loading/unloading time for an aerial rescue vehicle.

The following parameters are scenario-dependent, revealing specific information related to each scenario s , including the number and distribution of casualties across disaster areas, the identification of hospitals that have been compromised or destroyed, the total number of patients in damaged hospitals requiring transfer, and the percentage of patients discharged from functional hospitals during each period after the disaster.

- d_{tjls} : Number of casualties with injury level j occurring at location l in time period t under scenario s .
- g_{hs} : Equals 1 if hospital h is disrupted under scenario s ; 0 otherwise.
- Δ_{jhs} : Number of patients with injury level j at hospital h requiring evacuation under scenario s .
- λ_{tjus} : Percentage of patients with injury level j discharged after t periods from medical facility $u \in \mathcal{U}$ under scenario s .

In addition to the above parameters, we define the *risk* parameter as the probability that a generic adverse event occurs — such as a life-threatening change in patient condition or death. The total evacuation risk consists of two components: (1) *threat risk*, representing

the risk to patients who remain in disrupted hospitals while awaiting evacuation, and (2) *transportation risk*, representing the risk incurred during the transfer process (Bish et al., 2014). The threat risk parameter, α_{tjh} , denotes the probability that an adverse event occurs for a patient with injury level j who remains at disrupted hospital h during period t . Based on this, the *cumulative threat risk*, Λ_{tjh} , defined in Equation (107), gives the probability that a patient with injury level j at hospital h will experience the adverse event by the end of period t , if not yet evacuated. This cumulative risk increases over time as patients continue to be exposed to hazardous conditions. The underlying idea is that the longer a patient remains in a disrupted hospital, the greater the risk becomes — growing as a function of time periods.

$$\Lambda_{tjh} = 1 - \prod_{f=1}^t (1 - \alpha_{fjh}) \quad (107)$$

Regarding the transportation risk, we distinguish between two modes of transport, land-based and aerial, and define the associated risk parameters as β_{jm}^L and β_j^A , respectively. The parameter β_{jm}^L represents the probability that an adverse event occurs during land transportation using a rescue vehicle of type m for a patient with injury level j . Similarly, β_j^A denotes the probability of an adverse event occurring during aerial transportation for a patient with injury level j . The *cumulative land transportation risk*, denoted by Θ_{jhum}^L and defined in Equation (108), captures the risk faced by a patient with injury level j when evacuated from a disrupted hospital h to a backup medical facility u using a land-based rescue vehicle of type m . It is important to note that the total time exposure to risk includes both travel and handling durations. Specifically, transporting a patient by land between two locations i and j incurs a time of $(\pi_{ij}^L + 2\gamma_m^L)$, accounting for travel and loading/unloading times.

$$\Theta_{jhum}^L = 1 - (1 - \beta_{jm}^L)^{(\pi_{hu}^L + 2\gamma_m^L)} \quad (108)$$

Likewise, the *cumulative aerial transportation risk*, $\Theta_{jih'm}^A$, defined in Equation (109), refers to the total transportation risk for a patient with injury level j evacuated from a damaged hospital h to a functional backup hospital h' , via an ACF $i \in \mathcal{J}^A$, using both a land rescue vehicle of type m and an aerial rescue vehicle. For aerial transportation, the total time is given by $(\pi_{ij}^A + 2\gamma^A)$.

$$\Theta_{jih'm}^A = 1 - \left[(1 - \beta_{jm}^L)^{(\pi_{hi}^L + 2\gamma_m^L)} \cdot (1 - \beta_j^A)^{(\pi_{ih'}^A + 2\gamma^A)} \right] \quad (109)$$

The overall evacuation risk, whether for land-based or aerial transportation, is denoted by R_{tjhum}^L and $R_{tjih'm}^A$, respectively. These represent the risk associated with evacuating

a patient with injury level j from a disrupted hospital h in period t : either to a medical facility u using a land-based rescue vehicle of type m , or to a designated backup hospital h' via an ACF $i \in \mathcal{J}^A$ using both land and aerial transport. These evacuation risk values are calculated in Equations (110) and (111), respectively.

$$R_{tjhum}^L = 1 - [(1 - \Lambda_{tjh})(1 - \Theta_{jhum}^L)] \quad (110)$$

$$R_{tjih'h'm}^A = 1 - [(1 - \Lambda_{tjh})(1 - \Theta_{jih'h'm}^A)] \quad (111)$$

Each evacuation risk metric combines two components: the *cumulative threat risk* Λ_{tjh} , which captures the probability of an adverse event occurring before evacuation (i.e., while the patient remains in the disrupted hospital during period t), and the *transportation risk* — either Θ_{jhum}^L for land-based evacuation or $\Theta_{jih'h'm}^A$ for aerial evacuation — which accounts for the risk during the transfer based on injury level, vehicle type, and travel distance. We define the second-stage variables as follows:

- q_{tjlums} : Number of casualties with injury level j transferred from disaster area l to functional medical facility u using a land-based rescue vehicle of type m in period t under scenario s .
- μ_{tjls} : Number of unmet casualty demands with injury level j remaining at disaster area l in period t under scenario s .
- ς_{tjus} : Number of patients with injury level j discharged from medical facility u in period t under scenario s .
- u_{tjhum}^L : Number of patients with injury level j evacuated from non-functional hospital h to medical facility u using a land-based rescue vehicle of type m in period t under scenario s .
- $u_{tjih'h'ms}^A$: Number of patients with injury level j evacuated from non-functional hospital h to backup hospital h' via ACF $i \in \mathcal{J}^A$: first transported to ACF i using a land-based rescue vehicle of type m , then airlifted to hospital h' using an aerial rescue vehicle, all in period t under scenario s .
- ϖ_{tjhs} : Number of unevacuated patients with injury level j remaining at non-functional hospital h in period t under scenario s .
- ζ_{tus} : Available treatment capacity of medical facility u in period t under scenario s .

Based on the given notation, we formulate the second-stage model for each scenario

$s \in \mathcal{S}$ as follows:

$$Q(\mathbf{x}, \boldsymbol{\vartheta}, \mathbf{w}, \xi(s)) = \min \sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} \left(\sum_{l \in \mathcal{L}} \sum_{u \in \mathcal{U}} \sum_{m \in \mathcal{M}} \pi_{lu}^L q_{tjlm} + \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} \left(\sum_{u \in \mathcal{U}} \nu_j R_{tjhum}^L u_{tjhum}^L \right. \right. \\ \left. \left. + \sum_{i \in \mathcal{J}^A} \sum_{h' \in \mathcal{BC}_h} \nu_j R_{tjhih'm}^A u_{tjhih'm}^A \right) + \sum_{l \in \mathcal{L}} \rho_j \mu_{tjl} + \sum_{\substack{h \in \mathcal{H}: \\ g_{sh}=1}} \nu_j \Lambda_{|T|jh} \varpi_{tjhs} \right) \quad (112)$$

Subject to:

$$\sum_{\substack{m \in \mathcal{M}: \\ j \in \mathcal{J}_m \\ g_{hs} \neq 1}} q_{tjlm} + \sum_{\substack{m \in \mathcal{M}: \\ j \in \mathcal{J}_m \\ i \in \mathcal{J}}} q_{tjlim} + \mu_{tjl} = \mu_{(t-1)jl} + d_{tjl} \quad t \in \mathcal{T}, j \in \mathcal{J}, l \in \mathcal{L} \quad (113)$$

$$\sum_{\substack{h' \in \mathcal{K}_h: \\ g_{h's}=0 \& \\ j \in \mathcal{J}_m}} \sum_{m \in \mathcal{M}} u_{tjhh'm}^L + \sum_{\substack{i \in \mathcal{J}^A \\ h' \in \mathcal{K}_h: \\ g_{h's}=0 \& \\ j \in \mathcal{J}_m}} \sum_{m \in \mathcal{M}} u_{tjhih'm}^A \\ + \sum_{\substack{m \in \mathcal{M}: \\ j \in \mathcal{J}_m \\ i \in \mathcal{J}}} u_{tjhims}^L + \varpi_{tjhs} = \varpi_{(t-1)jhs} \cdot \mathbf{1}_{(t>1)} + \Delta_{jhs} \cdot \mathbf{1}_{(t=1)} \quad t \in \mathcal{T}, j \in \mathcal{J}, h \in \mathcal{H}, g_{hs} = 1 \quad (114)$$

$$\sum_{t'=1}^t \sum_{j \in \mathcal{J}_h} \left(\sum_{\substack{m \in \mathcal{M}: \\ j \in \mathcal{J}_m}} \left(\sum_{h' \in \mathcal{H}} \sum_{i \in \mathcal{J}^A} u_{t'jh'ihms}^A + \sum_{l \in \mathcal{L}} q_{t'jlm} + \sum_{h' \in \mathcal{H}} u_{t'jh'hms}^L \right) \right. \\ \left. - \varsigma_{t'jhs} \right) \leq \zeta_{ths} \quad t \in \mathcal{T}, h \in \mathcal{H} \quad (115)$$

$$\zeta_{ths} \leq c_h^{Hos.} \cdot \mathbf{1}_{(g_{hs} \neq 1)} \quad t \in \mathcal{T}, h \in \mathcal{H} \quad (116)$$

$$\varsigma_{tjhs} = \sum_{k=1}^t \lambda_{kjh} \left(\sum_{\substack{m \in \mathcal{M}: \\ j \in \mathcal{J}_m}} \left(\sum_{l \in \mathcal{L}} q_{(t-k)jlm} + \sum_{h' \in \mathcal{H}} u_{(t-k)jh'hms}^L \right. \right. \\ \left. \left. + \sum_{i \in \mathcal{J}^A} \sum_{h' \in \mathcal{H}} u_{(t-k)jh'ihms}^A \right) \right) \quad t \in \mathcal{T}, h \in \mathcal{H}, j \in \mathcal{J}_h \quad (117)$$

$$\sum_{t'=1}^t \sum_{j \in \mathcal{J}_i} \left(\sum_{\substack{m \in \mathcal{M}: \\ j \in \mathcal{J}_m}} \left(\sum_{l \in \mathcal{L}} q_{t'jlim} + \sum_{h \in \mathcal{H}} u_{t'jhims}^L \right) - \varsigma_{t'jis} \right) \leq \zeta_{tis} \quad t \in \mathcal{T}, i \in \mathcal{J} \quad (118)$$

$$\zeta_{tis} \leq \zeta_i^{ACF} x_i \quad t \in \mathcal{T}, i \in \mathcal{J} \quad (119)$$

$$\varsigma_{tjis} = \sum_{k=1}^t \lambda_{kjis} \left(\sum_{m \in \mathcal{M}} \left(\sum_{l \in \mathcal{L}} q_{(t-k)jlim} + \sum_{h \in \mathcal{H}} u_{(t-k)jhims}^L \right) \right) \quad t \in \mathcal{T}, i \in \mathcal{J}, j \in \mathcal{J}_i \quad (120)$$

$$\sum_{j \in \mathcal{J}_h} \left[\sum_{l \in \mathcal{L}} \pi_{lh}^L q_{tjlm} \cdot \mathbf{1}_{(g_{hs} \neq 1)} + \left(\sum_{\substack{h' \in \mathcal{K}_h: \\ g_{h's} \neq 1}} \pi_{hh'}^L u_{tjhh'ms}^L \right. \right. \\ \left. \left. + \sum_{\substack{i \in \mathcal{J}^A \\ h' \in \mathcal{K}_h: \\ g_{h's} \neq 1}} \pi_{hi}^L u_{tjhih'ms}^A + \sum_{i \in \mathcal{J}} \pi_{hi}^L u_{tjhims}^L \right) \cdot \mathbf{1}_{(g_{hs}=1)} \right] \\ \leq c_m^{\text{rescue}} n_{mh}^{LHos.} \quad t \in \mathcal{T}, h \in \mathcal{H}, m \in \mathcal{M} \quad (121)$$

$$\sum_{j \in \mathcal{J}_i} \sum_{l \in \mathcal{L}} \pi_{li}^L q_{tjlims} \leq c_m^{rescue} \vartheta_{im} \quad t \in \mathcal{T}, i \in \mathcal{I}, m \in \mathcal{M} \quad (122)$$

$$\sum_{j \in \mathcal{J}_{h'}} \sum_{i \in \mathcal{J}^A} \sum_{h' \in \mathcal{K}_h: g_{h's} \neq 1} \sum_{m \in \mathcal{M}} \pi_{ih'}^A u_{tjhih'ms}^A \leq c^{AR} \psi_h \quad t \in \mathcal{T}, h \in \mathcal{H}, \text{ s.t. } g_{hs} = 1 \quad (123)$$

$$\sum_{j \in \mathcal{J}_{h'}} \sum_{m \in \mathcal{M}} \left(u_{tjhh'ms}^L + \sum_{i \in \mathcal{J}^A} u_{tjhih'ms}^A \cdot \mathbf{1}_{(h' \in \mathcal{K}_h)} \right) \leq c_{h'}^{Hos} w_{hh'} \quad t \in \mathcal{T}, h, h' \in \mathcal{H}, h \neq h' \quad (124)$$

$q_{tjlims}, \mu_{tjls}, \varsigma_{tjus}, u_{tjhumts}^L,$

$$u_{tjhih'ms}^A, \varpi_{tjhs}, \zeta_{tus} \geq 0 \quad t \in \mathcal{T}, h, h' \in \mathcal{H}, l \in \mathcal{L}, i \in \mathcal{I}, j \in \mathcal{J}, u \in \mathcal{U}, m \in \mathcal{M} \quad (125)$$

The objective function (112) captures the second-stage cost in scenario s , which includes: the transportation cost of casualties from disaster areas to medical facilities, the evacuation risk costs associated with both land-based and aerial transfers, penalties for unmet casualty demand due to limited resources, and the cumulative threat risk for patients who remain unevacuated by the end of the planning horizon. Constraint (113) defines the allocation of casualties to ACFs and functional hospitals based on their injury level. On the right-hand side of this constraint, unmet casualty demand from the previous period is carried forward and added to the current period's demand. Constraint (114) handles the allocation of patients from disrupted hospitals to either ACFs or designated backup hospitals, using land-based or aerial rescue vehicles. In this constraint, the number of patients requiring evacuation is determined in the first period; in subsequent periods, only unevacuated patients from earlier periods remain to be handled. Constraints (115) and (116) enforce the treatment capacity limits of functional hospitals, while constraints (118) and (119) do the same for ACFs. The number of patients discharged in each period from hospitals and ACFs is determined by constraints (117) and (120), respectively. Rescue vehicle capacities are addressed in constraints (121) and (122), where hospitals operate their own land-based vehicles, while ACFs are assigned a specified number of vehicles. Constraint (123) ensures that aerial rescue vehicle use remains within available limits. Constraint (124) guarantees that patient transfers between hospitals can only occur if the receiving hospital has been pre-designated as a backup facility. Finally, constraint (125) enforces the non-negativity of all second-stage decision variables.

4.4 Deep-Reinforcement-Learning-Based Matheuristic

To solve our complex stochastic optimization problem, we propose a hybrid approach integrating deep reinforcement learning (DRL) with the adaptive large neighborhood search

(ALNS) algorithm, referred to as *DRL-ALNS*. This approach enhances ALNS by improving the flexibility and guidance of the search process, addressing key limitations of classical ALNS. ALNS, a heuristic for efficiently solving NP-hard problems, was introduced by Ropke and Pisinger (2006) and Pisinger and Ropke (2007) as an extension of the large neighborhood search (LNS) proposed by Shaw (1998). Both LNS and ALNS use two operators — destroy and repair — that iteratively deconstruct and reconstruct solutions. ALNS’s strength lies in its ability to adaptively explore diverse neighborhoods by tracking operator performance and adjusting selection probabilities, balancing intensification and diversification in the search.

However, the effectiveness of ALNS heavily depends on the design of its adaptive operator selection (AOS) mechanism and the configuration of its acceptance parameters. Traditional acceptance and AOS techniques, such as the Metropolis acceptance criterion and the roulette wheel selection, are commonly employed to guide the search and select promising operators. These mechanisms form the core of our baseline ALNS implementation, which we use as a benchmark for evaluating the performance of our proposed learning-based approach. Despite its widespread use, standard ALNS has several limitations. As highlighted by Turkeš et al. (2021), simple weight-based AOS strategies often fail to exploit meaningful patterns in the search trajectory and can struggle to prioritize truly effective operators. Moreover, traditional operator selection is typically based solely on historical performance, without accounting for the current context or state of the search. This reactive, rather than proactive, strategy may lead to suboptimal decisions — especially in highly dynamic and stochastic environments (Reijnen et al., 2024). To overcome these limitations, we incorporate a deep reinforcement learning agent into the ALNS framework. DRL enables the algorithm to learn from the evolving state of the search process and proactively configure operator selection and acceptance strategies in each iteration. This integration enables more effective guidance in the heuristic search choices made throughout the solution process, as the agent dynamically adapts to the structure of the current solution rather than relying solely on the past performance of the operators used. As suggested in Li et al. (2024), the ability of DRL to capture complex patterns and respond quickly to environmental changes makes it a natural fit for enhancing metaheuristic frameworks like ALNS.

Before delving into the specifics of the algorithm, it is important to first clarify the overall search strategy of our approach. The ALNS search focuses on finding the optimal first-stage decisions. The decision space is defined as the concatenation of the decision vectors x , ϑ , and w . The ALNS mechanics, which iteratively destroy and reconstruct

solutions, explore this search space to find improvements in the first-stage decisions. Importantly, the stochastic nature of the problem — represented by the second-stage cost — is not directly addressed within the search space but is incorporated through the evaluation of first-stage solutions. This evaluation is done at each iteration and relies on the expected second-stage cost, which is handled through the integration of the Deep Q-Learning (DQL) agent. The DQL agent is trained to efficiently select the most effective operators, enhancing the solution quality at each step. This approach contrasts with alternative methods that could include second-stage decisions within the search space, which we deliberately chose not to pursue to maintain focus on improving the first-stage decisions.

A major algorithmic complexity in our proposed approach lies in the evaluation of candidate solutions, which requires computing expected second-stage costs over the full set of possible scenarios at each ALNS iteration. To address this computational challenge efficiently, we implement a two-step preprocessing strategy. First, we apply a randomized quasi-Monte Carlo scenario generation method to better cover the full breadth of probabilistic possibilities for stochastic parameter realizations. Second, we employ a decision-based scenario clustering approach to reduce the number of scenarios used in expectation computations by retaining only those with the greatest impact on model solutions. This combined strategy enables more efficient execution of the optimization phase required at each DRL-ALNS iteration — specifically, the evaluation of incumbent first-stage solutions through second-stage problem solving. The remainder of this section presents our complete solution methodology. We first describe the scenario preprocessing steps in Sections 4.4.1 and 4.4.2, then outline the structure of the baseline ALNS algorithm in Section 4.4.3, and finally explain how the DRL agent is integrated into the search process in Section 4.4.4.

4.4.1 Advanced Scenario Sampling

In many stochastic optimization problems, the complete set of scenarios representing uncertain parameters is too large to handle computationally. To address this, a representative subset is often selected, but improving its quality is crucial for accurate results. One promising approach is using more sophisticated sampling strategies, such as Quasi-Monte Carlo (QMC) methods, which employ low-discrepancy sequences to fill the sample space more evenly. This reduces estimator variance, leading to more stable evaluations of expectations. However, QMC methods lack randomness, limiting their ability to generate unbiased estimates or confidence intervals. To overcome this, Randomized Quasi-Monte Carlo (RQMC) methods (Cranley and Patterson, 1976) introduce controlled randomness,

maintaining uniform coverage while enabling unbiased estimates. This hybrid approach reduces variance and allows for confidence interval assessment, making it ideal for generating efficient scenario sets with fewer samples.

4.4.2 Decision-Based Scenario Clustering

This section introduces the proposed decision-based scenario clustering method, which is applied in two stages: (i) computing an opportunity cost distance between scenarios, and (ii) selecting representative scenarios. Before delving into these steps, we present the general formulation of the optimization model introduced in Section 4.3. Suppose a decision-maker is faced with an optimization problem under uncertainty — such as selecting locations for ACFs given uncertain casualty levels. The decision variables are represented by the vector $\mathbf{x}$, and the uncertainty is described by a random vector ξ . Let $\phi(\mathbf{x}, \xi)$ denote the objective function that evaluates the quality of decision $\mathbf{x}$ under a specific realization ξ . Assuming the decision-maker aims to minimize this objective, the problem can be expressed as:

$$\min_{\mathbf{x} \in A} \mathbb{E}_{\xi}[\phi(\mathbf{x}, \xi)] \quad (126)$$

where A is the feasible region and $\phi(\mathbf{x}, \xi)$ denotes the cost when $\mathbf{x}$ is implemented under realization ξ . In practice, a finite scenario set $\mathcal{S} = \{1, \dots, S\}$ with probabilities $\{p_s\}$ is drawn from a probability measure $\mathbb{P}$, and one solves the sample-average approximation:

$$\min_{\mathbf{x} \in A} \sum_{s \in \mathcal{S}} p_s \phi(\mathbf{x}, \xi_s) \quad (127)$$

While this scenario-based formulation is widely used in stochastic optimization (Birge and Louveaux, 2011), a key practical challenge is selecting which scenarios to retain. The number of potential realizations across all data sources is typically very large, making it computationally prohibitive to include all of them. To address this, we propose a clustering approach that preserves scenario diversity based on decision relevance.

Step 1: Opportunity Cost Distance The goal is to interpret the information contained in each scenario ξ_s through the lens of its associated optimal decision. Therefore, we solve the deterministic version of the problem for each scenario $s \in \mathcal{S}$, obtaining:

$$\mathbf{x}(s) = \arg \min_{\mathbf{x} \in A} \phi(\mathbf{x}, \xi_s) \quad s \in \mathcal{S} \quad (128)$$

Solution $\mathbf{x}(s)$ represents the optimal decision assuming scenario s will occur with certainty. To quantify how “similar” two scenarios are from a decision-making standpoint, we use the *opportunity cost distance* proposed by Hewitt et al. (2022). For any pair of scenarios $s_1, s_2 \in \mathcal{S}$, this distance measures the cost incurred when using the decision optimized for one scenario in the context of the other:

$$d(s_1, s_2) = \phi(\mathbf{x}(s_1), \xi_{s_2}) - \phi(\mathbf{x}(s_2), \xi_{s_2}) + \phi(\mathbf{x}(s_2), \xi_{s_1}) - \phi(\mathbf{x}(s_1), \xi_{s_1}) \quad (129)$$

A small value of $d(s_1, s_2)$ implies that the two scenarios yield similar decisions, in the sense that their corresponding solutions perform well under both realizations. Using Equation (129), we compute the full pairwise distance matrix across all scenarios in $\mathcal{S}$.

Step 2: Scenario Selection Once the opportunity cost distance matrix has been constructed, the next step is to identify a representative subset of scenarios that reflects the diversity of the entire scenario space. Traditionally, this task is framed as a clustering problem over the scenario set $\mathcal{S}$, aiming to group scenarios that are similar in terms of their decision impact. This grouping is performed using the distance matrix constructed in the previous step. We employ a greedy Max–Min diversity selection algorithm (see Algorithm 2) inspired by farthest-point sampling. The algorithm begins by identifying the pair of scenarios with the greatest opportunity cost distance (line 3 of Algorithm 2), and then iteratively adds the scenario that is farthest from its closest counterpart in the set of already selected scenarios (lines 5–9 of Algorithm 2). This ensures that each additional scenario brings in the most novel decision information relative to the current set, making the approach computationally efficient, scalable, and well-suited for scenario reduction in stochastic optimization.

4.4.3 Basic Adaptive Large Neighborhood Search (ALNS)

The ALNS algorithm iteratively applies a destroy and a repair operator to improve the solution. In each iteration, a destroy operator removes part of the current solution, followed by a repair operator that reconstructs the solution by adding or modifying elements. The algorithm then evaluates whether the new solution improves the objective function.

In Basic ALNS, the selection of destroy and repair operators is performed using the roulette wheel selection method (see line 5 of Algorithm 2 in Windras Mara et al. (2022)). The key distinction in our approach is that we replace this mechanism with a deep reinforcement learning (DRL) agent (presented in Section 4.4.4) to dynamically select the

Algorithm 2 Greedy Max–Min Scenario Selection

```
1: Input: Distance matrix ( $D \in \mathbb{R}^{S \times S}$ ), number of scenarios to select ( $k$ )
2: Output: Set  $S'$  of  $k$  diverse scenario indices
3: Initialize  $S' \leftarrow \{i, j\}$ , where  $(i, j) = \arg \max_{i,j} D[i, j]$      $\triangleright$  Select most distant initial pair
4: while  $|S'| < k$  do
5:   Let  $R \leftarrow S \setminus S'$      $\triangleright$  Remaining unselected scenarios
6:   for each  $r \in R$  do
7:     Compute  $d_{\min}(r) = \min_{s \in S'} D[r, s]$ 
8:   end for
9:   Select  $r^* = \arg \max_{r \in R} d_{\min}(r)$ 
10:  Add  $r^*$  to  $S'$ 
11: end while
12: return  $S'$ 
```

appropriate destroy-repair combinations. The roulette wheel assigns probabilities to each operator based on its historical performance, with more successful operators having higher selection probabilities. These probabilities are updated iteratively, allowing the algorithm to adaptively explore various neighborhoods and enhance solution quality. This selection mechanism balances intensification in promising regions and diversification to avoid local optima.

4.4.4 Deep Reinforcement Learning-Based ALNS (DRL-ALNS)

The two-stage stochastic MIP presented in Section 4.3 is solved through a sequential decision-making process, where a reinforcement learning (RL) agent iteratively guides the ALNS destroy and repair operators. In each iteration, the agent treats the current first-stage solution as the observed state, selects a destroy/repair action to perturb this solution, and receives a reward based on the improvement in the objective after re-optimizing the MIP. A simulated annealing acceptance criterion then decides whether the new solution is accepted into the search. By embedding an RL agent within the ALNS framework, the method learns to adaptively select operators and escape local optima more effectively than a classical ALNS with fixed weight adaptation. We first formulate the DRL-ALNS as a Markov Decision Process (MDP), which is represented as a tuple $(\hat{S}, \mathcal{A}, R)$, wherein $\hat{S}$ represents the set of states, $\mathcal{A}$ represents the set of actions, and R represents the reward function. The components of our formulation are:

State ($\hat{s}$): A representation of the current first-stage solution (decision vector) for the two-stage MIP. This includes all binary and integer first-stage decision variables (i.e. x_i, ϑ_{mi} ,

and, $w_{hh'}$). At the start ($k = 0$), an initial state $\hat{s}_0$ is obtained by solving a linear relaxation version of the model to generate a high-quality starting solution for the first-stage decisions. It is important to note that, although the linear relaxation is solved, the solution is not necessarily feasible. Therefore, we round down the first-stage decision variables to the nearest integers and then check all relevant constraints to ensure that the solution satisfies all feasibility conditions. If any constraints are violated, adjustments are made to make the solution completely feasible. This initial optimization also provides an initial estimate of the value function.

Action (a): A composite destroy-and-repair operator applied to the current solution. We define a discrete action space of heuristic operators for first-stage decisions. For binary variables, actions include either removing some currently selected ACFs or taking no action (NoDestroy), followed by either adding new ACFs or swapping between opened and closed ones. For integer variables, actions include either decreasing some values or taking no action (NoChange), followed by either increasing others or again applying NoChange. Each action a thus specifies a pair of destroy/ repair heuristics for binary decisions and another pair for integer decisions. For example, an action might be $a = (\text{Remove}, \text{Add}, \text{Decrease}, \text{Increase})$, meaning we remove a subset of currently open ACFs, then add new ACFs, decrease some integer allocations, and then increase others. The magnitude of destruction is adaptive: at each iteration we randomly choose the number of elements to remove or change (e.g., between 10% and 50% of ACFs). This ensures a variable neighborhood size (“destruction severity”). The RL agent’s policy will learn which operator combination to apply based on the current state.

Reward (r): The reward reflects the improvement in the objective value. Let $f(\mathbf{x})$ denote the total cost of the current first-stage solution $\mathbf{x}$, calculated after solving all second-stage scenarios:

$$f(\mathbf{x}) = \min \left(\varrho_{hh'} w_{hh'} + \sum_{s \in \mathcal{S}} p_s \left(\sum_{t \in \mathcal{T}} \sum_{j \in \mathcal{J}} \left[\sum_{l \in \mathcal{L}} \sum_{u \in \mathcal{U}} \sum_{m \in \mathcal{M}} \pi_{lu} q_{tjlums} + \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} \left(\sum_{u \in \mathcal{U}} \nu_j R_{tjhum}^L u_{tjhum}^L \right. \right. \right. \right. \\ \left. \left. \left. + \sum_{i \in \mathcal{I}^A} \sum_{h' \in \text{BC}_h} \nu_j R_{tjihim}^A u_{tjihim}^A \right) + \sum_{l \in \mathcal{L}} \rho_j \mu_{tjls} + \sum_{\substack{h \in \mathcal{H}: \\ g_{sh}=1}} \nu_j \Lambda_{Tjh} \varpi_{tjhs} \right] \right) \right) \quad (130)$$

At each iteration, the reward is computed as the difference between the cost of the current solution $f(\mathbf{x})$ and the best cost found so far, i.e., $f(\mathbf{x}^*)$:

$$r = f(\mathbf{x}) - f(\mathbf{x}^*) \quad (131)$$

This reward encourages the agent to take actions that reduce the objective function. Since the goal of ALNS is cost minimization, using the cost difference as the reward aligns with the overall search objective. The RL agent learns a value function $Q(\hat{s}, a)$, estimating the long-term improvement in objective when action a is applied in state $\hat{s}$. The search proceeds iteratively. Each iteration k generates a new state $\hat{s}_{k+1}$ and reward r_k , and continues either for a fixed number of iterations or until a stopping criterion is met (e.g., no improvement over a given number of steps).

Deep Q-Learning for Operator Selection We employ a Deep Q-Learning (DQL) agent to learn the value of state-action pairs $(\hat{s}, a)$, i.e. the quality of applying each operator in a given solution state. Classic Q-learning uses a lookup table to store $Q(\hat{s}, a)$ values, but in our setting the state $\hat{s}$ is high-dimensional (hundreds of binary/integer variables) and the number of possible states is enormous. We therefore approximate the Q-value function with a deep neural network, which takes as input the state (encoded, for example, as a binary/continuous feature vector of the first-stage decisions) and returns Q-values for all possible actions. The network is trained online during the ALNS iterations via incremental Q-learning updates.

At each iteration k , given state s_k (i.e., the current first-stage solution), the agent selects an action a_k using an ϵ -greedy policy: with probability ϵ , it chooses a random operator (exploration), and with probability $1 - \epsilon$, it chooses the action with the highest estimated $Q(\hat{s}_k, a)$ (exploitation). In this context, each action a_k consists of a pair of destroy and repair operators. We initialize $\epsilon = 1.0$ and gradually decay it to a small final value (e.g., 0.1) over the first few thousand iterations. This allows the search to start with broad exploration and become increasingly greedy as the agent gains confidence in its learned Q-values.

To stabilize the learning process, we maintain a replay buffer that stores past experiences $(\hat{s}, a, r)$. These experiences are sampled in mini-batches from the buffer to update the Q-network. This helps break the correlation between consecutive experiences, which is crucial for improving the generalization capability of the model. By training the Q-network on a diverse set of experiences, the agent avoids overfitting to a limited set of states and actions, thereby improving learning efficiency in dynamic environments. A separate target network is used and updated periodically (every 500 steps) to improve the stability of Q-value estimates. The target network is a copy of the Q-network, but its parameters are updated less frequently, which helps to prevent rapid fluctuations in Q-value estimates during training. This slow update process ensures more stable learning by providing a consistent target for Q-value computations, avoiding the instability that could arise from using

the current Q-network as both the estimator and the target.

The DQL agent is trained to satisfy the Bellman optimality condition. After observing reward r_k and next state s_{k+1} , the network minimizes the temporal difference (TD) error:

$$Q(\hat{s}_k, a_k) \leftarrow Q(\hat{s}_k, a_k) + \alpha \left[r_k + \gamma \max_{a'} Q(\hat{s}_{k+1}, a') - Q(\hat{s}_k, a_k) \right] \quad (132)$$

where α is the learning rate, and γ is a discount factor (we use $\gamma = 0.9$ to mildly prioritize immediate rewards). The network parameters are updated via stochastic gradient descent on the mean-squared TD error, enabling the agent to learn which destroy/repair operators tend to produce better long-term improvements in the objective function.

In summary, the RL agent automates the operator selection scheme of ALNS by learning from experience rather than relying on hand-tuned weights. A traditional ALNS would adjust operator weights based only on recent performance (e.g., increasing weights for successful operators). In contrast, our DQL agent can condition its choices on the current state of the solution and predict which operator is most effective in that context. This addresses a known limitation of classical ALNS weight updating, which “cannot take advantage of short-term dependencies between the current state of the search and the selection of operators”. By using deep Q-learning, the algorithm can capture complex patterns (for example, recognizing that certain destroy-repair moves work well when specific structures are present in the solution) and thus adapt more intelligently as the search progresses. A summary of the algorithm is provided in Algorithm 3.

4.5 Computational Results

All computational experiments were conducted on a server equipped with two 2.40 GHz Intel Xeon Gold 6148 processors, using Python and Gurobi 11.0.1. Each run was executed on a single core. The proposed algorithm was tested on instances with planning horizons of 4 and 5 time periods, where each time period corresponds to 12 hours. This variation allows us to evaluate the performance of the solution methods under different post-disaster timeframes. The results are organized as follows: Section 4.5.1 introduces the test instances, while Sections 4.5.2 and 4.5.3 assess the performance of the proposed method in out-of-sample and in-sample evaluations, respectively.

Each solution approach was terminated either when a time limit of one hour was reached for matheuristic-based approaches, or after 100 consecutive iterations without improvement. For the exact MIP and progressive hedging algorithm (PHA), the time limit was set

Algorithm 3 DRL-ALNS Matheuristic Algorithm

- 1: **Input:** Two-stage stochastic MIP model; operator set $\mathcal{A}$; maximum number of iterations N .
 - 2: Solve the linear relaxation of the model (Section 4.3) to obtain initial first-stage solution $\mathbf{x}^0$.
 - 3: Set current solution $\mathbf{x} \leftarrow \mathbf{x}^0$, best solution $\mathbf{x}^* \leftarrow \mathbf{x}^0$, and best cost $f^* \leftarrow f(\mathbf{x}^0)$.
 - 4: Initialize the DQL agent with Q-network $Q_\theta(\hat{s}, a)$ (state dimension = $|\mathbf{x}|$, action dimension = $|\mathcal{A}|$).
 - 5: Initialize simulated annealing temperature $T \leftarrow T_0$.
 - 6: **for** $k = 1$ **to** N **do**
 - 7: Observe current state $\hat{s}$ based on features of $\mathbf{x}$.
 - 8: Select action $a \in A$ using ϵ -greedy policy based on $Q_\theta(\hat{s}, a)$.
 - 9: Apply the destroy operator in a to $\mathbf{x}$, then apply the repair operator to construct a new solution $\mathbf{x}'$.
 - 10: Fix $\mathbf{x}'$ as the first-stage decision, and solve the second-stage model (112)–(125) to compute $f(\mathbf{x}')$.
 - 11: Compute reward $r \leftarrow f(\mathbf{x}) - f(\mathbf{x}')$.
 - 12: Update the Q-network using the tuple $(\hat{s}, a, r)$.
 - 13: **if** $f(\mathbf{x}') < f^*$ **then**
 - 14: $\mathbf{x}^* \leftarrow \mathbf{x}'$ ▷ Update best solution
 - 15: $f^* \leftarrow f(\mathbf{x}')$
 - 16: **else**
 - 17: Compute acceptance probability $p \leftarrow \min\{1, \exp[-(f(\mathbf{x}') - f(\mathbf{x}))/T]\}$
 - 18: **if** $p > \text{Uniform}(0, 1)$ **then**
 - 19: $\mathbf{x} \leftarrow \mathbf{x}'$ ▷ Accept $\mathbf{x}'$ as current solution
 - 20: $f(\mathbf{x}) \leftarrow f(\mathbf{x}')$
 - 21: **end if**
 - 22: **end if**
 - 23: $T \leftarrow \rho \cdot T$ ▷ Cool down temperature
 - 24: **if** stopping criterion is met **then**
 - 25: **break**
 - 26: **end if**
 - 27: **end for**
 - 28: **Output:** Best solution $\mathbf{x}^*$ with objective value f^* .
-

to 10 hours. For training, we use 350 in-sample scenarios (or 50, if scenario reduction is applied). To evaluate out-of-sample performance, we simulate 500 independent scenarios, representing potential realizations of uncertain parameters over the planning horizon. These scenario paths are used consistently across all methods and instances to ensure comparability.

4.5.1 Implementation Details

This section outlines the procedure for generating CRP instances, incorporating uncertainties in post-disaster conditions such as casualty distribution across affected areas, hospital functionality, evacuation needs, and patient discharge rates over time. Most of the input data are adapted from prior studies including Sun et al. (2022), Oksuz and Satoglu (2020), and Bish et al. (2014). However, since our formulation introduces several novel components not previously addressed in the literature, additional parameters were generated specifically for this study. These new elements are described below.

We consider multiple problem scales by varying the number of potential ACFs ($i \in \{15, 20, 25, 30\}$), where 50% of ACFs are equipped with a staging area. The number of demand locations is also varied ($l \in \{15, 20, 25, 30\}$), and five hospitals are included in each instance, with capacities ranging from 400 to 600 beds. Three types of land-based rescue vehicles are modeled, with respective capacities of 1, 2, and 20 patients per trip (Rambha et al., 2021; Texas Emergency Medical Task Force, 2025). In addition, we consider a single aerial rescue vehicle type with a capacity of 10 patients per trip (Sun et al., 2021).

Hospitals, ACFs, and demand locations are uniformly distributed within a $50 \text{ km} \times 50 \text{ km}$ area. Casualty numbers at each disaster location are generated from a uniform distribution $U[0, 100]$ for each time period. Of these casualties, 40% are classified as high-priority and must be transported directly to hospitals, while the remaining 60% are low-priority and may be sent to either hospitals or ACFs (Sun et al., 2022).

Initial hospital occupancy rates are sampled from a uniform distribution $U[0.5, 0.6]$, reflecting moderate to high utilization at the time of the disaster (Statista Research Department, 2024). Patient discharge is modeled over a 20-period horizon. Based on Chan et al. (2019), we assume that 50% of high-priority patients are discharged after the average stay duration of 108.89 hours, while the remaining 50% are discharged earlier or later, symmetrically distributed across ± 6 periods. This forms a bell-shaped discharge pattern centered around the average stay duration. Each hospital has a 45% probability of being disrupted following a disaster (Liel and Deierlein, 2012). Following a hospital disruption, two types

of risk must be considered: *threat risk*, arising from delays in evacuation, and *transportation risk*, associated with transferring patients under hazardous conditions. The parameters associated with these risks are summarized in Table 10. As discussed in Section 4.2, the exponential form of the threat risk function is intended to capture the compounding danger caused by prolonged loss of critical utilities (e.g., electricity or water). We study the impact of different types of threat risk functions, including constant, linear, and exponential, in the case study in Section 4.6.

Table 10: Risk Parameter Settings for Threat and Transportation (taken from Bish et al. (2014))

Injury type	Threat risk param. (α_{tj})			Land-based trans. risk param. (β_{jm}^L)			Aerial trans. risk param. (β_j^A)
	Constant	Linear	Exponential	BLS	ALS	AMBUS	
High-Priority	0.0020	$0.00180t$	$6.25 \cdot 10^{-5} e^{(t/0.444)}$	0.0002	0.0001	0.0005	0.00020
Low-Priority	0.0012	$0.00108t$	$3.75 \cdot 10^{-5} e^{(t/0.472)}$	$5E-05$	$5E-05$	0.0001	0.00008

Further parameter settings are summarized in Table 11. The complete implementation code, test instances, and computational results are publicly available at: https://github.com/Pedram-Fdi/2Stage_HospitalEvacuation.git

Table 11: Parameters Used for Instance Generation

<i>Parameter</i>	<i>Value</i>	<i>Parameter</i>	<i>Value</i>
B	85% of total required budget for all ACFs	ν_j	$\{100000, 50000\}; \forall j \in \{1, 2\}$
c_i^{ACF}	$U[250, 500]; \forall i \in \mathcal{I}$	f_i	$c_i^{\text{ACF}} \times 1.3 \times 25; \forall i \in \mathcal{I}$
$q_{hh'}$	$U[200, 900]; \forall h, h' \in \mathcal{H}$	ρ_j	$\{150000, 75000\}; \forall j \in \{1, 2\}$
γ_m^L	$\{10, 10, 20\}; \forall m \in \{1, 2, 3\}$	γ^A	20

4.5.2 Performance on Out-of-Sample Scenarios

This section evaluates the performance of various solution approaches, including the proposed matheuristics, scenario sampling techniques, and clustering methods, on 500 out-of-sample scenarios. As a baseline, we include the PHA, a widely adopted decomposition method in disaster response optimization (Kamyabniya et al., 2024). We have provided additional details on the PHA implementation in the Appendix. Table 12 compares the performance of several algorithmic configurations, each combining different sampling and clustering strategies. The following notation is used to represent the tested variants:

- MIP_{RN} : MIP with RQMC sampling and no clustering.

- $ALNS_{RN}$: Vanilla ALNS with RQMC sampling and no clustering.
- $DRL-ALNS_{RN}$: DRL-ALNS using RQMC sampling and no clustering.
- $DRL-ALNS_{RD}$: DRL-ALNS with RQMC sampling and decision-based clustering.
- $DRL-ALNS_{RK++}$: DRL-ALNS with RQMC sampling and K-means++ clustering.
- PHA_{RN} : PHA with RQMC sampling and no clustering.

Note: The MIP_{QN} variant (MIP with QMC sampling and no clustering) was excluded from Table 12 due to space constraints but is still included in the performance analysis below.

These configurations are evaluated based on their ability to generalize to unseen post-disaster scenarios, balancing solution quality and computational efficiency. The first three columns of Table 12 describe the characteristics of each instance, including the number of planning periods ($\mathcal{T}$), ACF locations ($\mathcal{J}$), and disaster sites ($\mathcal{L}$). Each row presents the average values across five problem instances. The key performance indicators (KPIs) are as follows:

- **DT (Delayed Transfers)**: Average number of casualties not transferred on time.
- **DE (Delayed Evacuation)**: Average number of patients not evacuated on time.
- **NE (Not Evacuated)**: Average number of patients who are not evacuated at all during the planning horizon.

The Gap (%) column reports the performance gap, calculated as:

$$Gap = 100 \times \left(\frac{UB(m) - UB^*}{UB^*} \right),$$

where $UB(m)$ is the upper bound of a given method and UB^* is the best-known upper bound among all methods. The VSS (%) column shows the value of the stochastic solution, computed as $VSS = 100 \times \left(\frac{EEV - UB}{EEV} \right)$, where EEV is the expected value solution. Bold values in the table indicate the best performance for each metric.

According to Table 12, $DRL-ALNS_{RD}$ demonstrates the best overall performance, achieving the lowest average gap (0.75%) from the best upper bound and the most favorable KPIs — 9 delayed transfers (DT), 55 delayed evacuations (DE), and only 32 patients not evacuated (NE) during the planning horizon. $DRL-ALNS_{RN}$ also performed well, with a slightly higher gap of 0.99% and 2 additional delayed evacuations, highlighting the

benefits of decision-based clustering. This clustering method reduces the number of scenarios while preserving diversity in the scenario space. This leads to better out-of-sample generalization compared to $DRL-ALNS_{RK++}$, which relies solely on clustering based on scenario parameter values rather than the resulting objective values, thereby limiting its effectiveness.

Among the methods that successfully solved all instances without memory issues, $DRL-ALNS_{RK++}$ showed the weakest performance in terms of both optimality gap and KPI outcomes. The impact of intelligent action selection is evident when comparing $ALNS_{RN}$ and $DRL-ALNS_{RN}$, where the only difference is the use of deep reinforcement learning (DRL) in the latter. Incorporating DRL reduced the optimality gap from 2.17% to 0.99%, and improved the number of delayed evacuations from 60 to 57, underscoring the life-saving potential of such enhancements. In healthcare operations, even saving a single additional life is critical, highlighting the importance of adopting high-performing solution approaches.

In contrast, the baseline methods MIP_{RN} , MIP_{QN} , and PHA_{RN} failed to produce feasible solutions for medium to large instances within the 10-hour time limit. Even in the few cases where PHA_{RN} did converge, the average optimality gap remained high at 83.24%, indicating poor overall performance. Finally, the value of using RQMC in scenario generation is evident from the comparison between MIP_{RN} and MIP_{QN} . The former, which uses RQMC, achieved a 27.33% improvement in optimality gap over the latter (which uses QMC), based on the subset of instances successfully solved by both approaches. This demonstrates the advantage of introducing randomness to improve scenario coverage and solution quality.

4.5.3 Performance of the Algorithms on In-Sample Scenarios

In this section, we evaluate whether the learning-based matheuristic we developed can produce acceptable solutions compared to an exact method — specifically, solving the MIP using Gurobi. Table 13 presents a comparison of the in-sample performance of four methods: exact MIP (solved by Gurobi), PHA, ALNS, and DRL-ALNS. For a fair comparison, all methods are tested using RQMC-generated scenarios without any clustering. This ensures that differences in performance are due solely to the algorithms themselves, not to variations in scenario generation or clustering, which could otherwise distort the results.

The comparison is conducted only for instances where the exact solution could be obtained using Gurobi within a 10-hour time limit. As expected, the exact MIP achieves

Table 12: Out-of-Sample ALNS Performance with Various Scenario Generation and Clustering Methods.

$\mathcal{T}$	$\mathcal{I}$	$\mathcal{L}$	MIP_{RN}				$ALNS_{RN}$				$DRL - ALNS_{RN}$				$DRL - ALNS_{RD}$				$DRL - ALNS_{RK++}$				PHA_{RN}				VSS (%)	UB*
			DT	DE	NE	Gap (%)	DT	DE	NE	Gap (%)	DT	DE	NE	Gap (%)	DT	DE	NE	Gap (%)	DT	DE	NE	Gap (%)	DT	DE	NE	Gap (%)		
4	15	15	0	21	4	24.02	0	24	2	7.83	0	23	2	2.58	0	20	3	9.74	0	33	15	62.70	0	260	259	97.61	97.35	107,659
		20	0	46	18	1.56	0	48	17	0.69	0	47	17	0.53	0	45	17	1.11	0	52	24	18.92	0	259	259	90.94	91.24	410,879
		25	5	58	34	0.00	5	59	34	0.43	5	60	34	0.37	5	60	34	0.16	5	64	40	8.58	5	259	259	76.36	79.58	1,176,818
		30	14	57	40	0.00	14	56	40	0.21	14	58	40	0.18	14	56	40	0.03	14	61	47	6.20	14	331	207	62.68	69.57	2,258,734
	20	15	0	44	5	32.88	0	43	2	4.52	0	44	2	11.14	0	41	2	3.42	0	48	7	35.57	0	259	259	97.60	97.52	111,448
		20	0	50	13	3.61	0	51	13	1.66	0	56	13	1.06	0	51	13	1.58	0	58	24	36.63	0	259	259	92.72	92.68	348,182
		25	3	54	29	0.17	3	54	29	0.47	3	56	29	0.41	3	56	29	0.22	3	55	31	2.12	3	259	259	80.76	81.24	989,847
		30	13	65	39	0.00	13	67	39	0.23	13	66	39	0.20	13	64	39	0.03	13	70	46	5.65	13	259	259	63.84	69.59	2,250,289
	25	15	0	61	3	11.64	0	61	2	1.73	0	61	2	1.33	0	61	2	2.13	0	68	15	61.00	0	261	259	97.51	97.46	114,984
		20	-	-	-	-	0	45	17	1.39	0	47	17	0.36	0	44	17	0.00	0	49	21	11.48	-	-	-	-	91.75	390,176
		25	-	-	-	-	2	65	30	0.39	2	68	30	0.12	2	65	30	0.04	2	72	35	11.02	-	-	-	-	82.66	918,155
		30	-	-	-	-	12	69	38	0.21	12	68	38	0.13	12	69	38	0.03	12	75	45	5.70	-	-	-	-	69.77	1,943,036
	30	15	-	-	-	-	0	56	2	13.57	0	47	1	2.81	0	45	1	0.51	0	53	10	50.60	-	-	-	-	97.49	116,401
		20	-	-	-	-	0	30	19	0.83	0	32	20	1.42	0	31	19	0.11	0	33	23	9.93	-	-	-	-	89.86	449,465
		25	-	-	-	-	4	43	32	0.39	4	45	32	0.25	4	42	32	0.08	4	50	41	12.96	-	-	-	-	80.34	1,031,911
		30	-	-	-	-	15	53	44	0.84	15	52	43	0.07	15	49	43	0.02	15	55	47	2.73	-	-	-	-	64.69	2,325,538
	4-T Avg.	-	-	-	-	-	4	52	23	2.21	4	52	22	1.44	4	50	22	1.20	4	56	29	21.36	-	-	-	-	84.55	933,970
5	15	15	0	34	10	8.64	0	35	8	0.17	0	33	8	0.46	0	31	9	1.47	0	43	18	42.80	0	260	259	95.77	95.71	353,268
		20	2	56	32	0.03	2	63	33	2.93	2	57	32	0.31	2	56	32	0.14	2	63	39	15.77	2	259	259	85.08	85.80	1,266,668
		25	15	66	41	0.01	15	74	42	0.46	15	67	41	0.13	15	67	41	0.02	15	74	51	10.97	15	259	259	70.67	74.51	2,974,258
		30	35	105	81	0.00	35	110	81	0.24	35	106	81	0.05	35	105	81	0.01	35	105	83	1.08	35	259	259	48.44	59.36	6,559,045
	20	15	0	29	7	10.26	0	29	6	4.59	0	29	6	2.02	0	28	6	1.00	0	33	13	40.16	0	259	259	96.67	96.64	289,754
		20	1	47	26	0.64	1	61	25	0.43	1	57	25	0.35	1	47	25	0.08	1	51	33	17.30	1	260	259	88.61	88.91	1,034,082
		25	-	-	-	-	14	65	39	0.20	14	61	39	0.12	14	60	39	0.00	14	63	45	6.91	-	-	-	-	75.89	2,765,077
		30	-	-	-	-	33	119	81	0.05	33	123	81	0.04	33	120	81	0.00	33	122	82	0.74	-	-	-	-	52.28	6,726,980
	25	15	-	-	-	-	1	37	10	18.06	0	32	8	3.07	0	24	8	2.01	0	31	16	31.24	-	-	-	-	95.84	342,496
		20	-	-	-	-	1	66	29	2.11	1	50	28	0.18	1	50	28	0.02	1	55	36	18.11	-	-	-	-	87.60	1,110,926
		25	-	-	-	-	14	65	40	0.13	14	65	40	0.08	14	65	40	0.00	14	71	49	10.41	-	-	-	-	76.04	2,684,269
		30	-	-	-	-	34	143	84	0.61	34	97	84	0.04	34	96	84	0.02	34	97	85	0.32	-	-	-	-	53.52	6,858,238
	30	15	-	-	-	-	0	23	9	3.44	0	21	9	1.57	0	21	9	0.00	0	33	21	49.65	-	-	-	-	95.75	351,563
		20	-	-	-	-	3	48	30	0.35	3	46	30	0.23	3	43	30	0.00	3	47	35	9.45	-	-	-	-	84.62	1,335,683
		25	-	-	-	-	18	60	46	0.32	18	51	46	0.07	18	52	46	0.01	18	60	56	8.47	-	-	-	-	67.74	3,392,458
		30	-	-	-	-	40	100	96	0.06	40	100	96	0.05	40	100	96	0.01	40	101	97	0.56	-	-	-	-	46.84	7,726,144
	5-T Avg.	-	-	-	-	-	13	69	41	2.13	13	62	41	0.55	13	60	41	0.30	13	66	47	16.50	-	-	-	-	77.32	2,860,682
	Total Avg.	-	-	-	-	-	9	60	32	2.17	9	57	32	0.99	9	55	32	0.75	9	61	38	18.93	-	-	-	-	80.93	1,897,326

Note. “—” indicates that the solver ran out of memory.

Table 13: In-Sample Algorithm Performance.

Problem Size			MIP_{RN}		PHA_{RN}	$ALNS_{RN}$	$DRL - ALNS_{RN}$
$ \mathcal{T} $	$ \mathcal{J} $	$ \mathcal{L} $	Cost	Gap(%)	Gap(%)	Gap(%)	Gap(%)
4	15	15	248,911	0	94.48	1.38	1.33
		20	720,631	0	84.74	0.51	0.44
		25	1,403,768	0	73.08	0.35	0.31
		30	2,404,389	0	61.54	0.20	0.17
	20	15	266,899	0	94.20	1.35	1.31
		20	681,055	0	86.40	0.54	0.59
		25	1,258,363	0	76.67	0.41	0.39
		30	2,408,945	0	61.93	0.22	0.19
	25	15	274,713	0	94.06	1.41	1.17
	4-T Avg.		1,074,186	0	81.15	0.71	0.66
5	15	15	622,373	0	92.49	0.67	0.63
		20	1,675,184	0	81.07	2.68	0.25
		25	3,113,694	0	69.61	0.32	0.13
		30	6,732,868	0	47.38	0.40	0.05
	20	15	645,502	0	92.68	1.89	0.47
		20	1,433,189	0	84.60	0.43	0.37
	5-T Avg.		2,370,468	0	77.97	1.06	0.32
Total Average			1,595,517	0	79.86	0.85	0.52

the best performance on in-sample scenarios across all tested instances. Notably, DRL-ALNS comes closest to the MIP solution, with only a 0.52% average gap. This strong and consistent performance suggests that DRL-ALNS is a robust approach we can rely on for real-life instances where exact solutions are no longer feasible. In contrast, when deep reinforcement learning is not used, i.e., in the standard ALNS, the performance gap increases from 0.52% to 0.85%. This highlights the significant impact of intelligent, state-based decision-making in ALNS. PHA performs poorly even on small to medium instances, with an average gap of approximately 80% from the optimal solution.

4.6 Case Study

To assess the performance of the proposed algorithm, we applied it to a CRP problem inspired by real-world data from the 2011 earthquake in Van Province, Turkey. This region, known for its seismic vulnerability, was struck by a powerful magnitude 7.2 earthquake on

October 23, 2011, leading to over 600 deaths and approximately 2,500 injuries (International Federation of Red Cross and Red Crescent Societies, 2012). Our case study centers on identifying optimal locations for ACFs and coordinating evacuation plans for hospitals potentially impacted by such a disaster in this earthquake-prone area.

The earthquake impacted 94 distinct zones, indicated by red stars in Figure B.7a, where marker size is proportional to demand, and 31 potential ACF locations, represented by tents. Based on this setting, Noyan et al. (2016) developed clustered versions of the problem by aggregating the 94 demand points into 30- and 60-location sets using a p -median approach to minimize total travel time weighted by demand. We evaluate our model on both clustered and original instances. Each neighborhood is classified into one of three seismic intensity categories, *destructive*, *damaging*, or *very strong*, based on proximity to the fault and ground shaking severity. These categories correspond to different damage state distributions, ranging from *no damage* to *medium damage* to *heavy damage or collapse*. To determine baseline demand at each location, we incorporate neighborhood population, seismic intensity level, building damage state, and the overall injury count from the 2011 event (see Appendix in Noyan et al. (2016) for details). Scenario-specific demand is generated by multiplying these base values by random deviation factors, which are drawn from uniform distributions: $[0.75, 1.10]$ for destructive, $[0.75, 1.20]$ for damaging, and $[0.75, 1.30]$ for very strong areas.

The region contains four major hospitals, shown as plus-marked circles in Figure B.7a, that are designated for emergency response. However, each hospital faces a 45% chance of being non-operational during a seismic event, as reported in Liel and Deierlein (2012). In the case of such disruptions, occupied beds, assumed to range uniformly between 50% and 60% of the hospital's capacity, must be evacuated. All travel distances used in this study are computed based on geographic coordinates.

We evaluate KPI performance using four modeling approaches: (1) deterministic parameters (**Det**), (2) stochastic with exponential threat risk (**StoExpRisk**, see Figure 6b), (3) stochastic with linear threat risk (**StoLinRisk**), and (4) stochastic with constant threat risk (**StoConsRisk**).

Impact on ACF Establishment and Resource Allocation: Compared to the deterministic model, the stochastic approaches recommend establishing one additional ACF to better accommodate demand fluctuations under uncertainty. This results in over 300 extra available beds, enhancing surge capacity and improving the system's resilience to scenario variability.




Figure 6: Van Network: (a) Full layout, (b) StoExpRisk layout, (c) StoConsRisk Hos., (d) StoExpRisk Hos..

Comparison of Patient Evacuation Delays: Stochastic models, particularly **StoExpRisk**, substantially reduce patient evacuation delays compared to the deterministic baseline. Specifically, the number of delayed evacuations in **Det** is approximately 3.06 times higher than in both **StoLinRisk** and **StoConsRisk**, and 3.78 times higher than in **StoExpRisk**. These findings underscore the importance of explicitly modeling uncertainties such as casualty numbers, hospital disruptions, and discharge rates. The average value of the stochastic solution (VSS) across instances is 62.53%, further emphasizing the value of incorporating stochasticity into disaster response planning.

Additional insights emerge when comparing how the models behave under different threat risk penalization strategies. When an exponential penalty is applied to patients remaining in disrupted hospitals, as in **StoExpRisk**, the model proactively assigns backup hospitals to ensure continuity of care. For instance, hospital #4, although geographically remote, is selected as a backup for hospitals #1 and #2 (see Figure 6d). In contrast, no such coordination is observed in **StoConsRisk** (see Figure 6c), and in **Det**, no hospital serves as a backup for others, which results in the worst KPI performance. Compared to **Det**, the number of non-evacuated patients is reduced by 73.31% in **StoConsRisk**. Applying exponential penalties in **StoExpRisk** further reduces this number by an additional 6.32%, highlighting the importance of using appropriately scaled risk penalties to enhance life-saving outcomes in CHEs. It is also worth noting that the number of non-evacuated patients increases by an average of one person in **StoLinRisk** compared to **StoExpRisk**. Evacuating even one additional patient can be vital in disaster response, as it may lead to saving a life and reducing preventable harm.

4.7 Conclusion

This paper addressed the complex challenge of CRP, with a particular focus on hospital evacuation during large-scale disasters. We developed a learning-based hybrid algorithm that combines deep reinforcement learning (DRL) with adaptive large neighborhood search (ALNS), referred to as *DRL-ALNS*, and benchmarked its performance against the widely used PHA in disaster operations. To improve generalization on out-of-sample scenarios, we incorporated a decision-based scenario clustering technique that groups scenarios based on decision impact rather than parameter similarity, and compared it with traditional K-means++ clustering. We also enhanced scenario generation using randomized quasi-Monte Carlo (RQMC) sampling, which outperformed standard QMC by reducing the average optimality gap by 27.33% across instances solvable by both methods. Among all tested configurations, the variant using both RQMC and decision-based clustering, denoted as *DRL-ALNS_{RD}*, achieved the best overall performance. It successfully solved all instances with an average optimality gap of just 0.75%, compared to 0.99% for *DRL-ALNS_{RN}* and 83.24% for PHA on larger instances. Moreover, it outperformed other approaches across all key performance indicators on 500 out-of-sample scenarios, with only 9 delayed transfers (DT), 55 delayed evacuations (DE), and 32 non-evacuated patients (NE). In contrast, PHA and other exact methods (e.g., solving the extensive form) failed to produce feasible solutions for medium to large instances within the 10-hour time limit. Even in the few cases where PHA did converge, its solutions were highly suboptimal, with an average gap of 83.24%.

A case study based on the 2011 Van earthquake in Turkey further validated the real-world applicability of our approach. The *DRL-ALNS_{RD}* model with exponential threat risk performed best across all KPIs: average delay in casualty transport (DT), patient evacuation delay (DE), and number of non-evacuated patients (NE). Specifically, it reduced NE by up to 73.31% compared to the deterministic model, and outperformed the stochastic models with constant and linear threat risk.

Future research could extend this work in two promising directions. First, integrating road clearance and debris removal into the CRP framework would provide a more realistic representation of post-disaster logistics, as blocked or damaged roads can significantly delay casualty and patient transport. Second, considering the possibility of secondary disasters, such as aftershocks or cascading events that render designated backup hospitals non-operational, would require re-optimizing the network dynamically. Modeling such

evolving disruptions would further strengthen the resilience and adaptability of the proposed framework in real-world disaster settings.

Chapter 5

Conclusion

This thesis tackled three important challenges in casualty response planning (CRP) and the establishment of alternative care facilities (ACFs) immediately following catastrophic health events (CHEs): managing casualties with multiple injuries, optimizing platelet inventory and the deployment of portable apheresis machines, and incorporating hospital evacuation into CRP strategies.

In Chapter 2, we addressed the challenge of managing patients with multiple injuries by formulating two equivalent two-stage stochastic mixed-integer programming models that account for uncertain casualty demand and hospital treatment capacities. The second model significantly reduces solution time while maintaining mathematical equivalence with the first. To solve these models efficiently, we developed two decomposition-based algorithms: an L-shaped method and a branch-and-Benders-cut (B&BC) algorithm. We enhanced their performance through a series of advanced acceleration techniques, including Benders and Lagrangian dual decompositions, multi-cut reformulations, Pareto-optimal cuts, and lower bounding functional valid inequalities (LBFVIs). Our computational results showed a substantial reduction in the optimality gap — from over 7,900% in some instances down to about 1% — and a 42.5% improvement in runtime. The average value of the stochastic solution (VSS) was 95.7%, highlighting the critical role of uncertainty modeling. A case study based on an earthquake scenario in Turkey confirmed the practical applicability of our methods, with our best-performing algorithm, BBC_I , outperforming the progressive

hedging algorithm (PHA) by 90.76% in solution quality and with lower computational burden.

In Chapter 3, we investigated a novel CRP model that incorporates platelet inventory management and portable apheresis machine deployment, both of which are vital for surgical operations in disaster response. To solve this complex multi-stage stochastic problem, we proposed and benchmarked two algorithms — Stochastic Dual Dynamic Programming (SDDP) and Nested Benders Decomposition (NBD). To improve computational efficiency, we introduced strong cuts, multiple cuts, warm-up strategies, and LBFVIs. Additionally, we conducted a systematic analysis of scenario sampling and clustering strategies by combining quasi-Monte Carlo (QMC) and randomized QMC (RQMC) sampling with different clustering methods, including self-organizing maps (SOM) and K-means++. The best results were achieved by the $RKSDDP^{++}$ algorithm, which utilized RQMC sampling and K-means++ clustering. Unlike the NBD method, which ran into memory issues beyond three stages, $RKSDDP^{++}$ successfully solved all tested instances, achieving an average optimality gap of 0.86%.

In Chapter 4, we extended the CRP framework to incorporate hospital evacuation, addressing one of the most critical yet underexplored aspects of disaster response. We proposed a novel learning-based matheuristic that combines deep reinforcement learning with adaptive large neighborhood search (DRL-ALNS). This algorithm was further improved by integrating a decision-based scenario clustering method, which groups scenarios based on their decision impact, and by using RQMC sampling for more diverse scenario generation. The combination of RQMC and decision-based clustering, denoted as $DRL-ALNS_{RD}$, yielded the best performance across all tested configurations. In contrast to exact methods like Gurobi and PHA, which failed for instances with more than 25 potential ACFs due to memory issues, $DRL-ALNS_{RD}$ consistently solved all instances with an average optimality gap of just 0.75%.

Building upon the findings of this thesis, several promising research directions can be pursued to further advance casualty response planning (CRP) under uncertainty:

- **Integration of Road Conditions and Debris Clearance:** Post-disaster road accessibility is often compromised by debris and infrastructure damage, significantly affecting patient transport and resource delivery. Future models could incorporate debris clearance as an integrated subproblem within the CRP framework. This would enhance the realism of logistical constraints and allow for optimized routing that reflects evolving road conditions.
- **Modeling of Secondary Disasters and Cascading Failures:** Disasters often trigger secondary events such as aftershocks, fires, or hazardous material leaks, which can disrupt pre-planned medical facilities or evacuation paths. Incorporating secondary disasters into CRP models would allow for dynamic re-optimization of resource allocation and facility usage, thereby improving adaptability in unstable environments. Understanding the probability and impact of such cascading disruptions is essential for enhancing the resilience of the healthcare response network.
- **Incorporation of Travel Time and Service Time Uncertainties:** While this thesis considers uncertainty in demand and treatment capacities, future work could also model the variability in travel times (e.g., due to blocked roads or traffic congestion) and service times (e.g., due to staffing shortages or equipment delays). These uncertainties influence the efficiency of casualty transport and treatment and could be modeled through scenario-based approaches or robust optimization techniques to better capture the dynamics of post-disaster operations.
- **Integrating Crowdsourcing into Casualty Response Planning:** Crowdsourcing refers to the method of recruiting and organizing ad-hoc labor through open calls for participation. In disaster settings, it can be used to mobilize spontaneous volunteers, collect real-time information, and enhance the flexibility of response strategies. The crowd has the freedom to choose when, where, and how to contribute, making it essential for model designers to develop robust coordination mechanisms. Future research could investigate how to incorporate crowdsourced data into CRP models, allocate spontaneous volunteers efficiently under uncertainty, and leverage

incentive-based or learning-based strategies to guide participation. Challenges such as information verification and alignment with official response protocols also warrant further exploration.

Bibliography

- Alizadeh, M., Amiri-Aref, M., Mustafee, N., and Matilal, S. (2019). A robust stochastic casualty collection points location problem. *European Journal of Operational Research*, 279(3):965–983.
- Amani, M., Asumadu Sarkodie, S., Sheu, J.-B., Mahdi Nasiri, M., and Tavakkoli-Moghaddam, R. (2025). A data-driven hybrid scenario-based robust optimization method for relief logistics network design. *Transportation Research Part E: Logistics and Transportation Review*, 194:103931.
- Apte, A., Heidtke, C., and Salmerón, J. (2015). Casualty collection points optimization: A study for the district of columbia. *Interfaces*, 45(2):149–165.
- Arthur, D. and Vassilvitskii, S. (2006). k-means++: The advantages of careful seeding. Technical report, Stanford.
- Bagaria, J., Heggie, C., Abrahams, J., and Murray, V. (2009). Evacuation and sheltering of hospitals in emergencies: A review of international experience. *Prehospital and Disaster Medicine*, 24(5):461–467.
- Balcik, B. and İhsan Yanıkoğlu (2020). A robust optimization approach for humanitarian needs assessment planning under travel time uncertainty. *European Journal of Operational Research*, 282(1):40–57.
- Batun, S., Denton, B. T., Huschka, T. R., and Schaefer, A. J. (2011). Operating room pooling and parallel surgery processing under uncertainty. *INFORMS Journal on Computing*, 23(2):220–237.
- Bayram, V. and Yaman, H. (2018). Shelter location and evacuation route assignment under uncertainty: A benders decomposition approach. *Transportation Science*, 52(2):416–436.
- Birge, J. R. and Louveaux, F. (2011). *Introduction to Stochastic Programming*. Springer-Verlag, New York, NY, USA, 2nd edition.
- Birge, J. R. and Louveaux, F. V. (1988). A multicut algorithm for two-stage stochastic linear programs. *European Journal of Operational Research*, 34(3):384–392.

- Bish, D. R., Agca, E., and Glick, R. (2014). Decision support for hospital evacuation and emergency response. *Annals of Operations Research*, 221:89–106.
- Blumberg, N., Refaai, M., and Heal, J. (2015). Abo matching of platelet transfusions-“start making sense”“as we get older, and stop making sense...”-the talking heads (1984). *Blood Transfusion*, 13(3):347.
- Bruno, S., Ahmed, S., Shapiro, A., and Street, A. (2016). Risk neutral and risk averse approaches to multistage renewable investment planning under uncertainty. *European Journal of Operational Research*, 250(3):979–989.
- Canadian Blood Services (2023). Blood types in canada: how common or rare are they? Available at: <https://www.blood.ca/en/stories/blood-types-canada-how-common-or-rare-are-they>.
- Caunhye, A. M., Li, M., and Nie, X. (2015). A location-allocation model for casualty response planning during catastrophic radiological incidents. *Socio-Economic Planning Sciences*, 50:32–44.
- Caunhye, A. M. and Nie, X. (2018). A stochastic programming model for casualty response planning during catastrophic health events. *Transportation Science*, 52(2):437–453.
- Chan, C. W., Green, L. V., Lekwijit, S., Lu, L., and Escobar, G. (2019). Assessing the impact of service level when customer needs are uncertain: An empirical investigation of hospital step-down units. *Management Science*, 65(2):751–775.
- Chang, K.-H., Chen, T.-L., Yang, F.-H., and Chang, T.-Y. (2023). Simulation optimization for stochastic casualty collection point location and resource allocation problem in a mass casualty incident. *European Journal of Operational Research*, 309(3):1237–1262.
- Childers, A. and Taaffe, K. M. (2010). Healthcare facility evacuations: Lessons learned, research activity, and the need for engineering contributions. *Journal of Healthcare Engineering*, 1(1):125–140.
- Clingerman, J. (2019). *Estimating Mortality after Natural Disasters in Low- and Middle-Income Countries: A Systematic Review*. PhD thesis.
- Codato, G. and Fischetti, M. (2006). Combinatorial benders’ cuts for mixed-integer linear programming. *Operations Research*, 54(4):756–766.
- Costa, A. d. S. (2017). Assessment of operative times of multiple surgical specialties in a public university hospital. *Einstein (Sao Paulo)*, 15:200–205.
- Crainic, T. G., Hewitt, M., Maggioni, F., and Rei, W. (2021). Partial benders decomposition: General methodology and application to stochastic network design. *Transportation Science*, 55(2):414–435.

- Cranley, R. and Patterson, T. N. (1976). Randomization of number theoretic methods for multiple integration. *SIAM Journal on Numerical Analysis*, 13(6):904–914.
- Dalal, J. and Üster, H. (2018). Combining worst case and average case considerations in an integrated emergency response network design problem. *Transportation Science*, 52(1):171–188.
- Del Papa, J., Vittorini, P., D’Aloisio, F., Muselli, M., Giuliani, A. R., Mascitelli, A., and Fabiani, L. (2019). Retrospective analysis of injuries and hospitalizations of patients following the 2009 earthquake of l’aquila city. *International Journal of Environmental Research and Public Health*, 16(10):1675.
- Dong, J.-X., Lee, C.-Y., and Song, D.-P. (2015). Joint service capacity planning and dynamic container routing in shipping network with uncertain demands. *Transportation Research Part B: Methodological*, 78:404–421.
- Doocy, S., Daniels, A., Packer, C., Dick, A., and Kirsch, T. D. (2013). The human impact of earthquakes: A historical review of events 1980-2009 and systematic literature review. *PLoS currents*, 5.
- Dönmez, Z., Kara, B. Y., Özlem Karsu, and da Gama, F. S. (2021). Humanitarian facility location under uncertainty: Critical review and future prospects. *Omega*, 102:102393.
- Elçi, Ö., Noyan, N., and Bülbül, K. (2018). Chance-constrained stochastic programming under variable reliability levels with an application to humanitarian relief network design. *Computers & Operations Research*, 96:91–107.
- Emelogu, A., Chowdhury, S., Marufuzzaman, M., Bian, L., and Eksioglu, B. (2016). An enhanced sample average approximation method for stochastic optimization. *International Journal of Production Economics*, 182:230–252.
- Farahani, R. Z., Lotfi, M., Baghaian, A., Ruiz, R., and Rezapour, S. (2020). Mass casualty management in disaster scene: A systematic review of or&ms research in humanitarian operations. *European Journal of Operational Research*, 287(3):787–819.
- Ghasemi, P., Khalili-Damghani, K., Hafezalkotob, A., and Raissi, S. (2019). Uncertain multi-objective multi-commodity multi-period multi-vehicle location-allocation model for earthquake evacuation planning. *Applied Mathematics and Computation*, 350:105–132.
- Grass, E., Ortmann, J., Balcik, B., and Rei, W. (2023). A machine learning approach to deal with ambiguity in the humanitarian decision-making. *Production and Operations Management*, 32(9):2956–2974.
- Gu, J., Zhou, Y., Das, A., Moon, I., and Lee, G. M. (2018). Medical relief shelter location problem with patient severity under a limited relief budget. *Computers & Industrial Engineering*, 125:720–728.

- Gul, S., Denton, B. T., and Fowler, J. W. (2015). A progressive hedging approach for surgery planning under uncertainty. *INFORMS Journal on Computing*, 27(4):755–772.
- Helminen, M., Saarela, E., and Salmela, J. (2006). Characterisation of patients treated at the red cross field hospital in kashmir during the first three weeks of operation. *Emergency Medicine Journal*, 23(8):654–656.
- Hewitt, M., Ortmann, J., and Rei, W. (2022). Decision-based scenario clustering for decision-making under uncertainty. *Annals of Operations Research*, 315(2):747–771.
- Hosseinifard, Z., Abbasi, B., Fadaki, M., and Clay, N. M. (2020). Postdisaster volatility of blood donations in an unsteady blood supply chain. *Decision Sciences*, 51(2):255–281.
- Hu, Q.-M., Zhao, L., Li, H., and Huang, R. (2019a). Integrated design of emergency shelter and medical networks considering diurnal population shifts in urban areas. *IIE Transactions*, 51(6):614–637.
- Hu, S., Han, C., Dong, Z. S., and Meng, L. (2019b). A multi-stage stochastic programming model for relief distribution considering the state of road network. *Transportation Research Part B: Methodological*, 123:64–87.
- International Federation of Red Cross and Red Crescent Societies (2012). Emergency appeal operation update turkey: Van earthquake. Online.
- Jensen, J. L. W. V. (1906). Sur les fonctions convexes et les inégalités entre les valeurs moyennes. *Acta Mathematica*, 30(none):175 – 193.
- Jr, E. J. L. and Taskin, S. (2008). An insurance risk management framework for disaster relief and supply chain disruption inventory planning. *Journal of the Operational Research Society*, 59(5):674–684.
- Kamyabniya, A., Noormohammadzadeh, Z., Sauré, A., and Patrick, J. (2021). A robust integrated logistics model for age-based multi-group platelets in disaster relief operations. *Transportation Research Part E: Logistics and Transportation Review*, 152:102371.
- Kamyabniya, A., Sauré, A., Salman, F. S., Bénichou, N., and Patrick, J. (2024). Optimization models for disaster response operations: a literature review. *OR Spectrum*, 46(3):737–783.
- Kang, P., Tang, B., Liu, Y., Liu, X., Liu, Z., Lv, Y., and Zhang, L. (2015). Medical efforts and injury patterns of military hospital patients following the 2013 lushan earthquake in china: a retrospective study. *International Journal of Environmental Research and Public Health*, 12(9):10723–10738.
- Kohonen, T. (1982). Self-organized formation of topologically correct feature maps. *Biological cybernetics*, 43(1):59–69.

- Kılıcı, F., Kara, B. Y., and Bozkaya, B. (2015). Locating temporary shelter areas after an earthquake: A case for turkey. *European Journal of Operational Research*, 243(1):323–332.
- Larson, R. C., Metzger, M. D., and Cahn, M. F. (2006). Responding to emergencies: Lessons learned and the need for analysis. *Interfaces*, 36(6):486–501.
- Lee, G. and Loh, C.-H. (1999). Preliminary report from mceer-ncree workshop on the 921 taiwan earthquake. Accessed 11 Feb 2025.
- Li, Y., Mohammadi, M., Zhang, X., Lan, Y., and van Jaarsveld, W. (2024). Integrated trucks assignment and scheduling problem with mixed service mode docks: A q-learning based adaptive large neighborhood search algorithm. *arXiv preprint arXiv:2412.09090*.
- Li, Y., Zhang, J., and Yu, G. (2020). A scenario-based hybrid robust and stochastic approach for joint planning of relief logistics and casualty distribution considering secondary disasters. *Transportation Research Part E: Logistics and Transportation Review*, 141:102029.
- Liel, A. B. and Deierlein, G. G. (2012). Using collapse risk assessments to inform seismic safety policy for older concrete buildings. *Earthquake Spectra*, 28(4):1495–1521.
- Liu, K., Yang, L., Zhao, Y., and Zhang, Z.-H. (2023). Multi-period stochastic programming for relief delivery considering evolving transportation network and temporary facility relocation/closure. *Transportation Research Part E: Logistics and Transportation Review*, 180:103357.
- Liu, Y., Cui, N., and Zhang, J. (2019). Integrated temporary facility location and casualty allocation planning for post-disaster humanitarian medical service. *Transportation Research Part E: Logistics and Transportation Review*, 128:1–16.
- Lloyd, S. (1982). Least squares quantization in pcm. *IEEE Transactions on Information Theory*, 28(2):129–137.
- Long, Y., Lee, L. H., and Chew, E. P. (2012). The sample average approximation method for empty container repositioning with uncertainties. *European Journal of Operational Research*, 222(1):65–75.
- Lu-Ping, Z., Rodriguez-Llanes, J. M., Qi, W., van den Oever, B., Westman, L., Albela, M., Liang, P., Gao, C., De-Sheng, Z., Hughes, M., et al. (2012). Multiple injuries after earthquakes: a retrospective analysis on 1,871 injured patients from the 2008 wenchuan earthquake. *Critical Care*, 16:1–9.
- Ma, Z.-J., Wang, K.-M., and Dai, Y. (2019). An emergency blood allocation approach considering blood group compatibility in disaster relief operations. *International Journal of Disaster Risk Science*, 10:74–88.

- Magnanti, T. L. and Wong, R. T. (1981). Accelerating benders decomposition: Algorithmic enhancement and model selection criteria. *Operations Research*, 29(3):464–484.
- McDaniel, D. and Devine, M. (1977). A modified benders’ partitioning algorithm for mixed integer programming. *Management Science*, 24(3):312–319.
- Mousavi, S., Sajadi, S. M., AlemTabriz, A., and Najafi, S. E. (2024). Hybrid mathematical and simulation model for designing a hierarchical network of temporary medical centers in a disaster. *Journal of Simulation*, 18(2):119–135.
- Muylle, L., Wouters, E., De Bock, R., and Peetermans, M. E. (1992). Reactions to platelet transfusion: the effect of the storage time of the concentrate. *Transfusion Medicine*, 2(4):289–293.
- Noyan, N., Balcik, B., and Atakan, S. (2016). A stochastic optimization model for designing last mile relief networks. *Transportation Science*, 50(3):1092–1113.
- Noyan, N. and Kahvecioğlu, G. (2018). Stochastic last mile relief network design with resource reallocation. *OR Spectrum*, 40:187–231.
- Oksuz, M. K. and Satoglu, S. I. (2020). A two-stage stochastic model for location planning of temporary medical centers for disaster response. *International Journal of Disaster Risk Reduction*, 44:101426.
- Oksuz, M. K. and Satoglu, S. I. (2024). Integrated optimization of facility location, casualty allocation and medical staff planning for post-disaster emergency response. *Journal of Humanitarian Logistics and Supply Chain Management*, 14(3):285–303.
- Osorio, A. F., Brailsford, S. C., and Smith, H. K. (2018). Whole blood or apheresis donations? a multi-objective stochastic optimization approach. *European Journal of Operational Research*, 266(1):193–204.
- Papadakos, N. (2008). Practical enhancements to the magnanti–wong method. *Operations Research Letters*, 36(4):444–449.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., Blondel, M., Prettenhofer, P., Weiss, R., Dubourg, V., et al. (2011). Scikit-learn: Machine learning in python. *Journal of Machine Learning Research*, 12:2825–2830.
- Peleg, K. and Kellermann, A. L. (2009). Enhancing hospital surge capacity for mass casualty events. *JAMA*, 302(5):565–567.
- Pereira, M. and Pinto, L. (1991). Multi-stage stochastic optimization applied to energy planning. *Mathematical programming*, 52(1-3):359 – 375.
- Pisinger, D. and Ropke, S. (2007). A general heuristic for vehicle routing problems. *Computers & Operations Research*, 34(8):2403–2435.

- Rahmaniani, R., Ahmed, S., Crainic, T. G., Gendreau, M., and Rei, W. (2020). The benders dual decomposition method. *Operations Research*, 68(3):878–895.
- Rahmati, R., Neghabi, H., Bashiri, M., and Salari, M. (2023). Stochastic regional-based profit-maximizing hub location problem: A sustainable overview. *Omega*, 121:102921.
- Rambha, T., Nozick, L. K., Davidson, R., Yi, W., and Yang, K. (2021). A stochastic optimization model for staged hospital evacuation during hurricanes. *Transportation Research Part E: Logistics and Transportation Review*, 151:102321.
- Reijnen, R., Zhang, Y., Lau, H. C., and Bukhsh, Z. (2024). Online control of adaptive large neighborhood search using deep reinforcement learning. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 34, pages 475–483.
- Rockafellar, R. T. and Wets, R. J.-B. (1991). Scenarios and policy aggregation in optimization under uncertainty. *Mathematics of Operations Research*, 16(1):119–147.
- Ropke, S. and Pisinger, D. (2006). An adaptive large neighborhood search heuristic for the pickup and delivery problem with time windows. *Transportation Science*, 40(4):455–472.
- Runge, J. W. and Buddemeier, B. R. (2009). Explosions and radioactive material: A primer for responders. *Prehospital Emergency Care*, 13(4):407–419. PMID: 19731151.
- Saha, P. (2022). The best time to donate blood for a disaster is before it happens. Posted on May 28, 2022, 1:30 PM EDT.
- Salman, F. S. and Gül, S. (2014). Deployment of field hospitals in mass casualty incidents. *Computers & Industrial Engineering*, 74:37–51.
- Samani, M. R. G. and Hosseini-Motlagh, S.-M. (2023). Collaborative activities for matching supply and demand in the platelet network. *Expert Systems With Applications*, 231:120629.
- Sawano, T., Senoo, Y., Nonaka, S., Ozaki, A., Nishikawa, Y., Hori, A., Kotera, Y., Murakami, M., Zhao, T., and Tsubokura, M. (2023). Mortality risk associated with nuclear disasters depends on the time during and following evacuation of hospitals near nuclear power plants: An observational and qualitative study. *International Journal of Disaster Risk Reduction*, 85:103514.
- Setiawan, E., Liu, J., and French, A. (2019). Resource location for relief distribution and victim evacuation after a sudden-onset disaster. *IIE Transactions*, 51(8):830–846.
- Sever, M. S., Vanholder, R., and of ISN Work Group on Recommendations for the Management of Crush Victims in Mass Disasters, R. (2012). Recommendations for the management of crush victims in mass disasters. *Nephrology Dialysis Transplantation*, 27(Suppl_1):i1–i67.

- Shapiro, A. (2011). Analysis of stochastic dual dynamic programming method. *European Journal of Operational Research*, 209(1):63–72.
- Shaw, P. (1998). Using constraint programming and local search methods to solve vehicle routing problems. In *International conference on principles and practice of constraint programming*, pages 417–431. Springer.
- Sherali, H. D. and Lunday, B. J. (2013). On generating maximal nondominated benders cuts. *Annals of Operations Research*, 210:57–72.
- Squillace, N. (2010). Hospital evacuations: Historical precedence and modern preparedness. Master’s thesis, Wright State University Boonshoft School of Medicine.
- Statista Research Department (2024). Hospital acute care bed occupancy rates in select oecd countries worldwide in 2022. <https://www.statista.com/statistics/1116612/oecd-hospital-acute-care-occupancy-rates-select-countries-worldwide/>.
- Sun, H., Li, J., Wang, T., and Xue, Y. (2022). A novel scenario-based robust bi-objective optimization model for humanitarian logistics network under risk of disruptions. *Transportation Research Part E: Logistics and Transportation Review*, 157:102578.
- Sun, H., Wang, Y., and Xue, Y. (2021). A bi-objective robust optimization model for disaster response planning under uncertainties. *Computers & Industrial Engineering*, 155:107213.
- Tanaka, H., Oda, J., Iwai, A., Kuwagata, Y., Matsuoka, T., Takaoka, M., Kishi, M., Morimoto, F., Ishikawa, K., Mizushima, Y., et al. (1999). Morbidity and mortality of hospitalized patients after the 1995 hanshin-awaji earthquake. *The American Journal of Emergency Medicine*, 17(2):186–191.
- Texas Emergency Medical Task Force (2025). AMBUS - Ambulance Bus Program. https://txemtf.org/avada_portfolio/ambus/. Accessed: May 15, 2025.
- Thevenin, S., Adulyasak, Y., and Cordeau, J.-F. (2022). Stochastic dual dynamic programming for multiechelon lot sizing with component substitution. *INFORMS Journal on Computing*, 34(6):3151–3169.
- Tippong, D., Petrovic, S., and Akbari, V. (2022). A review of applications of operational research in healthcare coordination in disaster management. *European Journal of Operational Research*, 301(1):1–17.
- Turkeš, R., Sörensen, K., and Hvattum, L. M. (2021). Meta-analysis of metaheuristics: Quantifying the effect of adaptiveness in adaptive large neighborhood search. *European Journal of Operational Research*, 292(2):423–442.

- Van Slyke, R. M. and Wets, R. (1969). L-shaped linear programs with applications to optimal control and stochastic programming. *SIAM Journal on Applied Mathematics*, 17(4):638–663.
- Wang, Q. and Nie, X. (2022). A stochastic programming model for emergency supply planning considering transportation network mitigation and traffic congestion. *Socio-Economic Planning Sciences*, 79:101119.
- Windras Mara, S. T., Norcahyo, R., Jodiawan, P., Lusiantoro, L., and Rifai, A. P. (2022). A survey of adaptive large neighborhood search algorithms and applications. *Computers & Operations Research*, 146:105903.
- World Health Organization (2019). Trauma and emergency surgery kit (TESK). <https://www.who.int/emergencies/emergency-health-kits/trauma-emergency-surgery-kit-who-tesk-2019>. Accessed: 2023-08-08.
- Wu, C., Eftekhar, M., and Uichanco, J. (2024). Workforce configuration in charity settings: A forward-looking approach. *Manufacturing & Service Operations Management*, 26(6):2102–2120.
- Xing, X., Song, D., Qiu, C., Drake, P. R., and Zhan, Y. (2023). Joint tank container demurrage policy and flow optimisation using a progressive hedging algorithm with expanded time-space network. *European Journal of Operational Research*, 307(2):663–679.
- Xu, Y. and Szmerekovsky, J. (2023). The impact of transshipment on an integrated platelet supply chain: A multi-stage stochastic programming approach. *Computers & Industrial Engineering*, 176:108991.
- Yi, W., Nozick, L., Davidson, R., Blanton, B., and Colle, B. (2017). Optimization of the issuance of evacuation orders under evolving hurricane conditions. *Transportation Research Part B: Methodological*, 95:285–304.
- Yin, Y., Wang, J., Chu, F., and Wang, D. (2024). Distributionally robust multi-period humanitarian relief network design integrating facility location, supply inventory and allocation, and evacuation planning. *International Journal of Production Research*, 62(1-2):45–70.
- Yu, X., Ahmed, S., and Shen, S. (2019). On the value of multistage stochastic facility location with (or without) risk aversion. Technical report, Technical report. University of Michigan, Department of Industrial and Operations Engineering.
- Zhang, W., Huang, C., Gao, J., and Hou, X. (2025). Robust location-allocation decision considering casualty prioritization in multi-echelon humanitarian logistics network. *Information Sciences*, 695:121731.
- Zou, J., Ahmed, S., and Sun, X. A. (2019). Stochastic dual dynamic integer programming. *Mathematical Programming*, 175:461–502.

Appendix A

Appendices for Paper 1

A.1 Proof of Theorem 1

In order to prove the equality of the two models, it is sufficient to demonstrate that an algorithm exists, such as Algorithm 4 and Algorithm 5, which can transform any feasible solution of model **(P1)** into a feasible solution of model **(P2)**, and vice versa, while preserving the same objective function value. In these algorithms, we introduce a one-to-one function, $\mathcal{F}(p)$, which maps the value of p to the corresponding set of injuries. For example, $\mathcal{F}(3)$ returns the set injuries $\{1, 2\}$, $\mathcal{F}(15)$ returns the injury set $\{1, 2, 3, 4\}$, and $\mathcal{F}(0)$ corresponds to the empty set. To maintain our investigation into model equivalence, we introduce Algorithm 4, which takes the decision variable values from model **(P1)** and transforms them into equivalent variables for model **(P2)**.

It is worth highlighting that model **(P1)** and model **(P2)** are nearly identical except for three variables: $u_{pp'ihms}$, $\eta_{pp'ii'ms}$, and σ_{pis} in model **(P1)**, and u_{pihms} , $\eta_{pii'ms}$, and σ_{pjis} in model **(P2)**. Due to this similarity, our focus will be solely on these specific variables. This leads us to the subsequent Algorithm 5, where our attention turns to the decision variable values of model **(P2)**. In Algorithm 5, we aim to find feasible solutions for model **(P1)** using the variables obtained from model **(P2)**. In this algorithm we have:

- $\mathcal{Z}_{ips}$: The number of patients of class p at ACF i in scenario s .
- $\mathcal{N}_{ip'ps}$: The number of class p patients previously classified as class p' , who are now in a state of readiness for transfer from ACF i to either hospital h or another ACF in scenario s .

The objective function of model **(P1)** consists of objectives (1) and (10), while the objective function of model **(P2)** comprises objectives (1) and (21). With the understanding that $\sum_{p' \in \mathcal{P}} u_{p'pihms} = u_{pihms}$ for all $p \in \mathcal{P}, i \in \mathcal{I}, h \in \mathcal{H}, m \in \mathcal{M}, \text{ and } s \in \mathcal{S}$, and $\sum_{p' \in \mathcal{P}} \eta_{p'pii'ms} = \eta_{pii'ms}$ for all $p \in \mathcal{P}, i \in \mathcal{I}, i' \in \mathcal{I}, i' \neq i, m \in \mathcal{M}, \text{ and } s \in \mathcal{S}$, we can confidently assert that the objective functions of both models are effectively aligned. These insights establish a firm basis for confirming the equality of the objective functions, thus

providing a strong foundation for the conversion process between the two models.

Algorithm 4 Model (P1) to Model (P2)

```

1: for  $s \leftarrow 1$  to  $S$  do
2:    $\sigma_{pjis} := 0$  for  $p \in \mathcal{P}, j \in \mathcal{J}_i, i \in \mathcal{I}$ 
3:   for  $i \leftarrow 1$  to  $|\mathcal{I}|$  do
4:      $p^{\text{new}} := 0$ 
5:     for  $p \leftarrow 1$  to  $|\mathcal{P}|$  do
6:        $\mathcal{Q} := \emptyset$ 
7:       for  $p' \leftarrow 1$  to  $|\mathcal{P}|$  do
8:         if  $\mathcal{F}(p') \subset \mathcal{F}(p)$  then
9:            $\mathcal{Q} := \mathcal{F}(p) \setminus \mathcal{F}(p')$ 
10:           $t := 1$ 
11:           $p^{\text{new}} := p$ 
12:          for  $t \leftarrow 1$  to  $|\mathcal{Q}|$  do
13:             $\sigma_{p^{\text{new}} \mathcal{Q}_{\{t\}} is} += \left( \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} u_{pp'ihms} + \sum_{\substack{i' \in \mathcal{I} \\ i' \neq i}} \sum_{m \in \mathcal{M}} \eta_{pp'ii'ms} \right)$ 
14:             $p^{\text{new}} := \mathcal{F}^{-1}(\mathcal{F}(p^{\text{new}}) \setminus \mathcal{Q}_{\{t\}})$ 
15:             $\mathcal{Q}_{\{t\}} := \mathcal{Q}_{\{t+1\}}$ 
16:          end for
17:        end if
18:      end for
19:       $p^{\text{new}} := 0$ 
20:       $\mathcal{Q} := \emptyset$ 
21:       $p^{\text{new}} := p$ 
22:       $\mathcal{Q} := \mathcal{F}(p)$ 
23:       $t := 1$ 
24:      for  $t \leftarrow 1$  to  $|\mathcal{Q}|$  do
25:         $\sigma_{p^{\text{new}} \mathcal{Q}_{\{t\}} is} += \sigma_{pis}$ 
26:         $p^{\text{new}} := \mathcal{F}^{-1}(\mathcal{F}(p^{\text{new}}) \setminus \mathcal{Q}_{\{t\}})$ 
27:         $\mathcal{Q}_{\{t\}} := \mathcal{Q}_{\{t+1\}}$ 
28:      end for
29:    end for
30:  end for
31: end for

```

A.2 Proof of Proposition 1

For each scenario $s \in S$ the dualization of the subproblem (21)–(32) is as follows.

$$\begin{aligned}
\max \sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} & \left(d_{pls}^{\text{low}} w_{spl}^{(1)} + d_{pls}^{\text{high}} w_{spl}^{(2)} \right) - \sum_{h \in \mathcal{H}} c_{hs}^{\text{hospital}} w_{sh}^{(5)} - \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} c_m^{\text{rescue}} n_{mh}^{\text{hospital}} w_{smh}^{(9)} \\
& - \sum_{i \in \mathcal{I}} c_i^{\text{ACF}} \hat{x}_i w_{si}^{(6)} - \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \tau_j \hat{v}_{ij} w_{sij}^{(7)} - \sum_{i \in \mathcal{I}} \sum_{m \in \mathcal{M}} c_m^{\text{rescue}} \hat{\vartheta}_{mi} w_{sim}^{(10)} - \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}} \left(\sum_{h \in \mathcal{H}} \hat{r}_{ihk} \right) w_{sik}^{(14)} \quad (133)
\end{aligned}$$

Subject to:

Algorithm 5 Model (P2) to Model (P1)

```

1: for  $s \leftarrow 1$  to  $S$  do
2:    $\mathcal{N}_{ip'ps} := 0$ 
3:    $\mathcal{Z}_{ips} := 0$ 
4:   for  $i \leftarrow 1$  to  $|\mathcal{I}|$  do
5:     for  $p \leftarrow 1$  to  $|\mathcal{P}|$  do
6:        $\mathcal{Z}_{ips} := - \left( \sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} q_{plims} + \sum_{\substack{i' \in \mathcal{I} \\ i' \neq i}} \sum_{m \in \mathcal{M}} \eta_{pi'ims} - \sum_{j \in \mathcal{J}_i} \sigma_{pjis} \right)$ 
7:       for  $p' \leftarrow 1$  to  $|\mathcal{P}|$  do
8:         if  $\mathcal{F}(p) \subset \mathcal{F}(p')$  then
9:           for  $j \leftarrow 1$  to  $|\mathcal{J}|$  do
10:            if  $(j \in \mathcal{F}(p'))$  and  $(\mathcal{F}(p') \setminus \{j\} = \mathcal{F}(p))$  then
11:               $\mathcal{N}_{ip'ps} += \sigma_{pjis} - \mathcal{Z}_{ips}$ 
12:            end if
13:          end for
14:        end if
15:         $\sum_{m \in \mathcal{M}} \sum_{h \in \mathcal{H}} u_{p'pihms} + \sum_{m \in \mathcal{M}} \sum_{\substack{i' \in \mathcal{I} \\ i' \neq i}} \eta_{p'pii'ms} = \mathcal{N}_{ip'ps}$ 
16:      end for
17:       $\sigma_{pis} := \sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} q_{plims} + \sum_{m \in \mathcal{M}} \sum_{\substack{i' \in \mathcal{I} \\ i' \neq i}} \eta_{pi'ims} - \sum_{p' \in \mathcal{P}} \mathcal{N}_{ip'ps}$ 
18:    end for
19:  end for
20: end for

```

$$w_{pl}^{(1)} - w_{ip}^{(3)} - w_{ip}^{(4)} - w_i^{(6)} - \pi_{li} w_{im}^{(10)} \leq \pi_{li} \quad p \in \mathcal{P}, l \in \mathcal{L}, i \in \mathcal{I}, m \in \mathcal{M} \quad (134)$$

$$w_{ip}^{(3)} - w_h^{(5)} - \pi_{ih} w_{im}^{(10)} \leq \pi_{ih} \quad p \in \mathcal{P}, i \in \mathcal{I}, h \in \mathcal{H}, m \in \mathcal{M} \quad (135)$$

$$w_{pl}^{(2)} - w_h^{(5)} - \pi_{lh} w_{mh}^{(9)} \leq \pi_{lh} \quad p \in \mathcal{P}, l \in \mathcal{L}, h \in \mathcal{H}, m \in \mathcal{M} \quad (136)$$

$$w_{ip}^{(3)} - w_{i'p}^{(3)} - w_{i'p}^{(4)} - w_{i'}^{(6)} - \pi_{ii'} w_{im}^{(10)} \leq \pi_{ii'} \quad p \in \mathcal{P}, i \in \mathcal{I}, i' \in \mathcal{I}, i \neq i', m \in \mathcal{M} \quad (137)$$

$$w_{pl}^{(1)} \leq \rho_{pl}^{low} \quad p \in \mathcal{P}, l \in \mathcal{L} \quad (138)$$

$$w_{pl}^{(2)} \leq \rho_{pl}^{high} \quad p \in \mathcal{P}, l \in \mathcal{L} \quad (139)$$

$$w_{ip}^{(3)} - \sum_{\substack{p' \in \mathcal{P}: \\ j \notin p' \\ p \setminus \{j\} = p' \\ p \setminus \{j\} \neq \emptyset}} w_{ip'}^{(3)} + w_{ip}^{(4)} - \tau_j w_{ij}^{(7)} - \sum_{k \in \mathcal{K}} g_{jk} w_{ik}^{(14)} \leq 0 \quad p \in \mathcal{P}, i \in \mathcal{I}, j \in \{p \cap \mathcal{J}_i\} \quad (140)$$

$$w_{ip}^{(4)}, w_h^{(5)}, w_i^{(6)}, w_{ij}^{(7)}, w_{mh}^{(9)}, w_{im}^{(10)}, w_{ik}^{(14)} \geq 0 \quad p \in \mathcal{P}, i \in \mathcal{I}, h \in \mathcal{H}, m \in \mathcal{M}, k \in \mathcal{K}, j \in \mathcal{J}_i \quad (141)$$

We also know that we can write model (133) – (141) in the following format.

$$\begin{aligned}
\min \sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} & \left(-d_{pls}^{low} w_{spl}^{(1)} - d_{pls}^{high} w_{spl}^{(2)} \right) + \sum_{h \in \mathcal{H}} c_{hs}^{hospital} w_{sh}^{(5)} + \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} c_m^{rescue} n_{mh}^{hospital} w_{smh}^{(9)} \\
& + \sum_{i \in \mathcal{I}} c_i^{ACF} \hat{x}_i w_{si}^{(6)} + \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \tau_j \hat{v}_{ij} w_{sij}^{(7)} + \sum_{i \in \mathcal{I}} \sum_{m \in \mathcal{M}} c_m^{rescue} \hat{v}_{mi} w_{sim}^{(10)} + \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}} \left(\sum_{h \in \mathcal{H}} \hat{r}_{ihk} \right) w_{sik}^{(14)}
\end{aligned} \quad (142)$$

Subject to:

$$\text{Constraints (134) – (141)} \quad (143)$$

If x^0, ϑ^0, v^0 , and r^0 are core points within the feasible solution set, we can achieve the

Pareto optimal solution using model (144) – (146).

$$\begin{aligned} \min \sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} & \left(-d_{pls}^{low} w_{spl}^{(1)} - d_{pls}^{high} w_{spl}^{(2)} \right) + \sum_{h \in \mathcal{H}} c_{hs}^{hospital} w_{sh}^{(5)} + \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} c_m^{rescue} n_{mh}^{hospital} w_{smh}^{(9)} \\ & + \sum_{i \in \mathcal{I}} c_i^{ACF} x_i^0 w_{si}^{(6)} + \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \tau_j v_{ij}^0 w_{sij}^{(7)} + \sum_{i \in \mathcal{I}} \sum_{m \in \mathcal{M}} c_m^{rescue} \vartheta_{mi}^0 w_{sim}^{(10)} + \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}} \left(\sum_{h \in \mathcal{H}} r_{ihk}^0 \right) w_{sik}^{(14)} \end{aligned} \quad (144)$$

Subject to:

$$\begin{aligned} & \sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \left(-d_{pls}^{low} w_{spl}^{(1)} - d_{pls}^{high} w_{spl}^{(2)} \right) + \sum_{h \in \mathcal{H}} c_{hs}^{hospital} w_{sh}^{(5)} + \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} c_m^{rescue} n_{mh}^{hospital} w_{smh}^{(9)} \\ & + \sum_{i \in \mathcal{I}} c_i^{ACF} \hat{x}_i w_{si}^{(6)} + \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \tau_j \hat{v}_{ij} w_{sij}^{(7)} + \sum_{i \in \mathcal{I}} \sum_{m \in \mathcal{M}} c_m^{rescue} \hat{\vartheta}_{mi} w_{sim}^{(10)} \\ & + \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}} \left(\sum_{h \in \mathcal{H}} \hat{r}_{ihk} \right) w_{sik}^{(14)} = opt(\hat{x}, \hat{\vartheta}, \hat{v}, \hat{r}) \end{aligned} \quad (145)$$

$$\text{Constraints (134) – (141)} \quad (146)$$

For the sake of illustration, consider $\mathcal{D}$ as a vector containing dual variables:

$$\mathcal{D} = \{w_{pl}^{(1)}, w_{pl}^{(2)}, w_{ip}^{(3)}, w_{ip}^{(4)}, w_h^{(5)}, w_i^{(6)}, w_{ij}^{(7)}, w_{mh}^{(9)}, w_{im}^{(10)}, w_{ik}^{(14)}\}.$$

We establish a function denoted as

$$\begin{aligned} F^s(\mathcal{D}, x, \vartheta, v, r) = & \sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \left(-d_{pls}^{low} w_{spl}^{(1)} - d_{pls}^{high} w_{spl}^{(2)} \right) + \sum_{h \in \mathcal{H}} c_{hs}^{hospital} w_{sh}^{(5)} \\ & + \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} c_m^{rescue} n_{mh}^{hospital} w_{smh}^{(9)} + \sum_{i \in \mathcal{I}} c_i^{ACF} x_i w_{si}^{(6)} \\ & + \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \tau_j v_{ij} w_{sij}^{(7)} + \sum_{i \in \mathcal{I}} \sum_{m \in \mathcal{M}} c_m^{rescue} \vartheta_{mi} w_{sim}^{(10)} + \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}} \left(\sum_{h \in \mathcal{H}} r_{ihk} \right) w_{sik}^{(14)}. \end{aligned}$$

We can categorize all feasible solutions of the constraints (134) – (141) into three groups:

- $\mathcal{D}'$: These solutions are neither optimal for model (142) – (143) nor feasible for model (144) – (146).
- $\mathcal{D}^+$: These solutions are optimal for model (142) – (143), but they are not optimal for model (144) – (146).
- $\mathcal{D}^*$: These solutions are optimal for both models (142) – (143) and (144) – (146).

The equality $F(\mathcal{D}^+, \hat{x}, \hat{\vartheta}, \hat{v}, \hat{r}) = F(\mathcal{D}^*, \hat{x}, \hat{\vartheta}, \hat{v}, \hat{r})$ arises from both function values being equivalent to the optimal objective value within model (142) – (143). In the case of $F(\mathcal{D}^+, x^0, \vartheta^0, v^0, r^0)$ and $F(\mathcal{D}^*, x^0, \vartheta^0, v^0, r^0)$, the former surpasses the latter due to $\mathcal{D}^*$ emerging as the optimal solution for model (144) – (146), while $\mathcal{D}^+$ does not achieve optimality, i.e., $F(\mathcal{D}^+, x^0, \vartheta^0, v^0, r^0) > F(\mathcal{D}^*, x^0, \vartheta^0, v^0, r^0)$. Moreover, $F(\mathcal{D}', \hat{x}, \hat{\vartheta}, \hat{v}, \hat{r}) \geq \text{sec}(\hat{x}, \hat{\vartheta}, \hat{v}, \hat{r})$, as $\mathcal{D}'$ does not represent an optimal solution for model (142) – (143).

When considering $\mathcal{D}^+$, it is obvious that $F(\mathcal{D}^+, \hat{x}, \hat{\vartheta}, \hat{v}, \hat{r}) + \mathcal{E}F(\mathcal{D}^+, x^0, \vartheta^0, v^0, r^0) >$

$F(\mathcal{D}^*, \hat{x}, \hat{\vartheta}, \hat{v}, \hat{r}) + \mathcal{E}F(\mathcal{D}^*, x^0, \vartheta^0, v^0, r^0)$ due to the equality $F(\mathcal{D}^+, \hat{x}, \hat{\vartheta}, \hat{v}, \hat{r}) = F(\mathcal{D}^*, \hat{x}, \hat{\vartheta}, \hat{v}, \hat{r})$, along with the fact that $F(\mathcal{D}^+, x^0, \vartheta^0, v^0, r^0) > F(\mathcal{D}^*, x^0, \vartheta^0, v^0, r^0)$ and $\mathcal{E} > 0$. As a result, $\mathcal{D}^+$ does not achieve optimality in the subsequent model, which is equivalent to model (48) – (49).

$$\begin{aligned}
& \min \sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \left(-d_{pls}^{low} w_{spl}^{(1)} - d_{pls}^{high} w_{spl}^{(2)} \right) + \sum_{h \in \mathcal{H}} c_{hs}^{hospital} w_{sh}^{(5)} + \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} c_m^{rescue} n_{mh}^{hospital} w_{smh}^{(9)} \\
& + \sum_{i \in \mathcal{I}} c_i^{ACF} \hat{x}_i w_{si}^{(6)} + \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \tau_j \hat{v}_{ij} w_{sij}^{(7)} + \sum_{i \in \mathcal{I}} \sum_{m \in \mathcal{M}} c_m^{rescue} \hat{\vartheta}_{mi} w_{sim}^{(10)} + \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}} \left(\sum_{h \in \mathcal{H}} \hat{r}_{ihk} \right) w_{sik}^{(14)} \\
& + \mathcal{E} \left(\sum_{p \in \mathcal{P}} \sum_{l \in \mathcal{L}} \left(-d_{pls}^{low} w_{spl}^{(1)} - d_{pls}^{high} w_{spl}^{(2)} \right) + \sum_{h \in \mathcal{H}} c_{hs}^{hospital} w_{sh}^{(5)} + \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} c_m^{rescue} n_{mh}^{hospital} w_{smh}^{(9)} \right. \\
& \left. + \sum_{i \in \mathcal{I}} c_i^{ACF} x_i^0 w_{si}^{(6)} + \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \tau_j v_{ij}^0 w_{sij}^{(7)} + \sum_{i \in \mathcal{I}} \sum_{m \in \mathcal{M}} c_m^{rescue} \vartheta_{mi}^0 w_{sim}^{(10)} + \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}} \left(\sum_{h \in \mathcal{H}} r_{ihk}^0 \right) w_{sik}^{(14)} \right) \quad (147)
\end{aligned}$$

Subject to:

$$\text{Constraints (134) to (141)} \quad (148)$$

When considering $\mathcal{D}'$, it is evident that $\text{sec}(\hat{x}, \hat{\vartheta}, \hat{v}, \hat{r}) > F(\mathcal{D}^*, \hat{x}, \hat{\vartheta}, \hat{v}, \hat{r}) + \mathcal{E}F(\mathcal{D}^*, x^0, \vartheta^0, v^0, r^0)$ due to the equality $F(\mathcal{D}^*, x^0, \vartheta^0, v^0, r^0) = \text{opt}(x^0, \vartheta^0, v^0, r^0)$, $F(\mathcal{D}^*, \hat{x}, \hat{\vartheta}, \hat{v}, \hat{r}) = \text{opt}(\hat{x}, \hat{\vartheta}, \hat{v}, \hat{r})$, and the condition $0 < \mathcal{E} < \frac{\text{sec}(\hat{x}, \hat{\vartheta}, \hat{v}, \hat{r}) - \text{opt}(\hat{x}, \hat{\vartheta}, \hat{v}, \hat{r})}{\text{opt}(x^0, \vartheta^0, v^0, r^0)}$. Since $F(\mathcal{D}', \hat{x}, \hat{\vartheta}, \hat{v}, \hat{r}) + \mathcal{E}F(\mathcal{D}', x^0, \vartheta^0, v^0, r^0) > \text{sec}(\hat{x}, \hat{\vartheta}, \hat{v}, \hat{r})$, it follows that $F(\mathcal{D}', \hat{x}, \hat{\vartheta}, \hat{v}, \hat{r}) + \mathcal{E}F(\mathcal{D}', x^0, \vartheta^0, v^0, r^0) > F(\mathcal{D}^*, \hat{x}, \hat{\vartheta}, \hat{v}, \hat{r}) + \mathcal{E}F(\mathcal{D}^*, x^0, \vartheta^0, v^0, r^0)$. Consequently, $\mathcal{D}'$ also does not represent an optimal solution for model (48) – (49).

In the case where a solution does not meet optimality criteria for model (144) – (146), it consequently fails to attain optimality in model (48) – (49) as well. This implies that a solution achieving optimality in model (48) – (49) is an optimal solution in model (144) – (146), representing a Pareto-optimal solution.

A.3 Flowchart of the BBC_I Algorithm

To provide a clearer understanding of the implementation of our accelerated branch-and-Benders-cut algorithm with various enhancements, we present a flowchart in Figure A.1 that outlines the key steps of our best proposed algorithm, i.e., BBC_I .

A.4 Progressive Hedging Algorithm (PHA)

PHA, introduced by Rockafellar and Wets (1991), is a scenario-based decomposition method for solving two-stage and multi-stage stochastic programs. It has gained significant popularity in disaster response operations, as highlighted in the review by Kamyabniya et al. (2021). Several studies, including Yi et al. (2017), Hu et al. (2019b), and Li et al. (2020),




Figure A.1: Flowchart of the BBC_I algorithm.

have successfully applied PHA to solve large-scale stochastic optimization problems, reinforcing its effectiveness in handling complex decision-making under uncertainty. Given its frequent use in the literature and its relevance to our problem setting, we adopt PHA as a benchmark to evaluate the performance of our proposed approach. It is important to note that PHA operates as a heuristic search method when applied to our problem, in contrast to our exact approach.

PHA achieves this by decomposing the original two-stage stochastic program into independent deterministic subproblems, each corresponding to a single scenario. In this decomposition, the complete stochastic formulation—comprising the integrated objective function (1) and (21) along with the constraints (2)–(9) and (22)–(32)—is rewritten separately for each scenario. The first-stage decision variables are indexed by scenario $s \in S$, i.e., x_{is} , ϑ_{mis} , v_{ijs} , and r_{ihks} , while the second-stage variables remain unchanged. Each scenario-specific subproblem is then solved independently by substituting the uncertain parameters with their realized values. Since first-stage decisions must remain consistent across scenarios, PHA enforces nonanticipativity constraints (149)–(152), ensuring that these decisions are identical across all scenarios:

$$x_{is'} = x_{is''} \quad \forall s', s'' \in S, i \in \mathcal{I} \quad (149)$$

$$v_{ijs'} = v_{ijs''} \quad \forall s', s'' \in S, i \in \mathcal{I}, j \in \mathcal{J}_i \quad (150)$$

$$\vartheta_{mis'} = \vartheta_{mis''} \quad \forall s', s'' \in S, m \in \mathcal{M}, i \in \mathcal{I} \quad (151)$$

$$r_{ihks'} = r_{ihks''} \quad \forall s', s'' \in S, i \in \mathcal{I}, h \in \mathcal{H}, k \in \mathcal{K} \quad (152)$$

However, enforcing these constraints directly would prevent scenario-wise decomposition. Instead, PHA relaxes the nonanticipativity constraints and introduces penalty terms in the objective function. After solving each subproblem independently, PHA iteratively updates first-stage decisions, guiding them toward a single implementable solution. The final first-stage decisions are obtained by computing a weighted average across all scenarios:

$$\tilde{x}_i = \sum_{s \in S} \phi_s x_{is} \quad \forall i \in \mathcal{I} \quad (153)$$

$$\tilde{v}_{ij} = \sum_{s \in S} \phi_s v_{ijs} \quad \forall i \in \mathcal{I}, j \in \mathcal{J}_i \quad (154)$$

$$\tilde{\vartheta}_{mi} = \sum_{s \in S} \phi_s \vartheta_{mis} \quad \forall m \in \mathcal{M}, i \in \mathcal{I} \quad (155)$$

$$\tilde{r}_{ihk} = \sum_{s \in S} \phi_s r_{ihks} \quad \forall i \in \mathcal{I}, h \in \mathcal{H}, k \in \mathcal{K} \quad (156)$$

To drive convergence, PHA introduces a quadratic penalty term in the objective function:

$$\sum_{m \in \mathcal{M}} \lambda_m^k (m - \tilde{m}) + \frac{\rho}{2} (m - \tilde{m})^2. \quad (157)$$

Here, $\Upsilon = \{x_{is}, v_{ijs}, \vartheta_{mis}, r_{ihks}\}$ represents the set of first-stage decision variables subject to nonanticipativity constraints. The term $\tilde{m}$ denotes the implementable solution obtained by averaging across scenarios, ρ is a predefined penalty parameter, and λ_m^k is the Lagrange multiplier, updated in each iteration as:

$$\lambda_m^k = \lambda_m^{k-1} + \rho(m^k - \tilde{m}^k). \quad (158)$$

While the quadratic penalty term enforces nonanticipativity, solving MIP models with quadratic terms is computationally expensive (Hu et al., 2019b). Xing et al. (2023) demonstrated that incorporating a quadratic penalty in PHA can slow down convergence by up to a factor of six compared to a linear penalty formulation, even for relatively small-scale problems. To address this computational burden, we adopt a linearization technique inspired by Long et al. (2012) and Dong et al. (2015). Specifically, we replace the quadratic penalty term in (157) with an equivalent linear term using absolute values, significantly reducing computational complexity while preserving the effectiveness of PHA. For a detailed discussion on the linearization technique, we refer interested readers to Long et al. (2012).

Finally, after careful evaluation, we set the penalty parameter ρ to 0.1 in our implementation. The algorithm terminates when either the first-stage variables converge within a predefined tolerance ($\epsilon = 0.01$) or the maximum computational time (24 hours) is reached. This ensures consistency with the stopping criteria used for our B&BC algorithms.

A.5 Performance Analysis of Benders Decomposition Enhancements

In this section, we analyze the individual contributions of various strategies used in combination to accelerate optimization performance. By understanding the role of each strategy, we can categorize them into distinct groups and combine enhancements from different categories to improve efficiency.

One key strategy is the use of *Pareto-optimal cuts*, which focuses on improving the quality of each cut. Similarly, techniques such as *BDD (Benders Dual Decomposition)* and *LDD (Lagrangian Dual Decomposition)* also aim to generate stronger, higher-quality cuts. Combining Pareto-optimal cuts and BDD targets enhancing the cuts added to the master problem (MP), making the solution process more efficient. In contrast, *partial Benders decomposition* emphasizes cut diversity by incorporating *LBF (lower bounding functional valid inequalities)* in addition to optimality cuts. By doing so, LBF shifts the focus toward improving the overall lower bound, rather than solely concentrating on cut strength. Another important strategy is the *warm-up process*, which accelerates the initial stages of the solution process by rapidly generating a set of cuts. Without the warm-up phase, the initial steps would involve little information about the SPs, resulting in a blind search of the MP's feasible domain. The warm-up process, by solving the linear relaxation

of the MP, provides a more informed starting point for the optimization. To summarize, we classify these strategies into three main categories:

- **Strengthening the cuts** (e.g., Pareto-optimal cuts, BDD, LDD)
- **Diversifying the cuts** (e.g., LBF valid inequalities)
- **Improving the quality of the MP** (e.g., warm-up process)

We evaluate different combinations of strategies from these categories. In Table A.5.1, we present the problem setup, including the number of Alternative Care Facilities (ACFs) and demand locations. The table compares three combinations of strategies:

- *Warm-up, LDD, and LBF*
- *Warm-up, BDD, and LBF*
- *Warm-up, Pareto-optimal cuts, and LBF*

For each algorithm, we provide metrics such as the number of evaluated nodes, solution time, and the optimality gap. Our analysis shows that both algorithms— $BBC_{(WarmUp+BDD+LBF)}$ and $BBC_{(WarmUp+PO+LBF)}$ —performed similarly in terms of optimality gap, with both outperforming the algorithm that includes LDD. However, $BBC_{(WarmUp+PO+LBF)}$ achieved faster solution times than $BBC_{(WarmUp+BDD+LBF)}$, likely due to the use of linear SPs instead of mixed-integer ones.

Table A.5.1: Performance Comparison of Different Enhancement Combinations in the Algorithm.

<i>Problem Size</i>		$BBC_{(WarmUp+LDD+LBF)}$			$BBC_{(WarmUp+BDD+LBF)}$			$BBC_{(WarmUp+PO+LBF)}$		
# ACF	# De.Loc.	<i>Nodes No.</i>	<i>Time (sec)</i>	<i>Gap (%)</i>	<i>Nodes No.</i>	<i>Time (sec)</i>	<i>Gap (%)</i>	<i>Nodes No.</i>	<i>Time (sec)</i>	<i>Gap (%)</i>
10	10	6,974	7,147	0.00	4,570	1,535	0.00	5,072	1,775	0.00
	15	3,467	7,281	0.00	5,302	2,570	0.00	6,815	2,749	0.00
	20	10,422	10,206	0.00	6,272	4,785	0.00	4,579	3,493	0.00
	25	4,655	13,400	0.00	6,292	5,415	0.00	797	3,754	0.00
	30	2,111	15,907	0.00	16,040	6,785	0.00	2,057	4,866	0.00
15	10	9,191	23,562	0.00	8,353	5,080	0.00	3,922	4,451	0.00
	15	4,257	20,994	0.00	8,235	4,927	0.00	2,999	7,116	0.00
	20	5,956	38,658	0.00	10,847	7,695	0.00	3,295	5,442	0.00
	25	6,469	47,246	0.00	3,704	6,888	0.00	3,042	8,784	0.00
	30	7,485	47,985	0.00	9,468	12,254	0.00	15,529	14,335	0.00
20	10	6,611	50,970	0.00	8,095	18,033	0.00	12,284	11,257	0.00
	15	2,713	45,820	0.00	2,421	5,874	0.00	3,355	3,745	0.00
	20	8,194	52,630	0.00	9,684	13,134	0.00	6,234	13,233	0.00
	25	11,638	60,283	0.28	18,116	31,036	1.21	4,538	16,099	0.00
	30	20,788	65,284	1.62	15,700	41,793	1.80	4,280	21,340	0.00

25	10	4,384	48,712	0.00	3,413	13,028	0.00	4,179	7,597	0.00
	15	3,993	49,962	0.00	6,959	13,093	0.00	3,660	8,319	0.00
	20	11,694	53,187	0.00	6,753	15,672	0.00	3,261	12,422	0.00
	25	6,708	66,920	1.60	7,569	32,113	0.00	7,515	23,200	0.00
	30	9,558	74,476	3.27	10,849	66,497	2.63	11,297	38,854	0.00
30	10	1,743	52,610	0.00	5,596	20,728	0.00	925	5,936	0.00
	15	6,396	62,358	0.00	7,981	37,651	0.00	4,922	22,669	0.00
	20	5,870	66,098	1.16	12,093	50,627	0.73	11,231	23,049	0.00
	25	9,974	86,400	8.77	12,022	83,913	3.26	22,388	77,403	3.81
	30	8,935	86,400	5.80	15,136	81,232	2.68	14,488	67,580	2.73
Average		7,207	46,180	0.90	8,859	23,294	0.49	6,507	16,379	0.26

A.6 L-Shaped Algorithm Performance

In this section, we assess the performance of the L-shaped algorithm with and without the inclusion of LBF valid inequalities. It is important to highlight that both versions of the algorithm incorporate the generation of PO cuts and the warm-up process. The results for the CRP solved using the L-shaped algorithm are presented in Table A.2. Each row in this table represents the average results across five instances. As discussed in Section 2.6, the key distinction between the L-shaped and B&BC algorithms lies in how they handle the implementation of the master problem. The L-shaped algorithm, which is a pure cutting-plane method applied to our stochastic mathematical model, solves the master problem optimally at each iteration, while the B&BC algorithm requires solving the master problem only once.

Table A.2 presents the instance information in the “Problem Size” column, followed by the performance results for the L-shaped algorithm without and with the application of the LBF, listed under “L-shaped Algorithm without LBF” and “L-shaped Algorithm with LBF,” respectively. The table also includes additional columns such as “Iter.,” “Time (sec),” “LB,” “UB,” and “Gap (%)” presenting the number of iterations, solution time, lower bound, upper bound, and optimal gap in percentage, respectively, for each algorithm. The “# ACF” and “# De.Loc.” columns provide the count of potential locations for the ACFs and the number of demand locations, respectively.

While the inclusion of LBF valid inequalities within the L-shaped algorithm led to a significant reduction in the optimality gap and the required number of iterations, from approximately 722% to nearly 478% and from 170 to 37, respectively, the remaining gap still falls outside the acceptable range. Consequently, we chose to proceed with the B&BC algorithm for the subsequent sections of this paper.

Table A.2: Computational Results of the L-Shaped Algorithm.

<i>Problem Size</i>		<i>L-shaped Algorithm without LBF</i>					<i>L-shaped Algorithm with LBF</i>				
# ACF	# De.Loc.	<i>Iter.</i>	<i>Time</i> (sec)	<i>LB</i>	<i>UB</i>	<i>Gap</i> (%)	<i>Iter.</i>	<i>Time</i> (sec)	<i>LB</i>	<i>UB</i>	<i>Gap</i> (%)
10	10	118	31,651	70,937	71,408	0.67	24	6,501	70,955	71,352	0.57
	15	165	52,910	83,894	85,007	1.32	22	20,268	84,657	84,914	0.30
	20	197	61,351	109,528	116,237	6.50	44	48,111	111,973	112,658	0.61
	25	230	82,553	125,755	156,037	24.43	33	18,783	133,927	134,644	0.53
	30	236	86,400	131,754	189,684	45.47	51	30,581	152,403	153,188	0.51
15	10	163	66,833	67,844	73,866	8.84	45	33,884	69,680	70,468	1.07
	15	189	76,518	79,395	115,183	48.33	33	27,985	86,138	86,644	0.59
	20	194	86,400	87,910	206,204	143.06	41	45,071	105,777	110,158	4.12
	25	196	86,400	99,072	203,782	107.13	60	74,365	127,714	135,059	5.67
	30	188	86,400	113,773	372,035	229.89	54	70,857	147,322	173,758	17.86
20	10	171	83,708	81,175	190,021	137.21	51	64,163	88,361	94,837	6.66
	15	190	86,400	73,362	96,632	31.06	26	37,361	81,463	81,907	0.54
	20	167	86,400	101,853	392,503	289.21	41	76,871	132,324	843,526	530.06
	25	172	86,400	92,637	887,438	868.07	48	83,757	127,519	253,845	96.10
	30	169	86,400	134,911	1,463,184	976.90	52	86,400	184,122	1,811,799	896.17
25	10	157	81,332	57,259	102,664	78.46	19	41,402	61,679	64,004	3.30
	15	150	86,400	65,684	133,010	106.01	24	57,574	81,535	81,916	0.47
	20	159	86,400	75,021	392,523	426.89	35	75,965	101,121	1,537,159	1415.06
	25	157	86,400	90,058	811,981	800.93	34	83,613	124,002	2,295,585	1733.92
	30	139	86,400	105,158	1,632,509	1441.97	38	86,400	142,597	2,570,940	1723.28
30	10	146	86,400	56,300	92,383	64.40	26	48,779	64,778	103,855	60.10
	15	146	86,400	60,945	285,785	371.37	28	77,372	80,299	800,762	921.07
	20	153	86,400	74,639	939,314	1178.10	32	83,797	99,874	942,987	827.60
	25	139	86,400	87,624	4,629,815	5216.96	33	86,400	120,730	1,867,882	1453.81
	30	152	86,400	114,394	6,097,654	5460.21	35	86,400	154,719	3,613,042	2270.51
Average		170	80,226	89,635	789,474	722.54	37	58,106	109,427	723,876	478.82

Appendix B

Appendices for Paper 2

B.1 Hyperplane Formulation in the Backward Pass

This section provides the formulation of the first hyperplane added in the backward pass. After solving the dual of the subproblems at stage $T + 1$, the first hyperplane is introduced to $\mathcal{V}_T$, incorporating the probability-weighted sum of the dual variables associated with constraints (74)–(91). The complete formulation of this hyperplane is given by:

$$F_{T+1} \geq$$

$$\begin{aligned} & \sum_{\omega^T \in \Omega_T} p(\omega^T) \left(\sum_{j \in \mathcal{J}} \sum_{l \in \mathcal{L}} \sum_{c \in \mathcal{C}} w_{jcl}^{(2)\alpha}(\omega^{[T]}) \left(\mu_{(T-1)jcl}(\omega^{[T-1]}) + d_{Tjcl}(\omega^{[T]}) \right) + \sum_{j \in \mathcal{J} \setminus \{1\}} \sum_{c \in \mathcal{C}} \sum_{u \in \mathcal{U}} w_{jcu}^{(3)\alpha}(\omega^{[T]}) \left(\zeta_{(T-1)jcu}(\omega^{[T-1]}) \right) \right. \\ & \left. + \sum_{j \in \{1\}} \sum_{c \in \mathcal{C}} \sum_{h \in \mathcal{H}} w_{jch}^{(4)\alpha}(\omega^{[T]}) \left(\zeta_{(T-1)jch}(\omega^{[T-1]}) \right) + \sum_{i \in \mathcal{I}} w_i^{(5)\alpha}(\omega^{[T]}) \left(\sum_{j \in \mathcal{J}} \sum_{c \in \mathcal{C}} \zeta_{(T-1)jci}(\omega^{[T-1]}) - c_i^{ACF} \hat{x}_i \right) \right. \\ & \left. + \sum_{h \in \mathcal{H}} w_h^{(6)\alpha}(\omega^{[T]}) \left(\sum_{j \in \mathcal{J}} \sum_{c \in \mathcal{C}} \zeta_{(T-1)jch}(\omega^{[T-1]}) - c_h^{Hos.}(\omega^{[T]}) \right) + \sum_{c' \in \mathcal{C}} \sum_{r \in \mathcal{R}} \sum_{h \in \mathcal{H}} w_{c'rh}^{(7)\alpha}(\omega^{[T]}) \left(- \left(\eta_{(T-1)c'(r-1)h}(\omega^{[T-1]}) \right) \right. \\ & \left. + \sum_{\substack{h' \in \mathcal{H} \\ h' \neq h}} (b_{Tc'rh'h}''(\omega^{[T]}) - b_{Tc'rhh'}''(\omega^{[T]}) - \sum_{i \in \mathcal{I}} b_{Tc'rhi}(\omega^{[T]}) \mathbf{1}_{(r \neq 1)} - \hat{\varrho}_{(T-e)c'h} \mathbf{1}_{(r \geq e)}) \right) + \sum_{c' \in \mathcal{C}} \sum_{r \in \mathcal{R}} \sum_{i \in \mathcal{I}} w_{c'ri}^{(8)\alpha}(\omega^{[T]}) \\ & \left(- \left(\eta_{(T-1)c'(r-1)i}(\omega^{[T-1]}) + \sum_{h \in \mathcal{H}} b_{Tc'rhi}(\omega^{[T]}) + \sum_{\substack{i' \in \mathcal{I} \\ i' \neq i}} (b_{Tc'ri'i'}(\omega^{[T]}) - b_{Tc'rri'i'}(\omega^{[T]})) \mathbf{1}_{(r \neq 1)} \right) + \sum_{h \in \mathcal{H}} \sum_{m \in \mathcal{M}} w_{hm}^{(10)\alpha}(\omega^{[T]}) \right. \\ & \left(- c_m^{rescue} n_{mh}^{Hos.} \right) + \sum_{i \in \mathcal{I}} \sum_{m \in \mathcal{M}} w_{im}^{(11)\alpha}(\omega^{[T]}) \left(- c_m^{rescue} \hat{\vartheta}_{mi} \right) \left. \right) + \sum_{i \in \mathcal{I}} w_i^{(14)\alpha}(\omega^{[T]}) \left(- c^{Aph.} y_{Ti}(\omega^{[T]}) \right) + \sum_{h \in \mathcal{H}} w_h^{(15)\alpha}(\omega^{[T]}) \\ & \left(- c^{Aph.} n_h^{Aph.} \right) + \sum_{c \in \mathcal{C}} \sum_{u \in \mathcal{U}} w_{cu}^{(16)\alpha}(\omega^{[T]}) \left(- \tau \theta_{Tcu}^{Aph.}(\omega^{[T]}) \right) + \sum_{c \in \mathcal{C}} \sum_{h \in \mathcal{H}} w_{ch}^{(17)\alpha}(\omega^{[T]}) \left(- \theta_{Tch}^{Whole}(\omega^{[T]}) \right) \end{aligned} \quad (159)$$

Here, $p(\omega^T)$ represents the probability of scenario ω^T . The variables $w^{(k)\alpha}(\omega^{[T]})$ are the optimal dual variables for scenario $\omega^{[T]}$ at iteration α . For the following indices:

$$k \in \{2, 3, 4, 5, 6, 10, 11, 14, 15, 16, 17\},$$

the variables correspond to constraints (74)–(84), while for $k \in \{7, 8\}$, they are associated with constraints (90)–(91). In these constraints, t is replaced by $T + 1$ to adapt to the final stage formulation. If $e = 1$, $\hat{\varrho}_{(T-e)c'h}$ belongs to the stage $T - 1$ subproblem and does not appear in the hyperplane (159). For $e > 1$, $\hat{\varrho}_{(T-e)c'h}$ is fixed at its previous stage value.

B.2 Illustration of Scenario Trees

This section provides illustrations of a stagewise independent scenario tree and its corresponding recombining tree. Figures B.2 and B.3 visually represent the structure of the discretized stochastic process used in the SDDP algorithm, as described in Section 3.4 of the main text.




Figure B.2: Symmetric Scenario Tree




Figure B.3: Recombining Tree

B.3 Illustration of Single and Multiple Cuts Strategies

Figures B.4 and B.5 illustrate the single cut and multiple cuts strategies, respectively. In Figure B.4, the single cut aggregates the information from all scenarios at stage $t + 1$ into one cut applied to stage t . In Figure B.5, the multiple cuts strategy computes the cost-to-go function $F_{t+1}(\omega^t)$ separately for each scenario ω^t at stage t , with each scenario-specific cut added to stage t . Grey circles represent the scenario nodes sampled during the forward pass, and dashed arrows indicate the cuts added at each stage during the backward pass in both strategies.




Figure B.4: Single Cut Strategy




Figure B.5: Multiple Cut Strategy

B.4 Formulations for LBFVIs

As described in Section 3.4.5, the LBFVIs are added to enhance the lower bound achieved by the first-stage model. This principle is in line with the “*Scenario Creation*” technique proposed in the partial BD by Crainic et al. (2021). The following constraints are incorporated into each subproblem of stages 1 to T . Here, $\Lambda = \{\nu \mid \nu \in \{t+1, \dots, T\}\}$, and all variables are the same as the previous ones, except they are denoted with a bar to indicate that they are defined for the average scenario path.

$$\begin{aligned}
 F_{t+1} \geq & \sum_{\nu \in \Lambda} \left(\sum_{j \in \mathcal{J}} \left(\sum_{c \in \mathcal{C}} \left(\sum_{l \in \mathcal{L}} \sum_{u \in \mathcal{U}} \sum_{m \in \mathcal{M}} \pi_{lu} \bar{q}_{\nu jclm} + \sum_{u \in \mathcal{U}} P_j \bar{\zeta}_{\nu jc u} + \sum_{l \in \mathcal{L}} \rho_{jl} \bar{\mu}_{\nu jcl} \right. \right. \right. \\
 & \left. \left. + \sum_{r \in \mathcal{R}} \sum_{u \in \mathcal{U}} \Delta_u \bar{\eta}_{\nu cr u} \right) \right) + \sum_{u \in \mathcal{U}} \left(\psi_u \bar{\zeta}_{\nu u} + \sum_{\substack{j \in \mathcal{J}: \\ j \in \mathcal{J}_u}} \sum_{c' \in \mathcal{C}} \sum_{c \in \mathcal{C}} \sum_{r \in \mathcal{R}} W_{c'c} \bar{v}_{\nu jc' cru} \right) \\
 & + \sum_{r \in \mathcal{R}} \left(\sum_{h \in \mathcal{H}} \sum_{i \in \mathcal{I}} \pi_{hi} \bar{b}_{\nu cr hi} + \sum_{h \in \mathcal{H}} \sum_{\substack{h' \in \mathcal{H}: \\ h' \neq h}} \pi_{hh'} \bar{b}'_{\nu cr hh'} + \sum_{i \in \mathcal{I}} \sum_{\substack{i' \in \mathcal{I}: \\ i' \neq i}} \pi_{ii'} \bar{b}'_{\nu crii'} \right) \Bigg) \quad (160)
 \end{aligned}$$

$$\sum_{\substack{u \in \mathcal{U}: m \in \mathcal{M} \\ j \in \mathcal{J}_u}} \bar{q}_{\nu jclum} + \bar{\mu}_{\nu jcl} = \bar{d}_{\nu jcl} + \hat{\mu}_{(\nu-1)jcl} \quad j \in \mathcal{J}, l \in \mathcal{L}, c \in \mathcal{C}, \nu \in \Lambda \quad (161)$$

$$\sum_{c' \in \mathcal{G}_c} \sum_{r \in \mathcal{R}_j} \bar{v}_{\nu jc' cru} + \bar{\zeta}_{\nu jcu} = \sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} \bar{q}_{\nu jclum} + \hat{\zeta}_{(\nu-1)jcu} \quad j \in \mathcal{J} \setminus \{1\}, c \in \mathcal{C}, u \in \mathcal{U}, \nu \in \Lambda \quad (162)$$

$$\sum_{c' \in \mathcal{G}_c} \bar{v}_{\nu jc' c1h} + \bar{\zeta}_{\nu ch} = \sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} \bar{q}_{\nu jclhm} + \hat{\zeta}_{(\nu-1)jch} \quad j = 1, c \in \mathcal{C}, h \in \mathcal{H}, \nu \in \Lambda \quad (163)$$

$$\sum_{j \in \mathcal{J}_i} \sum_{c \in \mathcal{C}} \left(\sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} \bar{q}_{\nu jclim} + \hat{\zeta}_{(\nu-1)jci} \right) \leq c_i^{ACF} \hat{x}_i \quad i \in \mathcal{J}, \nu \in \Lambda \quad (164)$$

$$\sum_{j \in \mathcal{J}_h} \sum_{c \in \mathcal{C}} \left(\sum_{l \in \mathcal{L}} \sum_{m \in \mathcal{M}} \bar{q}_{\nu jclhm} + \hat{\zeta}_{(\nu-1)jch} \right) \leq \bar{c}_h^{Hos.} \quad h \in \mathcal{H}, \nu \in \Lambda \quad (165)$$

$$\bar{\varsigma}_{\nu u} = \sum_{c \in \mathcal{C}} \bar{\eta}_{\nu c|R|u} \quad u \in \mathcal{U}, \nu \in \Lambda \quad (166)$$

$$\sum_{j \in \mathcal{J}_i} \sum_{c \in \mathcal{C}} \sum_{l \in \mathcal{L}} \pi_{li} \bar{q}_{\nu jclim} \leq c_m^{rescue} \hat{\vartheta}_{mi} \quad i \in \mathcal{J}, m \in \mathcal{M}, \nu \in \Lambda \quad (167)$$

$$\sum_{c \in \mathcal{C}} \bar{\lambda}_{\nu ci} \leq c^{Aph.} \bar{y}_{\nu i} \quad i \in \mathcal{J}, \nu \in \Lambda \quad (168)$$

$$\sum_{c \in \mathcal{C}} \bar{\lambda}_{\nu ch} \leq c^{Aph.} n_h^{Aph.} \quad h \in \mathcal{H}, \nu \in \Lambda \quad (169)$$

$$\bar{\lambda}_{\nu cu} \leq \tau \bar{\theta}_{\nu cu}^{Aph.} \quad c \in \mathcal{C}, u \in \mathcal{U}, \nu \in \Lambda \quad (170)$$

$$\bar{\varrho}_{\nu ch} \leq \bar{\theta}_{\nu ch}^{Whole} \quad c \in \mathcal{C}, h \in \mathcal{H}, \nu \in \Lambda \quad (171)$$

$$\sum_{i \in \mathcal{J}} \bar{y}_{\nu i} \leq E_t \quad \nu \in \Lambda \quad (172)$$

$$\bar{y}_{\nu i} \leq E_t \hat{x}_i \quad i \in \mathcal{J}, \nu \in \Lambda \quad (173)$$

$$\sum_{i \in \mathcal{J}} \bar{b}_{\nu crhi} + \sum_{\substack{h' \in \mathcal{H}: \\ h' \neq h}} \bar{b}_{\nu crhh'} \leq \hat{\eta}_{(\nu-1)c(r-1)h} \quad c \in \mathcal{C}, r \in \mathcal{R} \setminus \{1\}, h \in \mathcal{H}, \nu \in \Lambda \quad (174)$$

$$\sum_{\substack{i' \in \mathcal{J}: \\ i' \neq i}} \bar{b}'_{\nu crii'} \leq \hat{\eta}_{(\nu-1)c(r-1)i} \quad c \in \mathcal{C}, r \in \mathcal{R} \setminus \{1\}, i \in \mathcal{J}, \nu \in \Lambda \quad (175)$$

$$\bar{\lambda}_{\nu c'h} \mathbf{1}_{(r=1)} + \sum_{\substack{j \in \mathcal{J}_h \\ c \in \mathcal{C} \\ l \in \mathcal{L}}} \pi_{lh} \bar{q}_{\nu jclhm} \leq c_m^{rescue} n_{mh}^{Hos.} \quad h \in \mathcal{H}, m \in \mathcal{M}, \nu \in \Lambda \quad (176)$$

$$\left(\hat{\eta}_{(\nu-1)c'(r-1)h} + \sum_{\substack{h' \in \mathcal{H}: \\ h' \neq h}} (\bar{b}''_{\nu c'rh'h} - \bar{b}''_{\nu c'rh'h'}) - \sum_{i \in \mathcal{J}} \bar{b}_{\nu c' rhi} \right) *$$

$$\mathbf{1}_{(r \neq 1)} + \hat{\varrho}_{(\nu-e)c'h} \mathbf{1}_{(r \geq e)} + A_{c'rh} \mathbf{1}_{(\nu=2)} = \sum_{\substack{j \in \mathcal{J}_h \\ j \in \mathcal{J}_r}} \sum_{\substack{c \in \mathcal{C}: \\ c' \in \mathcal{G}_c}} \beta_j \bar{v}_{\nu jc' crh} + \bar{\eta}_{\nu c' rh} \quad c' \in \mathcal{C}, r \in \mathcal{R}, h \in \mathcal{H}, \nu \in \Lambda$$

$$\bar{\lambda}_{\nu c'i} \mathbf{1}_{(r=1)} + \left(\hat{\eta}_{(\nu-1)c'(r-1)i} + \sum_{\substack{i' \in \mathcal{J}: \\ i' \neq i}} (\bar{b}'_{\nu c' ri'i} - \bar{b}'_{\nu c' ri'i'}) \right) \quad (177)$$

$$+ \sum_{h \in \mathcal{H}} \bar{b}_{\nu c' rhi} \mathbf{1}_{(r \neq 1)} = \sum_{\substack{j \in \mathcal{J}_i \\ j \in \mathcal{J}_r}} \sum_{\substack{c \in \mathcal{C}: \\ c' \in \mathcal{G}_c}} \beta_j \bar{v}_{\nu jc' cri} + \bar{\eta}_{\nu c' ri} \quad c' \in \mathcal{C}, r \in \mathcal{R}, i \in \mathcal{J}, \nu \in \Lambda \quad (178)$$

$$\bar{y}_{\nu i}, \bar{b}_{\nu crhi}, \bar{b}'_{\nu crii'}, \bar{b}''_{\nu crhh'}, \bar{q}_{\nu jclum}, \bar{\mu}_{\nu jcl}, \bar{\zeta}_{\nu jcu}, \quad i, i' \in \mathcal{J}, j \in \mathcal{J}, c, c' \in \mathcal{C}, l \in \mathcal{L}, \quad (179)$$


$$\bar{\lambda}_{\nu cu}, \bar{\varrho}_{\nu ch}, \bar{\eta}_{\nu cru}, \bar{\varsigma}_{\nu u}, \bar{v}_{\nu jc' cru} \geq 0 \quad u \in \mathcal{U}, m \in \mathcal{M}, h, h' \in \mathcal{H}, r \in \mathcal{R}, \nu \in \Lambda$$

Constraint (160) connects the lower bound with the expected cost-to-go, while constraints (161) to (179) are similar to constraints (74) to (92), except that these constraints focus on the average scenario path.


B.5 Comparison of Lower Bounds for SDDP Variants

Figures B.6a and B.6b present the LB values over the first 200 iterations for various SDDP variants applied to the instance with 4 stages, $|\mathcal{J}| = 5$, $|\mathcal{H}| = 5$, and $|\mathcal{L}| = 5$. The instance generation process follows the methodology outlined in Section 3.5.1 of the main text. Figure B.6a provides a comprehensive comparison across all algorithm variants, while Figure B.6b highlights three key variants: **AllEnhancements**, **NoStrongCut**, and **NoWarmUp**, showcasing their differences in convergence behavior.

The results confirm that each enhancement contributes to accelerating the SDDP algorithm’s convergence. Notably, the absence of LBFVIs results in significantly smaller LB values for a substantial number of initial iterations, keeping the lower bounds far from the upper bounds and delaying convergence. This highlights the critical role of LBFVIs in establishing stronger bounds early on, which supports faster overall convergence. Similarly, the absence of the warm-up strategy leads to delayed convergence for the first 50 iterations on this instance, with even more pronounced delays observed in larger instances. Strong cuts also play a vital role, as their removal slows convergence and results in less efficient lower-bound progress. Among the enhancements, the Multi-Cuts strategy stands out for its steep initial slope, demonstrating its effectiveness in improving early-stage convergence. Overall, the absence of any single enhancement negatively impacts convergence, as evident in the **NoEnhancements** variant, where the algorithm struggles to reach comparable LBs within the same number of iterations.



(a) Full comparison of all implemented variants.



(b) Highlighted variants.

Figure B.6: (Color online) LB progression across different SDDP variants.




Figure B.7: Van Network: (a) Full layout, (b) Det, (c) StoNoT, (d) StoT.

B.6 Case Study Details

This study focuses on identifying optimal locations for ACFs in the seismically active Van region, which has a history of earthquakes. On October 23, 2011, a devastating earthquake struck the eastern province of Van, measuring 7.2 on the Richter scale and causing significant loss of life. The earthquake resulted in over 600 fatalities and more than 2,500 injuries (International Federation of Red Cross and Red Crescent Societies, 2012). The 2011 Van earthquake has been the subject of extensive research in areas such as last-mile relief network design (Noyan et al., 2016), humanitarian relief networks (Elçi et al., 2018), shelter location planning (Kılıcı et al., 2015), and assessment planning (Balcık and İhsan Yanıkoğlu, 2020).

The earthquake affected 94 demand neighborhoods (represented as red star polygons in Figure B.7a, where the size indicates demand) and identified 31 potential locations for ACFs (depicted as black tents in Figure B.7a, with size indicating capacity). Noyan et al. (2016) grouped these 94 neighborhoods into clusters of 30 and 60 points using a p -median model to minimize total demand-weighted travel time. Each neighborhood falls into one of three damage intensity levels—"destructive," "damaging," or "strong or very strong"—based on earthquake intensity and proximity to the fault line. These intensity levels correlate with various damage states for neighborhood buildings, such as "no damage," "slight and medium damage," or "heavy damage and collapse." The base demand for each neighborhood was calculated by considering its population, damage intensity, damage state, and the total number of injuries from the 2011 Van earthquake. To generate demand for each scenario, we adjust the base demand by a deviation factor. These factors are sampled from uniform distributions with ranges of $[0.75, 1.25]$, $[0.75, 1.35]$, and $[0.75, 1.45]$ for "destructive," "damaging," and "strong or very strong" intensity levels, respectively.

In the Van district, there are four primary hospitals designated to support patients in the event of a disaster (illustrated as circles with a plus sign in Figure B.7a, where the size represents capacity). However, during an earthquake, the capacity of these hospitals becomes uncertain. To account for this, we assume that each hospital's capacity may be

randomly reduced by up to 21%, based on a uniform distribution from 0% to 21% (Ok-suz and Satoglu, 2020). Additionally, we calculate the distances based on their geographic coordinates. Figures B.7b, B.7c, and B.7d show the optimal number and locations of established ACFs under three models: the deterministic model using the average values of uncertain parameters, the stochastic model with transshipment disabled, and the stochastic model with transshipment enabled, respectively. Detailed explanations of these models are provided in the main text.