RESOURCE ALLOCATION PROBLEMS UNDER UNCERTAINTY IN HUMANITARIAN SUPPLY CHAINS

A Thesis
Presented to
The Academic Faculty

by

Melih Celik

In Partial Fulfillment
of the Requirements for the Degree
Doctor of Philosophy in the
School of Industrial and Systems Engineering

Georgia Institute of Technology
August 2014

RESOURCE ALLOCATION PROBLEMS UNDER UNCERTAINTY IN HUMANITARIAN SUPPLY CHAINS

Approved by:

Professor Ozlem Ergun, Advisor
School of Industrial and Systems Engineering
Georgia Institute of Technology

Professor Pinar Keskinocak, Advisor
School of Industrial and Systems Engineering
Georgia Institute of Technology

Professor Alan Erera
School of Industrial and Systems Engineering
Georgia Institute of Technology

Professor Eric Feron
School of Aerospace Engineering
Georgia Institute of Technology

Professor Julie Swann
School of Industrial and Systems Engineering
Georgia Institute of Technology

Date Approved: June 19, 2014

To Mom, Dad, and Simi

ACKNOWLEDGEMENTS

This thesis work could not have been complete without the help and support of many people. It is a great pleasure to express my gratitude to those who have helped me in achieving this goal, and I apologize in advance to the ones whose names are not mentioned in this short space.

My advisors, Dr. Özlem Ergun and Dr. Pınar Keskinocak, have played a significant role not only in supervising this thesis work, but also in shaping up my views on conducting research, teaching, and academic life in general. Their advices and the examples they have set for me will always shed light into my future career. I would also like to acknowledge the contributions of Drs. Julie Swann and Dima Nazzal to various parts of this thesis work. Along with them, the members of my reading committee, Dr. Alan Erera and Dr. Eric Feron have provided useful comments and feedback, which have helped improve the quality of the work significantly and provided guidance for future research directions.

Working at the Center for Health and Humanitarian Logistics over the last five years has allowed me to collaborate with a number of colleagues. In particular, the experience I gained from working with Álvaro Lorca, Ben Johnson, and Kael Stilp has been very valuable. Mallory Soldner has provided significant input for the fourth chapter of the thesis. I am also indebted to the faculty and staff of the Georgia Tech School of Industrial and Systems Engineering for providing a high quality research environment and the opportunity to interact with a diverse group of researchers. I am especially privileged to have met Drs. Dave Goldsman and Craig Tovey, to have been able to share their wisdom and experience, and to have gotten to know them on a personal level.

My decision to pursue a Ph.D. and an academic career was shaped during my undergraduate and master's level studies at the Industrial Engineering Department of the Middle East Technical University. It gives me great pleasure and excitement to be joining their faculty at the end of my doctoral work. I can never pay back the encouragement and mentoring that my master's thesis advisor, Dr. Haldun Süral, has provided ever since my undergraduate sophomore year. Any verbal description of his impact on my academic life would be an understatement.

I am also grateful to the fellow Ph.D. students at Georgia Tech ISyE, where I have been able to meet a diverse group of good friends from all over the world and extend my friendship with the ones I had known from before. While mentioning all of their names will not be possible here, I would particularly like to thank Yao-Hsuan Chen, who has been my harshest critic, biggest source of encouragement, and the most patient and understanding listener at the same time; Oğuzhan Özlü and Evren Gül, who have always been ready to share even completely random things about life; and Can Özüretmen, with whom I have endured long but mostly enjoyable working hours during the three years we shared an office together.

Last but not the least, I would like to express my gratitude to the members of my beloved family. My parents Meral and Salih Çelik have been able to endure the distance we have been separated with over the last five years, and have never ceased to make me feel their love and support along the way. My brother Semih Çelik, who is on the verge of obtaining his own doctoral degree, has been my best friend for almost thirty years. Every single achievement I have gained in life is completely devoted to them.

TABLE OF CONTENTS

LIST OF TABLES

LIST OF FIGURES

SUMMARY

In this thesis, we consider three applications of humanitarian, emergency response and public health supply chains, namely (i) debris management, (ii) allocation of perishable commodities, and (iii) specialized nutritious foods supply chain design, where the main decisions relate to the allocation of scarce resources and there is inherent uncertainty in the environment. Our particular focus is on incorporating stochasticity into the models that address these decisions and making use of information updates as they become available.

Debris management is one of the most time consuming and complicated activities among post-disaster operations. Debris clearance is aimed at pushing the debris to the sides of the roads so that relief distribution and search-and-rescue operations can be maintained in a timely manner. In the first part of the thesis, we define the stochastic debris clearance problem (SDCP), which captures post-disaster situations where there is limited information on the debris amounts on the roads and the information is updated as clearance activities proceed. We propose a partially observable Markov decision process (POMDP) model to solve SDCP to optimality. Due to the extensive computational requirements of the POMDP model, we propose a heuristic based on a continuous time approximation, and further reduce the computational burden by applying a limited look-ahead on the search tree and using heuristic pruning. The performance of these approaches is tested on randomly generated instances that reflect various geographical and information settings and instances based on a real world earthquake scenario.

Motivated by a public health supply chain problem, we consider the allocation of a perishable item in a decentralized supply chain in the second part of the thesis. The

supply chain consists of two levels: a capacitated warehouse upstream and multiple distribution locations downstream. Each distribution location serves multiple demand types with varying unit benefits accrued from satisfying the demand, and applies a certain usage policy to satisfy the uncertain demand. In particular, we focus on prioritization of higher-class demand and first-come, first-served (FCFS) usage policies. Under the assumption of no reallocation, we develop myopic algorithms that find the optimal allocation that maximizes the expected net benefit for the system. When demand arrivals follow a Poisson process, we establish a worst-case bound on the relative performance of optimal allocation under FCFS compared to prioritization. We also show that centralization, i.e., serving a certain demand class or all demand from a single location, can guarantee the improvement of expected net benefit over a decentralized system only if demand is prioritized and completely served from a single location. We compare the relative performances of the two usage policies, assess the effects of centralization, and compare our approaches to simpler allocation schemes using computational experiments from a dataset based on the 2009 H1N1 vaccination campaign in the state of Georgia.

Specialized nutritious foods (SNF) are aimed at the prevention and treatment of undernutrition, which leads to irreversible and long-term consequences such as hampering of body growth and reduction in productivity. Certain aspects of SNF, such as high product value, low shelf life, geographical disparity of supply and demand, limited visibility into the supply chain, and the need for refrigeration pose important challenges for efficient supply chain management. In the last part of this thesis, we present a case study on the SNF supply chain of a large scale humanitarian organization, where the aim is to determine the network structure and optimal allocation of these commodities to the distribution locations. We develop mixed integer programming models for determining the facility locations, incorporation of information updates, and equity of allocations. Using computational experiments based

on data from the organization, we determine aspects of the system that have the most significant impact on the selection of network structure, and show the impact of incorporating uncertainty and information updates to our models.

CHAPTER I

INTRODUCTION

Over the recent years, natural and man-made disasters have continued to impact the humanity. Recent examples include Typhoon Haiyan (2013), which affected the Philippines, the 2011 cascading disasters (earthquake, followed by tsunami and nuclear leakage) in Japan, and the 2010 earthquake in Haiti. The upward trends in urbanization and population density have increased the devastating effects of these disasters on human life and global economy. Compared to the previous decade, the annual average number of affected people and economic damage between 2001 and 2010 increased by more than 25% and 100%, respectively [37]. In addition to disasters, long-term humanitarian issues such as poverty, conflicts, pandemics, child mortality, maternal health and food security pose significant threats on human well-being. These issues have sustained effects and require significant time to be resolved. The ongoing conflict in Syria, and high HIV prevalence and child mortality rates in Sub-Saharan Africa provide recent examples to long-term issues. Among these, hunger and malnutrition are the greatest risks to global health, even greater than AIDS, malaria, and tuberculosis combined [171].

Humanitarian supply chain management includes activities related to the flow and storage of goods and information from the point of origin to the point of consumption to satisfy the requirements of beneficiaries [33]. These activities take place while mitigating, preparing for, responding to, and recovering from humanitarian challenges. The effectiveness of humanitarian supply chain management has substantial impacts on responding to and recovering from disasters and long-term issues; it has been observed that around 80% of humanitarian relief efforts relate to supply chain and

logistics activities [159]. The theoretical aspects and practical impacts of humanitarian supply chain management problems have led to an increasing number of studies in the literature over the recent years [3, 33].

Compared to the management of for-profit supply chains, humanitarian environments present a number of similar and different challenges. The existence of multiple stakeholders, such as governments, military organizations, non-governmental organizations, commercial institutions and donors results in multiple and conflicting objectives, decentralized decision making, and difficulty of coordination and collaboration. The relief items and services are generally funded by donors or originate as in-kind donations. The uncertainty on when and how much these funds or donations will be received creates an important challenge on material flow planning. Challenges with material flow may be exacerbated by disruption of facilities and road networks. Particularly following disasters, a quick and dynamic response is essential, and some relief items might be highly perishable, allowing a very short time window for satisfying beneficiary needs. On the demand side, due to the combination of event type, local conditions, local preparedness, insufficiency/inconsistency of data collection systems, demand estimation is difficult. Altogether, these factors generally limit the applicability of methods for seemingly similar problems in traditional for-profit supply chains, requiring novel modeling and solution approaches.

Uncertainty in humanitarian supply chains can arise in a number of contexts. Aside from the timing and quantities of funding, supply, and demand, the characteristics of the disaster or long-term issue are often uncertain. For example, the epicenter and magnitude of an earthquake, category and path of a hurricane, or the extent of the spread of an infectious disease are generally uncertain until the outbreak or the aftermath of these events. These characteristics determine not only the needs of the beneficiaries, but also the network structure and resource requirements for pre-

and post-disaster operations such as evacuation or debris removal. While the operations proceed, information updates need to be incorporated into the decision making mechanism as well.

In this thesis, we consider resource allocation problems in humanitarian supply chains and incorporate the underlying uncertainty in the respective systems into the models and solution approaches we develop. Our contributions include the following: (1) We introduce and formally define three resource allocation problems faced in humanitarian supply chains. (2) We develop novel models and solution approaches that consider the uncertainty in the system and devise methods for incorporating the information updates as events unfold. (3) Using structural analysis of the problems and computational experiments, we assess the value of considering the uncertainty and information updates by comparing our approaches to those that ignore these two aspects. (4) By means of computational experiments based on real data, we quantify the improvements resulting from our proposed approaches by comparison to those that mimic the methods used in practice.

The remainder of the thesis is organized as follows. In Chapter 2, we address decisions regarding the clearance of post-disaster debris from the roads so that disaster response activities can proceed. Chapter 3 focuses on the allocation of a perishable commodity from an allocation center to multiple distribution locations that serve multiple demand classes. Lastly, Chapter 4 presents a case study on the specialized nutritious foods supply chain of a large humanitarian organization, and the thesis is concluded in Chapter 5.

In the remainder of this chapter, we briefly summarize the three problems we consider, our solution approaches, and the main contributions of each chapter.

1.1 Stochastic Post-Disaster Debris Clearance

Large-scale disasters often generate substantial amounts of debris that equal years of regular solid waste that would otherwise be generated by the community. Timely removal of debris from the disaster area is important, since debris blocks roads and hampers disaster response activities such as search-and-rescue or relief transportation. Furthermore, in the long term, debris poses threats to human health and environment. Management of post-disaster debris is costly, accounting for an average of 27% of all disaster management costs [67]. *Debris clearance* refers to opening debris-blocked roads by pushing debris to road sides using specialized equipment, so that connectivity between supply and demand locations for relief commodities or services can be re-established in a timely manner.

In a multi-period planning horizon of operations, the main decision in debris clearance is to determine a sequence of roads to be cleared in each period, subject to estimations of debris amounts on the roads (and the corresponding estimates of clearance capacity needs) and a limited clearance capacity in each period. Estimating debris amounts on the roads is a challenging task, requiring significant amount of geographical and demographic data, which are not available for non-US based disasters. Even for the US, estimates of debris can be made only for aggregate levels (e.g., census tracts or counties) [68]. Debris amounts on the individual roads are observed as clearance activities proceed. Therefore, both uncertainty of these debris amounts and the information updates must be considered while making the clearance decisions.

The contributions in Chapter 2 are as follows:

- We formally define the Stochastic Debris Clearance Problem (SDCP), in which exact debris amount information is known only for roads that are immediately reachable by the clearance teams at any point in time. For the remainder of the debris-blocked roads, the decision maker has a discrete set of initial beliefs on the debris amounts with corresponding probabilities. As clearance proceeds

and new roads become reachable, these beliefs are updated.

- To solve SDCP to optimality, we propose a partially observable Markov decision process (POMDP) model with debris amounts as the state vector, which is only partially observable.

- The POMDP approach can only solve small-sized instances of SDCP due to the substantial size of action and observation spaces. To overcome this issue, we propose a heuristic based on a *continuous-time approximation*, where we relax the problem by allowing updates of clearance decisions upon clearing each edge, as opposed to the beginning of each period. Using simple network, supply and demand structures, we show that this approach provides a good approximation for the original POMDP model and can find solutions faster.

- We improve the running time of the continuous time approximation by using heuristic pruning and limited lookahead on the search tree. Our computational results show that significant time savings are obtained by these two heuristics at the expense of slight decrease in the solution quality.

- For learning problems in stochastic networks, we introduce the concept of *regional updates*. This approach uses a Bayesian information update mechanism to update the probabilities in a given region as actual information (debris amounts) of a subset of the edges (roads) in that region are observed. Using computational experiments, we specify cases under which such an update mechanism is useful.

- Our computational experiments show that proposed approaches significantly outperform those applied in practice. We also determine the settings where ignoring uncertainty in debris amounts is more detrimental in terms of affecting the timeliness of establishing connectivity in the network.

1.2 Perishable Commodity Allocation with Multiple Distribution Locations and Demand Types

Allocation and distribution of health/humanitarian relief commodities are challenging as the commodity supply is often insufficient to meet the demand, returns and reallocations of the commodity may not possible due to cost or perishability (e.g., in supply chains of donated blood or vaccines), and the beneficiaries may have different levels of needs (for example, newborn infants and elderly people might be prioritized for vaccine allocation). Such challenges arise in donated blood distribution supply chains and distribution of vaccines, especially when transport and information infrastructure is weak. Consequently, an efficient resource allocation scheme must take into account multiple classes of beneficiaries and the order they are served.

In Chapter 3, we consider the allocation of a perishable commodity from a capacitated warehouse to multiple distribution locations serving multiple classes of beneficiaries. In satisfying the demand, each distribution location applies a usage policy, which determines the order in which the demand is served. In particular, we focus on two usage policies: (i) prioritization, where a lower-class demand can be served only if there is remaining supply after the demand of higher classes has been served, and (ii) first-come, first-served (FCFS), where beneficiaries are served based on the order of arrival, regardless of class.

Given the unit benefits of satisfying demand from each class, usage policies, probability distributions of arrival for each demand class in each distribution location, and the capacity at the warehouse, the main decision is to allocate the on-hand inventory to the distribution locations to maximize the expected total net benefit. Since reallocations are impossible, the problem can be considered as an extension of the newsvendor problem with two echelons, multiple locations, multiple demand classes, and limited capacity at the upstream level.

Chapter 3 makes the following contributions:

- We consider two cases of usage policies at the distribution locations: (i) prioritization when demand arrivals of each class follow a general probability distribution, and (ii) FCFS with demand arrivals following a Poisson process. For both cases, we develop myopic algorithms that find the optimal allocation. To the best of our knowledge, this is the first study that considers a newsvendor problem with multiple demand classes in a multiple echelon supply chain with a capacitated warehouse upstream. Another contribution of this work is the consideration of a FCFS usage policy in a multiple echelon newsvendor setting.

- With two demand classes and Poisson arrivals, we analytically derive the worst-case relative performance ratio of the optimal allocation under FCFS usage to the optimal allocation under prioritization, in terms of the unit benefits and demand arrival rates.

- We compare serving part or all of the demand classes from a single distribution location to the original setting under Poisson demand arrivals, and obtain the following results:

  - Under FCFS usage, when capacity is sufficient, serving all demand classes from a single distribution location always improves expected benefits. This result extends the findings in [65] and [150] for the newsvendor problem to Poisson demand and multiple demand classes.

  - Suppose that capacity is limited and the demand for each class follows an identical distribution in all distribution locations in the original setting. In this case, serving all demand classes from a single distribution location still improves expected benefits under FCFS usage. This result is analogous to pooling in queues, where pooling identical servers into a single aggregate one improves the waiting times and throughput [164, 165].

  - Under prioritization, serving all demand classes from a single distribution location always improves expected benefits even when capacity is limited.

- We conduct computational experiments based on the 2009 H1N1 vaccination campaign with the aim of assessing (i) the relative performances of the optimal allocations under prioritization and FCFS usage policies to each other and to simple proportional allocation mechanisms, (ii) information asymmetry between assumed and actual usage policy, i.e., assuming prioritized usage to make the allocations, whereas distribution locations apply FCFS usage, (iii) the effect of serving demand from a decreased number of distribution locations facing aggregate demand, and (iv) equity under both usage policies.

1.3 Network Design for the Supply Chain of Specialized Nutritious Foods

Hunger and undernutrition (insufficient food intake) pose the greatest risks to human health globally. Undernutrition leads to irreversible and long-term consequences such as hampering of body growth and reduction in productivity. Currently, 870 million people around the world suffer from undernutrition, with 98 percent living in developing countries [73]. Prevention and treatment of undernutrition requires access to nutrient-dense foods. Over the recent years, specialized nutritious foods (SNF), which refer to a set of blended foods that contain a broad set of macro- and micro-nutrients in dense amounts, have proven to be effective in the prevention and treatment of undernutrition. SNF have higher commodity value and lower shelf life compared to other food aid commodities, and remote locations of suppliers result in longer lead times, increasing the risk of expirations. There are mismatches between supply, which arrives in irregular intervals and large batches due to availability of funding, and demand. Therefore, an efficient allocation of these commodities to distribution locations in a country needs to consider commodity age.

In Chapter 4, we consider the SNF distribution supply chain of a large humanitarian organization. In particular, the organization aims to assess the performance of two different network structures: (i) district warehouse-based, where intermediate district warehouses are located closer to the distribution locations and can hold stocks of the commodities.; (ii) hub-based, where intermediate hubs are located close to the central warehouse and act as crossdocks. We analyze the performance of these two network structures and the corresponding allocation strategies under various conditions.

The contributions of Chapter 4 include the following:

- We quantify the effects of various factors such as network structure, demand uncertainty, incorporation of actual demand observations, and commodity perishability by means of a set of computational experiments based on actual data from the organization.

- We develop optimization models for the allocation of (i) a perishable commodity, (ii) under uncertain demand, (iii) with transportation capacities, and (iv) under limited supply. These models can be solved using commercial software within a few hours.

- In addition to finding a cost-effective allocation scheme, we also consider the effects of maximizing equity as the objective instead of minimizing costs, and develop novel heuristic approaches for this case.

CHAPTER II

THE POST-DISASTER DEBRIS CLEARANCE PROBLEM UNDER INCOMPLETE INFORMATION

In recent years, there has been a significant increase in the number and impact of natural and man-made disasters worldwide [37]. An important issue faced by local and federal authorities is the management of post-disaster debris, which can be defined as any type of waste generated by disasters, including vegetation (e.g., trees, limbs), construction and demolition waste (e.g., concrete, roof material), municipal solid waste (e.g., household garbage), white goods, motor vehicles, dead animals, and hazardous chemicals. Management of post-disaster debris deserves significant attention due to the size, cost, and complicated structure of the activities involved.

Volume of post-disaster debris may sometimes be equivalent to multiple years of solid waste that would normally be generated. For instance, the debris generated by Hurricane Katrina in 2005 is estimated to be more than 100 million cubic yards (CY), which is enough to fill a football field up to 16 miles high [45]. Similarly, Hurricane Andrew in 1992 generated around 43 million CY of debris, the March 2011 cascading disasters in Japan produced more than 35 million CY [88], and the 2010 Haiti earthquake resulted in no less than 20 million CY [168]. Debris also contributes significantly to disaster management costs by accounting for an estimated 27% [67].

In the aftermath of a disaster, debris hampers effective search-and-rescue and relief distribution efforts by blocking roads and preventing access to disaster sites, highlighting the significance of timely debris clearance and removal efforts. Motivated by this, the Federal Emergency Management Agency (FEMA) [67] provides a Public Assistance Debris Management Guide to help public and authorities deal with the

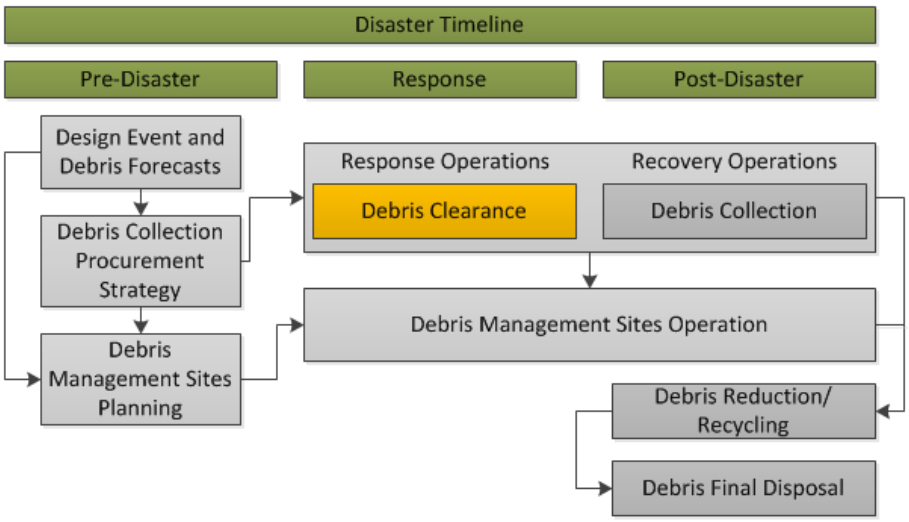

Figure 1: Elements of a debris management plan during a disaster timeline

impact of post-disaster debris. Adapted from [67], Figure 1 gives the elements of a debris management plan during different phases of the disaster. Before the disaster occurs, a set of potential disaster scenarios are considered to forecast the resulting amount and distribution of debris. Equipment procurement and workforce allocation plans are made based on these expectations. Locations and capacities of facilities such as debris processing sites and landfills are also determined in this stage.

In the first few days following the disaster (response phase), the resources are often extremely limited and the urgency of delivering aid to those in need is very high; hence, the main objective for the clearance activities is to unblock the roads by pushing the debris to the side so that search-and-rescue activities as well as relief transportation can proceed. These activities, on which this chapter focuses, are referred to as *debris clearance*. In the recovery phase, *debris collection* activities consist of transportation of the debris from the disaster area to collection sites for further processing or disposal. Lastly, *debris recycling/disposal* consists of processing (e.g., burning, crushing, grinding) debris at predetermined facilities followed by further processing for recycling and/or disposal by transportation to landfills.

FEMA [67] provides guidelines related to post-disaster debris management by outlining the activities and resources as well as assigning responsibilities. However, it takes a descriptive rather than a prescriptive approach. In particular, for the debris clearance stage, it lacks guidance on prioritizing roads in a systematic way in order to complete the clearance activities in a timely manner. The lack of operational and tactical guidance on how to carry out debris management activities motivates our research towards developing mathematical models for decision support.

In the debris clearance stage, road prioritization decisions are generally made by local or federal emergency management authorities. Current practice for debris clearance mainly determines a list (sequence) of roads to clear with the goal of enabling immediate response efforts [67]. More detailed scheduling decisions, such as clearance team assignments and routing of clearance equipment, are made a posteriori based on these priorities.

In determining the clearance schedule, it is important to consider the fact that the resources required to clear the roads depends on the amount of debris on each road segment. Debris amounts on different roads might be hard to estimate. For instance, while Jean-Louis and Hilaire [90] provided an aggregate estimate of 20-25 million CY of debris following the Haiti earthquake, no estimates on individual roads or regions were available. In certain cases, even obtaining aggregate estimates might be a challenge. As an example, aggregate debris estimations were deemed impossible for the 2013 tornadoes in Oklahoma due to the lack of debris estimation models for tornadoes. Given these challenges in estimating the debris amounts, our model assumes that these amounts are stochastic (for a subset of roads). By considering this uncertainty, we define the *stochastic debris clearance problem* (SDCP), in which the main decision is to dynamically create a clearance sequence for debris-covered roads in the network so that demand and supply for relief can be connected in a timely manner. We model SDCP as a multi-period problem where clearance decisions are

made at the beginning of each period (e.g., a day, or a shift of eight or 12 hours) and there is limited clearance capacity (resources) in each period, determined by workforce and equipment constraints. To the best of our knowledge, our work is the first in the literature that provides a systematic approach for prioritizing road clearance considering the uncertainty in the resources required to clear the debris.

The consideration of uncertainty along with the multi-period nature of the problem add significant complexity and make exact methods impractical for real-sized instances. To overcome this, we propose a heuristic that relaxes the multi-period nature of the decision making process. We also develop solution methods that incorporate regional information updates, in which case the clearance resource requirements on the roads that become reachable and/or observable in a region lead to adjustments of those on the remaining roads in the same region. We test the performances of these approaches and compare them to a set of "benchmark" approaches (motivated by the current practice) on structured networks and various instances of a disaster scenario. Our computational results demonstrate the effectiveness of the models and heuristics, and provide insights on how road prioritization decisions are affected by factors such as the structure of the network and the extent of uncertainty. In particular, we show that our approach provides significant improvements over those ignoring uncertainty or assigning road priorities without considering the network structure or debris amounts. We also characterize the situations where employing regional updates improve the solutions.

The remainder of the chapter is organized as follows: §2.1 provides an overview of the literature relevant to the problem and underlines the main contributions of this work. §2.2 defines SDCP formally and introduces the partially observable Markov decision process (POMDP) model to solve it. §2.4 describes the developed heuristics to solve larger-sized problems, whereas §2.5 discusses structural results for the POMDP model and the proposed continuous-time heuristic. §2.6 provides the results of the

computational experiments performed on randomly generated grid and ring networks and instances generated based on an earthquake scenario. We conclude the chapter in §2.7 and discuss potential future research directions.

2.1 Literature Review

Significant contributions have been made by the operations research and management science literature to the field of disaster management. Altay and Green [3] provide a detailed review and classification of these studies and regard debris removal as a major disaster response/recovery problem and an important future research direction. Caunhye et al. [31] focus on preparedness and response stages of disasters, also pointing out debris removal among future research directions. For recent reviews on humanitarian logistics, please see Apte [8] and Çelik et al. [33].

The work in this chapter is related to a number of different streams in the literature including policy-related qualitative studies on debris management, network repair and relief distribution, online and stochastic networks, and partially observable Markov decision processes (POMDPs).

Brown et al. [28] review the post-disaster waste management literature and conclude that the literature is mainly comprised of disaster waste management plans or guidelines and isolated case studies. Examples of such guidelines include FEMA [67] for the United States, Solis et al. [147] for Canada, and United Nations [161] for the global level. There are also case studies regarding debris management decisions in specific disasters. These include Mendonça and Hu [113] and Luther [109] on Hurricane Katrina and Srinivas and Nakegava [148] on the 2004 Indian Ocean Tsunami. Swan [152] and Ekici et al. [63] take a location-independent perspective and provide analysis on issues such as allocation of responsibilities and administrative procedures. As Fetter and Falasca [71] also point out, despite the increasing interest in the policy-related aspects of debris management, very few studies propose decision

support systems and quantitative tools. The work in this chapter tries to bridge this gap.

The debris clearance problem has been modeled by Stilp et al. [149] considering deterministic clearance resource requirements. The authors show the NP-hardness of the problem even under very simple network structures, such as a network with a single supply and a single demand node. They formulate a mixed integer programming model and propose heuristic methods. Aksu and Özdamar [2] present a path-based mixed integer programming formulation for the same problem. The work in this chapter considers uncertainty in debris volumes over the network and the corresponding clearance times. Through structural analysis and computational experiments, we show that policies that ignore uncertainty by assuming expected clearance times and solving a deterministic model can lead to poor solution quality.

A problem relevant to debris clearance is that of restoring a disrupted infrastructure network after a disaster. Perrier et al. [125] consider the problem of routing and scheduling of snow plowing equipment in urban areas where each road segment has a pre-assigned priority level. Similarly, Şahin et al. [139] consider debris clearance in the response stage using predetermined road priorities. In general, the decision of prioritizing roads is not trivial due to possible differences between type, severity, and location of the disaster as well as the requirements of connectivity among communities and critical facilities. Therefore, in our model, instead of taking the priorities as inputs, we include them as decisions. Using computational experiments, we also show that policies that use pre-assigned road priorities may lead to significantly worse results.

A number of studies consider the network repair and relief distribution decisions simultaneously [32, 40, 99, 112, 121, 173]. In these models, it is possible to repair edges that are not physically reachable by the repair team due to damaged edges. This is not the case in SDCP, as debris clearance teams can only reach those roads (edges

in the network) which are connected (i.e., through a sequence of clear roads) to their current location. Furthermore, as opposed to the deterministic repair (clearance) times in these studies, the work in this chapter assumes uncertainty in edge repair times.

In most online/stochastic network problems, edge travel times are known and one or more servers provide service to demand nodes with known demand arrival times; however, some uncertain characteristics of the demand (location, quantity, destination to transport the load, etc.), are revealed only at the time of demand occurrence. Studies on online problems include Jaillet and Wagner [86, 87], Angelelli et al. [6], and Jaillet and Lu [85] on the traveling salesman problem (TSP), Ausiello et al. [11] on the prize-collecting TSP, Anagnostopoulos et al. [5] on the facility location problem, Bent and Van Hentenryck [21] on the vehicle routing problem (VRP), and Ascheuer et al. [10] and Feuerstein and Stougie [72] on the dial-a-ride problem. While how uncertainty is revealed in these problems is similar to SDCP, the source of uncertainty being nodes rather than edges causes significant differences in the decisions made, hence, also in the solution methods. The main decisions in these problems involve where and how long to wait, to accept or reject an incoming demand, etc., as opposed to edge sequencing decisions in SDCP.

An aspect of the stochastic network problems relevant to SDCP is the incorporation of learning as uncertainty is revealed. Learning algorithms applicable to various problems on networks are proposed by Littlestone and Warmuth [106], Bent and Van Hentenryck [22, 23], Kalai and Vempala [92], Powell and Ryzhov [128], Ryzhov and Powell [138], and Van Hentenryck et al. [84]. These approaches are quite general and can be applied to a wide range of stochastic network problems. However, they either involve problems that are encountered by the decision maker multiple times (therefore learning occurs each time a new problem is encountered) or assume that no a priori information exists.

Stochastic network problems in which uncertainty arises from edge lengths (travel times) include robotic mapping and the Canadian traveller problem. In robotic mapping, a robot uses its error-prone sensors to learn about and map the network structure of an unknown environment as quickly and accurately as possible. A detailed review of this problem is given by Thrun [156]. While the learning aspect of the problem bears similarities to SDCP, robotic mapping aims to learn the network as quickly as possible, while SDCP also needs to consider demand satisfaction (by connecting nodes of relief demand and supply).

Various versions of the Canadian traveller problem exist, of which the most relevant to SDCP is studied by Karger and Nikolova [95], where an agent travels from a source node to a destination node through a shortest path. Edge lengths are uncertain and can have different values, each with an associated probability. Uncertainty is revealed when the agent reaches the edge, which is also the case in SDCP. Despite the similarity of uncertainty and how it is revealed, consideration of a single source and sink node as opposed to the case of multiple supply and demand nodes in SDCP, as well as the focus on minimizing the travel distance instead of relief distribution results in significant differences in the modeling, algorithms, and structure of the optimal solutions.

Partially observable Markov decision processes (POMDP) provide a general framework to model sequential decision problems in environments where system state is not completely observable. A review of POMDP models is given by Ross et al. [136], with a focus on online planning approaches, which find a good local policy for the current belief state, usually based on a limited look-ahead, update information as new observations are received, and repeat the process. In contrast, offline POMDP approaches find optimal actions to perform in every belief state, requiring considerable computation time. While our exact POMDP model uses an offline approach, we make use of online approaches in our heuristics for the SCDP.

The contributions of this chapter are three-fold. First, we introduce SDCP, a problem applicable in many real-life post-disaster situations. As in the case of most online and stochastic problems, there is a key trade-off between taking the actions that bring the most immediate benefit and learning the network structure to take more informed actions in the future. In SDCP, the trade-off is between clearing the most beneficial roads in the next period and learning the debris volume throughout the network. Second, to resolve this trade-off and overcome the challenges of uncertainty and the periodic nature of the problem, we propose a heuristic procedure that relaxes the periodicity in the decision making scheme to reduce the number of possible actions. Our approach can also make use of a regional information update scheme, which is motivated by the fact that the disaster is expected to have a similar impact on geographically close areas. In this update scheme, debris information obtained for roads in a certain region leads to updated belief probabilities for the clearance resource requirements in the same region. The proposed solution approaches are not problem-specific and can be applied to online stochastic problems on networks where decisions need to be made at the beginning of each period. Examples of such problems include power systems repair and post-disaster relief distribution with stochastic demands. Lastly, we compare the proposed approaches to their counterparts, which (i) ignore uncertainty in clearance resource requirements and (ii) mimic current debris clearance practice. Through computational experiments, we classify the cases where the difference between the benefits obtained by a stochastic approach is more significant.

2.2 Problem Description and Model

2.2.1 The Stochastic Debris Clearance Problem (SDCP)

The debris clearance problem aims to develop a prioritization scheme (dynamically) for the clearance of debris-blocked roads with the goal of establishing connectivity

between supply and demand points to enable the flow of relief commodities or activities between them. The resources to carry out debris clearance are very limited in the first few days following the disaster, hence efficient and effective management of these resources is important.

We model SDCP on an undirected network $G = (N, E)$. Clearance activities may start from a number of different locations in the network, depending on where clearance resources are initially stored or prepositioned. The planning horizon is divided into discrete time periods, given by the set T. The node set is partitioned into subsets N_S (supply), N_D (demand), and N_T (transshipment). Supply points may consist of critical facilities such as fire stations, hospitals, shelters, distribution centers for relief supplies, etc. Demand points are population centers consisting of people in need of relief commodities or access to services offered by the supply points. At the beginning of each period $t \in T$, the edge set is partitioned into two subsets, B_t and U_t, which denote the sets of blocked and unblocked edges, respectively. The notation we use throughout the remainder of this section is given in Table 1.

At the beginning of each period $t \geq 2$ in T, the parameters rd_{it} and rs_{it} denote the remaining unsatisfied demand of node $i \in N_D$ and remaining supply capacity of node $i \in N_S$, respectively. rd_{i1} and rs_{i1} denote the initial demand and supply at node i, respectively. Once a demand node becomes connected to a supply node, it starts accruing benefits proportional to the amount of demand satisfied in the remainder of the planning horizon. We assume that once resources (e.g., commodities or service) are assigned to satisfy the demand in a certain node, that assignment remains the same for the remainder of the horizon, i.e., those resources cannot be assigned to a different demand node in a future period. Per each unit of demand satisfied, demand node i accrues p_i units of benefit in each period. The objective in SDCP is to maximize the total (expected) benefit during the planning horizon. The benefits being accrued over time motivates connecting the supply nodes to the demand nodes as quickly as

Table 1: Notation

Index sets

T	Periods
N	Node set
N_S	Supply nodes
N_T	Transshipment nodes
N_D	Demand nodes
E	Edge set
U_t	Unblocked edges at beginning of period t
B_t	Blocked edges at beginning of period t
RB_t	Blocked edges that are reachable by a clearance team at beginning of period t
UB_t	Blocked edges that are not reachable by any team at beginning of period t
DL_{ij}	Set of possible clearance resource requirement levels for blocked edge (i, j)

Parameters

p_i	Benefit accrued in node $i \in N_D$ per unit demand satisfied in each period
k_t	Total clearance capacity in period $t \in T$
FC_{ij}	Flow capacity on edge (i, j) if unblocked
W_{ij}^t	Amount of resources required to clear edge $(i, j) \in B_t$ (Random variable for $(i, j) \in UB_t$, known for $(i, j) \in RB_t$)
$\overrightarrow{W_t}$	Vector of resource requirements on each edge in B_t in period $t \in T$
β_{ij}^{dt}	Probability of observing debris level $d \in DL_{ij}$ on edge $(i, j) \in UB_t$ based on belief at beginning of period $t \in T$
$\overrightarrow{Y_t}$	Vector consisting of probabilities of observing each debris level on each edge in UB_t based on belief at beginning of period $tinT$
RC_t	Remaining clearance capacity at the beginning of period $t \in T$ after ongoing clearance from period $t-1$ is completed
rs_{it}	Remaining supply capacity in node $i \in N_S$ at beginning of period $t \in T$
$\overrightarrow{RS_t}$	Vector consisting of remaining supply in each supply node at beginning of period $t \in T$
rd_{it}	Remaining unsatisfied demand in node $i \in N_D$ at beginning of period $t \in T$
$\overrightarrow{RD_t}$	Vector consisting of remaining demand in each demand node at beginning of period $t \in T$
γ^t	Remaining resource requirement for the partially cleared edge (if any) at the end of period $t \in T$

possible. The problem can be equivalently modeled using a penalty for not satisfying the demand in a given period. It is also possible to extend the models to incorporate changes in supply capacity (e.g., due to new commodity arrivals or service capacity increase) and demand (e.g., due to population movements) by introducing additional variables. Throughout this chapter, we assume a single relief commodity or service to be offered. However, all the models and solution methods can be extended to the case of multiple commodities by introducing a commodity index and incorporating it into supply and demand parameters.

There are two important distinctions between SDCP and the network repair and snow plowing problems. First, we assume that clearance units cannot traverse debris-blocked roads, i.e., for a clearance team to reach a debris-blocked edge, there needs to be a cleared path to that edge from the team's current location. Second, we assume that clearance activities take significantly more time than traveling on cleared roads. These assumptions are in line with the current practice, where road prioritization is taken as the main decision in the response stage and routing decisions are secondary and depend on the prioritization decisions.

We incorporate uncertainty into the problem by assuming resource requirements to clear the roads (that is, the debris amounts) to be stochastic for a subset of blocked roads. While the assumption of deterministic clearance times is valid in cases where it is possible to get good debris amount estimations (e.g., smaller hurricanes), estimating debris amounts and corresponding resource requirements require a significant amount of demographic and structural data as well as proper damage forecast models on the disaster area, which may not be easily available. Even when such estimates are available (e.g., [68]), they are in aggregate form (e.g., for a region) and accuracy might be lost when they are disaggregated to road segments. To model uncertainty, let $RB_t \subseteq B_t$ denote the blocked edges that can be reached by a clearance unit using edges in U_t and $UB_t = B_t \setminus RB_t$. We denote by W_{ij}^t the amount of resources required

to clear edge $(i,j) \in B_t$. We assume that W_{ij}^t are known for $(i,j) \in RB_t$ and are random variables following a categorical distribution for $(i,j) \in UB_t$. For $(i,j) \in UB_t$, a discrete belief vector $\beta_{ij}^t = \left\{\beta_{ij}^{lt} : l \in DL_{ij}\right\}$ gives the set of the probabilities under the current belief state for each possible requirement level $l \in DL_{ij}$. The decision maker assumes that resource requirements on blocked edges are independent of each other. Later, we relax this assumption by introducing regional updates, where the dependence between the resource requirements in the same region is handled using an information update mechanism. Due to constraints on workforce and equipment capacity, at most k_t units of resources can be allocated to debris clearance in each period $t \in T$.

Clearance decisions change the sets RB_t and UB_t dynamically, as illustrated on a simple example in Figure 2. Upon the clearance of edges $(0,3)$ and $(3,6)$ in period t, edges $(2,6)$, $(3,7)$, and $(6,8)$ become reachable at the beginning of period $t+1$. We assume that once the clearance decisions for a period are made, they are fixed until the beginning of the next period. Therefore information updates cannot be incorporated before the period ends.

The main decision in SDCP is to determine a sequence of roads to be cleared. Since the clearance resource requirements of a subset of edges are not known in advance, this might result in partial clearance of an edge at the end of a period. In this case, the clearance of this edge has to be continued at the beginning of the next period. This implies that if there is ongoing clearance from period $t-1$ in period t, part or all of k_t will be spent on this ongoing clearance before any new decisions for period t can be made. To incorporate this into the modeling process, RC_t denotes the remaining clearance capacity in period t after the ongoing clearance is completed. Lastly, only an entirely cleared edge provides information about the amount of resources required to clear its adjacent edges.

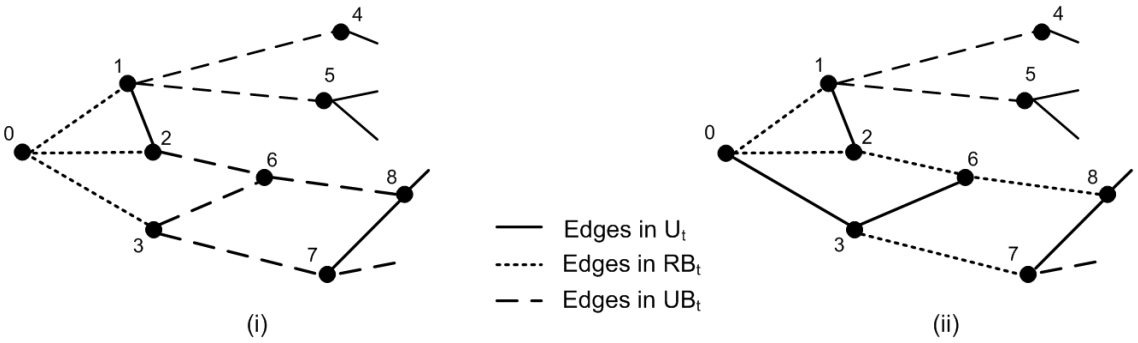

Figure 2: The unblocked, reachable blocked, and unreachable blocked edges (i) before edges $(0, 3)$ and $(0, 6)$ are cleared in period t, (ii) after the edges are cleared in the period t

2.2.2 The POMDP Model

SDCP can be formulated as a stochastic network problem. The system state s_t in period $t \in T$ is denoted by the quartet $s_t = (\overrightarrow{W_t}, RC_t, \overrightarrow{RS_t}, \overrightarrow{RD_t})$, where $\overrightarrow{W_t} = \{W_{ij}^t : (i,j) \in B_t\}$, $\overrightarrow{RS_t} = \{rs_{it} : i \in N_S\}$, and $\overrightarrow{RD_t} = \{rd_{it} : i \in N_D\}$. The size of the state vector is $|B_t|+|N_S|+|N_D|+1$.

Throughout the remainder of this section, we describe the details of the POMDP model to solve SDCP.

The state space S_t in period $t \in T$ is given by:

$$S_t = \Bigg\{ (\overrightarrow{W_t}, RC_t, \overrightarrow{RS_t}, \overrightarrow{RD_t}) : W_{ij}^t \in DL_{ij} \ \forall (i,j) \in B_t;\ 0 \le RC_t \le k_t; \\ 0 \le rs_{it} \le rs_{i1} \ \forall i \in N_S;\ 0 \le rd_{it} \le rd_{i1} \ \forall i \in N_D \Bigg\}. \quad (1)$$

Let $\overrightarrow{Y_t} = \{\beta_{ij}^{dt} : (i,j) \in B_t,\ d \in DL_{ij}\}$ be the probability vector of observing different resource requirement levels on each blocked edge under the current belief. More formally, given the belief probabilities β_{ij}^{dt}, the belief state can be represented as $bs_t = (\overrightarrow{Y_t}, RC_t, \overrightarrow{RS_t}, \overrightarrow{RD_t})$, where $\beta_{ij}^{lt} \ge 0 \ \forall l \in DL_{ij}$ and $\sum_{l \in DL_{ij}} \beta_{ij}^{lt} = 1$ for all (i,j) in B_t, and $|bs_t| = \sum_{ij \in B_t} |DL_{ij}| + |N_S| + |N_D| + 1$.

The set of possible actions $\mathcal{A}_t$ in period t is any potentially feasible sequence of edges in B_t that can be cleared under the current belief state bs_t. Let $\mathcal{E}_t$ be the set of sequences of edges that can be cleared without violating connectivity. Then, the action space consists of all sequences $A \subseteq \mathcal{E}_t$ such that when the last edge, $e'(A)$, in the ordered set is excluded, clearance capacity constraint is strictly satisfied under minimum possible clearance requirement levels. Let $W_{ij}^{min,t}$ be the minimum possible amounts of resources required to clear the debris on edge (i,j) based on the current belief, i.e., $W_{ij}^{min,t} = \min\{l : \beta_{ij}^{lt} > 0\}$. In more formal terms:

$$\mathcal{A}_t = \left\{ A \subseteq \mathcal{E}_t : \sum_{(i,j) \in A \setminus e'(A)} W_{ij}^{min,t} < RC_t \right\}. \tag{2}$$

If $RC_t = 0$, the action space consists only of the edge for which there is ongoing clearance from the previous period.

Given the current belief state bs_t and a feasible action $A \in \mathcal{A}_t$, the observation space consists of: (i) all possible $O \subseteq UB_t$ that can be cleared or become reachable as a result of the action, and (ii) all possible combinations of possible resource requirements on each of these subsets. For instance, if $A = \{(0,3),(3,6)\}$ is a feasible action in Figure 2, and if none of the edges is completely cleared at the end of the current period, $O = \emptyset$. Otherwise, if only $(3,6)$ is completely cleared, $O = \{(3,6),(3,7)\}$. If both edges are completely cleared, $O = \{(2,6),(3,6),(3,7),(6,8)\}$. In more formal terms, the observation space $\mathcal{O}(bs_t,\, A)$ is defined by:

$$\mathcal{O}(bs_t,\, A) = \Bigg\{ \overrightarrow{w_t} : w_{ij} \geq 0 \ \forall (i,j) \in O,\ O \subseteq UB_t,\ \beta_{ij}^{w_{ij},t} > 0 \ \forall (i,j) \in O; \\ \sum_{(i,j) \in E'} w_{ij} \leq RC_t;\ RB_{t+1}(E') \subseteq O \Bigg\}, \tag{3}$$

where $E' \subseteq E_{B_t}^*$ is the set of edges that are completely cleared and $RB_t'(E')$ is the set of blocked edges that became reachable after the edges in E' have been cleared.

Since the decision maker assumes that clearance resource requirements on different blocked edges are independent, the probability of an observation $\overrightarrow{w_t}$, given current

belief bs_t and action A is:

$$P(\overrightarrow{w_t} \,|\, bs_t,\, A) = \prod_{(i,j)\in O} \beta_{ij}^{w_{ij},t}. \tag{4}$$

Let $U_t'(\overrightarrow{w_t})$ be the set of edges that are observed to be clear of debris at the end of period t. When an action A under the current belief bs_t is taken, the resulting reward also depends on the resulting observation $\overrightarrow{w_t} \in \mathcal{O}(bs_t,\, A)$. Calculating the reward requires solving a maximum weighted flow problem for the subgraph of G induced by $U_{t+1}(\overrightarrow{w_t})$. In more formal terms, the reward is calculated from:

$$R(\overrightarrow{w_t} \,|\, bs_t, A) = \sum_{i\in N_D} p_i \,(rd_{i1} - rd_{it}) + z^*(f^*, x^* \,|\, \overrightarrow{w_t}, bs_t, A), \tag{5}$$

Demand satisfied in each period keeps accruing benefits throughout the horizon. Therefore the first term in (5) captures these benefits for demand satisfied in periods prior to t. The second term $z^*(f^*, x^* \,|\, \overrightarrow{w_t}, bs_t, A)$ is determined by solving the linear program (LP1) corresponding to the maximum weighted flow problem on the subgraph of G induced by $U_t'(\overrightarrow{w_t})$, where f_{ij} denotes the amount of flow on edge (i, j) and x_i denotes the amount of demand at node i satisfied in the current period:

$$(\textbf{LP1}) \qquad \max \quad \sum_{i\in DN} p_i\, x_i \tag{6}$$

$$s.t. \quad \sum_{j:(j,i)\in U_t'(\overrightarrow{w_t})} f_{ji} - \sum_{j:(i,j)\in U_t'(\overrightarrow{w_t})} f_{ij} = x_i \qquad \forall i \in DN \tag{7}$$

$$\sum_{j:(i,j)\in U_t'(\overrightarrow{w_t})} f_{ij} - \sum_{j:(j,i)\in U_t'(\overrightarrow{w_t})} f_{ji} \le rs_{it} \qquad \forall i \in SN \tag{8}$$

$$\sum_{j:(j,i)\in U_{t+1}'(\overrightarrow{w_t})} f_{ji} - \sum_{j:(i,j)\in U_t'(\overrightarrow{w_t})} f_{ij} = 0 \qquad \forall i \in TN \tag{9}$$

$$x_i \le rd_{it} \qquad \forall i \in DN \tag{10}$$

$$f_{ij} \le FC_{ij} \qquad \forall (i,j) \in U_t'(\overrightarrow{w_t}) \tag{11}$$

$$x_i \ge 0 \qquad \forall i \in DN \tag{12}$$

$$f_{ij} \ge 0 \qquad \forall (i,j) \in U_t'(\overrightarrow{w_t}). \tag{13}$$

Objective function (6) maximizes the total weighted flow from supply nodes to demand nodes. Constraints (7) through (9) ensure that the flow balance is conserved for demand, supply, and transshipment nodes, respectively. Constraints (10) bound the net flow into a demand node by the total demand of the node. Constraints (11) are the road capacity constraints, whereas (12) and (13) are the non-negativity constraints.

The expected reward under each action A is the expected value of the possible rewards resulting from the observations emanating from it:

$$ER(bs_t,\, A) = \sum_{\overrightarrow{w_t} \in \mathcal{O}(bs_t,\, A)} \left(\prod_{(i,j) \in O} \beta_{ij}^{w_{ij},t}\, R(\overrightarrow{w_t} \,|\, bs_t, A) \right). \tag{14}$$

Based on the current belief state bs_t, under action A, and observation $\overrightarrow{w_t}$, the edge sets are updated as follows:

$$U_{t+1} \leftarrow U'_t(\overrightarrow{w_t}), \quad B_{t+1} \leftarrow E \setminus U_{t+1} \tag{15}$$

$$RB_{t+1} \leftarrow RB'_t(E'(A, \overrightarrow{w_t})), \quad UB_{t+1} \leftarrow B_{t+1} \setminus RB_{t+1}, \tag{16}$$

where $E'(A, \overrightarrow{w_t})$ is the set of edges that are completely cleared based on action A and observation $\overrightarrow{w_t}$. Given these updated sets, the updated belief $bs_{t+1}(bs_t,\, A,\, \overrightarrow{w_t}) = (\overrightarrow{Y_{t+1}},\, RC_{t+1},\, \overrightarrow{RS_{t+1}},\, \overrightarrow{RD_{t+1}})$ is obtained by:

$$\beta_{ij}^{d,t+1} = \begin{cases} \beta_{ij}^{lt}, & \text{if } (i,j) \in UB_{t+1} \\ 1, & \text{if } (i,j) \in RB_{t+1},\ l = W_{ij}^t \\ 0, & \text{if } (i,j) \in RB_{t+1},\ l \neq W_{ij}^t \end{cases} \tag{17}$$

$$RC_{t+1} = \max\{k_{t+1} - \gamma_t,\, 0\} \tag{18}$$

$$rs_{i,t+1} = rs_{it} - \left(\sum_{j:(i,j) \in U'_{t+1}(\overrightarrow{w_t})} f_{ij} - \sum_{j:(j,i) \in U'_t(\overrightarrow{W})} f_{ji} \right) \tag{19}$$

$$rd_{i,t+1} = rd_{it} - x_i^*, \tag{20}$$

where γ_t is the remaining clearance resource requirement for the partially cleared edge (if any) at the end of period t and x_i^* is the amount of demand satisfied at node i, found by solving the linear program defined by (6)-(13). Equation (17) states that if a blocked edge stays unreachable after period t, then its belief probabilities do not change. If a blocked edge is reachable, then the belief probability is 1 for its actual amount and 0 for all other amounts. Equation (18) adjusts the starting capacity at the beginning of period $t+1$ by considering potential ongoing clearance from period t. Equations (19) and (20) update the remaining supply and demand quantities in the supply and demand nodes, respectively.

We set $RC_1 = k_1$ and $\gamma^0 = 0$. Based on these, γ_t is updated as follows:

$$\gamma_t = \begin{cases} 0, & \text{if } RC_t > \sum_{(i,j)\in E'(A,\overrightarrow{w_t})} W_{ij}^t \text{ and } A = E'(A,\overrightarrow{w_t}) \\ \gamma_{t-1} - k_t, & \text{if } RC_t = 0 \\ \sum_{(i,j)\in E''(A,\overrightarrow{w_t})} W_{ij} - RC_t, & \text{otherwise,} \end{cases}$$

where the set $E''(A,\overrightarrow{w_t})$ includes edges in $E'(A,\overrightarrow{w_t})$ and the edge that is partially cleared in period t. In the first case, there is no partial clearance at the end of period t and hence $\gamma_t = 0$. In the second case, the partial clearance from period $t-1$ consumes all the clearance capacity in period t, therefore γ_t equals the remaining requirement on the partially cleared edge at the end of period t. In the last case, a number of edges are completely cleared, and γ_t equals the remaining resource requirement on the partially cleared edge at the end of the period.

The POMDP model for SDCP can be solved by using Bellman equations and treating the belief states as states of an MDP model. Let $v^*(bs_t)$ denote the optimal benefit-to-go of belief state bs_t. Then, the following set of Bellman equations solve

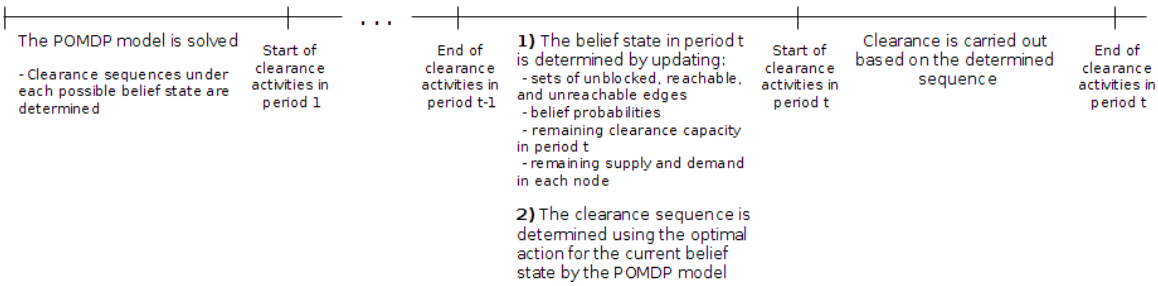

Figure 3: Timeline of decisions and events in the POMDP model before and during period t

the POMDP model for SDCP:

$$v^*(bs_t) = \max_{A \in \mathcal{A}_t} \left\{ \sum_{\overrightarrow{w_t} \in \mathcal{O}(bs_t,\, A)} \prod_{(i,j) \in O} \beta_{ij}^{w_{ij},t} \left(R(\overrightarrow{w_t} \,|\, bs_t,\, A) + v^*(bs_{t+1} \,|\, bs_t,\, A,\, \overrightarrow{w_t}) \right) \right\}. \tag{21}$$

The solution process for the POMDP can be represented by a search tree in which each belief is represented by a belief node and connected to these are alternative action nodes with their corresponding rewards. Observations are represented by arcs emanating from the action nodes, each with their corresponding probabilities. Each new observation leads to an updated belief node. The optimal policy can be found by backward recursion, taking expectations over observations and selecting the action with maximum benefit-to-go. Consequently, the POMDP model is an offline (or static) solution method, that is, it needs to be solved once and finds the optimal action for each possible belief state resulting from all possible observations. Figure 3 summarizes the timeline of decisions throughout a given period.

The solution process for the POMDP can be represented by a search tree in which each belief is represented by a belief node and connected to these are alternative action nodes with their corresponding rewards. Observations are represented by arcs emanating from the action nodes, each with their corresponding probabilities. Each new observation leads to an updated belief node. The optimal policy can be found by backward recursion, taking expectations over observations and selecting the action

with maximum benefit-to-go.

Since the action space consists of subsets of possible edges that can be (partially) cleared in one period and the observation space consists of possible combinations of actual debris amounts on the edges that become reachable, both action and observation spaces are substantially large, even for small-sized instances. Hence, the number of possible belief states increase significantly with the number of alternative actions and possible observations, resulting in a "curse of dimensionality" [136]. Nevertheless, this exact solution approach provides a framework for solving small- to mid-sized problems to optimality and for developing heuristics.

2.2.3 Extension to Regional Updates

In the above formulation, belief probabilities β_{ij}^{dt} for the edges (i, j) in UB_{t+1} remain the same (see Equation (17)). Furthermore, the decision maker assumes that the clearance resource requirements on the roads are independent of each other (Equation (4)). However, in real life disaster scenarios, the clearance resource requirements on the roads in the same region might be correlated due to similar geographic and demographic properties. One way to incorporate this is to start with an initial estimate of correlation matrix and apply Bayesian updates on these parameters as more blocked roads become reachable. However, the initial correlation estimates are very difficult or even impossible to obtain in the case of SDCP. Hence, we utilize an alternative framework for updating belief probabilities in a region concurrently based on previous observations in the same region [129]. The approach still uses Equation (4) to calculate the probabilities, but handles the inherent dependence by updating the individual belief probabilities in equation (17) using prior observations in the same region.

To incorporate potential within-region similarities in clearance resource requirements and information updates into our model, we use a Bayesian information update

approach similar to Powell and Ryzhov [129]. The disaster area, i.e., first we assume that the edge set E, is partitioned into regions $\{E_1, E_2, \ldots, E_R\}$ and let U_r^t, B_r^t, and UB_r^t denote the clear, blocked, and unreachable blocked edges in region r at the beginning of period t, respectively. Assume without loss of generality, that all roads in the set UB_r^1 in region r have the same resource requirement levels $D = \{1, 2, \ldots, N^r\}$. Suppose also that these follow a categorical distribution, that is, each resource requirement level $l \in D$ has an associated probability $\beta_{l,1}^r$ for all $(i,j) \in UB_1^r$. To preserve conjugacy while updating the parameters of the categorical distribution, we apply the Dirichlet distribution as prior [78], where:

$$\beta_{l,1}^r = \frac{\alpha_{l,1}^r}{\alpha_{l,1}^r + \ldots + \alpha_{N^r,1}^r} \text{ for all } l = 1, 2, \ldots, N^r, \tag{22}$$

and $\alpha_{l,1}^r$ represents the "weight" given to resource requirement level l relative to the other clearance resource requirements in the same region.

Assume that for debris level l, we have made c_{lt}^r observations in region r after clearance is complete in period t. Then, we let $\alpha_{l,t+1}^r = \alpha_{lt}^r + c_{lt}^r$ for all $l = 1, 2, \ldots, N^r$. In this update mechanism, the higher the initial α values, the lower is the effect of new resource requirement observations on the update of the belief probabilities. Therefore while the relative α values determine the belief probability for each level, the nominal values determine the strength of the initial beliefs.

To incorporate the regional update mechanism into the POMDP model, Equation (17) is replaced by:

$$\alpha_{l,t+1}^r = \alpha_{lt} + c_{lt}^r \qquad \forall l \in \{1, 2, \ldots, N^r\} \tag{23}$$

$$\xi_{l,t+1}^r = \frac{\alpha_{l,t}^r}{\alpha_{1,t}^r + \ldots + \alpha_{N^r,t}^r} \qquad \forall l \in \{1, 2, \ldots, N^r\} \tag{24}$$

$$\beta_{ij}^{l,t+1} = \begin{cases} \xi_{l,t+1}^r, & \text{if } (i,j) \in UB_{t+1} \\ 1, & \text{if } (i,j) \in RB_{t+1},\ l = W_{ij}^t \\ 0, & \text{if } (i,j) \in RB_{t+1},\ l \neq W_{ij}^t \end{cases} \qquad \forall l \in \{1, 2, \ldots, N^r\}. \tag{25}$$

This regional update mechanism can be used both in the POMDP approach described in this section and the heuristics presented in §2.4.

2.3 Illustrative Examples

While curse of dimensionality prevents finding optimal solution for large-sized instances, availability of an exact solution procedure may facilitate deriving structural results from smaller instances. In this section, we show that even on simple networks with unit demand, it is hard to derive structural results due to the periodic nature of the problem. We also use an example on a small illustrative network to show that solving SDCP using the POMDP model may provide significant benefits compared to solving a deterministic model with expected debris amounts, and that making regional updates may bring additional benefits in certain cases. The last part of this section introduces a continuous time approximation, which provides the baseline for the heuristic approach discussed in §2.4.

2.3.1 An Illustrative Example on Path Networks with Unit Demand

In this part, we will show that even for a simple graph structure (a path network) and simple assumptions on the supply, demand, and beliefs on the debris amounts, optimal policies do not yield an easily generalizable structure. The insights on the difficulty of deriving structural properties even under very restrictive and simple conditions motivates us to use a continuous time approximation and develop heuristic procedures based on it in §2.4.

Figure 4 presents an example with six demand nodes, each with unit demand, a single supply node with unlimited supply, and a constant clearance capacity of C units per period. Clearance starts from the supply node. The roads incident to the supply node have $\frac{C}{2}$ units of debris and for each unreachable road, the two possible debris amounts are βC and $(1-\beta)C$, where $0 < \beta < \frac{1}{2}$. For the unreachable roads to the left and right of the supply node, corresponding belief probabilities for βC are p

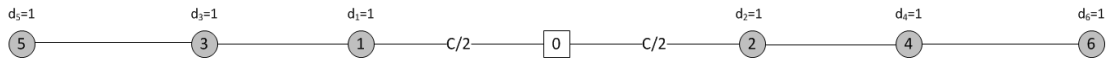

Figure 4: An instance on a path with six demand nodes, each with unit demand

and q respectively, and the planning horizon is four periods. Our aim is to determine the optimal clearance policy based on the combinations of values of β, p, and q.

In Figure 5, plots of the optimal policies based on varying values of the parameters are given. When $\beta \leq \frac{1}{4}$, four possible policies can be optimal:

- Policy 1 first clears the edges to the right of the supply node, then clears the edges to the left.
- Policy 2 first clears the edges to the left of the supply node, then clears the edges to the right.
- Policy 3 starts by clearing edges $(0, 1)$ and $(0, 2)$ first. Then it proceeds left unless there is a reachable low-debris edge on the right, in which case it clears this low-debris edge before proceeding to the left.
- Policy 4 starts by clearing edges $(0, 1)$ and $(0, 2)$ first. Then it proceeds right unless there is a reachable low-debris edge on the left, in which case it clears this low-debris edge before proceeding to the right.

The results in Figure 5 demonstrate that when β is small and $|p - q|$ is large, the optimal policy favors clearing the roads on the side with less expected debris. When the possible debris amount variability is high and belief probabilities are closer to 0.5, the optimal policy favors learning the debris structure over taking the actions that bring the most immediate benefit. While this behavior is intuitive, an interesting observation is the behavior of the optimal policy when β changes. Until the threshold value of $\beta = \frac{1}{4}$, policies 1 and 3 co-exist for different values of p and q, whereas after this threshold is passed, only policy 1 can be optimal.

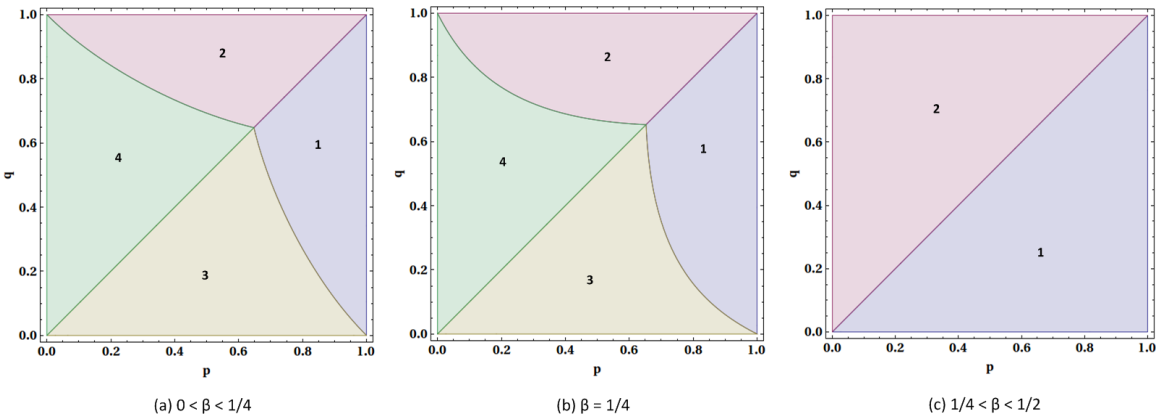

Figure 5: Optimal policies for the instance in Figure 4 for different β values and combinations of p and q

The length of the path also determines the possible optimal policies. When there are two demand nodes on each side, optimal policies resemble policies 3 and 4 only, regardless of the value of β, whereas with four demand nodes on each side, more complicated policies concurrently involving both learning and switching to the side with less expected debris to reach immediate benefits can be optimal.

The main difficulty in generalizing the optimal policies for more complicated network, supply, demand, and belief structures arises from the periodic nature of the problem, as this leads to an exponentially large action space, making many of these actions candidates to be included in the optimal policy. Furthermore, finiteness of the network prevents stationary policies from becoming optimal. The example given in this section has many restrictive simplifications such as unit demand, path network structure, and same belief amounts and probabilities for each side. In more general instances, the interdependency between these components of the problem make it even harder to generalize the results obtained in this section, adding another dimension of complexity to the problem.

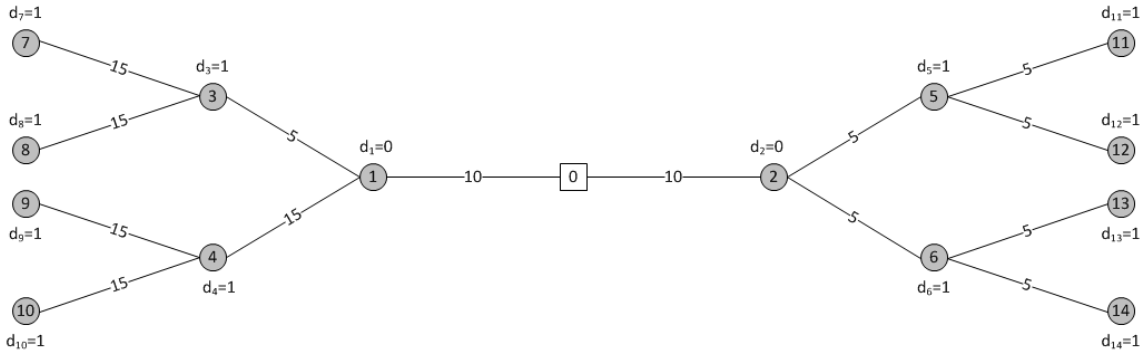

Figure 6: An example network with 12 demand nodes, each with unit demand. Quantities on the road segments show actual debris amounts.

2.3.2 An Illustrative Example to Compare Different Solution Approaches to SDCP

Figure 6 gives the network settings for a small illustrative example that we use to compare three different approaches for solving SDCP. The first of these approaches is to solve the problem using a deterministic algorithm in which we assume the debris volumes on all road segments are set to their expectations. The second approach solves SDCP using the POMDP model without the regional update mechanism described in §2.2.3, whereas the last approach adds on regional updates.

The example consists of a single supply node (given by node 0 in Figure 6), from which the clearance starts, and assumes unlimited supply capacity. There are 12 demand nodes for a single commodity, each with unit demand at the beginning of the first period and no additional demand in the coming periods. Debris amounts on each unreachable road segment can be either 5 or 15 units (actual debris amounts are given in Figure 6) and the clearance capacity in each period is 20 units. For each initially unreachable road to the left and right of the supply node, corresponding belief probabilities for 5 units of debris are 0.5 and 0.475, respectively. Planning horizon consists of seven periods. The regional update approach considers two regions (roads on the left and right of the supply node, respectively).

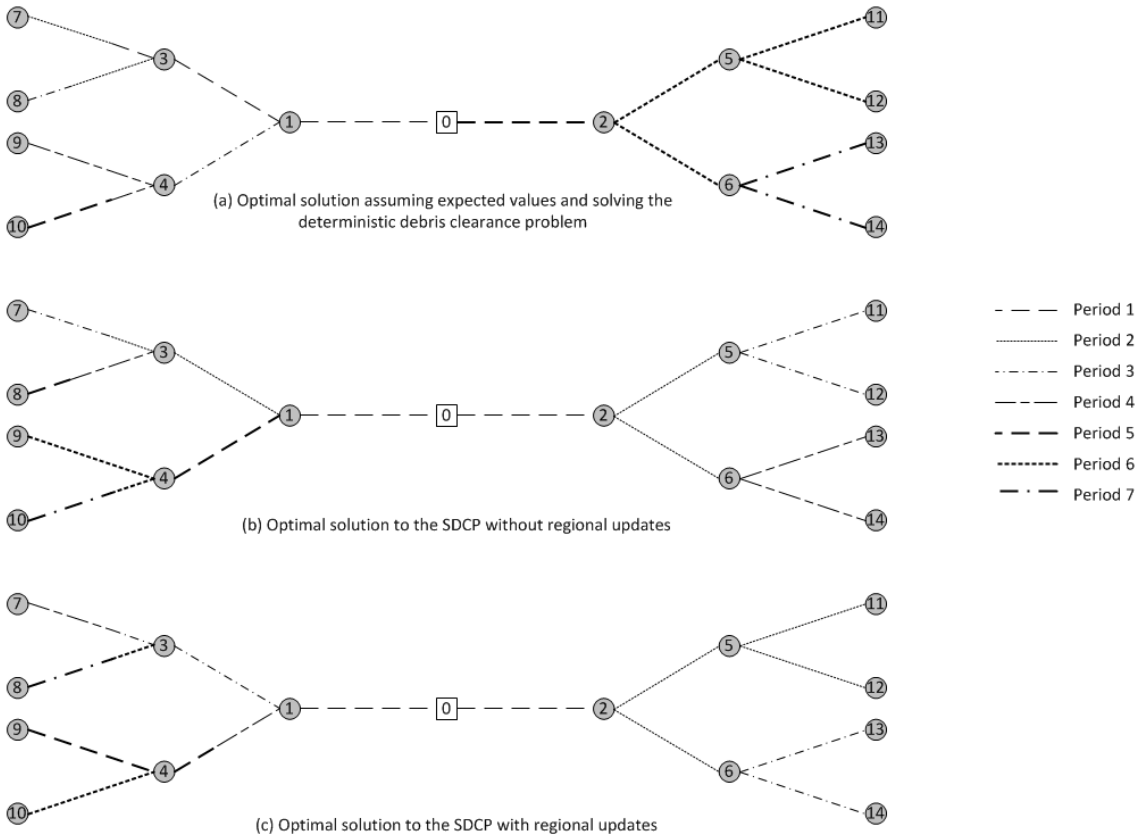

Figure 7: Optimal solutions for each of the three approaches to solving the SDCP instance in Figure 6

Figure 7 gives the clearance sequences found by solving the three approaches. Because of the homogeneous unit demand as well as identical potential debris amounts and belief probabilities within each side of the network, the expected value approach clears the roads on the left side first and proceeds with the roads on the right once the clearance on the left side is complete. This approach results in a total benefit of 40 units. The POMDP approach without regional updates starts by clearing the roads adjacent to the supply node, proceeding to clear the roads on the top-left, on the right, and lastly the ones on the bottom-left, resulting in a total benefit of 50 units.

The comparison between the first two approaches in this small example illustrates the trade-off between taking the actions that bring the most immediate benefit versus

learning the debris structure. While the deterministic approach favors the former by clearing first the roads that seemingly have less debris, the stochastic approach first collects information about the surrounding blocked roads and takes actions based on the results of this learning process. The difference between these two approaches is prominent when, as in this example, belief distribution is close to uniform and actual amounts have lower belief probability.

When regional updates are also taken into account, the POMDP approach clears the roads on the right side of the supply node after observing the debris amounts on edges $(1,3)$, $(1,4)$, $(2,5)$, and $(2,6)$. In this way, it avoids clearing the high-debris roads on the top-left, resulting in a total benefit of 52 units. In general, making regional updates tends to be beneficial compared to the POMDP approach without updates when the beliefs are close to uniform and within-region variability is low, as in this small example (The results in Tables 5 and 6 in §2.6 provide a more detailed analysis of the effects of initial beliefs on the difference between these two approaches).

This small example illustrates that although solving the POMDP model has more computational burden compared to the deterministic model with expected debris amounts, it improves the accrued benefit. Furthermore, under certain conditions, updating regional information may lead to additional benefits. A more detailed analysis of the results from applying different approaches to SDCP is made in §2.6.

2.3.3 Examples on Infinite Paths with Unit Demand

In the previous example, the main difficulty in deriving structural results for SDCP arises from the periodic structure of the problem, which results in the action space to be substantially large. In this part, we analyze SDCP on infinite path networks and attempt to observe the effect of stationary policies on the structure of the optimal solutions. The following analysis shows that for the original problem, the computational burden is still significant and the analytical results are not simple. We also

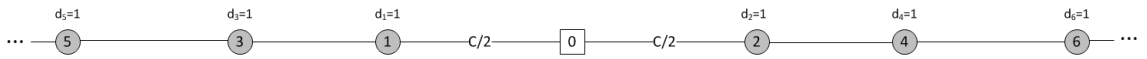

Figure 8: An instance on an infinite path with unit demand on all demand nodes

observe that a continuous time approximation decreases the computational burden significantly, allowing us to solve more general problems in shorter time. The results obtained in this part extend immediately to infinite stars and trees of the same demand and belief structure.

Figure 8 shows the infinite extension of the path network in Figure 4. As before, all nodes have unit demand, there is unlimited supply, and clearance starts from the supply node. There is a clearance capacity of C units per period, roads incident to the supply node have $\frac{C}{2}$ units of debris (without loss of generality), whereas the unreachable roads on the left and right of the supply node have βC units of debris with probability p and q, respectively. Due to the infinite nature of the problem, the objective is modified to maximize the discounted benefit obtained. Each demand node reached accrues a single unit of benefit and there is a per-period discount factor of λ.

The problem can be modeled as a MDP with the state definition consisting of four components: The first two components state the side (left or right) of ongoing clearance from the previous period at the beginning of the period and remaining capacity after ongoing clearance is complete. The last two components give the amount of debris on the nearest reachable road on both sides, unless there is ongoing clearance from previous period on that side. Action definition is identical to that given in the original POMDP model. The reward of each action is the expectation over all possible debris observations of the number of nodes reached in that period.

Our aim in solving the MDP is to find the set of parameters for which each stationary policy is optimal. To do so, we first pick a candidate stationary policy to be optimal for a subset of parameters and apply a policy iteration approach. That

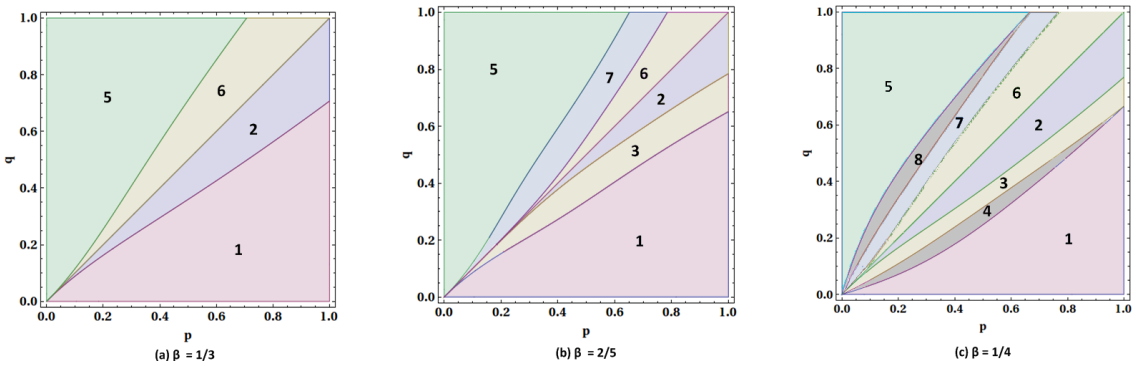

Figure 9: Values of p and q for which each policy becomes optimal for different β values when $\lambda = 0.5$

is, we first find the expected discounted reward under the given policy, and check if switching to another policy improves the reward or not.

Figure 9 gives the values of p and q that make different policies optimal for varying values of β when λ is 0.5. A common characteristic of all optimal policies is that in any given state, the optimal clearance sequence in each period starts by finishing the ongoing clearance from the previous period and clearing any low-debris roads that are reachable at the beginning of the period. After these initial clearance actions, Policy 1 consists of persistently clearing the roads on the left. Policy 2 persistently clears the roads on the right if there is high debris on the left and low debris on the right at the beginning of the period. Policies 3 and 4 first clear one and two road segments on the right, respectively, before going left, if the initial debris on the right is low (otherwise they keep going left). Policies 5 through 8 are symmetric versions of policies 1 through 4 when $p \leq q$. The results show that as the values of p and q are closer to each other, learning the debris structure is a good strategy, whereas when these values are farther apart, greedily moving to the side with less expected debris gives better expected benefit.

An interesting observation from Figure 9 is the fact that the number of potentially optimal policies is not monotonic with respect to the value of β. With β values of $\frac{1}{4}$, $\frac{1}{3}$, and $\frac{2}{5}$, the number of potentially optimal policies is eight, four, and six, respectively.

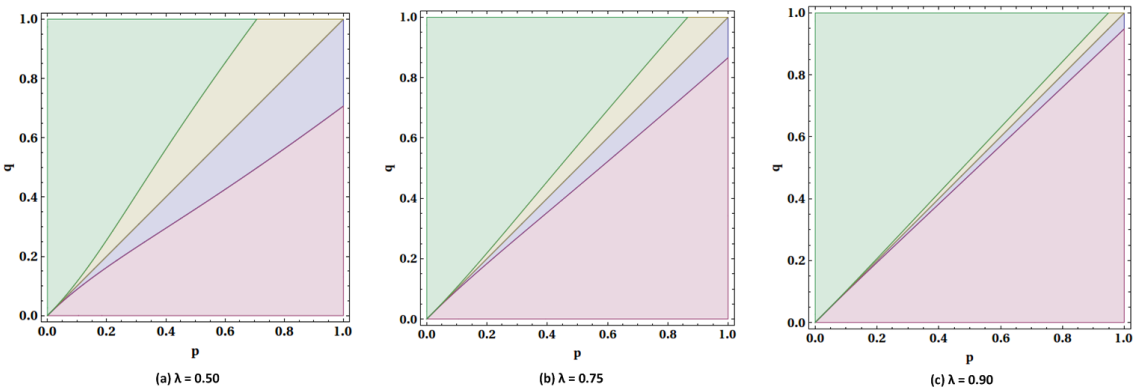

Figure 10: Values of p and q for which each policy becomes optimal for different λ values when $\beta = \frac{1}{3}$

The number of possible states is determined by where a partial clearance might start in the beginning of a period, which is not monotonic with respect to β. Since the number of potentially optimal policies is mainly determined by the number of possible states, the former is not monotonic with respect to β.

In Figure 10, the optimal policies for varying values of p and q are shown for different levels of λ when β is $\frac{1}{3}$. Changing the values of λ does not change the number of potentially optimal policies, since the difference between potential debris levels is constant, which results in the same set of potential states, and consequently, same set of potentially optimal policies. The only change is in the values of p and q for which the policies can be optimal. As expected, increasing the value of λ favors taking the actions that bring the most immediate benefit, since in the long run the side with less expected debris accrues more benefits. Consequently, letting $\lambda = 1$ leads to only policies 1 or 8 being potentially optimal, as the need to reach a demand node in a timely manner is completely ignored.

2.3.4 A Continuous Approximation

As in the case with finite networks, applying a periodic policy produces results that are hard to compute and generalize when the network structure is more complicated. We next study the case in which instead of making the clearance sequencing decisions at

the beginning of each period, we can update the debris information and decide which road to clear next after clearing each blocked road. Since this implies a continuous time model, we replace λ by a continuous discount factor $\alpha = -ln\,\lambda$. We also allow the number of possible debris levels on each side to be more than two. These possible debris levels are also allowed to be different on the left and right side. Let W_{sl} denote the debris volume corresponding to debris level l on side $s \in \{L, R\}$ and, without loss of generality, let $W_{s1} \leq \ldots \leq W_{sn}$, where n denotes the number of debris levels on side s. Let $p_s(l)$ be the corresponding probability of observing debris level l on side s.

Applying such an approximation scheme reduces the sizes of the state and action spaces drastically. In this case, state space S can be defined by (W_L, W_R), the amounts of debris on the left and right at the beginning of each clearance epoch. Action space reduces to going to either left or right. The corresponding Bellman equation is given by:

$$v^*(W_L, W_R) = max\Bigg\{1 + e^{-\alpha W_L} \sum_{(u, W_R \in S)} p_L(u)\, v^*(u, W_R), \\ 1 + e^{-\alpha W_R} \sum_{(W_L, u \in S)} p_R(u)\, v^*(W_L, u)\Bigg\}, \tag{26}$$

where $v^*(W_L, W_R)$ is the optimal value of the state (W_L, W_R). This set of equations can be easily solved by value or policy iteration. Since the number of possible states is significantly reduced and there are only two possible actions, this approach saves significant computational burden over applying a periodic approach.

For the special case of two potential debris levels on each side (W_{L1} and W_{L2} on the left and W_{R1} and W_{R2} on the right), the optimal policy is given by Algorithm 1. While the algorithm does not give insight into how the optimal policies compare to those for the above examples, it does illustrate the fact that the continuous approximation can handle any number of possible debris amounts in a more efficient way.

Algorithm 1 The algorithm to find the optimal policy with two possible debris levels on each side

if $P_L(W_{L1}) + P_L(W_{L2}) \leq P_R(W_{R1}) + P_R(W_{R2})$ **then**
  if $\frac{1}{1-e^{-\alpha W_{R1}}} \geq 1 + e^{-\alpha W_{R2}} \frac{P_L(W_{L1}+P_L(W_{L2}))}{1-P_L(W_{L1})e^{-\alpha W_{L1}}-P_L(W_{L2})e^{-\alpha W_{L2}}}$ **then**
    Always go left
  else
    Go right on (W_{L2}, W_{R1}) and go left otherwise
  end if
else
  if $\frac{1}{1-e^{-\alpha W_{L1}}} \geq 1 + e^{-\alpha W_{L2}} \frac{P_R(W_{R1}+P_R(W_{R2}))}{1-P_R(W_{R1})e^{-\alpha W_{R1}}-P_R(W_{R2})e^{-\alpha W_{R2}}}$ **then**
    Always go right
  else
    Go left on (W_{L1}, W_{R2}) and go right otherwise
  end if
end if

2.4 Heuristic Approaches

While the POMDP approach described in §2.2 can solve small to moderate sized instances of SDCP to optimality, the sizes of the action and observation spaces result in a significant computational burden. To solve larger sized problems, we propose three heuristics. The first uses a continuous time approach, ignoring the periodic nature of the clearance process (§2.4.1). The other two heuristics limit the action and observation spaces by reducing the size and depth of the search tree (§2.4.2).

2.4.1 A Heuristic Based On a Continuous-Time Approximation

In SDCP, the set of feasible actions is determined by subsets of edges that can be cleared in a period and the observation space consists of all feasible combinations of clearance resource requirements on the unreachable edges. The size of the action space poses an important challenge on solving the POMDP model. This challenge is exacerbated by the existence of multiple periods. The observations from §2.3 can be summarized as follows:

Observation 1. *Ignoring the stochastic aspect of the problem may lead to significant benefit losses unless the initial beliefs of the decision maker accurately estimate the*

actual debris amounts.

Observation 2. *Due to the periodic nature of the problem, optimal policies for simple network and demand structures are not easily generalizable.*

Particularly motivated by Observation 2, we propose a heuristic approach based on a continuous-time approximation of SDCP, which we denote as POMDP-C. The approximation assumes that once a debris-blocked edge is cleared, we select the next edge to be cleared immediately, without any consideration of periods or capacity. The main aim is to find an entire sequence of edges to be cleared. Once the sequence is determined, a transformation process results in a periodic policy that is capacity-feasible for each period. This relaxation reduces the size of the action space drastically, as it now consists only of blocked edges that are immediately reachable. It also shrinks the observation space significantly, since new observations are obtained only from the edges that become reachable, which are now smaller in number.

In the network optimization literature, there are several "continuous space" relaxations for finding near optimal routes for various versions of the TSP and the vehicle routing problem. Examples include Daganzo [47, 48, 49], Del Castillo [54], Daganzo and Erera [50], Francis and Smilowitz [74], and Zhong et al. [175]. In this chapter, we take a different approach by focusing on time, rather than space, in our approximation.

POMDP-C can be given as follows: While definitions of the state space and the belief state space do not change (except that index t now indicates that $t-1$ edges have been cleared so far), the action space now consists of blocked edges that are reachable. More formally, Equation (2) becomes:

$$\mathcal{A}_t = RB_t. \tag{27}$$

Since the action does not clear any unreachable blocked edges, new observations only arise from edges that become reachable after an edge is cleared as a consequence

of the action. If $A' \in \mathcal{A}_t$ is the edge that has just been cleared, then Equation (3) becomes:

$$\mathcal{O}(bs_t,\, A') = \{w_{ij}\ \forall (i,j) \in RB_{t+1}(A')\}. \tag{28}$$

Due to the modification of the observation space, the transition probabilities apply only to the edges that become reachable as a result of the action, that is, Equation (4) is replaced by:

$$P(\overrightarrow{W}|bs_t,\, A') = \prod_{(i,j)\in RB_{t+1}(A')} \beta_{ij}^{W_{ij},t}. \tag{29}$$

In order to adjust the reward function to the continuous time approximation, the benefit accrued must be multiplied by the amount of time that has passed since the beginning of the current period. Let set $CE_t(bs_t,\, A')$ include edge A' and the edges cleared until the t^{th} clearance. Then, Equation (14) needs to be multiplied by:

$$\begin{cases} \frac{W_{ij}}{C}, & \text{if } \left\lfloor \frac{\sum_{(i,j)\in CE(bs_t,A')} W_{ij}}{C} \right\rfloor = \left\lfloor \frac{\sum_{(i,j)\in CE(bs_t,A')\setminus A'} W_{ij}}{C} \right\rfloor \\ \frac{\sum_{(i,j)\in CE(bs_t,A')\setminus A'} W_{ij}}{C} - \left\lfloor \frac{\sum_{(i,j)\in CE(bs_t,A')\setminus A'} W_{ij}}{C} \right\rfloor, & \text{if } \left\lfloor \frac{\sum_{(i,j)\in CE(bs_t,A')} W_{ij}}{C} \right\rfloor \neq \left\lfloor \frac{\sum_{(i,j)\in CE(bs_t,A')\setminus A'} W_{ij}}{C} \right\rfloor. \end{cases} \tag{30}$$

Furthermore, the term $U_{t+1}(\overrightarrow{W})$ in Equations (7)-(9), (11), and (13) can now be replaced by $U_t \cup A'$. Due to the change in the observation space, Equation (14) is simplified to:

$$R(bs_t,\, A') = \sum_{\overrightarrow{W}\in\mathcal{O}(bs_t,\, A')} \left(\prod_{(i,j)\in RB_{t+1}(A')} \beta_{ij}^{w_{ij},t}\, R(\overrightarrow{W} \,|\, bs_t, A') \right). \tag{31}$$

The set updates in the belief update equations are also simplified. Update steps in expressions (15) and (16) change to:

$$U_{t+1} \leftarrow U_t \cup A', \quad B_{t+1} \leftarrow B_t \setminus A' \tag{32}$$

$$RB_{t+1} \leftarrow RB_{t+1}(A'), \quad UB_{t+1} \leftarrow B_{t+1} \setminus RB_{t+1}. \tag{33}$$

The remaining components of the belief state update are unchanged.

Based on the given simplifications, Bellman equations in Equation (21) become:

$$v^*(bs_t) = \max_{A' \in RB_t} \left\{ \sum_{\overrightarrow{W} \in \mathcal{O}(bs_t,\, A')} \prod_{(i,j) \in O} \beta_{ij}^{w_{ij},t} \left(R(\overrightarrow{W} \,|\, bs_t,\, A') + v^*(bs_{t+1} \,|\, bs_t,\, A',\, \overrightarrow{W}) \right) \right\}. \quad (34)$$

An important property of the continuous time approximation is given in the following observation.

Observation 3. *The continuous time approximation provides an upper bound on the optimal objective value of SDCP.*

Proof. Consider the optimal policy of the original problem. Suppose first that it does not involve idling at the end of any period. This optimal policy gives a sequencing of edges under each possible debris amount observation. Each sequence under each possible observation can be transformed into a continuous time policy by considering the edges in the sequence individually. By construction, the rewards obtained in both cases are identical. Suppose now that the optimal policy in the original problem involves idling. Then the transformed policy in the continuous model gives an expected reward that is strictly higher. Hence in both cases, the optimal policy of the original problem can be transformed to a feasible policy in the continuous approximation with an objective value that is no smaller than that in the original problem. □

To transform the optimal continuous policy into one that is feasible for the original problem, we use a sampling approach. We start from the root node of the search tree and follow the action sequence resulting from the continuous approximation. To move through observation arcs, we select the observations randomly based on their occurrence probabilities. Once sequencing is complete, we cluster the edges for the first period based on the smallest debris observations until capacity is exhausted. After receiving the actual observations, we move along the search tree according to the these observations and repeat the procedure. The resulting heuristic, is called the *continuous-time heuristic*.

Observation 4. *The expected reward obtained by POMDP-C yields a lower bound on the optimal objective value of SDCP.*

Proof. By construction, the transformation results in a policy whose actions are in the action space defined by Equation (2). Furthermore, the observations are selected from the search tree for the original problem. Thus, the resulting policy is feasible for SDCP, and the expected reward of the policy is therefore a lower bound on the optimal objective value. □

2.4.2 Heuristics for Approximate Tree Search

While the computational results in §2.6 show that the run-time of POMDP-C is significantly lower compared to the POMDP model discussed in §2.2, it still fails to solve large sized problems due to the depth of the tree. Hence, it is not practical to use the continuous-time heuristic on large instances. To overcome this issue and enhance the computational time performance, we propose two additional heuristics, namely *heuristic pruning* (HP) and *limited-lookahead* (LL), which are applicable on the search trees of both the original POMDP model and the continuous-time approximation.

HP is an "upper bounding" scheme on the search tree, which is commonly used for solving POMDP models in the literature (e.g., [124, 135]). The general approach for a maximization problem, which is analogous to branch-and-bound, is to find an upper bound for the belief node, compare it to the best known lower bound so far, and prune the search node if its upper bound is smaller than the lower bound. In the case of SDCP, to get an upper bound for a belief node, we assume that all remaining unreachable edges have the lowest possible clearance resource requirements and solve a deterministic version of SDCP. We then prune the belief nodes whose upper bounds are lower than the expected benefit from the best solution so far. For larger problems, due to significant computational requirements of solving the deterministic debris clearance problem to optimality, we use the heuristics proposed

by Stilp et al. [149]. As problem size is an important factor in the solution time of HP, we apply it only if the number of remaining blocked edges is smaller than a threshold. While the objective value resulting from the heuristic does not always give a valid upper bound, it can be used to heuristically prune some of the belief nodes.

LL aims to keep in check the depth of the search tree by applying limited look-ahead, which assumes expected debris amounts on all unreachable blocked edges and uses a greedy approach to select a subset of edges (or a single edge for the continuous-time approximation) that gives the highest ratio of benefit to cleared debris for each period. LL finds the approximate benefit-to-go values of the leaf nodes in the tree. By backward induction from these leaf nodes to the root, we can find the action with the highest benefit-to-go for the coming period and repeat for the next period after receiving the actual observations. Both HP and LL fall into the category of online algorithms for POMDP [136], as they dynamically update the decisions based on new information.

The procedures discussed in this section can be applied under both the original POMDP approach and the continuous time approximation. They can also easily incorporate regional updates, as the information update mechanism is not modified by any of these approaches.

2.5 Structural Results on Simple Networks

In this section, we consider SDCP on simple network structures and show that under certain settings, the continuous-time heuristic solves SDCP to optimality. Furthermore, under several more general network settings, solutions generated by the heuristic yields structural properties that are very similar those of the optimal solution.

Consider a network with a single supply and a single demand node connected with k node-disjoint paths (NDP). Furthermore, assume that for each initially unreachable edge, there are two possible resource requirement levels, l_1 and l_2, where $l_1 < l_2$. On

each path $j \in \{1, 2, \ldots, k\}$, the probability that an edge has l_1 units is Bernoulli distributed with probability p_j, denoted by $Bernoulli(l_1, l_2; \{p_j\}_{j=1}^k)$. We focus on two special cases: (i) $l_1 = 0$ and $l_2 > 0$ (without loss of generality, we assume $l_2 = a$), and (ii) $l_2 > l_1 > 0$ (without loss of generality, we assume $l_1 = 1$, $l_2 = K$).

Theorem 1. *On an NDP network with a single supply and demand node and resource requirements distributed with $Bernoulli(0, a; \{p_j\}_{j=1}^k)$, the optimal (continuous and periodic) policy proceeds on the path that has the least expected debris on its blocked edges until the demand node.*

Proof. For the sake of brevity, we only provide the proof for the continuous case. The proof is very similar for the periodic case.

At any step of the optimal algorithm, let $v^*\left(\{w_j, n_j, p_j\}_{j=1}^k\right)$ denote the expected amount of debris to be cleared until reaching the demand node, where w_j, n_j, and p_j denote the amount of the debris in the first reachable edge, the number of unreachable edges until the demand node, and the probability of observing no debris on an unreachable edge on path j. We first note that the optimal expected debris amount to be cleared is symmetric with regard to the paths, that is,

$$v^*\left(w_j, n_j, p_j, a_l, n_l, p_l, \{w_m, n_m, p_m\}_{m \neq j,l}\right) = v^*\left(w_l, n_l, p_l, w_j, n_j, p_j, \{w_m, n_m, p_m\}_{m \neq j,l}\right).$$

Next, we show that the expected debris amount to be cleared is monotone nondecreasing in n_j and monotone nonincreasing in p_j, $j = 1, \ldots, k$.

In order to show the monotone nondecreasing property for n_j, by symmetry, it suffices to show the property for n_1. Let $\mathcal{C}_{n_1} = (w_1, n_1, p_1, \{w_j, n_j, p_j\}_{j>1})$, $\mathcal{C}_{n_1+n} = (w_1, n_1 + n, p_1, \{w_j, n_j, p_j\}_{j>1})$ for some $n > 0$, and $\tilde{\mathcal{C}}_{n_1+n} = \mathcal{C}_{n_1+n}$ with the last n edges on the first path having zero debris. Let $\mathcal{A}$ be the optimal algorithm on $\mathcal{C}_{n_1+n}$. Define the algorithm $\mathcal{A}^*$ as follows: Run $\mathcal{A}$ until either the demand node or the last node before the zero-debris edges on the first path. If it reaches the former first, stop. Else, proceed to the demand node from the latter node. If $v^*_{A'}(\mathcal{C})$ denotes the

expected debris amount to be cleared on network configuration $\mathcal{C}$, we have:

$$\begin{aligned} v^*\left(w_1, n_1, p_1, \{w_j, n_j, p_j\}_{j>1}\right) &\leq v^*_{\mathcal{A}}\left(w_1, n_1, p_1, \{w_j, n_j, p_j\}_{j>1}\right) \\ &= v^*_{\mathcal{A}^*}(\tilde{\mathcal{C}}_{n_1+n}) \\ &\leq v^*\left(w_1, n_1+n, p_1, \{w_j, n_j, p_j\}_{j>1}\right). \end{aligned}$$

For p_1, let

$$\mathcal{C}_{p_1} = \left(w_1, n_1, p_1, \{w_j, n_j, p_j\}_{j>1}\right),$$

and

$$\mathcal{C}_{p_1-\epsilon} = \left(w_1, n_1, p_1-\epsilon, \{w_j, n_j, p_j\}_{j>1}\right).$$

Let $\mathcal{A}$ be the optimal algorithm on $\mathcal{C}_{p_1-\epsilon}$. Then:

$$\begin{aligned} v^*\left(w_1, n_1, p_1, \{w_j, n_j, p_j\}_{j>1}\right) &\leq v^*_{\mathcal{A}}(\mathcal{C}_{p_1}) \\ &\leq v^*_{\mathcal{A}}(\mathcal{C}_{p_1-\epsilon}) \\ &= v^*\left(a+a_1, n_1, p_1-\epsilon, \{w_j, n_j, p_j\}_{j>1}\right). \end{aligned}$$

At any step of the optimal algorithm, let ED_j be the total expected debris on the unreachable edges on path j. Since the optimal amount of debris to be cleared is nondecreasing in n_j and nonincreasing in p_j, it is also nondecreasing in ED_j. As the next step, we show that if $w_1 = \min_j w_j$ and $ED_1 = \min_j ED_j$, it is optimal to proceed along the first path until the first unreachable edge. If $w_1 = 0$, the claim is obviously true. Therefore assume that $w_1 = 1$, which requires that $w_j = 1$ for all paths j. Assume the optimal algorithm chooses to proceed along a path $l \neq 1$. If a zero-debris edge is encountered next, the optimal policy is to keep on this path. Otherwise, by the monotonicity of the expected debris volume function, it is still optimal to proceed on this path, since $v^*\left(\{1, ED_j\}_{j=1}^k\right)$ being minimized by choosing path j implies that $v^*\left(1, ED_l\frac{n_l-2}{n_l-1}, \{1, ED_j\}_{j\neq l}\right)$ is still minimized by choosing path j. The expected cost of this policy is $1+ED_j$. However, the expected cost of proceeding on path 1 from the beginning is $1+ED_1$, which is smaller.

This result implies that when all paths initially have 1 unit of debris, it is optimal to proceed with the one that has the least expected debris remaining. Hence, the optimal policy starts by traversing through all the zero-debris edges that are reachable and when all such edges are explored, proceeding with the path with the lowest remaining expected debris. □

In addition to characterizing the optimal policy, we also state the expected time to reach the demand node under the optimal policy in the following theorem.

Theorem 2. *On an NDP network with a single supply and demand node, the expected debris volume to be cleared until reaching the demand node under the optimal policy and Bernoulli*$(0, a; \{p_j\}_{j=1}^{k})$ *debris levels is given by:*

$$\sum_{i=0}^{n_{min}-1} \left[\prod_{j=1}^{k} \left(1 - p_j^{\,n_j - \left\lceil \frac{i(1-p_{min})}{1-p_j} \right\rceil} \right) - \prod_{j=1}^{k} \left(1 - p_j^{\,n_j - \left\lceil \frac{i(1-p_{min})}{1-p_j} \right\rceil - 1} \right) \right] \Big[a \left(1 + i(1-p_j) \right) \Big], \quad (35)$$

where n_{min} and p_{min} denote the number of edges and the probability of observing no debris on an edge on the path that minimizes the product $n_j\, p_j$ over all paths j, respectively.

Proof. The optimal policy under these settings is to traverse all reachable zero-debris edges until there is none left, and proceed with the path with the least expected debris on the unreachable edges until the demand node. In the case there is at least one path that reaches the demand node through zero-debris edges, the debris volume cleared under the optimal policy is zero. Therefore, we focus on the case where no path has all zero-debris edges until the demand node. Let n_{min} and p_{min} denote the number of edges and the probability of observing no debris on an edge on the path that minimizes the product $n_j\, p_j$ over all paths j, respectively. Furthermore, let $E(D)$ denote the total expected debris until reaching the demand node under the optimal continuous policy. We start by conditioning on the number of unreachable edges on the path that has the least expected debris on its unreachable edges after all reachable

zero-debris edges are cleared:

$$E(D) = \sum_{i=0}^{n_{min}-1} E(D \,|\, \text{path with least expected debris on unreachable edges has } i \text{ unreachable edges}) \times P(\text{path with least expected debris on unreachable edges has } i \text{ unreachable edges}).$$

We also have:

$$E(D \,|\, \text{path with least expected debris on unreachable edges has } i \text{ unreachable edges}) = a(1 + i(1-p)).$$

For $P(\text{path with least expected debris on unreachable edges has } i \text{ unreachable edges})$, we have:

$$\begin{aligned}
&P(\text{path with least expected debris on unreachable edges has } i \text{ unreachable edges}) \\
=\, &P(\text{path with least expected debris on unreachable edges has } \geq i \text{ unreachable edges}) \\
&- P(\text{path with least expected debris on unreachable edges has } \geq i+1 \text{ unreachable edges}) \\
=\, &P\left(\text{on each path } j, \text{ the first } n_j - \left\lceil \frac{i(1-p_{min})}{1-p_j} \right\rceil \text{ edges have at least one edge with } a \text{ units of debris}\right) \\
&- P\left(\text{on each path } j, \text{ the first } n_j - \left\lceil \frac{i(1-p_{min})}{1-p_j} \right\rceil - 1 \text{ edges have at least one edge with } a \text{ units of debris}\right) \\
=\, &\prod_{j=1}^{k} \left(1 - p_j^{\,n_j - \left\lceil \frac{i(1-p_{min})}{1-p_j} \right\rceil}\right) - \prod_{j=1}^{k} \left(1 - p_j^{\,n_j - \left\lceil \frac{i(1-p_{min})}{1-p_j} \right\rceil - 1}\right).
\end{aligned}$$

Altogether, these yield:

$$E(D) = \sum_{i=0}^{n_{min}-1} \left[\prod_{j=1}^{k} \left(1 - p_j^{\,n_j - \left\lceil \frac{i(1-p_{min})}{1-p_j} \right\rceil}\right) - \prod_{j=1}^{k} \left(1 - p_j^{\,n_j - \left\lceil \frac{i(1-p_{min})}{1-p_j} \right\rceil - 1}\right)\right] \Big[a\,(1 + i(1-p_j))\Big].$$

□

When clearance resource requirements can take two positive values, the optimal policies for the continuous and periodic models are more complicated and not identical, but exhibit similar properties, as stated in Theorem 3. Here, a 1-debris edge and a K-debris edge have 1 and K units of resource requirements, respectively.

Theorem 3. *Consider an NDP network with Bernoulli*$(1, K; \{p_j\}_{j=1}^{k})$ *debris levels.*

- *If a* 1*-debris edge is cleared last on a path, the optimal (continuous and periodic) policy is to proceed on that path until a* K*-debris edge is reached.*
- *If a* K*-debris edge is cleared last on a path,*
  - *the optimal continuous policy is to proceed on that path until the demand node.*
  - *the optimal periodic policy is to proceed on that path until the demand node, unless a* K*-debris edge is completely cleared at the end of a period, and another* K*-debris edge becomes reachable.*

Proof. (Sketch) Suppose the optimal policy traverses an edge with 1 unit of debris. This indicates that the sum of the cost of traversing this edge and the expected cost of optimal policy at the next node on this path is no more than that of any other action. Since 1 unit of debris at the next edge is the lowest possible amount and hence decreases the expected debris on the remainder of the path, it still dominates switching to another path. Thus, if a 1-debris edge is traversed, it is optimal to keep on the same path until a K-debris edge is encountered.

Suppose now that the optimal continuous policy traverses a K-edge on a path. By the above argument, if the next edge on this path has 1 unit of debris, it is optimal to remain on the same path. Suppose the next encountered edge is a K-debris edge. Since the action to choose this path with one more unreachable edge and a K-debris edge as its reachable edge dominated all remaining actions, with one

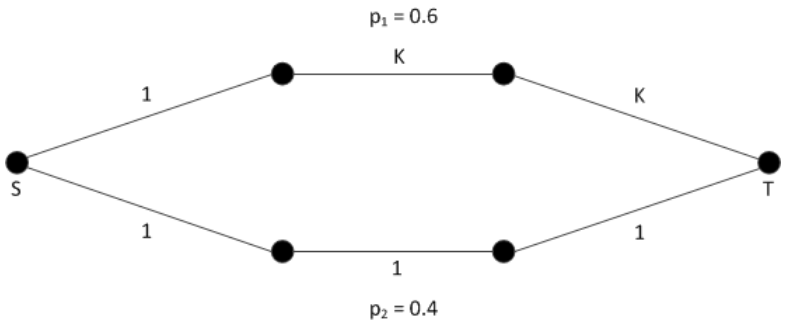

Figure 11: A counterexample with $C = K + 1$ on a network with two node-disjoint parallel paths. The periodic model proceeds with the three edges on the second path after clearing the first K-debris edge on the first path.

fewer unreachable edge, staying on this path still dominates switching to any of the remaining paths. Hence if a K-debris edge is traversed, it is optimal to keep on the same path until the demand node.

Suppose that under the optimal periodic policy, a K-debris edge is cleared in period t and the period is over at the same time the clearance is complete. If a 1-debris edge becomes reachable, by the above argument, it is optimal to proceed with it. Assume that a K-debris edge becomes reachable. An example showing that it may be optimal to switch to a different path is given in Figure 11. The network consists of two node-disjoint paths with three edges each, $p_1 = 0.6$, $p_2 = 0.4$, and $C = K + 1$. In the first period, the optimal periodic policy clears the first two edges on the first path and encounters another K-debris edge on the first path at the end of the period. In this situation, it is optimal to switch to the second path and traverse it until the demand node. □

While the optimal policy on NDPs with $Bernoulli(0, a; \{p_j\}_{j=1}^{k})$ debris levels cannot be fully characterized, finding the optimal solution is tractable for the continuous model. On the other hand, it takes exponential time to find the optimal policy under the same settings for the periodic model.

Theorem 4. *Consider an NDP network with* $Bernoulli(1, K; \{p_j\}_{j=1}^{k})$ *debris levels.*

The optimal policy for the continuous model can be found in polynomial time.

Proof. We start by showing that the continuous approximation with two paths is polynomially solvable. At any instant before committing to a path, let the graph knowledge (state) be represented by $(x_1, n_1; x_2, n_2)$, where x_j is the amount of debris on the first reachable edge on path j and n_j is the number of unreachable edges in that instant. Unlike [95], our state definition does not need to consider the number of 1-debris edges encountered so far on a path and the current location, as there is no cost of traversing an edge on the way back. Under this state definition, there are $O(|E|^2)$ states, as opposed to $O(|E|^4)$ states for CTP. Since the only actions are to proceed with the current path or to switch to the other, the optimal policy can be found in polynomial time.

Next, we show that the general problem is polynomially solvable by using the 2-path algorithm as a subroutine. First, order the paths starting with a 1-debris edge in increasing expected debris. Out of the ones starting with a K-debris edge, only keep the one with the least expected debris (committing to this path dominates committing to other paths starting with a K-debris edge, and once we commit to this path, the optimal policy is to traverse it until the demand node).

If all paths start with a K-debris edge, it is optimal to traverse the one with the least expected debris. Otherwise, formulate an MDP with the same state definition as defined for two paths, except that the index $j = 1$ for the path that has been traversed until its first K-debris edge and has the least expected debris remaining (the "best" path so far), and the index $j = 2$ for the current path. Since there is no need to store the number of 1-debris edges encountered so far on the paths and the current path, the number of states is $O(|E|^2)$, as opposed to $O(|E|^5)$ states for CTP. The only actions are to keep on the current path, to switch back to the best path so far, or to start exploring the next best unexplored path. Hence, the whole procedure can be completed in polynomial time. □

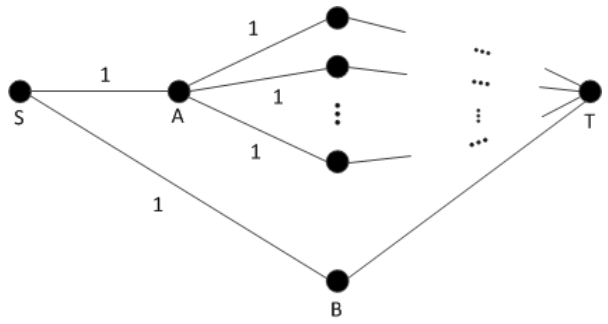

Figure 12: A counterexample on a series-parallel graph with $p = 0.5$. The optimal continuous policy proceeds to A, then returns to S, and traverses the path in the lower half to T.

The last result we show for the single supply and single demand node case shows the limitations of the continuous and periodic models when network structure is series-parallel, even when the only two possible debris levels are zero and a, with the belief distribution identical over the whole network.

Theorem 5. *The result in Theorem 1 does not apply to series-parallel graphs in general.*

Proof. The counterexample from Karger and Nikolova [95], shown in Figure 12, has n NDPs between nodes A and T, each with $log\, n$ edges, where $n \geq 5$. Using the result from Theorem 2 for NDPs with equal number of edges and belief distribution iid over the network with $p = 0.5$, the expected debris to be cleared on the NDP between nodes A and T is no more than 0.45. Since the expected cost after reaching B is 0.5, the optimal algorithm initially moves to A. However, since all edges incident to A on the parallel paths have 1 unit of debris, the expected optimal debris to be cleared from A to T is at least 3 units. Since the expected additional debris cleared by going back from A to S, and proceeding from S to T using B is 1.5 units, it is optimal to traverse the edge from S to A back after it is cleared. □

The main motivation in switching paths in this case arises when the optimal policy follows a path until a node that has a set of NDPs emanating from it that have the

highest debris possible on their first edge. Upon this realization, it may be preferable to switch to a path that has fewer edges until the destination node. For comparison, the deterministic version of the problem on a single path is NP-hard, as shown by Stilp et al. [149]. The results in Theorems 1 through 4 show that when potential debris amounts are limited to two fixed levels, the problem becomes substantially easier to solve.

The results we have provided in Theorems 1 through 4 extend the results by Karger and Nikolova [95] on CTP to the case where probabilities are path-dependent.

Next, we direct our attention to the case of a single supply node and multiple demand nodes. Assume there exist k_i NDPs leading to demand node i. We extend the notation to $Bernoulli(l_1, l_2; \{p_{ij}\}_{j=1}^{k_i})$ for each demand node i. Let I' denote the set of demand nodes not currently reached, $J(i)$ represent the set of paths leading to node i, and d_i denote the demand on node i. The aforementioned structural results can be extended to these cases as well for sufficiently large T.

Theorem 6. *Consider an NDP network with a single supply node, multiple demand nodes with varying per-unit benefits and demands, and* $Bernoulli(0, a; \{p_{ij}\}_{j=1}^{k_i})$ *debris levels. If* $|T|$ *is large enough to allow all demand nodes to become reachable when clearance is complete, the optimal (continuous and periodic) policy is as follows:*

1. *Proceed on path* $j^* \in J(i)$ *that maximizes the ratio* $\dfrac{p_i d_i}{n_{J(i)} p_i a}$ *over all* $i \in I'$ *and* $j \in J(i)$*, where* $n_{J(i)}$ *is the number of unreachable edges in path* $J(i)$*. Remove* i *from* I'*.*

2. *Repeat until* $I' = \emptyset$*.*

Proof. When T is sufficiently large, the problem is one of minimizing the time to reach the last demand node, which is identical to minimizing the makespan on a single machine where jobs have stochastic task times. As shown by [137], the optimal solution for this problem is to sequence the jobs (demand nodes) in the order of decreasing

ratio of job weights (per-unit benefit times demand on the nodes) to expected task times (expected debris on the remaining unreachable edges on the path). □

The partial characterization of the optimal policy in Theorem 3 applies to the case with $Bernoulli(1, K; \{p_{ij}\}_{j=1}^{k_i})$ debris levels. Furthermore, the optimal policy for the continuous model can be found in polynomial time. For comparison, the deterministic version of the problem with a single supply node and multiple demand nodes has been shown to be NP-hard, even on trees [149].

When there exists a single path between the supply node and each demand node, we obtain the following result:

Corollary 1. *For sufficiently large T, the results in Theorems 1 through 4 apply to the case of star networks where the supply node is in the center and demand nodes are on the leaves.*

It can be shown that the results also extend to the case of perfect k-ary trees, but not to general trees. Again, the main difficulty that arises for general trees is that it may be preferable to switch to a less deep subtree when a deeper subtree with less expected debris happens to have the highest debris level on all edges emanating from its root. It is also easy to show that the results do not apply when multiple supply nodes are present.

The results we present in this section show that on simple network settings, the optimal policy under the continuous approximation and the periodic model are either identical, or show similar structural properties. Hence, the results motivate our development of the continuous-time heuristic and exhibit under which network structures we can expect it to perform well. In the next section, we use computational experiments to investigate whether these results extend to more general settings.

2.6 Computational Experiments

The goals of our computational experiments are two-fold: to (i) determine the largest problems for which the POMDP approach can find optimal solutions, and (ii) analyze the performance of the heuristic discussed in §2.4. The quality of the heuristic solutions is evaluated by the gaps between the calculated upper bounds and the best objective values obtained during the computational experiments. We also report the run times of the heuristics, since timely decision making is crucial in this context.

[149] developed heuristics that run fast and find good quality solutions for the deterministic debris clearance problem, which is NP-hard. The stochastic nature of SDCP presents further computational challenges. The approaches we propose for solving SDCP take considerably more time and are useful only if they provide significant improvements for solution quality over solving the deterministic model with expected debris amounts. We compare the solutions generated by our proposed approaches for SDCP versus the deterministic approach, and determine when there are significant differences.

Current operations for debris clearance use predetermined road prioritization. To the best of our knowledge, a systematic procedure for determining road priorities does not exist in practice. We consider a number of "benchmark" heuristics to mimic current practice. We also determine the settings under which performance differences between benchmark heuristics and the SDCP heuristics we proposed in §2.4 are more significant.

The remainder of this section is organized as follows. We first describe the set of instances used in the experiments and how they are generated. Then, we describe the benchmark heuristics that ignore the stochastic aspect of the problem and/or apply predetermined road prioritization. Finally, we present the results of the experiments and derive conclusions.

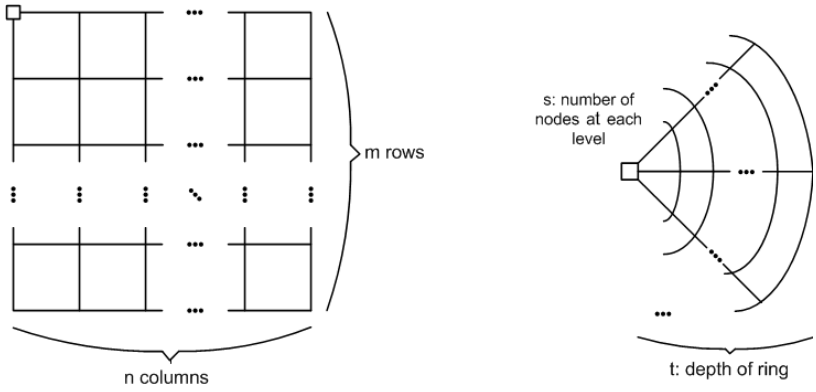

Figure 13: (a) $G(m, n)$, a grid network with m rows and n columns, (b) $R(s, t)$, a ring network with t levels and s nodes at each level (excluding the center)

2.6.1 Instance Generation

The instance sets we generate can be classified under two main groups: (i) structured grid and ring networks, which resemble urban city road networks, and (ii) a realistic earthquake scenario.

We denote a grid network with m rows and n columns by $G(m \times n)$, whereas $R(s \times t)$ represents a ring network with t levels and s nodes at each level, as illustrated in Figure 13. The sizes of both grid and ring networks in the experiments are (5×5), (10×10), and (15×15). By varying the size of the network and the number of blocked roads, we test the performances of both optimal and heuristic approaches regarding instance size.

Each structured instance has a single supply node, from which the clearance process starts, and there is no capacity for supply. For grid networks, the supply node is on the upper left corner, whereas for ring graphs, it is in the center. In further experiments, we randomly vary the location of the supply node, and also conduct experiments with two supply nodes. The results for these cases are given in Appendix A and are in line with those presented in this section. Demand on each node is randomly generated based on a uniform distribution with parameters $U(0, 5)$. Each unit

Table 2: Number of blocked edges for each instance

		% roads blocked		
Problem set		**25**	**50**	**100**
Random instances	$G(5 \times 5)$ and $R(5 \times 5)$	***	***	40
	$G(10 \times 10)$ and $R(10 \times 10)$	45	90	180
	$G(15 \times 15)$ and $R(15 \times 15)$	105	210	420
Earthquake scenario		151	302	604

of satisfied demand accrues one unit of benefit per period. The percentage of roads blocked is set to 25%, 50%, and 100%, resulting in instance sizes given in Table 2, where the blocked edges are chosen randomly and their number varies from 40 to 420.

The clearance resource requirements are generated based on a uniform distribution with parameters $U(0, 9)$, $U(0, 5)$, and $U(0, 1.2)$ for the cases with network sizes of (5×5), (10×10), and (15×15), respectively. These settings yield approximately equal total expected debris amounts for each instance. The clearance capacity is 30 units per period and there are eight, six, and four periods for the cases with 100%, 50%, and 25% blocked roads, respectively.

The initial beliefs of the decision maker on the clearance resource requirements on unreachable roads are constructed as follows: There are two possible beliefs regarding the resource requirement for each road. The first is the actual one, whereas the second is determined by multiplying the actual requirement by a random variable for each road. For half of the instances, the random variable follows a uniform distribution with parameters $U(0.5, 2)$, whereas for the remainder it follows a uniform distribution with parameters $U(0, 1.5)$. The corresponding initial beliefs, i.e., probabilities corresponding to each possible requirement level, are p_{ij} and $1 - p_{ij}$, respectively. To analyze the effect of belief accuracy, we let p_{ij} follow a uniform distribution with one of the following sets of parameters: $U(0, 0.5)$, $U(0, 1)$, $U(0.25, 0.75)$, and $U(0.5, 1)$. Note that a higher mean for the distribution implies higher *accuracy* of beliefs. Furthermore, when the mean is closer to zero or one, the *certainty* in beliefs is higher, regardless of the accuracy.

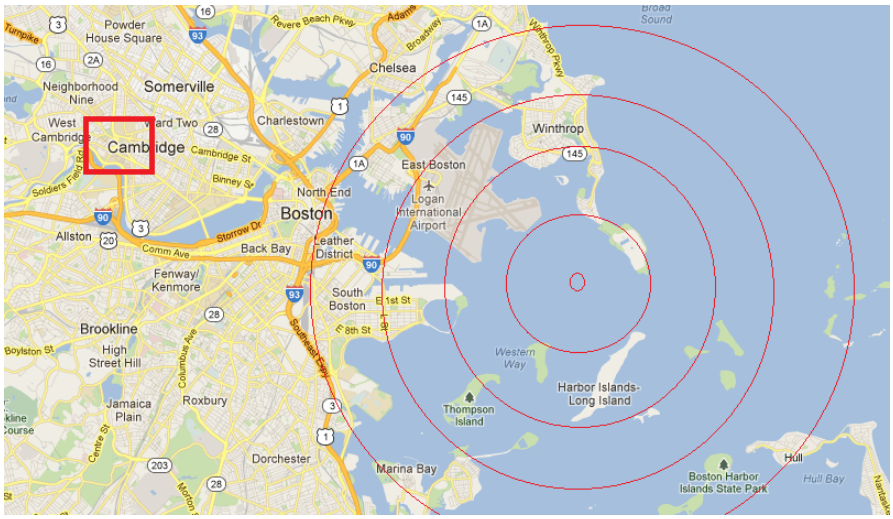

Figure 14: The epicenter of the earthquake scenario and the focus area

For each network structure and parameter settings, we generate 50 instances by varying the demand, blocked edges, debris amounts, and belief probabilities, leading to a total of 2,800 instances.

We use a second set of problem instances based on an earthquake scenario with an epicenter around 20 miles off the east coast of Boston, Massachusetts. The scenario is generated using the disaster simulation tool Hazus by [68]. The earthquake (with epicenter shown in Figure 14), is of magnitude 6.5 on the Richter scale. The disaster area we consider consists of the six census tracts in central Cambridge, MA (see Figure 14).

For this set, we assume that the clearance resource requirements are directly proportional to the debris amounts. We use a two-phase approach to set the actual debris amounts on the 604 road segments. First, we use Hazus [68] to estimate the aggregate debris amount for each census tract. Then, we determine the road segments in each census tract and distribute the total debris amount to these segments proportional to road length.

There are two supply nodes (hospitals) in the disaster region: a regional hospital

and the health center of a university, with capacities of 3,000 and 415 units, respectively, estimated proportional to their bed capacities. There are 446 demand nodes, with a total demand of 2,500 units, and the aggregate demand of each census tract is determined proportional to its population. The aggregate demand in each census tract is then distributed to the nodes in proportion to the total length of the roads emanating from the nodes. There is no additional supply/demand occurring throughout the planning horizon. Each unit of satisfied demand accrues one unit of benefit per period.

As in the case of random instances, the percentage of blocked roads is varied to be 25%, 50%, and 100%. The length of the planning horizon for these instances is three, four, and six periods, respectively. Per period, clearance capacity is 15,000 CY for the first two periods, 23,000 CY for the third and fourth periods, and 31,000 CY in the remainder of the horizon. With these capacities, more than 95% of all demand nodes can be reached in each instance.

The possible debris amounts and initial belief probabilities are determined in an identical way to those for random instances. When regional updates are considered, each census tract is taken as a separate region. For each case, we generate 50 instances by varying the choice of blocked edges, potential debris amounts, and initial belief probabilities, leading to a total of 600 instances in our second set of experiments.

2.6.2 Benchmark Approaches

To assess if and when the approaches we propose for SDCP are beneficial, we compare them to simpler approaches that ignore the uncertainty and/or apply predetermined road prioritization. The proposed and benchmark approaches are summarized in Table 3.

To assess the effect of uncertainty in our instances, we solve each instance under perfect information on the clearance resource requirements. One way to decrease the

Table 3: A summary and comparison of proposed and benchmark heuristics, with ⋆ indicating problem sizes solvable up to 40 blocked roads; ⋆ ⋆ up to 100 blocked roads, ⋆ ⋆ ⋆ up to 200 blocked roads, and ⋆ ⋆ ⋆ ⋆ more than 200 blocked roads

		Periodic nature	**Stochasticity**	**Instance size**	**Network structure**	**Regional updates**
Perfect information	Optimal	✓	✓	⋆	✓	✓
	Heuristic	✓	✓	⋆ ⋆ ⋆ ⋆	✓	✓
Deterministic with expected amounts	Optimal	✓		⋆	✓	
	Heuristic	✓		⋆ ⋆ ⋆ ⋆	✓	
Predetermined road prioritization	DPD			⋆ ⋆ ⋆ ⋆		
	DPDW			⋆ ⋆ ⋆ ⋆		
	MFP	✓		⋆ ⋆ ⋆ ⋆		
Stochastic approach	POMDP-C	✓	✓	⋆	✓	✓
Continuous approximation	POMDP		✓	⋆ ⋆	✓	✓
	HP		✓	⋆ ⋆ ⋆	✓	✓
	LL		✓	⋆ ⋆ ⋆ ⋆	✓	✓
	HP & LL		✓	⋆ ⋆ ⋆ ⋆	✓	✓

complexity due to the uncertainty is to solve a deterministic model using expected debris amounts on each road segment. For small- to medium-sized instances, the deterministic problem can be solved to optimality, whereas for larger instances, we use the heuristics developed by Stilp et al. [149]. All of the benchmark heuristics assume a deterministic environment with expected values as debris amounts.

Currently, to the best of our knowledge, there is no systematic procedure in practice for how the clearance activities are sequenced. Guidelines such as FEMA [67] suggest clearing roads of higher convenience with regard to disaster response operations (main arteries, evacuation routes, routes for care facilities, shelters, etc.) first, and clearing roads of lesser importance later, possibly during the recovery stage. In the Hazus database [68], there are mainly four types of roads, and each road is assigned a predetermined weight from 1 (most important) to 4 (least important). Figure 15 shows the road network and the pre-assigned road priorities for the earthquake scenario.

We develop and test three benchmark algorithms to calculate road priorities based on these road weights, mimicking the current practice. The first algorithm (*DPD*) works in a similar way to the continuous approximation in §2.4, i.e., it greedily chooses the road segment to clear one at a time, based on the highest ratio of demand-reached

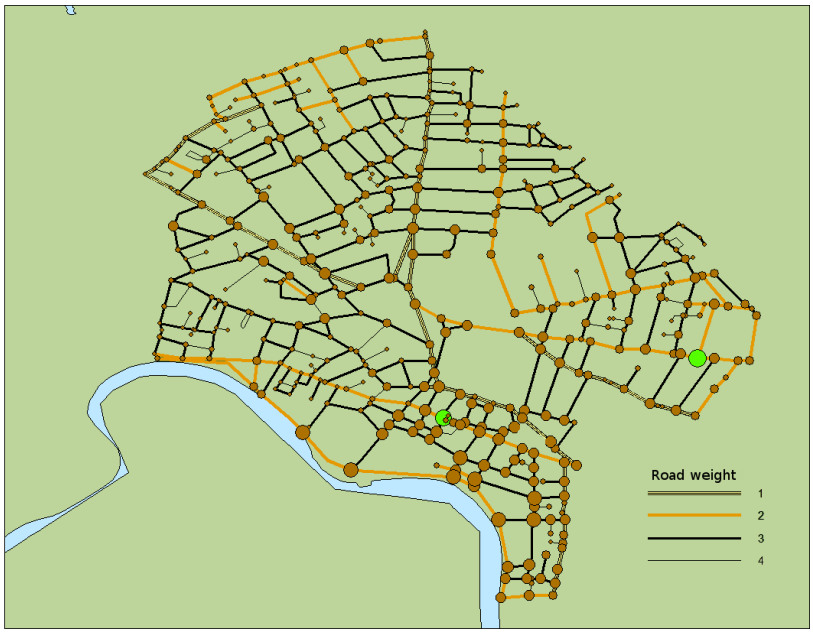

Figure 15: Road types for the earthquake scenario

to debris-amount among the immediately reachable roads with lowest weight (highest importance). The second approach ($DPDW$) is similar to the first one, except that it does not use road priorities. Upon each clearance, among immediately reachable roads, it clears the one with the highest ratio of demand-reached to debris-amount-per-road-weight. For both DPD and DPDW, the resulting solution is converted to a periodic solution using the transformation described in §2.4.2. The last approach (MFP) differs from the first two in the sense that it is applied on a periodic basis. In each period, among the roads with highest importance, it selects the roads to be cleared by maximizing the total flow at the end of the period.

2.6.3 Computational Results

In assessing the solution quality of the proposed and benchmark approaches, we consider a number of optimality gaps. Let z^*, z^p, z^u, and z^h denote the objective function values for the optimal POMDP approach, optimal/heuristic solution with perfect information, upper bound obtained by POMDP-C, and any proposed or benchmark heuristic solution, respectively. Then, we define the *optimality gap*, *upper bound gap*,

Table 4: Average upper bound gaps and heuristic gaps for POMDP-C along with CPU times (in seconds)

Distribution of p_{ij}	Network	% roads blocked	Upper bound gap	Optimality gap	CPU time for optimal	CPU time for approximation
$U(0, 0.5)$	G(5×5)	100	9.14	7.34	1,275	280
	G(10×10)	25	12.65	10.18	5,256	402
	R(5×5)	100	11.08	5.35	1,423	256
	R(10×10)	25	13.09	7.61	2,994	259
$U(0, 1)$	G(5×5)	100	8.24	7.50	1,340	277
	G(10×10)	25	13.20	11.00	5,406	447
	R(5×5)	100	12.13	6.72	1,304	236
	R(10×10)	25	12.90	8.62	3,270	268
$U(0.25, 0.75)$	G(5×5)	100	9.70	5.79	1,449	273
	G(10×10)	25	12.54	13.04	5,530	451
	R(5×5)	100	11.44	6.90	1,271	268
	R(10×10)	25	12.09	12.58	3,035	275
$U(0.5, 1)$	G(5×5)	100	10.12	8.27	1,430	294
	G(10×10)	25	12.56	12.07	6,103	462
	R(5×5)	100	11.69	7.22	1,083	204
	R(10×10)	25	11.46	12.74	3,197	237

and *perfect information gap* as $\frac{z^*-z^h}{z^*}$, $\frac{z^u-z^*}{z^*}$, and $\frac{z^p-z^h}{z^p}$, respectively. The experiments throughout this section have been conducted on a computer with 2.10 GHz Intel Core 2 Duo processor and 4 GB RAM.

Since all of our heuristics use a continuous-time approximation for SDCP, we compare the performance of POMDP-C to the optimal POMDP approach on random instances of size up to 45 blocked roads. For larger instances, POMDP fails to find a feasible solution due to reaching the memory limit. The average upper bound gaps and optimality gaps (over 50 instances) are reported in Table 4, along with the average CPU times for each approach.

First, we observe that for small instances, **upper bound and optimality gaps of POMDP-C are at most around 13% on average.** Furthermore, the variation of these gaps with changing belief structures is not statistically significant. Although the number of blocked roads is almost identical in all instances, with a fixed network type (grid or ring) and belief structure, increasing the density of the network increases the gap in almost all instances. Intuitively, when the number of roads to be cleared

per period increases, this results in a decrease in the accuracy of a continuous-time approximation (a continuous-time approximation is most accurate when there is a single road to be cleared in each period).

POMDP-C provides 86% CPU time savings on average, compared to POMDP. The average time savings range from 78% to 93%. The level of CPU time savings compared to relatively small gaps from optimal solutions underlines the significance of applying POMDP-C during the disaster response phase, where timely decision making is crucial.

Next, we test the performance of POMDP-C and the related heuristics on all structured and earthquake scenario instances. Tables 5 and 6 give the perfect information gaps for random and earthquake scenario instances, respectively. Since only small sized instances can be solved to optimality under perfect information, for larger instances we use the solutions obtained by the heuristic of Stilp et al. [149] to calculate the gaps. In Tables 5 and 6, the optimal and heuristic results are denoted by O or H, respectively. Tables 7 and 8 show the average CPU times of each approach under each setting. For brevity, Table 7 only gives the CPU times for the case where all p_{ij} follow a uniform distribution with parameters $U(0.5, 1)$. For the remaining scenarios, the CPU times are provided in Table 17 in the Appendix. A blank entry in each table indicates that the instances under the given setting cannot be solved due to insufficient memory.

The largest instances solvable by POMDP-C have 105 blocked roads. HP is able to solve instances with at most 180 blocked roads, while LL and the hybrid application of HP & LL (HP-LL) can solve all instances we generated. HP-LL also provides significant time savings, solving the largest instances within 2.5 hours on average, compared to the CPU times of HP or LL, which vary between two and eight hours.

Table 5: Average perfect information gaps of the proposed and benchmark approaches on random instances

			Perfect information	Deterministic with expected amounts		No regional updates					With regional updates				
						Periodic	Continuous approximation				Periodic	Continuous approximation			
Distribution of p_{ij}	Instance	% roads blocked	O - optimal H - heuristic	Optimal	Heuristic	POMDP	POMDP	HP	LL	HP & LL	POMDP	POMDP	HP	LL	HP & LL
$U(0, 0.5)$	$G(5 \times 5)$	100	O	41.41	44.30	22.75	24.43	25.14	25.46	26.37	23.77	24.45	25.34	26.40	28.57
	$G(10 \times 10)$	25	O	45.47	49.99	26.17	28.77	30.16	30.35	32.19	26.84	28.91	29.88	31.43	33.00
	$G(10 \times 10)$	50	H	***	56.22	***	31.41	33.09	33.23	33.94	***	31.30	32.28	33.07	33.74
	$G(10 \times 10)$	100	H	***	57.21	***	***	34.52	35.08	35.51	***	***	32.88	34.98	36.05
	$G(15 \times 15)$	25	H	***	48.29	***	28.15	30.35	30.92	32.84	***	28.27	30.97	31.20	32.58
	$G(15 \times 15)$	50	H	***	55.07	***	***	***	32.27	33.19	***	***	***	31.72	34.61
	$G(15 \times 15)$	100	H	***	63.23	***	***	***	33.68	34.93	***	***	***	33.60	34.83
	$R(5 \times 5)$	100	O	39.22	42.92	25.76	27.02	29.01	29.11	30.31	26.03	27.36	28.22	28.96	32.41
	$R(10 \times 10)$	25	O	45.52	48.62	28.48	28.91	29.37	29.93	31.49	29.40	30.54	32.47	32.57	33.81
	$R(10 \times 10)$	50	H	***	61.03	***	30.89	31.30	32.47	33.01	***	31.18	33.17	33.78	34.71
	$R(10 \times 10)$	100	H	***	64.06	***	***	33.07	34.20	35.63	***	***	32.23	34.47	35.60
	$R(15 \times 15)$	25	H	***	57.12	***	30.38	32.59	33.50	34.00	***	30.28	33.10	33.82	35.20
	$R(15 \times 15)$	50	H	***	60.96	***	***	***	35.01	36.22	***	***	***	35.78	36.93
	$R(15 \times 15)$	100	H	***	63.93	***	***	***	39.14	41.17	***	***	***	37.82	39.68
$U(0, 1)$	$G(5 \times 5)$	100	O	33.09	34.32	29.16	31.08	31.63	33.38	34.78	26.83	26.93	28.41	29.66	29.82
	$G(10 \times 10)$	25	O	37.38	38.27	33.03	36.64	37.87	39.86	41.68	30.33	32.83	33.91	34.72	36.54
	$G(10 \times 10)$	50	H	***	43.31	***	39.53	41.74	41.86	43.75	***	35.65	38.27	39.46	41.34
	$G(10 \times 10)$	100	H	***	46.05	***	***	43.19	42.47	45.89	***	***	38.80	39.29	42.37
	$G(15 \times 15)$	25	H	***	38.94	***	36.09	38.92	39.70	41.53	***	34.62	35.35	34.78	37.65
	$G(15 \times 15)$	50	H	***	43.29	***	***	***	40.92	42.65	***	***	***	37.51	40.01
	$G(15 \times 15)$	100	H	***	50.40	***	***	***	43.37	43.86	***	***	***	39.50	41.38
	$R(5 \times 5)$	100	O	33.36	34.11	30.69	33.06	34.95	34.60	35.28	29.19	30.79	32.00	34.07	35.32
	$R(10 \times 10)$	25	O	38.20	39.36	35.23	36.63	38.21	39.67	42.63	33.00	32.80	34.80	36.26	38.78
	$R(10 \times 10)$	50	H	***	48.01	***	37.84	39.51	41.53	42.53	***	34.67	35.51	37.37	39.78
	$R(10 \times 10)$	100	H	***	50.04	***	***	41.13	41.63	44.23	***	***	36.94	37.64	41.51
	$R(15 \times 15)$	25	H	***	45.40	***	39.81	41.43	42.71	45.81	***	36.19	37.49	38.24	39.97
	$R(15 \times 15)$	50	H	***	48.73	***	***	***	44.61	46.07	***	***	***	38.53	41.57
	$R(15 \times 15)$	100	H	***	49.38	***	***	***	47.70	49.93	***	***	***	43.33	43.75
$U(0.25, 0.75)$	$G(5 \times 5)$	100	O	33.09	34.32	25.43	27.19	29.84	30.94	31.24	23.20	25.16	26.73	27.92	30.19
	$G(10 \times 10)$	25	O	37.38	38.27	28.91	32.78	34.70	35.51	36.79	26.31	28.88	31.17	33.48	33.99
	$G(10 \times 10)$	50	H	***	43.31	***	35.52	35.88	37.45	37.64	***	32.78	33.79	33.52	34.21
	$G(10 \times 10)$	100	H	***	46.05	***	***	37.87	38.22	39.09	***	***	34.45	35.19	35.48
	$G(15 \times 15)$	25	H	***	38.94	***	31.70	33.22	34.01	36.23	***	29.59	30.64	31.15	32.63
	$G(15 \times 15)$	50	H	***	43.29	***	***	***	37.65	38.07	***	***	***	34.10	34.41
	$G(15 \times 15)$	100	H	***	50.40	***	***	***	37.40	37.97	***	***	***	33.90	35.94
	$R(5 \times 5)$	100	O	33.36	34.11	28.03	29.87	31.60	30.61	31.80	25.67	28.12	27.91	29.55	31.43
	$R(10 \times 10)$	25	O	38.20	39.36	30.99	33.01	33.45	34.39	34.67	30.31	31.41	31.72	31.87	33.07
	$R(10 \times 10)$	50	H	***	48.01	***	33.69	34.21	34.44	35.94	***	32.47	32.60	32.27	34.03
	$R(10 \times 10)$	100	H	***	50.04	***	***	34.59	35.86	35.99	***	***	32.76	33.31	34.45
	$R(15 \times 15)$	25	H	***	45.40	***	35.25	36.89	37.02	38.09	***	32.50	33.68	34.54	35.82
	$R(15 \times 15)$	50	H	***	48.73	***	***	***	37.93	39.44	***	***	***	34.84	37.21
	$R(15 \times 15)$	100	H	***	49.38	***	***	***	44.28	45.89	***	***	***	37.69	39.87
$U(0.5, 1)$	$G(5 \times 5)$	100	O	27.02	29.60	19.37	20.80	21.35	21.46	21.80	18.95	19.25	20.14	20.99	20.58
	$G(10 \times 10)$	25	O	30.96	34.14	21.13	21.52	22.02	22.27	22.48	21.04	22.03	22.92	22.49	23.35
	$G(10 \times 10)$	50	H	***	37.25	***	24.52	25.07	25.73	25.60	***	23.49	23.77	24.33	24.35
	$G(10 \times 10)$	100	H	***	38.69	***	***	27.25	27.36	27.51	***	***	25.39	25.68	26.41
	$G(15 \times 15)$	25	H	***	33.55	***	23.95	25.49	25.62	25.62	***	22.80	24.16	24.36	24.66
	$G(15 \times 15)$	50	H	***	35.74	***	***	***	26.77	26.86	***	***	***	25.20	25.50
	$G(15 \times 15)$	100	H	***	39.84	***	***	***	27.82	27.83	***	***	***	25.96	26.69
	$R(5 \times 5)$	100	O	28.74	30.37	21.88	23.30	23.19	23.29	23.45	20.94	22.19	23.58	23.52	23.56
	$R(10 \times 10)$	25	O	31.58	33.15	22.98	24.81	24.99	25.21	25.54	22.72	24.58	25.06	25.74	25.69
	$R(10 \times 10)$	50	H	***	39.12	***	24.90	25.00	25.77	26.43	***	24.48	25.70	25.94	26.12
	$R(10 \times 10)$	100	H	***	40.72	***	***	26.12	26.44	26.70	***	***	26.33	26.67	26.56
	$R(15 \times 15)$	25	H	***	37.13	***	26.27	27.61	28.02	28.03	***	24.84	26.09	26.69	26.49
	$R(15 \times 15)$	50	H	***	40.86	***	***	***	28.70	29.01	***	***	***	26.75	27.24
	$R(15 \times 15)$	100	H	***	42.32	***	***	***	30.51	30.81	***	***	***	28.96	29.18

Table 6: Average perfect information gaps of the proposed and benchmark approaches on earthquake scenario instances

		Perfect information	Deterministic with expected amounts		Predetermined road priorities			No regional updates					With regional updates				
								Periodic	Continuous approximation				Periodic	Continuous approximation			
Distribution of p_{ij}	% roads blocked	O - optimal H - heuristic	Optimal	Heuristic	DPD	DPDW	MFP	POMDP	POMDP	HP	LL	HP & LL	POMDP	POMDP	HP	LL	HP & LL
$U(0, 0.5)$	25	H	***	49.85	61.58	59.10	51.22	***	***	25.47	28.38	31.29	***	***	23.82	25.58	27.52
	50	H	***	52.04	66.69	64.54	60.46	***	***	***	33.52	37.16	***	***	***	29.57	32.41
	100	H	***	57.56	75.18	73.61	68.24	***	***	***	39.24	42.90	***	***	***	32.65	36.13
$U(0, 1)$	25	H	***	38.20	56.15	54.81	49.81	***	***	31.35	33.75	35.99	***	***	25.45	27.78	29.12
	50	H	***	45.68	58.93	56.65	52.15	***	***	***	42.46	45.01	***	***	***	34.67	37.14
	100	H	***	47.36	64.82	61.31	58.84	***	***	***	47.51	49.12	***	***	***	36.30	38.17
$U(0.25, 0.75)$	25	H	***	38.20	56.15	54.81	49.81	***	***	26.86	29.53	32.50	***	***	26.38	28.35	31.08
	50	H	***	45.68	58.93	56.65	52.15	***	***	***	36.76	39.90	***	***	***	30.18	33.21
	100	H	***	47.36	64.82	61.31	58.84	***	***	***	38.90	41.56	***	***	***	31.62	34.70
$U(0.5, 1)$	25	H	***	33.72	51.86	48.68	42.69	***	***	20.13	21.10	21.74	***	***	20.19	21.34	22.11
	50	H	***	35.23	58.55	55.02	44.30	***	***	***	21.30	23.71	***	***	***	23.40	26.00
	100	H	***	39.14	61.74	58.52	53.11	***	***	***	25.12	27.71	***	***	***	26.67	29.93

Table 7: Average CPU times (in seconds) for the proposed and benchmark approaches on random instances

		Deterministic with perfect information		Deterministic with expected amounts		No regional updates					With regional updates				
						Periodic	Continuous approximation				Periodic	Continuous approximation			
Instance	% roads blocked	Optimal	Heuristic	Optimal	Heuristic	POMDP	POMDP	HP	LL	HP & LL	POMDP	POMDP	HP	LL	HP & LL
$G(5 \times 5)$	100	26.9	5.6	25.8	3.9	1,430.1	294.4	175.1	229.4	69.3	1,537.7	313.7	188.5	241.1	72.4
$G(10 \times 10)$	25	78.6	12.9	72.8	11.5	6,103.2	462.3	238.2	314.4	92.7	6,221.5	481.5	249.6	323.8	97.3
$G(10 \times 10)$	50	***	36.8	***	20.8	***	2,262.4	1,119.8	1,433.7	429.6	***	2,326.4	1,173.4	1,513.8	446.6
$G(10 \times 10)$	100	***	101.4	***	76.7	***	***	1,990.9	2,645.7	783.2	***	***	2,114.9	2,786.9	817.7
$G(15 \times 15)$	25	***	25.9	***	25.1	***	482.4	241.2	321.7	96.0	***	499.1	250.7	339.7	101.6
$G(15 \times 15)$	50	***	66.8	***	28.6	***	***	***	2,095.7	627.7	***	***	***	2,236.5	674.2
$G(15 \times 15)$	100	***	132.5	***	129.5	***	***	***	3,900.5	1,175.2	***	***	***	4,097.9	1,200.5
$R(5 \times 5)$	100	34.8	6.1	29.5	4.2	1,083.2	203.8	118.4	152.7	46.0	1,166.4	212.2	123.6	159.3	47.9
$R(10 \times 10)$	25	72.4	14.0	63.9	11.2	3,196.5	237.3	132.4	178.5	82.3	3,294.5	252.4	138.1	188.4	85.6
$R(10 \times 10)$	50	***	42.4	***	26.6	***	2,017.8	950.9	1,256.8	372.8	***	2,096.2	1,008.9	1,335.5	387.5
$R(10 \times 10)$	100	***	116.1	***	88.4	***	***	1,800.4	2,345.2	701.4	***	***	1,850.9	2,437.2	732.9
$R(15 \times 15)$	25	***	28.0	***	13.9	***	394.7	202.7	274.0	81.3	***	409.7	213.7	288.0	83.9
$R(15 \times 15)$	50	***	70.5	***	70.8	***	***	***	2,009.7	604.5	***	***	***	2,099.7	652.6
$R(15 \times 15)$	100	***	147.8	***	129.3	***	***	***	3,467.8	1,049.7	***	***	***	3,639.5	1,116.3

Table 8: Average CPU times (in seconds) for the proposed and benchmark approaches on earthquake scenario instances

		Deterministic with perfect information		Deterministic with expected amounts		Predetermined road priorities			No regional updates					With regional updates				
									Periodic	Continuous approximation				Periodic	Continuous approximation			
Distribution of p_{ij}	% roads blocked	Optimal	Heuristic	Optimal	Heuristic	DPD	DPDW	MFP	POMDP	POMDP	HP	LL	HP & LL	POMDP	POMDP	HP	LL	HP & LL
$U(0, 0.5)$	25	***	126	***	139	2.6	2.7	104	***	***	20,431	8,503	1,980	***	***	***	8,612	2,088
	50	***	249	***	274	3.8	3.6	205	***	***	***	16,657	3,529	***	***	***	16,742	3,663
	100	***	333	***	366	6.5	6.5	285	***	***	***	23,980	7,731	***	***	***	24,097	7,832
$U(0, 1)$	25	***	113	***	124	2.6	2.3	102	***	***	19,931	8,203	1,942	***	***	***	8,257	2,018
	50	***	204	***	224	3.9	3.9	160	***	***	***	15,294	3,971	***	***	***	15,374	4,044
	100	***	340	***	375	6.4	6.1	160	***	***	***	25,302	7,712	***	***	***	25,426	7,786
$U(0.25, 0.75)$	25	***	138	***	152	2.6	2.5	114	***	***	20,317	7,709	2,171	***	***	***	7,784	2,293
	50	***	303	***	333	4.2	3.9	244	***	***	***	16,256	4,203	***	***	***	16,350	4,327
	100	***	472	***	519	6.3	6.4	437	***	***	***	26,046	7,803	***	***	***	26,158	7,915
$U(0.5, 1)$	25	***	173	***	190	2.7	2.7	144	***	***	20,481	7,439	2,350	***	***	***	7,549	2,415
	50	***	365	***	401	4.2	4.3	293	***	***	***	15,540	4,753	***	***	***	15,622	4,877
	100	***	704	***	774	6.3	6.4	625	***	***	***	28,952	8,781	***	***	***	29,059	8,258

Table 9: Summary of the effects of belief structure on solution quality

	Belief variance	
Beliefs	**Low**	**High**
Accurate, certain	Accuracy and certainty of information yields good solution quality	High variance results in more need for exploration and therefore a loss of quality in solutions
Uncertain	Significant time spent for exploration, resulting in loss of benefits	
Inaccurate, certain	Inaccuracy corrected as new information is obtained	

Without regional updates, the proposed heuristics provide an average of more than 86% CPU time savings at the expense of at most 20% solution quality loss. For random instances, when the instances can be solved to optimality, the average perfect information gaps of POMDP-C, HP, LL, and hybrid approaches are 7.6%, 12.6%, 15.1%, and 19.4%, respectively, with average computational time savings of 86.5%, 92.2%, 89.8%, and 96.7%, respectively over the optimal POMDP model. In particular, HP-LL takes less than 93 seconds on average, POMDP takes around 102 minutes. Furthermore, the performance of the proposed heuristics are robust with respect to the uncertainty in the environment, as the heuristic performance does not depend on the distribution of the belief parameters.

Proposed heuristics perform better when beliefs are more certain (and not necessarily accurate). For the random instances solved by HP & LL with no regional updates, the average perfect information gaps for p_{ij} uniformly distributed with parameters $U(0,5)$, $U(0,1)$, $U(0.25,0.75)$, and $U(0.5,1)$ are 33.6%, 42.9%, 37.1%, and 26.3%, respectively. For the earthquake scenario instances, the gaps are 37.1%, 43.4%, 38.0%, and 24.4%, respectively. Hence, the best results are obtained when beliefs are accurate, i.e., distribution mean is higher. When beliefs are inaccurate but certain (distribution mean close to zero), the gaps are better compared to the cases where they are uncertain (distribution mean centered around 0.5). A summary of findings related to the effects of belief structure on solution quality is given in Table 9.

Regional updates improve solution quality when beliefs are uncertain (belief distribution mean is around 0.5). For random instances with regional updates, the average perfect information gaps are 32.3%, 39.3%, 34.5%, and 25.3%, for p_{ij} following a uniform distribution with parameters $U(0, 5)$, $U(0, 1)$, $U(0.25, 0.75)$, and $U(0.5, 1)$, respectively. For the earthquake scenario, the average perfect information gaps for the same belief structures are 32.0%, 34.8%, 33.0%, and 26.0%. These results imply that regional updates are most beneficial when initial beliefs are uncertain, since they provide a mechanism for more effective exploration.

Without regional updates, when belief certainty is higher, proposed heuristics outperform the deterministic approach. When debris amounts are assumed to be deterministic and set to their expected values, the heuristic approach by Stilp et al. [149] solves all the instances in no more than 13 minutes on average. However, the average perfect information gaps with p_{ij} following a uniform distribution with parameters $U(0, 5)$, $U(0, 1)$, $U(0.25, 0.75)$, and $U(0.5, 1)$ for the random instances are 54.2%, 43.4%, 43.4%, and 36.0%, respectively, whereas the gaps are 53.2%, 43.7%, 43.7%, and 36.0%, respectively for the earthquake scenario instances. The hybrid heuristic outperforms the deterministic approach for the first and last cases, that is, when beliefs are certain. For the first case, the deterministic approach performs poorly due to inaccurate beliefs. For the last, while both approaches perform well due to accurate beliefs, the stochasticity in the problem environment (e.g., roads that actually have the less likely amount of debris) is captured better by the proposed heuristics. For the remaining belief settings, the average performances of the deterministic approach and the proposed heuristics are close, and the deterministic approach can find better solutions. However, **when regional updates are incorporated, proposed heuristics always outperform the deterministic approach**. A practical implication of these results is the insight that it is worth using a stochastic approach when beliefs are certain. **If there is regional information, using**

a stochastic approach results in significant improvement. When beliefs are uncertain and regional information is not available, the deterministic approach with expected resource requirements can be used.

Proposed heuristics outperform predetermined road prioritization. The results in Table 8 indicate that for our instances, DPD and DPDW approaches find solutions within seven seconds on average, whereas MFP can take up to ten minutes. However, there is a significant difference between the gaps of these approaches and proposed heuristics from perfect information. For belief distributions of $U(0, 5)$, $U(0, 1)$, $U(0.25, 0.75)$, and $U(0.5, 1)$, the average gaps for DPD are 67.8%, 60.0%, 60.0%, and 57.4%, respectively. For the same belief settings, the gaps are 65.8%, 57.6%, 57.6%, and 54.1% for DPDW, and 60.0%, 53.6%, 53.6%, and 46.7% for MFP, respectively. **These results underline the significance of considering road priorities as dynamic decisions rather than taking them as predetermined inputs.**

Due to its myopic nature, DPD prioritizes the clearance of high-importance roads incident to low-demand nodes over low-importance roads incident to high-demand nodes. This may lead to significant loss of benefits. An example for this is given in Figure 16, where the high demand nodes in the south of the region are reached later than the low demand ones located north and west of the region. Similarly, the myopic nature of DPDW may require significant time to reach high-demand nodes that are not easily reachable from roads of high importance. Such a case is shown in Figure 17, where the low demand nodes located close to high demand nodes in the west are favored over high demand nodes in the south, which are secluded as they can only be reached using low importance roads.

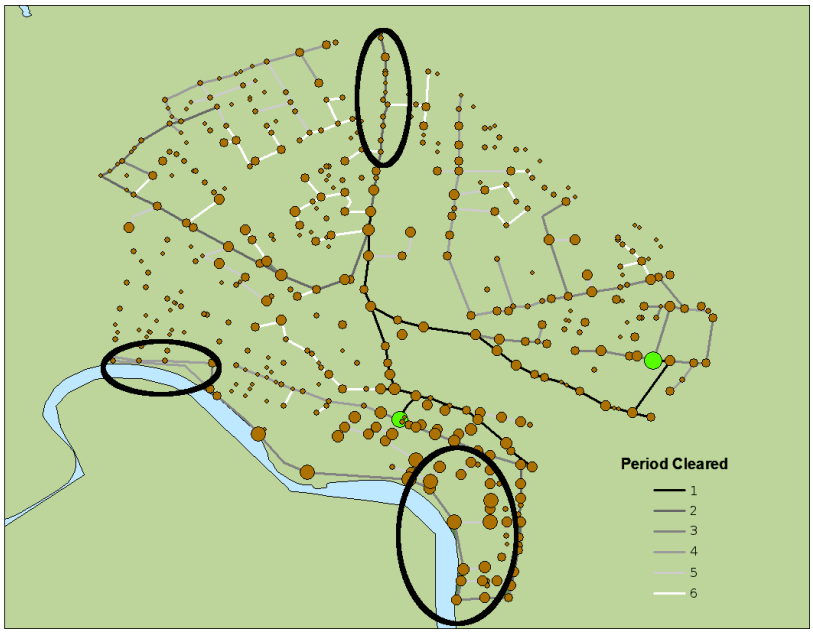

Figure 16: Clearance sequence for DPD on an earthquake scenario instance with 604 blocked roads. For DPD, high demand nodes on low priority roads and low demand nodes on main roads decrease the performance of the approach.

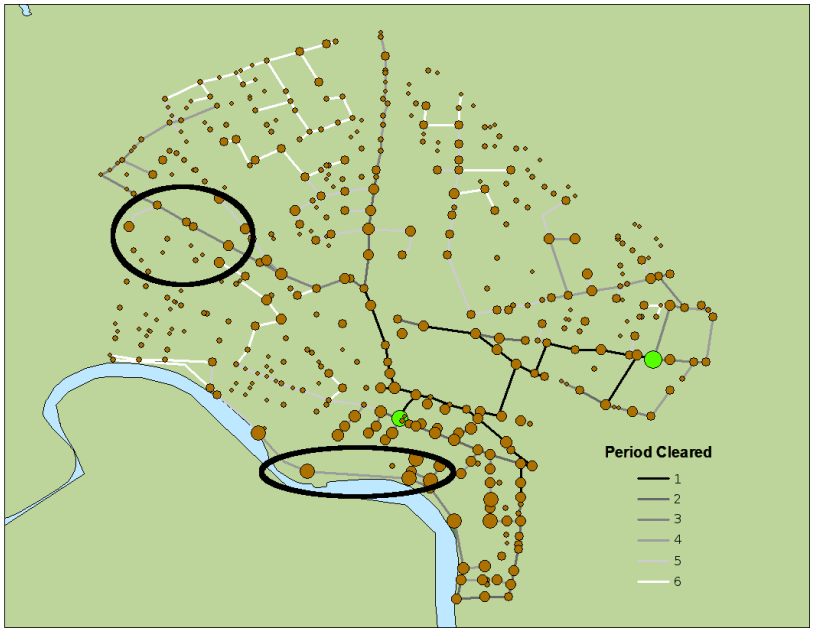

Figure 17: Clearance sequence for DPDW on an earthquake scenario instance with 604 blocked roads. For DPDW, clusters of high and low demand nodes and secluded high demand nodes decrease the performance of the approach

2.7 Conclusions and Further Research Directions

In this paper, we studied the debris clearance problem under incomplete information on clearance resource requirements, a practical problem faced especially in post-disaster situations where information on the disaster region is scarce. While a POMDP approach can solve the problem optimally, it does so with significant computational burden. Our proposed heuristics reduce the sizes of the action and observation spaces by employing a continuous-time approximation for the POMDP and applying heuristic pruning and/or limited look-ahead on the search tree. All of our approaches allow for updating information on a regional basis.

Our computational experiments on random instances as well as an earthquake scenario reveal that considering the stochasticity of the problem is useful especially when beliefs on the possible debris amounts are certain, regardless of accuracy. For the remaining cases, application of regional updates always improves the quality of the solutions. Furthermore, the current practice of predetermined road prioritization or ignoring stochasticity of the disaster environment may significantly reduce solution quality.

While deciding on road clearance prioritization, one can also consider the travel times of clearance units. This requires incorporating routing decisions and hence significantly increases the complexity of the problem. When travel times are considered, road clearance sequences for individual clearance teams might also need to determined.

SDCP addresses decisions in an environment where computational resources may be limited and decisions need to be made quickly. An important research direction is to further improve the computational time required for the solution of the problem. Another practical concern is the usability of the models by debris management authorities by means of user-friendly decision tools, which is addressed in a parallel work [107].

Finally, future work could focus on the problems of clearance, collection, disposal, and recycling in a more holistic way, although this could lead to substantial complexity.

CHAPTER III

PERISHABLE COMMODITY ALLOCATION IN HUMANITARIAN SUPPLY CHAINS WITH MULTIPLE DEMAND TYPES: EFFECTS OF USAGE POLICY AND CENTRALIZATION

3.1 Introduction

In many health and humanitarian supply chains, beneficiaries might need the commodities at different levels of urgency. For example, each patient receiving donated blood units have varying levels of need, as in the case of emergency patients with class IV hemorrhage versus those receiving routine elective surgeries. In certain cases, such as vaccination campaigns, groups including children under a certain age, elderly people, and pregnant women can be determined as prioritized groups, since they are more vulnerable to the disease [38]. When supply of the commodity is scarce, this differentiation of demand classes, along with the uncertainty of demand, pose an important challenge on the allocation of the commodity to the locations that distribute these commodities.

Allocations of commodities in humanitarian supply chains in developing countries are usually made from a central warehouse. For the case of donated blood distribution, this may be a regional collection center, which, after a donation campaign, has a certain inventory of blood units it needs to allocate to local hospitals that are assigned to it a priori. Due to the perishability of the commodity and/or because of the lack of appropriate transportation infrastructure, reallocation of the commodity may not be possible, which further increases the emphasis on efficient allocation.

In this chapter, we consider the problem of allocating a perishable commodity from a capacitated warehouse to multiple distribution locations facing demand from multiple classes. While our focus is on items to be distributed in health and humanitarian supply chains with emphasis on developing country settings, our models are also applicable in for-profit supply chains with price-differentiated customer types and perishable items (e.g., fashion industry, newspaper or magazine sales). Demand satisfied from each type accrues a certain level of unit benefit. The item that is distributed either has a too short usable life (e.g., daily newspapers or blood platelets) to allow for redistribution, or there is lack of transportation or information infrastructure (e.g., distribution of health items in developing countries), which deems redistribution impossible. Since supply of the commodity is scarce, any allocated commodity that is unused at the end of the planning horizon incurs a certain benefit loss. The demand from each class is stochastic with distributions that are known by the warehouse, and are independent of each other. Furthermore, each distribution location applies a certain policy that determines in which order to satisfy the demand from different classes, which in turn affects the allocation policy at the warehouse. The aim is to determine the optimal allocation policy based on the demand distributions and usage policies so that the expected net benefit is maximized.

The distribution locations may follow different policies to satisfy the demand from multiple classes. One way is to prioritize higher-class demand over the lower classes. This can be implemented by satisfying the demand of the highest class in a certain period, the second highest class in the next period, and so on. Such applications of prioritization of higher-class demand can also be observed in for-profit supply chains, such as the fashion apparel industry, where products that are not sold to the higher-class customers are later sold to regular customers at discounted prices. However, especially in health and humanitarian settings, such prioritization schemes may not

be as easy to apply. For instance, for the 2009 vaccination campaign, although priority groups were determined by the Centers for Disease Control and Prevention (CDC), lack of sub-prioritization schemes among the priority groups resulted in non-standard procedures throughout the country, including demand satisfaction on a first-come, first-served basis without regard to prioritization levels [163]. In this chapter, we particularly focus on prioritization and first-come, first-served usage at the distribution locations.

Aside from determining the optimal allocation schemes under prioritized and first-come, first-served usage policies at the distribution location level, we also aim to determine the structure of these allocations. In particular, we analyze the parameters that influence the allocation scheme under each of the optimal strategies, and how these two usage policies perform with respect to each other in extreme cases and in more realistic scenarios.

One way to overcome the adverse effects of demand variability is to aggregate the demand by centralizing the distribution locations. We show both analytical and computational results for the effects of centralization on the performances of the two usage policies. Most of our analysis builds on the assumption that the warehouse has complete information about the usage policies at the distribution locations. In addition, using computational experiments, we show the effects of information asymmetry between the usage policy at the distribution locations and the assumed policy to make the allocations at the warehouse level on the overall expected benefits. The effects of usage policies and centralization on the equity of distributions are also analyzed computationally.

The remainder of this chapter is organized as follows: We present a review of the related literature, particularly on the newsvendor problem with multiple locations and/or demand classes, inventory rationing, blood or perishable product supply chains, and pooling in §3.2. In §3.3, we formally define the problem and present the

models that find the optimal allocation policies under prioritization and first-come, first-served usage. Analytical results on these optimal allocation schemes are given in §3.4. This is followed by the computational experiments in §3.5, which use a data set based on the 2009 H1N1 vaccination campaign in the state of Georgia. We conclude the chapter and discuss potential research directions in §3.6.

3.2 Literature Review

The work we present in this chapter is related mainly to four streams of work, namely newsvendor models in which multiple locations and/or multiple demand types exist, inventory allocation/rationing and resource allocation in decentralized supply chains, perishable product supply chains, and pooling of resources in supply chains and queues.

Since allocation decisions can be made only once and reallocations are not possible, the problem considered in this chapter relates to the well-known newsvendor problem. Due to its theoretical importance and wide application areas, the newsvendor problem has attracted a vast amount of literature. Khouja [98] and Qin et al. [133] provide detailed literature reviews on this problem, while Petruzzi and Dada [126] present a review focusing on the pricing aspects. Our particular focus is on the cases where multiple locations and/or multiple demand classes exist.

Within the context of the newsvendor problem with multiple customer types, Khouja [96] addresses the problem with decreasing selling price over time, with the demand under each price as a function of demand on the initial price. Khouja [97] extends this model to the case of multiple discounts at the supplier and retailer level. With decreasing selling costs, the problem resembles the one we consider under prioritization. However, the existence of perfect correlation among the demand classes significantly changes the structure of the optimal solution. The most relevant work in this stream to the one discussed in this chapter is due to Şen and Zhang [140], in which

multiple demand classes with independent demands are served from an uncapacitated single distribution location at decreasing or increasing prices. The former case is similar to our problem, with the difference that there is a unit production/purchase cost at the warehouse, rather than an overage cost when the product is not sold. For a single location under prioritization, we derive a similar optimal allocation scheme to that by Şen and Zhang [140]. We also extend this approach to first-come, first-served usage policy at the distribution locations, as well as to the case of multiple distribution locations under a capacitated warehouse. Other studies on the newsvendor problem with multiple demand types include Bassok et al. [16] and Shumsky and Zhang [145], where multiple products are involved with full or limited substitutability. These studies consider the allocation of limited capacity to customer classes, rather than to distribution locations. In our case, the allocation problem is exacerbated by the fact the allocation needs to be made to multiple locations, which further serve multiple demand types.

For the newsvendor problem with multiple locations, Kouvelis and Gutierrez [100] consider a multi-period extension with two locations, non-overlapping selling seasons and differing risk, production and transportation costs. Our work mainly differs by the consideration of a single allocation period and a general number of distribution locations. Taaffe et al. [153] focus on a single-period version with multiple distribution locations with Bernoulli demand and aim to allocate a limited resource among these locations. Our work focuses on Poisson demand arrivals, which is a more widely used demand arrival assumption in health and humanitarian environments. Furthermore, in both [100] and [153], demand is satisfied on a first-come, first-served basis. We also extend our results to the cases where demand classes are prioritized.

The second stream of research relevant to our work is the problem of allocating capacity or inventory among multiple product/customer classes with stochastic demand, also called the *rationing* problem. Generally, the studies in this stream prove

optimality of simple-to-apply policies such as order-up-to or threshold. Some examples include Balakrishnan et al. [14] for make-to-order (MTO) systems with two demand classes, Ha [81] for a single-item make-to-stock (MTS) system with Poisson demand and exponential service times, Ha [82] for a MTS system with two classes, Poisson demand, and Erlang-k service times, De Véricourt et al. [53] for a general number of classes, Duran et al. [61] for a two-class customer system under prioritization, and Gayon et al. [76] for a MTS system with Erlang-k service times considering interruption or restarting of service. While allocating a scarce capacity to multiple demand classes, these studies assume an infinite planning horizon, a different cost mechanism, and inclusion of service rates. These assumptions allow simple inventory control policies, such as using order-up-to levels, to be applicable. Focusing on a single allocation decision results in a significantly different allocation scheme that depends on marginal benefits of the distribution locations. A similar stream considers the allocation of a product to multiple retailers in a decentralized two-echelon supply chain. At the beginning of each period, a supplier makes inventory replenishment decisions and allocates on-hand stocks to the retailers. Examples for this include Bernstein and Federgruen [24] with deterministic demand and competing retailers, Bernstein and Federgruen [25] for the case of stochastic demand, and Benjaafar et al. [20] under congestion. The main focus for this stream of research is on finding a coordinating mechanism for the supply chain, rather than an efficient allocation scheme.

A third relevant stream considers the allocation of perishable products to distribution locations. For a literature review on the supply chains of perishable products, the interested reader is referred to Karaesmen et al. [94]. Our particular focus in this stream is perishable commodities in health and humanitarian supply chains, of which the most widely studied commodity is donated blood units. A recent review on blood supply chain management is given by Bëlien and Farcé [17]. The seminal work

in this stream is due to Prastacos [130], where a central distribution system with a single warehouse and multiple depots is considered with possibility of redistribution, focusing on a single period. Prastacos [131] extends this work by observing that the 2-period optimal policy is not significantly different from that of a single period. In a recent work, Qin et al. [132] further extend this study to include transshipments and redistribution from the depot. Federgruen et al. [69] consider distribution of a perishable product involving transportation decisions and Deniz et al. [56] propose a discrete-time model for a supply chain with demands for different ages of the product. Carden and DelliFraine [30] perform a qualitative study on the benefits and disadvantages to centralization in blood centers. One important aspect of perishable product systems, common to all the aforementioned studies except [130] and [131], is product age. With the assumption that the perishable item is fresh in the beginning of the period, and due to the fact that the system we consider does not allow redistribution, we do not consider the age of the items in this study. Furthermore, we assume that the planning horizon is shorter than the useful life of the item. Hence, any demand satisfied during the period is within the useful life of the item. As opposed to the studies in this stream, considering a simpler setting allows finding the optimal allocation scheme under different usage policies rather than resorting to heuristics.

The consideration of the effect of centralization in this chapter deems the literature on pooling relevant. In particular, we focus on centralization in supply chains and pooling in queues. For supply chains, Eppen [65] has shown that under normal demand and identical holding and penalty costs, a centralized system reduces costs compared to its decentralized counterpart. Stulman [150] extend this to the case of Poisson demand and when service level requirements exist. The centralization results we obtain in this chapter further extend these to the case of multiple demand classes and different demand satisfaction policies. Other research on centralization in supply chains includes Gazis [77] for a linear program and structural results for the overall

supply chain, Gerchak and He [79], who show that benefits from centralization are directly proportional to demand variability, Benjaafar et al. [19] on the benefits of centralization when multiple products and load-dependent lead times are present, and Demirağ and Swann [55] for a sea cargo network application.

A second dimension of interest in the pooling literature is the pooling of queues. The first well-known result in this area is the fact that an $M/M/1$ queue with a service rate of $m\mu$ performs better than an $M/M/m$ queue with a service rate of μ for each server if arrival rates are the same. However, this result does not extend to general queues. Benjaafar [18] provides performance bounds for various queueing disciplines and pooling scenarios. More recent work in this area includes Argon and Andradóttir [9], who consider partial pooling in tandem lines and derive conditions and settings where pooling improves the system, Tekin et al. [155] on cross-training in call centers, and van Dijk and van der Sluis [164, 165] on call centers with various queueing and demand satisfaction settings. While inclusion of service times deem queue pooling structurally different from the problem we consider, certain analogies can be driven for the two systems in that the more variability there is for the servers in the former and the demand rates in the latter, the more beneficial centralization is.

The contributions of the work in this chapter can be summarized as follows: First, we derive an optimal allocation algorithm in a newsvendor setting with multiple locations and demand types and a capacitated warehouse. Furthermore, to the best of our knowledge, this is the first study that considers such an allocation scheme in a newsvendor setting when demand is satisfied on a first-come, first-served basis. Second, we derive structural results on the relative performances of optimal allocations under prioritization and first-come, first-served usage, and show when the differences between two usage policies are more significant. Third, to the best of our knowledge, we establish the first results on satisfying part or all of the demand from a single

distribution location in a newsvendor setting where multiple demand types and multiple locations exist, and the supplier is capacitated. Lastly, apart from providing theoretical results and bounds for the approaches we consider, we also use computational experiments on a real-life based setting to analyze the effects of information asymmetry and equity in the supply chain under the given settings.

3.3 Problem Definition and Models

We consider a two-level supply chain with a single *warehouse* in the upstream level and multiple *distribution locations* downstream. The distribution locations are indexed by the set I. Within the context of allocating blood units in a developing country, the warehouse can be considered as the blood collection center and the distribution locations as hospitals that need the blood units. In a classical newsvendor setting, the warehouse may be a regional distribution center, and the distribution locations can be newsstands or kiosks that sell the newspapers. We assume a single commodity to be distributed, which results in no loss of generality as long as the demands for the commodities are independent of each other under these settings. The commodity to be distributed either has a short shelf life (e.g., distributing blood units in a developing country setting) or is a periodical (e.g., newspapers) so that there is no time for redistribution once the commodities are allocated to the distribution locations.

The warehouse has C units of the commodity on hand to be allocated. For blood units or relief commodities, this is the number of units that have been donated. In the case where the allocation center produces the commodity, it can be interpreted as the production capacity. At the downstream level, each distribution location faces demand from multiple *demand classes*, indexed by J. When health/humanitarian resources are to be distributed, the main difference in the demand classes is in terms of how urgently the beneficiaries need the commodity. In the case of a for-profit supply chain, the demand classes are determined by the price they are willing to pay.

Each unit of demand satisfied from class j accrues w_j units of benefit. Without loss of generality, we assume $w_1 \geq w_2 \geq \ldots \geq w_{|J|}$. For each unit that is allocated to a distribution location, but is not used for satisfying demand until expiration or the end of the planning horizon, there is a benefit loss of w_0. This benefit loss can also be considered as the "opportunity cost" of the units not being assigned to another distribution location where they could actually be used. We also assume that there is no benefit loss at the warehouse for unallocated units.

We assume that the demand for each class j in each distribution location i is a random variable D_j^i that follows a distribution function $f_j^i(.)$. Based on historical demand or forecasts, it is assumed that the warehouse knows this distribution. The demand distribution of each class is assumed to be independent of each other. In the case where the demand is assumed to arrive under a Poisson process, the arrival rate for each demand class j in distribution location i is denoted by λ_j^i. Once demand starts arriving at a location, it is satisfied based on a *usage policy*. The set of alternative usage policies is denoted by K. Unless otherwise stated, the warehouse is assumed to know the usage policy of the distribution locations.

The sequence of events is given as follows:

1. The warehouse observes the amount of commodities on hand.

2. Based on the capacity, demand distributions, and usage policies at the distribution locations, the warehouse determines the amount of units (x_i) to be allocated to each distribution location i.

3. Demand from different classes is dynamically observed in each distribution location and satisfied according to the predetermined usage policies.

    - In the case where higher classes are prioritized, there are two alternative ways of how demand is observed and satisfied. In the first one, each class

Table 10: Notation

Index sets	
I	Distribution locations
J	Demand classes
K	Usage policies; $K = \{P, F\}$, where P and F denote prioritization and first-come-first-served usage, respectively

Decision variables	
x_i	Units allocated to distribution location $i \in I$
$(x_i^k)^*$	Units allocated to distribution location $i \in I$ under optimal allocation for usage policy $k \in K$

Parameters	
C	Capacity of the central allocation unit
w_j	Benefit obtained by satisfying one unit of demand of class $j \in J$
w_0	Benefit loss due to one unit of commodity assigned to, but unused at a distribution location
D_j^i	Random variable denoting units of demand class $j \in J$ of distribution location $i \in I$
$f_j^i(.)$	Probability distribution function for demand class $j \in J$ of distribution location $i \in I$
$F_j^i(.)$	Cumulative distribution function for demand class $j \in J$ of distribution location $i \in I$
λ_j^i	Rate of the probability distribution function when demand of class $j \in J$ at location $i \in I$ follows a Poisson distribution

Additional notation	
$E_k(z_i\|x_i)$	Expected net benefit at demand location $i \in I$ when it applies usage policy $k \in K$ and x_i units are allocated to it
$MB_k(x_i)$	Marginal expected net benefit obtained when $x_i + 1$ units are allocated to distribution location $i \in I$ instead of x_i units
π_k	Optimal expected total net benefit under usage policy $k \in K$ at each distribution location
π_k^{cen}	Optimal expected total net benefit under a centralized system with usage policy $k \in K$ at each distribution location
$\pi_{k_1\|k_2}$	Optimal expected total net benefit when each demand location applies usage policy $k_2 \in K$ and the allocation is made assuming usage policy $k_1 \in K$, $k_1 \neq k_2$
X	Vector of size $\|I\|$ denoting allocations to each distribution location

of demand is served in decreasing order of priority in subsequent sub-periods. In the second alternative, each distribution location takes orders while commodities are being shipped to these locations, and upon arrival demand is served based on class priority.

- In the case of first-come first-served usage, demand is served based on the order of arrivals.

4. Based on the demand satisfied during the horizon, benefits are accrued and losses are incurred based on the demand satisfied and unused units expire.

The notation we use throughout the rest of the chapter is summarized in Table 10. The expected benefit under usage policy $k \in K$ when x_i units are allocated to distribution location i is given by $E_k(z_i|x_i)$. The marginal expected benefit of allocating an additional unit to location i under policy $k \in K$ when x_i units are already allocated is denoted by $MB_k(x_i)$. In mathematical terms,

$$MB_k(x_i) = E_k(z_i|x_i + 1) - E_k(z_i|x_i).$$

Furthermore, we denote by π_k the optimal expected net benefit obtained under policy k.

Whereas the system we consider consists of multiple distribution locations, we also analyze the case where part or all of the demand classes are served from a single distribution location (a centralized system). In this case, we use π_k^C to denote the optimal centralized expected net benefit. Lastly, when the warehouse assumes a usage policy of $k_1 \in K$, whereas the actual usage policy is $k_2 \neq k_1$ at the distribution locations, the optimal expected net benefit under this assumption is given by $\pi_{k_1|k_2}$.

Throughout the remainder of this section, we characterize the optimal allocation schemes under two different usage policies at the distribution locations. In the first one, called *prioritized usage*, or simply *prioritization*, demand is satisfied by prioritizing higher-value demand classes over lower-value ones. In the second one, demand is

satisfied based on the order of arrivals, i.e., on a *first-come-first-served (FCFS)* basis.

3.3.1 Characterization of Optimal Allocation Under Prioritization

The first policy we consider prioritizes the satisfaction of demand from higher-value classes over lower-value ones. Under this policy, demand from a lower-value class can be satisfied only if there are leftovers after all the demand from higher-value classes are satisfied.

There are two ways through which prioritization can be applied. The first, which is particularly easier to apply when there are two demand classes, requires that the lower-class wait until the end of the planning horizon, until when the demand from the higher-class is satisfied. Any leftover items can be used to serve the lower-class at the end of the horizon. When blood units or relief commodities are distributed within a developing country setting, this can be achieved by serving only the most urgent patients/beneficiaries throughout the horizon and serving the remaining demand in decreasing order of value at the end. A second way is to divide the horizon into distinct sub-periods and satisfy the demand of demand classes in decreasing order of value in each subsequent sub-period. This type of demand satisfaction has been used in the 2009 H1N1 vaccination campaign, where only children under a certain age were vaccinated in the first few weeks, and the remaining demand was satisfied once this period was over. Such a method of prioritization is also applied in the fashion apparel industry, where the products are sold at decreasing unit prices in subsequent sub-periods.

To determine the optimal allocation scheme under prioritization, we first need to focus on a single distribution location and assume that there is no capacity at the warehouse. In this case, x_i denotes the "order quantity" for distribution location i. Without loss of generality, we assume that the demand takes continuous values. For the sake of brevity, we also drop the location subscript i from the calculations. The

expected net benefit function is given by:

$$\begin{aligned}E_P(z|x) = & E_P(z|D_1 > x)P(D_1 > x) \\ & + \sum_{j=1}^{|J|} \left[E_P \left(z \,\middle|\, \sum_{i=1}^{j-1} D_i \leq x, \sum_{i=1}^{j} D_i > x \right) P \left(\sum_{i=1}^{j-1} D_i \leq x, \sum_{i=1}^{j} D_i > x \right) \right] \\ & - E_P \left(z \,\middle|\, \sum_{i=1}^{|J|} D_i \leq x \right) P \left(\sum_{i=1}^{|J|} D_i \leq x \right).\end{aligned}$$

The first term in the equation gives the expected benefit when demand of the first class exceeds the order quantity x. In this case, all ordered units are used for the highest demand class. The second term considers the case where the total demand of classes 1 to $j-1$ is less than the order quantity, but the total demand of classes 1 to j exceeds x. In this case, demand of the first $j-1$ classes is completely satisfied, whereas leftovers are used to satisfy the demand of the j^{th} class, and the demand from $j+1^{st}$ to $|J|^{th}$ classes are completely unsatisfied. The last term considers the case where total demand is less than the order quantity, in which case all demand is satisfied, but a benefit loss is also incurred. Expanding the probability terms, we get:

$$\begin{aligned}E_P(z|x) = & \int_{x}^{\infty} w_1 \, x f_1(u_1) du_1 + \int_{0}^{x} \int_{x-u_1}^{\infty} \left(w_1 u_1 + w_2 \left(x - u_1 \right) \right) f_2(u_2) du_2 \, f_1(u_1) du_1 + \ldots \\ & + \int_{0}^{x} \int_{0}^{x-u_1} \ldots \int_{0}^{x-U_{n-2}} \int_{x-U_{n-1}}^{\infty} \left(w_1 u_1 + \ldots + w_{n-1} u_{n-1} + w_n \left(x - U_{n-1} \right) \right) f_n(u_n) du_n \ldots f_1(u_1) du_1 \\ & + \int_{0}^{x} \int_{0}^{x-u_1} \ldots \int_{0}^{x-U_{n-1}} \left(w_1 u_1 + \ldots + w_n u_n - w_0 \left(x - U_n \right) \right) f_n(u_n) du_n \ldots f_1(u_1) du_1, \end{aligned} \tag{36}$$

where $U_j = \sum_{l=1}^{j} u_l$. In the next proposition, we state the optimality equation, and provide a proof by induction. The optimality equation has an intuitive structure, which we make use of when we consider multiple distribution locations and a capacitated warehouse.

Proposition 1. *The optimal quantity x that maximizes Equation (36) is given by:*

$$w_1\, P(D_1 > x) + \sum_{j=2}^{|J|} w_j \left[P\left(\sum_{l=1}^{j-1} D_l \leq x \right) - P\left(\sum_{l=1}^{j} D_l \leq x \right) \right] = w_0\, P\left(\sum_{l=1}^{|J|} D_l \leq x \right), \tag{37}$$

or equivalently,

$$\sum_{j=1}^{|J|-1} (w_j - w_{j+1}) P\left(\sum_{l=1}^{j} D_l > x \right) + w_n P\left(\sum_{l=1}^{|J|} D_l > x \right) = w_0\, P\left(\sum_{l=1}^{|J|} D_l \leq x \right). \tag{38}$$

Proof. See Appendix B. □

The marginal benefit function defined by Equations (37) and (38) has an intuitive interpretation. For each demand class, it calculates the probability that the $x + 1^{st}$ unit satisfies the demand of class j and multiplies this probability with the unit benefit of this class (w_j). In the case that there are no more than x units of demand, the $x + 1^{st}$ unit results in a loss of w_0, which happens with the probability of demand not exceeding x units. The optimality equation is similar for a discrete demand distribution, with the main difference that the optimal quantity x is the smallest possible value that has the left-hand side of Equation (37) less than or equal to the right-hand side. for which case the proof is similar. Under these settings, the expected benefit and the optimality equation are similar to the case in Şen and Zhang [140], where no capacity on the upstream unit is assumed, and the overage cost (w_0) is incurred at the upstream level.

Next, we turn our attention to the case in which there are multiple distribution locations and the warehouse has a capacity of C units. Without loss of generality, we assume discrete demand and the support of each demand distribution function is the set of nonnegative integers. We have:

$$\begin{aligned} \pi_P = \max & \sum_{i=1}^{|I|} E_P(z_i|x_i) \\ s.t. & \sum_{i=1}^{|I|} x_i \leq C \qquad\qquad \forall i \in I \end{aligned}$$

$$x_i \geq 0 \text{ and integer} \qquad \forall i \in I.$$

The objective function π_P is the maximum of the sum of $|I|$ expected net benefit functions that have nonincreasing marginal benefits. Therefore, the marginal benefits obtained by the overall system are also nonincreasing. This observation, along with the fact that the marginal benefits of each distribution location i under a given allocation x_i can be calculated from Equation (37), immediately suggests a greedy allocation algorithm, summarized in Algorithm 2.

Algorithm 2 Optimal allocation for prioritization usage policy at the distribution locations

Initialize $X = (0, 0, \ldots, 0)$ and $C' = C$.
for $i = 1, 2, \ldots, |I|$ **do**
  Calculate
$$MB_P^i = w_1 P(D_1^i > x_i) + \sum_{j=2}^{|J|} w_j \left[P\left(\sum_{l=1}^{j-1} D_l^i \leq x_i \right) - P\left(\sum_{l=1}^{j} D_l^i \leq x_i \right) \right] - w_0 P\left(\sum_{l=1}^{j} D_l^i \leq x_i \right).$$
end for
while $C' > 0$ **do**
  if $\max_{i \in I} \{MB_P^i\} \leq 0$ **then**
    Set $C' = 0$ and terminate. Current allocation is optimal.
  else
    $i^* \leftarrow \arg\max_{i \in I} \{MB_P^i\}$.
    $x_{i^*} \leftarrow x_{i^*} + 1$.
    $$MB_P^{i^*} = w_1 P(D_1^i > x_{i^*}) + \sum_{j=2}^{|J|} w_j \left[P\left(\sum_{l=1}^{j-1} D_l^i \leq x_{i^*} \right) - P\left(\sum_{l=1}^{j} D_l^i \leq x_{i^*} \right) \right] - w_0 P\left(\sum_{l=1}^{j} D_l^i \leq x_{i^*} \right).$$
    $C' \leftarrow C' - 1$
  end if
end while

The greedy allocation algorithm initializes by finding the marginal benefit of allocating the first unit to each distribution location. Until there is no more distribution location with a positive marginal benefit left, it keeps allocating units in a greedy fashion and updates the marginal benefit value of the distribution location to which the last allocation was made. Since the *while* loop requires C steps, the algorithm runs in pseudopolynomial time.

3.3.2 Characterization of Optimal Allocation Under First-Come, First-Served Usage

The second policy we consider satisfies the incoming demand based on a first-come, first-served (FCFS) basis, without regard to the class to which the arriving demand belongs. In health and humanitarian settings (e.g., blood allocation), this may be applied when the beneficiaries have to travel to the distribution locations (e.g., hospitals), and hence cannot wait until the end of the planning horizon for the demand of beneficiaries with more urgent demand to be satisfied.

As in the case of our analysis for prioritization, we first focus on a single distribution location and an uncapacitated warehouse, and characterize the optimal order quantity. As before, we drop the location index i for this part. Since the order of the arrivals is important under FCFS usage, we assume that the arrivals of demand from each class j follow a Poisson process with rate λ_j to make use of the uniformity of arrivals.

The expected net benefit function is given as follows:

$$
\begin{aligned}
E_F(z|x) &= E_F\left(z\middle|\sum_{j=1}^{|J|} D_j \le x\right) P\left(\sum_{j=1}^{|J|} D_j \le x\right) + E_F\left(z\middle|\sum_{j=1}^{|J|} D_j \ge x+1\right) P\left(\sum_{j=1}^{|J|} D_j \ge x+1\right) \\
&= \sum_{u_1=0}^{x}\sum_{u_2=0}^{x-u_1}\cdots\sum_{u_{|J|}=0}^{x-U_{|J|-1}}\left[\sum_{j=1}^{|J|}\left(w_j\,u_j - w_0\,(x-U_{|J|})\right)\prod_{j=1}^{|J|}\frac{e^{-\lambda_j}(\lambda_j)^{u_j}}{u_j!}\right] + \\
&\quad \sum_{u_1=0}^{\infty}\sum_{u_2=0}^{\infty}\cdots\sum_{u_{|J|-1}=0}^{\infty}\sum_{u_{|J|}=(0,x-U_{|J|-1})^+}^{\infty}\sum_{k_1=(0,x-U_{-1})^+}^{(x,u_1)^-}\sum_{k_2=(0,x-U_{-2})^+}^{(x-u_1,u_2)^-}\cdots\sum_{k_{|J|-1}=(0,x-U_{-|J|-1})^+}^{(x-U_{|J|-2},u_{|J|-1})^-} \\
&\quad \left[\left(w_1\,k_1 + \sum_{j=2}^{|J|} w_j\left(x - \sum_{l=1}^{j-1} k_l\right)\right)\frac{\binom{u_1}{k_1}\binom{u_2}{k_2}\cdots\binom{u_{|J|-1}}{k_{|J|-1}}\binom{u_{|J|}}{x-\sum_{j=1}^{|J|-1}k_j}}{\binom{U_{|J|}}{x}}\prod_{j=1}^{|J|}\frac{e^{-\lambda_j}(\lambda_j)^{u_j}}{u_j!}\right],
\end{aligned}
\tag{39}
$$

where $U_j = \sum_{l=1}^{j} u_l$, $U_{-j} = \sum_{l=1}^{|J|} u_l - u_j$, $(a,b)^+ = \max\{a,b\}$, and $(a,b)^- = \min\{a,b\}$.

The first part of the expected net benefit function considers the case where the order amount exceeds total demand by conditioning on the number of demand arrivals from each class, the total of which cannot exceed x. In the second part, we focus on the case where total demand exceeds x and condition on the demand arrivals from each class. Since total demand exceeds the order quantity, we also need to condition

on the number of arrivals from each demand class within the first x occurrences of demand.

Since arrivals of each demand class follow a Poisson process, in the case where a total of u_j arrivals have been observed during the planning horizon for class j, the probability of observing k_j arrivals of class j within the first x arrivals follows multinomial distribution. The minimum and maximum possible values of k_j are determined such that (i) $\sum_{j=1}^{|J|} k_j = x$, and (ii) if $U_{-j} < x$, then k_j is at least $x - U_{-j}$.

Next, we determine the marginal benefit function under FCFS usage and show that the differences between consecutive marginal benefits are nonincreasing.

Proposition 2. *The marginal benefit of allocating an extra unit of commodity when x units have already been allocated to a set of distribution locations applying FCFS usage is given by:*

$$MB_F(x) = \frac{\sum_{j=1}^{|J|} w_j \lambda_j}{\sum_{j=1}^{|J|} \lambda_j} \sum_{t=x+1}^{\infty} \frac{e^{-\sum_{j=1}^{|J|} \lambda_j} \left(\sum_{j=1}^{|J|} \lambda_j\right)^t}{t!} - w_0 \sum_{t=0}^{x} \frac{e^{-\sum_{j=1}^{|J|} \lambda_j} \left(\sum_{j=1}^{|J|} \lambda_j\right)^t}{t!}. \tag{40}$$

Furthermore, the differences between consecutive marginal benefits is nonincreasing in x.

Proof. See Appendix B. □

The marginal benefit function for FCFS usage has the identical intuition to that of prioritization; for each demand class, we find the probability that the $x + 1^{st}$ unit satisfies the demand of class j, and we multiply this probability with the unit benefit of this class (w_j). From this sum, we subtract the probability that demand is less than the allocated $x + 1$ units, multiplied by the unit loss of w_0.

Having determined the marginal benefit function for FCFS usage and having established that the marginal benefits are nonincreasing, we direct our attention to the

case of multiple distribution locations and a capacity of C units at the warehouse. The optimization problem we are trying to solve is given by:

$$\begin{aligned} \pi_F = \max & \sum_{i=1}^{|I|} E_F(z_i|x_i) & \\ s.t. & \sum_{i=1}^{|I|} x_i \leq C & \forall i \in I \\ & x_i \geq 0 \text{ and integer} & \forall i \in I. \end{aligned}$$

As in the case of prioritization, π_F is the maximum of the sum of $|I|$ functions that have nonincreasing marginal benefits, which implies that the optimal allocation under FCFS usage policy also has nonincreasing marginal benefits. Since the marginal benefits under a given allocation at a distribution location can be calculated by Equation (40), we can use a similar greedy algorithm to determine the optimal allocation quantities for each location. The algorithm is summarized in Algorithm 3.

Algorithm 3 Optimal allocation for FCFS usage policy at the distribution locations

1: Initialize $X = (0, 0, \ldots, 0)$ and $C' = C$.
2: **for** $i = 1, 2, \ldots, |I|$ **do**
3: Calculate

$$MB_F^i = \frac{\sum_{j=1}^{|J|} w_j \lambda_j}{\sum_{j=1}^{|J|} \lambda_j} \sum_{t=x_i+1}^{\infty} \frac{e^{-\sum_{j=1}^{|J|} \lambda_j} \left(\sum_{j=1}^{|J|} \lambda_j\right)^t}{t!} - w_0 \sum_{t=0}^{x_i} \frac{e^{-\sum_{j=1}^{|J|} \lambda_j} \left(\sum_{j=1}^{|J|} \lambda_j\right)^t}{t!}.$$

4: **end for**
5: **while** $C' > 0$ **do**
6: **if** $\max_{i \in I} \{MB_F^i\} \leq 0$ **then**
7: Set $C' = 0$ and terminate. Current allocation is optimal.
8: **else**
9: $i^* \leftarrow \arg\max_{i \in I} \{MB_F^i\}$.
10: $x_{i^*} \leftarrow x_{i^*} + 1$.
11: $MB_F^{i^*} = \frac{\sum_{j=1}^{|J|} w_j \lambda_j}{\sum_{j=1}^{|J|} \lambda_j} \sum_{t=x_i+1}^{\infty} \frac{e^{-\sum_{j=1}^{|J|} \lambda_j} \left(\sum_{j=1}^{|J|} \lambda_j\right)^t}{t!} - w_0 \sum_{t=0}^{x_i} \frac{e^{-\sum_{j=1}^{|J|} \lambda_j} \left(\sum_{j=1}^{|J|} \lambda_j\right)^t}{t!}.$
12: $C' \leftarrow C' - 1$
13: **end if**
14: **end while**

Similar to Algorithm 2, Algorithm 3 starts by initializing the marginal benefits using Equation (40). It then finds the distribution location with the largest nonnegative marginal benefit, allocates one more unit for the location, and updates its marginal benefit. The algorithm terminates either when capacity is exhausted or when there is no distribution location with a positive marginal benefit left. Because of the number of operations in the *while* loop, the algorithm is pseudopolynomial.

To the best of our knowledge, Algorithms 2 and 3 are the first ones proposed for the newsvendor problem with a capacitated distribution center, multiple distribution locations, and multiple demand classes.

3.4 Structural Results

In the previous section, we have established the optimal allocation schemes under prioritization and FCFS usage in the distribution locations. In this section, we first analyze the properties of the optimal order quantities when there is a single distribution location and an uncapacitated warehouse. Second, we compare the relative performances of optimal order quantities (when there is a single distribution location) or optimal allocations (when there are multiple distribution locations) under prioritization and FCFS usage. Lastly, we consider serving part or all of the demand classes from a single location by centralization of the inventory, and compare the resulting optimal allocations to those for the original decentralized case under both usage policies.

3.4.1 Structural Results for a Single Distribution Location and Uncapacitated Warehouse

Under prioritization, Equation (37) does not have a closed-form result for the optimal quantity x. The first result we show in this section states that under prioritization at a single distribution location and uncapacitated warehouse, the optimal quantity can be bounded by the optimal single-class quantities for the highest- and lowest-priority

demand classes. Let $G(.)$ denote the cumulative distribution function of the total demand over all classes, that is,

$$G(x) = P(D_1 + D_2 + \ldots + D_{|J|}) \leq x.$$

For the discrete case, $G^{-1}(p)$ maps p to the smallest integer value x for which $G(x) = p$.

Proposition 3. *Under prioritization at a single distribution location and with an uncapacitated warehouse, we have:*

$$G^{-1}\left(\frac{w_{|J|}}{w_{|J|} + w_0}\right) \leq \left(x^P\right)^* \leq G^{-1}\left(\frac{w_1}{w_1 + w_0}\right).$$

Proof. See Appendix B. □

This result is in line with that obtained by Şen and Zhang [140], where the cost of overage is incurred at the warehouse and the optimal allocation quantity is bounded between $G^{-1}\left(\frac{w_{|J|}-w_0}{w_{|J|}}\right)$ and $G^{-1}\left(\frac{w_1-w_0}{w_1}\right)$, Proposition 3 gives a tighter bound, as the difference in the threshold probabilities at $\frac{w_0(w_1-w_{|J|})}{w_1\, w_{|J|}}$ for the former is larger than that for the latter, which is $\frac{w_0(w_1-w_{|J|})}{(w_1+w_0)(w_{|J|}+w_0)}$.

By comparison, the optimal order quantity for FCFS usage policy can be exactly given by:

$$\left(x^F\right)^* = G^{-1}\left(\frac{\sum_{j=1}^{|J|} w_j\, \lambda_j}{\sum_{j=1}^{|J|} (w_0 + w_j)\, \lambda_j}\right).$$

The optimal quantity given by this equation has a similar structure to that for the traditional newsvendor model. The numerator of the inverse distribution function is the *weighted underage cost* per unit, where the weights for each demand class correspond to the demand arrival rate. Similarly, the denominator is the sum of per unit *weighted overage* and underage costs.

Next, we present a lemma for two demand classes that states that for a given allocation x, if the marginal benefits for both prioritization and FCFS usage are of the same sign, then the ratio of $MB_F(x)$ to $MB_P(x)$ is nondecreasing in x. Using this fact helps us establish a number of identities throughout this section.

Lemma 1. *Assume that* $|J|= 2$ *and* $MB_F(x-1) \leq 0$ *and* $MB_P(x-1) \leq 0$, *or* $MB_P(x) \geq 0$ *and* $MB_P(x) \geq 0$ *for a given* x. *Suppose also that* $min\{\lambda_1, \lambda_2\} \geq \frac{w_0}{w_2}$. *Then:*

$$\frac{MB_F(x)}{MB_P(x)} - \frac{MB_F(x-1)}{MB_P(x-1)} \geq 0.$$

Proof. See Appendix B. □

Using Lemma 1, we can establish the fact that the optimal order quantity under prioritization is no larger than that under FCFS usage policy.

Theorem 7. *Assume that* $min\{\lambda_1, \lambda_2\} \geq \frac{w_0}{w_2}$. *Then* $(x^P)^* \leq (x^F)^*$.

Proof. We have:

$$\lim_{x\to\infty}(MB_F(x) - MB_P(x)) = \lim_{x\to\infty}\left(\frac{w_1\,\lambda_1 + w_2\,\lambda_2}{\lambda_1+\lambda_2}\,P(D_1+D_2 \geq x+1)\right.$$
$$\left. -\,(w_1 - w_2)\,P(D_1 \geq x+1) - w_2\,P(D_1+D_2 \geq x+1)\right) = 0.$$

Furthermore, since $P(D_1 \geq x+1)$ tends to zero more quickly than $P(D_1+D_2 \geq x+1)$ and $w_2 \leq \frac{w_1\,\lambda_1+w_2\,\lambda_2}{\lambda_1+\lambda_2}$, the limit approaches zero from above, implying that the $MB_F(x)$ is slightly larger in the limit. By Lemma 1, the ratio of the marginal benefit of FCFS to prioritization usage is nondecreasing when both marginal benefits are nonpositive. This implies that the marginal benefit function for prioritization has to become negative no later than that of the FCFS marginal benefit function, implying that $(x^P)^* \leq (x^F)^*$. □

The condition $\min\{\lambda_1, \lambda_2\} \geq \frac{w_0}{w_2}$ in Theorem 7 is not very restrictive; it basically ensures that the optimal order quantity in this distribution location is nonzero. Therefore, the conclusion $(x^P)^* \leq (x^F)^*$ holds in all nontrivial cases.

3.4.2 Relative Performance Comparison for Prioritization and FCFS Usage Policies

In this part, we compare the performances of optimal allocation schemes under prioritization and FCFS usage at the distribution locations. We focus on the case with two demand classes with Poisson arrivals. First, we start with the case of a single distribution location and derive the worst-case performance ratio of FCFS performance relative to that of prioritization under a given set of demand rate and unit benefit parameters.

Theorem 8. *In the case where:* $\frac{\lambda_2}{\lambda_1} = K$ *for some constant* $K \geq 0$, $\frac{w_1}{w_2} = L$ *for some constant* $L \geq 1$, *and* $min\{\lambda_1, \lambda_2\} \geq T = \frac{w_0}{w_2}$ *for some constant* $T \geq 0$,

$$\frac{\pi_F}{\pi_P} \geq \frac{K+L}{L(K+1)}, \tag{41}$$

and the bound is asymptotically tight.

Proof. See Appendix B. □

By Theorem 8, we observe that there are two main factors that determine the ratio between the optimal performances of the two usage policies. These are the ratio of the unit benefits of the two classes, the increase of which increases the difference in performances; and the ratio of demand rates for each class. Figure 18 shows the change in the worst-case ratio for $\frac{\pi_F}{\pi_P}$ for changing values of K and L.

An immediate corollary follows from Theorem 8:

Corollary 2. *When* K *or* L *are not fixed, it is possible that* $\frac{\pi_F}{\pi_P} \to 0$.

Next, we turn our attention to the case where there are multiple distribution locations, and the ratio of the demand rate for the two classes (K) can take values

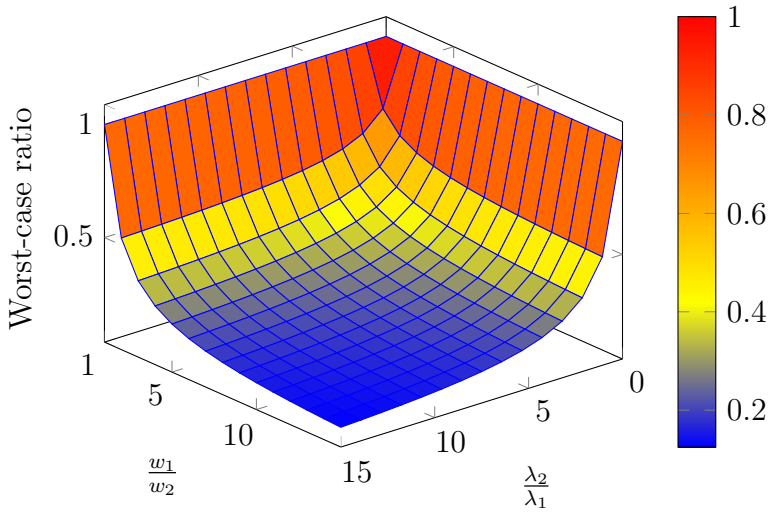

Figure 18: Plot of the worst-case ratio for $\frac{\pi_F}{\pi_P}$ for the ranges $K = \frac{\lambda_2}{\lambda_1} = [0, 15]$ and $L = \frac{w_1}{w_2} = (1, 15]$

between a lower bound l and an upper bound u. The following theorem states that the worst-case bound depends only on the value of u.

Theorem 9. *With $|I|$ distribution locations and $l \leq \frac{\lambda_2^i}{\lambda_1^i} \leq u$ for all $i = 1, 2, \ldots, |I|$, $\frac{w_u}{w_n} = L \geq 1$, and $\frac{w_o}{w_n} = T \geq 0$, with $\min_i\{\lambda_1^i, \lambda_2^i\} \geq T$,*

$$\frac{\pi_F}{\pi_P} \geq \frac{u + L}{L(u + 1)},$$

and the bound is asymptotically tight.

Proof. If both FCFS and prioritization usage result in the same allocations for each distribution location, by the result of Theorem 8, we know that the ratio between the two allocations will never be less than $\frac{u+L}{L(u+1)}$, as each distribution location has a rate ratio less than u. Hence, for the whole system, this ratio is a lower bound.

To see that the ratio is asymptotically tight, consider the case where λ_1^i is sufficiently large for each distribution location i so that the condition in Equation (125) is satisfied. Then for each distribution location i, we have

$$\frac{\pi_F}{\pi_P} \to \frac{u + L}{L(u + 1)},$$

which obviously holds for the whole system. □

Theorem 9 states that the worst-case performance of the system depends on the demand rate ratio of a single distribution location, namely the one with the highest demand rate ratio of lower to higher class demand.

When the demand rate ratio K is constant throughout the system, the worst-case performance is identical to that for a single distribution location.

Corollary 3. *With $|I|$ distribution locations, when $l = u = K$,*

$$\frac{\pi_F}{\pi_P} \geq \frac{K+L}{L(K+1)},$$

and the bound is tight when λ_1^i is large enough to guarantee that Equation (125) holds for each distribution location.

As in the case of a single distribution location, FCFS usage might perform infinitely worse than prioritization if the demand rate and unit benefit ratios are not fixed.

Corollary 4. *With $|I|$ distribution locations, it is possible to have $\frac{\pi^{F|F}}{\pi^{P|P}} \to 0$.*

Hence, the results from the single distribution location case extend to multiple locations when the demand rate ratio is either fixed or unbounded.

The last theorem in this part gives the intuitive result that the optimal performance of FCFS usage can never exceed that of prioritization.

Theorem 10. *For the case with $|I|$ distribution locations and $|J|$ demand classes, $\pi_F \leq \pi_P$.*

Proof. Let $(x_F^i)^*$ be the amount of units allocated to location i under the optimal allocation for FCFS usage policy. Define a heuristic allocation scheme $\mathcal{H}$ as one which allocates $(x_F^i)^*$ units to each location i. Let $\pi_P(\mathcal{H})$ denote the benefit of this allocation scheme under prioritization. It is obvious that under each possible observation of the number of arrivals in each distribution location, $\mathcal{H}$ will achieve at least equal total net

benefit under prioritization compared to the FCFS allocation, since demand can be "sorted" in terms of classes and satisfied in that order. Since net benefits are at least equal under each observation, this holds true for the expectation as well. Therefore, we have: $\pi_P \geq \pi_P(\mathcal{H}) \geq \pi_F$. □

3.4.3 Structural Results for Centralized Systems

In this part, we focus on the case where, instead of a set of distribution locations, (i) there exists a single distribution location that faces the aggregate demand faced by the individual locations (complete centralization), or (ii) a certain demand class is served from a single and separate distribution location that specifically serves this type of demand (partial centralization). The objective in both cases is to compare the performances of the (partially or completely) centralized system with respect to the decentralized counterpart. For this part, we use the additional notation $\Lambda_j = \sum_{i=1}^{|I|} \lambda_j^i$ to denote the total rate of demand of class j throughout the system. Unless otherwise stated, the results are valid for a general number of demand classes. The remaining demand rate and unit benefit settings are identical to those in the preceding part.

We start with the case of a system employing FCFS usage. Our first conclusion is that if there is sufficient capacity, the centralized system performs at least as well as the decentralized one.

Theorem 11. *If the warehouse is uncapacitated,* $\pi_F \leq \pi_F^{cen}$.

Proof. Denoting the quantity of units allocated to distribution location i by $(x_F^i)^*$ under the decentralized setting, define a heuristic $\mathcal{H}$ for the centralized system which allocates $\sum_{i=1}^{|I|}(x_F^i)^*$ units to the centralized location. As before, we denote by $\pi_C(\mathcal{H})$ the expected net benefit of heuristic $\mathcal{H}$ under the centralized system. For every realization of demand quantities for each class in each location, the benefit obtained by $\mathcal{H}$ will be equal, as it has the same total number of units on hand and faced the exact

demand faced by the decentralized system. Therefore, it follows that the expected net benefit under $\mathcal{H}$ is also equal. Consequently, we have $\pi_F^{cen} \geq \pi_C(\mathcal{H}) = \pi_F$. $\square$

While the centralized system always performs at least as well as the decentralized one under FCFS usage, this may not be the case if the system is capacitated. The following proposition gives a sufficient condition.

Proposition 4. *Under limited allocation capacity, a decentralized system under FCFS usage policy may perform better than complete centralization. For two demand classes, a sufficient condition for the demand rate and unit benefit settings is such that there is no distribution location i satisfying:*

$$\begin{aligned}(w_1 - w_2)\,(\lambda_1^i \Lambda_2 - \lambda_1^i \Lambda_1) \leq e^{-(\lambda_1^i + \lambda_2^i)} \Big[& (1 - e^{-(\Lambda_1^{-i} + \Lambda_2^{-i})}) \left((w_1 + w_0)\,\lambda_1^i \Lambda_1 + (w_2 + w_0)\,\lambda_2^i \Lambda_2 \right) \\ & + \lambda_1^i \Lambda_2 \left((w_1 + w_0) - (w_2 + w_0)\, e^{-(\Lambda_1^{-i} + \Lambda_2^{-i})} \right) \\ & - \lambda_2^i \Lambda_1 \left((w_1 + w_0)\, e^{-(\Lambda_1^{-i} + \Lambda_2^{-i})} - (w_2 + w_0) \right) \Big], \end{aligned} \tag{42}$$

where:

$$\Lambda_1^{-i} = \sum_{\substack{j=1 \\ j \neq i}}^{|I|} \lambda_1^j \quad \textit{and} \quad \Lambda_2^{-i} = \sum_{\substack{j=1 \\ j \neq i}}^{|I|} \lambda_2^j.$$

Proof. The sufficient condition is derived from the case where the first allocation to the decentralized system has a higher expected net marginal benefit than that of the completely centralized system. For this, we need:

$$\max_{i \in I} \left\{ \left(\frac{w_1\,\lambda_1^i + w_2\,\lambda_2^i}{\lambda_1^i + \lambda_2^i} + w_0 \right) P_{dec}(D_1 + D_2 \geq 1) - w_0 \right\} \geq \left(\frac{w_1\,\Lambda_1 + w_2\,\Lambda_2}{\lambda_1^i + \Lambda_2} + w_0 \right) P_{cen}(D_1 + D_2 \geq 1) - w_0,$$

Where $P_{cen}(.)$ and $P_{cen}(.)$ denote the probability functions for the decentralized and completely centralized systems, respectively. This condition holds only if there exists no distribution location i such that:

$$\left(\frac{w_1\,\lambda_1^i + w_2\,\lambda_2^i}{\lambda_1^i + \lambda_2^i} + w_0 \right) P(D_1 + D_2 \geq 1) \leq \left(\frac{w_1\,\Lambda_1 + w_2\,\Lambda_2}{\lambda_1^i + \Lambda_2} + w_0 \right) P(D_1 + D_2 \geq 1),$$

that is, for at least one of the distribution locations, the probability of observing at least one unit of demand is sufficiently large to exceed the expected benefit of the first unit of demand in the completely centralized system. The rest of the proof follows from direct algebra. □

If all distribution locations have identical arrival rates in the decentralized system, the condition given by Equation (42) does not hold. However, this is not sufficient to prove that distribution locations with identically distributed arrivals imply a completely centralized system being better at all capacity levels, since Equation (42) is not a necessary condition. Experimenting over a large set of parameters, we observe that complete centralization with identically distributed arrivals does improve the expected benefits over the original decentralized system.

Observation 5. *Suppose* $\lambda_1^i = \lambda_1$ *and* $\lambda_2^i = \lambda_2$ *are identical for all* i*. Then, expected benefit under the completely centralized system is at least as good as that of the decentralized system, regardless of the capacity.*

Figure 19 shows the plot of the maximum performance difference $\pi_F^{cen} - \pi_F$ over all capacity levels with varying λ_1, λ_2 and $|J|$ values. In all such experiments, we have observed the minimum performance difference at the minimum possible capacity level ($C = 1$ unit). Under increasing capacity, the difference increases to a certain point, and decreases again near the optimal allocation amount, but never reaches the difference level for $C = 1$. The importance of Observation 5 is that the result is analogous to that from queuing theory: For servers with identical service rates, it is always beneficial to pool the servers into one that has the aggregate rate of the individual servers and a single queue. As the difference between the service rates of the servers increases, pooling is no longer beneficial to system performance [164, 165]. The experimental findings form a basis for structurally proving the fact stated in Observation 5.

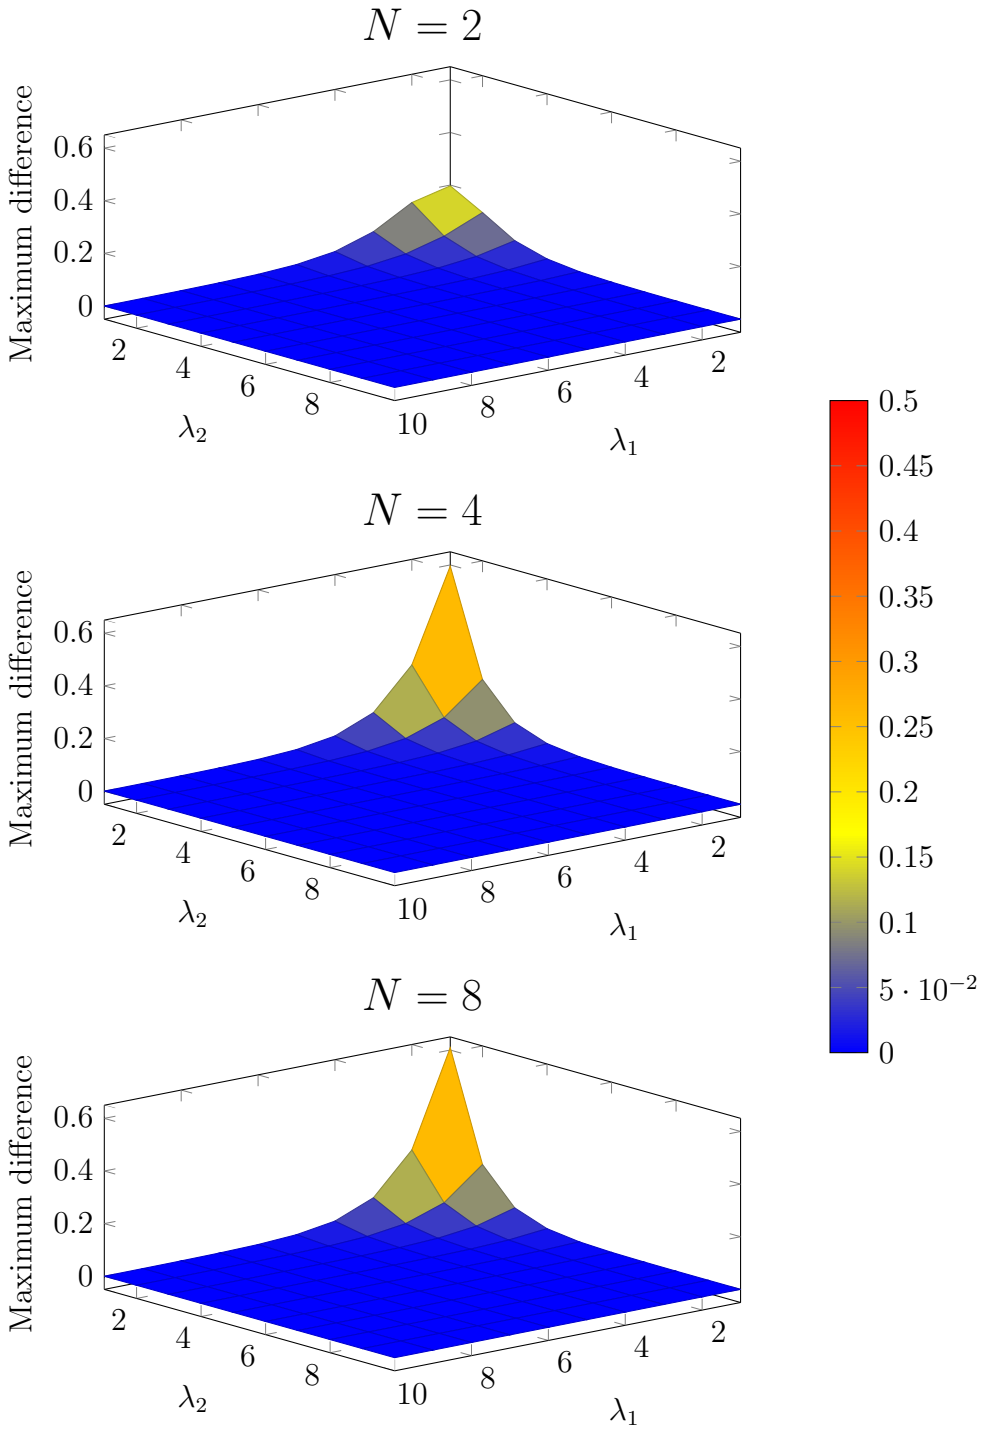

Figure 19: Plots of the maximum difference between the centralized and decentralized system performances over all capacity levels with λ_1 and λ_2 in the range $[1, 10]$ and (i) $|J|= 2$, (ii) $|J|= 4$, and (iii) $|J|= 8$ distribution locations with identical arrival rates

While it is possible to have the decentralized system performing better than a completely centralized system under FCFS usage, we show in the following theorem that a completely centralized system is always at least as good under prioritization. To the best of our knowledge, this is the first result on centralization in newsvendor models with multiple distribution locations and demand types when demand is prioritized.

Theorem 12. *Under prioritization usage policy at the distribution locations, a completely centralized system always performs at least as well as the decentralized one, regardless of the capacity level.*

Proof. The proof we provide for this is similar to that for Theorem 11. We use $(x_P^i)^*$ to denote the number of units allocated to distribution location i under a decentralized system. Let heuristic $\mathcal{H}$ allocate $\sum_{i=1}^{|I|}(x_P^i)^*$ units to the completely centralized system. Since usage is prioritized, the centralized system can take advantage of prioritizing the higher class demands that would not be prioritized over the lower class demands at other locations in the decentralized system, and hence can result in a better net benefit. As the net benefit of the centralized system is at least as good as the decentralized benefit under $\mathcal{H}$ for each possible observation of demand, the next expected benefit is also at least as high. This yields:

$$\pi_P^{cen} \geq \pi_C(\mathcal{H}) = \pi_P,$$

regardless of the capacity level at the warehouse. □

Since complete centralization may not be beneficial when demand is satisfied on a first-come, first-served basis, an alternative way to improve the system might involve partial centralization by demand classes. For example, within the context of blood allocation, more urgent patients might be served by a specialized hospital. In the following proposition, we show that under both FCFS and prioritization policies, partial centralization may not always lead to beneficial results.

Proposition 5. *Consider the following two centralization schemes: (1) satisfaction of higher-class demand in a single location separate from the remaining locations, and (2) satisfaction of lower-class demand in a single location. For both cases, under FCFS or prioritization, the original decentralized system may have a better expected net benefit, even when capacity is sufficient to supply the optimal amounts.*

Proof. Consider the following setting with $w_1 = 5$, $w_2 = 2$, $w_0 = 1$, and $|J| = 2$: $\lambda_1^1 = 5$, $\lambda_2^1 = 2$, $\lambda_1^2 = 6$, and $\lambda_2^2 = 3$. Under FCFS usage, the optimal decentralized allocations are $(x_1^F)^* = 9$ and $(x_2^F)^* = 11$.

Under case (1), optimal centralized FCFS allocations are $(x_1^F)^* = 2$, $(x_2^F)^* = 4$, and $(x_C^F)^* = 14$ for locations 1, 2, and the central location that serves higher-class demand, respectively. If $C \leq 13$, the partially centralized system performs better. However, if $C \geq 14$ (therefore also when the system is uncapacitated), the expected net benefit under the decentralized system is better than that of the partially centralized system. For the same example under case (2), optimal centralized FCFS allocations are $(x_1^F)^* = 7$, $(x_2^F)^* = 8$, and $(x_C^F)^* = 6$ for locations 1, 2, and the central location that serves lower-class demand, respectively. In this case, the partially centralized system is better up to and including $C = 14$, whereas the decentralized system has a better expected performance otherwise. For the uncapacitated system, the decentralized and partially centralized systems have expected net benefits of 56.697 and 55.103 units, respectively.

For the same example, under a decentralized system applying prioritization, $(x_1^P)^* = 9$ and $(x_2^P)^* = 11$. For case (1), partial centralization results in 2, 4, and 14 units for locations 1, 2, and the center that serves higher-class demand. The partially centralized system is better only if $C \leq 10$. Under case (2), optimal allocations are 7, 8, and 6 for the same locations, which outperforms the decentralized system for $C \leq 2$.

In the example given, higher-class demand is more prevalent in both locations. In the opposite case, where $\lambda_1^1 = 2$, $\lambda_2^1 = 12$, $\lambda_1^2 = 1$, and $\lambda_2^2 = 10$, we obtain similar

Table 11: Summary of findings for decentralized, partially centralized, and completely centralized systems under FCFS and prioritization policies

Usage policy	Complete centralization		Partial centralization	
	Suff. cap.	**Insuff. cap.**	**Highest class**	**Lowest class**
FCFS	Centralized system always performs better	Decentralized system may perform better	Centralized system may perform better only up to a certain capacity	
Prioritization	Centralized system always performs better		Centralized system may perform better only up to a certain capacity	

results. The optimal decentralized FCFS allocations are $(x_1^F)^* = 16$ and $(x_2^F)^* = 13$. For case (1), optimal partially centralized FCFS allocations are 13, 11, and 5 units for the first, second, and central locations. This partially centralized system performs better only if $C \leq 10$. For case (2), optimal partially centralized FCFS allocations are 3, 2, and 24 units for the same locations. Again, this can only outperform the decentralized system for $C \leq 9$ units.

For the same example, optimal decentralized prioritization allocations are $(x_1^P)^* = 15$ and $(x_2^P)^* = 11$. For case (1), optimal allocations under partial centralization are $(x_1^P)^* = 13$, $(x_2^P)^* = 11$, and $(x_C^P)^* = 5$. Up to $C = 2$, the partially centralized system has better expected net benefit, whereas for $C \geq 3$ the opposite holds. For case (2), partially centralized system yields $(x_1^P)^* = 3$, $(x_2^P)^* = 2$, and $(x_C^P)^* = 24$. In this case, the decentralized system always outperforms partial centralization, regardless of the capacity. □

Hence, when the structural results for centralization are considered, we can conclude that the only way to guarantee an improved expected benefit over the original decentralized system regardless of the capacity level is (i) by applying prioritization, and (ii) by complete centralization into a single location. Our structural results on centralization are summarized in Table 11.

3.5 Computational Experiments

While the results in §3.4 provide insights into the structure of the optimal order or allocation schemes under prioritization and FCFS usage, we also would like to observe how these would perform under more realistic scenarios. Additionally, we also would like to analyze the effects of changing the problem parameters (e.g., demand rate, capacity, and unit benefits) on the overall expected benefits. In the preceding section, we have established the fact that under sufficient capacity, complete or partial centralization of the distribution locations always performs at least as well as the original decentralized system under sufficient capacity. The extent at which the centralized system performs better will also be analyzed in this section. Furthermore, we have shown in Section 3.4 that a decentralized system can perform better if capacity is scarce when FCFS usage is applied. Another goal of our computational experiments is to observe if this is the case in general under limited capacity.

So far in our analysis we have assumed that the warehouse has complete knowledge about the usage policies in each distribution location. Due to lack of information infrastructure, this may not always be the case. For example, the warehouse may assume prioritization at the distribution locations while making the allocations, whereas the actual usage is FCFS. Therefore, the effect of information asymmetry between the actual usage at the distribution locations and the assumed usage at the warehouse is also of interest in this section.

In humanitarian supply chains, one of the main concerns of resource allocation is the *equity* of the distribution. While measuring the equity of the system is context-dependent, an equitable system in the context given in this chapter can be defined as one in which all beneficiaries in a given demand class has an equal probability of access to the service. We also extend this definition to include the whole beneficiaries as a single class and define a second equity measure based on this.

The remainder of this section is organized as follows: First, we present the set of

instances we have used, which is based on the H1N1 vaccination campaign in the state of Georgia in 2009. Next, we define two benchmark allocation approaches to mimic the actual application schemes. The last part of the section analyzes the results of the experiments.

3.5.1 Instance Generation

The problem instances we use for our computational experiments are based on the H1N1 vaccination campaign in the state of Georgia. Upon the possibility of an H1N1 breakout, a national campaign was initiated to mitigate against the effects of the disease by vaccination. Each state was allocated a certain number of vaccine units based on population. There were mainly two different classes of demand: The Advisory Committee on Immunization Practices (ACIP) at the CDC determined children and elderly citizens as the priority groups. The vaccination campaign was set up in two stages: In the first phase, vaccines would be provided to only the prioritized group. In the second stage, they would be made available to both demand classes. However, due to difficulties in implementation (especially failure to reject the demand from the non-prioritized class in the first phase) and extreme political pressure to use all the vaccine units, prioritization was not fully applied.

As mentioned, we limit our analysis to the state of Georgia. The population data used in the analysis is obtained from the results of the 2000 Census [162]. There are two demand classes: The higher-priority class consists of the children under the age of 18, which consisted of approximately 25% of the whole population. The remainder of the population was considered as the lower-priority class. Demand arrivals were assumed to follow a Poisson process, and in the base case, the rates of the Poisson process for each class were assumed to be 30% of the population size of the class. To analyze the effects of demand, these rates were varied to 10% and 50% of their corresponding population.

Because of the substantial amount of political pressure to allocate all the vaccine units to the fullest extent possible, the unit benefits of satisfying demand were taken as significantly higher compared to the loss of benefit for unused units. For the base case, w_1 and w_2 were taken at 4,000 and 2,000 per unit, whereas w_0 was assumed to be 100 per unit. To analyze the effects of these unit benefits, we also use a (w_1, w_2) pair of $(3500, 2500)$ and $(5000, 1000)$. Our main aim here is to see the relative difference between prioritization and FCFS when the ratio $\frac{w_1}{w_2}$ changes. We also consider a case where benefit loss due to allocated but unused units is more significant compared to the benefit of satisfying demand. This case better represents allocation of blood units in a developing country setting, and hence is called the *blood* case throughout the rest of the section. The unit benefit for the lower priority class is 1,000, and the benefit loss due to allocated but unused units is 8,000 per unit in this case. The capacity of the warehouse was determined as 11.5% of the total population based on the national capacity of vaccines. To measure the effect of changing capacity, we also increase the capacity to 40% of the total population.

Two main levels of centralization of inventory were considered for allocation. The first assumes that allocation is made to locations at the census tract level, where $|I| = 1,618$, and for the second, the allocation is assumed to be made to locations at the county level, in which case $|I| = 159$. In practice, demand can be satisfied in this way as follows: The warehouse at the state level makes allocations to the census tracts or counties based on the overall demand levels for each class and the usage policies applied. Each census tract or state (depending on the level of centralization) has a central warehouse storing the commodities, which are in turn transported to local hospitals on an as-needed basis. Figure 20 shows the percentage of children under 18 at the county and census tract level. It can be observed from the figure that there is significant variance in the percentages when the system is more decentralized, which contributes to the loss in system performance.

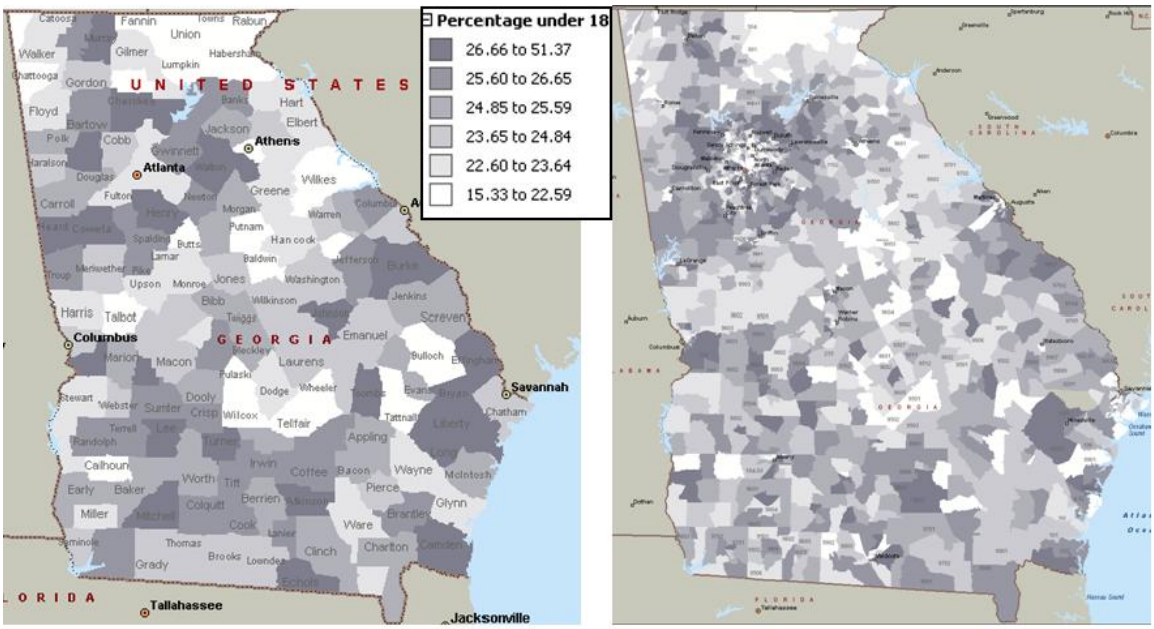

Figure 20: Percentages of the prioritized class (children under the age of 18) over the population of (i) the county, and (ii) the census tract

Table 12: Summary of parameters used in the computational experiments, with base case values given in bold

Parameter	Possible values
Mean demand as % of total population	10, **30**, 50
Overall capacity as % of total population	**11.5**, 40
Unit benefits and losses (w_1, w_2, w_0)	$(\mathbf{4000}, \mathbf{2000}, \mathbf{100})$, $(3500, 2500, 100)$, $(5000, 1000, 100)$, $(4000, 1000, 8000)$, $(3500, 2500, 8000)$, $(5000, 1000, 8000)$
Number of distribution locations	1618 (census tracts), 159 (counties)

Table 12 summarizes the parameters used in the computational experiments. The base case parameters are shown in bold.

3.5.2 Benchmark Approaches

The optimal allocation of commodities from the warehouse requires solving the problem using one of the Algorithms 2 or 3. Since this requires computational power and involves technical complexity, it may be more appealing for the decision makers to make the allocations using a simpler scheme. For instance, the actual allocation

during the vaccination campaign was made based on the total demand rate in each location. In order to use them as benchmarks to compare to our models, we consider two simple-to-apply allocation policies.

The first policy, which is called the λ_1-*policy*, allocates the units on hand to the distribution locations proportionally based on the prioritized demand rate in each location i (λ_1^i). Since the units to be allocated are discrete, this leaves a small part of the units unallocated. Therefore, the remainder is allocated based on the rate of total demand ($\lambda_1^i + \lambda_2^i$) to the distribution locations. While allocating all of the units on hand may not necessarily be optimal, within the context of the vaccine allocation campaign, this is motivated by the political pressure to use all units on hand.

The second policy, which we will refer to as the λ_t-*policy*, makes the allocation proportionally based on the total demand rate in each location. Once again, any remaining units on hand after this allocation is allocated using the rate of total demand.

3.5.3 Computational Results

We begin analyzing the computational results by considering the effect of varying unit benefits for each demand class. Figures 21 and 22 show the total net benefit under prioritization, FCFS usage, and under prioritization allocation with FCFS actual usage.

As was also observed in Section 3.4, **the relative performance gap between prioritization and FCFS usage policies is mainly determined by the ratio of unit benefits for each class**. From the results in Figure 21, the ratio of the expected benefits under FCFS to expected benefits under prioritization is around 89% when (w_1, w_2) = (3500, 2500), whereas the ratio decreases to 60% when (w_1, w_2) = (5000, 1000). By comparison, the theoretical lower bounds for these ratios from Section 3.4 are approximately 78% and 40%, respectively.

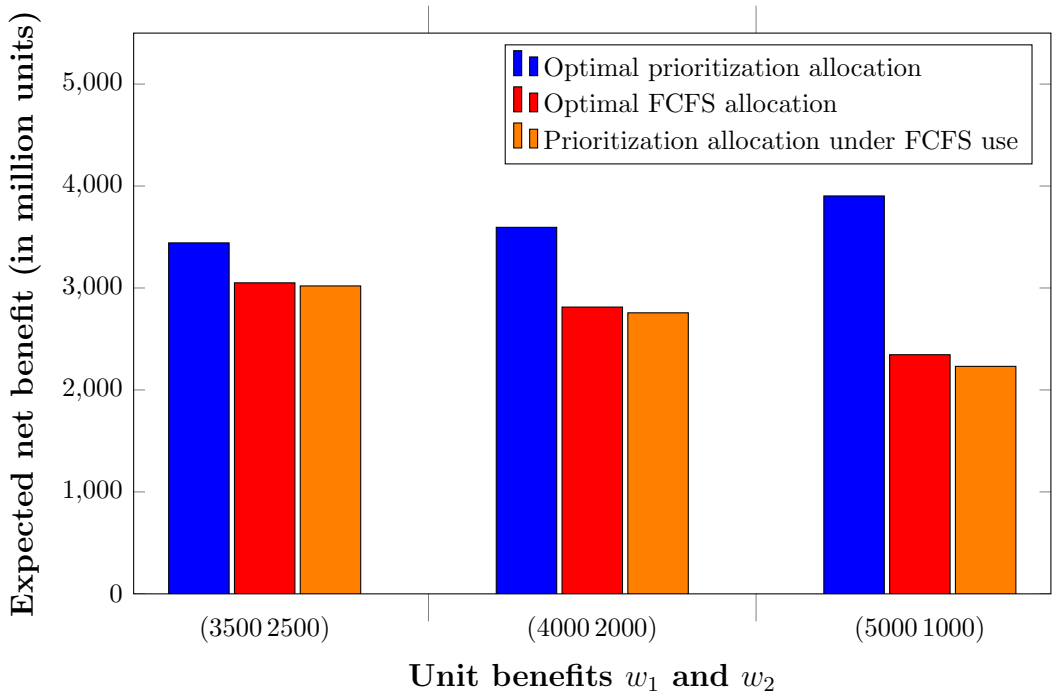

Figure 21: Total net benefit for the base case with varying unit benefits (w_1, w_2) and demand at census tract levels

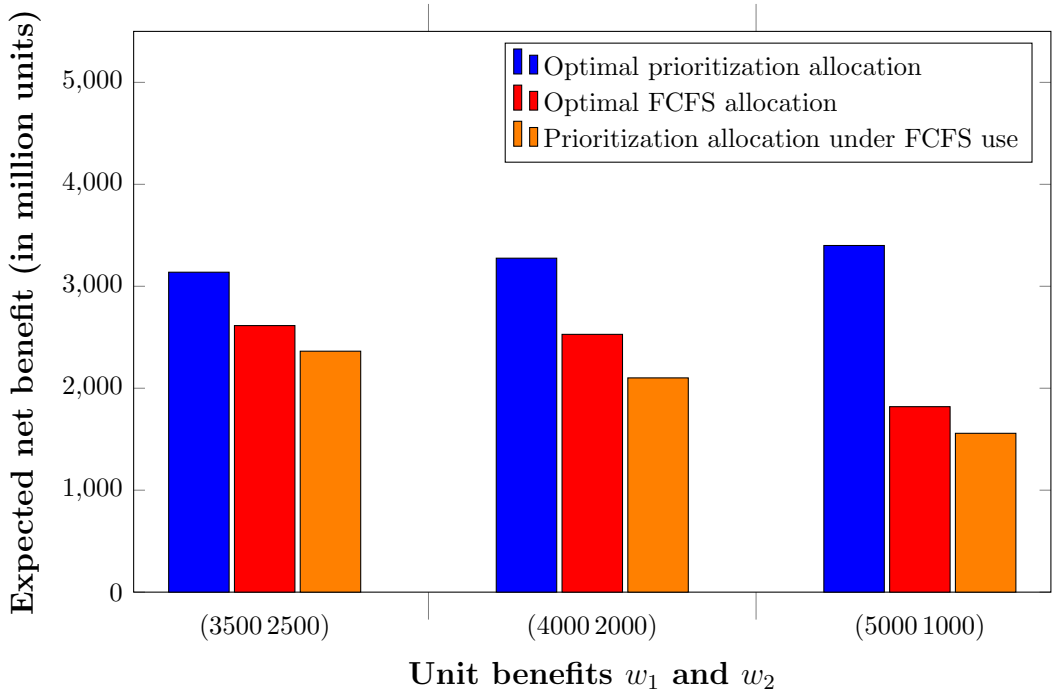

Figure 22: Total net benefit for the blood case with varying unit benefits (w_1, w_2) and demand at census tract levels

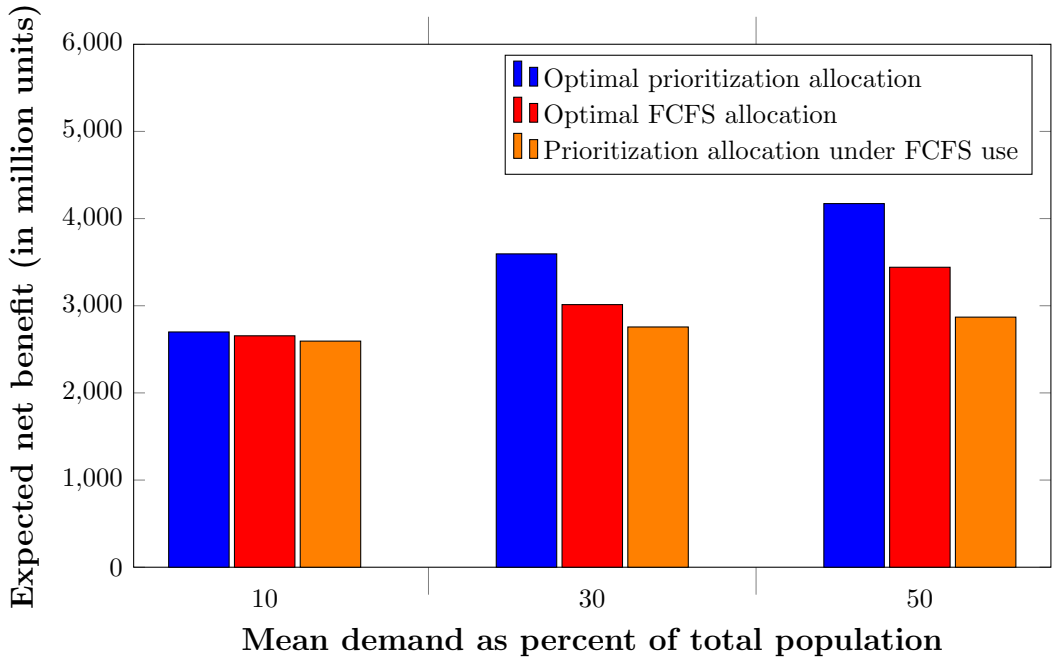

Figure 23: Total net benefit for the vaccine case with demand at census tract level and varying ratio of the rate of arrivals to total population

The results for the blood case in Figure 22 are similar to those obtained for the vaccine case. As the ratio of the unit benefits increases, the performance of FCFS usage policy relative to prioritization usage decreases. An important observation here is that **the difference due to information asymmetry is higher when the unit benefit loss due to unused units is higher**. This is mainly because as w_0 increases, the allocations made to distribution locations under prioritization and FCFS usage policies become more and more different, which results in a more significant benefit loss when the usage information is incorrect.

Next, we direct our attention to the case where the rate of the Poisson process for demand arrivals from each class is varied. We also consider the effect of centralization of inventory on the overall system performance. Figures 23 and 24 show the total net benefit under prioritization usage, FCFS usage, and prioritization allocation under FCFS usage for varying demand rates at the census tract and county levels, respectively.

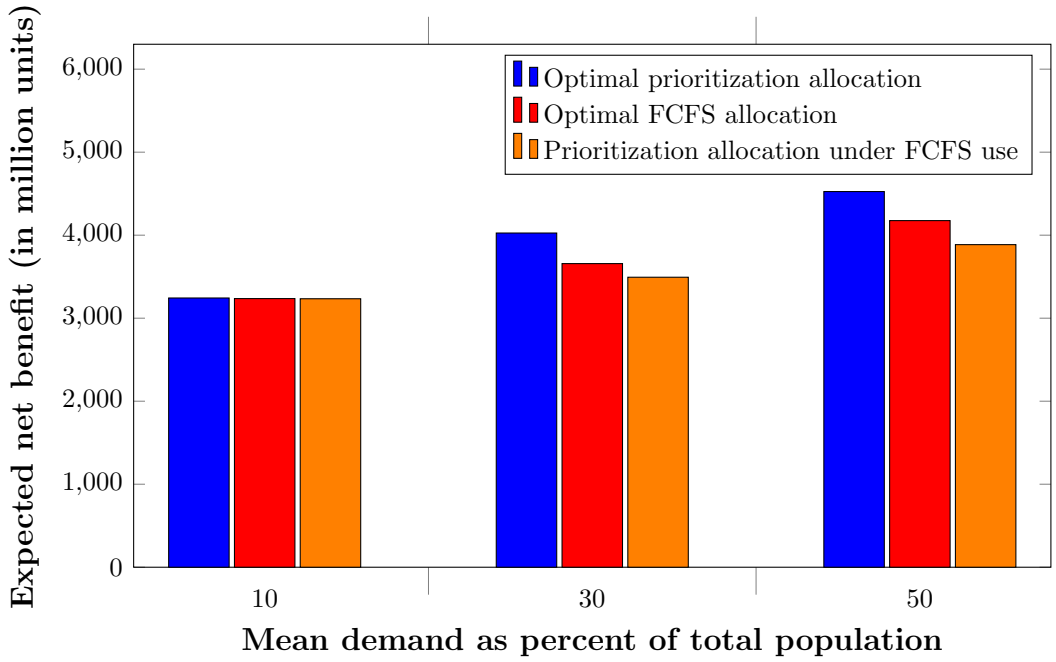

Figure 24: Total net benefit for the vaccine case with demand at county level and varying ratio of the rate of arrivals to total population

The first observation from Figures 23 and 24 is that **as the demand rate increases, so does the difference between the performances of prioritization and FCFS allocations**. When the demand rate is low (around 10% of the corresponding population), the capacity level, which is around 11.5% of the total state population, allows for allocating units to the distribution locations almost up to their optimal levels under the uncapacitated case. In this situation, the allocations under the optimal FCFS and prioritization usage policies are similar to each other, and result in comparable expected net benefits. However, as the demand rate is increased, capacity becomes an issue, and the advantage of prioritizing higher-class demand becomes more apparent. When demand is at the census tract level, the ratio of expected net benefit under FCFS usage to that under prioritization is more than 98.4% when demand rates are 10% of the population, and decreases to 82.5% when demand rate is 50% of the population.

When distribution locations are at the county level, expected benefits

under prioritization and FCFS are both higher, compared to those when distribution is at the census tract level. The expected benefit difference between the two policies decreases under centralization. The demand variance reduction due to centralization of inventory improves prioritization performance by an average of around 14% and improves FCFS performance by an average of 21%. An interesting observation is that although improvement in the prioritization performance decreases with increasing demand rate, it stays fairly constant for FCFS. This is mainly because prioritization allocations do not change significantly with increasing demand rate, and capacity restrictions prevent it from taking advantage of higher demand when the rate is higher. In contrast, FCFS allocations are significantly altered by the variance reduction, and hence even when capacity is binding, the benefits of centralization can still be reaped. When the demand rate is 10% of the population, the expected benefit performance ratio of optimal allocation under FCFS usage to prioritization is 99.8%, and the ratio only decreases to 93% when demand rate increases to 50%.

Similar results are observed when capacity is varied. Figure 25 shows the expected net benefits for optimal allocations under prioritization and FCFS usage policies, as well as prioritization allocation under FCFS usage when capacity is varied as 11.5% and 40% of the total population. As before, *when capacity is increased, the optimal prioritization and FCFS allocations are more similar to each other*. In particular, when capacity is 40% of the total population, the expected benefit performance ratio of optimal allocation under FCFS usage to prioritization is around 99.7%.

The next set of experiments we consider relate to the use of different methods at different centralization levels. Figure 27 show the expected net benefits under the optimal method, the λ_1- and the λ_t-policies at the census tract and county levels for the blood case.

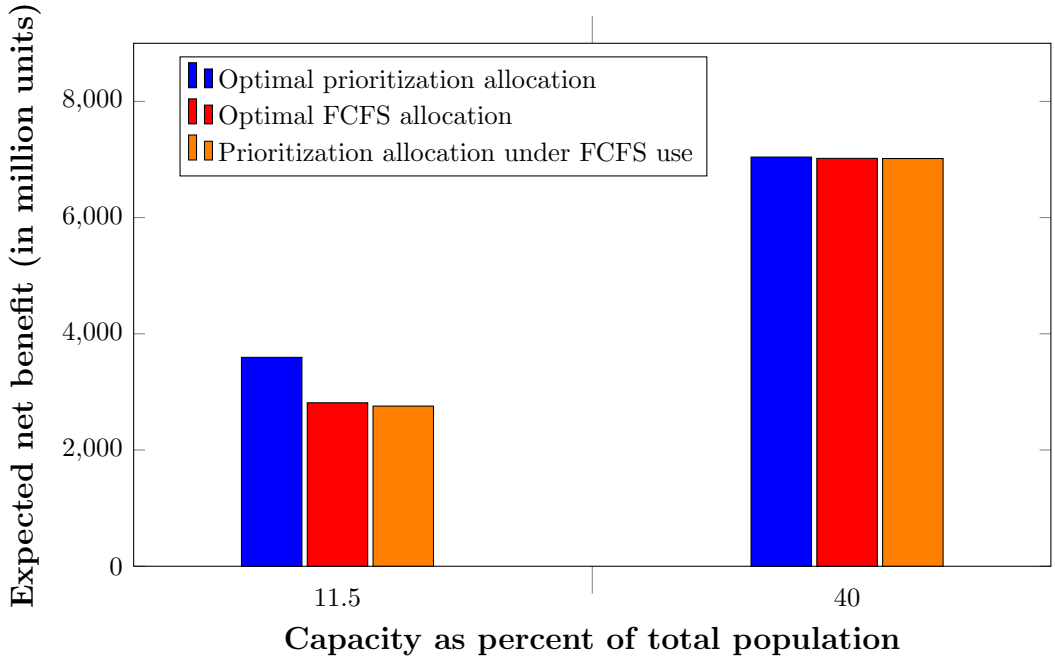

Figure 25: Total net benefit for the vaccine case with demand at census tract level and varying capacity as percent of the total population

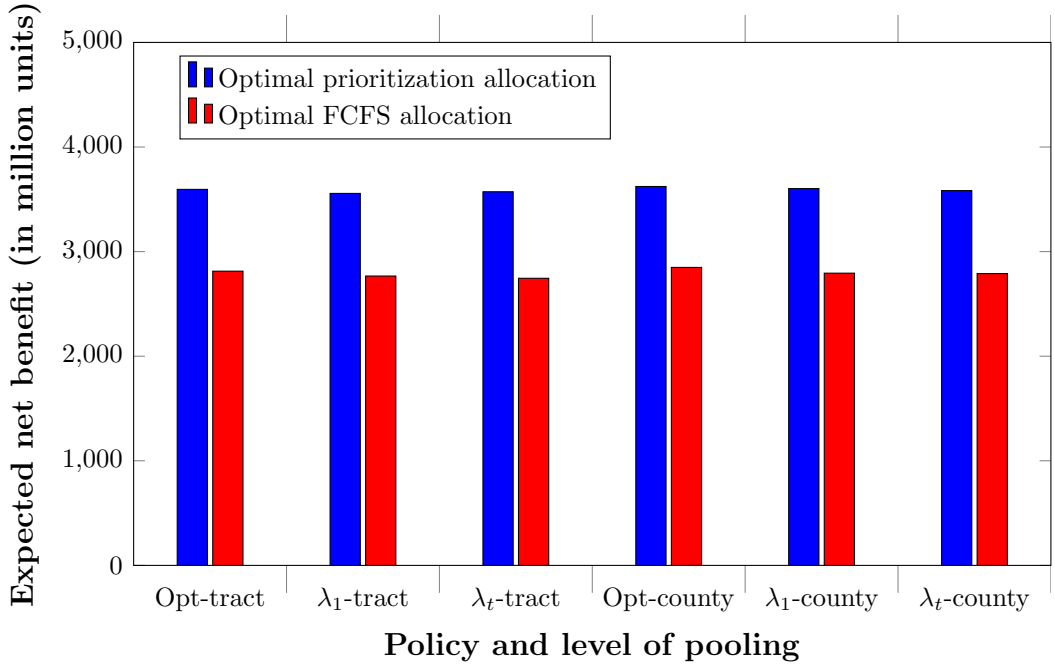

Figure 26: Total net benefits for the baseline case with varying methods and centralization level

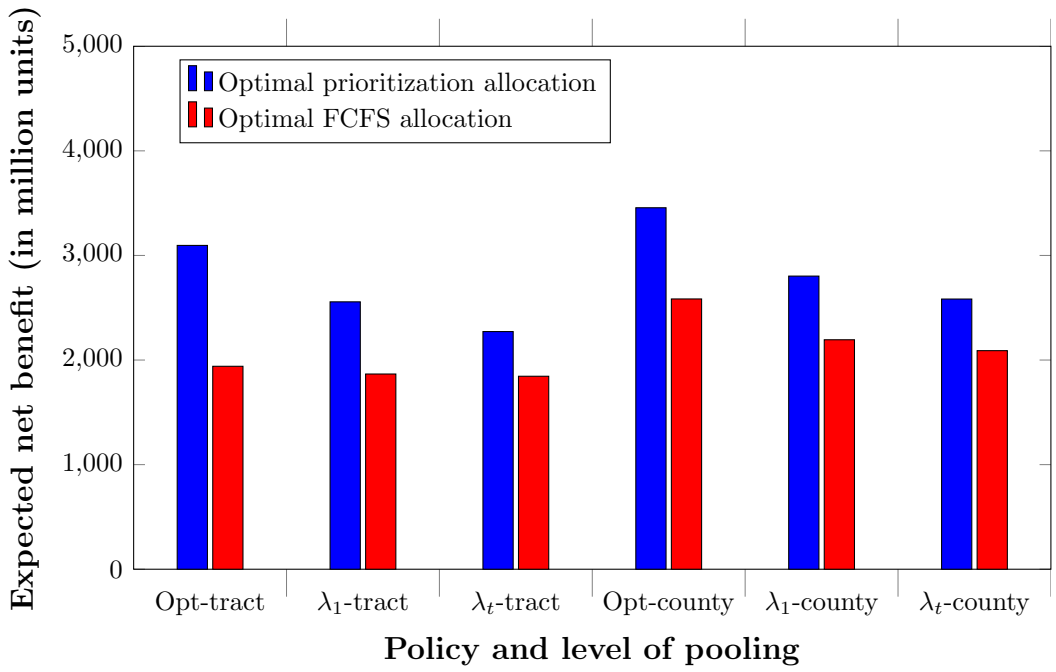

Figure 27: Total net benefits for the blood case with varying methods and centralization level

The first observation we have on the effect of different methods on system performance is that for the base case, that is, **when the benefit loss due to unused units is low, the λ_1- and λ_t-policies perform very close to optimal under both prioritization and FCFS usage**, as can be observed from Figure 26. This is mainly due to the fact that both optimal and benchmark policies allocate almost all of the units on hand due to low w_0. Furthermore, as also observed in Section 3.4, the optimal prioritization allocation is mainly driven by the λ_1 values at the distribution locations, and the optimal FCFS allocation is driven by the relative ratio of λ_1 to λ_t. Hence, the benchmark policies result in very similar allocations to those by the optimal policies.

The performances of benchmark policies start to decrease when unit benefit loss due to unused units is higher, which can also be observed from Figure 27. This is because of the increase in w_0 resulting in more conservative allocations for the optimal policies, whereas the benchmark policies allocate all units

on hand, which leads to significant expected benefit loss. As expected, prioritization performance of the λ_1-policy is higher compared to that of λ_t (around 82% of optimal benefit for λ_1 compared to around 74% for λ_t on average). For FCFS, the performances are closer (around 90% of optimal benefit for λ_1 compared to around 88% for λ_t on average).

The last discussion on our computational experiments considers the issue of equity of allocations under varying allocation policies and centralization levels. As mentioned before, an equitable system within the context of this study is defined as one in which all beneficiaries in a given class have equal probability of being units allocated to them. Hence, as a first measure of equity, we consider the percentage of higher-class beneficiaries whose demand is expected to be satisfied over all distribution locations. In Figure 28, one standard-deviation intervals for percentage of λ_1^i satisfied over all distribution locations i are given for varying centralization levels and allocation policies. These intervals are calculated as follows: For each location i, we divide the allocated quantity by the corresponding λ_1^i value. The intervals reported in the chart are the average over all locations and one-standard deviation intervals from this average. Similarly, Figure 29 shows the one-standard-deviation intervals for percentage of $\lambda_1^i + \lambda_2^i$ satisfied over all distribution locations i for varying centralization levels and allocation policies.

By our definition of equity, cases with higher averages have a better overall performance, whereas lower standard deviation implies better equity. From Figures 28 and 29, we observe that due to capacity being tight, almost all units on hand are allocated the optimal policies under both usage policies. Consequently, the average percent of λ_1^i and $\lambda_1^i + \lambda_2^i$ is almost identical for all allocation policies and centralization levels. However, there is a significant difference between (i) applying prioritization as opposed to FCFS, and (ii) satisfying demand at the county level as opposed to at the census tract level. For λ_1, switching from FCFS usage to prioritization reduces the

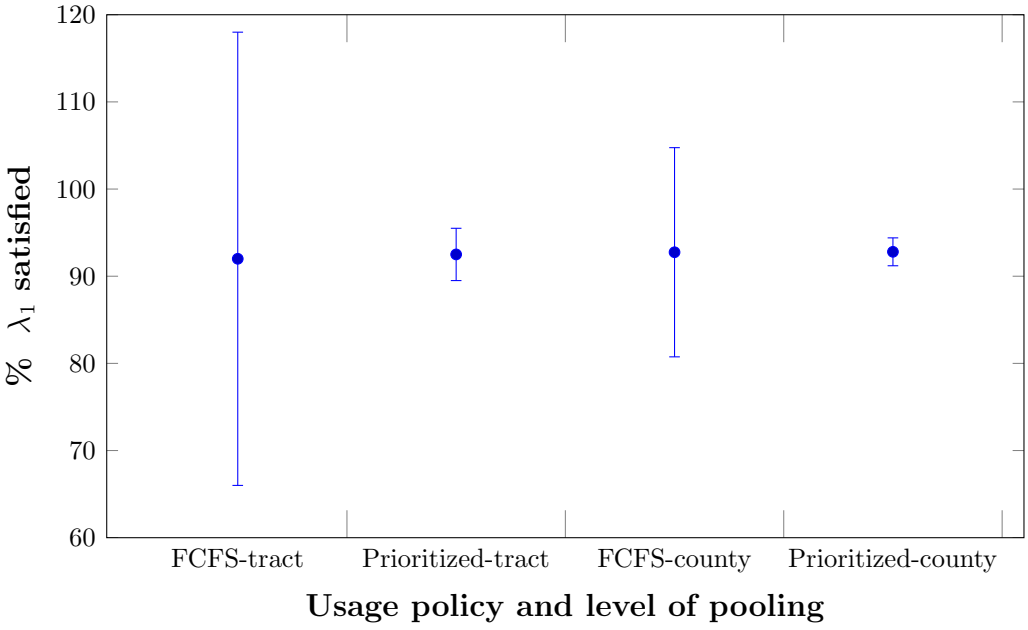

Figure 28: One-standard-deviation intervals for percentage of λ_1^i satisfied over all distribution locations for varying centralization levels and allocation policies

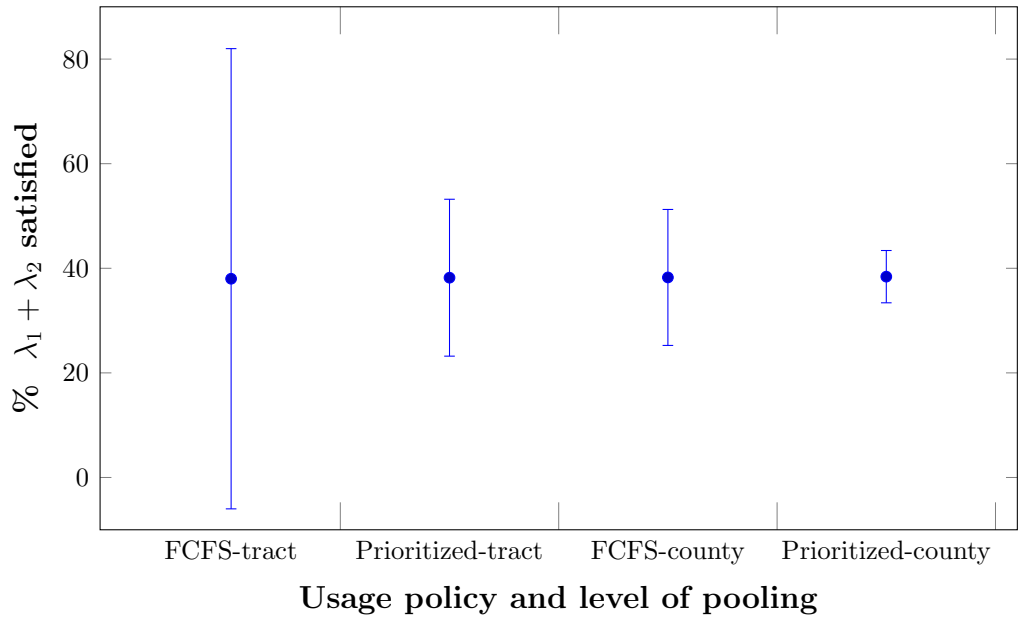

Figure 29: One-standard-deviation intervals for percentage of $\lambda_1^i + \lambda_2^i$ satisfied over all distribution locations for varying centralization levels and allocation policies

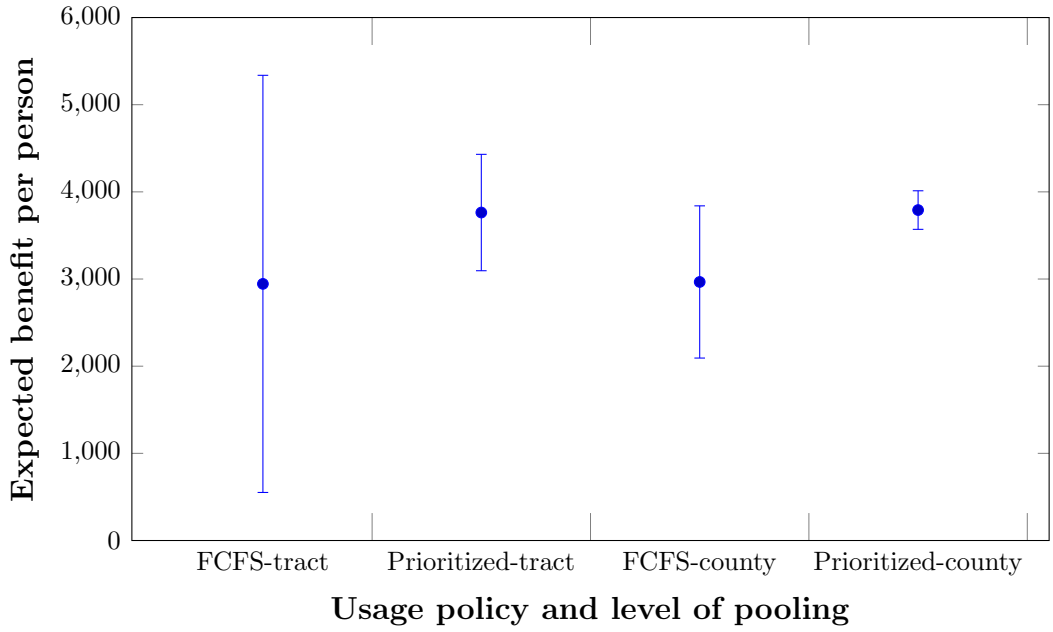

Figure 30: One-standard-deviation intervals for the expected benefit per unit demand satisfied

standard deviation from around 25% of λ_1 to less than 2.5% at the census tract level, and from around 11% to less than 0.5% at the county level. Similarly, centralizing demand while keeping usage policy constant reduces standard deviation from around 25% to 11% for FCFS usage, and from around 2.5% to less than 0.5% under prioritization. For this specific case, change of usage policy is more effective in increasing equity than centralizing demand. Similar results apply for $\lambda_1 + \lambda_2$, except that for this case, the difference between changing usage policy and changing centralization level is not that significant.

The way we define our first equity measure can be misleading especially for the case of satisfying higher priority demand, as covering the same percent of λ_1^i under prioritization and FCFS does not imply that they will satisfy an equal amount of urgent demand. Motivated by this, a second equity measure we define focuses on the expected benefit per unit demand satisfied over all locations. Figure 30 shows the one-standard deviation intervals for each policy and centralization level pair. In this case, average performance is almost identical when switching from census tract level

to county level. However, **the expected benefit per person increases** by around 1,000 units **when switching from FCFS usage to prioritization**. Furthermore, **equity increases both when switching from FCFS usage to prioritization and from census tract level to the county level**.

3.6 Conclusions and Further Research Directions

In this chapter, we have considered the allocation of a perishable commodity from a capacitated warehouse to multiple distribution locations that serve different classes of beneficiaries, based on the urgency of need for the commodity. The usage policies have significant influence on determining the allocation schemes from the upstream level. In particular, we have focused on two different ways of satisfying demand, namely prioritizing higher class demand, and serving the demand based on the order of arrival. For both cases, we have shown that myopic algorithms that make use of structural results from the uncapacitated case can make the optimal allocations.

Systems that prioritize higher-class demand result in higher expected net benefits compared to those under FCFS usage. Given the unit benefits and demand rates, we have established worst-case bounds on the relative performance of FCFS to that of prioritization, for both single-and multiple-location cases. To the best of our knowledge, this is the first such comparison of these two usage policies for the newsvendor problem with multiple echelons and multiple demand classes. These results are also useful for determining whether it is worthwhile to switch to a prioritization scheme given the system parameters.

Our structural analysis on serving part or all of the demand classes from a single location has shown that if capacity is sufficient, serving all demand from a single location always benefits the system, both under FCFS and prioritization usage. For capacitated systems, it might be possible for a decentralized system to perform better under FCFS usage, while this is not the case with prioritization. When part of the

demand classes are served from a single location, the system performance may get worse, even under prioritization or sufficient capacity. Our results bear resemblance to those under pooling in queues, particularly under FCFS usage. Among other findings, our computational experiments have quantified the effects of applying a FCFS usage during the 2009 H1N1 vaccination campaign as well as the effects of information asymmetry about the usage policies on the allocation decisions. We have also shown the importance of centralization and prioritization of demand for improving the equity of allocations.

Our work in this chapter leads to a number of interesting future research directions. The structural proof of the experimental findings in Observation 5 will strengthen the connection of our results to pooling in queuing theory. While our structural results on centralization shed light into when a centralized or the original decentralized system performs better, the extent of this performance difference is open to analysis. Such an analysis can also be performed for the analysis of information asymmetry on the usage policy between the warehouse and the distribution location. Lastly, real-time allocation of the commodity in multiple periods can be modeled using Markov decision processes. This also allows the modeling of commodity age, which will extend the analysis in this chapter to cases where benefits are age-dependent. An application area for such a case is the distribution of fresh foods, where serving more fresh commodities to the beneficiaries results in higher benefits.

CHAPTER IV

A CASE STUDY ON THE SUPPLY CHAIN OF SPECIALIZED NUTRITIOUS FOODS

4.1 Introduction

Globally, the greatest risks on human health are posed by hunger and malnutrition, at a rate that is greater than those of AIDS, malaria, and tuberculosis combined [171]. Undernutrition, which refers to inadequate intake of nutrients for a sustained period of time, leads to irreversible and long-term consequences such as hampering of body growth and reduction in productivity. Currently, 870 million people around the world suffer from undernutrition, 98 percent of whom live in the developing countries [73].

Undernutrition mainly occurs in three different forms: *wasting*, *stunting*, and *micronutrient deficiencies*. Wasting (also referred to as acute malnutrition) is associated with a low weight-to-height ratio. If this ratio is below the level of three standard deviations less than the global average (defined by the World Health Organization standards), an individual is *severely wasted*. A ratio between two and three standard deviations less than average indicates *moderate wasting*. Similarly, stunting (also known as chronic malnutrition) is indicated by a low weight-to-age ratio, which can be moderate or severe, defined in the same way as for wasting. Micronutrient deficiency occurs when the intake of vitamins or minerals is below acceptable levels. Undernutrition can be caused by factors such as illness, inadequate food intake, food insecurity, and poor health environment [172]. Newborn infants, young children and lactating mothers in developing regions of the world are particularly vulnerable to undernutrition, prevention and/or treatment of which requires access to nutrient-dense foods.

In this chapter, we focus on the activities of a humanitarian organization that aims to fight hunger and malnutrition worldwide. Currently, the organization works in and manages the food distribution supply chains of more than 90 countries, and reaches millions of beneficiaries in need by using trucks, planes and ships. To treat or prevent undernutrition, it works with the Food and Agriculture Organization of the United Nations (FAO) and the United Nations Children's Fund (UNICEF) as well as smaller non-governmental organizations. Within this collaboration, it particularly focuses on the prevention and treatment of moderate wasting, prevention of stunting, and prevention of micronutrient deficiencies. To this end, the organization uses a set of blended foods that contain a broad set of macro- and micro-nutrients in dense amounts, generally referred to as *specialized nutritious foods (SNF)*.

The SNF products of the organization mainly consist of three categories of foods. *Fortified blended foods (FBF)* are prepared by mixing cereals (generally corn, wheat, or rice) with other ingredients (mainly soya), and blended, cooked and fortified with certain vitamins and minerals. *Lipid-based nutrient supplements (LNS)*, more popularly known as *ready-to-use therapeutic foods (RUTF)* are ready-to-eat lipid-pastes that provide energy, nutrients, and fatty acids. *Micronutrient powders (MNP)* are vitamin-mineral mixes packed in small doses. Each of these categories differ in terms of shelf life, commodity value, addressed undernutrition category, and required dosage. More detailed information about the SNF products of the organization is given in Appendix C.

Specialized nutritious foods have relatively higher value and lower shelf life compared to the other commodities provided by the organization. This deems the expiration of these commodities very costly. Furthermore, the remoteness of the suppliers from the destination countries increases the lead times, and hence the risk of expirations. Commodity procurement is mainly dependent on the availability of funding from donors and governments, which is usually made in irregular intervals and in

large amounts. In turn, commodities become available for distribution irregularly and in large batches, which, coupled with limited shelf life, might lead to mismatches of supply and usage. Considering all these factors, an efficient allocation scheme that considers commodity age and the mismatches between supply and demand is important to resolve the trade-off between satisfying as much demand as possible by storing the commodity in larger amounts and avoiding expirations by storing at the demand locations as little as possible.

SNF supply chains in various countries might involve different network structures depending on factors such as the types of SNF distributed, proximity of supply, available corridors for transportation, demand structure, and so on. In particular, decisions such as whether or not to store the commodities in intermediate storage locations or use intermediate hubs as crossdocks play an important role in determining effective allocation and transportation strategies, and determine the level of shortages and expirations. Therefore, the overall performance of the supply chain depends on the interdependence between the network structure and allocation/transportation decisions.

Unlike a for-profit food distribution supply chain, where supply arrivals are more regular, the irregularity of SNF arrivals due to funding availability might lead to supply shortages during certain periods in the planning horizon. Under such circumstances, an important concern is to distribute the commodities to the beneficiaries in an *equitable* manner, since favoring some of the beneficiaries over the others to maintain cost-effectiveness may not be socially desirable. Considering equity of distribution as the main objective has two main implications by increasing (i) the system-wide costs to maintain the same overall system level, and (ii) the complexity of the modeling process.

The organization has observed a steady increase in the usage of SNF, with a more than 100% increase in annual shipments from 2009 to 2012. Given the challenges

presented by the characteristics of SNF on the supply chain, it aims to determine an efficient and effective distribution system by considering network design, commodity allocation, and transportation decisions and equity as an additional consideration.

The increased global usage of SNF has resulted in the publication of guidelines for the supply chain management of these commodities by various organizations. Recently, the United Nations World Food Programme (WFP) has published a set of guidelines to provide comprehensive guidance for "improving planning, procurement, and delivery of nutritious foods to beneficiaries; and maintaining the nutritional quality of SNF by promoting optimal control during transport, storage, stock management and distribution" [172]. While these guidelines present a detailed set of recommendations throughout all stages of the supply chain, there is lack of quantitative models and decision support for network design, allocation, and transportation decisions of specialized nutritious foods. As we shall describe in detail in the next section, the lack of quantitative decision support results in inefficiencies in the supply chain performance, which include increased lead times until reaching the beneficiaries, expirations of commodities both at the upstream and downstream locations, and concurrent shortages. In this chapter, we present a case study on the supply chain management of the SNF that aims to bridge this gap by developing mathematical models and analyzing the effects of various aspects of the supply chain on these decisions. In particular, our aim is to analyze the combined effects of network structure and allocation strategies on the cost-effectiveness and equity of distribution considering the perishability of commodities, transportation capacities, and uncertainty of demand.

The remainder of this chapter is organized as follows: In the next section, we describe a typical SNF supply chain along with a detailed analysis of the aforementioned challenges, and specify the aims of this case study in more detail. §4.3 reviews the relevant literature and provides the contributions of this chapter, whereas §4.4

presents the mathematical models addressing the network design and commodity allocation decisions, along with proposed solution approaches. In §4.5, we discuss our computational experiments before concluding the chapter in §4.6 and pointing to further relevant research directions.

4.2 A Typical SNF Supply Chain and Challenges in Supply Chain Management

4.2.1 A Typical SNF Supply Chain Structure and Flow of Events

While country-specific conditions may result in slight differences in the processes, the sequence of events in all SNF supply chains managed by the organization is similar. Demand can arise in two different situations. *Country development programs* aim to feed the undernourished population under non-emergency situations and have relatively stable demand. *Emergency programs* aim to satisfy food needs resulting from extreme events such as famines, droughts, or earthquakes, and may cause significant surges in usage. For both cases, *country offices* need to determine forecasts on a rolling horizon basis and inform the headquarters (we use the terms *demand* and *usage* interchangeably to refer to the needs of a distribution location in a given period). Procurement for a country needs depends on these forecasts as well as the availability of funding. Most common sources of funding are governments and donations from international non-governmental organizations. In general, funds are made available in irregular intervals and in bulk amounts. When need arises and funding is available, a *procurement requisition* is issued to initiate the procurement of commodities addressed for a specific country. Once funds are released and procurement is authorized, an appropriate supplier is selected and notified using a *production order*, and manufacturing starts.

Most suppliers of SNF are international companies with plants located in Europe and North America. For RUTF, the commodities are patented and can only be produced in certified plants, all of which are located overseas from the target countries.

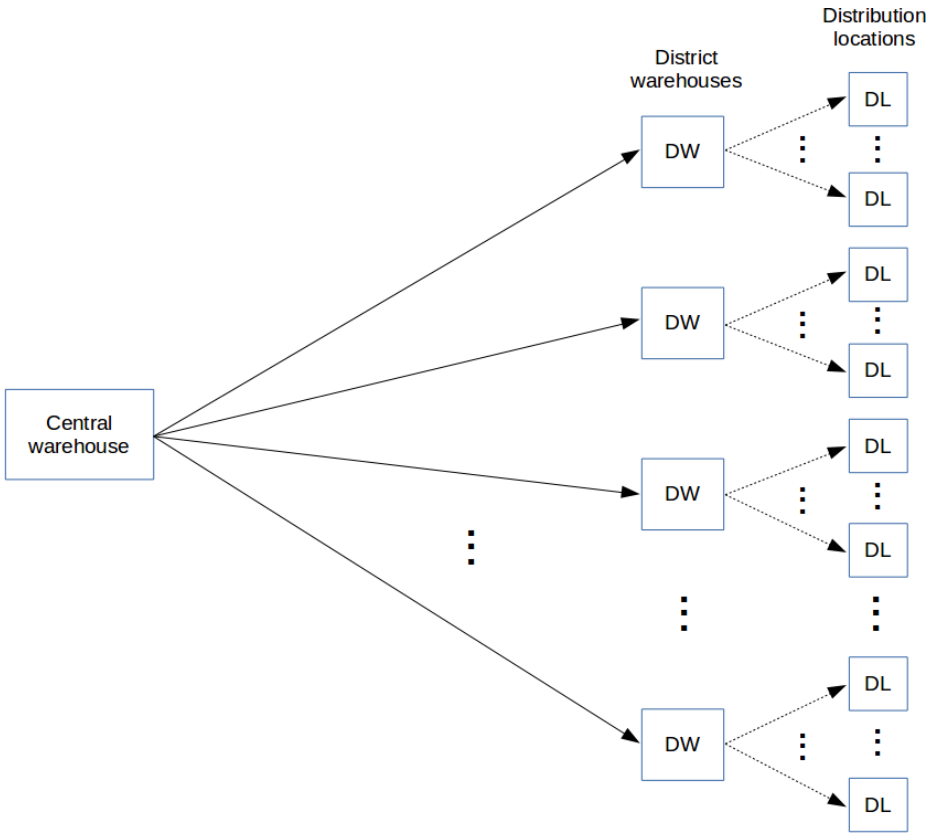

Figure 31: The SNF supply chain structure in a recipient country

For FBF, only limited local production is available. Reasons for lack of local production include lack of economies of scale to justify cost-efficiency, the need to import certain nutrients from overseas, lack of arable land to cultivate the ingredients, and unsuitability of the climate. Overall, more than 90% of all historical SNF deliveries have been made from overseas. In the case of an overseas delivery, the supplier is responsible for on-time manufacturing and quality maintenance of the commodities, as well as initiating the shipment from the *port of origin*. In general, the ownership of the commodities is transferred to the buyer at the *port of arrival*. Here, the customs process takes place, and if the port of arrival is not in the recipient country, an *overland shipment* takes place to transport the commodity to the recipient country.

Figure 31 shows the general SNF supply chain structure in a recipient country. In general, shipments to a recipient country are first transported to a *central warehouse*

(CW), usually located in the capital city or the largest port. The central warehouse makes allocations and transfers the commodities to the *district warehouses (DW)* through a predetermined *corridor* based on stock availability at the CW and DWs, forecasts at the DW level, lead times, and availability of transportation capacity. Each district warehouse then allocates these commodities to *distribution locations (DL)* based on the same factors at the DW and DL levels. District warehouses can store SNF and are located in close proximity of distribution locations. As a result, deliveries from DWs to DLs are more frequent with shorter lead times compared to those from the CW to DWs.

Distribution locations are generally owned and operated by partners. Upon arrival at the distribution locations, the responsibility of the organization in the logistics process ends. Hence, distribution of the commodities to the beneficiaries and storage of the commodities at the DLs is carried out by the partner organizations. However, the organization may provide assistance in transportation, regularly receives forecasts from the DLs, and receives usage data from some of the DLs.

4.2.2 Challenges in SNF Supply Chain Management

Compared to the other foods provided by the organization, specialized nutritious foods have higher product value. As can be observed from Figures 46 and 47 in Appendix C, the cost of fortified blended foods ranges from $570 to $630 per metric ton, and RUTF cost around $3,700. Furthermore, the commodities have no alternative use in the case of an expiration, and are disposed of using special procedures, which incur additional costs. Therefore, it is important that the commodities are delivered to beneficiaries on time to avoid the costs of expiration.

SNF have shelf life ranging from 12 months for FBF to 24 months for RUTF, which are much shorter compared to the other foods delivered by the organization

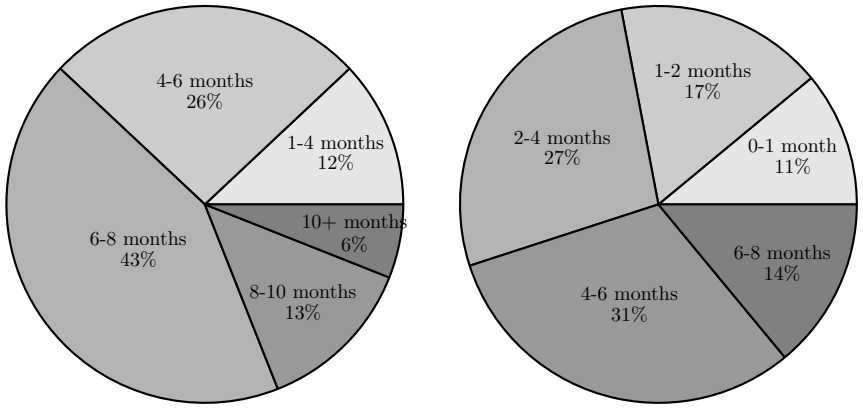

Figure 32: Remaining shelf life of corn-soya blend shipments upon arrival to the central warehouse (left) and distribution locations (right) in a landlocked country

(for example, wheat and maize can be stored for more than 10 years under appropriate conditions). This, together with the fact that most commodities are mostly supplied from overseas, result in long lead times and tight time windows to deliver the commodities to the beneficiaries. Figure 32 provides an example for the case of corn-soya blend shipments (with a shelf life of 12 months) to the CW (left) and DLs (right) in a landlocked country from January 2012 to July 2013. The data for the central warehouse shows that on average, the lead time to the CW is around five months, which includes the shipment lead time until reaching the port of arrival, customs delays at the arrival port, and the lead time for overland shipment. Delays of the customs procedures at the entry ports sometimes require the commodity to await approval for nearly a month. Transport lead times increase throughout the year during the rainy season. Furthermore, for countries undergoing conflicts, lack of corridor security may lead to transportation over longer routes, further increasing the lead times. Current transport operations do not involve tracking of best before dates of the commodities, which consequently leads to some shipments arriving at the DL very close to their expiration date. Consequently, around 11% of commodities arriving at the distribution locations have less than a month before being delivered to beneficiaries, and the average lead time until reaching the DL is almost eight months.

Long lead times lead to requirement that the commodities be consumed within a short period of time after reaching the DL to avoid expirations. However, this is only possible if there is sufficient demand for usage at the DL. To determine if this is the case, Figure 33 compares the monthly corn-soya blend arrivals into the central warehouse in the same recipient country to monthly aggregate forecasts through the period January 2012 - July 2013, whereas Figure 34 makes the same comparison using cumulative values. The general pattern for arrivals and usage can be observed from Figure 33. Supply arrivals are usually in bulk amounts and arrival frequency is irregular over time. This is in line with the pattern for availability of funding, which is needed to initiate the procurement and production processes. Furthermore, suppliers also tend to produce the SNF in bulk amounts to benefit from economies of scale. Usage forecasts, on the other hand, follow a regular seasonal pattern. These increase particularly during the lean season between November and May, since agricultural output is not sufficient to satisfy nutritional needs of the country during that period. Figure 34 shows the mismatch between commodity arrivals and usage, which implies that there will be periods where supply arrivals exceed usage and vice versa. As a result, there is risk of observing shortages and expirations throughout the planning horizon.

Figure 35 shows the fraction of on-hand corn-soya blend stocks expiring at each distribution location and the fraction of usage amounts that cannot be fulfilled for DLs that report actual usage in the same recipient country during the period January 2012 - July 2013. As a consequence of long lead times and supply-usage mismatches, almost half of the distribution locations have units expiring. While the average expiration rate is around 4.5%, for some DLs the rate goes up to around 15%, which results in significant commodity loss. The average shortage level for the DLs that report actual usage is around 10% and only eight of the 165 usage reporting DLs can completely satisfy the demand. These findings underline the importance of an effective allocation

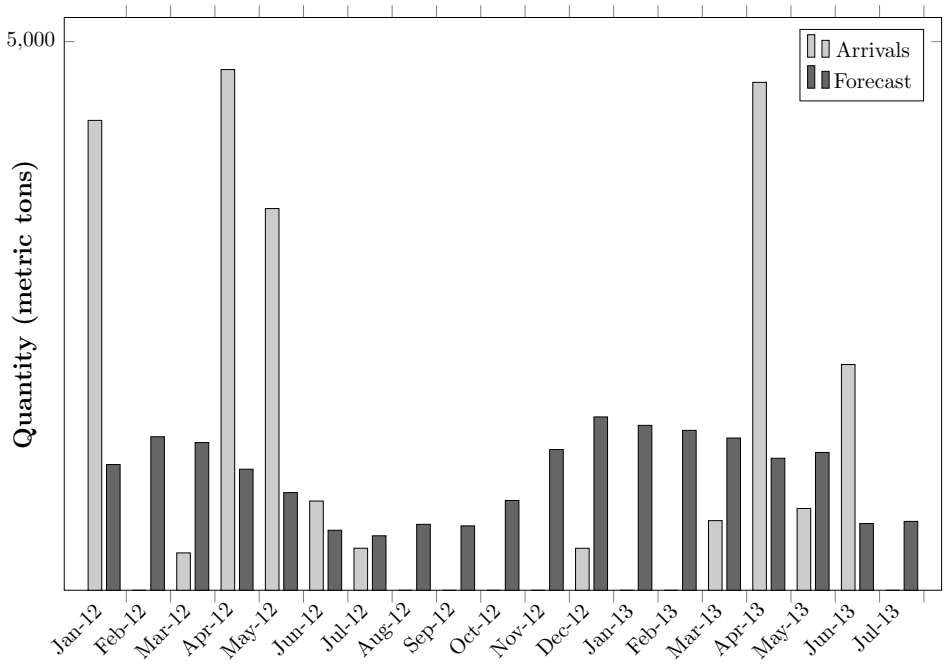

Figure 33: Comparison of monthly arrivals into the central warehouse to aggregate monthly forecasts at district warehouse level for corn-soya blend, data from January 2012 to July 2013

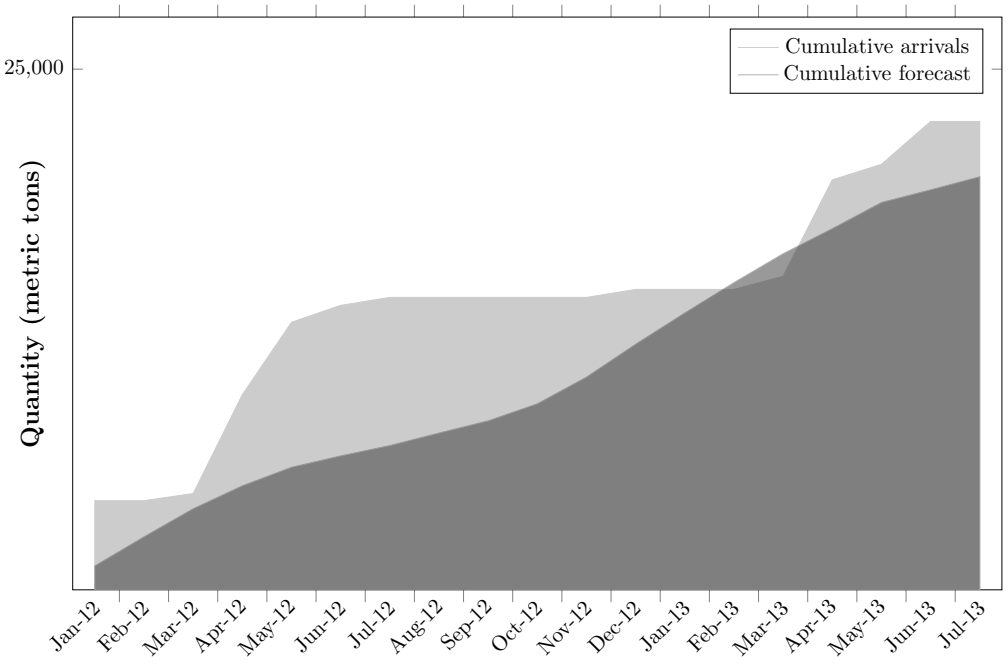

Figure 34: Comparison of cumulative monthly arrivals into the central warehouse to cumulative aggregate monthly forecasts at district warehouse level for corn-soya blend, data from January 2012 to July 2013

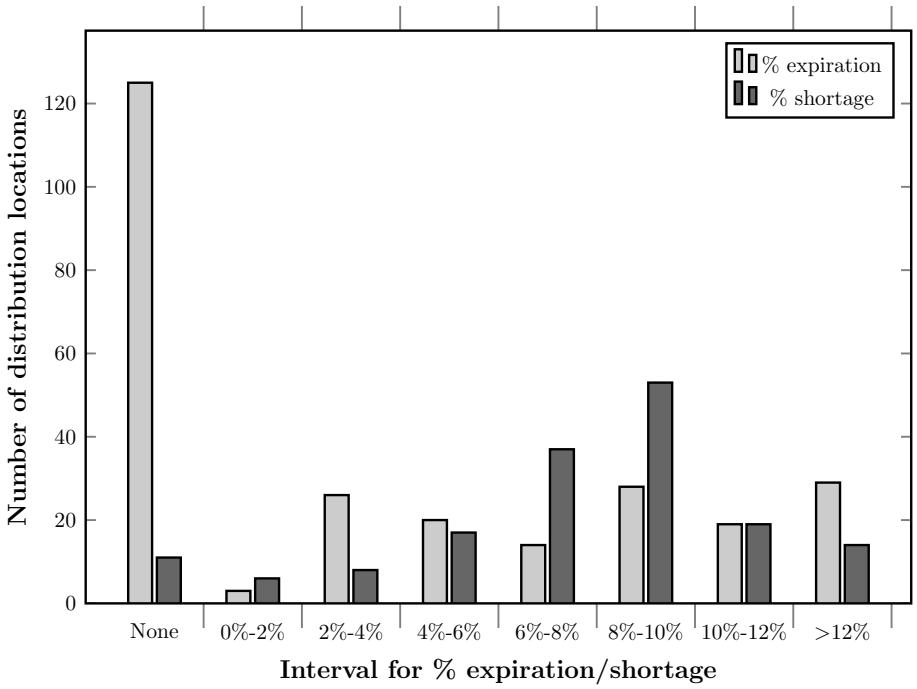

Figure 35: Histogram of expirations as % of on-hand CSB inventory over all distribution locations and shortages as % of demand for DLs that report usage

scheme which takes into account the commodity ages and considers the trade-off resulting from the supply-demand mismatch between satisfying as much demand as possible by storing the commodity in larger amounts and avoiding expirations by storing at the demand locations as little as possible.

Usage forecasts at the distribution location level are made by the partners, and the organization provides guidance to the partners on how to forecast demand. However, there is limited historical data to evaluate the forecasts, as many distribution locations do not report actual usage levels. Figure 36 shows the mean absolute percent error (MAPE) values for the 165 DLs that report actual corn-soya blend usage amounts for the period January 2012 - July 2013. The MAPE values are calculated by dividing the absolute value of the difference between the actual and forecast amounts by the actual usage value, and taking the average over all periods. The average MAPE over all DLs is around 22%, which implies that usage forecasts are not necessarily accurate. Ignoring the potential errors in forecasts is likely to result in allocations that do not

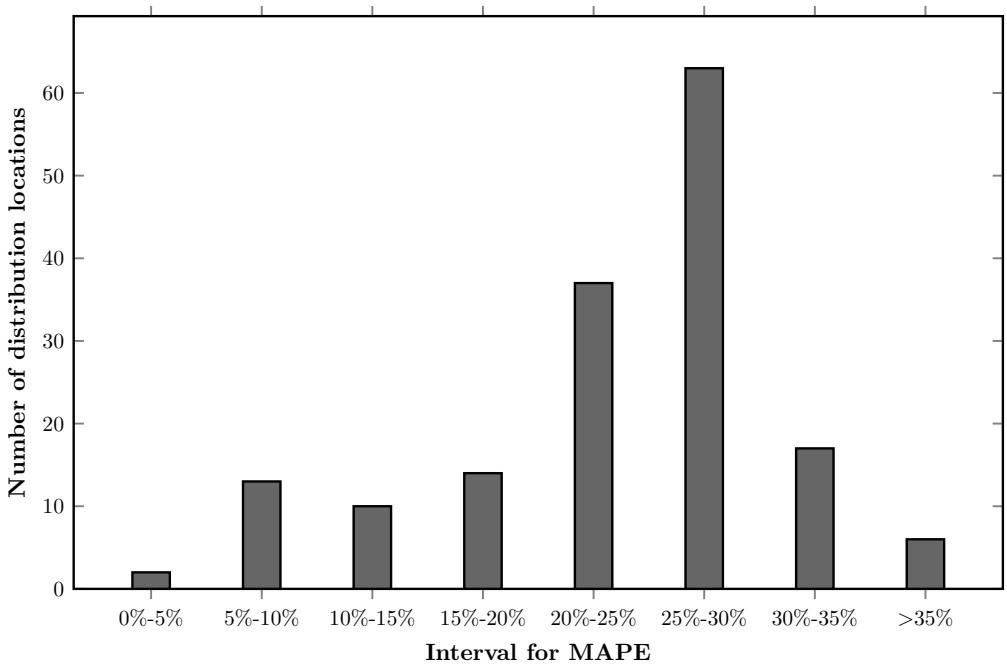

Figure 36: Mean absolute percent error intervals for usage forecasts by DLs that report actual usage

accurately match the demand, as exemplified by Figure 35. Therefore allocations of SNF must take into account the uncertainty in demand resulting from these forecast errors.

An important factor that determines the level of shortages and expirations at the DLs is how the commodities are allocated. Currently, allocations of the commodities from the central warehouse to distribution locations do not consider specific commodity age and potential errors in the forecasts. There are mainly two reasons for this. First, while the information system potentially allows for batch tracking, such entries are not made, and hence it is impossible to track individual batches. Second, forecasts are made and evaluated by the partners that operate the DLs, but only part of the DLs report these evaluations to the organization. As a result, instead of evaluating historical data to assess the accuracy of future usage forecasts, allocations are made assuming these forecasts are completely accurate. To underline the importance of applying a more systematic approach for commodity allocation, we use a

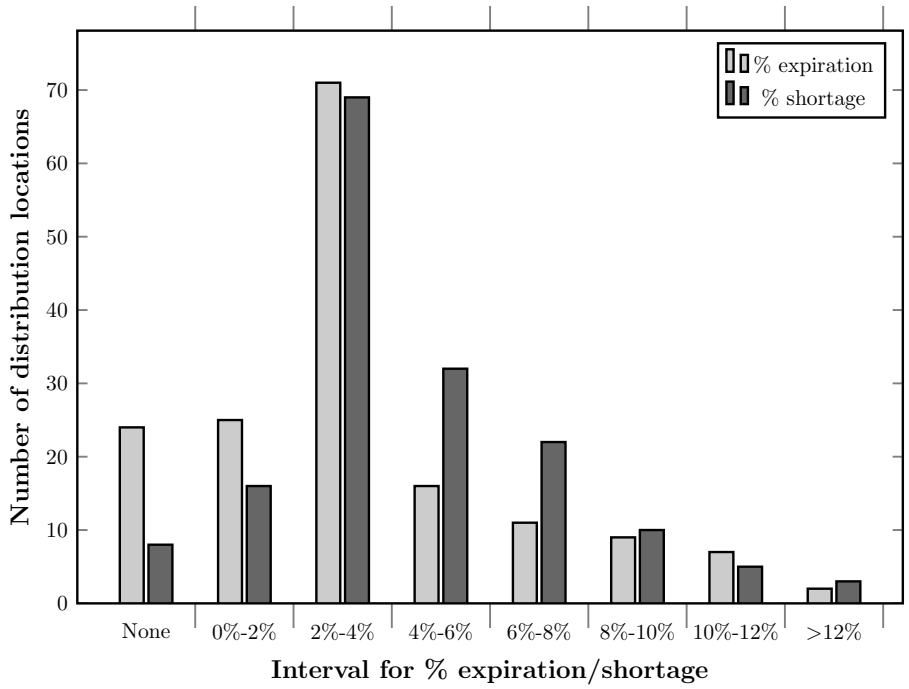

Figure 37: Histogram of expirations as % of on-hand CSB inventory under the simple allocation policy over all distribution locations and shortages as % of demand for DLs that report usage

simple allocation procedure that assumes deterministic demand equal to the forecast amounts at each DL. Under this scheme we use average historical lead times to calculate the time to transport the commodities from the central warehouse through DWs to the DLs, ignoring any vehicle capacities. We also assume a first-expiring first-out (FEFO) allocation. If there is sufficient supply, each DL receives shipments equal to the forecast amounts. Otherwise, allocation is made proportional to these quantities.

The resulting fractions of expirations and shortages under the proportional allocation scheme are given in Figure 37. Even under a simple systematic allocation approach and utilizing FEFO shipments, the average expiration and shortage rates decrease to 3.3% and 5.9% from the current values of 4.5% and 10%, respectively. Another important result from this analysis is that, as in the current situation, the mismatch between supply and usage leads to expirations and shortages occurring in the same DLs. This further underlines the importance of tracking commodity age

when making the allocation decisions.

4.2.3 An Alternative Network Structure with Intermediate Hubs

The effectiveness of the allocation decisions is significantly affected by the structure of the distribution network. In the existing network structure, commodity storage at the district warehouse level increases the lead time before reaching the distribution locations, since the commodities spend additional time while being stored at the intermediate DW level. Furthermore, allocations from the central warehouse only consider the aggregate demand forecasts from the district warehouse level, which in turn make allocations to the distribution location level based on these allocations. An alternative strategy, applied in various countries by other organizations, is to replace the district warehouses with hubs that do not store any commodities. Since these hubs act as crossdock points, the increases in the lead time due to storage at the intermediate level are reduced. Furthermore, this system allows allocation of commodities from the central warehouse directly to the DL level by incorporating complete visibility through the whole supply chain into the allocation process.

Figure 38 illustrates an alternative SNF distribution network with intermediate hubs that act as crossdock points between the central warehouse and distribution locations. When shipments arrive at the hubs, they are immediately separated by destination and loaded into smaller trucks that serve each distribution location. By avoiding storage at the district warehouse level, transportation lead time is expected to decrease compared to the existing system. However, since the hubs are located closer to the CW than the DLs compared to the existing system, this increased distance between the intermediate level and distribution locations implies that each of these smaller trucks have to travel longer distances. In addition, while a hub system allows direct allocation to the DLs, these allocations will be made based on disaggregated forecasts, which are less accurate compared to the aggregate forecasts at the

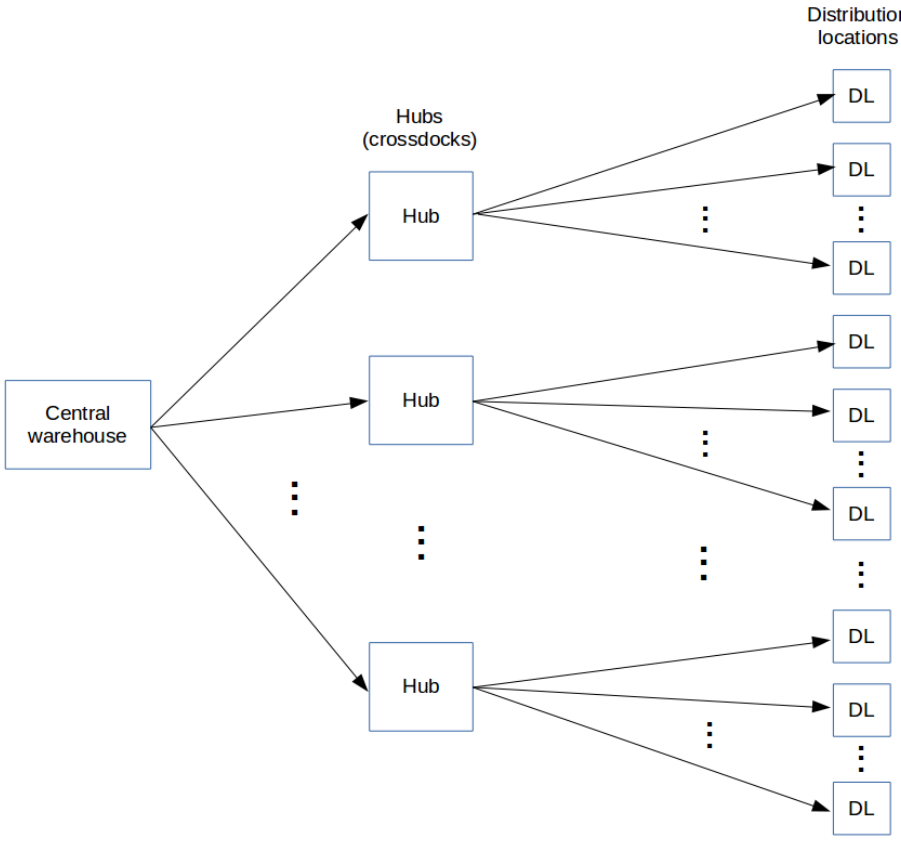

Figure 38: An alternative SNF distribution network with intermediate hubs that act as crossdocks between the central warehouse and distribution locations

district warehouse level. Intermediate storage at the district warehouses also allow hedging against uncertainty of SNF usage at the DL level. While it is possible to send more frequent shipments from the central warehouse to the intermediate level in a hub system, the existing system allows more frequent dispatches from the intermediate level to distribution locations. Hence, there is a trade-off between the complete visibility and reduced lead times of the hub system versus the more accurate forecasts, more frequent shipments, and the availability of buffer stocks at the intermediate level to hedge against uncertainty in the existing system.

4.2.4 Objectives of the Case Study

In this case study, we aim to quantify the trade-offs between the two network structures using mathematical models that incorporate the limited shelf life of the commodities, uncertainty in the usage amounts, and transportation capacities. Our main objective is to analyze the effects of network design (in particular, the existing system with district warehouses vs. intermediate hubs) and commodity allocation decisions on (i) the extent at which the demand is satisfied, (ii) amount of expirations, and (iii) system-wide costs. We also analyze the aspects of the system that lead to one network structure outperforming the other. By varying certain parameters of the system, we observe the effect of demand uncertainty (forecast errors) and commodity perishability on the performance measures. For the hub system, we assess the improvements resulting from choosing hub locations (as opposed to treating hub locations as given).

In many societal applications, cost-minimizing decisions may result in part of the beneficiaries (e.g., ones that are easier to reach) being favored over others. Therefore, a desirable aspect of the decisions in such applications is the level of equity at which the beneficiaries are served. Within our context, one measure of equity is maximizing the minimum fraction of demand served over all distribution locations. In general, models with the objective of maximizing equity are harder to solve. To overcome

this issue, we develop novel heuristic approaches that find good quality solutions and compare these to those under a cost minimization objective.

The research questions we aim to answer in this case study can be summarized as follows:

- How do the different network structures affect allocation of the commodities?
- Which aspects of the system play an important role in determining the allocations under each network structure?
- Under what circumstances does one network structure perform better compared to the other?
- How do usage uncertainty and commodity perishability affect the system costs?
- How does the hub system perform relative to the existing system if hub locations are part of decisions? How does it perform if hub locations are given?
- If the main objective is to provide equitable distribution instead of minimizing costs, how do the performances change?

We aim to answer these questions by proposing models that simultaneously address the network design, commodity allocation, and transportation decisions, and applying these models in computational experiments that are based on historical SNF shipment data.

4.3 Literature Review

In this chapter, we consider a perishable commodity allocation problem with emphasis on the design of the distribution network. The supply chain consists of multiple echelons, and commodity demand is stochastic. Considering these aspects, our work is relevant to a number of streams in the literature. These include case studies in

commercial and humanitarian food distribution, inventory and transportation models in multiple echelons, location-inventory models in multiple echelons, perishable inventory models in single and multiple echelons, and studies on equitable commodity allocation in humanitarian supply chains.

Food distribution in commercial supply chains has been a widely studied area over the last decades (for a recent review of the literature in this field, see Akkerman et al. [1]). With the growing interest in health and humanitarian supply chains over the recent years, distribution of donated food in this context has been attracting an increasing level of attention. Recent studies in this stream include De Angelis et al. [52], who consider the weekly planning of emergency deliveries of the WFP in Angola by planes, Epstein et al. [66], where a combinatorial auction algorithm is used to determine the allocation of food to schools in Chile, Alvarenga et al. [4], in which a decision support tool is developed to model the operations of the WFP in the Horn of Africa, Consuelos Salas et al. [46], where the focus is on supplying food for a shelter following a hurricane by determining the order quantities, and Rancourt et al. [134], who consider the tactical operations of the distribution centers of the WFP in Kenya. A vast majority of the studies in this field consider a location- or disaster-specific situation, whereas as Altay and Green [3] also point out, there is also need for determining the characteristics of humanitarian food aid collection, allocation, and distribution problems in general. In this chapter, we aim to analyze distribution strategies on different network structures and aim to generalize our findings by considering data from multiple locations.

4.3.1 Case Studies on Commercial and Humanitarian Food Distribution

To the best of our knowledge, there exists a limited amount of studies focusing on case studies regarding humanitarian food distribution. An example is Jaspars [89], who presents a review of food distribution methods under conflicts in different regions of

the world, including the Somalia, Balkans, Zaire, and South Sudan. The paper mainly makes recommendations on how to develop and implement principled approaches to minimize food-related problems during conflicts. Tomasini and Van Wassenhove [157] consider the 2002 food crisis in Southern Africa and mainly discuss issues related to donations, in which case the food donations being genetically modified resulted in most of the donations being rejected. In a separate case study, Tomasini and Van Wassenhove [158] focus on the partnership between the WFP and TNT, a Dutch express delivery and logistics services firm, where TNT provided logistics services for the operations of the WFP in Africa, Asia, and Haiti. The most relevant work in this stream is the case study on a specific type of RUTF (Plumpy' Nut) by Swaminathan [151]. Our case study mainly differs from the work in this stream by comparing alternative network structures using quantitative models.

4.3.2 Multi-Echelon Inventory Management

There exists a vast amount of literature on inventory management with multiple echelons. Recent reviews of this stream are presented by Axsäter [12] and van Houtum [166]. Earlier work in this area mainly aims to structure the optimal inventory policies under simpler conditions such as unlimited supply and no transportation capacity, starting with the seminal work of Clark and Scarf [41], which proves the optimality of base stock policies for a serial system with a finite planning horizon. This result is extended to the case of infinite planning horizon by Federgruen and Zipkin [70] under normal demand. For general distribution systems, applying base stock policies may result in negative allocation to a subset of locations in certain periods [64]. For the cases where no negative allocation is made, i.e., when the *balance assumption* holds, these policies are still optimal [57]. When the balance assumption does not hold, the base stock policy gives a lower bound on the overall cost of the system [39]. Heuristics for the case when the balance assumption is not valid are proposed and analyzed by

Axsäter et al. [13] and Doğru et al. [59]. The main difference of our work from these studies is the perishability of the commodities and capacities on supply as well as transportation, which disqualify the optimality of policies proposed in this stream.

4.3.3 Joint Transportation and Inventory Management

Joint inventory and transportation decisions in multiple echelon supply chains have been addressed in a number of different ways. When demand is deterministic, the problem is an extension of the lot sizing problem. Pochet and Wolsey [127] propose a dynamic programming model to find the optimal solution for a single item, capacitated supplier, single retailer, capacitated vehicles, unlimited fleet size, and constant batch sizes. Lee et al. [101] and Anily and Tzur [7] partially characterize the optimal policy for the multiple retailer and multiple item cases, respectively. When demand is stochastic, an important assumption is that the supplier is uncapacitated. Çetinkaya and Lee [36] consider a system with intermediate hubs and propose an exact and approximate approach for the cases when the facilities are capacitated and uncapacitated. Çetinkaya and Bookbinder [35], Gürbüz et al. [80], and Çetinkaya et al. [34] propose heuristics for various extensions of this system under different cost structures. These include installation-based (s, S), echelon-based (order up to S when total demand exceeds Q), and constant time (order every T periods) policies, which are likely to be infeasible in our case due to limited supply.

4.3.4 Joint Facility Location and Inventory Management

We consider the joint decisions of facility location and inventory management when designing the network for the hub system. Daskin et al. [51] address the joint inventory-location problem under stochastic demand and fixed location, transportation, and inventory holding costs. They propose a Lagrangian relaxation for the case with constant coefficient of variation of demand, and a heuristic based on retailer reassignment, distribution center exchanges, and variable fixing. Snyder et al. [146]

and Özsen et al. [123] extend this approach to scenario-based demand and capacitated warehouses, respectively. Shen et al. [142] formulate a set covering model for the problem in [51], and use a branch-and-price approach. Shu et al. [144] use a similar set covering model for the non-constant coefficient variation case. Considering service level as a separate objective, Shen and Daskin [143] develop a weighted objective approach and an evolutionary algorithm. For networks with more than two echelons, Tancrez et al. [154] use a simple flow-fixing heuristic and Shahabi et al. [141] formulate a nonlinear program, transform it into a conic program and use commercial solvers to find the optimal solution.

4.3.5 Perishable Inventory Management

Inventory management of perishable commodities under stochastic demand poses additional challenges even in a single echelon, mainly since the optimality of simple policies such as base stock policies are no longer valid. Earlier work in this area (for a detailed review, see Nahmias [118]) aims to establish the structure of the optimal policies under different cost structures; Van Zyl [167] and Nahmias and Pierskalla [119] for two periods of shelf life, and Fries [75] and Nahmias [116] for general shelf life. Due to the complicated nature of these optimal policies, subsequent work in this area focuses on a heuristic policy and aims to find the optimal parameters of these policies. Brodheim et al. [26] and Cohen [43] find bounds on the stationary probabilities of the optimal order-up-to and constant batch size policies, respectively. Nahmias [115] shows the inferiority of the constant batch size policies, which directed focus to finding the optimal parameters of the order-up-to policies [102, 103, 117, 120]. More recent work in this area considers lead times [114, 169, 170], service level guarantees [114], and batch order sizes [27]. As for non-perishable commodities, limited supply in our case deems these optimal and heuristic policies infeasible. Furthermore, the work in this stream does not consider the costs and capacities of transportation.

Under multiple echelons, inventory models for perishable commodities focus completely on heuristic policies. Yen [174] and Cohen et al. [44] propose proportional and fixed allocation rules and show that order-up-to policies work well for two retailers. Prastacos [130, 131] characterizes the optimal policies under myopic allocation, which only considers the costs in the next period. Federgruen et al. [69] extend these ideas to include transportation costs and use a Lagrangian approach. Lin and Chen [104] propose a genetic algorithm to solve the more general allocation problem heuristically. Kanchanasuntorn and Techanitisawad [93] and Lystad et al. [110] use approximate approaches to find the parameters for order-up-to policies. Assuming deterministic demand and capacitated supply, Hemmelmayr et al. [83] and Coelho and Laporte [42] formulate models to find the optimal replenishment, routing, and inventory management policies in a finite horizon. The former uses a variable neighborhood search approach, whereas the latter proposes a branch-and-cut algorithm to find close-to-optimal solutions. Our work is similar to [83] and [42] in terms of limited supply and transportation considerations, but we focus on uncertain demand and do not consider specific routing of vehicles.

4.3.6 Optimizing Equity of Allocations

An important objective in planning for delivery of food commodities is providing equitable service to the beneficiaries. In the facility location and commodity allocation literature, equity has been modeled in a number of ways. Extensive reviews of these models are provided in Eiselt and Laporte [62] and Marsh and Schilling [111]. The most commonly used equity objective in commodity allocation is to maximize the minimum fraction of demand satisfied over all locations (e.g., [15, 105, 160]). Other equity measures in this literature include lexicographic approaches (e.g., [108, 122]), minimizing the Gini coefficient [60], and minimizing customer disutility [91]. Cao et al. [29] consider maximizing minimum fraction of demand satisfied in different

spatial levels (national, province, and municipality-based) and observe the effects on facility location and commodity allocation. To the best of our knowledge, our work is the first to consider equity for allocation of perishable commodities under stochastic demand and multiple periods.

4.3.7 Contributions of the Chapter

The contributions of this chapter are mainly three-fold: First, we contribute to the literature on food distribution case studies by quantifying the effects of various factors such as network structure, demand uncertainty, commodity perishability, and high commodity value on demand coverage, commodity expiration, lead time, and system-wide costs. Second, to the best of our knowledge, this study is the first to consider the allocation of (i) a perishable commodity, (ii) under uncertain demand, (iii) transportation capacities, and (iv) limited supply. Lastly, in addition to focusing on cost minimization as the main objective, we also consider the effects of prioritizing equity of distribution on the allocations, and develop a heuristic approach.

4.4 Mathematical Models and Solution Approaches

In this section, we discuss the details of how events unfold under the hub system and the existing district warehouse network structure, and present the mathematical models developed for each case.

It has been discussed in Section 4.2 that forecast errors (and hence demand uncertainty) are inherent in the SNF distribution system of the organization. To incorporate uncertainty into our models, we consider a set of *demand scenarios*, where each scenario specifies the amount of demand in each period of the planning horizon. Hence, our models consider two-stage decisions, where first stage decisions involve commodity allocations in the first period (and possibly hub location decisions), and the second stage decisions consist of allocations for the remaining periods depending

on which scenario occurs. Uncertainty (which demand scenario actually occurs) is revealed at the end of the first stage, i.e., first period of the planning horizon. The main aim in implementing this two-stage model is to determine the first period decisions considering the remainder of the planning horizon.

Revelation of the actual demand scenario at the end of the first period is a simplification of reality, because in practice demand forecasts and scenarios may change from one period to the next as new information dynamically becomes available. Hence, to incorporate the information updates, our models are solved on a *rolling horizon* basis as follows: At the beginning of the first period, each model is solved for a fixed planning horizon (e.g., 12 months), determining the allocations in each period (and hub locations, if applicable). The allocation decisions for the first period (and hub locations) are implemented based on the output of the model. Next, actual demand information becomes available for the first period and demand forecasts (scenarios) for the (remainder of the) planning horizon are updated, the model is solved with the updated input at the beginning of the second period. The allocation decisions for the second period are implemented and the procedure is repeated for each period until the end of the planning horizon.

In the first part of this section, we present our cost minimization models for the hub system and the existing structure. Next, we describe the models to prioritize equity of allocations, and our heuristic approaches to solve these equity-based models. The sets, parameters, and decision variables used throughout this section are given in Table 13.

4.4.1 Models for the Hub System

The hub system consists of a central warehouse that receives SNF stocks and allocates these to distribution locations via hubs that act as crossdock points. The locations of these crossdock points can be predetermined or part of decisions. The flow of events

Table 13: The sets, parameters, and decision variables used in the models

Sets	
R	Set of regions indexed by r
D	Set of distribution locations indexed by i
D_r	Set of distribution locations in region $r \in R$
J	Set of commodity ages indexed by j
S	Set of demand scenarios indexed by s
T	Set of periods indexed by t
Parameters	
p	Unit cost of shortage
w	Unit cost of expiry
γ_i	Per-vehicle cost of transporting commodities from the central warehouse to a hub located in $i \in D$
γ_{uv}	Fixed charge for transporting commodities from a hub located in $u \in D$ to distribution location $v \in D$ in the same region
τ_i	Transportation lead time from the central warehouse to distribution location $i \in D$
d_{is}^t	Demand at distribution location $i \in D$ under scenario $s \in S$ in period $t \in T$
k	Number of vehicles available in the central warehouse in each period
ξ_j^t	Amount of new arrivals of age $j \in J$ to the central warehouse from the supplier at the beginning of period $t \in T$
c_1	Capacity of each vehicle located at the central warehouse
c_2	Maximum amount of commodity that can be transported from a hub to a distribution location
q_s	Probability of observing demand scenario $s \in S$
Decision variables	
y_{uv}^t	$= \begin{cases} 1, & \text{if there is a shipment from a hub in } u \in D \text{ to distribution location } v \in D \text{ in period } t \in T \\ 0, & \text{otherwise} \end{cases}$
z_i	$= \begin{cases} 1, & \text{if a hub is open in } i \in D \\ 0, & \text{otherwise} \end{cases}$
r_i^t	Number of vehicles dispatched for distribution location $i \in D$ in period $t \in T$
v_{ij}^t	Amount of commodity of age j allocated to distribution location $i \in D$ at the beginning of period $t \in T$
l_{ijs}^t	Amount of demand at distribution location $i \in D$ satisfied from the stocks of age $j \in J$ in period $t \in T$ under scenario $s \in S$
I_{ijs}^t	Inventory of commodity of age $j \in J$ in distribution location $i \in D$ at the beginning of period $t \in T$ under scenario $s \in S$
X_j^t	Inventory of commodity of age j in the central warehouse at the beginning of period $t \in T$
e_{is}^t	Number of units expired in distribution location $i \in D$ at the end of period $t \in T$ under scenario $s \in S$
e_0^t	Number of units expired in the central warehouse at the end of period $t \in T$
b_{is}^t	Amount of unmet demand in distribution location $i \in D$ at the end of period $t \in T$ under scenario $s \in S$

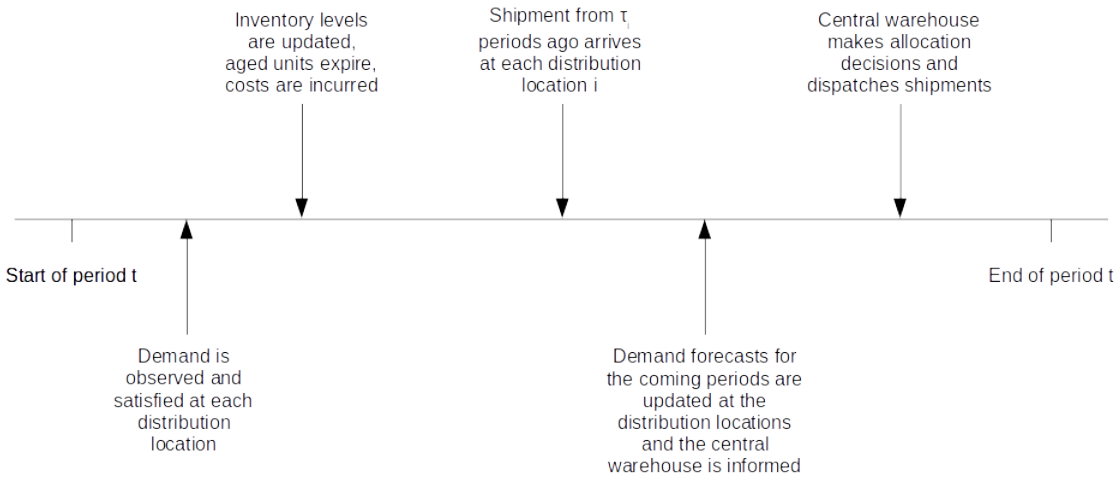

Figure 39: Sequence of events in a given period for the hub system

in each period is presented in Figure 39.

At the start of each period, the central warehouse receives shipments of SNFs from the suppliers and shipments dispatched from the central warehouse τ_i periods ago arrive at each distribution location i. The total amount of shipments from the central warehouse is limited by the number of vehicles available and their capacity. We assume that a single dispatch is sufficient for transporting the SNFs between each hub-DL pair.

Once the shipments arrive at the distribution locations, demand forecasts (scenarios and associated probabilities) are updated for the upcoming periods. The central warehouse then makes decisions on how much SNF to allocate to each distribution location and dispatches the shipments, which will arrive τ_i periods later in each distribution location i. Meanwhile, demand is observed at each distribution location, which results in updates of inventory levels. Costs of expiration and shortage (if any) are also incurred. It should be noted here that while the flow of events in Figure 39 presents a dynamic environment, our two-stage models to determine the allocations (and possibly hub location decisions) are static in the sense that they are solved assuming one of the current set of scenarios will hold for the reminder of the planning

horizon. Implementing a rolling horizon procedure by inputting a set of updated demand scenarios (reflecting forecast updates) allows incorporating the dynamic nature of the environment into our models.

Incorporating hub location decisions and forecast and demand information updates involve the trade-off between solution quality (in terms of the expected cost of the system) and computational time requirement. For the hub system, we consider three different models. Of these two aspects, (HL0) only incorporates information updates, (HL1-N) incorporates only location decisions, and (HL1-R) considers both.

In (HL0), hub locations are considered as inputs to the model. Therefore, the main decisions consist of allocation of the SNFs to the distribution locations. Let H denote the set of hubs, indexed by h, and $h(i)$ denote the hub of the region distribution location i is located in. The allocations in each period are determined by solving the following mixed integer program:

$$(\mathbf{HL0}) = \min \sum_{i\in D}\sum_{t\in T}\sum_{s\in S} q_s \left(pb_{is}^t + we_{is}^t\right) + \sum_{h\in H}\sum_{t\in T}\gamma_h r_h^t + \sum_{h\in H}\sum_{v\in D}\sum_{t\in T}\gamma_{hv} y_{hv}^t + \sum_{t\in T} we_0^t \tag{43}$$

$$\text{s.t.} \sum_{i\in D} v_{ij}^t \le X_j^t \qquad j\in J,\, t\in T \tag{44}$$

$$X_0^t = \xi_0^t \qquad t\in T \tag{45}$$

$$X_j^{t+1} = X_{j-1}^t + \xi_j^t - \sum_{i\in D} v_{i,j-1}^t \qquad j\in J\setminus\{0\},\, t\in T\setminus\{|T|\} \tag{46}$$

$$e_0^t \ge X_{|J|-1}^t - \sum_{i\in D} v_{i,|J|-1}^t \qquad t\in T \tag{47}$$

$$v_{ij}^t = 0 \qquad i\in D,\, j\in\{|J|-\tau_i,\ldots,|J|\},\, t\in T \tag{48}$$

$$v_{ij}^t = 0 \qquad i\in D,\, j\in J,\, t\in\{|T|-\tau_i+1,\ldots,|T|\} \tag{49}$$

$$\sum_{i\in D_r}\sum_{j\in J} v_{ij}^t \le \sum_{h\in H} c_1 r_h^t \qquad r\in R,\, t\in T \tag{50}$$

$$\sum_{i \in D} r_i^t \leq k \qquad t \in T \tag{51}$$

$$e_{is}^t = I_{i,|J|-1,s}^t + v_{i,|J|-\tau_i-1}^{t-\tau_i} - l_{i,|J|-1,s}^t \qquad i \in D,\ s \in S,\ t \in \{\tau_i + 1, \ldots, |T|\} \tag{52}$$

$$e_{is}^t = I_{i,|J|-1,s}^t - l_{i,|J|-1,s}^t \qquad i \in D,\ s \in S,\ t \in \{1, \ldots, \tau_i\} \tag{53}$$

$$b_{is}^t \geq d_{is}^t - \sum_{j \in J} l_{ijs}^t \qquad i \in D,\ s \in S,\ t \in T \tag{54}$$

$$I_{ijs}^{t+1} = I_{i,j-1,s}^t + v_{i,j-\tau_i-1}^{t-\tau_i} - l_{i,j-1,s}^t \qquad i \in D,\ j \in \{\tau_i + 1, \ldots, |J|\},\ s \in S,\ t \in \{\tau_i + 1, \ldots, |T|-1\} \tag{55}$$

$$I_{ijs}^{t+1} = I_{i,j-1,s}^t - l_{i,j-1,s}^t \qquad i \in D,\ j \in \{1, \ldots, \tau_i\},\ s \in S,\ t \in \{\tau_i + 1, \ldots, |T|\} \tag{56}$$

$$I_{ijs}^{t+1} = I_{i,j-1,s}^t - l_{i,j-1,s}^t \qquad i \in D,\ j \in J \setminus \{0\},\ s \in S,\ t \in \{1, \ldots, \tau_i\} \tag{57}$$

$$I_{i0s}^{t+1} = 0 \qquad i \in D,\ s \in S,\ t \setminus \{|T|\} \tag{58}$$

$$\sum_{j \in J} l_{ijs}^t \leq d_{is}^t \qquad i \in D,\ s \in S,\ t \in T \tag{59}$$

$$\sum_{j \in J} v_{ij}^t \leq c_2\, y_{h(i),i}^t \qquad i \in D,\ t \in T \tag{60}$$

$$v_{ij}^t \geq 0 \qquad i \in D,\ j \in J,\ t \in T \tag{61}$$

$$I_{ijs}^t \geq 0 \qquad i \in D,\ j \in J,\ s \in S,\ t \in T \tag{62}$$

$$e_{is}^t \geq 0 \qquad i \in D \cup \{0\},\ s \in S,\ t \in T \tag{63}$$

$$b_{is}^t \geq 0 \qquad i \in D,\ s \in S,\ t \in T \tag{64}$$

$$X_j^t \geq 0 \qquad j \in J,\ t \in T \tag{65}$$

$$l_{ijs}^t \geq 0 \qquad i \in D,\ j \in J,\ s \in S,\ t \in T \tag{66}$$

$$y_{hv}^t \in \{0, 1\} \qquad h \in H,\ v \in D,\ t \in T \tag{67}$$

$$r_i^t \in \mathbb{Z}_+ \qquad i \in D,\ t \in T. \tag{68}$$

The objective function minimizes the sum of expected shortage cost, expiration

cost, and transportation cost. Constraint set (44) limits the total allocation quantities for each age to those available in the central warehouse. Constraint (45) updates the fresh commodities at the central warehouse, whereas Constraint (46) updates the age of the inventory at the central warehouse. Constraint set (47) calculates the amount of expiring commodities at the central warehouse at the end of a period after allocation is made. Constraints (48) and (49) prevent allocation of commodities that will expire before reaching the distribution location and those that will reach the distribution location after the planning horizon is over, respectively. Constraints (50) and (51) make the vehicle assignments to hubs and limits shipments by vehicle capacity, respectively.

Constraints (52) and (53) determine the amount of expired commodities in each distribution location, while Constraint (54) determines the amount of shortage. Using Constraints (55)-(58), the inventory levels for each age and demand level are updated for each distribution location. Constraint (59) bounds the amount of satisfied demand by the total demand in each distribution location, scenario, and period. Constraint (60) requires that the shipment indicator variable between a hub and a distribution location to be equal to 1 if there is a shipment from the central warehouse to the latter. Constraints (61)-(68) define the set constraints for the non-negativity, integrability, and binary restrictions for the decision variables.

The algorithm to implement the (HL0) model on a rolling horizon basis works in the following way: At the beginning of the first period, the mixed integer program (HL0) is solved for $|T|= |J|$. The allocations of the first period are made based on the decision variables corresponding to the first period. Then, the inventory levels are updated based on the actual demand of the period, potential demand levels for each scenario are updated based on forecast updates, and costs due to expiration and shortage (if any) are incurred. The model is then solved for the next $|J|$ periods starting from the second period. This results in the allocation decisions of the second

period, and the process is repeated until the end of the planning horizon.

In (HL1-N) and (HL1-R), both location and allocation decisions are included. Without loss of generality, we assume that hubs can only be located at the existing distribution locations. At the beginning of the first period, we solve the following mixed integer program:

$$(\mathbf{HL1}) = \min \sum_{i\in D}\sum_{t\in T}\sum_{s\in S} q_s \left(pb_{is}^t + we_{is}^t\right) + \sum_{i\in D}\sum_{t\in T}\gamma_i r_i^t + \sum_{u\in D}\sum_{v\in D}\sum_{t\in T}\gamma_{uv} y_{uv}^t + \sum_{t\in T} we_0^t \tag{69}$$

s.t. (44)-(49), (51)-(59), (61)-(66), (68)

$$\sum_{i\in D_r}\sum_{j\in J} v_{ij}^t \le \sum_{i\in D_r} c_1 r_i^t \qquad r\in R,\, t\in T \tag{70}$$

$$r_i^t \le z_i \qquad i\in D,\, t\in T \tag{71}$$

$$\sum_{i\in D_r} z_i = 1 \qquad r\in R \tag{72}$$

$$\sum_{j\in J} v_{ij}^t \le \sum_{u:R(u)=R(i)} c_2\, y_{ui}^t \qquad i\in D,\, t\in T \tag{73}$$

$$y_{li}^t \le z_l \qquad l\in D, i\in D,\, t\in T \tag{74}$$

$$y_{uv}^t \in \{0,1\} \qquad u\in D,\, v\in D,\, t\in T \tag{75}$$

$$z_i \in \{0,1\} \qquad i\in D \tag{76}$$

The main difference between (HL1) and (HL0) arises from the inclusion of the location variables. Constraint (71) ensures that there can be a vehicle dispatched for a hub only if the hub is open. A single hub is enforced to be opened in each region by Constraint (72). Constraint (74) imposes the requirement that a hub must be open in order to have incoming and outgoing shipments.

We use the following algorithm to implement the (HL1-N) model: The mixed integer program (HL1) is solved for the whole planning horizon, and the resulting location and commodity allocation decisions are taken throughout the whole planning

horizon. In this sense, (HL1-N) can be considered as a benchmark to analyze the effects of ignoring information updates. The algorithm to implement (HL1-R) on a rolling horizon basis starts by solving (HL1) for $|J|$ periods at the beginning of the first period. The hub location and first period commodity allocation decisions are made based on the results of this model. Starting from the second period and throughout the remainder of the planning horizon, the model (HL0) is solved on a rolling horizon basis.

4.4.2 Models for the District Warehouse System

In this case, district warehouses act as intermediate storage locations and hold stocks of the commodity. Another important distinction in this case is that the central warehouse observes the aggregate demand at the district warehouses, rather than the demand at the individual distribution locations. Therefore, there are two decision levels in this case: At the higher level, the central warehouse decides on the allocation quantities to the district warehouses by minimizing the expected total shortage and expiry costs at the district level and transportation cost to the districts. The districts, in turn, make their allocation decisions based on these quantities.

Figure 40 summarizes the sequence of events in a given period. At the beginning of the period, each district warehouse i receives shipments dispatched from the central warehouse τ_i periods ago. The distribution locations update demand forecasts for the upcoming periods. Upon receiving these updates, the district warehouses update the aggregate forecasts and inform the central warehouse, which, in turn, makes the allocation decisions and dispatches the shipments to the district warehouses. By taking into account the incoming shipments and on-hand inventory, district warehouses make the allocation decisions. There are multiple shipments from a DW to the DLs between consecutive shipments from the CW to the DW, since the downstream shipments are made more frequently. Shipments from the district warehouses reach distribution

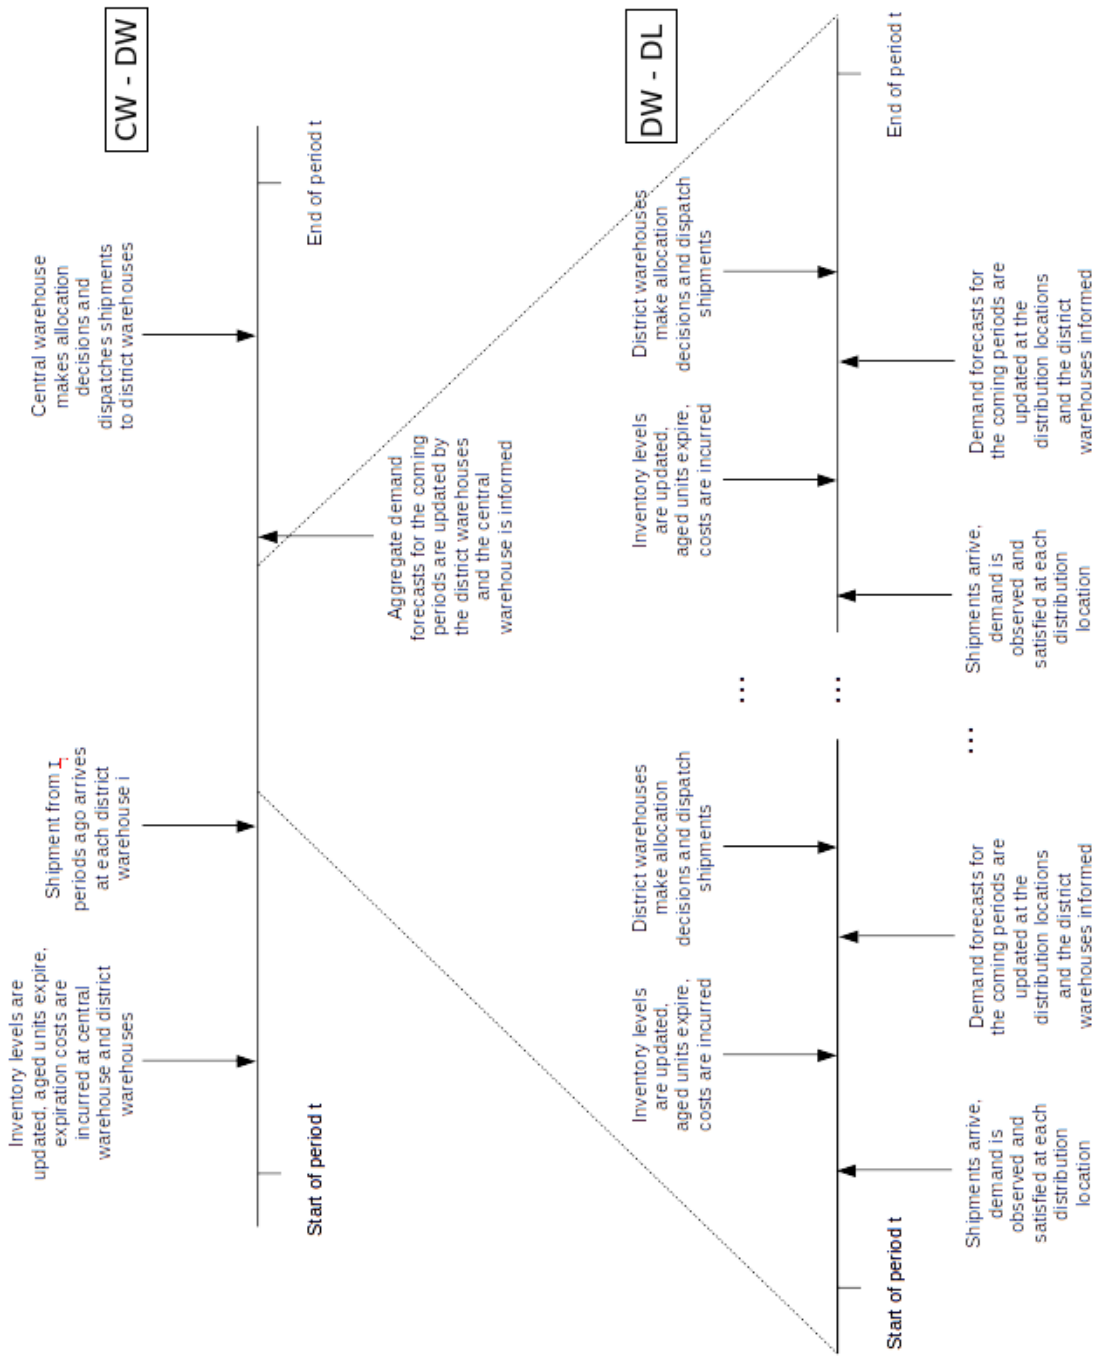

Figure 40: Sequence of events in a given period for the district warehouse system, where the top and bottom parts describe flow of events the CW-DW and DW-DL portions of the network, respectively

locations before demand is observed. At the end of the period, inventory levels are updated, and costs of expiration and shortage (if any) are incurred.

We slightly change the notation from the previous section. The set H denotes the set of district warehouses. Similarly, D_h denotes the set of distribution locations assigned to district warehouse h, and the inventory level of age j at the beginning of the period at district warehouse h is denoted by I^t_{hjs}. In this case, the lead time is broken into two components τ_h and τ_i, which denote the lead time from the central warehouse to district warehouse h, and the lead time from district warehouse to distribution location i. Since district warehouse locations are fixed, the shipment variables are represented by y^t_i, which takes a value of 1 if there is incoming shipment for distribution location i. As before, we assume individual shipments to the district warehouses and distribution locations due to geographical diversity. Let $d^t_{hs} = \sum_{i \in D_h} d^t_{is}$ denote the aggregate demand level of the distribution locations assigned to district warehouse h and v^t_{hj} denote the quantity assigned to district warehouse h from the central warehouse. Then, we have the following model for the central warehouse:

$$(\mathbf{DW0}) = \min \sum_{h \in H} \sum_{t \in T} \sum_{s \in S} q_s \left(pb^t_{hs} + we^t_{hs} \right) + \sum_{h \in H} \sum_{t \in T} \gamma_h r^t_h + \sum_{t \in T} we^t_0 \tag{77}$$

$$\text{s.t.} \sum_{h \in H} v^t_{hj} \leq X^t_j \qquad j \in J,\, t \in T \tag{78}$$

$$X^t_0 = \xi^t \qquad t \in T \tag{79}$$

$$X^{t+1}_j = X^t_{j-1} - \sum_{h \in H} v^t_{h,j-1} \qquad j \in J \setminus \{0\},\, t \in T \setminus \{|T|\} \tag{80}$$

$$e^t_0 \geq X^t_{|J|-1} - \sum_{h \in H} v^t_{h,|J|-1} \qquad t \in T \tag{81}$$

$$v^t_{hj} = 0 \qquad h \in H,\, j \in \{|J|-\tau_i, \ldots, |J|\},\, t \in T \tag{82}$$

$$v^t_{hj} = 0 \qquad h \in H,\, j \in J,\, t \in \{|T|-\tau_i+1, \ldots, |T|\} \tag{83}$$

$$\sum_{j \in J} v^t_{hj} \leq c_1 r^t_h \qquad h \in H,\, t \in T \tag{84}$$

$$\sum_{h\in H} r_h^t \le k \qquad t \in T \tag{85}$$

$$e_{hs}^t = I_{h,|J|-1,s}^t + v_{h,|J|-\tau_h-1}^{t-\tau_h} - l_{h,|J|-1,s}^t \qquad h \in H,\ s \in S,\ t \in \{\tau_h+1,\dots,|T|\} \tag{86}$$

$$e_{hs}^t = I_{h,|J|-1,s}^t - l_{h,|J|-1,s}^t \qquad h \in H,\ s \in S,\ t \in \{1,\dots,\tau_h\} \tag{87}$$

$$b_{hs}^t \ge d_{hs}^t - \sum_{j\in J} l_{hjs}^t \qquad h \in H,\ s \in S,\ t \in T \tag{88}$$

$$I_{hjs}^{t+1} = I_{h,j-1,s}^t + v_{h,j-\tau_h-1}^{t-\tau_h} - l_{h,j-1,s}^t \qquad h \in H,\ j \in \{\tau_h+1,\dots,|J|\},\ s \in S,\ t \in \{\tau_h+1,\dots,|T|-1\} \tag{89}$$

$$I_{hjs}^{t+1} = I_{h,j-1,s}^t - l_{h,j-1,s}^t \qquad h \in H,\ j \in \{1,\dots,\tau_h\},\ s \in S,\ t \in \{\tau_h+1,\dots,|T|\} \tag{90}$$

$$I_{hjs}^{t+1} = I_{h,j-1,s}^t - l_{h,j-1,s}^t \qquad h \in H,\ j \in J\setminus\{0\},\ s \in S,\ t \in \{1,\dots,\tau_h\} \tag{91}$$

$$I_{h0s}^{t+1} = 0 \qquad h \in H,\ s \in S,\ t \setminus \{|T|\} \tag{92}$$

$$\sum_{j\in J} l_{hjs}^t \le d_{hs}^t \qquad h \in H,\ s \in S,\ t \in T \tag{93}$$

$$v_{hj}^t \ge 0 \qquad h \in H,\ j \in J,\ t \in T \tag{94}$$

$$I_{hjs}^t \ge 0 \qquad h \in H,\ j \in J,\ s \in S,\ t \in T \tag{95}$$

$$e_{hs}^t \ge 0 \qquad h \in H \cup \{0\},\ s \in S,\ t \in T \tag{96}$$

$$b_{hs}^t \ge 0 \qquad h \in H,\ s \in S,\ t \in T \tag{97}$$

$$X_j^t \ge 0 \qquad j \in J,\ t \in T \tag{98}$$

$$l_{hjs}^t \ge 0 \qquad h \in H,\ j \in J,\ s \in S,\ t \in T \tag{99}$$

$$r_h^t \in \mathbb{Z}_+ \qquad h \in H,\ t \in T. \tag{100}$$

The constraints are similar to those under (HL0), except the transportation costs from the district warehouses to distribution locations not being included, as well as the routing constraints at the district level being disregarded. Once the v_{hj}^t values (represented in the following model as $\bar{v}_{hj}^t$ to avoid confusion) are determined by the

CW, these are used at each district as inputs to determine the allocations to the distribution locations in the following mixed integer program:

$$(\mathbf{DW1}) = \min \sum_{i\in D}\sum_{t\in T}\sum_{s\in S} q_{is}^t \left(pb_{is}^t + we_{is}^t\right) + \sum_{u\in D}\sum_{v\in D}\sum_{t\in T}\gamma_{uv}y_{uv}^t + \sum_{h\in H}\sum_{t\in T} we_h^t \tag{101}$$

s.t. (48), (49), (52)-(59), (61)-(64), (66), (67)

$$I_{hj}^{t+1} = I_{h,j-1,s}^t + \bar{v}_{h,j-\tau_h-1}^{t-\tau_h} - \sum_{i\in D_h} v_{i,j-1}^t \qquad h\in H,\ j\in\{\tau_h+1,\ldots,|J|\},\ t\in\{\tau_h+1,\ldots,|T|-1\} \tag{102}$$

$$I_{hj}^{t+1} = I_{h,j-1}^t - \sum_{i\in D_h} v_{i,j-1}^t \qquad h\in H,\ j\in\{1,\ldots,\tau_h\},\ t\in\{\tau_h+1,\ldots,|T|\} \tag{103}$$

$$I_{hj}^{t+1} = I_{h,j-1}^t - \sum_{i\in D_h} v_{i,j-1}^t \qquad h\in H,\ j\in J\setminus\{0\},\ t\in\{1,\ldots,\tau_h\} \tag{104}$$

$$I_{h0}^{t+1} = 0 \qquad h\in D,\ t\setminus\{|T|\} \tag{105}$$

$$e_h^t \geq I_{h,|J|-1}^t - \sum_{i\in D_h} v_{h,|J|-1}^t \qquad h\in H,\ t\in T \tag{106}$$

$$\sum_{j\in J} v_{ij}^t \leq c_2\, y_i^t \qquad i\in D,\ t\in T \tag{107}$$

$$I_{hj}^t \geq 0 \qquad h\in H,\ j\in J,\ s\in S,\ t\in T \tag{108}$$

$$e_h^t \geq 0 \qquad h\in H,\ s\in S,\ t\in T \tag{109}$$

$$y_i^t \in \{0,1\} \qquad i\in D,\ t\in T \tag{110}$$

Again, the objective function and constraints are similar to those under (HL0). The objective function for (DW1) does not include transportation costs from the central warehouse to the districts, and the initial inventory at the district level also includes allocations from the central warehouse in the same period.

As in the algorithms that implement (HL0) and (HL1-R), the algorithms for both (DW0) and (DW1) are implemented on a rolling horizon basis. The resulting decisions corresponding to the first period are taken, updates are made, and the model is solved starting from the next period, and so on.

4.4.3 Models for Equity of Distribution

It is highly likely that under models focusing on cost minimization, convenient distribution locations that have shorter lead times and/or lower transportation costs from hubs/district warehouses might be favored for distribution. A desirable aspect in the distribution of commodities in the health/humanitarian context is that all beneficiaries have equal access to the commodity [29]. This requirement leads to the need to minimize the deviation of satisfied demand over all distribution locations.

Minimizing the deviation of satisfied demand leads to many alternative objective function definitions, including minimizing variance, mean absolute deviation, the Gini coefficient, range, coefficient of variation, or maximizing the minimum fraction of demand satisfied [111]. In this section, we will introduce models and solution approaches that aim to achieve the latter objective.

We apply an identical framework for all approaches discussed in the preceding section by (i) limiting the total cost for each iteration of the approaches to a certain total cost limit, (ii) taking the minimum fraction of demand satisfied over all DLs in each period, and (iii) maximizing the sum of these fractions over all periods. For (HL0), this translates into the following mixed integer program:

$$\textbf{(HL0-E)} = \max \sum_{t \in T} \psi^t \tag{111}$$

s.t. (44)-(68)

$$\psi^t \leq \sum_{s \in S} q_s \left(\frac{\sum_{j \in J} l^t_{ijs}}{d^t_{is}} \right) \qquad i \in D,\, t \in T \tag{112}$$

$$\sum_{i \in D} \sum_{t \in T} \sum_{s \in S} q^t_{is} \left(pb^t_{is} + we^t_{is} \right) + \sum_{h \in H} \sum_{t \in T} \gamma_h r^t_h$$

$$+\sum_{h\in H}\sum_{v\in D}\sum_{t\in T}\gamma_{hv}y^t_{hv}+\sum_{t\in T}we^t_0\le(1+\eta)B \tag{113}$$

$$\psi^t\ge 0 \qquad\qquad t\in T \tag{114}$$

Here, the decision variable ψ^t, defined in constraint (112), tracks the minimum expected fraction of demand satisfied over all locations for each period. The objective function (111) maximizes the sum of these variables over all periods. Constraint (113) limits the total expected shortage, expiration, and transportation costs. Here, B is the optimal cost obtained by (HL0) for the given horizon, and η is a nonnegative factor to analyze the effect of increasing cost limit on equity of allocations.

For (HL1), we obtain the following formulation for equity:

$$\textbf{(HL1-E)} = \max\sum_{t\in T}\psi^t \tag{115}$$

$$\text{s.t. (44)-(49), (51)-(59), (61)-(66), (68), (70)-(76), (112), 114)}$$

$$\sum_{i\in D}\sum_{t\in T}\sum_{s\in S}q_s\left(pb^t_{is}+we^t_{is}\right)+\sum_{i\in D}\sum_{t\in T}\gamma_i r^t_i+\sum_{u\in D}\sum_{v\in D}\sum_{t\in T}\gamma_{uv}y^t_{uv}$$
$$+\sum_{t\in T}we^t_0\le(1+\eta)B \tag{116}$$

Similar to (HL0-E), Constraint (116) limits the total expected cost. Formulations (DW0-E) and (DW1-E) provide the equity-based formulations for the central warehouse and district warehouse levels of the district warehouse system, respectively:

$$\textbf{(DW0-E)} = \max\sum_{t\in T}\psi^t \tag{117}$$

$$\text{s.t. (78)-(100), (114)} \tag{118}$$

$$\psi^t\le\sum_{s\in S}q_s\left(\frac{\sum_{j\in J}l^t_{ijs}}{d^t_{hs}}\right)\quad h\in H,\, t\in T \tag{119}$$

$$\sum_{h\in H}\sum_{t\in T}\sum_{s\in S}q_s\left(pb^t_{hs}+we^t_{hs}\right)+\sum_{h\in H}\sum_{t\in T}\gamma_h r^t_h+\sum_{t\in T}we^t_0\le(1+\eta)B \tag{120}$$

In this case, Constraints (119) and (120) take the expectation over aggregate demand at the district warehouses, instead of the distribution locations.

$$\textbf{(DW1-E)} = \max \sum_{t \in T} \psi^t \tag{121}$$

$$\text{s.t. (48), (49), (52)-(59), (61)-(64), (66), (67), (102)-(110), (112), 114)}$$

$$\sum_{i \in D} \sum_{t \in T} \sum_{s \in S} q_s \left(pb_{is}^t + we_{is}^t \right) + \sum_{i \in D} \sum_{t \in T} \gamma_i r_i^t + \sum_{u \in D} \sum_{v \in D} \sum_{t \in T} \gamma_{uv} y_{uv}^t$$

$$+ \sum_{t \in T} we_0^t \leq (1 + \eta) B \tag{122}$$

In general, the equity models take significantly longer time to solve compared to the cost-minimizing counterparts. This is mainly due to two reasons. First, finding a good initial feasible solution that adheres to the budget constraint takes significant time. Second, many alternative solutions may lead to the same minimum demand satisfaction level. Throughout our computational experiments, the approaches that use (HL1-E) for the first period encounter insufficient memory, particularly caused by the existence of the binary location variables. To solve this model, we resort to a variable fixing heuristic. As a first step, we solve the (HL1-E) by relaxing the location variables to be between 0 and 1. In our preliminary experiments, we have observed that a large majority of these variables take values between 0-0.2 or 0.8-1. Our heuristic fixes the variables in the former and latter groups to 0 and 1, respectively. Following this, we solve the integer program and obtain the heuristic solution. In computational experiments, we assess the quality of this heuristic using the resulting optimality gaps.

4.5 Computational Experiments

The main objective of the computational experiments we present in this section is to compare performances of the two network structures (the existing network with district warehouses and the alternative network with intermediate hubs) considering various factors.

It has been shown in §4.2 that one of the most prominent attributes of the SNF is perishability, and using historical data from a landlocked country, it was concluded that ignoring perishability might increase the level of shortages and expirations throughout the supply chain. Since incorporating perishability requires that commodity age be tracked, it increases the level of complexity of the models. In our computational experiments, we aim to analyze the cases where incorporating perishability into the models provides more significant savings in expiration and shortage costs. The analysis in §4.2 also revealed that significant lead time may be required before commodities reach the central warehouse in a recipient country, especially if the country is landlocked. One of our goals in this section is to observe the effect of decreasing these lead times, which in turn decreases the risk of expirations.

Another important result from §4.2 is that usage forecasts may not always be accurate, as shown by the usage data from a subset of the distribution locations in a landlocked country. As a result, ignoring uncertainty when making the allocation decisions might be detrimental for the overall performance of the system. The uncertainty in demand is incorporated into the models in §4.4 and is another aspect that increases model complexity. Through computational experiments, we aim to show in what cases solving these more complex models improves the system performance more significantly. Since the quality of usage forecasts plays an important role as well, we also assess the effects of the variability of demand distribution on the allocations.

Regarding how actual usage is incorporated, the models presented in §4.4 can be solved in two different ways. In the first one, the model is solved once without considering the actual usage occurrence. While this has the advantage of significantly reduced running times, solution quality may be significantly affected, particularly when actual occurrences are unlikely and/or not close to the mean of the distribution. As an alternative, these models are solved on a rolling horizon basis, requiring multiple runs of the same model, but with increased solution quality. When designing the hub

Table 14: Parameter settings for the computational experiments

Parameters	**Values or distributions**
Number of regions	12
Number of districts	39
Number of distribution locations	264
Shelf life	12 months
Planning horizon	36 months
Number of possible demand levels	10
Factors for determining possible demand levels	(0.75, 1.25), (0.90, 1.10), (0.50, 1.50)
Unit costs (Shortage, expiration, transportation)	(100, 50, 0.3), (100, 100, 0.3), (100, 50, 0.6)
Number of vehicles	50
Additional cost limit factor for equity models	0, 0.25, 0.50, 0.75

network, locations of these hubs can be taken as given or as part of decisions. A similar trade-off between model complexity and solution quality also exists in this situation. Analyzing these trade-offs is another goal of our computational experiments.

Even though one of the main drivers of the optimal allocations in the SNF supply chain are the costs of shortage, expiration and transportation, it is generally not straightforward to quantify these in societal applications. To mitigate against the bias created by our choice of these parameters, we analyze the sensitivity of our conclusions by varying the relative values of these costs. We also analyze the effect of prioritizing equity of the allocations rather than minimizing system-wide costs. The parameter settings we consider throughout this section are summarized in Table 14.

The remainder of this section is organized as follows: First, we describe the instances we have generated and used based on data from the organization. Following this, we introduce a set of benchmark approaches that ignore the perishability of commodities and the uncertainty in the system. Lastly, we present the results of our experiments and make conclusions on selecting the network structure of the SNF supply chain.

4.5.1 Instance Generation

Our computational experiments are performed on instances based on the operations of the organization in a landlocked country. While we use part of the network data,

confidential data such as supply arrivals, usage forecasts, and lead times are randomized.

We focus on the delivery of corn-soya blend (with a commodity shelf life of 12 months) for a period of two years from the central warehouse to 264 distribution locations. 39 district warehouses and 12 potential hubs. The distances between these facilities are determined using the actual road network data (publicly available in [58]). Lead times are randomly generated based on distances between the central warehouse and district warehouses/hubs, and the hubs and the distribution locations (lead time from DWs to DLs is assumed negligible). For the sake of simplicity, changes in lead time throughout the year due to factors such as rainy seasons are ignored. Supply arrivals and demand forecasts are generated based on the actual values, and are multiplied by a positive factor uniformly generated between 0.75 and 1.25. Age distribution of arriving shipments is based on Figure 32. The total initial inventory at each distribution location i is calculated by multiplying the $\tau_i + 1$ with the average demand at this DL. The initial inventory for each age from 0 to $|J|-1$ is assumed as equal. Similar initial inventory calculations are made for the district warehouses and (potential) hubs by aggregating the average demand at the downstream DLs and using the lead time from the CW to these facilities.

10 potential usage levels are assumed for each facility in each period. For the baseline case, these are generated randomly between 0.75 and 1.25 times the forecast amount, and are assumed equally likely. Actual usage amounts are randomly chosen among these levels. To analyze the effect of forecast variance, these factors are varied as 0.9 and 1.1 for the low variability, and 0.5 and 1.5 for the high variability cases, respectively.

For both network structures, 50 vehicles can be dispatched from the central warehouse to the DW/hubs in each period. The vehicle capacities are determined by the

equation:

$$\left\lfloor 1.25 \, \frac{\sum_{i \in D} \sum_{t \in T} d_i^t}{|J||K|} \right\rfloor,$$

as also used by Coelho and Laporte [42]. Here, d_i^t is replaced by the usage forecast in distribution location i for each period t, $|J|$ is 12 months, and $|K|= 50$. For the district warehouse system, shipments from the CW to the DWs and from the DWs to the DLs occur once a month and four times a month, respectively. For the hub system, shipments from the CW to the hubs and from the hubs to the DLs are twice a month.

In addition to the confidentiality of the cost data, unit costs of expiration and shortage are harder to quantify. For the baseline case, we assume \$100 per unit of shortage and \$50 per unit of expiration. fixed cost of using a corridor is calculated based on a unit cost of \$0.30 per kilometer. For the equity-based models, we vary the total cost limit as 100%, 125%, 150%, and 175% of the cost resulting from the objective value of the corresponding cost minimization model by setting the factor η to 0, 0.25, 0.50, and 0.75, respectively.

4.5.2 Benchmark Approaches

To assess the performance of our proposed approaches, we compare the results to two benchmark approaches, the first of which ignores the forecast errors and uncertainty of usage, whereas the second one ignores the perishability of the commodity.

With the first approach, all demand is assumed to be deterministic and equal to the expected usage. In mathematical terms, the scenario-based demand terms d_{is}^t in the models described in §4.4 are replaced by $d_i^t = \sum_{s \in S} q_{is}^t d_{is}^t$. Since demand is deterministic, the index s in the decision variables l_{ijs}^t, I_{ijs}^t, e_{is}^t, and b_{is}^t can be dropped, which reduces the complexity of solving the models significantly. The terms

$$\sum_{i \in D} \sum_{t \in T} \sum_{s \in S} q_{is}^t (p \, b_{is}^t + w \, e_{is}^t)$$

in the objective functions of the models are replaced by

$$\sum_{i \in D} \sum_{t \in T} (p\, b_i^t + w\, e_i^t).$$

The main aim in using this benchmark approach is to determine the improvement in system performance when stochasticity is incorporated into the models.

In the second approach, commodity shelf life is assumed to be infinite. In this case, the decision variable e_{is}^t and the constraints defining this variable are removed from the models. Furthermore, the index set J that keeps track of the commodity age is also dropped, which results in the removal of constraints that update the commodity age. As in the case of assuming deterministic usage amounts, the aim in this case is to assess the value of tracking commodity age and analyze the cases when age tracking results in significant performance improvement. In our computational experiments, we show the results of the allocation decisions resulting from these approaches on the actual demand data. Therefore, even though the latter model ignores perishability, the resulting costs include those due to commodity expiration.

4.5.3 Computational Results

Our computational experiments have been performed on a computer with Intel Core i7-4500U CPU with quad-core 1.80 GHz processor and 8 GB RAM using the commercial solver IBM ILOG CPLEX 12.6 with Concert technology. The CPU time limit for (HL0), (HL1-R), (HL1-N) and (DW) models has been set at six hours, whereas a limit of 12 hours has been set for (HL1-E) and (DW-E). For the latter set of models, the resulting solutions from (HL1-R) and (DW) instances have been fed in as initial feasible solutions, and the variable fixing heuristic described in §4.4 has been implemented. For the cost minimizing models, the largest observed optimality gap is below 5%, whereas the largest optimality gap for the equity models is below 8%.

In the first set of our experiments, we aim to compare the two network structures among themselves under different models in the baseline case. Furthermore, we also

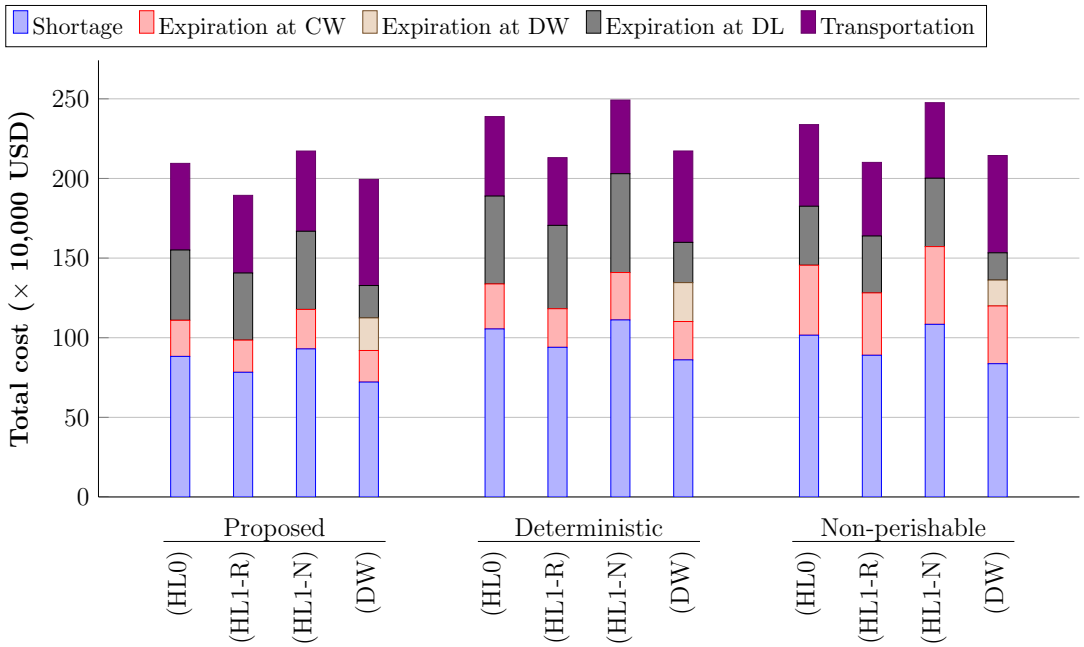

Figure 41: Breakdown of costs for the baseline settings resulting from the proposed models in §4.4 (i) incorporating usage uncertainty and commodity perishability, (ii) assuming deterministic usage with accurate forecasts, and (iii) ignoring perishability of the commodities

analyze the effect of considering forecast errors and perishability of the commodities by comparing the results from these models to those under models that ignore these two aspects. Figure 41 breaks down the costs under the baseline settings using the four models with the proposed approaches in §4.4, the deterministic approach with expected usage levels, and the non-perishable approach that assumes infinite shelf life and ignores expirations.

When the two network structures are compared among each other, we observe that with the network, cost, supply and usage parameters in the baseline settings **the hub system outperforms the district warehouse system only if hub location decisions are taken into account and a rolling horizon approach is used**. Increased accuracy of forecasts and the possibility of more frequent shipments allows for better satisfaction of demand and results in lower expiration and shortage costs. However, the increase in the transportation costs (particularly from the DWs to DLs)

offsets these savings and results in a 5.4% increase in the total cost compared to the (HL1-R) model. When hub locations are not included in the decisions, the main cost increase is due to transportation (due to the hubs being in less convenient locations) and shortage (since some far-off shipments resulting from the locations of the hubs are not performed). When actual usage information is not incorporated, shortage costs increase by more than 18%, leading to an overall cost increase of 14.7%.

When forecast errors are ignored, the performances of hub and district warehouse systems are comparable. The main reason behind this is the fact that when forecast errors are not considered, the frequency of shipments from the DW level to the DL level decreases, resulting in savings in transportation costs. Since the DW system makes use of aggregate forecasts from the district warehouses to make the allocation decisions from the central warehouse, ignoring forecast errors affects the DW system less in terms of shortages, compared to the hub system. Consequently, the overall performances of the (HL1-R) and (DW) become closer, with the gap decreasing to less than 2%. When hub location decisions or actual usage information updates are not incorporated into the models, the gap between these models and (HL1-R) is higher as the level of shortages increases more, due to far-off shipments not being carried out for the (HL0) model and inaccurate usage information for (HL1-N).

The overall performances of hub and district warehouse systems are also similar when commodity perishability is ignored. For all four models, ignoring perishability leads to a lower amount of commodities allocated to the distribution locations, which leads to an increased number of expirations at the central warehouse level, decreased expirations downstream, and lower transportation costs. As in the case of ignoring uncertainty, the accuracy of forecasts with the DW system results in a lower increase in the expiration costs at the CW level. As a result, the performances of (HL1-R) and (DW) are within 2% of each other in terms of overall costs.

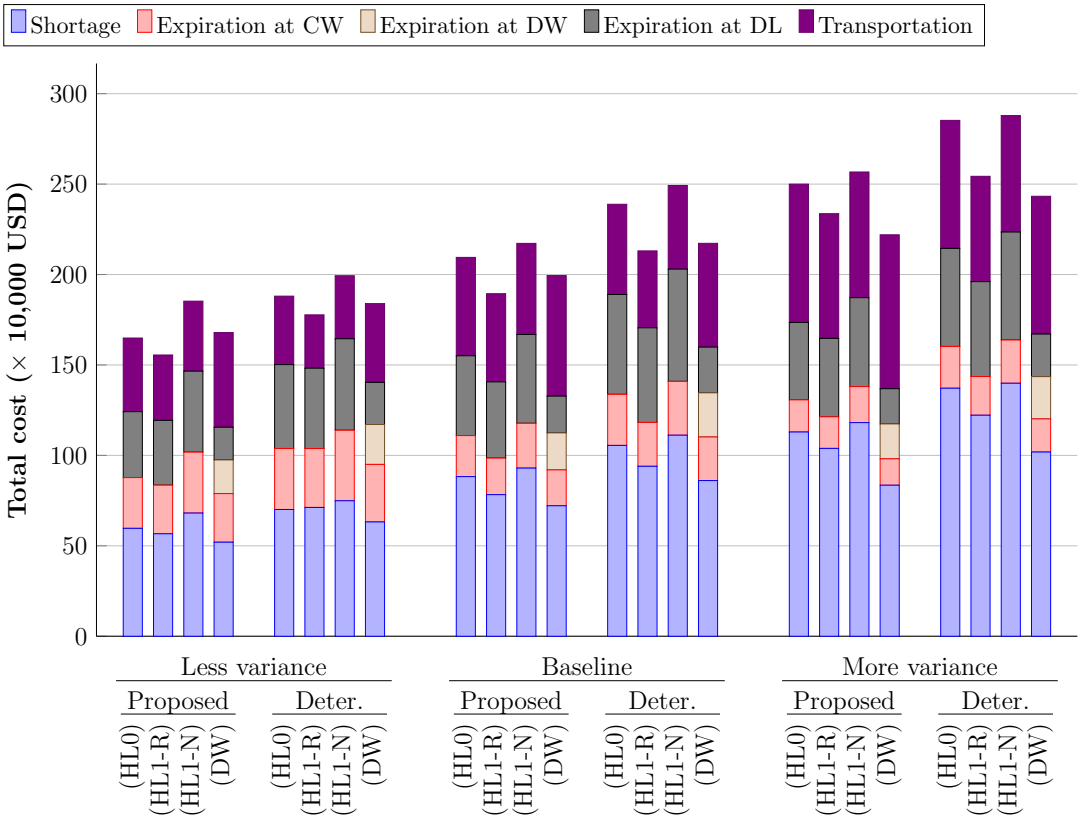

Figure 42: Breakdown of costs with changing forecast variability resulting from the proposed models in §4.4 (i) incorporating usage uncertainty and commodity perishability, and (ii) assuming deterministic usage with accurate forecasts

Forecast accuracy is an important factor which, combined with the network structure, plays an important role in determining the allocations. To assess the effect of the quality of forecasts, we vary the usage forecast range at each distribution location for each period. In the baseline case, the potential usage levels are randomly chosen from the values between 0.75 and 1.25 times the actual value. For the lower and higher variability cases, these coefficients are 0.9 and 1.1, and 0.5 and 1.5, respectively. As before, the number of possible levels is kept at 10 and are assumed equally likely. figure 42 shows the cost breakdown of the four models corresponding to the two network structures under the baseline, lower variance, and higher variance cases under the proposed approaches and those that ignore uncertainty in usage.

The first observation from Figure 42 is that under the proposed approaches, **when forecast accuracy is higher, the hub structure outperforms the DW system, even when hub locations are not part of the decisions**. The main reason leading to this result is the fact that with increased forecast accuracy, the need to buffer against usage uncertainty is decreased, and the shortage and expiration costs are significantly lower at the downstream level, which benefits the hub system more significantly as this system has higher shortage and expiration costs in the baseline case. Furthermore, with higher forecast accuracy, relative benefits of aggregate forecasts at the DW level are no longer significant compared to the hub system. Altogether, these lead to the hub system being more favorable when forecast variability is lower. As expected, the gap between the deterministic approach ignoring forecast variability and proposed approaches is smaller, since the expected usage better reflects the actual usage in this case.

Exactly opposite results are observed **when forecast variability is higher**, in which case **the district warehouse system outperforms the hub system, even when hub locations are part of decisions**. With the proposed models, the district warehouse system performs 5.3% and 12.6% better than the hub system with

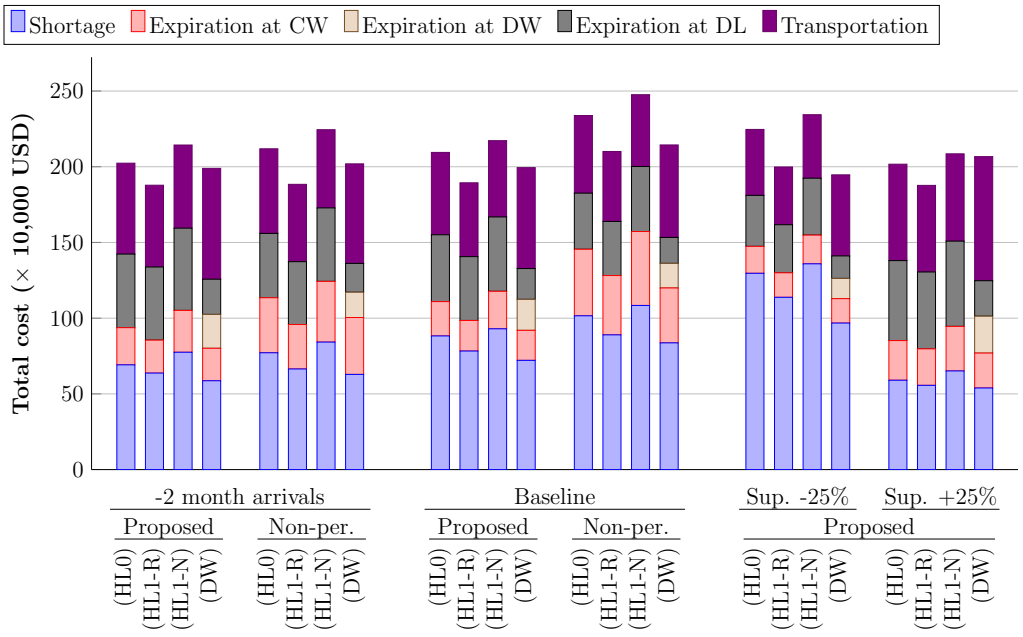

Figure 43: Breakdown of costs (i) when shipments arrive two months earlier and (ii) the baseline settings under the proposed models and the benchmark approach that ignore perishability, and (iii) when supply amounts are decreased and increased by 25% under the proposed models

and without location decisions, respectively. Similarly, the main reason behind this difference is that costs of shortage significantly increase from the baseline case, and more so for the hub system. The accuracy of aggregate forecasts for the DW system buffers against the risk of expirations and shortages better and the savings in the increase of these costs offset the increase in the transportation cost. Similar results hold under the deterministic model, except that the performance difference is more significant in that case.

An important issue that complicates the allocation and transportation decisions is the fact that significant lead time might be required until commodities arrive at the central warehouse. In our next set of experiments, we quantify the effects of decreased lead times on system performance. Lead time reduction can be achieved in a number of ways, including adoption of local production of the foods and improvement of the customs procedures at the ports and at the borders, as well as better tracking of

commodities and employing a first expiring, first out system. In the first part of Figure 43, the breakdown of the costs are shown when supply arrivals are two months earlier, under the proposed models and those that ignore the perishability of the commodities.

Under the proposed model, the results are in parallel to that under the baseline case, that is, **when commodities arrive earlier, a hub system performs better only if information updates and hub location decisions are incorporated**. In this case, increased availability due to arrival of fresher commodities of commodities decreases the costs of shortage, which favors the hub system with higher shortage costs. On the other hand, Figure 34 shows that cumulative demand and supply are not significantly different from each other in general. Therefore, increased availability of commodities leads to an increase in commodity expirations as well, which affects the hub system more since the expiration costs are higher for this system due to inaccuracy of forecasts. Since the increase in transportation costs is not significant, these two factors offset each other and the relative performances of the two systems are roughly the same. Because of the buffer created by increased effective commodity shelf life upon arrival at the CW, importance of considering perishability decreases and the models ignoring perishability provide compatible results, as shown in Figure 42.

Supply availability depends on the availability of funding, which may vary substantially from time to time. In the last part of Figure 42, we show the breakdown of costs under the proposed models when the supply amounts are decreased and increased by 25%. Considering the former case in Figure 42, we observe that **when supply is scarce, district warehouses play an important role in buffering against demand uncertainty by holding stocks of the commodity at the intermediate level**. The hub system, on the other hand, has to resolve the trade-off between the risk of overcommitment by immediately sending shipments to DLs and

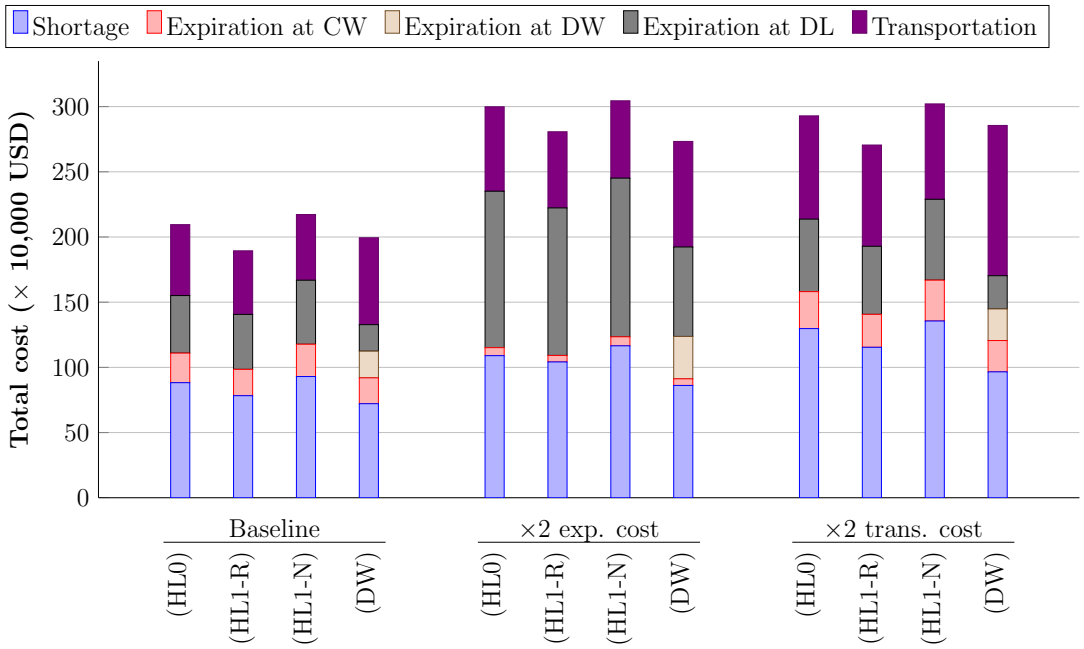

Figure 44: Breakdown of costs (i) for the baseline case, (ii) when unit expiration costs are increased by 100%, and (iii) when unit transportation costs are increased by 100%

the risk of shortages when more stock is kept at the CW level due to lead times. Consequently, increased shortage costs result in the (HL1-R) model being outperformed by the DW system. Indeed, the DW system actually incurs less cost overall compared to the case where there are more supplies arriving, unlike the hub system.

When supply amounts are increased by 25%, we observe opposite effects. In this case, the savings in the shortage costs have the biggest influence on the results. **For the hub system with a rolling horizon and location decisions, a 25% increase in supply results in more than 35% of shortage costs**. Since more demand can be satisfied in this case, transportation costs increase, more so for the DW system as DW-DL shipments are more frequent. Whereas models for the hub system incur comparable (in the case of HL1-R) to or better costs (in the case of HL0 and HL1-N) than the baseline, the opposite is true for the DW system.

An important determinant of the relative performance of the two network structures is the relative values of unit shortage, expiration, and transportation costs. To

assess the effect of changing unit costs, we extend our experiments to include cases where (i) unit expiration costs and (ii) unit transportation costs are doubled, with all else being equal. The results for these experiments are shown in Figure 44.

With increasing unit expiration costs, a hub system is no longer convenient. Under this scenario, the optimal solutions under all models tend to push the commodities to the distribution locations as quickly as possible, to avoid expirations at the central warehouse. Due to uncertainty of demand, this early commitment leads to increased level of shortages. Furthermore, due to the increase in the unit cost, expiration costs also increase. Due to better forecast accuracy, the district warehouse system is less affected by these increases. Despite the increased transportation costs due to the increased number of downstream shipments, the district warehouse system outperforms the hub system under the given settings.

The intuition with increased unit transportation costs is that the district warehouse system might be affected more, and the hub system is expected to outperform this network type. However, the computational results indicate that **even with increased unit transportation costs, a hub system outperforms the DW system only with hub locations as decision variables and incorporation of actual usage information**. This is mainly because both systems tend to make shipments in larger batches to decrease the transportation costs, which leads to early commitment as in the case of increased unit expiration costs. As a result, shortage and expiration costs increase, offsetting the relative performance loss for the district warehouse system.

In the last set of our experiments, we assess the equity of the allocations resulting from the proposed models, and analyze the results of the models based on equity and compare the performances of the two network structures when equity is the primary concern. Figure 45 shows a box-whisker plot of the percent demand satisfied (*coverage*) over the DLs. The first pair corresponds to the cost minimization solutions

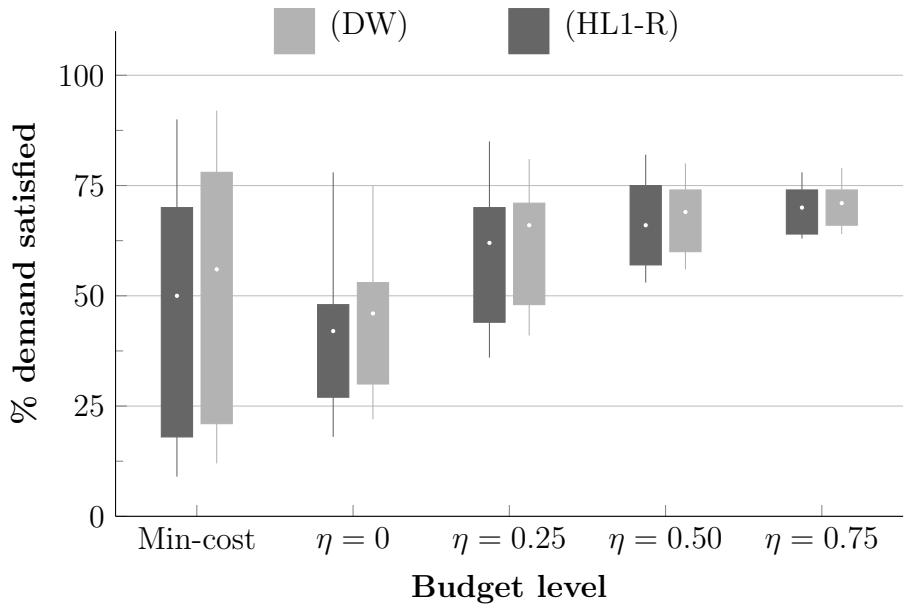

Figure 45: Box-whisker plots for the fraction of demand satisfied over all DLs for (i) the minimum cost solution, and (ii) max-min coverage solution with varying budget levels under the (HL1-R) and (DW) models

by (HL1-R) and (DW) for the baseline case, whereas the remaining pairs correspond to the equity maximizing solutions (HL1-RE) and (DW-E) with varying total cost limits.

When the coverage levels **under the cost minimizing models** are considered, we observe that **the district warehouse system outperforms the hub system in terms of minimum and average coverage** (11% vs. 8% for minimum coverage, and 56% vs. 50% for average coverage), which can be attributed to the accuracy of aggregate forecasts and the increased frequency of shipments from the DWs to DLs. When the objective is to maximize the minimum coverage over the network, the minimum coverage can be improved by 9% and 10% for the DW and hub systems, respectively. However, these savings can only be achieved at the expense of 10% and 8% of loss on the average coverage for these two network structures.

For the district warehouse system, consecutive 25% increments in the total cost limit result in minimum coverage increases of 19%, 15%, and 8%, implying diminishing

marginal returns. This also applies to the hub network, where the incremental gains are 18%, 17%, and 10%. A similar diminishing returns exists for the average coverage (20%, 3% and 2% increases for the DW system; 20%, 5% and 4% for the hub system). For all levels, the minimum and average coverage under the DW system is higher than that of the hub system (the differences are 13% and 4%, respectively for $\eta = 0$, and decrease as η increases), which is attributable to similar reasons as in the cost minimization case. At $\eta = 0.75$, the two networks result in almost identical coverage levels, which can be considered as the limit coverage that can be achieved, regardless of the network type. This is due to the fact that when the total cost constraint is not tight, the allocations are limited by the transportation capacity and supply availability.

Another important observation from our analysis of equity is that while the equity performance difference between the two network structures is statistically significant for the cost minimization models, the same conclusion holds only with the same cost limit under the equity maximizing models.

4.6 Conclusions and Further Research Directions

The usage of specialized nutritious foods is becoming increasingly popular throughout the globe due to the effectiveness of these foods in prevention and treatment of moderate and severe acute malnutrition. SNF exhibit challenging features such as high commodity value, low shelf life, long lead times, and highly uncertain demand, which pose important challenges for the management of SNF supply chains.

In this chapter, we have presented a case study which mainly addresses the decision of choosing a network structure for the SNF supply chain. To this end, we developed multi-period commodity allocation and transportation models that consider the perishability of the commodities, choosing the locations of hubs, potential

inaccuracy of usage forecasts, and incorporation of information updates as actual usage amounts become known. To the best of our knowledge, our models are the first in the literature to consider perishable commodities in a multiechelon supply chain with uncertain demand, limited supply and transportation capacities.

Our analysis of the alternative network structures through computational experiments reveals that the choice is not straightforward, and needs to be made carefully considering various aspects of the network, supply, demand, and cost structure. In particular, our results indicate that a hub system is preferred in systems where demand variance is low, forecast accuracy is high, supply is not scarce, transportation costs are relatively high compared to the costs of expiration and shortage, and when hub locations and information updates are incorporated into the decision making process. If equity is the primary objective, the district warehouse system performs slightly better, unless the total cost limit is sufficiently high. It is important to note that our conclusions also depend on a set of important assumptions, such as independent demand over time and locations, fixed shipment frequencies at each level, and no fixed cost of changing to a hub structure and opening the facilities.

The analysis in this chapter can be extended in a number of different ways. First, the experiments are performed assuming a fixed network structure. The sensitivity of the results in §4.6 needs to be analyzed using different network structures in different countries or regions. The choice of cost parameters also influences the results of the analysis. While the changes in the results from varying the cost parameters can be analyzed to a certain extent from our experiments, a more detailed sensitivity analysis sheds light into the strength of these influences. Supply arrivals have been assumed as given due to the dependence on funding. the system performance can be improved by decisions on the timing and quantity of orders from the suppliers. In comparing the two systems when hub locations are part of decisions, the fixed costs of these hubs were ignored. For a more realistic comparison, it needs to be kept in mind that

locating these facilities will incur additional costs.

An important extension of this work is on how demand uncertainty is handled. While assuming a discrete distribution for demand in each period, assuming more structured distributions may reveal certain structural results, which in turn can be used for developing heuristic solution approaches. A multi-stage stochastic programming approach, although with increased complexity, also provides handling other aspects of the system such as dependence of demand through periods and locations.

CHAPTER V

CONCLUSIONS AND FURTHER RESEARCH DIRECTIONS

In this thesis, we considered three applications of humanitarian, emergency response and public health supply chains where the main decisions relate to the allocation of scarce resources and there is inherent uncertainty in the problem environment. We considered the stochastic nature of the problems and the use of information updates (i) analytically, by analyzing simple special cases or performance under worst-case scenarios, and (ii) through computational experiments based on real-life data.

Chapter 2 has presented models and heuristics to sequence the clearance of debris-blocked roads in the aftermath of a disaster to establish connectivity between relief supply and demand. Given the uncertainty of debris amounts due to lack of geographical/demographic data or information systems, there is a trade-off between "learning" the uncertain aspects of the network by exploration and greedily aiming to establish connectivity by exploiting the existing information on the network. An interesting aspect of this problem is that learning and establishing connectivity are interdependent, i.e., the decisions of which roads to clear in each period to establish connectivity determine which roads will become accessible, and hence "what will be learned." Through computational experiments, we have shown that incorporating stochasticity is important when initial beliefs on the debris amounts are certain. When these beliefs are accurate, the quality of estimates result in improved solution quality. With inaccurate but certain beliefs, the proposed algorithms make better use of learning compared to those that ignore the stochasticity. When beliefs are uncertain, a regional update mechanism enhances learning and produces better results.

Our work in Chapter 2 leads to many interesting future research directions. The proposed continuous-time approximation, which relaxes the requirement to make decisions at the beginning of each period and allows new decisions after the clearance of each road, is applicable to many stochastic network problems that consist of multiple periods and involve the trade-off between exploring the uncertainties of the network and exploiting existing information to greedily establishing connectivity. Examples of these problems include infrastructure repair and network surveillance. The learning structures can be enhanced in a number of ways. For example, the frequency of regional updates can be made part of the decisions. Especially when variability of estimates is high, it might be more useful to make the updates less frequently, as aggregate forecasts will be more accurate. In Chapter 2, we have assumed that once a clearance decision has been made for a road segment, this road has to be cleared completely before proceeding. In other practical applications, there may be flexibility of not finishing the recovery of an edge if the actual clearance resource requirement turns out to be larger than expected. An alternative update approach is to observe remote edges by sending part of the clearance resources to these.

In Chapter 3, we focused on the allocation of a perishable commodity to multiple distribution locations that serve multiple classes of demand. We have justified using stochastic newsvendor-type allocation models by comparing these to benchmark approaches that make proportional allocations by assuming deterministic demand equal to the demand rates, and showed that considering stochasticity is especially important when unit loss due to overage is higher. Additional areas of focus in this chapter were the effects of usage policy and centralization. For the former, we have shown analytical results that demonstrate the worst-case performance of first-come, first-served (FCFS) usage relative to that under prioritization. Through computational experiments, we have also analyzed the effect of making allocations assuming a certain usage policy (e.g., prioritization) at the distribution locations, whereas actual

usage policy is different (e.g., FCFS). For the latter, we have analytically shown cases where centralization may or may not improve expected system performance.

Future work regarding perishable commodity allocation to multiple locations and demand types involves a number of directions. When multiple allocations throughout the shelf life of a commodity are possible, a multi-period stochastic model is useful to determine the optimal allocation under the two usage policies. Furthermore, such a model allows consideration of age-dependent benefits and the value of information updates. While computational results have shown the detrimental effects of assuming a usage policy different from the actual one when making the allocations, a worst-case analysis is still open to exploration. A similar analysis can be performed for the effects of centralization as well.

Chapter 4 has considered the network design and allocation of high value, low shelf life commodities to prevent and treat undernutrition. Demand uncertainty has been incorporated into the models by means of discrete scenarios, and information updates have been considered by solving the proposed models on a rolling horizon basis, as actual demand information becomes available. We have shown the importance of stochasticity and information updates using computational experiments based on data from the operations of a large humanitarian organization. Equity has been considered as a second objective, and separate models have been developed to analyze the equity of proposed cost minimization models.

The work in Chapter 4 can be extended by considering different network structures in other countries and by other organizations to generalize the lessons learned. Another direction is to derive structural results assuming more structured demand distributions and possibly a simpler network structure, and compare these results to those in the literature on perishable and multiple echelon commodity allocation.

APPENDIX A

FURTHER COMPUTATIONAL RESULTS FOR RANDOM INSTANCES

Table 15 gives the average percent gaps from perfect information solutions for random instances when p_{ij} follows a uniform distribution with parameters $U(0.5, 1)$ and there are two supply nodes. Similarly, Table 16 shows the average gaps for the same belief settings with a single supply node, where clearance starts from outside the disaster area, instead of the supply node. For both cases, the results are in parallel with the baseline case we present in §2.6, where there is a single supply node and the clearance starts from this supply node.

Table 17 extends the results given in Table 7 to the cases where p_{ij} follows a uniform distribution with parameters $U(0, 0.5)$, $U(0, 1)$, $U(0.25, 0.75)$, and $U(0.5, 1)$.

Table 15: Average perfect information gaps for random instances under p_{ij} following a uniform distribution with parameters $U(0.5, 1)$ and two supply nodes

		Deterministic with perfect information		Deterministic with expected amounts		No regional updates					With regional updates				
						Periodic	Continuous approximation				Periodic	Continuous approximation			
Instance	% roads blocked	Optimal	Heuristic	Optimal	Heuristic	POMDP	POMDP	HP	LL	HP & LL	POMDP	POMDP	HP	LL	HP & LL
$G(5 \times 5)$	100	0.00	1.61	23.13	22.33	15.16	15.29	15.53	15.49	16.11	14.12	14.49	16.11	16.44	15.41
$G(10 \times 10)$	25	0.00	3.12	25.55	25.98	16.01	17.79	18.57	18.97	18.62	15.86	17.86	18.00	18.95	17.78
$G(10 \times 10)$	50	***	4.01	***	31.35	***	20.02	19.82	21.25	20.77	***	17.69	18.66	19.88	20.72
$G(10 \times 10)$	100	***	4.64	***	33.09	***	***	21.10	20.81	21.38	***	***	20.30	21.13	19.76
$G(15 \times 15)$	25	***	4.45	***	27.70	***	18.30	19.10	19.83	19.21	***	16.57	18.01	17.87	19.65
$G(15 \times 15)$	50	***	5.33	***	29.23	***	19.46	19.81	20.82	20.47	***	18.17	19.99	18.79	19.79
$G(15 \times 15)$	100	***	6.10	***	34.99	***	***	19.90	21.87	21.79	***	***	20.49	20.56	20.73
$R(5 \times 5)$	100	0.00	2.46	25.14	25.32	15.86	16.37	17.78	18.18	18.65	15.29	15.95	16.93	17.17	16.58
$R(10 \times 10)$	25	0.00	3.74	27.78	28.23	18.98	19.71	18.88	19.28	18.66	16.06	16.98	17.25	18.10	18.54
$R(10 \times 10)$	50	***	3.78	***	33.07	***	19.90	19.40	19.52	20.46	***	18.47	17.85	18.18	19.49
$R(10 \times 10)$	100	***	4.09	***	34.71	***	***	20.05	20.34	20.86	***	***	19.08	18.90	20.29
$R(15 \times 15)$	25	***	4.66	***	31.88	***	19.19	20.18	20.39	22.43	***	19.45	20.32	20.60	19.59
$R(15 \times 15)$	50	***	5.01	***	35.80	***	19.64	21.90	20.65	20.71	***	19.55	20.61	21.00	20.47
$R(15 \times 15)$	100	***	6.27	***	35.87	***	***	22.13	22.64	24.15	***	***	21.39	21.29	21.85

Table 16: Average perfect information gaps for random instances under p_{ij} following a uniform distribution with parameters $U(0.5, 1)$ and clearance starting from outside the disaster area

		Deterministic with perfect information		Deterministic with expected amounts		No regional updates					With regional updates				
						Periodic	Continuous approximation				Periodic	Continuous approximation			
Instance	% roads blocked	Optimal	Heuristic	Optimal	Heuristic	POMDP	POMDP	HP	LL	HP & LL	POMDP	POMDP	HP	LL	HP & LL
$G(5 \times 5)$	100	0.00	1.58	43.03	44.60	23.45	23.77	24.46	25.67	25.99	22.84	24.30	26.49	27.31	29.14
$G(10 \times 10)$	25	0.00	3.61	46.05	50.46	26.41	27.49	30.65	31.44	31.61	26.95	28.23	29.02	30.71	33.00
$G(10 \times 10)$	50	***	5.18	***	55.37	***	31.23	31.37	34.04	36.24	***	29.49	31.67	32.25	32.19
$G(10 \times 10)$	100	***	3.64	***	56.55	***	***	35.69	34.82	35.85	***	***	32.55	34.97	38.47
$G(15 \times 15)$	25	***	4.89	***	49.06	***	27.12	29.44	31.07	33.35	***	29.11	29.27	31.65	32.49
$G(15 \times 15)$	50	***	5.60	***	54.41	***	***	***	33.28	31.94	***	***	***	33.19	35.01
$G(15 \times 15)$	100	***	6.06	***	65.81	***	***	***	32.55	36.15	***	***	***	32.84	36.76
$R(5 \times 5)$	100	0.00	2.68	36.78	41.27	25.38	26.63	29.14	29.69	29.98	25.60	27.13	28.82	28.03	32.89
$R(10 \times 10)$	25	0.00	3.25	43.49	49.90	27.92	29.11	28.55	29.25	29.64	29.12	30.20	32.39	33.29	33.54
$R(10 \times 10)$	50	***	3.82	***	58.93	***	31.73	32.69	33.25	32.53	***	31.16	32.84	32.35	35.98
$R(10 \times 10)$	100	***	4.62	***	64.04	***	***	33.23	36.57	34.66	***	***	31.17	34.03	36.24
$R(15 \times 15)$	25	***	4.91	***	57.23	***	29.15	32.51	33.23	34.88	***	30.63	33.81	32.94	35.26
$R(15 \times 15)$	50	***	4.79	***	58.49	***	***	***	33.14	36.18	***	***	***	34.48	36.80
$R(15 \times 15)$	100	***	6.16	***	61.33	***	***	***	39.27	39.30	***	***	***	39.51	39.37

Table 17: Average CPU times (in seconds) for the proposed and benchmark approaches on random instances

			Deterministic with perfect information		Deterministic with expected amounts		No regional updates					With regional updates				
							Periodic	Continuous approximation				Periodic	Continuous approximation			
Distribution of p_{ij}	Instance	% roads blocked	Optimal	Heuristic	Optimal	Heuristic	POMDP	POMDP	HP	LL	HP & LL	POMDP	POMDP	HP	LL	HP & LL
	$G(5 \times 5)$	100	24.89	6.10	27.50	3.55	1479.61	271.72	193.15	226.33	72.14	1343.72	294.66	190.78	212.89	66.38
	$G(10 \times 10)$	25	81.77	13.89	74.31	11.83	6284.43	523.71	262.86	273.89	88.99	6181.25	453.48	275.34	312.10	86.57
	$G(10 \times 10)$	50	***	36.21	***	21.47	***	2173.06	1117.96	1320.89	407.74	***	1978.03	1067.74	1770.20	389.96
	$G(10 \times 10)$	100	***	99.38	***	78.30	***	***	2217.19	2404.34	900.66	***	***	2132.19	2854.77	870.75
	$G(15 \times 15)$	25	***	29.19	***	20.43	***	430.59	214.16	304.70	100.35	***	542.03	274.36	356.63	99.71
	$G(15 \times 15)$	50	***	71.34	***	28.93	***	***	***	1819.47	525.24	***	***	***	2215.93	583.79
$U(0, 0.5)$	$G(15 \times 15)$	100	***	147.71	***	135.58	***	***	***	4104.47	1136.02	***	***	***	3710.65	1397.90
	$R(5 \times 5)$	100	32.53	6.94	31.79	3.80	1120.56	173.04	119.53	166.68	52.02	1210.25	225.86	136.07	154.87	48.25
	$R(10 \times 10)$	25	83.21	13.64	64.03	12.24	3077.16	249.37	125.54	180.15	75.07	3462.87	263.68	129.87	154.24	94.95
	$R(10 \times 10)$	50	***	45.31	***	28.22	***	1957.23	1094.72	1147.77	417.82	***	2145.04	1085.43	1331.45	456.09
	$R(10 \times 10)$	100	***	116.86	***	93.91	***	***	1579.95	2784.09	822.13	***	***	1713.58	2665.49	669.99
	$R(15 \times 15)$	25	***	25.38	***	13.36	***	422.78	191.89	267.59	81.61	***	435.98	185.66	278.40	95.08
	$R(15 \times 15)$	50	***	67.04	***	77.15	***	***	***	1735.73	616.41	***	***	***	1845.61	623.39
	$R(15 \times 15)$	100	***	163.47	***	138.77	***	***	***	3615.20	896.01	***	***	***	3822.38	913.98
	$G(5 \times 5)$	100	24.00	5.95	29.22	3.85	1387.71	237.10	178.11	211.97	68.47	1374.70	260.85	167.14	230.64	65.20
	$G(10 \times 10)$	25	87.41	14.76	72.45	12.54	5773.54	437.68	290.58	303.68	104.22	7190.76	464.54	254.20	286.27	81.49
	$G(10 \times 10)$	50	***	33.83	***	23.06	***	2036.47	1208.23	1334.54	440.85	***	1907.62	1038.68	1705.58	350.35
	$G(10 \times 10)$	100	***	107.47	***	71.14	***	***	2033.88	2408.62	988.70	***	***	2198.97	2783.91	770.08
	$G(15 \times 15)$	25	***	29.76	***	17.63	***	462.49	204.33	297.90	99.02	***	537.46	236.89	375.10	101.54
	$G(15 \times 15)$	50	***	68.89	***	28.76	***	***	***	1696.07	541.06	***	***	***	2165.10	585.89
$U(0, 1)$	$G(15 \times 15)$	100	***	155.96	***	129.78	***	***	***	4494.28	1114.35	***	***	***	3697.79	1268.11
	$R(5 \times 5)$	100	35.58	6.95	36.01	3.46	985.74	184.76	121.97	163.12	51.95	1048.04	248.61	134.81	179.42	41.69
	$R(10 \times 10)$	25	95.52	12.66	69.28	12.18	3044.39	267.35	122.53	172.08	84.52	3551.31	265.64	117.45	150.67	108.80
	$R(10 \times 10)$	50	***	49.40	***	25.50	***	1718.12	1101.08	1157.62	371.76	***	2354.54	1020.73	1192.11	462.15
	$R(10 \times 10)$	100	***	117.74	***	104.74	***	***	1449.64	2802.91	856.10	***	***	1622.34	2898.14	657.88
	$R(15 \times 15)$	25	***	26.77	***	15.03	***	454.08	203.30	252.75	79.89	***	421.44	176.62	303.16	78.21
	$R(15 \times 15)$	50	***	64.28	***	76.30	***	***	***	1932.85	640.71	***	***	***	1656.13	699.35
	$R(15 \times 15)$	100	***	190.02	***	137.42	***	***	***	3062.59	1014.86	***	***	***	3617.74	917.88
	$G(5 \times 5)$	100	24.46	6.13	25.44	3.77	1545.64	229.58	169.12	182.22	72.04	1500.60	278.45	159.91	232.78	59.63
	$G(10 \times 10)$	25	85.33	14.76	75.71	11.31	5901.14	409.24	247.70	299.01	101.57	7514.09	441.79	257.08	312.39	73.73
	$G(10 \times 10)$	50	***	35.73	***	25.06	***	1829.71	1052.16	1392.42	501.18	***	2045.46	974.34	1887.58	364.99
	$G(10 \times 10)$	100	***	118.05	***	65.75	***	***	2033.24	2416.52	1007.44	***	***	1907.02	2475.25	747.13
	$G(15 \times 15)$	25	***	29.57	***	15.95	***	475.63	212.53	283.93	109.00	***	539.11	254.80	383.75	83.73
	$G(15 \times 15)$	50	***	76.33	***	28.04	***	***	***	1523.40	559.23	***	***	***	2172.45	520.58
$U(0.25, 0.75)$	$G(15 \times 15)$	100	***	177.58	***	119.87	***	***	***	4419.80	1212.60	***	***	***	3419.55	1206.58
	$R(5 \times 5)$	100	37.02	6.64	41.52	3.93	1123.76	180.33	116.59	168.19	53.31	1070.69	272.51	139.68	194.41	40.43
	$R(10 \times 10)$	25	95.77	11.86	66.60	11.65	2996.91	260.83	140.53	174.32	85.19	3682.58	228.08	111.39	138.12	107.05
	$R(10 \times 10)$	50	***	49.83	***	28.62	***	1633.16	1065.56	1093.98	400.69	***	2359.42	1142.92	1088.27	469.28
	$R(10 \times 10)$	100	***	121.26	***	105.65	***	***	1463.05	2916.12	805.22	***	***	1466.72	3371.43	661.93
	$R(15 \times 15)$	25	***	29.60	***	16.14	***	454.73	206.06	247.07	70.27	***	422.84	148.12	303.27	69.78
	$R(15 \times 15)$	50	***	66.85	***	76.51	***	***	***	1958.45	628.68	***	***	***	1799.60	670.48
	$R(15 \times 15)$	100	***	199.66	***	121.94	***	***	***	2980.52	1056.07	***	***	***	3940.47	915.90

APPENDIX B

PROOFS OF STATEMENTS IN CHAPTER 3

Proof of Proposition 1. First, we show that the stationary point of Equation (36) is indeed given by Equation (37). We prove this by induction on $|J|$. When there is only a single class, we have:

$$E_P(z|x) = \int_x^\infty w_1\, x\, f_1(u_1)du_1 + \int_0^x (w_1\, u_1 - w_o(x-u_1))f_1(u_1)du_1.$$

Taking the derivative and setting equal to zero, we obtain:

$$\int_x^\infty w_1\, f_1(u_1)du_1 - w_1\, x\, f_1(x) - \int_0^x w_o\, f_1(u_1)du_1 + w_1\, x\, f_1(x) = 0,$$

which yields:

$$w_1\, P(D_1 > x) = w_0\, P(D_1 \leq x),$$

which is the desired equality. Assume now that when $|J|= n$, the stationary point of Equation (36) is given by Equation (37), that is, the stationary point x of Equation (36) satisfies:

$$w_1\, P(D_1 > x) + \sum_{j=2}^{n} w_j \left[P\left(\sum_{l=1}^{j-1} D_l \leq x \right) - P\left(\sum_{l=1}^{j} D_l \leq x \right) \right] = w_0\, P\left(\sum_{l=1}^{j} D_l \leq x \right).$$

Before proceeding, it should be noted that Equation (36) can be rewritten as follows:

$$\begin{aligned} E_P(z|x) =& w_1 \left(\int_x^\infty x\, f_1(u_1)du_1 \right) + \sum_{j=2}^{n} w_j \left[\int_0^x \cdots \int_0^{x-U_{n-1}} \left(\int_{x-U_{j-1}}^\infty (x - U_{j-1}) f_j(u_j)du_j \right.\right. \\ &\left.\left. + \int_0^{x-U_{j-1}} u_j f_j(u_j)du_j \right) f_{j-1}(u_{j-1})du_{j-1} \ldots f_2(u_2)du_2\, f_1(u_1)du_1 \right] \\ &- w_0 \left(\int_0^x \cdots \int_0^{x-U_{n-1}} (x - U_n)\, f_n(u_n)du_n \ldots f_2(u_2)du_2\, f_1(u_1)du_1 \right). \end{aligned} \tag{123}$$

For $|J|= n+1$, the expression for the expected net benefit is given by:

$$E(z|x) = \int_x^\infty w_1\, x f_1(u_1)du_1 + \int_0^x \int_{x-u_1}^\infty (w_1u_1 + w_2\,(x-u_1))\, f_2(u_2)du_2\, f_1(u_1)du_1 + \dots$$
$$+ \int_0^x \int_0^{x-u_1} \dots \int_0^{x-U_{n-1}} \int_{x-U_n}^\infty (w_1u_1 + \dots w_nu_n + w_n\,(x-U_n))\, f_{n+1}(u_{n+1})du_{n+1} \dots f_1(u_1)du_1$$
$$+ \int_0^x \int_0^{x-u_1} \dots \int_0^{x-U_n} (w_1u_1 + \dots + w_nu_n - w_0\,(x-U_n))\, f_{n+1}(u_{n+1})du_{n+1} \dots f(u_1)du_1,$$

which can be expanded as:

$$w_1\Bigg(\int_x^\infty x\, f_1(u_1)du_1 + \int_0^x \int_{x-u_1}^\infty u_1\, f_2(u_2)du_2\, f_1(u_1)du_1 + \dots$$
$$+ \int_0^x \int_0^{x-u_1} \dots \int_0^{x-U_{n-1}} \int_{x-U_n}^\infty u_1\, f_{n+1}(u_{n+1})du_{n+1}\, f_n(u_n)du_n \dots f_2(u_2)du_2\, f_1(u_1)du_1$$
$$+ \int_0^x \int_0^{x-u_1} \dots \int_0^{x-U_{n-1}} \int_0^{x-U_n} u_1\, f_{n+1}(u_{n+1})du_{n+1}\, f_n(u_n)du_n \dots f_2(u_2)du_2\, f_1(u_1)du_1 \Bigg)$$
$$+w_2\Bigg(\int_0^x \int_{x-u_1}^\infty (x-u_1)\ f_2(u_2)du_2\, f_1(u_1)du_1$$
$$+ \int_0^x \int_0^{x-u_1} \int_{x-U_2}^\infty u_2\, f_3(u_3)du_3\, f_2(u_2)du_2\, f_1(u_1)du_1 + \dots$$
$$+ \int_0^x \int_0^{x-u_1} \dots \int_0^{x-U_{n-1}} \int_{x-U_n}^\infty u_2\, f_{n+1}(u_{n+1})du_{n+1}\, f_n(u_n)du_n \dots f_2(u_2)du_2\, f_1(u_1)du_1$$
$$+ \int_0^x \int_0^{x-u_1} \dots \int_0^{x-U_{n-1}} \int_0^{x-U_n} u_2\, f_{n+1}(u_{n+1})du_{n+1}\, f_n(u_n)du_n \dots f_2(u_2)du_2\, f_1(u_1)du_1 \Bigg)$$
$$+ \dots$$
$$+w_n\Bigg(\int_0^x \int_0^{x-u_1} \dots \int_0^{x-U_{n-2}} \int_{x-U_{n-1}}^\infty (x-U_{n-1})\, f_n(u_n)du_n \dots f_2(u_2)du_2\, f_1(u_1)du_1$$
$$+ \int_0^x \int_0^{x-u_1} \dots \int_0^{x-U_{n-2}} \int_0^{x-U_{n-1}} u_n f_n(u_n)du_n \dots f_2(u_2)du_2\, f_1(u_1)du_1$$

$$+\int_0^x \int_0^{x-u_1} \cdots \int_0^{x-U_{n-1}} \int_{x-U_n}^{\infty} u_n\, f_{n+1}(u_{n+1})du_{n+1} \ldots f_2(u_2)du_2\, f_1(u_1)du_1$$
$$+\int_0^x \int_0^{x-u_1} \cdots \int_0^{x-U_{n-1}} \int_0^{x-U_n} u_n\, f_{n+1}(u_{n+1})du_{n+1} \ldots f_2(u_2)du_2\, f_1(u_1)du_1\Bigg)$$
$$+w_{n+1}\Bigg(\int_0^x \int_0^{x-u_1} \cdots \int_0^{x-U_{n-1}} \int_{x-U_n}^{\infty} (x-U_n) f_{n+1}(u_{n+1})du_n \ldots f_2(u_2)du_2\, f_1(u_1)du_1$$
$$+\int_0^x \int_0^{x-u_1} \cdots \int_0^{x-U_{n-1}} \int_0^{x-U_n} u_{n+1}\, f_{n+1}(u_{n+1}du_{n+1}) \ldots f_2(u_2)du_2\, f_1(u_1)du_1\Bigg)$$
$$-w_0\Bigg(\int_0^x \int_0^{x-u_1} \cdots \int_0^{x-U_{n-1}} \int_0^{x-U_n} (x-U_{n+1})\, f_{n+1}(u_{n+1})du_{n+1} \ldots f_2(u_2)du_2\, f_1(u_1)du_1\Bigg).$$

This simplifies to:

$$w_1\left(\int_x^{\infty} x\, f_1(u_1)du_1\right) + \sum_{j=2}^{n} w_j\Bigg[\int_0^x \cdots \int_0^{x-U_{n-1}} \Bigg(\int_{x-U_{j-1}}^{\infty} (x-U_{j-1}) f_j(u_j)du_j$$
$$+\int_0^{x-U_{j-1}} u_j f_j(u_j)du_j\Bigg) f_{j-1}(u_{j-1})du_{j-1} \ldots f_2(u_2)du_2\, f_1(u_1)du_1\Bigg]$$
$$+\, w_{n+1}\Bigg(\int_0^x \int_0^{x-u_1} \cdots \int_0^{x-U_{n-2}} \int_{x-U_n}^{\infty} (x-U_n) f_{n+1}(u_{n+1})du_n \ldots f_2(u_2)du_2\, f_1(u_1)du_1$$
$$+\int_0^x \int_0^{x-u_1} \cdots \int_0^{x-U_{n-1}} \int_0^{x-U_n} u_{n+1}\, f_{n+1}(u_{n+1}du_{n+1}) \ldots f_2(u_2)du_2\, f_1(u_1)du_1\Bigg)$$
$$-\, w_0\Bigg(\int_0^x \int_0^{x-u_1} \cdots \int_0^{x-U_{n-1}} \int_0^{x-U_n} (x-U_{n+1})\, f_{n+1}(u_{n+1})du_{n+1} \ldots f_2(u_2)du_2\, f_1(u_1)du_1\Bigg).$$

The first n terms of the sum are identical to the first n terms of the sum in Equation (123). Since there are no terms including w_j, $j \leq n$ in the last two terms of the above sum, by the induction hypothesis, the derivative of the first n terms gives:

$$w_1\, P(D_1 > x) + \sum_{j=2}^{n} w_j \left[P\left(\sum_{l=1}^{j-1} D_l \leq x\right) - P\left(\sum_{l=1}^{j} D_l \leq x\right)\right].$$

Hence, to complete the proof that x given by Equation (37) is the stationary point of the expected net benefit function, we need two more steps. First, we need to show that the derivative of the $n+1^{st}$ term in the sum is $P\Big(\sum_{l=1}^{n} D_l\Big) - P\Big(\sum_{l=1}^{n+1} D_l\Big)$. Indeed:

$$\begin{aligned}
&\frac{d}{dx}\Bigg(\int_0^x \int_0^{x-u_1} \dots \int_0^{x-U_{n-2}} \int_{x-U_n}^{\infty} (x-U_n) f_{n+1}(u_{n+1})du_n \dots f_2(u_2)du_2\, f_1(u_1)du_1 \\
&\quad + \int_0^x \int_0^{x-u_1} \dots \int_0^{x-U_{n-1}} \int_0^{x-U_n} u_{n+1}\, f_{n+1}(u_{n+1}du_{n+1}) \dots f_2(u_2)du_2\, f_1(u_1)du_1\Bigg) \\
&= \int_0^x \int_0^{x-u_1} \dots \int_0^{x-U_{n-1}} \int_{x-U_n}^{\infty} f_{n+1}u_{n+1}du_{n+1}\, f_n u_n du_n \dots f_2(u_2)du_2\, f_1(u_1)du_1 \\
&= \int_0^x \int_0^{x-u_1} \dots \int_0^{x-U_{n-1}} P\Bigg(D_{n+1} > x - U_n \Big| D_1 = u_1, \dots, D_n = u_n\Bigg) \\
&= P\Bigg(\sum_{l=1}^{n} D_l \le x, \sum_{l=1}^{n+1} D_l > x\Bigg) = P\Bigg(\sum_{l=1}^{n} D_l\Bigg) - \Bigg(\sum_{l=1}^{n+1} D_l\Bigg).
\end{aligned}$$

Second, we need to show that the derivative of the last term in the sum gives $P\Big(\sum_{l=1}^{n+1} D_l \le x\Big)$.

$$\begin{aligned}
&\frac{d}{dx}\Bigg(\int_0^x \int_0^{x-u_1} \dots \int_0^{x-U_{n-1}} \int_0^{x-U_n} (x-U_{n+1})\, f_{n+1}(u_{n+1})du_{n+1} \dots f_2(u_2)du_2\, f_1(u_1)du_1\Bigg) \\
&= \int_0^x \int_0^{x-u_1} \dots \int_0^{x-U_{n-1}} P\Bigg(D_{n+1} \le x - U_n \Big| D_1 = u_1, \dots, D_n = u_n\Bigg) = P\Bigg(\sum_{l=1}^{n+1} D_l \le x\Bigg).
\end{aligned}$$

To show that x given by Equation (37) is the global minimizer of the expected net benefit equation (36), we take its second derivative:

$$\begin{aligned}
\frac{d^2}{dx^2}E_P(z|x) &= \frac{d}{dx}\Bigg[\sum_{j=1}^{|J|-1}(w_j - w_{j+1})P\left(\sum_{l=1}^{j} D_l > x\right) + w_n P\left(\sum_{l=1}^{|J|} D_l > x\right) - w_0\, P\left(\sum_{l=1}^{j} D_l \le x\right)\Bigg] \\
&= -\sum_{j=1}^{|J-1|}(w_j - w_{j+1})\, f_{D_1+\dots+D_j}(x) - w_0\, f_{D_1+\dots+D_{|J|}}(x).
\end{aligned}$$

Since $w_j \ge w_{j+1}$ and $w_0 \ge 0$, the second derivative is always nonpositive, implying $E_P(z|x)$ is concave and the stationary point x defined by Equation (37) is the global

minimizer. □

Proof of Proposition 2. For brevity, we will show a sketch of the proof for the marginal benefit equation when $|J|= 2$. The expected net benefit equation under x units of allocation is given by:

$$
\begin{aligned}
E_F(z|x) = &\sum_{u_1=0}^{x-1}\sum_{u_2=0}^{x-u_1-1}(w_1\,u_1 + w_2\,u_2 - w_0\,(x-u_1-u_2))\,\frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!}\frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!}\\
&+\sum_{u_1=0}^{\infty}\sum_{u_2=(0,x-u_1)^+}^{\infty}\sum_{k=(0,x-u_2)^+}^{(x,u_1)^-}\frac{\binom{u_1}{k}\binom{u_2}{x-k}}{\binom{u_1+u_2}{x}}(w_1\,k + w_2(x-k))\frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!}\frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!}\\
=&\sum_{u_1=0}^{x-1}\sum_{u_2=0}^{x-u_1-1}(w_1\,u_1 + w_2\,u_2 - w_0\,(x-u_1-u_2))\,\frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!}\frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!}\\
&+\sum_{u_1=0}^{x}\sum_{u_2=x-u_1}^{x}\sum_{k=x-u_2}^{u_1}\frac{\binom{u_1}{k}\binom{u_2}{x-k}}{\binom{u_1+u_2}{x}}(w_1\,k + w_2(x-k))\frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!}\frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!}\\
&+\sum_{u_1=0}^{x}\sum_{u_2=0}^{x}\sum_{k=0}^{u_1}\frac{\binom{u_1}{k}\binom{u_2}{x-k}}{\binom{u_1+u_2}{x}}(w_1\,k + w_2(x-k))\frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!}\frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!}\\
&+\sum_{u_1=x+1}^{\infty}\sum_{u_2=0}^{x}\sum_{k=x-u_2}^{x}\frac{\binom{u_1}{k}\binom{u_2}{x-k}}{\binom{u_1+u_2}{x}}(w_1\,k + w_2(x-k))\frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!}\frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!}\\
&+\sum_{u_1=x+1}^{\infty}\sum_{u_2=x-u_1}^{\infty}\sum_{k=0}^{x}\frac{\binom{u_1}{k}\binom{u_2}{x-k}}{\binom{u_1+u_2}{x}}(w_1\,k + w_2(x-k))\frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!}\frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!}
\end{aligned}
$$

The marginal benefit function at x units of allocation is given as:

$$
\begin{aligned}
MB_F(x) =&E_F(z|x+1) - E_F(z|x)\\
=&\sum_{u_1=0}^{x}\sum_{u_2=0}^{x-u_1}(w_1\,u_1 + w_2\,u_2 - w_0\,(x+1-u_1-u_2))\,\frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!}\frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!}\\
&+\sum_{u_1=0}^{x+1}\sum_{u_2=x+1-u_1}^{x+1}\sum_{k=x+1-u_2}^{u_1}\frac{\binom{u_1}{k}\binom{u_2}{x+1-k}}{\binom{u_1+u_2}{x+1}}(w_1\,k + w_2(x+1-k))\frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!}\frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!}\\
&+\sum_{u_1=0}^{x+1}\sum_{u_2=0}^{x+1}\sum_{k=0}^{u_1}\frac{\binom{u_1}{k}\binom{u_2}{x+1-k}}{\binom{u_1+u_2}{x+1}}(w_1\,k + w_2(x+1-k))\frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!}\frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!}\\
&+\sum_{u_1=x+2}^{\infty}\sum_{u_2=0}^{x+1}\sum_{k=x-u_2}^{x+1}\frac{\binom{u_1}{k}\binom{u_2}{x+1-k}}{\binom{u_1+u_2}{x+1}}(w_1\,k + w_2(x+1-k))\frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!}\frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!}\\
&+\sum_{u_1=x+2}^{\infty}\sum_{u_2=x+1-u_1}^{\infty}\sum_{k=0}^{x+1}\frac{\binom{u_1}{k}\binom{u_2}{x+1-k}}{\binom{u_1+u_2}{x+1}}(w_1\,k + w_2(x+1-k))\frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!}\frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!}\\
&-\sum_{u_1=0}^{x-1}\sum_{u_2=0}^{x-u_1-1}(w_1\,u_1 + w_2\,u_2 - w_0\,(x-u_1-u_2))\,\frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!}\frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!}\\
&-\sum_{u_1=0}^{x}\sum_{u_2=x-u_1}^{x}\sum_{k=x-u_2}^{u_1}\frac{\binom{u_1}{k}\binom{u_2}{x-k}}{\binom{u_1+u_2}{x}}(w_1\,k + w_2(x-k))\frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!}\frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!}\\
&-\sum_{u_1=0}^{x}\sum_{u_2=0}^{x}\sum_{k=0}^{u_1}\frac{\binom{u_1}{k}\binom{u_2}{x-k}}{\binom{u_1+u_2}{x}}(w_1\,k + w_2(x-k))\frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!}\frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!}
\end{aligned}
$$

$$
\begin{aligned}
&- \sum_{u_1=x+1}^{\infty} \sum_{u_2=0}^{x} \sum_{k=x-u_2}^{x} \frac{\binom{u_1}{k}\binom{u_2}{x-k}}{\binom{u_1+u_2}{x}} (w_1\, k + w_2(x-k)) \frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!} \frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!} \\
&- \sum_{u_1=x+1}^{\infty} \sum_{u_2=x-u_1}^{\infty} \sum_{k=0}^{x} \frac{\binom{u_1}{k}\binom{u_2}{x-k}}{\binom{u_1+u_2}{x}} (w_1\, k + w_2(x-k)) \frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!} \frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!} \\
&= \sum_{u_1=0}^{x} \left(w_1\, u_1 + w_2\,(x-u_1)\right) \frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!} \frac{e^{-\lambda_2}(\lambda_2)^{x-u_1}}{(x-u_1)!} \\
&+ \sum_{u_2=0}^{x+1} \sum_{k=x-u_2-1}^{x+1} \frac{\binom{x+1}{k}\binom{u_2}{x+1-k}}{\binom{x+1+u_2}{x+1}} (w_1\, k + w_2(x+1-k)) \frac{e^{-\lambda_1}(\lambda_1)^{x+1}}{x+1!} \frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!} \\
&+ \sum_{u_1=0}^{x} \left(\sum_{u_2=x-u_1}^{x} \sum_{k=x-u_2}^{u_1} \frac{\binom{u_1}{k}\binom{u_2+1}{x+1-k}}{\binom{u_1+u_2+1}{x+1}} (w_1\, k + w_2(x+1-k)) \frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!} \frac{e^{-\lambda_2}(\lambda_2)^{u_2+1}}{(u_2+1)!} \right. \\
&\qquad \left. - \sum_{u_2=x-u_1}^{x} \sum_{k=x-u_2}^{u_1} \frac{\binom{u_1}{k}\binom{u_2}{x-k}}{\binom{u_1+u_2}{x}} (w_1\, k + w_2(x-k)) \frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!} \frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!} \right) \\
&+ \sum_{u_2=x+2}^{\infty} \sum_{k=0}^{x+1} \frac{\binom{x+1}{k}\binom{u_2}{x+1-k}}{\binom{x+1+u_2}{x+1}} (w_1\, k + w_2(x+1-k)) \frac{e^{-\lambda_1}(\lambda_1)^{x+1}}{x+1!} \frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!} \\
&+ \sum_{u_1=0}^{x} \left(\sum_{u_2=x+1}^{\infty} \sum_{k=0}^{u_1} \frac{\binom{u_1}{k}\binom{u_2+1}{x+1-k}}{\binom{u_1+u_2+1}{x+1}} (w_1\, k + w_2(x+1-k)) \frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!} \frac{e^{-\lambda_2}(\lambda_2)^{u_2+1}}{(u_2+1)!} \right. \\
&\qquad \left. - \sum_{u_2=x+1}^{\infty} \sum_{k=0}^{u_1} \frac{\binom{u_1}{k}\binom{u_2}{x-k}}{\binom{u_1+u_2}{x}} (w_1\, k + w_2(x-k)) \frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!} \frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!} \right) \\
&- \sum_{u_2=0}^{x} \sum_{k=x-n}^{x} \frac{\binom{x+1}{k}\binom{u_2}{x-k}}{\binom{x+1+u_2}{x}} (w_1\, k + w_2(x-k)) \frac{e^{-\lambda_1}(\lambda_1)^{x+1}}{x+1!} \frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!} \\
&+ \sum_{u_1=x+2}^{\infty} \sum_{k=0}^{x+1} \frac{\binom{u_1}{k}\binom{x+1}{x+1-k}}{\binom{x+1+u_1}{x+1}} (w_1\, k + w_2(x+1-k)) \frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!} \frac{e^{-\lambda_2}(\lambda_2)^{x+1}}{(x+1)!} \\
&+ \sum_{u_1=x+2}^{\infty} \left(\sum_{u_2=0}^{x} \sum_{k=x-u_2}^{x} \frac{\binom{u_1}{k+1}\binom{u_2}{x-k}}{\binom{u_1+u_2}{x+1}} (w_1\,(k+1) + w_2(x-k)) \frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!} \frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{(u_2)!} \right. \\
&\qquad \left. - \sum_{u_2=0}^{x} \sum_{k=x-u_2}^{x} \frac{\binom{u_1}{k}\binom{u_2}{x-k}}{\binom{u_1+u_2}{x}} (w_1\, k + w_2(x-k)) \frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!} \frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!} \right) \\
&- \sum_{u_2=x+1}^{\infty} \sum_{k=0}^{x} \frac{\binom{x+1}{k}\binom{u_2}{x-k}}{\binom{x+1+u_2}{x}} (w_1\, k + w_2(x-k)) \frac{e^{-\lambda_1}(\lambda_1)^{x+1}}{x+1!} \frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!} \\
&+ \sum_{u_1=x+2}^{\infty} \left[\sum_{k=0}^{x} \frac{\binom{u_1}{k}\binom{x+1}{x-k}}{\binom{u_1+x+1}{x}} (w_1\, k + w_2(x-k)) \frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!} \frac{e^{-\lambda_2}(\lambda_2)^{x+1}}{(x+1)!} \right. \\
&\qquad \sum_{u_2=x+2}^{\infty} \left(\frac{\binom{u_1}{x+1}}{\binom{u_1+u_2}{x+1}} w_1\,(x+1) \frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!} \frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!} \right. \\
&\qquad\quad + \sum_{k=0}^{x} \frac{\binom{u_1}{k}\binom{u_2}{x+1-k}}{\binom{u_1+u_2}{x+1}} (w_1\, k + w_2(x+1-k)) \frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!} \frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!} \\
&\qquad\quad \left. \left. - \sum_{k=0}^{x} \frac{\binom{u_1}{k}\binom{u_2}{x-k}}{\binom{u_1+u_2}{x}} (w_1\, k + w_2(x-k)) \frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!} \frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!} \right) \right]
\end{aligned}
$$

$$- \sum_{u_1=0}^{x} \sum_{u_2=0}^{x-u_1} w_0 \frac{e^{-\lambda_1}(\lambda_1)^{u_1}}{u_1!} \frac{e^{-\lambda_2}(\lambda_2)^{u_2}}{u_2!}.$$

After cancelling out the duplicate terms of opposite signs, using probability identities, and considerable algebraic simplification, we obtain:

$$MB_F(x) = \frac{w_1\,\lambda_1 + w_2\,\lambda_2}{\lambda_1 + \lambda_2} \sum_{t=x+1}^{\infty} \frac{e^{-(\lambda_1+\lambda_2)}\,(\lambda_1+\lambda_2)^t}{t!} - w_0 \sum_{t=0}^{x} \frac{e^{-(\lambda_1+\lambda_2)}\,(\lambda_1+\lambda_2)^t}{t!}.$$

For the more general case of $|J|$ demand classes, the difference between the marginal benefits at x and $x-1$ units is calculated as:

$$\begin{aligned}
MB_F(x) - MB_F(x-1) =& \frac{\sum_{j=1}^{|J|} w_j\,\lambda_j}{\sum_{j=1}^{|J|} \lambda_j} \sum_{t=x+1}^{\infty} \frac{e^{-\sum_{j=1}^{|J|}\lambda_j}\left(\sum_{j=1}^{|J|}\lambda_j\right)^t}{t!} - w_0 \sum_{t=0}^{x} \frac{e^{-\sum_{j=1}^{|J|}\lambda_j}\left(\sum_{j=1}^{|J|}\lambda_j\right)^t}{t!} \\
& - \frac{\sum_{j=1}^{|J|} w_j\,\lambda_j}{\sum_{j=1}^{|J|} \lambda_j} \sum_{t=x}^{\infty} \frac{e^{-\sum_{j=1}^{|J|}\lambda_j}\left(\sum_{j=1}^{|J|}\lambda_j\right)^t}{t!} - w_0 \sum_{t=0}^{x-1} \frac{e^{-\sum_{j=1}^{|J|}\lambda_j}\left(\sum_{j=1}^{|J|}\lambda_j\right)^t}{t!} \\
=& - \frac{\sum_{j=1}^{|J|} w_j\,\lambda_j}{\sum_{j=1}^{|J|} \lambda_j} \frac{e^{-\sum_{j=1}^{|J|}\lambda_j}\left(\sum_{j=1}^{|J|}\lambda_j\right)^x}{x!} - w_0 \frac{e^{-\sum_{j=1}^{|J|}\lambda_j}\left(\sum_{j=1}^{|J|}\lambda_j\right)^x}{x!} \\
=& - \left(\frac{\sum_{j=1}^{|J|} w_j\,\lambda_j}{\sum_{j=1}^{|J|} \lambda_j} + w_0 \right) P\left(\sum_{j=1}^{|J|} D_j = x \right),
\end{aligned}$$

which is nonpositive due to the nonnegativity of w_j, $j \geq 0$. □

Proof of Proposition 3. For the continuous case, by Equation (38), we have:

$$\sum_{j=1}^{|J|} (w_j - w_{j+1})\, P\left(\sum_{l=1}^{j} D_l > x^* \right) + w_{|J|}\, P\left(\sum_{l=1}^{|J|} D_l > x^* \right) = w_0\, P\left(\sum_{l=1}^{|J|} D_l \leq x^* \right).$$

This gives:

$$w_{|J|}\, P\left(\sum_{l=1}^{|J|} D_l > x^* \right) \leq w_0\, P\left(\sum_{l=1}^{|J|} D_l \leq x^* \right),$$

and hence

$$P\left(\sum_{l=1}^{|J|} D_l \leq x^* \right) \geq \frac{w_{|J|}}{w_{|J|} + w_0}.$$

Therefore, $x^* \geq G^{-1}\left(\frac{w_{|J|}}{w_{|J|}+w_0}\right)$. We also have, using Equation (38):

$$w_1\, P\left(\sum_{l=1}^{|J|} D_l > x^*\right) \geq w_0\, P\left(\sum_{l=1}^{|J|} D_l \leq x^*\right),$$

and hence:

$$P\left(\sum_{l=1}^{|J|} D_l \leq x^*\right) \leq \frac{w_1}{w_1+w_0},$$

implying that $x^* \leq G^{-1}\left(\frac{w_1}{w_1+w_0}\right)$. For the discrete case, we use:

$$\sum_{j=1}^{|J|}(w_j - w_{j+1})\, P\left(\sum_{l=1}^{j} D_l > x^*\right) + w_{|J|}\, P\left(\sum_{l=1}^{|J|} D_l > x^*\right) \leq w_0\, P\left(\sum_{l=1}^{|J|} D_l \leq x^*\right)$$

to establish $x^* \geq G^{-1}\left(\frac{w_{|J|}}{w_{|J|}+w_0}\right)$. We also have:

$$\sum_{j=1}^{|J|}(w_j - w_{j+1})\, P\left(\sum_{l=1}^{j} D_l > x^* - 1\right) + w_{|J|}\, P\left(\sum_{l=1}^{|J|} D_l > x^* - 1\right) > w_0\, P\left(\sum_{l=1}^{|J|} D_l \geq x^* - 1\right),$$

thus $x^* - 1 < G^{-1}\left(\frac{w_1}{w_1+w_0}\right)$. Since the inequality is strict, by the definition of G^{-1}, this gives $x^* \leq G^{-1}\left(\frac{w_1}{w_1+w_0}\right)$. □

Proof of Conjecture 1. The left-hand side of the inequality in the lemma yields:

$$\frac{MB_F(x)\, MB_P(x-1) - MB_P(x)\, MB_F(x-1)}{MB_P(x)\, MB_P(x-1)}.$$

By the fact that the marginal benefits are nonincreasing in x, the denominator is always nonnegative. Hence, it suffices to show the nonnegativity of the numerator to prove the lemma.

By plugging in the previously derived expressions, the numerator expands to:

$$\left(\frac{w_1\lambda_1 + w_2\lambda_2}{\lambda_1+\lambda_2}\, P(D_1+D_2 \geq x+1) - w_0\, P(D_1+D_2 \leq x)\right)\left(w_1\, P(D_1 \geq x) + w_2\left(P(D_1 \leq x-1)\right.\right.$$
$$\left.\left. - P(D_1+D_2 \leq x-1)\right) - w_0\, P(D_1+D_2 \leq x-1)\right) - \left(\frac{w_1\lambda_1 + w_2\lambda_2}{\lambda_1+\lambda_2}\, P(D_1+D_2 \geq x)\right.$$
$$\left. - w_0\, P(D_1+D_2 \leq x-1)\right)\left(w_1\, P(D_1 \geq x+1) + w_2\left(P(D_1 \leq x) - P(D_1+D_2 \leq x)\right)\right.$$

$$- w_0\, P(D_1 + D_2 \le x)\Big).$$

Using the identities:

$$P(D_1 + D_2 \ge x + 1) = P(D_1 + D_2 \ge x) - P(D_1 + D_2 = x),$$

$$P(D_1 + D_2 \le x) = P(D_1 + D_2 \le x - 1) + P(D_1 + D_2 = x),$$

$$P(D_1 \ge x + 1) = P(D_1 \ge x) - P(D_1 = x),$$

$$P(D_1 \le x) = P(D_1 \le x - 1) + P(D_1 = x), \text{ and}$$

$$(P(D_1 \le x - 1) - P(D_1 + D_2 \le x - 1)) - (P(D_1 \le x) - P(D_1 + D_2 \le x)) = P(D_1 = x)$$

$$- P(D_1 + D_2 = x),$$

the numerator can be rewritten as:

$$\begin{aligned}
&\left(\frac{w_1\lambda_1 + w_2\lambda_2}{\lambda_1 + \lambda_2}\, P(D_1 + D_2 \ge x) - w_0\, P(D_1 + D_2 \le x - 1)\right)\left((w_1 - w_2)\, P(D_1 = x)\right. \\
&\left. + (w_2 + w_0)\, P(D_1 + D_2 = x)\right) - \left(\left(\frac{w_1\,\lambda_1 \,+\, w_2\,\lambda_2}{\lambda_1 \,+\, \lambda_2} \,+ w_0\right) P(D_1 + D_2 = x)\right) \\
&\left(w_1\, P(D_1 \ge x) + w_2\Big(P(D_1 \le x - 1) - P(D_1 + D_2 \le x - 1)\Big) - w_0\, P(D_1 + D_2 \le x - 1)\Big)\right. \\
&= \frac{w_1\lambda_1 + w_2\lambda_2}{\lambda_1 + \lambda_2}\,(w_1 - w_2)\, P(D_1 + D_2 \ge x)\, P(D_1 = x) \\
&\quad + \frac{w_1\lambda_1 + w_2\lambda_2}{\lambda_1 + \lambda_2}\,(w_2 + w_0)\, P(D_1 + D_2 \ge x)\, P(D_1 + D_2 = x) \\
&\quad - w_0(w_1 - w_2)\, P(D_1 + D_2 \le x - 1)\, P(D_1 = x) \;-\; w_0\, w_2\, P(D_1 + D_2 \le x - 1)\, P(D_1 + D_2 = x) \\
&\quad - \left(\frac{w_1\lambda_1 + w_2\lambda_2}{\lambda_1 + \lambda_2} \,+ w_0\right) w_1\, P(D_1 \ge x)\, P(D_1 + D_2 = x) \\
&\quad - \left(\frac{w_1\lambda_1 + w_2\lambda_2}{\lambda_1 + \lambda_2} \,+ w_0\right) w_2\, \Big(P(D_1 \le x - 1) - P(D_1 + D_2 \le x - 1)\Big)\, P(D_1 + D_2 = x) \\
&\quad - \frac{w_1\lambda_1 + w_2\lambda_2}{\lambda_1 + \lambda_2}\, w_0\, P(D_1 + D_2 \le x - 1)\, P(D_1 + D_2 = x) \\
&= (w_1 - w_2)\left[\left(\frac{w_1\lambda_1 + w_2\lambda_2}{\lambda_1 + \lambda_2} \,+ w_0\right)\Big(P(D_1 + D_2 \ge x)\, P(D_1 = x) - P(D_1 \ge x)\, P(D_1 + D_2 = x)\Big)\right. \\
&\quad \left. + \; w_0\left(\frac{\lambda_1}{\lambda_1 + \lambda_2}\, P(D_1 + D_2 = x) \;-\; P(D_1 = x)\right)\right].
\end{aligned}$$

Since $w_1 \ge w_2$, it suffices to show that:

$$\begin{aligned}
&\left[\left(\frac{w_1\lambda_1 + w_2\lambda_2}{\lambda_1 + \lambda_2} \,+ w_0\right)\Big(P(D_1 + D_2 \ge x)\, P(D_1 = x) - P(D_1 \ge x)\, P(D_1 + D_2 = x)\Big)\right. \\
&\quad \left. + \; w_0\left(\frac{\lambda_1}{\lambda_1 + \lambda_2}\, P(D_1 + D_2 = x) \;-\; P(D_1 = x)\right)\right] \ge 0. \qquad (124)
\end{aligned}$$

The next step in the proof of the conjecture is to show that given $\lambda_1 \geq \frac{w_0}{w_2}$ and $\lambda_1 \geq \frac{w_0}{w_2}$, the inequality in (124) always holds. $\square$

Proof of Theorem 8. First, we show that the inequality in (41) gives a valid bound. By Theorem 7, we have $(x^P)^* \leq (x^F)^*$. Furthermore, by Conjecture 1, the ratio of FCFS marginal benefit to that of prioritization is nondecreasing with respect to x. Combining these two results, it suffices to show that if (41) holds for the first unit allocated, it is valid for any $(x^P)^*$, $(x^F)^*$ pair. Therefore, we need to show that:

$$\frac{MB_F(0)}{MB_P(0)} = \frac{E_F(z\,|\,1) - E_F(z\,|\,0)}{E_P(z\,|\,1) - E_P(z\,|\,0)} \geq \frac{K+L}{L(K+1)}.$$

Using the aforementioned identities and ratios,

$$\begin{aligned}
MB_P(0) &= w_1\,P(D_1 \geq 1) + w_2\,P(D_1 = 0, D_2 \geq 1) - w_0\,P(D_1 + D_2 = 0) \\
&= L\,w_2\,(1 - e^{-\lambda_1}) + w_2\,e^{-\lambda_1}\,(1 - e^{-K\,\lambda_1}) - w_0\,e^{-(K+1)\,\lambda_1} \\
&= L\,w_2 - \frac{L\,w_2}{e^{\lambda_1}} + \frac{w_2}{e^{\lambda_1}} - \frac{w_2}{e^{(K+1)\,\lambda_1}} - \frac{w_0}{e^{(K+1)\,\lambda_1}} \\
MB_F(0) &= \frac{w_1\,\lambda_1 + w_2\,\lambda_2}{\lambda_1 + \lambda_2}\,P(D_1 + D_2 \geq 1) - w_0\,P(D_1 + D_2 = 0) \\
&= \frac{K+L}{K+1}\,w_2\,(1 - e^{-(K+1)\,\lambda_1}) - w_0\,e^{-(K+1)\,\lambda_1} \\
&= \frac{K+L}{K+1}\,w_2 - \frac{K+L}{K+1}\,\frac{w_2}{e^{(K+1)\,\lambda_1}} - \frac{w_0}{e^{(K+1)\,\lambda_1}}.
\end{aligned}$$

Based on these:

$$\begin{aligned}
\frac{L\,(K+1)}{K+L}\,MB_F(0) &= L\,w_2 - \frac{L\,w_2}{e^{(K+1)\,\lambda_1}} - \frac{L\,(K+1)}{K+L}\,\frac{w_0}{e^{(K+1)\,\lambda_1}} \\
&= L\,w_2 - \frac{L\,w_2}{e^{(K+1)\,\lambda_1}} - \frac{L\,K}{K+L}\,\frac{w_0}{e^{(K+1)\,\lambda_1}}\,\frac{L}{K+L}\,\frac{w_0}{e^{(K+1)\,\lambda_1}},
\end{aligned}$$

which gives:

$$\frac{L\,(K+1)}{K+L}\,MB_F(0) - MB_P(0) = \frac{(K+L)\,(L-1)\,(e^{K\,\lambda_1} - 1)\,w_2 + K\,(1-L)\,w_0}{(K+L)\,e^{(K+1)\,\lambda_1}}.$$

The denominator is always nonnegative. Therefore, to show the difference between $\frac{L\,(K+1)}{K+L}\,MB_F(0)$ and $MB_P(0)$ is nonnegative, it suffices to show the nonnegativity of

the numerator. For this, we need:

$$(K+L)\,(L-1)\,(e^{K\,\lambda_1}-1)\,w_2\;-\;K\,(L-1)\;\geq\;0 \Leftrightarrow\; w_0\;\leq\;\frac{K+L}{K}\,(e^{K\,\lambda_1}-1)\,w_2.$$

To show the validity of this inequality, we start with the following fact:

$$ln\left(\frac{T\,K}{K+L}\;+1\right)^{\frac{1}{K}}\leq T.$$

This is due to the function being nonincreasing with $L \geq 1$ and

$$\lim_{K\to 0}\left(\frac{T\,K}{K+1}\;+1\right)^{\frac{1}{K}}=e^{T}.$$

Using this fact, we obtain:

$$\begin{aligned}
&\lambda_1 \geq T \geq ln\left(\frac{T\,K}{K+L}\;+1\right)^{\frac{1}{K}}=\frac{1}{K}\,ln\left(\frac{T\,K}{K+L}\;+1\right)\\
\Leftrightarrow\; &e^{K\lambda_1}\geq\frac{T\,K}{K+L}+1\;\Leftrightarrow\;\left(1+\frac{L}{K}\right)(e^{K\lambda_1}-1)\geq T\\
\Leftrightarrow\; &T\,w_2\leq\left(1+\frac{L}{K}\right)(e^{K\lambda_1}-1)\,w_2\\
\Leftrightarrow\; &w_0\leq\left(1+\frac{L}{K}\right)(e^{K\lambda_1}-1)\,w_2,
\end{aligned}$$

which is the desired inequality.

To see that the bound is asymptotically tight, consider the marginal benefits of allocating one more unit under FCFS and prioritization usage when x units are already allocated in each case:

$$\begin{aligned}
MB_F(x) &= \frac{w_1\,\lambda_1\;+\;w_2\,\lambda_2}{\lambda_1\;+\;\lambda_2}\,P(D_1+D_2>x)\;-\;w_0\,P(D_1+D_2\leq x)\\
MB_P(x) &= w_1\,P(D_1\geq x+1)\;+w_2\left(P(D_1+D_2\geq x+1)-P(D_1\geq x+1)\right)\\
&\quad -\,w_0\,P(D_1+D_2\leq x).
\end{aligned}$$

Now, given capacity C, set λ_1 sufficiently large so that:

$$P(D_1\leq C)=\sum_{t=0}^{C}\frac{e^{-\lambda_1}\,(\lambda_1)^t}{t!}\to 0. \tag{125}$$

Then, obviously,

$$P(D_1 + D_2 \leq C) = \sum_{t=0}^{C} \frac{e^{-(\lambda_1+\lambda_2)}\,(\lambda_1+\lambda_2)^t}{t!} \to 0,$$

which gives:

$$\pi_P = \sum_{i=0}^{C-1} MB_P(x) \to C\,w_1, \quad \text{and} \tag{126}$$

$$\pi_F = \sum_{i=0}^{C-1} MB_F(x) \to C\,\frac{w_1\,\lambda_1\,+\,w_2\,\lambda_2}{\lambda_1\,+\,\lambda_2}. \tag{127}$$

Substituting $\lambda_2 = K\lambda_1$ and $w_1 = Lw_2$, we obtain:

$$\frac{\pi_F}{\pi_P} \to \frac{C\,\frac{(K+L)\,w_2\,\lambda_1}{\lambda_1+K\,\lambda_1}}{C\,L\,w_2} = \frac{\frac{(K+L)\,\lambda_1\,w_2}{(K+1)\,\lambda_1}}{L\,w_2} = \frac{K+L}{L(K+1)}.$$

□

APPENDIX C

SPECIFICATIONS FOR SPECIALIZED NUTRITIOUS FOODS

Programme	TREATMENT of Moderate Acute Malnutrition (MAM)				
Generic Product Term	Lipid-based Nutrient Supplement (LNS) Large Quantity (92-100 g)			Fortified Blended Foods (FBF) (200-250g)	
Current Nutrition Products	 Plumpy'sup™[1] (Peanut-based)	 eeZeeRUSF™ (Peanut-based)	 AchaMum™ (Chickpea based)	 Supercereal Plus	 Supercereal[2]
Target Group	Children 6-59 months	Children 6-59 months	Children 6-59 months	Children 6-59 months	Pregnant and Lactating Women (PLW) Malnourished individuals on ART/DOTS
Key Ingredients	Peanuts, sugar, whey, vegetable oil, milk, soy protein, cocoa, V&M	Peanut, sugar, milk solids, vegetable oil, V&M	Chickpeas, vegetable oil, milk powder, sugar, V&M, soya lecithin	Corn or wheat, soya, milk powder, sugar, oil, V&M	Corn or wheat, soya, V&M
Daily ration	92g sachet	92g sachet	100g sachet	200g (includes provision for sharing)	200-250g (includes provision for sharing)
Nutrient profile	500 kcal, 13g protein (10%), 31g fat (56%). Contains EFA, meets RNI and PDCAAS.	500 kcal, 13g protein (10%), 31g fat (56%). Contains EFA, meets RNI and PDCAAS.	520 kcal, 13g protein (10%), 29g fat (50%). Contains EFA, meets RNI and PDCAAS.	788 kcal, 33g protein (17%), 20g fat (23%). Contains EFA, meets RNI and PDCAAS.	750-950 kcal, 31-38g protein (16%), 16-20g fat (19%). Meets RNI and PDCAAS.
Duration of Intervention[3]	60-90 days	60-90 days	60-90 days	60-90 days	PLW: 180 days, ART & DOTS: 180 days (estimated)
Shelf life	24 months	24 months	6 months	12 months	12 months
Packaging details	Carton: 14.7kg (gross) and 13.8kg (net) has 150 sachets	Carton: 14.9kg (gross) and 13.8kg (net) has 150 sachets	Carton: 10.5kg (net) has 105 sachets	Primary: 1.5kg (net) bag; Secondary: 15kg (net) carton has 10 bags; or 18kg sack has 12 bags	25kg (net) bags
Product cost (USD) *Oct11	$ 3,196/MT (0.29/daily ration)	$ 2,771/MT (0.25/daily ration)	$ 3,039/MT (0.30/daily ration)	Supercereal Plus (corn or wheat based): $ 1,188/MT (0.24/daily ration)	Supercereal (corn): $ 570/MT (0.11-0.14/daily ration)[4]. Supercereal (wheat): $ 630/MT (0.13-0.16/daily ration)[4]

Figure 46: Specifications for specialized nutritious foods used to treat moderate acute malnutrition

Programme	PREVENTION of Stunting: All products listed below can be used for the prevention of stunting					
	PREVENTION of Acute Malnutrition: Four products specifically for the prevention of acute malnutrition				TREATMENT and PREVENTION of Micronutrient Deficiencies[1]:	
Generic Product Term	Lipid-based Nutrient Supplement (LNS) Medium Quantity (20-50g)		Fortified Blended Food (FBF) (100-200g)		LNS Small Quantity (≤ 20g)	Micronutrient Powders (1g)
Current Nutrition Products	 Plumpy'doz™ (Peanut-based)	 WawaMum™ (Chickpea-based)	 Supercereal Plus	 Supercereal[2]	 Nutributter™ (Peanut-based)	 Micronutrient Powders (MNP)
Target Group	Children 6-23 months	Children 6-23 months	Children 6-23 months	Pregnant and Lactating Women	Children 6-23 months	Children 6-59 months
Key Ingredients	Vegetable fat, peanut, sugar, milk powder, whey, sugar, V&M, cocoa	Chickpeas, vegetable oil, milk powder, sugar, V&M	Corn or wheat, soya, milk powder, sugar, oil, V&M	Corn or wheat, soya, V&M	Peanuts, vegetable fat, sugar, skim milk powder, whey, V&M	Vitamins and minerals (V&M)
Daily ration	47g portion (1/7 portion of a pot)	50g sachet	100-200g (200g includes provision for sharing)	100-200g (200g includes provision for sharing).	20g sachet	1g sachet every second day
Nutrient profile	247 kcal, 5.9g protein (10%), 16g fat (58%). Contains EFA, meets RNI and PDCAAS.	260 kcal, 6.5g protein (10%), 14.5g fat (50%). Contains EFA, meets RNI and PDCAAS.	394-788 kcal, 16-33g protein (17%), 10-20g fat (23%). Contains EFA, meets RNI and PDCAAS.	380-750 kcal, 15-31g protein (16%), 8-16g fat (19%). Meets RNI and PDCAAS.	108 kcal, 2.6g protein (10%), 7g fat (58%). Contains EFA, meets RNI and PDCAAS.	Meets RNI. (No energy, fat or protein content)
Duration of Intervention[2]	90-180 days	90-180 days	90-180 days	90-180 days	180-545 days	180-545 days
Shelf life	24 months	6 months	12 months	12 months	18 months	24 months
Packaging details	Primary packaging: 325g pots. Carton: 12.7kg (gross) and 11.7kg (net) has 36 pots	Carton: 10.5kg (net) has 210 sachets	Primary: 1.5kg (net) bag; Secondary: 15kg (net) carton has 10 bags; or 18kg sack has 12 bags	25kg (net) bags	Carton: 11.95kg (gross) and 10.92kg (net) has 546 sachets	Carton: 14kg (gross) has 240 boxes; 30 sachet in each box. *Packaging varies with supplier
Product Cost (USD) *Oct11	$ 3,703/MT (0.17/daily ration)	$ 3,389/MT (0.17/daily ration)	Supercereal Plus (corn or wheat): $ 1,188/MT (0.12-0.24/daily ration)	Supercereal (corn): $ 570/MT (0.06-0.11/daily ration)[4] Supercereal (wheat): $ 630/MT (0.06-0.13/daily ration)[4]	$ 4,077/MT (0.08/daily ration)	$ 25,178/MT (0.025/daily ration)

Figure 47: Specifications for specialized nutritious foods used to prevent stunting

REFERENCES

[1] Akkerman, R., Farahani, P., and Grunow, M., "Quality, safety and sustainability in food distribution: A review of quantitative operations management approaches and challenges," *OR Spectrum*, vol. 32, no. 4, pp. 863–904, 2010.

[2] Aksu, D. T. and Özdamar, L., "A mathematical model for post-disaster road restoration: Enabling accessibility and evacuation," *Transportation Research Part E: Logistics and Transportation Review*, vol. 61, no. 1, pp. 56–67, 2014.

[3] Altay, N. and Green, W. G., "OR/MS research in disaster operations management," *European Journal of Operational Research*, vol. 175, no. 1, pp. 475–493, 2006.

[4] Alvarenga, R., Ergun, Ö., Li, J., Mata, F., Shekhani, N., Slaton, D., Stone, J., Vasudevan, A., and Yang, E., "World Food Programme East African corridor optimization," Working Paper, H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA, 2010.

[5] Anagnostopoulos, A., Bent, R., Upfal, E., and Hentenryck, P. V., "A simple and deterministic competitive algorithm for online facility location," *Information and Computation*, vol. 194, no. 2, pp. 175–202, 2004.

[6] Angelelli, E., Speranza, M. G., and Savelsbergh, M. W. P., "Competitive analysis for dynamic multiperiod uncapacitated routing problems," *Networks*, vol. 49, no. 4, pp. 308–317, 2007.

[7] Anily, S. and Tzur, M., "Shipping multiple items by capacitated vehicles: An optimal dynamic programming approach," *Transportation Science*, vol. 39, no. 2, pp. 233–248, 2005.

[8] Apte, A., "Humanitarian logistics: A new field of research and action," *Foundations and Trends in Technology, Information and Operations Management*, vol. 3, no. 1, pp. 1–100, 2009.

[9] Argon, N. T. and Andradóttir, S., "Partial pooling in tandem lines with cooperation and blocking," *Queueing Systems*, vol. 52, no. 1, pp. 5–30, 2006.

[10] Ascheuer, N., Krumke, S., and Rambau, J., "Online dial-a-ride problems: Minimizing the completion time," in *STACS 2000* (Reichel, H. and Tison, S., eds.), vol. 1770 of *Lecture Notes in Computer Science*, pp. 639–650, Springer Berlin Heidelberg, 2000.

[11] Ausiello, G., Bonifaci, V., and Laura, L., "The online prize-collecting traveling salesman problem," *Information Processing Letters*, vol. 107, no. 6, pp. 199–204, 2008.

[12] Axsäter, S., "Supply chain operations: Serial and distribution inventory systems," in *Handbooks in Operations Research and Management Science* (de Kok, A. G. and Graves, S. C., eds.), vol. 11, pp. 525–560, Elsevier B. V., 2003.

[13] Axsäter, S., Marklund, J., and Silver, E. A., "Heuristic methods for centralized control of one-warehouse, n-retailer inventory systems," *Manufacturing & Service Operations Management*, vol. 4, no. 1, pp. 75–97, 2002.

[14] Balakrishnan, N., Sridharan, S. V., and Patterson, J. W., "Rationing capacity between two product classes," *Decision Sciences*, vol. 27, no. 2, pp. 185–214, 1996.

[15] Balcik, B., Beamon, B. M., and Smilowitz, K., "Last mile distribution in humanitarian relief," *Journal of Intelligent Transportation Systems*, vol. 12, no. 2, pp. 51–63, 2008.

[16] Bassok, Y., Anupindi, R., and Akella, R., "Single-period multiproduct inventory models with substitution," *Operations Research*, vol. 47, no. 4, pp. 632–642, 1999.

[17] Beliën, J. and Forcé, H., "Supply chain management of blood products: A literature review," *European Journal of Operational Research*, vol. 217, no. 1, pp. 1–16, 2012.

[18] Benjaafar, S., "Performance bounds for the effectiveness of pooling in multiprocessing systems," *European Journal of Operational Research*, vol. 87, no. 2, pp. 375–388, 1995.

[19] Benjaafar, S., ElHafsi, M., and De Véricourt, F., "Demand allocation in multiple-product, multiple-facility, make-to-stock systems," *Management Science*, vol. 50, no. 10, pp. 1431–1448, 2004.

[20] Benjaafar, S., Li, Y., Xu, D., and Elhedhli, S., "Demand allocation in systems with multiple inventory locations and multiple demand sources," *Manufacturing & Service Operations Management*, vol. 10, no. 1, pp. 43–60, 2008.

[21] Bent, R. and Hentenryck, P. V., "Waiting and relocation strategies in online stochastic vehicle routing," in *Proceedings of the 20th International Joint Conference on Artifical Intelligence*, IJCAI'07, pp. 1816–1821, 2007.

[22] Bent, R. W. and Hentenryck, P. V., "Online stochastic and robust optimization," in *Advances in Computer Science - ASIAN 2004. Higher-Level*

Decision Making (Maher, M. J., ed.), vol. 3321 of *Lecture Notes in Computer Science*, pp. 286–300, Springer Berlin Heidelberg, 2004.

[23] Bent, R. W. and Hentenryck, P. V., "Scenario-based planning for partially dynamic vehicle routing with stochastic customers," *Operations Research*, vol. 52, no. 6, pp. 977–987, 2004.

[24] Bernstein, F. and Federgruen, A., "Pricing and replenishment strategies in a distribution system with competing retailers," *Operations Research*, vol. 51, no. 3, pp. 409–426, 2003.

[25] Bernstein, F. and Federgruen, A., "Decentralized supply chains with competing retailers under demand uncertainty," *Management Science*, vol. 51, no. 1, pp. 409–426, 2005.

[26] Brodheim, E., Derman, C., and Prastacos, G., "On the evaluation of a class of inventory policies for perishable products such as blood," *Management Science*, vol. 21, no. 11, pp. 1320–1325, 1975.

[27] Broekmeulen, R. A. C. M. and van Donselaar, K. H., "A heuristic to manage perishable inventory with batch ordering, positive lead-times, and time-varying demand," *Computers & Operations Research*, vol. 36, no. 11, pp. 3013–3018, 2009.

[28] Brown, C., Milke, M., and Seville, E., "Disaster waste management: A review article," *International Journal of Logistics: Research and Applications*, vol. 31, no. 6, pp. 1085–1098, 2011.

[29] Cao, W., Çelik, M., Ergun, Ö., Swann, J., and Viljoen, N., "A multi-level location-allocation problem in human milk banking network expansion," working paper, H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA, 2012.

[30] Carden, R. and DelliFraine, J. L., "An examination of blood center structure and hospital customer satisfaction: What can centralized and decentralized blood centers learn from each other?," *Health Marketing Quarterly*, vol. 22, no. 3, pp. 21–42, 2005.

[31] Caunhye, A. M., Nie, X., and Pokharel, S., "Optimization models in emergency logistics: A literature review," *Socio-Economic Planning Sciences*, vol. 46, no. 1, pp. 4–13, 2012.

[32] Çavdaroğlu, B., Hammel, E., Mitchell, J. E., Sharkey, T. C., and Wallace, W. A., "Integrating restoration and scheduling decisions for disrupted interdependent infrastructure systems," *Annals of Operations Research*, vol. 203, no. 1, pp. 279–294, 2011.

[33] Çelik, M., Ergun, Ö., Johnson, B., Keskinocak, P., Lorca, Á., Pekgün, P., and Swann, J., "Humanitarian logistics," in *Tutorials in Operations Research* (Mirchandani, P., ed.), vol. 9, pp. 18–49, Hanover, MD: INFORMS, 2012.

[34] Çetinkaya, S., Tekin, E., and Lee, C.-Y., "A stochastic model for joint inventory and outbound shipment decisions," *IIE Transactions*, vol. 40, no. 3, pp. 324–340, 2008.

[35] Çetinkaya, S. and Bookbinder, J. H., "Stochastic models for the dispatch of consolidated shipments," *Transportation Research Part B: Methodological*, vol. 37, no. 8, pp. 747–768, 2003.

[36] Çetinkaya, S. and Lee, C.-Y., "Optimal outbound dispatch policies: Modeling inventory and cargo capacity," *Naval Research Logistics*, vol. 49, no. 6, pp. 531–556, 2002.

[37] Center for Research on the Epidemiology of Disasters, "EM-DAT: The international disaster database," 2013. Available online at http://www.emdat.be, accessed February 2013.

[38] Centers for Disease Control and Prevention (CDC), "2009 H1N1 Vaccination Recommendations," 2009. Available online at http://www.cdc.gov/h1n1flu/vaccination/acip.htm, accessed September 2013.

[39] Chen, F. and Zheng, Y. S., "Lower bounds for multi-echelon stochastic inventory systems," *Management Science*, vol. 40, no. 11, pp. 1426–1443, 1994.

[40] Chen, S., Chen, Y., Long, Y., and Shen, Y., "Optimized road network emergency restoration scheduling for relief distribution," in *International Conference of E-Business and E-Government (ICEE)*, (Shanghai, China), pp. 1–4, 2011.

[41] Clark, A. J. and Scarf, H., "Optimal policies for a multi-echelon inventory problem," *Management Science*, vol. 6, no. 4, pp. 475–490, 1960.

[42] Coelho, L. C. and Laporte, G., "Optimal joint replenishment, delivery and inventory management policies for perishable products," working paper, CIRRELT, Montreal, Canada, 2013.

[43] Cohen, M. A., "Analysis of single critical number ordering policies for perishable inventories," *Operations Research*, vol. 24, no. 4, pp. 726–741, 1976.

[44] Cohen, M. A., Pierskalla, W. P., and Yen, H.-C., "Analysis of ordering and allocation policies for multi-echelon, age-differentiated inventory systems," *TIMS Studies in the Management Sciences*, vol. 16, no. 1, pp. 353–378, 1981.

[45] Congressional Research Service, "Managing disaster debris: Overview of regulatory requirements, agency roles, and selected challenges," 2011. Available online at http://www.fas.org/sgp/crs/misc/RL34576.pdf, accessed January 2013.

[46] Consuelos Salas, L., Robles Cárdenas, M., and Zhang, M., "Inventory policies for humanitarian aid during hurricanes," *Socio-Economic Planning Sciences*, vol. 46, no. 4, pp. 272–280, 2012.

[47] Daganzo, C. F., "The distance traveled to visit n points with a maximum of c stops per vehicle: An analytic model and an application," *Transportation Science*, vol. 18, no. 4, pp. 331–350, 1984.

[48] Daganzo, C. F., "Modeling distribution problems with time windows: Part I," *Transportation Science*, vol. 21, no. 3, pp. 171–179, 1987.

[49] Daganzo, C. F., "Modeling distribution problems with time windows. Part II: Two customer types," *Transportation Science*, vol. 21, no. 3, pp. 180–187, 1987.

[50] Daganzo, C. F. and Erera, A. L., "On planning and design of logistics systems for uncertain environments," in *Lecture Notes in Economics and Mathematical Systems: New Trends in Distribution Logistics* (Speranza, M. G. and Stahly, P., eds.), vol. 480, pp. 3–21, Berlin: Springer-Verlag, 1999.

[51] Daskin, M. S., Coullard, C. R., and Shen, Z.-J. M., "An inventory-location model: Formulation, solution algorithm and computational results," *Annals of Operations Research*, vol. 110, no. 1–4, pp. 83–106, 2002.

[52] De Angelis, V., Mecoli, M., Nikoi, C., and Storchi, G., "Multiperiod integrated routing and scheduling of World Food Programme cargo planes in Angola," *Computers & Operations Research*, vol. 34, no. 6, pp. 1601–1615, 2007.

[53] De Véricourt, F., Karaesmen, F., and Dallery, Y., "Optimal stock allocation for a capacitated supply system," *Management Science*, vol. 48, no. 11, pp. 1486–1501, 2002.

[54] Del Castillo, J. M., "A heuristic for the traveling salesman problem based on a continuous approximation," *Transportation Research Part B: Methodological*, vol. 33, no. 2, pp. 123–152, 1998.

[55] Demirağ, Ö. Ç. and Swann, J. L., "Capacity allocation to sales agents in a decentralized logistics network," *Naval Research Logistics*, vol. 54, no. 7, pp. 796–810, 2007.

[56] Deniz, B., Karaesmen, I., and Scheller-Wolf, A., "Managing perishables with substitution: Inventory issuance and replenishment heuristics," *Manufacturing & Service Operations Management*, vol. 12, no. 2, pp. 319–329, 2010.

[57] Diks, E. B. and de Kok, A. G., "Optimal control of a divergent multi-echelon inventory system," *European Journal of Operational Research*, vol. 111, no. 1, pp. 75–97, 1998.

[58] DIVA-GIS, "Chad road network shapefiles," 2014. Available online at http://biogeo.ucdavis.edu/data/ diva/rds/TCD_rds.zip, accessed: May 2014.

[59] Doğru, M. K., Kok, A. G., and van Houtum, G. J., "A numerical study on the effect of the balance assumption in one-warehouse multi-retailer inventory systems," *Flexible Services and Manufacturing Journal*, vol. 21, no. 3–4, pp. 114–147, 2009.

[60] Drezner, T., Drezner, Z., and Guyse, J., "Equitable service by a facility: Minimizing the Gini coefficient," *Computers & Operations Research*, vol. 36, no. 12, pp. 3240–3246, 2009. New developments on hub location.

[61] Duran, S., Liu, T., Simchi-Levi, D., and Swann, J. L., "Optimal production and inventory policies of priority and price-differentiated customers," *IIE Transactions*, vol. 39, no. 1, pp. 845–861, 2007.

[62] Eiselt, H. A. and Laporte, G., "Objectives in location problems," working paper, CIRRELT, Montreal, Canada, 1995.

[63] Ekici, S., McEntire, D. A., and Afedzie, R., "Transforming debris management: Considering new essentials," *Disaster Prevention and Management*, vol. 18, no. 5, pp. 511–522, 2009.

[64] Eppen, G. and Schrage, L., "Centralized ordering policies in a multi-warehouse system with lead times and random demand," *TIMS Studies in the Management Sciences*, vol. 16, no. 1, pp. 51–67, 1981.

[65] Eppen, G. D., "Effects of centralization on expected costs in a multi-location newsboy problem," *Management Science*, vol. 25, no. 5, pp. 498–501, 1979.

[66] Epstein, R., Henríquez, L., Catalán, J., Weintraub, G. Y., and Martínez, C., "A combinational auction improves school meals in Chile," *Interfaces*, vol. 32, no. 6, pp. 1–14, 2002.

[67] Federal Emergency Management Agency (FEMA), "Public assistance debris management guide," 2007. Available online at http://www.fema.gov/pdf/government/grant/pa/demagde.pdf, accessed January 2013.

[68] Federal Emergency Management Agency (FEMA), "Hazus: FEMA's methodology for estimating potential losses from disasters," 2012. Available online at http://www.fema.gov/hazus, accessed January 2013.

[69] Federgruen, A., Prastacos, G., and Zipkin, P. H., "An allocation and distribution model for perishable products," *Operations Research*, vol. 34, no. 1, pp. 75–82, 1986.

[70] Federgruen, A. and Zipkin, P., "Computational issues in an infinite-horizon, multiechelon inventory model," *Operations Research*, vol. 32, no. 4, pp. 818–836, 1984.

[71] Fetter, G. and Falasca, M., "Establishing the need for decision support in disaster debris disposal," in *Proceedings of the 8th International ISCRAM Conference*, (Lisbon, Portugal), pp. 1–4, ISCRAM, 2011.

[72] Feuerstein, E. and Stougie, L., "On-line single-server dial-a-ride problems," *Theoretical Computer Science*, vol. 268, no. 1, pp. 91–105, 2001.

[73] Food and Agriculture Organization of the United Nations (FAO), "The state of food insecurity in the world," 2012. Available online at http://www.fao.org/publications/sofi/en/, accessed September 2013.

[74] Francis, P. and Smilowitz, K., "Modeling techniques for periodic vehicle routing problems," *Transportation Research Part B: Methodological*, vol. 40, no. 10, pp. 872–884, 2006.

[75] Fries, B. E., "Optimal ordering policy for a perishable commodity with fixed lifetime," *Operations Research*, vol. 23, no. 1, pp. 46–61, 1975.

[76] Gayon, J. P., De Véricourt, F., and Karaesmen, F., "Stock rationing in an M/Er/1 multi-class make-to-stock queue with backorders," *IIE Transactions*, vol. 41, no. 1, pp. 1096–1109, 2009.

[77] Gazis, D. C., "Resource allocation in a large decentralized enterprise," *European Journal of Operational Research*, vol. 30, no. 3, pp. 339–343, 1987.

[78] Gelman, A., Carlin, J. B., Stern, H. S., and Rubin, D. B., *Bayesian Data Analysis*. New York, NY: Chapman & Hall, 2004.

[79] Gerchak, Y. and He, Q. M., "On the relation between the benefits of risk pooling and the variability of demand," *IIE Transactions*, vol. 35, no. 11, pp. 1027–1031, 2003.

[80] Gürbüz, M. C., Moinzadeh, K., and Zhou, Y.-P., "Coordinated replenishment strategies in inventory/distribution systems," *Management Science*, vol. 53, no. 2, pp. 293–307, 2007.

[81] Ha, A., "Inventory rationing in a make-to-stock production system with several demand classes and lost sales," *Management Science*, vol. 43, no. 8, pp. 1093–1103, 1997.

[82] Ha, A., "Stock rationing in an M-Ek-1 make-to-stock queue," *Management Science*, vol. 46, no. 1, pp. 77–87, 2000.

[83] Hemmelmayr, V., Doerner, K. F., Hartl, R. F., and Savelsbergh, M. W., "Delivery strategies for blood products supplies," *OR Spectrum*, vol. 31, no. 4, pp. 707–725, 2009.

[84] Hentenryck, P. V., Bent, R., and Upfal, E., "Online stochastic optimization under time constraints," *Annals of Operations Research*, vol. 177, pp. 151–183, 2010.

[85] Jaillet, P. and Lu, X., "Online traveling salesman problems with service flexibility," *Networks*, vol. 58, no. 2, pp. 137–146, 2011.

[86] Jaillet, P. and Wagner, M. R., "Online routing problems: Value of advanced information as improved competitive ratios," *Transportation Science*, vol. 40, no. 2, pp. 200–210, 2006.

[87] Jaillet, P. and Wagner, M. R., "Generalized online routing: New competitive ratios, resource augmentation, and asymptotic analyses," *Operations Research*, vol. 56, no. 3, pp. 745–757, 2008.

[88] Japan Times, "Radiation, legalities complicate cleanup efforts," 2011. Available online at http://english.kyodonews.jp/news/2011/04/85506.html, accessed January 2013.

[89] Jaspars, S., *Solidarity and soup kitchens: a review of principles and practice for food distribution in conflict*. Overseas Development Institute, 2000.

[90] Jean-Louis, M. and Hilaire, E., "Debris removal situation, context and recommendations," 2010. Available online at http://haiti-ecap.com/Websites/dig/images/H-Debris-SF3.pdf, accessed January 2013.

[91] Johnson, M. P., Gorr, W. L., and Roehrig, S., "Location of service facilities for the elderly," *Annals of Operations Research*, vol. 136, no. 1, pp. 329–349, 2005.

[92] Kalai, A. and Vempala, S., "Efficient algorithms for online decision problems," *Journal of Computer and System Sciences*, vol. 71, no. 3, pp. 291–307, 2005.

[93] Kanchanasuntorn, K. and Techanitisawad, A., "An approximate periodic model for fixed-life perishable products in a two-echelon inventory-distribution system," *International Journal of Production Economics*, vol. 100, no. 1, pp. 101–115, 2006.

[94] Karaesmen, I., Scheller-Wolf, A., and Deniz, B., "Managing perishable and aging inventories: Review and future research directions," in *Planning*

Production and Inventories in the Extended Enterprise (Kempf, K. G., Keskinocak, P., and Uzsoy, R., eds.), vol. 151, pp. 393–436, Springer US, 2011.

[95] Karger, D. and Nikolova, E., "Exact algorithms for the Canadian traveller problem on paths and trees," Technical Report, Computer Science and Artificial Intelligence Laboratory, Massachusetts Institute of Technology, Cambridge, MA, 2008.

[96] Khouja, M., "The newsboy problem under progressive multiple discounts," *European Journal of Operational Research*, vol. 84, no. 1, pp. 458–466, 1995.

[97] Khouja, M., "The newsboy problem with multiple discounts offered by suppliers and retailers," *Decision Sciences*, vol. 27, no. 3, pp. 589–599, 1996.

[98] Khouja, M., "The single-period (news-vendor) problem: literature review and suggestions for future research," *Omega: International Journal of Management Science*, vol. 27, no. 1, pp. 537–553, 1999.

[99] Kibar, A. N. and Salman, S., "Logistics planning for restoration of network connectivity after a disaster," Working Paper, Department of Industrial Engineering, Koç University, Ankara, Turkey, 2013.

[100] Kouvelis, P. and Gutierrez, G., "The newsvendor problem in a global market: Optimal centralized and decentralized policies for a two-market stochastic inventory system," *Management Science*, vol. 43, no. 1, pp. 571–585, 1997.

[101] Lee, C.-Y., Çetinkaya, S., and Jaruphongsa, W., "A dynamic model for inventory lot sizing and outbound shipment scheduling at a third-party warehouse," *Operations Research*, vol. 51, no. 5, pp. 735–747, 2003.

[102] Lian, Z. and Liu, L., "A discrete-time model for perishable inventory systems," *Annals of Operations Research*, vol. 87, no. 1, pp. 103–116, 1999.

[103] Lian, Z., Liu, L., and Neuts, M. F., "A discrete-time model for common lifetime inventory systems," *Mathematics of Operations Research*, vol. 30, no. 3, pp. 718–732, 2005.

[104] Lin, C.-W. R. and Chen, H.-Y. S., "Dynamic allocation of uncertain supply for the perishable commodity supply chain," *International Journal of Production Research*, vol. 41, no. 13, pp. 3119–3138, 2003.

[105] Lin, Y. H., Batta, R., Rogerson, P. A., Blatt, A., and Flanigan, M., "A logistics model for delivery of prioritized items: Application to a disaster relief effort," working paper, School of Industrial and Systems Engineering, University of Buffalo (SUNY), Buffalo, NY, 2009.

[106] Littlestone, N. and Warmuth, M. K., "The weighted majority algorithm," in *30th Annual Symposium on Foundations of Computer Science*, (Research Triangle Park, NC), pp. 256–261, 1989.

[107] Lorca, Á., Çelik, M., Ergun, Ö., and Keskinocak, P., "A decision support tool for post-disaster debris operations," Working Paper, H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA, 2013.

[108] Luss, H., "On equitable resource allocation problems: A lexicographic minimax approach," *Operations Research*, vol. 47, no. 3, pp. 361–378, 1999.

[109] Luther, L., "Disaster debris removal after Hurricane Katrina: Status and associated issues," 2008. Avalable online at http://gactaern.org/Unit%20Plan/Healthcare%20Science_JM/Emergency_and_Disaster_Preparedness_25.564/EDP-17_Katrina_Debris_Report_JM.pdf, accessed: April 2013.

[110] Lystad, E. D., Ferguson, M. E., and Alexopoulos, C., "Single stage heuristic for perishable inventory control in two-echelon supply chains," working paper, Georgia Institute of Technology, Atlanta, GA, 2006.

[111] Marsh, M. T. and Schilling, D. A., "Equity measurement in facility location analysis: A review and framework," *European Journal of Operational Research*, vol. 74, no. 1, pp. 1–17, 1994.

[112] Matisziw, T. C., Murray, A. T., and Grubesic, T. H., "Strategic network restoration," *Networks and Spatial Economics*, vol. 10, no. 3, pp. 345–361, 2010.

[113] Mendonça, D. and Hu, Y., "Hurricane Katrina debris removal operations: The role of communication and computing technologies," *Learning from Catastrophe: Quick Response Research in the Wake of Hurricane Katrina*, vol. 40, pp. 283–306, 2006.

[114] Minner, S. and Transchel, S., "Periodic review inventory-control for perishable products under service-level constraints," *OR Spectrum*, vol. 32, no. 4, pp. 979–996, 2010.

[115] Nahmias, S., "A comparison of alternative approximations for ordering perishable inventory," *INFOR*, vol. 13, no. 1, pp. 175–184, 1975.

[116] Nahmias, S., "Optimal ordering policies for perishable inventory-II," *Operations Research*, vol. 23, no. 4, pp. 735–749, 1975.

[117] Nahmias, S., "The fixed-charge perishable inventory problem," *Operations Research*, vol. 26, no. 3, pp. 464–481, 1978.

[118] Nahmias, S., "Perishable inventory theory: A review," *Operations Research*, vol. 30, no. 4, pp. 680–708, 1982.

[119] Nahmias, S. and Pierskalla, W. P., "Optimal ordering policies for a product that perishes in two periods subject to stochastic demand," *Naval Research Logistics Quarterly*, vol. 20, no. 2, pp. 207–229, 1973.

[120] Nandakumar, P. and Morton, T. E., "Near myopic heuristics for the fixed-life perishability problem," *Management Science*, vol. 39, no. 12, pp. 1490–1498, 1993.

[121] Nurre, S. G., Çavdaroğlu, B., Mitchell, J. E., Sharkey, T. C., and Wallace, W. A., "Restoring infrastructure systems: An integrated network design and scheduling problem," *European Journal of Operational Research*, vol. 223, no. 3, pp. 794–806, 2012.

[122] Ogryczak, W., "On the lexicographic minimax approach to location problems," *European Journal of Operational Research*, vol. 100, no. 3, pp. 566–585, 1997.

[123] Özsen, L., Coullard, C. R., and Daskin, M. S., "Capacitated warehouse location model with risk pooling," *Naval Research Logistics*, vol. 55, no. 4, pp. 295–312, 2008.

[124] Paquet, S., Chaib-draa, B., and Ross, S., "Hybrid POMDP algorithms," in *Proceedings of the Workshop on Multi-Agent Sequential Decision Making in Uncertain Domains (MSDM-06)*, (Hakodate, Japan), pp. 133–147, MSDM, 2006.

[125] Perrier, N., Langevin, A., and Amaya, C. A., "Vehicle routing for urban snow plowing operations," *Transportation Science*, vol. 42, no. 1, pp. 44–56, 2008.

[126] Petruzzi, N. C. and Dada, M., "Pricing and the newsvendor problem: A review with extensions," *Operations Research*, vol. 47, no. 2, pp. 183–194, 1999.

[127] Pochet, Y. and Wolsey, L. A., "Lot-sizing with constant batches: Formulation and valid inequalities," *Mathematics of Operations Research*, vol. 18, no. 4, pp. 767–785, 1993.

[128] Powell, W. B. and Frazier, P., "Optimal learning," in *Tutorials in Operations Research* (Mirchandani, P., ed.), vol. 5, pp. 213–246, Hanover, MD: INFORMS, 2012.

[129] Powell, W. B. and Ryzhov, I. O., *Optimal Learning*. Hoboken, NJ: John Wiley & Sons, 2012.

[130] Prastacos, G., "Optimal myopic allocation of a product with fixed lifetime," *Journal of the Operational Research Society*, vol. 29, no. 9, pp. 905–913, 1978.

[131] Prastacos, G., "Allocation of a perishable product inventory," *Operations Research*, vol. 29, no. 1, pp. 95–107, 1981.

[132] Qin, Y., Paul, A., and Alptekinoğlu, A., "Optimal distribution of blood products," Working Paper, Cox School of Business, Southern Methodist University, Dallas, TX, 2012.

[133] Qin, Y., Wang, R., Vakharia, A. J., Chen, Y., and Seref, M. M. H., "The newsvendor problem: Review and directions for future research," *European Journal of Operational Research*, vol. 213, no. 1, pp. 361–374, 2011.

[134] Rancourt, M. E., Cordeau, J. F., Laporte, G., and Watkins, B., "Tactical network planning for food aid distribution in Kenya," Working Paper, Interuniversity Research Centre on Enterprise Networks, Logistics and Transportation (CIRRELT), Montreal, QC, Canada, 2013.

[135] Ross, S. and Chaib-draa, B., "AEMS: An anytime search algorithm for approximate policy refinement in large POMDPs," in *Proceedings of the 20th International Joint Conference on Artificial Intelligence (IJCAI-07)*, (Hyderabad, India), pp. 2592–2598, 2007.

[136] Ross, S., Pineau, J., Paquet, S., and Chaib-draa, B., "Online planning algorithms for POMDPs," *Journal of Artificial Intelligence Research*, vol. 32, no. 1, pp. 663–704, 2008.

[137] Rothkopf, M. S., "Scheduling with random service times," *Management Science*, vol. 12, no. 1, pp. 707–713, 1966.

[138] Ryzhov, I. O. and Powell, W. B., "Information collection on a graph," *Operations Research*, vol. 59, no. 1, pp. 188–201, 2011.

[139] Şahin, H., Kara, B. Y., and Karaşan, O., "On debris removal during the response phase," Working Paper, Department of Industrial Engineering, Bilkent University, Ankara, Turkey, 2013.

[140] Şen, A. and Zhang, A., "The newsboy problem with multiple demand classes," *IIE Transactions*, vol. 31, no. 5, pp. 431–444, 1999.

[141] Shahabi, M., Akbarinasaji, S., Unnikrishnan, A., and James, R., "Integrated inventory control and facility location decisions in a multi-echelon supply chain network with hubs," *Networks and Spatial Economics*, vol. 13, no. 4, pp. 497–514, 2013.

[142] Shen, Z.-J. M., Coullard, C., and Daskin, M. S., "A joint location-inventory model," *Transportation Science*, vol. 37, no. 1, pp. 40–55, 2003.

[143] Shen, Z.-J. M. and Daskin, M. S., "Trade-offs between customer service and cost in integrated supply chain design," *Manufacturing & Service Operations Management*, vol. 7, no. 3, pp. 188–207, 2005.

[144] Shu, J., Teo, C.-P., and Shen, Z.-J. M., "Stochastic transportation-inventory network design problem," *Operations Research*, vol. 53, no. 1, pp. 48–60, 2005.

[145] Shumsky, R. A. and Zhang, F., "Dynamic capacity management with substitution," *Operations Research*, vol. 57, no. 3, pp. 671–684, 2009.

[146] Snyder, L. V., Daskin, M. S., and Teo, C.-P., "The stochastic location model with risk pooling," *European Journal of Operational Research*, vol. 179, no. 3, pp. 1221–1238, 2007.

[147] Solis, G. Y., Hightower, H. C., Sussex, J., and Kawaguchi, J., *Disaster Debris Management*. Vancouver, BC: Disaster Preparedness Resources Centre, University of British Columbia, 1996.

[148] Srinivas, H. and Nakegawa, Y., "Environmental aspects of the Indian Ocean tsunami recovery," *Journal of Environmental Management*, vol. 89, no. 1, pp. 4–13, 2009.

[149] Stilp, M., Ergun, Ö., and Keskinocak, P., "Management of debris operations," Working Paper, H. Milton Stewart School of Industrial and Systems Engineering, Georgia Institute of Technology, Atlanta, GA, 2012.

[150] Stulman, A., "Benefits of centralized stocking for the multi-centre newsboy problem with first come, first served allocation," *Journal of the Operational Research Society*, vol. 38, no. 9, pp. 827–832, 1987.

[151] Swaminathan, J. M., "UNICEF's Plumpy' Nut supply chain," tech. rep., Teaching Case. Kenan–Flagler Business School, University of North Carolina, Chapel Hill, NC, 2009.

[152] Swan, R. C., "Debris management planning for the 21st century," *Natural Hazards Review*, vol. 1, no. 4, pp. 222–225, 2000.

[153] Taaffe, K., Romeijn, E., and Tirumalasetty, D., "A selective newsvendor approach to order management," *Naval Research Logistics*, vol. 55, no. 8, pp. 769–784, 2009.

[154] Tancrez, J.-S., Lange, J.-C., and Semal, P., "A location-inventory model for large three-level supply chains," *Transportation Research Part E: Logistics and Transportation Review*, vol. 48, no. 2, pp. 485–502, 2012.

[155] Tekin, E., Hopp, W. J., and Van Oyen, M. P., "Pooling strategies for call center agent cross-training," *IIE Transactions*, vol. 41, no. 6, pp. 546–561, 2009.

[156] Thrun, S., "Robotic mapping: A survey," in *Exploring Artificial Intelligence in the New Millenium* (Lakemeyer, G. and Nebel, B., eds.), pp. 1–36, Burlington, MA: Morgan Kaufmann, 2002.

[157] Tomasini, R. M. and Van Wassenhove, L. N., "Genetically modified food donations and the cost of neutrality: Logistics response to the 2002 food crisis in Southern Africa," tech. rep., No 03/2004-5169. INSEAD, Fontainebleau, France, 2004.

[158] Tomasini, R. and Van Wassenhove, L., "Moving the World: The TPG-WFP Partnership–Looking for a Partner," tech. rep., No 02/2004-5187. INSEAD, Fontainebleau, France, 2004.

[159] Trunick, J. A., "Logistics when it counts," *Logistics Today*, vol. 46, no. 2, p. 38, 2005.

[160] Tzeng, G. H., Cheng, H. J., and Huang, T. D., "Multi-objective optimal planning for designing relief delivery systems," *Transportation Research Part E: Logistics and Transportation Review*, vol. 43, no. 6, pp. 673–686, 2007.

[161] United Nations Office for the Coordination of Humanitarian Affairs, "Disaster waste management guidelines," 2011. Available online at http://ochanet.unocha.org/p/Documents/DWG%20Annex%20XII.Disaster%20waste%20management%20contingency%20planning.pdf, accessed January 2013.

[162] U.S. Census Bureau, "Census 2000 gateway," 2001. Available online at http://www.census.gov/main/www/cen2000.html, accessed September 2013.

[163] U.S. Department of Health and Human Services, "An HHS retrospective on the 2009 H1N1 influenza pandemic to advance all hazards preparedness," 2012. Available online at http://www.phe.gov/Preparedness/mcm/h1n1-retrospective/Documents/h1n1-retrospective.pdf, accessed September 2013.

[164] van Dijk, N. M. and van der Sluis, E., "To pool or not to pool in call centers," *Production and Operations Management*, vol. 17, no. 3, pp. 296–305, 2008.

[165] van Dijk, N. M. and van der Sluis, E., "Pooling is not the answer," *European Journal of Operational Research*, vol. 197, no. 1, pp. 415–421, 2009.

[166] van Houtum, G. J., "Multiechelon production/inventory systems: Optimal policies, heuristics, and algorithms," in *Tutorials in Operations Research* (Johnson, M. P., Norman, B., and Secomandi, N., eds.), pp. 163–199, INFORMS, Hanover, MD, 2006.

[167] Van Zyl, G. J. J., *Inventory Control for Perishable Commodities*. Unpublished Ph.D. thesis, University of North Carolina, 1964, Chapel Hill, NC, 1964.

[168] Washington Post, "Haiti faces colossal and costly cleanup before it can rebuild," 2010. Available online at http://www.washingtonpost.com/wp-dyn/content/article/2010/03/06/AR2010030602544.html, accessed January 2013.

[169] Williams, C. L. and Patuwo, B. E., "A perishable inventory model with positive order lead times," *European Journal of Operational Research*, vol. 116, no. 2, pp. 352–373, 1999.

[170] Williams, C. L. and Patuwo, B. E., "Analysis of the effect of various unit costs on the optimal incoming quantity in a perishable inventory model," *European Journal of Operational Research*, vol. 156, no. 1, pp. 140–147, 2004.

[171] World Food Programme, "World hunger," 2012. Available online at http://www.wfp.org/hunger, accessed September 2013.

[172] World Food Programme, "Managing the supply chain of specialized nutritious foods," 2013. Available online at http://documents.wfp.org/stellent/groups/public/documents/manual_guide_proced/wfp 259937.pdf, accessed May 2014.

[173] Yan, S. and Shih, Y., "Optimal scheduling of emergency roadway repair and subsequent relief distribution," *Computers & Operations Research*, vol. 36, no. 6, pp. 2049–2065, 2009.

[174] Yen, H.-C., *Inventory management for a perishable product multi-echelon system.* Unpublished Ph.D. thesis, Northwestern University, 1975, Evanston, IL, 1975.

[175] Zhong, H., Hall, R. W., and Dessouky, M., "Territory planning and vehicle dispatching with driver learning," *Transportation Science*, vol. 41, no. 1, pp. 74–89, 2007.