Title
Dynamic Facility Relocation and Inventory Management for Disaster Relief

Author
Richter, Amber Rae

Publication Date
2016

Peer reviewed|Thesis/dissertation

Dynamic Facility Relocation and Inventory Management
for Disaster Relief

By

Amber Rae Richter

A dissertation submitted in partial satisfaction of the

requirements for the degree of

Doctor of Philosophy

in

Engineering - Industrial Engineering and Operations Research

in the

Graduate Division

of the

University of California, Berkeley

Committee in charge:

Professor Zuo-Jun (Max) Shen, Chair
Professor Philip M. Kaminsky
Professor Terry A. Taylor

Spring 2016

Abstract

Dynamic Facility Relocation and Inventory Management
for Disaster Relief

by

Amber Rae Richter

Doctor of Philosophy in Industrial Engineering and Operations Research

University of California, Berkeley

Professor Zuo-Jun (Max) Shen, Chair

Disasters strike suddenly and cause destruction which disrupts the availability of basic survival supplies for people living in affected areas. The efficiency of humanitarian organizations in providing relief has a direct and crucial impact on the survival, health, and recovery of affected people and their communities. To better prepare to respond to disasters, many relief organizations use supply pre-positioning. However, the real and potential needs of different locations change over time and when an organization uses traditional warehouse pre-positioning, relief operations are limited by set inventory locations that are difficult to alter. For this reason, a well known organization recently considered including a large supply holding ship in its operations. By holding inventory on a ship, the organization would be able to dynamically relocate its inventory over time in response to changing relief supply demand forecasts.

To our knowledge, the research contained herein is the first to examine dynamic inventory relocation for responding to disasters over time. Specifically, we examine how to optimally relocate and manage inventory for a single mobile inventory to serve stochastic demand at a number of potential disaster sites over time. While we keep in mind the motivating example of a supply holding ship in the disaster relief setting throughout this dissertation, the model and most of the results are applicable to any type of mobile inventory, facility, or server in any setting.

We first examine the dynamic relocation problem. We model the problem using dynamic programming and develop analytical and numerical results regarding optimal relocation policies, the optimal path and speed of relocation decisions, and the value of inventory mobility over traditional warehouse pre-positioning. To help overcome the computational complexity of the problem, we develop a heuristic which solves relatively large problem instances in our numerical experiments within 0.5% of optimality in less than 0.1% of the time required by an exact algorithm.

As it is suboptimal to consider relocation decisions and inventory management decisions separately, we also examine the joint dynamic relocation and inventory management problem. To our knowledge, we are the first to examine the dynamic relocation and inventory management problem with stochastic demand. Similarly to the dynamic relocation prob-

lem, we model this problem using dynamic programming. We develop a number of analytical results characterizing the optimal relocation and inventory management policies.

As the first to examine these problems, we hope this research serves as a catalyst for other research in this area; accordingly, we conclude this dissertation by discussing a number of areas for future research.

To my family,
Michael, Chris, Ryan, and Carolyn Richter
and Jessica, Mike, and Benjamin Wolf

Contents

List of Figures

List of Tables

Acknowledgments

Throughout my PhD, I have received a tremendous amount of support and encouragement from many different people. My words here are not enough to fully express my sincere and heartfelt appreciation to all those who have helped me.

I would first like to thank my advisor, Professor Max Shen. If it were not for his encouragement when I was an undergraduate student, I may never have applied to a PhD program. Throughout my PhD, he gave me the freedom to choose my own research topic, structure my own time, and explore other interests and ventures alongside my dissertation for which I am incredibly grateful. He always made himself available whenever I needed guidance and supported me through the many different phases of my PhD program and internship and job search. Throughout the time I have known him, he has served as a mentor to me and has always had my professional and personal interests and happiness in mind.

I am very grateful to Phil Kaminsky, Terry Taylor, and Ying-Ju Chen who have also served on my committees throughout the years. It is their interest, enthusiasm, and insights that made my qualifying exam and dissertation workshops exceptionally useful and fun. They have provided invaluable feedback and guidance that has helped to shape the work detailed in this dissertation. In addition to direct feedback, Phil and Ying-Ju taught courses in supply chain management and dynamic programming which provided me with some of the key fundamental technical knowledge I needed to approach this research.

I would like to thank Professor J. George Shanthikumar for discussing some of the research problems in this dissertation with Max and me and sharing his insights and feedback. I would also like to thank the many other people who have provided feedback on this work, especially Dr. Paula Lipka and Dr. Chen Chen of UC Berkeley and Professor Gemma Berenguer of Purdue University. I am also very grateful to the National Science Foundation and UC Berkeley for supporting the work upon which this dissertation is based through the National Science Foundation Graduate Research Fellowship under Grant No. DGE 1106400 and the UC Berkeley Chancellor's Fellowship.

I would also like to thank the many other members of the Industrial Engineering and Operations Research Department who have supported me throughout the years. I am especially thankful to Professor Rhonda Righter, who gave me the opportunity and taught me to teach and has been especially kind and supportive, to Candi Yano, who has always been eager to provide encouragement and useful advice and support, and to the many members of the department staff, especially Michael Campbell, Rochelle Niccolls, Brionna Garner, and Anayancy Paz, who provided administrative support and friendship throughout key points in my work.

I am incredibly grateful to the many officemates and colleagues who have supported me as friends throughout my time here and made it an enjoyable and fun experience. I am especially grateful to many of them of which I will list only a few here: Dr. Paula Lipka, Dr. Nguyen Truong, Dr. Evan Davidson, Dr. Yifen Chen, Dr. Wei Qi, Dr. Tony Ke, Stewart Liu, Rebecca Sarto, Quico Spaen, Cheng Lyu, Birce Tezel, Tugce Gurek, Kevin Li, Jiung Lee, and Dr. Giulia Pedrielli.

Finally, I would like to thank my family and friends who have given me so much love,

support, and happiness and who have been my cheerleaders throughout my PhD. It is your love and support that has kept me motivated and strong. I am forever grateful to you all.

1 Introduction

In 2014, there were 324 major natural disasters which left 140.7 million people in need of relief in the form of basic survival supplies (e.g. food, water, and shelter) and basic services (e.g. medical care) (Guha-Sapir et al. 2015). Natural and man-made disasters, such as earthquakes, tropical cyclones, tsunamis, and some complex humanitarian emergencies, strike suddenly and leave communities in a state of emergency. They cause destruction which disrupts the availability of medical and basic survival supplies for people living in the affected areas.

Humanitarian organizations, such as the National Red Cross and Red Crescent Societies, World Vision International, and CARE International, work to provide relief to people affected by disasters. Their efficiency in providing relief has a direct and crucial impact on the survival, health, and recovery of affected people and their communities. Relief organizations must respond as quickly and effectively as possible to minimize human suffering and loss of life. Despite this, relief organizations have only relatively recently begun to view their logistics as a strategic component of their relief efforts rather than simply as a necessary expense (Beamon and Kotleba 2006). Accordingly, relief efforts often lack efficiency and effectiveness. For example, Secretary of Homeland Security Michael Chertoff testified: "FEMA's logistics systems simply were not up to the task" of providing supplies to the tens of thousands of people in need of relief from Hurricane Katrina (U.S. Senate 2006).

Planning relief operations when a disaster strikes without existing infrastructure or established procedures is extremely time consuming. Furthermore, procuring supplies after a disaster has occurred is often prohibitively expensive or impossible to organize in the critically short relief period necessitated by major disasters due to stockouts or long lead times. Thus, it is imperative for relief organizations to prepare to respond to disasters before disasters happen.

One important strategy that many organizations use to prepare is procuring supplies ahead of time and pre-positioning them in warehouses in areas which are close to or have access to major disaster-prone regions. For example, World Vision International operates a global pre-positioning system which pre-positions supplies in the US, Italy, Germany, and Dubai and the World Food Programme manages the United Nations Humanitarian Response Depot which is able to send supplies anywhere in the world within 24-48 hours (Balcik and Beamon). Inventory pre-positioning allows organizations to potentially reduce the lead time to affected individuals, spend less of their limited budgets on procuring and transporting supplies, and have more control over the availability and quality of their supplies.

However, the real and potential needs of different locations change over time. For example, hurricanes are more likely during hurricane season, complex emergencies are more likely in a time of unrest rather than in a time of peace, and tsunamis are more likely following major earthquakes. When an organization uses traditional pre-positioning, relief operations are limited by set inventory locations that are difficult to alter. For this reason, a well known organization recently considered including a large supply holding ship in its operations along with smaller ships or helicopters to transport supplies to affected communities. By holding inventory on a ship, the organization would be able to dynamically relocate its inventory over time in response to changing relief supply demand forecasts. For example,

- during hurricane season, the inventory could be moved closer to hurricane zones
- if there are rising tensions in a particular location, the inventory could be moved closer to where conflict may erupt
- if a disaster just occurred and supplies will be needed over the next couple of periods, the inventory could be moved closer to the ongoing relief operations

Inventory mobility can potentially decrease operational costs and response times over traditional pre-positioning as the inventory may be closer to affected communities. Additionally, the flexibility afforded by inventory mobility may reduce the need for inventory duplication in multiple locations as the same inventory may be able to serve a larger area or different areas at different times. This would result in lower operational and holding costs and less inventory lost to expiration. Furthermore, a supply holding ship does not require a working airport to reach an affected community and may be more easily secured from theft or disaster-related damage.

For these reasons, the organization motivating this work is not the first to consider using a ship or other type of mobile unit for holding and delivering relief supplies and services. For example, Floating Doctors, a nonprofit medical relief team, is primarily based on a ship which provides medical relief to coastal regions worldwide and responded to the Haiti earthquake in 2010. Project Hope, an international health care organization, received a retired U.S. Navy Hospital ship as a donation in 1958 and used it for a number of years in providing health care, education, and disaster relief. Many countries, including the US, Germany, China, and Mexico, use hospital or military ships for disaster response. Additionally, several organizations store hurricane relief supplies in containers which can be moved in response to updated hurricane forecasts.

To our knowledge, this research is the first to examine mobile inventory supply chains for disaster relief. As the first step to understanding and evaluating these systems, we examine how to optimally relocate and manage inventory for a single mobile inventory, such as a supply holding ship, to serve stochastic demand at a number of potential disaster sites over time. While we keep in mind the motivating example of a supply holding ship in the disaster relief setting throughout this research, the models and most of the results are applicable to any type of mobile inventory, facility, or server in any setting. Especially relevant may be the application to military sea basing (see Qiu and Sharkey (2013)) or to the Maritime Administration's National Defense Reserve Fleet as discussed in the next section. In the disaster relief context, mobile inventory may take the form of a ship, trailer, container, or mobile health clinic.

In Section 3, we examine the dynamic facility relocation decisions. We also examine the value of inventory mobility over traditional warehouse pre-positioning to provide managerial insights on when investing in or encouraging the donation of a mobile inventory aspect of a supply chain may be worthwhile. While there are many areas in which inventory mobility may provide value, such as decreased response time, we focus on potential operational cost advantages resulting from decreased transportation distances to affected communities.

It is suboptimal to consider relocation decisions and inventory management decisions separately as the inventory location affects how much it costs to restock. For example, if the inventory is far from a supplier, it will likely cost more to place and receive an order. Thus, in Section 4, we consider making dynamic relocation and inventory management decisions simultaneously by examining the joint relocation and inventory management problem.

The main contributions of this dissertation are as follows. Through examining the dynamic facility relocation problem in Section 3,

1. to our knowledge, we are the first in the literature to consider the dynamic alteration of a disaster relief supply chain in response to changing demand patterns over time

2. to our knowledge, we are the first in the literature to examine mobile inventory supply chains for disaster relief

3. to our knowledge, we are the first in the literature to study dynamic facility location with demand which evolves according to a non-stationary discrete time Markov chain (DTMC)

4. we develop a model to determine an optimal relocation plan for a single mobile inventory to serve stochastic demand over time

5. we derive analytical results regarding optimal movement and relocation policies, including results which allow us to reduce the size of some large problem instances

6. we develop analytical and numerical results regarding the value of inventory mobility over traditional warehouse pre-positioning and extract managerial insights on when investing in or encouraging the donation of a mobile inventory aspect of a supply chain may be worthwhile

7. we design a heuristic which solves relatively large problem instances in our numerical experiments within 0.5% of optimality in less than 0.1% of the time required by an exact algorithm and is optimal when the demands are temporally independent

Through examining the joint dynamic facility relocation and inventory management problem in Section 4,

8. to our knowledge, we are the first in the literature to consider both dynamic relocation and inventory management decisions with stochastic demand

9. we develop a model to determine an optimal relocation and inventory management policy for a single mobile inventory to serve stochastic demand over time

10. we develop analytical results characterizing the optimal relocation and inventory management policies; specifically, we show that:

    (a) a multiperiod (s, S) policy is optimal when we are restricted to movement policies that do not depend on the inventory level

(b) while a multiperiod (s, S) policy is not optimal for the general problem, in an optimal policy we will either order nothing or order up to some sum of potential future period demands

(c) counterintuitively, the optimal amount to order is not necessarily decreasing in our initial inventory level

(d) in an optimal policy, we will not place an order if our initial inventory level is weakly greater than the maximum possible demand for the rest of the horizon

(e) it is sufficient for optimality to consider a smaller feasible set of inventory locations, defined in Section 4.1.2 below, and thereby reduce the size of the problem

(f) in an optimal policy, we will not move the mobile inventory toward the supplier at the expense of moving it farther from all potential disaster sites in a period in which we do not place an order

(g) in an optimal policy, we will not place an order if the mobile inventory is in a location farthest from the supplier and our current inventory level is weakly greater than the maximum possible current period demand

This dissertation is organized as follows. Section 2 reviews the relevant literature. Section 3 examines the dynamic facility relocation problem and the value of inventory mobility in disaster relief. Section 4 examines the joint dynamic relocation and inventory management problem. Finally, Section 5 concludes the dissertation and suggests future research directions.

2 Literature Review and Motivation

In this section, we review the literature relevant to the problem of managing mobile inventory supply chains for disaster relief. The streams of literature most relevant to this research focus on locating and managing pre-positioned inventory for disaster relief and the dynamic facility location problem. Section 2.1 reviews the literature related to locating and managing pre-positioned inventory and Section 2.2 reviews the literature related to the dynamic facility location problem.

2.1 Pre-Positioning for Disaster Relief

When a disaster occurs, a relief organization needs to be able to quickly deliver relief supplies to the area affected by the disaster. If there are no supply inventories near the location of the disaster, or if these supplies are not well managed, it may be difficult for the relief organization to respond to the disaster effectively. Thus, it is imperative for pre-positioned inventories to be carefully positioned and managed. In this section, we review operations research literature in strategic emergency supply pre-positioning. This stream of literature can be further divided into that which focuses on location and stocking decisions and that which focuses on inventory management decisions for pre-positioned supplies; we review these focus areas in Sections 2.1.1 and 2.1.2, respectively.

2.1.1 Location and Stocking Decisions for Pre-Positioned Supplies

Most of the emergency supply pre-positioning literature focuses on making one-time decisions of supply locations and stocking levels to optimally respond to a single disaster period. Typically, papers in this stream model the pre-positioning problem as extensions of well known facility location models and their novelty is in what types of complications or constraints they consider. The theses of Akkihal (2006) and McCall (2006) appear to be among the first operations research papers to look at disaster relief pre-positioning. Balcik and Beamon (2008) use a maximal covering location model to determine the number, location, and stocking levels of capacitated distribution centers to maximize the total expected demand covered for one period. They consider several types of supplies with different levels of importance and response time requirements and take into account transportation capacity restrictions and pre- and post-disaster budgetary constraints. Their computational experiments point out the importance of pre-disaster investments which they say have been underrated compared to investments in post-disaster response activities. Verma and Gaukler (2011) develop a two-stage stochastic programming model based on the capacitated facility location model to determine the locations of warehouses with pre-determined capacities. Their model takes into account uncertainties in the functioning of warehouses and the subsequent availability of supplies following a disaster. They use a case study to show that their model places facilities at a safer distance from disaster epicenters rather than directly on top of high-risk areas.

Ukkusuri and Yushimito (2008) were the first to take into account disruptions in the transportation network. They develop a model to determine the locations of pre-positioned supplies to maximize the probability that demand points can be reached from at least one supply holding site for one period. Rawls and Turnquist (2010) develop a two-stage stochastic programming model which, in addition to determining the number, location and stocking levels, determines the size of supply holding warehouses to minimize cost, including penalties for unmet demand, for a single disaster period. Their model allows for uncertainties in the survival of inventories and disruptions in the transportation network. They develop a heuristic algorithm for solving their model using the Integer L-Shaped Method and test the robustness and applicability of their model on a test case. Rawls and Turnquist (2011) extend this model to account for service quality constraints. Noyan (2012) also extends the Rawls and Turnquist (2010) model. Noyan notes that, like the Rawls and Turnquist (2010) model, most two-stage stochastic programming models represent a risk neutral approach in that they consider the expectation of an event occurring as the preference criterion of the decision maker. Noyan thus extends the model by incorporating the conditional-value-at-risk (CVaR) as the risk measure on the total cost for the reason that "Considering only the expected values may not be good enough for rarely occurring disaster events." Noyan reformulates the model as a two-stage mean-risk stochastic programming model which incorporates the trade-off between the expected total cost and a risk measure on the random total cost. Similarly to Rawls and Turnquist (2010), Noyan develops a heuristic for solving the model using the L-Shaped Method.

The majority of the stochastic facility location models related to the pre-positioning of supplies for disaster response, including all those described above, involve a single echelon network. Döyen et al. (2012) develop a two-echelon two-stage stochastic programming model in which locations are determined for both regional rescue centers (RRCs) and local rescue centers (LRCs) where RRCs (the first echelon facilities) are located before an event occurs and LRCs (the second echelon facilities) are located after an event occurs and stocked by the RRCs to serve end demand. They develop a solution heuristic based on Lagrangian relaxation and augment it by local search to improve the efficiency of the solution technique.

Although most emergency supply pre-positioning papers take as inputs a set of scenarios of possible events that could occur and probabilities on each scenario's occurrence, not all do. Campbell and Jones (2011) were the first to consider both risk and inventory levels related to the pre-positioning of emergency supplies for disaster relief without the use of scenarios. Their motivations are that it is not always possible to build these scenarios as this requires having sufficient historical data and that a limited number of scenarios may not be sufficient to represent the potential set of outcomes.

Most research in this literature stream uses historical disaster data as proxies for future demand forecasts. Noting that historical data may not accurately represent expected effects of disasters on areas which have expanded in population and infrastructure over the years, Barzinpour and Esmaeili (2014) utilize software which predicts the impact of an earthquake on Iran's current infrastructure and population level to develop an earthquake pre-positioning plan for Iran. They develop a multi-objective mixed integer linear programming model based on the maximal covering model to locate warehouses and assign warehouses to urban regions.

Some models, like Barzinpour and Esmaeili (2014), are developed specifically for a particular organization to solve a specific problem instance. These models tend to have more assumptions and constraints which are specifically tailored to the situation in question. Duran, Gutierrez, and Keskinocak (2011) develop a mixed integer program (MIP) for CARE International, one of the world's largest disaster relief organizations, to evaluate the effect of pre-positioning on CARE's average response time to disasters worldwide. Their model finds the optimal number and location of pre-positioned warehouses while allowing for demand to be met by both the pre-positioned supplies as well as outside suppliers; this is in contrast to most other papers which assume that the pre-positioned supplies will be the only source of supplies when a disaster hits. The model takes as input a given specific initial investment, in terms of the maximum number of warehouses and total inventory to allocate, with an objective of minimizing CARE's average response time over a number of demand scenarios. Their results allowed CARE to determine their desired pre-positioning network and a plan for how to construct it as they acquired more funds. As an example outside of the disaster relief context, Amouzegar et. al (2006) and McGarvey et. al (2010) examine the optimal pre-positioning of war reserve material resources for the U.S. Air Force using an MIP model.

Bozkurt and Duran (2012) extend the work of Duran, Gutierrez, and Keskinocak (2011) by examining the expansion plan of CARE's pre-positioning network. They argue that pre-positioned inventory locations should be robust to changes in disaster types, locations, and magnitudes over the years in which they are in use. For this reason, in contrast to Duran, Gutierrez, and Keskinocak (2011) which uses historical disaster data from 1997-2006, the authors use historical disaster data from 1977-2006. They run the Duran et. al (2011) model separately with data from each of these three decades and show that differences in disaster occurrences between the three decades suggest different pre-positioned inventory locations. With these changing location recommendations over the decades, they infer trends in disaster occurrences and subsequently recommend a new location for CARE to add to its pre-positioning network.

Bozkurt and Duran (2012) is the first study, to our knowledge, which notes that the optimal configuration of a pre-positioned network may change over time. This suggests a long term view of the advantage of having a flexible, mobile inventory system for disaster relief as its configuration can be directly altered over time. Bozkurt and Duran (2012) approach this issue by considering one-time alterations of an existing network. To our knowledge, no existing work considers the design or dynamic alteration beyond a single change of a disaster relief inventory supply chain. This research on mobility inventory supply chains for disaster relief addresses this literature gap.

2.1.2 Inventory Management Decisions for Pre-Positioned Supplies

While most emergency supply pre-positioning literature focuses on making one-time decisions of warehouse locations and stocking levels, another stream of the literature focuses on inventory management decisions for pre-positioned emergency supplies. Ozbay and Ozguven (2007) consider the deliveries and demand of a certain type of supply at one distribution center designated as a gathering center for evacuees following a major disaster. They develop a two-stage stochastic programming model to determine the initial amount of safety stock of the supply needed so that the chance of a stock-out during response to a disaster is less than or equal to some given probability. The objective is to minimize inventory holding costs, shortage costs, surplus inventory costs, and the costs of adjusting the safety stock level after one or more periods of demand is realized. Ozbay and Ozguven (2013) extended this model as part of a complex humanitarian inventory management system and control model. They developed a stochastic programming model to serve as the off-line planning strategy which is called multiple times by the on-line inventory management system which uses Radio Frequency Identification Device (RFID) technology for commodity tracking and logistics to manage the coordination of relief activities before and after an event occurs. The stochastic programming model from their previous paper is extended to account for multiple supply types, where some supplies are substitutable for others, and multiple suppliers, where multiple suppliers may be needed for the same supply type. In this case, the desired safety stock levels of the different supply types are dependent on each other due to capacity constraints. As actual inventory levels can deviate from optimal levels during response to a disaster due to stochastic disruptions, the on-line inventory management system is developed using the stochastic programming model described to minimize the impacts of these disruptions.

Beamon and Kotleba (2006-1) develop a stochastic inventory management model that determines optimal order quantities and reorder points for long-term emergency relief response in complex humanitarian emergencies which have unpredictable demand patterns and long durations. They focus their paper on the second civil war in south Sudan (1983-2005). They consider a continuous review inventory management system under a (Q, r)-type policy. There are two options for re-supply, one a normal reorder option and the other an emergency reorder option, each with a constant lead time. They also note the need for a more thorough study on the implications of back order costs as the corresponding commercial logistics interpretation of lost sales is not appropriate in the context of humanitarian logistics where the implications of unmet demand may be the suffering or even the death of the "customer." Beamon and Kotleba (2006-2) further studies inventory management strategies for pre-positioned stocks of supplies intended to support relief related to the civil war in south Sudan. They compare the mathematical model developed in the paper just described to a heuristic model, which may be preferable due to computational time limitations, and to a simple 'naive' model whose advantage is its ease of implementation. They test the three strategies using a simulation to be able to analyze the affects of back order costs, service levels, and the demand distribution on total cost, response time, and flexibility and develop a performance measurement system to compare the strategies.

Rottkemper et al. (2011) focus on inventory management during on ongoing long-term humanitarian relief operation taking into account the possibility of overlapping disasters, such as an epidemic outbreak or warehouse fire, which could cause a sudden increase in demand or decrease in supply. In their linear multi-period MIP, established regional depots are served from a central depot which is, in turn, served by a global depot and inventory management decisions are made for each level. They consider a single type of supply and allow for the option of moving inventory between regional depots. The objective function is to minimize costs including penalty costs for unsatisfied demand. The model assumes that there is a deterministic portion of demand associated with the ongoing humanitarian operation as well as a stochastic future demand to allow for possible future disruptions. Penalty costs for the deterministic demand are higher than those for the stochastic demand and demand is backlogged. They develop a rolling horizon solution approach as well as an alternative solution approach based on a simple decision tree heuristic. Additionally, they present a number of numerical experiments including those which they used to determine the appropriate value for penalty costs.

As weather forecasts allow decision makers to predict the onset of hurricanes, and weather patterns during the pre-hurricane season allow for updates of these forecasts, there exists a subset of disaster response literature related to preparedness specifically for hurricanes. For example, Davis et al. (2013) examine reallocating inventory among open, capacitated relief supply warehouses to prepare for an impending hurricane using a two-stage stochastic programming model. This work utilizes short-term hurricane path forecasts to predict which warehouses and which populations will be most affected by the storm. It takes into account potential loss of supplies at warehouses and road congestion before and after the storm due to evacuations and storm related road damage. Taskin and Lodree (2010) develop a model which determines an inventory management policy for manufacturing and retail firms for the pre-hurricane season and the beginning of the hurricane season. Their stochastic programming methodologies allow inventory managers to alter inventory decisions as new information regarding the incoming hurricane season comes to view. The methodologies are designed in this way in hopes of reducing the amount of stock outs and thus improving the ability of people in the community to acquire the supplies they need in the pre-hurricane and hurricane seasons. The model divides the pre-hurricane season into several periods in which the inventory manager has the ability to determine inventory levels that account for demands in the subsequent periods and at the beginning of the hurricane season. They formulate the problem as a multi-stage stochastic programming model with recourse to determine order or production quantities which minimize total cost. They note that their research could be useful for not-for-profit disaster relief operations as well.

Rawls and Turnquist (2012) is unique in that it considers both warehouse location and short term inventory management. They develop a two-stage stochastic model which extends the models they developed in their works described above (Rawls and Turnquist (2010), Rawls and Turnquist (2012)) to also make inventory management decisions for meeting short term demands similar to those made in the model developed by Ozbay and Ozguven (2007). Their potential stocking locations include both potential warehouse locations as well as gathering centers, or shelters, for evacuees. Their model determines the location and size of warehouses, pre-positioned stocking quantities of various types of emergency supplies at these warehouses and at the shelters, and an inventory management policy for each shelter in response to potential events including a distribution plan from which warehouses will meet orders from that shelter over a finite number of time periods. Their model allows for differentiation between supplies for which demand is dependent on the duration of time for which evacuees reside in a shelter and supplies for which demand is not and requires that all demands be met in scenarios comprising a certain percentage of all outcomes. Their model takes into account limited shelter storage capacity, limited capacity of shipment travel nodes, transportation network availability, supply survival uncertainties, and the rate at which each warehouse can process and satisfy orders which is dependent on the warehouse's size. They test their model on a case study of hurricane preparedness in North Carolina and note the need for the development of a specialized algorithm to solve some instances of the problem.

Despite the growing number of papers on strategic decision making in emergency supply pre-positioning for disaster relief, to our knowledge, no existing work considers the design or dynamic alteration beyond a single change of a disaster relief inventory supply chain. Furthermore, to our knowledge, there does not exist any research on the operations of mobile inventory supply chains for disaster relief. With this work, we seek to address these literature gaps by examining how to optimally manage a mobile supply inventory to prepare for and respond to disasters over time.

2.2 Dynamic Facility Location Problem

Closely related to the problem of relocating a mobile inventory for responding to disasters over time is the stream of literature on the dynamic facility location problem. The dynamic facility location problem is the problem of locating, and possibly relocating, facilities over time in response to changing system parameters, such as demand or distribution costs. Finding an optimal solution to this problem requires finding a balance between the costs of establishing new facilities or changing existing ones and the benefits gained from making these decisions. Rather than finding decisions which are optimal for only the current state of the system, these models seek to find optimal decisions which are robust to changes in parameters over time. For more comprehensive reviews of dynamic facility location literature, the reader is referred to Arabani and Farahani (2012), Farahani et al. (2009-2), Snyder (2006), and Owen and Daskin (1998).

For our purposes, research on the dynamic facility location problem can be divided into two categories: that which considers relocation and that which does not.

There are many papers in the dynamic facility location literature that do not consider the possibility of relocating facilities. Among these are papers which only consider when and where to open facilities and do not allow for the closing or relocating of facilities (e.g. Ghaderi and Jabalameli (2013), Current et al. (1997), Melachrinoudis et al. (1995), Jornsten and Bjorndal (1994), Shulman (1991), and Erlenkotter (1981)). These models typically assume that once a facility is opened at a particular location, it remains open for the duration of the planning horizon. Also among the papers which do not consider the possibility of relocating are those which consider when and where to both open and close facilities but do not model these decisions in a way that can represent relocation. Some of these papers either require that once a facility is established, it cannot be removed, or once a facility that was established previous to the planning horizon is closed, it cannot be reopened, and thus the models described in these papers clearly cannot apply to relocation decisions (e.g. Hinojosa et al. (2008) and de Gama and Captivo (1998)). Other papers in this area use binary variables to model the capacity levels or operating status of each potential facility for each time period (e.g. Jean et al. (2015), Thanh et al. (2008), Dias et al. (2007), Romauch and

Hartl (2005), Canel et al. (2001), Melachrinoudis and Min (2000), and Sweeney and Tatham (1976)). The models considered in these papers do not connect the decisions of closing and opening facilities at different sites sufficiently to be able to represent facility relocation.

There are, however, many papers which do consider the relocation of facilities. As many authors have noted (e.g. Farahani et al. (2009-2), Owen and Daskin (1998), and Current et al. (1997)), dynamic facility location problems are computationally complex and thus most research in this area has been limited to deterministic problems. Thus, among the dynamic facility location papers which consider facility relocation, most assume demand is deterministic or that it is appropriate to use a deterministic proxy for uncertain demand (e.g. Qiu and Sharkey (2013), Halper and Raghavan (2011), Farahani et al. (2009-1), Melo et al. (2005), Gue (2003), Drezner and Wesolowsky (1991), Campbell (1990), Chand (1988), Wesolowsky and Truscott (1975), Wesolowsky (1973), Ballou (1968)). However, for our purposes, a deterministic proxy for uncertain demand is not appropriate as disasters and subsequent demand for relief supplies are far too unpredictable. A useful model must take into account uncertainties in when, where, and in what amounts demand for supplies will occur.

Qiu and Sharkey (2013) analyze a problem similar to those we consider in this dissertation in the military logistics setting. Motivated by recent military interest in the capability of sea basing to serve as the logistical hub during military operations, they develop a dynamic programming (DP) model for finding the optimal location plan and inventory plan for a single mobile facility to satisfy demand over a finite horizon. They propose algorithms for solving their model with the objective of minimizing costs with and without capacity and service constraints. However, they do not consider stochastic demand and thus their model is not applicable to the disaster relief setting. This paper does, though, point out another relevant application area of the models presented in this dissertation: the operations of sea bases for military operations. The reader is refereed to Qiu and Sharkey (2013) for further information regarding this application. A related application is to the operations of reserve military equipment between deployments such as the National Defense Reserve Fleet which serves as a reserve of ships for national defense and national emergencies.

A few papers in the dynamic facility location literature which consider relocation consider stochastic demand, however, to our knowledge, no paper models demand as evolving according to a non-stationary discrete time Markov chain (DTMC).

Berman and Odoni (1982) study a single-facility location problem similar to the p-median problem where travel times are stochastic and described by the state of the network at any given time and the facilities (e.g., an ambulance) can be relocated as travel times change. The state of the network changes according to a stationary ergodic Markov chain. Their model seeks to find the optimal set of server locations for each network state. They develop a heuristic to solve the problem which, in each iteration, goes through each system state and tries to improve the decision for this state while keeping all others unchanged. The heuristic terminates when the decisions do not change during an iteration. They assume that the relocation cost function is nondecreasing and concave and that demand occurs in each period at a certain location according to a Bernoulli distribution. For reasons noted above, this simple way of modeling potential demand does not fully represent the complexity of the problem at hand.

Chow and Regan (2011) study the dynamic facility location problem applied to air tanker response to wildland fires. They minimize the cost and time of deploying the tankers to fires. They assume that simultaneously occurring fires are unlikely but that a single fire may require multiple air tankers and thus their model requires that a certain number of the closest facilities be able to cover a demand node. They assume demand follows an autoregressive process and that demand for different nodes are independent of each other. They simplify the problem by making relocation decisions independent of future relocation decisions as the model would otherwise be “too complex” and require DP to solve. Thus, each relocation decision is modeled by a separate MIP p-median-type model that determines the optimal single relocation of the servers to serve at least 90% of the expected demand over a short rolling horizon and use branch and bound to solve it. Thus, their approach finds a myopic policy and does not consider the full complexity of the problem at hand.

Rosenthal et al. (1978) study the relocation over an infinite horizon and discrete set of locations of a facility which serves customers that relocate and change cost parameters according to a stationary Markov chain. They claim to be the first paper to introduce methods of stochastic decision processes into location analysis. They develop heuristics for finding the pure stationary policy for both the single-customer and the multi-customer cases.

As we discuss in Section 3.1, it is most appropriate in the disaster relief setting to model demand as evolving according to a non-stationary DTMC. To our knowledge, no paper on dynamic facility location models demand as evolving according to a non-stationary DTMC, nor has dynamic facility location with relocation been studied in the disaster relief setting. Furthermore, to our knowledge, no paper considers dynamic facility relocation and inventory management with stochastic demand. This work seeks to address these literature gaps as well as those described in Section 2.1.

3 Dynamic Facility Relocation and the Value of Inventory Mobility in Disaster Relief

In this section, we examine the dynamic facility relocation problem and the value of inventory mobility in disaster relief. This section is organized as follows. Section 3.1 presents the model and analytical results on optimal relocation policies and the value of inventory mobility. Section 3.2 presents results on the optimal path and speed of relocation decisions using a special case of the model. Section 3.3 describes the Base State Heuristic (BSH) designed to solve the problem as modeled in Section 3.1. Section 3.4 presents numerical results regarding the problem as modeled in Section 3.1. The proofs of the results in this section can be found in Section 3.5.

3.1 Dynamic Relocation Model

In relocating a mobile inventory, we make sequential decisions over time in response to changing disaster forecasts. Specifically, in each period we observe demand which informs our next period forecast and decide where to move the inventory for the following period. We model this sequential decision making problem using dynamic programming (DP). As noted in Section 1, the real and potential needs of different locations may change over time and thus are non-stationary. Furthermore, demand forecasts may depend on current demand realizations and thus exhibit the Markovian property. For example, tsunamis are more likely following a large earthquake, the current weather and weather forecast gives information on whether or not there will be a hurricane in the next period, and demand in a period immediately following a disaster and subsequent demand spike will likely also be high due to ongoing relief operations. Thus, it is most appropriate to model potential disaster site demands as evolving according to a non-stationary discrete time Markov chain (DTMC). We assume a finite horizon of T time periods as it unrealistic to forecast disasters infinitely into the future or to assume that forecast evolution will be stationary. Additionally, we assume that the inventory has infinite capacity, all demand must be satisfied, and our objective is to minimize cost. Let $d(\cdot, \cdot)$ be the Euclidean distance metric. Our notation is as follows:

J :	set of potential disaster sites
I :	set of potential inventory locations
D_j :	finite set of possible demand levels at site $j \in J$
D :	set of possible demand vectors $d = (d^1, ..., d^{\|J\|})$, where $d^j \in D_j \ \forall j \in J$, describing the demand at all $\|J\|$ potential disaster sites
l_j :	location of site $j \in J$
$f(d(i, i'))$:	cost to move the inventory from $i \in I$ to $i' \in I$ in one period
$g_j(d(i, l_j), d^j)$:	cost to serve $d^j \in D_j$ units of demand at site $j \in J$ from $i \in I$
$p_t^j(d_t^j, d_{t+1}^j)$:	probability site $j \in J$ demand will transition from $d_t^j \in D_j$ to $d_{t+1}^j \in D_j$ after period $t \in \{0, ..., T-1\}$

$p_t(d_t, d_{t+1})$: probability the demand vector will transition from $d_t \in D$ to $d_{t+1} \in D$ after period $t \in \{0, ..., T-1\}$

While a few of our results specify specific forms of I, in general, I may be continuous or discrete. Let $G(i, d) = \sum_{j \in J} g_j(d(i, l_j), d^j)$ be the cost to serve the demand described by demand vector $d \in D$ from $i \in I$. Note that each demand vector $d \in D$ describes the demand at all $|J|$ potential disaster sites in a potential demand realization, $D = D_1 \times D_2 \times ... \times D_{|J|}$, and f may include maintenance or operational costs. Assume that $g_j(\cdot, 0) = 0 \ \forall j \in J$ and $\sum_{d_{t+1} \in D} p_t(d_t, d_{t+1}) = 1 \ \ \forall \ d_t \in D$ and $t \in \{0, ..., T-1\}$. Furthermore, assume that $p_t(d_t, d_{t+1})$ is some function q of the individual transition probabilities, i.e. $p_t(d_t, d_{t+1}) = q\left(p_t^1(d_t^1, d_{t+1}^1), ..., p_t^{|J|}(d_t^{|J|}, d_{t+1}^{|J|})\right) \ \forall \ d_t, d_{t+1} \in D$ and $t \in \{0, ..., T-1\}$; while not necessary for the results presented in this dissertation, this assumption makes the problem more tractable by allowing solution algorithm implementations to store only a subset of the $T \times |D| \times |D|$ transition probabilities.

We can find the minimum cost-to-go of being in location i_t when the demand vector is d_t in time period $t \ \forall \ (i_t, d_t) \in I \times D, \ t \in \{0, 1, ..., T\}$ using the following DP equations:

$$V_t(i_t, d_t) = G(i_t, d_t) + \min_{i_{t+1} \in I} \left\{ f(d(i_t, i_{t+1})) + \sum_{d_{t+1} \in D} p_t(d_t, d_{t+1}) V_{t+1}(i_{t+1}, d_{t+1}) \right\} \tag{3.1}$$

$$V_T(i_T, d_T) = G(i_T, d_T)$$

The cost-to-go is the cost to serve the current period demand plus the cost to relocate the inventory for the following period plus the expected future period cost. To find the value of inventory mobility in disaster relief in terms of cost savings, we also need to define the cost of the traditional warehouse pre-positioning system. The DP equations to represent the total cost of responding to disasters over time using a traditional warehouse pre-positioning system with a single inventory are as follows:

$$\bar{V}_t(\bar{i}, d_t) = G(\bar{i}, d_t) + \sum_{d_{t+1} \in D} p_t(d_t, d_{t+1}) \bar{V}_{t+1}(\bar{i}, d_{t+1}) \tag{3.2}$$

$$\bar{V}_T(\bar{i}, d_T) = G(\bar{i}, d_T)$$

Note that in most cases, $\bar{i}$ will be chosen as the optimal stationary position, i.e. $\bar{i} = \arg\min_{i \in I} \bar{V}_0(i, d_0)$. Note that if $\bar{i}$ is subject to optimization in this way in the traditional inventory system, it may be most appropriate when calculating the value of inventory mobility to allow i_0 to be subject to optimization in the mobile inventory system as well. If we define some additional notation, we can express the expected cost of responding to disasters over time using a traditional warehouse pre-positioning system represented by the DP equations 3.2 using a single equation. Let P^t be the full transition matrix between demand states in period t and let P_i^t be the ith row of P^t. Without loss of generality, let d_0 be the first demand state, that is, the demand state that has its probabilities of transitioning to other demand states specified by the first row of $P^t \ \forall t$. Furthermore, let $\hat{d}$ be a column vector representing all possible demand states with elements in the same order as the rows and columns of P^t and $G(i, \hat{d})$ be the column vector for which each element is $G(i, d)$ where d is the corresponding element of $\hat{d}$. Furthermore, define $p^t(d_0, d)$ to be the probability that the demand vector is $d \in D$ in period $t \in \{1, ..., T\}$ given that the demand vector in period 0 is d_0, i.e. $p^1(d_0, d) = P^0_{1,d}$ and $p^t(d_0, d) = [P^0 \prod_{k=1}^{t-1} P^k]_{1,d}$ where Q_{ij} is the $(i, j)^{\text{th}}$ entry of the matrix Q.

Theorem 1. *The expected cost of responding to disasters over time using a traditional warehouse pre-positioning system with a single inventory located at $\bar{i} \in I$ represented by the DP equations 3.2 can be written as one equation as follows:*

$$\bar{V}_0(\bar{i}, d_0) = G(\bar{i}, d_0) + P^0_1 \left[I + \sum_{t=2}^{T} \prod_{k=1}^{t-1} P^k \right] G(\bar{i}, \hat{d}) \tag{3.3}$$

$$= G(\bar{i}, d_0) + \sum_{d \in D} \left[G(\bar{i}, d) \sum_{t=1}^{T} p^t(d_0, d) \right] \tag{3.4}$$

Note that this is a simple linear function of problem parameters. Also note that equation 3.4 gives an intuitive view of the expected cost of responding to disasters over time using traditional, non-mobile inventory. That is, the cost is simply the sum over all possible demand vectors of the cost to serve that demand vector from the stationary inventory location $\bar{i}$ times the expected amount of time that the demand will be as described by that demand vector.

Recall that we focus on the value of inventory mobility associated with the change in cost resulting from decreased transportation distances to affected communities in using mobile rather than traditional pre-positioned inventory. That is, we define the **value of inventory mobility** as the optimal cost of responding to disasters over time using a traditional pre-positioned inventory minus the optimal cost using a mobile inventory:

$$\min_{\bar{i} \in I_s} \bar{V}_0(\bar{i}, d_0) - \min_{i_0 \in I} V_0(i_0, d_0)$$

where $I_s \subseteq I$; for example, for a supply holding ship, I may be a large body of water while I_s is the coastline, or parts of the coastline, where a stationary warehouse can be located. The value we have defined here is the operating cost savings over T time periods; this value can be used to determine the payback period of a mobile inventory system investment. We can prove the following:

Theorem 2. *If the cost to allow the mobile inventory to remain in the same location is* 0*, i.e.* $f(0) = 0$*, then the value of inventory mobility is greater than or equal to* 0*.*

We can also prove the following regarding the optimal stationary position:

Theorem 3. *Assume $I = \mathbb{R}$, $l_j \in \mathbb{R}\ \forall j \in J$, and $g_j(d(\cdot, \cdot), d^j)$ is linear and non-decreasing in $d(\cdot, \cdot)\ \forall\ j \in J$. Then $\exists$ an optimal solution $\bar{i}^*$ minimizing $\bar{V}_0(\bar{i}, d_0)$ such that $\bar{i}^* = l_j$ for some $j \in J$.*

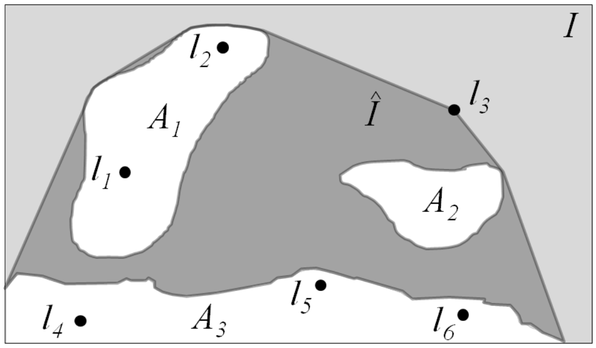

Figure 1: Example network satisfying Assumption 1. I consists of the shaded regions, the white regions A_1, A_2, and A_3 are infeasible regions, $\hat{B}$ is $A_1 \cup A_2 \cup A_3$, and $\hat{I}$ is the dark gray shaded region. Under Assumption 1, Theorem 5 assures that there exists an optimal solutions to the relocation problem such that the inventory is always located within $\hat{I}$.

In other words, in the case described, Theorem 3 assures that the stationary inventory will be optimally located at one of the potential disaster sites. Thus, the optimal location can be found by simply enumerating the cost for each potential disaster site and choosing the location with the smallest cost. In this way, Theorem 3 allows us to reduce the size of the stationary problem in the one-dimensional case.

Typically, realistic instances of the dynamic relocation problem will exist in $\mathbb{R}^2$ and have a large state space; thus, they will take a long time to solve. Under a few basic assumptions, we can show that it is sufficient to consider a smaller feasible set of inventory locations and thereby reduce the size of the stationary problem and the dynamic relocation problem. Let $conv(C)$ and $Int(C)$ denote the convex hull and interior of a set C, respectively, and define the following:

Assumption 1. *Assume $I \subseteq \mathbb{R}^2$ is connected and closed and $g_j(y, d^j)$ is non-decreasing in the distance y $\forall j \in J$. Let $B = \{A_1, ..., A_l\}$ be a set of finitely many mutually disjoint, connected, closed, and bounded subsets of $\mathbb{R}^2$ such that $i \notin I$ $\forall$ $i \in Int(A_n)$ and $A_n \in B$, $conv(I \cup \{l_j | j \in J\}) \subseteq I \cup B$, and $I \cup B$ is convex.*

Furthermore, let $\hat{B} = \{A_n \in B | \{A_n \cap conv(\{l_j | j \in J\}) \neq \emptyset\} \vee \{A_n \cap conv(\hat{B} \setminus A_n) \neq \emptyset\}\}$ and $\hat{I} = I \cap conv(\hat{B} \cup \{l_j | j \in J\})$. See Figure 1 for an example network satisfying Assumption 1. Assumption 1 defines the set B of all infeasible regions for the mobile inventory, or areas where the inventory cannot be located (e.g. land in the case of inventory on a ship). We prove that it is sufficient for both the stationary and mobile inventory problems to consider only inventory locations within $\hat{I}$, the feasible inventory locations within the convex hull of the potential disaster sites extended to include any overlapping infeasible regions (see the dark gray shaded region in Figure 1 for an example). Let $proj_C(x)$ be the projection of a point x onto a closed set C and note the following lemma:

Lemma 1. *Let $C \subset \mathbb{R}^2$ be a closed, convex set and $x \in \mathbb{R}^2$. Then $d(proj_C(x), z) \leq d(x, z)$ $\forall z \in C$.*

Using Lemma 1, we can prove the following which details the result for the stationary system:

Theorem 4. *Under Assumption 1, $\exists$ an $\bar{i}^* \in \hat{I}$ minimizing $\bar{V}_0(\bar{i}, d_0)$.*

Thus, under the assumptions listed, Theorem 4 allows us to reduce the size of the stationary problem by reducing the number of inventory locations that we must consider. If I is convex, then Theorem 4 assures that there exists an optimal stationary inventory location within the convex hull of the potential disaster sites.

Note that in the dynamic relocation problem, the state space is all possible inventory locations cross all possible demand vectors, $I \times D$. Even relatively small problem sizes have a large state space and thus take a long time to solve to optimality. For this reason, reducing the size of this state space for the dynamic relocation problem is worthwhile. Using the following lemma, we can prove a similar result to Theorem 4 for the mobile inventory system.

Lemma 2. *Let $C \subset \mathbb{R}^2$ be a closed, convex set and $x, y \in \mathbb{R}^2$. Then $d(proj_C(x), proj_C(y)) \leq d(x, y)$.*

We also need some additional notation. Let $\pi = \{i_0^\pi, ..., i_T^\pi\} \in \Pi$ be a policy consisting of an initial location i_0^π and a sequence of functions i_{t+1}^π that map states (i_t, d_t) into the decision of where to move the mobile inventory for the following period, Π be the set of all feasible policies for 3.1, and the cost of the mobile inventory system under policy π be represented by $V_0^\pi(i_0^\pi, d_0)$ where

$$V_t^\pi(i_t, d_t) = G(i_t, d_t) + f(d(i_t, i_{t+1}^\pi(i_t, d_t))) + \sum_{d_{t+1} \in D} p_t(d_t, d_{t+1}) V_{t+1}^\pi \left(i_{t+1}^\pi(i_t, d_t), d_{t+1}\right)$$

$$V_T^\pi(i_T, d_T) = G(i_T, d_T)$$

With this notation and Lemma 2, we can prove the following:

Theorem 5. *Assume the cost to move $f(y)$ is non-decreasing in the distance y. Under Assumption 1, $\exists$ an optimal policy $\pi^* = \{i_0^{\pi^*}, ..., i_T^{\pi^*}\} \in \Pi$ such that $i_0^{\pi^*} \in \hat{I}$ and $i_{t+1}^{\pi^*}(i_t, d_t) \in \hat{I}$ $\forall$ $i_t \in \hat{I}$, $d_t \in D$, and $t \in \{0, ..., T-1\}$.*

Theorem 5 assures that, under the assumptions listed, there exists an optimal policy such that the mobile inventory will always remain within $\hat{I}$. If I is convex, then Theorem 5 assures that there exists an optimal policy such that the mobile inventory will always remain within the convex hull of the potential disaster sites.

We now characterize the value of inventory mobility to provide managerial insights on when implementing a mobile inventory system is worthwhile. The following result concerns the effect of the movement cost function f on the value:

Theorem 6. *Assume $\hat{f}(y) \geq f(y)$ $\forall y \geq 0$. Then the value of inventory mobility is weakly greater in a system with movement cost f than in the same system with movement cost $\hat{f}$.*

Theorem 6 states that the value is greater when the movement cost is lower. This is intuitive as the cost of moving is lower while the benefit of moving remains the same. Thus, a mobile inventory system with high movement costs will likely not be worthwhile. Using similar logic, we may expect that the value is greater when the cost to serve function, G, is greater as this may increase the benefit of moving closer to potential disaster sites. However, Example 1 shows that even in comparing identical systems where one has a cost to serve function which is an increasing monotonic transformation of the other (meaning that ordering is preserved), the value of inventory mobility may be lower with a greater cost to serve function.

Example 1. *Consider two potential disaster sites with locations l_1 and l_2, respectively, and deterministic demand d_t, $t \in \{0, 1\}$. Let $T = 1$, $I = I_s = \{l_1, l_2\}$, $f(y) = 1\ \forall y$, and G and $\hat{G}$ be as follows:*

$G(i, d_t)$	$i = l_1$	$i = l_2$
$t = 0$	0	3
$t = 1$	3	0

$\hat{G}(i, d_t)$	$i = l_1$	$i = l_2$
$t = 0$	1	3
$t = 1$	3	1

With cost to serve G, a stationary inventory will be located at either l_1 or l_2 at cost 3 and a mobile inventory will be located at l_1 for $t = 0$ and l_2 for $t = 1$ at cost 1. The value of inventory mobility is $3 - 1 = 2$. With movement cost $\hat{G}$, the stationary and mobile inventory policies are the same, but the value is now $4 - 3 = 1 < 2$.

Example 1 suggests that for the value of inventory mobility to be greater when G is greater, G must be greater in a sense which not only preserves order but which also ensures relative differences between two function evaluations weakly increase. Theorem 7 characterizes such a transformation.

Theorem 7. *Assume $G(i, d) \geq 0\ \forall i \in I$ and $d \in D$, $f(0) = 0$, and $f(y) \geq 0\ \forall y \geq 0$. Let $\hat{G}(i, d) = aG(i, d) + b$ where $a \geq 1$, $b \geq 0$. Then the value of inventory mobility is weakly greater in a system with cost to serve $\hat{G}$ than in the same system with cost to serve G.*

Thus, if it is possible to decrease the cost to serve demand from a stationary inventory, then a mobile inventory system is less worthwhile; on the other hand, as the cost to serve increases, a mobile inventory system may become a more worthwhile investment as long as relocation costs are not similarly affected. We further explore the sensitivity of the value of inventory mobility to other parameters in Section 3.4.

We can also find bounds on the value of inventory mobility. Using equations 3.1, 3.3, and 3.4 we arrive at the following result.

Theorem 8. *Assuming $i^*(d) = \arg\min_{i \in I} G(i, d)$ exists $\forall\ d \in D$, the value of inventory mobility in disaster relief when the disaster relief system is as represented in 3.1 is bounded above by*

$$G(\bar{i}, d_0) - G(i^*(d_0), d_0) + P_1^0 \left[I + \sum_{t=2}^{T} \prod_{k=1}^{t-1} P^k \right] \left[G(\bar{i}, \hat{d}) - G(i^*(\hat{d}), \hat{d}) \right]$$

$$= G(\bar{i}, d_0) - G(i^*(d_0), d_0) + \sum_{d \in D} \left[\left[G(\bar{i}, d) - G(i^*(d), d) \right] \sum_{t=1}^{T} p^t(d_0, d) \right] \tag{3.5}$$

where $i^*(\hat{d})$ is the column vector for which each element is the value of i^* evaluated at the corresponding element of $\hat{d}$ and $G(i^*(\hat{d}), \hat{d})$ is the column vector for which each element is $G(i^*(d), d)$ where d is the corresponding element of $\hat{d}$.

This result is due to the fact that the value of inventory mobility is bounded above by the difference between the cost when using traditional warehousing and the cost of being able to be in the best position possible in each period. This bound will be tighter when movement costs are relatively low and when demands are relatively easier to predict. If the bound is less than or equal to the cost of investing in inventory mobility, then this bound can be used to justify the use of traditional warehousing over mobile inventory for disaster relief. Bound 3.5 has a similar interpretation to equation 3.4. That is, the value if inventory mobility is bounded above by the sum over all possible demand states of the difference in cost to serve that demand state from the stationary inventory position and the cost to serve that demand state from the best position possible times the expected proportion of time that the demand will be in that particular state.

We can also find a lower bound on the value of inventory mobility in disaster relief. We note that any policy feasible for the DP equations 3.1 will provide an upper bound on the optimal minimum cost for those DP equations. Since the value of inventory mobility is calculated by subtracting the optimal cost of the mobile inventory system (equations 3.1) from the optimal cost of the traditional inventory system (equations 3.3), any feasible policy will give us a lower bound on the value of inventory mobility. With this logic, we find the following theorem.

Theorem 9. *The value of inventory mobility in disaster relief when the disaster relief system is as represented in 3.1 is bounded below by*

$$\bar{V}_0(\bar{i}, d_0) - V_0^\pi(i_0, d_0) \tag{3.6}$$

where $\pi = \{i_0^\pi, ..., i_T^\pi\} \in \Pi$.

If $\pi \in \hat{\Pi}$ where $\hat{\Pi}$ is the set of all feasible policies for 3.1 that specify a single path which is independent of demand realizations and $\hat{i}_t^\pi$ is the position of the inventory in period t under policy π then 3.6 can be written explicitly as

$$\begin{aligned} &\bar{V}_0(\bar{i}, d_0) - G(\hat{i}_0^\pi, d_0) - P_1^0 G(\hat{i}_1^\pi, \hat{d}) - P_1^0 \sum_{t=1}^{T-1} \prod_{k=1}^{t-1} P^k G(\hat{i}_{t+1}^\pi, \hat{d}) - \sum_{t=0}^{T-1} f(d(\hat{i}_t^\pi, \hat{i}_{t+1}^\pi)) \\ &= \bar{V}_0(\bar{i}, d_0) - G(\hat{i}_0^\pi, d_0) - \sum_{t=0}^{T-1} \sum_{d \in D} p^t(d_0, d) G(\hat{i}_t^\pi, d) - \sum_{t=0}^{T-1} f(\hat{i}_t^\pi, \hat{i}_{t+1}^\pi) \end{aligned} \tag{3.7}$$

where $\bar{V}_0(\bar{i}, d_0)$ is given by equations 3.3 or 3.4.

For an example of a feasible policy for 3.1 that specifies a single path which is independent of demand realizations, consider the open-loop myopic policy π defined as

$$\hat{i}_0^\pi = \arg\min_{i_0 \in I} G(i_0, d_0) \tag{3.8}$$
$$\hat{i}_{t+1}^\pi = \arg\min_{i_{t+1} \in I} \left\{ f(d(\hat{i}_t^\pi, i_{t+1})) + \sum_{d \in D} p^t(d_0, d) G(i_{t+1}, d) \right\}, \quad t = 0, ..., T-1$$

Bound 3.6 will be tight if π is a policy which is close to optimal. Bound 3.7 will be tight when it is not especially beneficial to take into account future period costs when making movement decisions, for example, when movement costs are relatively low or when the demand forecasts do not have a lot of variation. If the lower bound found in Theorem 9 is greater than or equal to the cost of investing in inventory mobility, then this bound can be used to justify such an investment.

Recognizing that $i^*(d_0) = \arg\min_{i \in I} G(i, d) = \hat{i}_0^\pi$, the gap between the upper bound given by equation 3.5 and the lower bound given by equation 3.7 is

$$\sum_{t=1}^{T} \sum_{d \in D} \left[p^t(d_0, d) \left[G(\hat{i}_t^\pi, d) - G(\hat{i}^*(d), d) \right] \right] + \sum_{t=1}^{T-1} f(\hat{i}_t^\pi, \hat{i}_{t+1}^\pi)$$

This optimality gap will be tight when movement costs are relatively low or when future demand states are relatively easy to predict.

3.2 Path and Speed of Relocation Decisions

In Section 3.1, we introduced our model for optimally relocating a mobile inventory over time in response to changing demand forecasts. This model focuses on the high level decisions of where to locate the mobile inventory over time; also important in managing a mobile inventory are the lower level decisions of what path and at what speed to move the inventory between those locations. The general model does not address these decisions directly but rather assumes a given value of the cost to move function which depends on these decisions. In this section, we analyze these path and speed decisions using a special case of the relocation model. The results of this section can be used to define the cost to move function f in the general model. Furthermore, as we see at the end of the section, this special case also provides some intuition on the optimal relocation decisions.

Given a relocation policy and assuming an adequate stock of fuel, the path and speed decisions for how to move between a pair of sites can be determined independently from the corresponding decisions for another pair of sites. Thus, consider the problem of determining the best path and speed for moving between two particular sites. Let T be the length of a single period in the original relocation problem, let the inventory be located at the first site of the pair of sites in period '0,' fix the location to the second site in period 'T,' and let f be the cost to move the inventory in the corresponding smaller increments of time. With this set up, the optimal path and speed decisions can be determined by a special case of the model described by the DP equations

$$V_t(i_t) = \min_{i_{t+1} \in I} [\, f(d(i_t, i_{t+1})) + V_{t+1}(i_{t+1})]$$

$$V_{T-1}(i_{T-1}) = f(d(i_{T-1}, i_T)) + V_T(i_T)$$

$$V_T(i_T) = \sum_{j \in J} p_j g_j(d(i_T, l_j))$$

For simplicity, we let p_j be the probability a disaster will occur at site $j \in J$ at time T and let each potential disaster location have only one possible demand level. Accordingly, we omitted the demand level argument of the cost to serve functions g_j. These simplifications have no affect on the intermediate movement decisions. When f is convex and non-decreasing, assumptions that hold in most applications of the problem, we can prove the following:

Theorem 10. *If f is convex and non-decreasing in the distance moved and I is convex, then the minimum cost to move the inventory from i_0 to i_T in T periods is $Tf\,(d(i_0, i_T)/T)$ and it is optimal to move $\frac{1}{T}d(i_0, i_T)$ toward i_T in each period.*

Given the final location of the inventory i_T, Theorem 10 states that the optimal movement plan is to move the inventory as slowly as possible along the direct path from i_0 to i_T. In relocating the inventory between two locations in any time period, it is optimal to move the inventory at the slowest speed possible along the direct path between the two locations.

In some applications, there may be a fixed cost $\bar{f}$ (e.g. truck rental cost, operator wages, etc.) to move the inventory which is incurred only if the inventory is moved a positive distance. With a fixed cost, it is no longer necessarily optimal to move in every period; there is a trade off between moving as slowly as possible and not moving in every period to avoid the fixed cost. However, we can prove that once the number of periods in which to move is determined, the inventory will be moved as slowly as possible over that number of time periods. We will use the following lemma in the proofs of Theorems 11 and 13.

Lemma 3. $h(\tau) := \tau f\left(\frac{\bar{d}}{\tau}\right)$ *is convex in* τ

Theorem 11. *If f is convex and non-decreasing in the distance moved, I is convex, and there is a fixed cost $\bar{f}$ to move the inventory a positive distance in a given period, then the minimum cost to move the inventory from i_0 to i_T in T periods is*

$$C_f(\tau^*) := \tau^* \bar{f} + \tau^* f\,(d(i_0, i_T)/\tau^*) + (T - \tau^*) f(0)$$

where τ^* is the minimum between T and the smallest $\tau \in \mathbb{Z}^+$ such that

$$C_f(\tau) \le C_f(\tau + 1) \tag{3.9}$$

and it is optimal to move $\frac{1}{\tau^*}d(i_0, i_T)$ towards i_T in τ^* of the periods and to not move in the other periods.

For example, if $f(d(i,i')) = a_1(d(i,i'))^2 + a_3$ where $a_1 > 0$ and $I = \mathbb{R}^2$, then the cost to move is $C_f(\tau^*) = \tau^* \bar{f} + a_1[(x_T - x_0)^2 + (y_T - y_0)^2]/\tau^* + Ta_3$ where τ^* is the smallest $\tau \in \mathbb{Z}^+$ such that $\frac{1}{\tau(\tau+1)} \leq \bar{f}/[(x_T - x_0)^2 + (y_T - y_0)^2]$.

Theorems 10 and 11 state that, under a few basic assumptions, it is optimal to move the inventory as slowly as possible along the direct path from i_0 to i_T and that a fixed cost shortens the amount of time over which to move the inventory at a constant, slow rate. The following Theorems 12 and 13 state that these results also hold in the continuous time setting, the most appropriate setting for the operational level path and speed decisions. Define the variable cost to move the inventory at a rate of r for one unit of time to be $f_c(r)$ and the fixed cost per unit of time of moving the inventory at a positive rate to be $\bar{f}_c$. The following theorems are the continuous analogs of Theorems 10 and 11, respectively.

Theorem 12. *If f_c is convex and non-decreasing in the movement rate and I is convex, then the minimum cost to move the inventory from i_0 to i_T in T time units is $Tf_c\left(d(i_0, i_T)/T\right)$ and it is optimal to move directly toward i_T at a rate of $d(i_0, i_T)/T$ $\forall t \in [0, T]$.*

That is, given the final location of the inventory at the end of the planning horizon and a convex cost to move function, it is optimal to move the inventory as slowly as possible along the direct path from i_0 to i_T. Considering the fixed cost,

Theorem 13. *If f_c is convex and non-decreasing in the movement rate, I is convex, and there is a fixed cost $\bar{f}_c$ per unit of time of moving the inventory at a positive rate, then the minimum cost to move the inventory from i_0 to i_T in T time units is*

$$C_{f_c}(t^*) := t^*\bar{f}_c + t^* f_c\left(d(i_0, i_T)/t^*\right) + (T - t^*)f_c(0)$$

where t^* is the minimum of $C_{f_c}(t)$ over $[0, T]$ and it is optimal to move directly toward i_T at a rate of $d(i_0, i_T)/t^*$ $\forall t \in [0, t^*]$.

As in the discrete time setting, due to the fixed cost, it is no longer necessarily optimal to move continuously throughout the planning horizon; that is, the fixed cost introduces a trade off between moving as slowly as possible and not moving constantly to avoid the fixed cost. Once the length of time in which to move, t^*, is determined, however, the inventory will be moved as slowly as possible over t^* time units and remain stationary for the remaining $T - t^*$ time units.

To gain some intuition on the optimal relocation decisions of the original problem, let us further explore the utility of Theorem 10 through a few examples where we no longer fix the location of the mobile inventory in period T. First, consider the case where $I = \mathbb{R}$, that is, when the mobile inventory's movement is restricted to one dimension, and the $|J|$ potential disaster sites are also in one dimension. In this case, $i_t = x_t \in \mathbb{R}$ and $l_i \in \mathbb{R}$ and the distance measure is $d(x, x') = |x' - x|$. Furthermore, let us assume that the cost functions f and g_j take the following form:

$$f(d(x, x')) = a_1(d(x, x'))^2 + a_3 \tag{3.10}$$

$$g_j(d(x, l_j)) = b_{1j}(d(x, l_j))^2 + b_{3j} \quad \forall\ j \in J \tag{3.11}$$

where $a_1 > 0$ and $\frac{1}{T}a_1 + \sum_{j \in J} p_j b_{1j} > 0$. We assume f and g_j are quadratic in the distance to reflect possible nonlinear fuel costs or penalties for serving demand from long distances due to lead time considerations. The DP equations can be rewritten as follows,

$$V_t(x_t) = \min_{x_{t+1} \in \mathbb{R}} \left[\, a_1(x_{t+1} - x_t)^2 + a_3 + V_{t+1}(x_{t+1}) \right] \tag{3.12}$$

$$V_T(x_T) = \sum_{j \in J} p_j \left[b_{1j}(x_T - l_j)^2 + b_{3j} \right] \tag{3.13}$$

It is cumbersome to solve these DP equations directly using the DP algorithm. However, noting that there are only inventory movement costs until the end of the horizon, the problem can be simplified by Theorem 10:

Corollary 1. *If the cost function f takes the form of 3.10 where $a_1 > 0$ and $I = \mathbb{R}$, then the minimum cost to get from a position x_0 to a position x_T in T periods is*

$$\frac{1}{T}a_1(x_T - x_0)^2 + Ta_3$$

and it is optimal to move $\frac{1}{T}^{th}$ of the total distance $|x_T - x_0|$ in each period.

Using Corollary 1 and DP equations 3.12 and 3.13, the optimal cost of this one dimensional version of the model can be written as

$$V_0(x_0) = \min_{x_T \in \mathbb{R}} \left[\frac{1}{T}a_1(x_T - x_0)^2 + Ta_3 + \sum_{j \in J} p_j \left[b_{1j}(x_T - l_j)^2 + b_{3j} \right] \right]$$

Note that the problem has been reduced to the minimization of a quadratic function over a single variable, x_T. Solving the first and second order conditions, we arrive at the following result regarding the optimal ending position of the inventory, x_T^*. Note that together with the specification stated in Corollary 1 that it is optimal to move $\frac{1}{T}^{\text{th}}$ of the total distance $|x_T^* - x_0|$ in each period, the specification of x_T^* will sufficiently describe the optimal policy for this one dimensional version of the model.

Theorem 14. *If the cost functions f and g_j take the form of 3.10 and 3.11 where $a_1 > 0$ and $\frac{1}{T}a_1 + \sum_{j \in M} p_j b_{1j} > 0$ and $I = \mathbb{R}$, then the optimal ending position of the inventory for period T is*

$$x_T^* = \frac{\frac{1}{T}a_1 x_0 + \sum_{j \in J} p_j b_{1j} l_j}{\frac{1}{T}a_1 + \sum_{j \in J} p_j b_{1j}}$$

and it is optimal to move $\frac{1}{T}^{th}$ of the total distance $|x_T^* - x_0|$ towards x_T^* in each period.

Note that the optimal ending position x_T^* is a weighted average of the starting position x_0 and the $|J|$ potential disaster site locations l_j where the weight on x_0 is the quadratic coefficient term of the cost to move divided by the number of periods and the weight on l_j is the quadratic coefficient term of the cost to serve demand. If p_j or b_{1j} increases while all other problem parameters remain constant, then the weights on potential disaster site j will increase causing the optimal ending location to be closer to potential disaster site j as the coefficient on distance of the expected cost to serve demand at potential disaster site j will increase. Similarly, if T decreases or a_1 increases, then the cost to move increases and thus the optimal ending location will be closer to x_0; there is less incentive to move as the benefit from moving remains the same while the cost to move increases.

Let us now consider slightly more general cost functions. Consider the case where the cost functions f and g_j take the following form, where we now have included a linear cost term:

$$f(d(x, x_t)) = a_1(d(x, x_t))^2 + a_2 d(x, x_t) + a_3 \tag{3.14}$$

$$g_j(d(x, l_j)) = b_{1j}(d(x, l_j))^2 + b_{2j} d(x, l_j) + b_{3j} \quad \forall\, j \in J \tag{3.15}$$

where $a_1 > 0$, $a_2 > 0$, and $\frac{1}{T}a_1 + \sum_{j \in J} p_j b_{1j} > 0$. Assuming that the cost functions f and g_j take the form of 3.14 and 3.15, the DP equations can be rewritten as follows:

$$V_t(x_t) = \min_{x_{t+1}} \left[\, a_1(x_{t+1} - x_t)^2 + a_2|x_{t-1} - x_t| + a_3 + V_{t+1}(x_{t+1})\right] \tag{3.16}$$

$$V_T(x_T) = \sum_{j \in J} p_j \left[b_{1j}(x_T - l_j)^2 + b_{2j}|x_T - l_j| + b_{3j}\right] \tag{3.17}$$

The functions over which these DP equations minimize are not everywhere differentiable as DP equations 3.12 and 3.13 are. Thus, approaching the DP equations directly to solve the problem would be even more cumbersome. However, again by Theorem 10,

Corollary 2. *If the cost function f takes the form of 3.14 where $a_1 > 0$ and $a_2 > 0$ and $I = \mathbb{R}$, then the minimum cost to move the inventory from a position x_0 to a position x_T in T periods is*

$$\frac{1}{T}a_1(x_T - x_0)^2 + a_2|x_T - x_0| + Ta_3$$

and it is optimal to move $\frac{1}{T}^{th}$ of the total distance $|x_T - x_0|$ in each period.

Using Corollary 2 and DP equations 3.16 and 3.17, the optimal cost of this one-dimensional version of the model can be written as

$$V_0(x_0) = \min_{x_T} \left[\frac{1}{T}a_1(x_T - x_0)^2 + a_2|x_T - x_0| + Ta_3 \right.$$
$$\left. + \sum_{j \in J} p_j[b_{1j}(x_T - l_j)^2 + b_{2j}|x_T - l_j| + b_{3j}]\right]$$

Note that, just as before, the problem has been reduced to the minimization of a continuous and convex function over the single variable x_T; however, due to the many linear absolute value terms, it is no longer possible to find the analytical solution for the minimizing x_T^*. It is possible, however, to solve this problem efficiently using one of the many numerical algorithms designed to efficiently find the minimum of a convex function.

We can find similar results to Corollary 1 and Theorem 14 for the two-dimensional case where $I = \mathbb{R}^2$, $l_j \in \mathbb{R}^2 \ \forall j \in J$, and $i_t = (x_t, y_t)$. For each of the $|J|$ potential disaster sites, let $l_j = (\bar{x}_j, \bar{y}_j)$. Consider the case where the cost functions f and g_j take the form of 3.10 and 3.11, respectively, where $a_1 > 0$ and $\frac{1}{T}a_1 + \sum_{j\in J} p_j b_{1j} > 0$. The DP equations can be rewritten as follows:

$$V_t(x_t, y_t) = \min_{x_{t+1}, y_{t+1}} \left[a_1 \left[(x_{t+1} - x_t)^2 + (y_{t+1} - y_t)^2 \right] + a_3 + V_{t+1}(x_{t+1}, y_{t+1}) \right] \tag{3.18}$$

$$V_T(x_T, y_T) = \sum_{j\in J} p_j \left[b_{1j} \left[(x_T - \bar{x}_j)^2 + (y_T - \bar{y}_j)^2 \right] + b_{3j} \right]$$

Notice that the cost-to-go functions are separable in the x_t and y_t variables. By Theorem 10,

Corollary 3. *If the cost function f takes the form of 3.10 where $a_1 > 0$ and $I = \mathbb{R}^2$, then the minimum cost to move the inventory from a position (x_0, y_0) to a position (x_T, y_T) in T periods is*

$$\frac{1}{T} a_1 \left[(x_T - x_0)^2 + (y_T - y_0)^2 \right] + T a_3$$

and it is optimal to move $\frac{1}{T}^{th}$ of the total distance $\sqrt{(x_T - x_0)^2 + (y_T - y_0)^2}$ in each period.

Using Corollary 3, the optimal cost of the model can be written as follows:

$$\begin{aligned} V_0(x_0, y_0) = \min_{x_T, y_T} \Bigg[& \frac{1}{T} a_1 \left[(x_T - x_0)^2 + (y_T - y_0)^2 \right] + T a_3 \\ & + \sum_{j\in J} p_j \left[b_{1j} \left[(x_T - \bar{x}_j)^2 + (y_T - \bar{y}_j)^2 \right] + b_{3j} \right] \Bigg] \end{aligned} \tag{3.19}$$

Note that the problem has been simplified to the minimization of a separable quadratic function of two variables. The separability is due to the simple quadratic structure of f and g_j which squares the Euclidean distance argument so that the resulting total cost function 3.19 is convex in x_T and y_T. Solving first and second order conditions, we arrive at the following result.

Theorem 15. *If the cost functions f and g_j take the form of 3.10 and 3.11 where $a_1 > 0$ and $\frac{1}{T}a_1 + \sum_{j\in J} p_j b_{1j} > 0$ and $I = \mathbb{R}^2$, then the optimal ending position of the inventory for period T is*

$$(x_T^*, y_T^*) = \left(\frac{\frac{1}{T}a_1 x_0 + \sum_{j\in J} p_j b_{1j} \bar{x}_j}{\frac{1}{T}a_1 + \sum_{j\in J} p_j b_{1j}}, \ \frac{\frac{1}{T}a_1 y_0 + \sum_{j\in J} p_j b_{1j} \bar{y}_j}{\frac{1}{T}a_1 + \sum_{j\in J} p_j b_{1j}} \right)$$

and it is optimal to move $\frac{1}{T}\sqrt{(x_0 - x_T^*)^2 + (y_0 - y_T^*)^2}$ towards (x_T^*, y_T^*) in each period.

Theorems 10 and 15 fully describe the optimal policy for this simple case. The optimal final location is a weighted average of the starting location and the $|J|$ potential disaster sites. If T decreases or a_1 (the scalable portion of the cost to move) increases while all other parameters remain constant, then the cost to move the same distance over the planning horizon increases and the optimal ending location will be closer to the starting position. Thus, if there is little time between relatively high expected demand periods or if the cost to move is high, we expect that the inventory may not cover as much distance in an optimal policy for the general problem as it would otherwise. If p_j or b_{1j} (the scalable portion of the expected cost to serve demand at site j) increase, then the expected cost to serve demand at site j increases causing the optimal ending location to be closer to site j. Thus, if the expected cost to serve demand at a particular site is high, the inventory will likely move closer or remain close to that site.

3.3 Base State Heuristic

In this section, we describe a solution heuristic for the general problem described in Section 3.1. We assume we have discretized I so that it is a finite set. It is reasonable to assume that the cost considerations of the problem will still be sufficiently captured as service and relocation costs do not change significantly with small changes in location.

The backwards DP algorithm can be used to calculate the cost-to-go value $V_t(i_t, d_t)$ exactly for all states and time periods. In this algorithm, one starts in period T and calculates the value of being in each state (i_T, d_T), for $i_T \in I$ and $d_T \in D$, using DP equations 3.1. Then one does the same for period $T-1$ and continues in this manner until one reaches the initial period. The DP equations 3.1 can then be used to find the optimal solution for any state and time period of where to relocate the inventory for the following period. However, the demand transition probabilities must be developed based on incomplete information and can be updated as new information becomes available; accordingly, a solution should only be used for the period in which the model was solved and the model should be re-solved each period using updated parameters. Recognizing that time periods may be relatively short, it is imperative to be able to solve the problem quickly. As we show in Section 3.4.1, however, calculating the cost-to-go values exactly using the backwards DP algorithm is prohibitively time consuming for large problem instances. Accordingly, we develop the Base State Heuristic (BSH) to estimate the cost-to-go values in a more appropriate amount of time.

Several factors are behind the excessive computational time required by the backward DP algorithm. In this algorithm, for every time period one must calculate the exact cost-to-go for every possible state (i_t, d_t). Calculating this value exactly is time consuming as it requires calculating a computationally expensive expected value over all possible transitions of the demand. Furthermore, calculating the expected value requires the transition probabilities $p_t(d_t, d_{t+1})$ for all $d_{t+1} \in D$. Since the $|D| \times |D|$ full transition matrix P^t is too big to store for any realistic problem sizes (about 34 GB of space is required to store P^t in MATLAB for a problem with 16 potential disaster sites with 2 possible demand levels each), one must calculate the transition probabilities for each cost-to-go calculation one at a time as they are needed.

The BSH is similar to the backwards DP algorithm in that it starts at the end of the horizon, estimates the cost-to-go for each state, and then steps back in time until it reaches the beginning of the horizon. However, instead of using the DP equations, which include computationally expensive expected value calculations, to calculate each cost-to-go exactly, the heuristic uses them to estimate the cost-to-go for only a subset of states, the "Base States." The heuristic then estimates the cost-to-go of the remaining states, or "non-Base States," using a function of the costs estimated for the Base States. Specifically, the BSH estimates the cost-to-go of each non-Base State (i_t, d_t) using some function Ψ of the estimated cost-to-go of some subset $S_{(i_t,d_t)}$ of the Base States which are deemed "similar" the non-Base State. Psudocode for the BSH can be found in the following algorithm.

for *all* $(i_T, d_T) \in I \times D$ **do**
  Initialize $\hat{V}_T(i_T, d_T) = G(i_T, d_T)$
end
for $t = T-1, ..., 0$ **do**
  Calculate costs for Base States:
  for *all* $(i_t, d_t) \in$ *Base States* **do**
    for *all* $d_{t+1} \in D$ **do**
      Calculate $p_t(d_t, d_{t+1}) = q(p_t^1(d_t^1, d_{t+1}^1), ..., p_t^{|J|}(d_t^{|J|}, d_{t+1}^{|J|}))$
    end
    Set $\hat{V}_t(i_t, d_t) =$
    $G(i_t, d_t) + min_{i_{t+1}} \left\{ f(d(i_t, i_{t+1})) + \sum_{d_{t+1} \in D} p_t(d_t, d_{t+1}) \hat{V}_{t+1}(i_{t+1}, d_{t+1}) \right\}$
  end
  Estimate costs for non-Base States:
  for *all* $(i_t, d_t) \in$ *non-Base States* **do**
    Set $\hat{V}_t(i_t, d_t) = \Psi(S_{(i_t,d_t)})$
  end
end
Return $\hat{V}_t(i_t, d_t) \ \forall \ (i_t, d_t) \in I \times D, \ t \in \{0, 1, ..., T\}$

Algorithm: Base State Heuristic

We now further describe which states are Base States and Ψ and the subsets $S_{(i_t,d_t)}$ over which Ψ operates. For the Base States, we choose a constant k as the maximum number of disasters (sites with positive demand) possible at the same time in a Base State. A given state is a Base State if the number of sites that are experiencing demand are less than or equal to k. If k is small, the number of Base States is a small percentage of the total number of states. The total number of Base States in the general case and for two specific examples are given in Table 1. If $|I|$ is large and contributing significantly to the computational time, then the number of Base States could be reduced by only including a subset of the potential inventory locations. Let $S_{(i_t,d_t)}$ be all of the Base States (i_t, d'_t) for which d'_t has k of the

k	General Case	$\|J\| = 16$, 2 possible demand levels each (% of total states)	$\|J\| = 8$, 4 possible demand levels each (% of total states)
1	$(1+\sum_{i=1}^{\|J\|}(\|D_i\|-1))\|I\|$	$17\|I\|$ (0.0259%)	$25\|I\|$ (0.0381%)
2	$(1+\sum_{i=1}^{\|J\|}(\|D_i\|-1)$ $+\sum_{i<j}((\|D_i\|-1)(\|D_j\|-1)))\|I\|$	$137\|I\|$ (0.209%)	$277\|I\|$ (0.423%)
3	$(1+\sum_{i=1}^{\|J\|}(\|D_i\|-1)$ $+\sum_{i<j}((\|D_i\|-1)(\|D_j\|-1))$ $+\sum_{i<j<l}((\|D_i\|-1)(\|D_j\|-1)(\|D_l\|-1)))\|I\|$	$697\|I\|$ (1.06%)	$1789\|I\|$ (2.73%)

Table 1: Total number of Base States for the general case and two examples. In parenthesis are the percentages of the total number of states.

same positive levels of demand as d_t. For example, in a case with $|J| = 3$ and $k = 1$,

$$S_{(i_t,(1,2,1))} = \{(i_t,(1,0,0)),(i_t,(0,2,0)),(i_t,(0,0,1))\}$$

Finally, define Ψ to be the maximum of the costs-to-go calculated for the base states in $S_{(i_t,d_t)}$ plus the cost to serve positive demand of the non-Base State not represented in the Base State. In our example, the estimate of the cost-to-go is then

$$\begin{aligned}\hat{V}_t(i_t,(1,2,1)) &= \Psi(S_{(i_t,(1,2,1))}) \\ &= max\{\ \hat{V}_t(i_t,(1,0,0)) + G(i_t,(0,2,1)),\ \hat{V}_t(i_t,(0,2,0)) + G(i_t,(1,0,1)), \\ &\qquad \hat{V}_t(i_t,(0,0,1)) + G(i_t,(1,2,0))\}\end{aligned}$$

These specifications, which we will refer to as **Specification 1**, were chosen because they are relatively easy to implement and perform well as it is likely that not many sites will experience demand at the same time. Another possible way to specify the Base States (Specification 2) is to choose the maximum number N_B of states to include in the Base States and then to choose the Base States as the N_B states (X_t, d_t) with the largest probabilities $p^t(d_0, d_t)$ of occurring for each time period. However, while upon first thought it seems that choosing Base States to be the most likely states will perform better than Specification 1, using this method does not guarantee that you will have useful states with which to estimate the non-Base States. Thus, it seems that the advantage gained out of more accurately estimating the most-likely states is outweighed by being unable to use these estimates to accurately estimate the non-Base States. Furthermore, since in each stage even the Base States' costs-to-go are calculated using the expected value of all future states (Base States *and* non-Base States), the inaccuracies in the non-Base State values are propagated throughout.

Furthermore, Specification 1 gives a concrete and straight forward way of choosing which Base States are used to estimate each non-Base State. On the contrary, Specification 2 has many decisions associated with its implementation. Specifically, one must choose whether to use different Base States for each time period or the same Base States for each time period (in our computational experiments, we found that implementations using the same Base States for each time period tended to perform better). Once the Base States are defined, one must choose which Base States should be used to estimate each non-Base State. One must either specify this one non-Base State at a time or come up with a general rule. We have yet to find another implementation which performs better than Specification 1. Furthermore, the following result holds:

Theorem 16. *Under Specification 1, the policy suggested by the BSH is optimal when the demands are temporally independent, i.e. when* $p_t(d_t, d_{t+1}) = p_t(d'_t, d_{t+1}) \ \forall \ d_t, d'_t, d_{t+1} \in D$, *and* $t \in \{0, 1, ..., T-1\}$.

Theorem 16 states that the policy suggested by the BSH is optimal when the demand realization in one period gives no additional information regarding the next period's demand. We describe performance results of the BSH under Specification 1 in Section 3.4.

3.4 Numerical Results

In this section, we present performance results of the BSH as well as numerical experiments and managerial insights on the model described in Section 3.1. We do not seek to estimate the value or simulate true operations for a specific case study. Rather, as the first to study inventory mobility for disaster relief, we seek to understand how the cost and value of inventory mobility depend on the parameter values and demand characteristics to gain insights on when mobile inventory systems are worthwhile. Accordingly, we consider a network which can be easily modified to represent different settings and solved to optimality quickly to enable running many test cases. A real potential disaster network is too complex and large to be modified for the purpose of our experiments. Furthermore, as we see in Section 3.4.1, large instances of the problem take a prohibitive amount of time to solve to optimality. Thus, we consider a simple and small generic network. We keep in mind the motivating example of a mobile inventory kept on a ship. The potential disaster sites are located randomly (except where otherwise noted) on one of two parallel coasts which are 2000 miles apart and 2000 miles long each; this network is illustrated in Figure 2. The ship's potential locations are restricted to points within these two coasts defined by a discrete grid with 200 mile spacing for a total of 121 nodes. We assume $T = 30$ and each time period is two days.

We utilize historical data from EM-DAT (Guha-Sapir et al. 2014), a database containing statistics on the occurrence and effects of major disasters, to generate the simple demand levels and transition probabilities of our experiments. We utilize data on earthquakes and tropical cyclones (e.g. hurricanes) which occurred in the years 1990 to 2014 in North and South America, Europe, and Africa. To be able to consider response to tropical cyclones and use historical data to generate problem parameters in a realistic manner, we assume our planning horizon is the 30 time periods from August 15 to October 13 which is during the Atlantic hurricane season. As detailed in Table 2, we identified five geographically compact regions which were active areas for earthquakes in the years considered and five which were active areas for cyclones in August 15 through September 13 of the years considered. Tables 3 and 4 contain summary statistics for each of these regions. For each region and each type

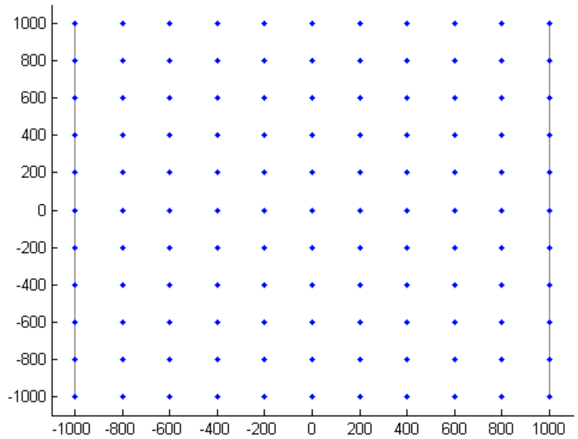

Figure 2: The network considered in our numerical experiments. Potential disaster sites are on one of the two coasts (the vertical black lines) and the potential inventory locations are the blue points between the coasts.

of disaster, we aggregated data points of disasters which occurred on the same day. We only considered disasters which affected at least 100 people and considered each disaster a "small" disaster if it affected between 100 and 30,000 people and a "large" disaster if it affected at least 30,000 people. In the database, affected people are those requiring basic survival supplies such as food, water, and shelter or immediate medical assistance.

The transition probabilities and demand levels for each test case were randomly generated to be in ranges similar to the statistics displayed in Tables 3 and 4. Except where otherwise noted, our numerical set up is as follows. Each potential disaster site has two potential demand levels (for tractability): 0 and a realization of an exponential random variable with distribution fitted to the average number of people affected by all disasters for the 10 regions using the method of moments (i.e. $1/\lambda = 227,092$, the average value over the 10 regions). For each potential disaster site and time period, the probability of transitioning from 0 demand to positive demand is a realization of a random variable which is distributed Uniform(0.002849, 0.052); the range of the distribution is the range of the single period disaster probabilities for the 10 regions. Once a disaster occurs, a relief organization will typically provide supplies for some period of time; here we assume that the pre-positioned inventory is used to provide supplies for an average of 3 periods while permanent supply lines are established. Thus, for each potential disaster site and each time period, the probability of remaining in the same level of positive demand is $2/3 \approx 0.6667$. We assume that demands are independent across sites (we relax this assumption in the correlation experiment of Section 3.4.2) and that all sites have zero demand at the start of the horizon.

The cost to move between nodes is determined by a function of the distance traveled. The function used is roughly based on a fuel consumption model for container ships developed by Notteboom and Cariou (2009) and is as follows:

$$Cost_f(x) = \begin{cases} 600\left(\frac{20.5}{24}\right)\left(\frac{x/\psi}{14}\right) & \text{if } x <= 672\psi \\ 600 * 260\left(\frac{x/\psi}{48*24.5}\right)^{3.3} & \text{if } x > 672\psi \end{cases}$$

Active earthquake regions		Active tropical cyclone regions	
E1:	Guatemala, Honduras, El Salvador	C1:	Haiti, the Dominican Republic, Puerto Rico
E2:	Greece, Albania	C2:	Florida
E3:	Southwestern Peru (Lima Province, regions of Ica, Ayacucho, Apurímac, Cuzco, Arequipa)	C3:	Southeastern Mexico (states of Veracruz, Tabasco, and Compeche)
E4:	Northern and western Columbia (departments of Nariño, Cauca, Valle del Cauca, Chocó, Risaralda, Antioquia, Cundinamarca, Boyacá, Arauca)	C4:	Eastern Caribbean islands (St Kitts and Nevis, Antigua and Barbuda, Dominica, Martinique, St Lucia, St Vincent and the Grenadines, Barbados, Grenada)
E5:	Nicaragua (west of Matagalpa), Costa Rica, Panama (west of Santiago)	C5:	Louisiana, southern Mississippi, southern Alabama

Table 2: Active regions for earthquakes and tropical cyclones.

Region	Average # of people affected by			Single period disaster probability	% of disasters that were "small"
	all	"small"	"large"		
E1	236204	5550	605251	0.0028	0.6154
E2	16555	3469	95071	0.0031	0.8571
E3	85616	2873	361426	0.0028	0.7692
E4	91570	3220	621675	0.0031	0.8571
E5	12057	4287	128618	0.0035	0.9375
C1	89106	7449	221799		
C2	1196227	2764	1673613		
C3	204127	6415	327696		
C4	8718	3885	61884		
C5	330738	4832	548009		

Table 3: Regional summary statistics including the average number of people affected by various disaster sizes. To calculate the single period disaster probability, we divided the total number of disasters that occurred by 25 years and divided the result by $365/2 = 182.5$*, the total number of time periods per year. Additional cyclone region statistics are detailed in Table 4.*

Region	Single period disaster probability			% of disasters that were "small"		
	8/15-9/3	9/4-9/23	9/24-10/13	8/15-9/3	9/4-9/23	9/24-10/13
C1	0.0520	0.0240	0.0080	0.6154	0.5000	1.0000
C2	0.0080	0.0120	0.0080	0.5000	0.3333	0.0000
C3	0.0160	0.0120	0.0240	0.5000	0.0000	0.5000
C4	0.0160	0.0160	0.0160	1.0000	0.7500	1.0000
C5	0.0200	0.0120	0.0080	0.2000	0.6667	0.5000

Table 4: Additional summary statistics for the tropical cyclones regions. To calculate the single period disaster probability, we divided the total number of disasters that occurred in the specified range of time by 25 and divided the result by 10, the number of time periods in each range of time.

where $\psi \approx 1.150779$ is the number of miles in a nautical mile and the cut off point, 672ψ, is approximately the number of miles that can be traveled in 48 hours at a speed of 14 knots. The function is of this form to model the fact that at a certain speed, less fuel efficiencies can be gained by traveling at lower speeds. \$600 is the price per ton of fuel. In the definition of the case where $x <= 672\psi$, $\left(\frac{20.5}{24}\right)$ is the tons of fuel used per hour and $\left(\frac{x/\psi}{14}\right)$ is the number of hours to travel at 14 knots. In the definition of the case where $x > 672\psi$, $260\left(\frac{x/\psi}{48*24.5}\right)^{3.3}$ is roughly the fuel consumption for the two-day time period. Thus, the cost of relocating the ship from position (x_1, y_1) to position (x_2, y_2) is

$$f(d((x_1, y_1), (x_2, y_2))) = Cost_f\left(\sqrt{(x_2 - x_1)^2 + (y_2 - y_1)^2}\right)$$

Following Rawls and Turnquist (2012), we assume the cost to serve demand at a potential disaster site is \$0.0015 per unit of demand per mile between the ship and the potential disaster site. All numerical experiments were performed on a 3.07GHz Windows 7 x64 laptop computer using MATLAB.

3.4.1 Base State Heuristic Numerical Results

Table 5 shows numerical results on the performance of the BSH for $k \in \{1, 2, 3\}$ for cases with different numbers of potential disaster sites, $|J| \in \{3, 4, ..., 16\}$. The optimal CPU time is the amount of time it takes to solve the model exactly using the backwards DP algorithm. Table 5 shows that the time it takes to solve the problem to optimality increases dramatically as $|J|$ increases. However, it also shows that the BSH, even with $k = 1$, solves the problem to near optimality in a small fraction of the time. For example, the case with 16 sites took over 31 hours to solve to optimality and only 86.06 seconds, or 0.08% of the time, to solve to within 0.21% of optimality using the BSH with $k = 1$. For all test cases, the BSH found a solution within 0.43% of optimality in less than 87 seconds. It appears that Ψ is a good estimate of the cost-to-go for the non-Base States; it is such that the resulting

		Optimal	BSH, $k = 1$		BSH, $k = 2$		BSH, $k = 3$	
$\|J\|$	n	time	time	% error	time	% error	time	% error
3	10	0.07	0.03	0.0967%	0.05	0.0008%	N/A	N/A
4	10	0.12	0.04	0.0560%	0.08	0.0001%	0.11	0.0000%
5	10	0.28	0.06	0.1760%	0.15	0.0009%	0.23	0.0000%
6	10	0.62	0.08	0.1913%	0.23	0.0007%	0.42	0.0000%
7	10	1.40	0.12	0.1974%	0.36	0.0011%	0.73	0.0000%
8	10	3.33	0.20	0.1999%	0.59	0.0011%	1.30	0.0000%
9	10	9.25	0.35	0.2286%	1.07	0.0021%	2.60	0.0001%
10	10	26.28	0.69	0.2487%	2.06	0.0025%	5.25	0.0001%
11	10	91.81	1.43	0.1915%	4.46	0.0031%	12.29	0.0001%
12	10	353.76	3.43	0.2620%	11.44	0.0038%	32.74	0.0001%
13	10	1700.27	7.73	0.2011%	29.54	0.0041%	95.63	0.0002%
14	5	6901.68	17.02	0.2688%	71.05	0.0037%	246.57	0.0002%
15	2	27819.90	36.82	0.2411%	164.41	0.0035%	611.58	0.0002%
16	1	112043.30	86.06	0.2144%	376.55	0.0029%	1487.37	0.0002%

Table 5: BSH performance for a varying number of potential disaster sites $|J|$. n is the number of runs, times are CPU times (in seconds), and CPU times and percent errors (percent deviations from optimal) are the average over the n runs.

$\hat{V}_t(i_t, d_t)\ \forall i_t \in I$ are in relatively the right order for each time period such that optimal, or close to optimal, decisions of where to relocate the inventory are made.

Table 5 also shows that increasing k weakly decreases the percent within optimality slightly at the expense of increased computation time. From this perspective, for large $|J|$, the percent error is likely low enough with $k = 1$ that using a larger k is not worth the additional computation time. However, Table 6 suggests another advantage of using a larger k. In using the BSH, we obtain an estimate of the true expected cost of the policy suggested by the BSH, $\hat{V}_0(i_0, d_0)$, rather than the true expected cost itself. For planning purposes, it may be worthwhile to have a more accurate estimate of the expected cost of using the policy. Table 6 shows that increasing k improves our estimate of the cost for all $|J|$ tested.

3.4.2 Managerial Insights

In this section, we present numerical experiments from which we extract further managerial insights. The numerical set up is as described above except where otherwise noted. We first sought to get a sense of how the concentration of risk, or the location configuration of the potential disaster sites, affects the value of inventory mobility. We ran 1000 test cases each of two different potential disaster site location configurations with six sites each: less concentrated risk with sites in the distant locations $(-1000, -1000)$, $(-1000, 0)$, $(-1000, 1000)$, $(1000, 1000)$, $(1000, 0)$, and $(1000, -1000)$, and more concentrated risk with sites in the nearby locations $(-1000, -200)$, $(-1000, -120)$, $(-1000, -40)$, $(-1000, 40)$, $(-1000, 120)$,

$\lvert J\rvert$	n	$k=1$	$k=2$	$k=3$	$\lvert J\rvert$	n	$k=1$	$k=2$	$k=3$
3	10	5.5%	0.2%	N/A	10	10	17.2%	2.5%	0.3%
4	10	5.5%	0.2%	0.0%	11	10	17.8%	3.1%	0.4%
5	10	10.7%	0.8%	0.0%	12	10	19.8%	3.8%	0.5%
6	10	12.5%	1.1%	0.1%	13	10	21.1%	4.5%	0.8%
7	10	12.6%	1.1%	0.1%	14	5	20.7%	4.4%	0.7%
8	10	14.2%	1.5%	0.1%	15	2	21.3%	4.6%	0.8%
9	10	16.4%	2.2%	0.2%	16	1	22.4%	4.7%	0.8%

Table 6: BSH expected cost estimate deviation from true expected cost of the policy suggested by the BSH for a varying number of potential disaster sites $|J|$. n is the number of runs. The percent deviation for each k is the average over these n runs. Note that, for all runs, the BSH cost estimate was below the true cost and thus the deviation shown is the percent below the true cost.

and $(-1000, 200)$. The results are displayed as box plots in Figure 3. Furthermore, a test with all potential disaster sites located at a single point resulted in a zero value of inventory mobility. From the figure, it appears that the value of inventory mobility is on average greater when the risk is less concentrated or when the potential disaster sites are farther from each other. When the risk is less concentrated, the optimal stationary inventory position, which must compromise between different optimal locations throughout the planning horizon, is farther from each potential disaster site. Thus, there is more value in being able to move closer, or shorten the large distance, to individual sites as demand states are realized. In this case, the cost to move the inventory is outweighed by the benefit of getting closer to potential disaster sites that are either experiencing or are likely to soon experience a disaster.

Especially when they are located near each other, the demand of two potential disaster sites may be correlated. To test the effect of varying levels of correlation on the value of inventory mobility, we ran 1000 test cases with two sites located at $(-1000, 0)$ and $(1000, 0)$, respectively. As each site has two potential demand levels, in a given period, the demand vector realization for the next period can be represented by a bivariate binomial random variable and a specific correlation between the sites can be easily incorporated. For each test case, we varied the correlation coefficient ρ between -1 and 1 in increments of 0.2. Figure 4 displays the value of inventory mobility for a random five of these test cases. We see that the value is greater when the demands are less correlated. This holds in the weak sense for all test cases and in the strong sense for all test cases with non-zero values of inventory mobility for all values of ρ. This is intuitive; if two potential disaster sites are perfectly correlated, then they both experience demand at the same time and there is less value in being able to move closer to one over the other. In this case, it may be worthwhile to establish duplicate stationary inventories near each site rather than invest in a mobile inventory system. On the other hand, when the demands are not correlated or are negatively correlated, then one site may experience demand while the other does not, adding value to the ability to move. This experiment suggests that the value will be greater for systems in which sites that are relatively far from each other have demands which are less correlated.

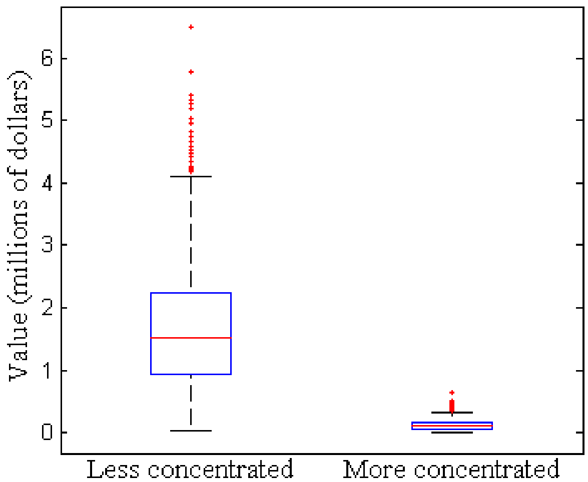

Figure 3: Value of inventory mobility for six potential disaster sites for varying risk concentrations.

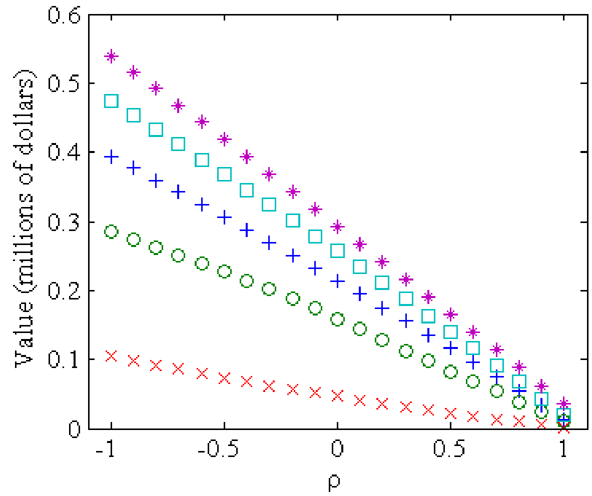

Figure 4: Value of inventory mobility for two potential disaster sites for varying levels of correlation. Each symbol represents a different test case.

Demands in disaster relief and other settings have varying levels of frequency and magnitudes. For example, most demands in the disaster relief setting occur rarely but are of high magnitude whereas demands in most other settings are frequent but of relatively low magnitude. To test the effect of demand frequency and magnitude on the value of inventory mobility, we ran 1000 test cases with six potential disaster sites each for each of two different settings: rare and high magnitude demand, with demand parameters assigned as described at the beginning of this section, and frequent and low magnitude demands, with parameters assigned similarly to as we described at the beginning of this section but with the exponential demand generation random variable parameter, λ, and the min and max of the uniform probability generation random variable multiplied by 10. Figure 5 displays the value for these two settings; it appears that the value is on average greater when demands are rare and of high magnitude rather than frequent and of low magnitude. When demands are rare and of high magnitude, when there is positive demand at a particular location, it is likely that no other or very few other sites also have positive demand. Thus, the mobile inventory will likely be used to serve only one or a small number of high demand sites at a time and can alter its location accordingly. When demands are frequent and of low magnitude, however, many sites will experience low levels of demand often. Thus, the mobile inventory will likely not deviate much from the optimal stationary position which is a compromise between serving all demand sites over the planning horizon. In this case, it may be worthwhile to establish duplicate stationary inventories rather than invest in a mobile inventory system.

In the following experiment, we further explore how the value depends on various demand characteristics by examining the value associated with serving different types of disasters. We focus on earthquakes and tropical cyclones and define the demand parameters differently than we have thus far to more accurately represent earthquake- and cyclone-related demand

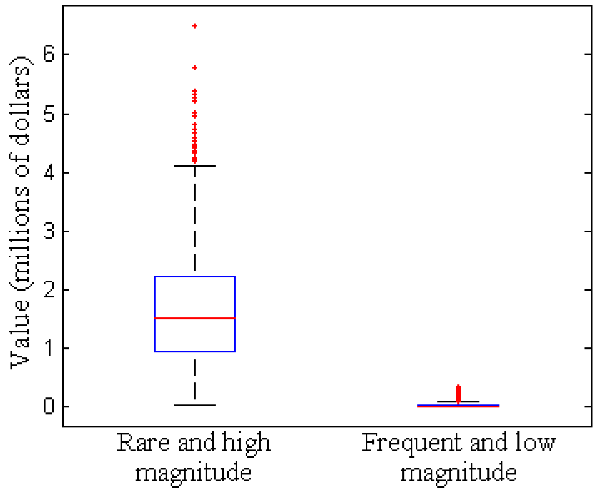

Figure 5: Value of inventory mobility for six potential disaster sites for varying demand frequencies and magnitudes.

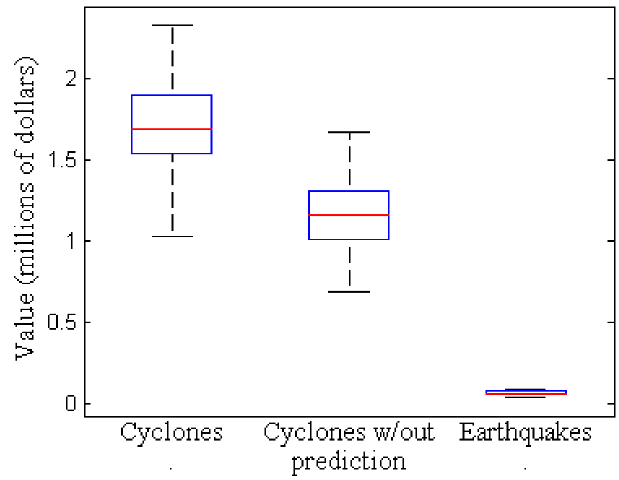

Figure 6: Value of inventory mobility for five potential disaster sites associated with different types and variations of disasters.

characteristics. Each test case has five potential disaster sites where all of the sites are active regions for earthquakes or all are active regions for cyclones. In each test case, the five sites are assigned demand parameters to be similar to the five regions associated with the same type of disaster. We define demand parameters for the regions using the statistics displayed in Tables 3 and 4. We assume that earthquake transition probabilities do not change over time while cyclone probabilities do. In addition, weather forecasts enable better predictions of when cyclones will occur. Thus, instead of a single zero-demand state, cyclones have two zero-demand states, one where a cyclone is likely to happen in the next period ("0-Likely" or "0-L") and one where a cyclone is not likely ("0-Unlikely" or "0-U"). For each disaster region, two positive demand levels were assigned to match the average number of people affected by "small" and "large" disasters, respectively. As above, we assume the probability of transitioning from positive demand to the same level of positive demand is ≈ 0.6667. For simplicity, we assume that the probability of transitioning from one level of positive demand to a different level is 0. For earthquake regions, the probability of transitioning from 0 to the "small" level of demand is the single period disaster probability times the percentage of disasters that are "small." Similarly, the probability of transitioning from 0 to the "large" level of demand is the single period disaster probability times the percentage of disasters that are "large." To ensure robustness, in each test case we assign the probabilities of transitioning to positive demand in this way with the single period disaster probability uniformly varied between $+/-10\%$ of the value listed in Table 3. As an example, the matrix describing the demand transition probabilities for earthquake region E1 is shown in Table 7a.

For cyclone regions, we assume the transition probabilities change over time. Specifically, we assign different probabilities for each of the 10-period time ranges August 15 - September 3, September 4 - September 23, and September 24 - October 13. On its website, the National Hurricane Center shows that from 2009 to 2013, its 48 hour forecasts of the *(a) Earthquake region E1*

	0	5,550	605,251
0	0.9972	0.0018	0.0011
5,550	0.3333	0.6667	0.0000
605,251	0.3333	0.0000	0.6667

(b) Cyclone region C1

	0-U	0-L	7,449	221,799
0-U	0.9257	0.0743	0	0
0-L	0.3	0	0.4308	0.2692
7,449	0.3333	0	0.6667	0
221,799	0.3333	0	0	0.6667

Table 7: Demand transition probabilities for regions E1 and C1 from August 15 - September 3 where we use the single period disaster probability listed in Tables 3 and 4 without variation for clarity.

tracks of tropical cyclones in the Atlantic Basin were accurate to within 100 miles about 70% of the time (National Hurricane Center, 2015). Thus, we assume that the probability of transitioning from the 0-Likely demand state to a positive level of demand is 0.7. To calculate the probability of transitioning to each specific level of positive demand for each region in each of the time ranges, we multiply 0.7 times the percentage of disasters of that size for the region in that time range as listed in Table 4. Then, to calculate the probability of transitioning from 0-Unlikely to 0-Likely, we take the single period disaster probability for that time range and divide by 0.7. To ensure robustness, in each test case we assign the probability of transitioning from 0-Unlikely to 0-Likely in this way with the single period disaster probability uniformly varied between $+/-10\%$ of the value listed in Table 4. As an example, the matrix describing the transition probabilities for cyclone region C1 from August 15 - September 3 is shown in Table 7b.

In each test case, the five sites are distributed among the five locations $(-1000, -1000)$, $(-1000, 0)$, $(-1000, 1000)$, $(1000, -500)$, and $(1000, 500)$; since each site is associated with one region, there are $5! = 120$ possible configurations. For each type of disaster and each of the 120 configurations, we ran 100 test cases. For further comparison, we also ran 100 test cases for each of the 120 configurations for tropical cyclones without weather forecasts predictions; that is, each potential disaster site has only one zero-demand state and we generated the demand transition probabilities using the same method we used for earthquakes. Figure 6 suggests that there is a greater value of inventory mobility when serving tropical cyclones compared to earthquakes. This observation may be explained by several factors. The first is that cyclones are easier to predict due to weather forecasts; from the figure, we see that the value is on average greater when we are able to predict one period ahead of time when a cyclone is likely to occur. Another possible factor may be our assumption that the cyclone demand transition probabilities change throughout the planning horizon; however, in the setting we consider here, similar experiments reveal that the value is on average only slightly greater (about 0.04% greater) when we use non-stationary probabilities throughout the planning horizon for cyclones rather than stationary probabilities as we do for earthquakes. Lastly, in this experiment, cyclone sites tend to have greater probabilities of positive demand and on average greater demand levels; this factor likely explains the differences in value between serving cyclones without weather forecast predictions and earthquakes. Thus, while mobile inventory has value in serving all types of disasters, the value is greater when it is used to provide relief to disasters which are easier to forecast accurately and which have relatively greater demand frequency and magnitude. Note that this experiment is a snapshot of the value in the middle of a cyclone season. In a time window outside of their cyclone season, cyclone sites will likely have a very small chance of experiencing positive demand and thus there may be little value associated with serving these sites. While this may be true, there is likely significant value in being able to change the location of the inventory as changes in demand forecasts such as these occur over time.

3.5 Proofs

3.5.1 Proof of Theorem 1

Proof. The DP equations 3.2 can be written as one equation as follows:

$$
\begin{aligned}
\bar{V}_0(\bar{i}, d_0) &= G(\bar{i}, d_0) + P_1^0\left[G(\bar{i}, \hat{d}) + P^1\left[G(\bar{i}, \hat{d}) + P^2\left[G(\bar{i}, \hat{d}) + P^3\left[... + P^{T-1}\left[G(\bar{i}, \hat{d})\right]\right]\right]\right]\right] \\
&= G(\bar{i}, d_0) + P_1^0[I + P^1 + P^1P^2 + ... + P^1P^2 \cdots P^{T-1}]G(\bar{i}, \hat{d}) \\
&= G(\bar{i}, d_0) + P_1^0\left[I + \sum_{t=2}^{T}\prod_{k=1}^{t-1} P^k\right] G(\bar{i}, \hat{d}) \\
&= G(\bar{i}, d_0) + \sum_{t=1}^{T}\sum_{d\in D} p^t(d_0, d)G(\bar{i}, d) \\
&= G(\bar{i}, d_0) + \sum_{d\in D} G(\bar{i}, d)\left[\sum_{t=1}^{T} p^t(d_0, d)\right] \\
&= \sum_{d\in D} G(\bar{i}, d)\left[\sum_{t=1}^{T} p^t(d_0, d) + I\{d = d_0\}\right] \qquad \square
\end{aligned}
$$

3.5.2 Proof of Theorem 2

Proof. Let $\bar{i}^* = \arg\min_{\bar{i}\in I_s} \bar{V}_0(\bar{i}, d_0)$. Consider the feasible policy π for the mobile inventory system represented by DP equations 3.1 in which the stationary inventory is kept at $\bar{i}^*$ for the length of the planning horizon. In this policy, movement costs are zero since the inventory never moves and $f(0) = 0$. Thus, π is feasible for the mobile inventory system and $V_0^\pi(\bar{i}^*, d_0) = \bar{V}_0(\bar{i}^*, d_0)$. Thus,

$$
\begin{aligned}
& \min_{i_0\in I} V_0(i_0, d_0) \leq \bar{V}_0(\bar{i}^*, d_0) = \min_{\bar{i}\in I_s} \bar{V}_0(\bar{i}, d_0) \\
\equiv\quad & \min_{\bar{i}\in I_s} \bar{V}_0(\bar{i}, d_0) - \min_{i_0\in I} V_0(i_0, d_0) \geq 0 \qquad \square
\end{aligned}
$$

3.5.3 Proof of Theorem 3

Proof. Let $J = \{1, ..., |J|\}$ and $g_j(|x-l_j|, d^j) = b_{2j}|x-l_j|+b_{3j}$ where b_{2j} and b_{3j} are functions of d_j. Without loss of generality, let $l_j > 0\ \forall j \in J$ and $l_j < l_{j+1}\ \forall j \in J \setminus |J|$. By Theorem

4, an optimal solution x^* exists such that $x^* \in [l_1, l_{|J|}]$. Assume $\neg\exists$ an optimal solution x^* minimizing $\bar{V}_0(x, d_0)$ such that $x^* = l_j$ for some $j \in J$. Then, for an optimal solution x^*, $\exists i \in J \setminus |J|$ such that $l_i < x^* < l_{i+1}$. The cost at this optimal solution is

$$\begin{aligned}
\bar{V}_0(x^*, d_0) &= \sum_{d\in D} \left[\sum_{j=1}^{|J|} (b_{2j}|x^* - l_j| + b_{3j}) \right] \left[\sum_{t=1}^{T} p^t(d_0, d) + I\{d = d_0\} \right] \\
&= \sum_{d\in D} K_d \left[\sum_{j=1}^{i} (b_{2j}(x^* - l_j) + b_{3j}) + \sum_{j=i+1}^{|J|} (b_{2j}(l_j - x^*) + b_{3j}) \right] \\
&= \sum_{j=1}^{i} (x^* - l_j) \sum_{d\in D} K_d b_{2j} + \sum_{j=i+1}^{|J|} (l_j - x^*) \sum_{d\in D} K_d b_{2j} + \sum_{j=1}^{|J|} \sum_{d\in D} K_d b_{3j}
\end{aligned}$$

where the constant $K_d = \sum_{t=1}^{T} p^t(d_0, d) + I\{d = d_0\}$ where $I\{\cdot\}$ is the indicator function.

Consider first the case where

$$\sum_{j=1}^{i} \sum_{d\in D} K_d b_{2j} \geq \sum_{j=i+1}^{|J|} \sum_{d\in D} K_d b_{2j}$$

Then

$$\begin{aligned}
\bar{V}_0(l_i, d_0) &= \bar{V}_0(x^* - (x^* - l_i), d_0) \\
&= \sum_{j=1}^{i} (x^* - (x^* - l_i) - l_j) \sum_{d\in D} K_d b_{2j} + \sum_{j=i+1}^{|J|} (l_j - (x^* - (x^* - l_i))) \sum_{d\in D} K_d b_{2j} \\
&\quad + \sum_{j=1}^{|J|} \sum_{d\in D} K_d b_{3j} \\
&= \left(\sum_{j=i+1}^{|J|} \sum_{d\in D} K_d b_{2j} - \sum_{j=1}^{i} \sum_{d\in D} K_d b_{2j} \right) (x^* - l_i) + \bar{V}_0(x^*, d_0) \\
&\leq \bar{J}_0(x^*, d_0)
\end{aligned} \tag{3.20}$$

If we are instead in the case where

$$\sum_{j=1}^{i} \sum_{d\in D} K_d b_{2j} < \sum_{j=i+1}^{|J|} \sum_{d\in D} K_d b_{2j}$$

then it can be similarly shown that $\bar{V}_0(l_{i+1}, d_0) \leq \bar{V}_0(x^*, d_0)$. Thus we have contradicted our assumption that $\neg\exists$ an optimal solution x^* minimizing $\bar{V}_0(x, d_0)$ such that $x^* = l_j$ for some $j \in J$. $\square$

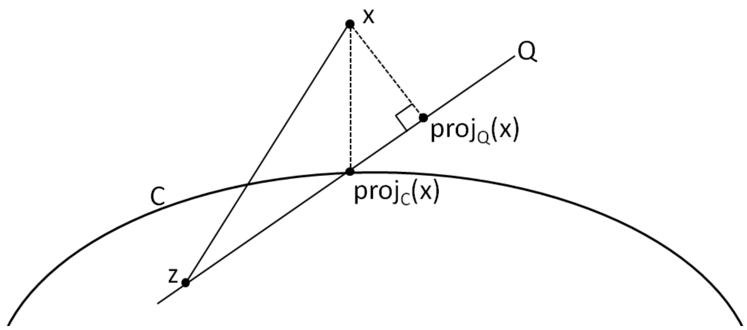

Figure 7

3.5.4 Proof of Lemma 1

Proof. Suppose $x \in C$. Then $proj_C(x) = x$ and thus $d(proj_C(x), z) = d(x, z)$. Assume $x \notin C$. Consider Figure 7. Consider some point $z \in C$ and the line Q passing through z and $proj_C(x)$. By definition, $proj_C(x) = \arg\min_{z \in C} d(x, z)$, that is, $proj_C(x)$ is the closest point in C to x. Additionally, the line segment $[proj_C(x), z] \subseteq C$. Since the Euclidean distance metric is continuous, the projection of x on Q, $proj_Q(x)$, must lie on Q such that $proj_C(x)$ lies between z and $proj_Q(x)$. Thus,

$$d(proj_C(x), z) \leq d(proj_Q(x), z) \tag{3.21}$$

By definition of a projection, the line segment $[x, proj_Q(x)]$ is orthogonal to Q. Thus, the triangle formed by the points x, $proj_Q(x)$, and z is a right triangle. By equation 3.21 and the triangle inequality,

$$\begin{aligned} d(proj_C(x), z) &\leq d(proj_Q(x), z) \\ &< d(x, z) \end{aligned}$$

□

3.5.5 Proof of Theorem 4

Proof. Let $L = conv(\hat{B} \cup \{l_j | j \in J\})$. Assume $\exists$ some $\tilde{i} \in I$ minimizing $\bar{V}_0(i, d_0)$ such that $\tilde{i} \notin \hat{I} = I \cap L$. Note that $\hat{B} \cup \{l_j | j \in J\}$ is a closed and bounded subset of $\mathbb{R}^2$ and thus is compact by the Heine-Borel Theorem. Thus, L is closed since the convex hull of a compact set in $\mathbb{R}^2$ is compact and every compact set is closed. Note that $l_j \in L\ \forall j \in J$. By Lemma 1, $\bar{i} = proj_L(\tilde{i}) \in L$ is such that $d(\bar{i}, l_j) \leq d(\tilde{i}, l_j)\ \forall j \in J$. Thus, $\bar{V}_0(\bar{i}, d_0) \leq \bar{V}_0(\tilde{i}, d_0)$ since $g_j(d(\cdot, \cdot), d^j)$ is non-decreasing in $d(\cdot, \cdot)\ \forall j \in J$. By construction, the boundary $L \backslash Int(L) \in I$ and thus $\bar{i} = proj_L(\tilde{i}) \in I$. Also $\bar{i} \in L$ and $\bar{i} \in I \Rightarrow \bar{i} \in \hat{I}$. Thus, $\bar{i} \in \hat{I}$ is also a minimizer of $\bar{V}_0(i, d_0)$ and satisfies the conditions of the theorem. □

3.5.6 Proof of Lemma 2

Proof. Suppose $x, y \in C$. Then $proj_C(x) = x$ and $proj_C(y) = y$ and thus $d(proj_C(x), proj_C(y)) = d(x, y)$. Without loss of generality, suppose $x \in C$ and $y \notin C$. Then,

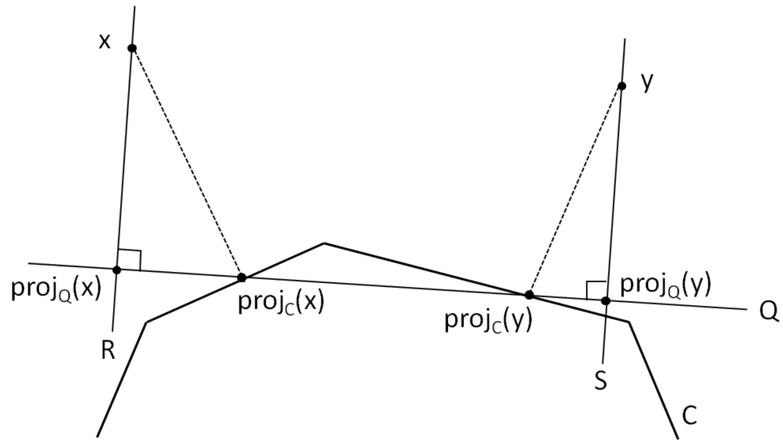

Figure 8

by Lemma 1,

$$\begin{aligned} d(proj_C(x), proj_C(y)) &= d(x, proj_C(y)) \\ &\leq d(x, y) \end{aligned}$$

Assume $x, y \notin C$. If $proj_C(x) = proj_C(y)$, then

$$\begin{aligned} d(proj_C(x), proj_C(y)) &= 0 \\ &\leq d(x, y) \end{aligned}$$

Assume $proj_C(x) \neq proj_C(y)$. Let Q be the line through $proj_C(x)$ and $proj_C(y)$. Consider Figure 8. The line segment $[proj_C(x), proj_C(y)] \subseteq C$ since C is convex and $proj_C(x)$, $proj_C(y) \in C$. Since the Euclidean distance function is continuous, the projection of x on Q, $proj_Q(x)$, must lie on Q such that $proj_C(x)$ lies between $proj_Q(x)$ and $proj_C(y)$ and the projection of y on Q, $proj_Q(y)$, must lie on Q such that $proj_C(y)$ lies between $proj_Q(y)$ and $proj_C(x)$, as shown in Figure 8. Thus,

$$d(proj_C(x), proj_C(y)) \leq d(proj_Q(x), proj_Q(y)) \tag{3.22}$$

By definition, the line connecting x and $proj_Q(x)$, R, and the line connecting y and $proj_Q(y)$, S, are orthogonal to Q. Thus, $R||S$ and $d(proj_Q(x), proj_Q(y)) \leq d(a, b)$ $\forall a \in R$ and $b \in S$. Thus, using equation 3.22 and noting that $x \in R$ and $y \in S$,

$$\begin{aligned} d(proj_C(x), proj_C(y)) &\leq d(proj_Q(x), proj_Q(y)) \\ &\leq d(x, y) \end{aligned}$$ □

3.5.7 Proof of Theorem 5

Proof. Let $L = conv(\hat{B} \cup \{l_j | j \in J\})$. Note that $\hat{B} \cup \{l_j | j \in J\}$ is a closed and bounded subset of $\mathbb{R}^2$ and thus is compact by the Heine-Borel Theorem. Thus, L is closed since the convex hull of a compact set in $\mathbb{R}^2$ is compact and every compact set is closed.

For $i \in I$, suppose $proj_L(i) \notin I$. Thus, $d(i, proj_L(i)) > 0$. By definition of the projection, $proj_L(i) \in L$. Consider the line segment R joining i and $proj_L(i)$. Let $q \in R \cap L$ be such that q is on the boundary of L (q exists since L is closed). Then, since $\hat{B} \subseteq B$, by the definitions of B and $\hat{B}$, $q \in I$. Thus, $q \neq proj_L(i)$ and $q \in R \Rightarrow d(i, q) < d(i, proj_L(i))$ which is a contradiction since $q \in L$. Thus, $proj_L(i) \in I \ \forall i \in I \Rightarrow proj_L(i) \in I \cap L = \hat{I} \ \forall i \in I$.

Let $\pi \in \Pi$ be an optimal policy such that $i_0^\pi \notin \hat{I}$ or $i_{t+1}^\pi(i_t, d_t) \notin \hat{I}$ for some $i_t \in \hat{I}$, $d_t \in D$, and $t \in \{0, ..., T-1\}$. Let

$$Z_t(i_t, d_t, i_{t+1}) = f(d(i_t, i_{t+1})) + \sum_{d_{t+1} \in D} p_t(d_t, d_{t+1}) V_{t+1}(i_{t+1}, d_{t+1})$$

Then $i_{t+1}^\pi(i_t, d_t) = \arg\min_{i_{t+1} \in I} Z_t(i_t, d_t, i_{t+1})$. Consider the policy $\pi_p = \{i_0^{\pi_p}, ..., i_T^{\pi_p}\}$ such that

$$\begin{aligned}
&i_0^{\pi_p} = proj_L(i_0^\pi) \\
&i_{t+1}^{\pi_p}(i_t, d_t) = proj_L(i_{t+1}^\pi(i_t, d_t)) \quad \forall\ i_t \in \hat{I},\ d_t \in D,\ t \in \{0, ..., T-1\} \\
&i_{t+1}^{\pi_p}(i_t, d_t) = i_{t+1}^\pi(i_t, d_t) \quad \forall\ i_t \in I \setminus \hat{I},\ d_t \in D,\ t \in \{0, ..., T-1\}
\end{aligned}$$

By the argument above, π_p is a feasible policy such that $i_0^{\pi_p} \in \hat{I}$ and $i_{t+1}^{\pi_p}(i_t, d_t) \in \hat{I} \ \forall\ i_t \in \hat{I}$, $d_t \in D$, and $t \in \{0, ..., T-1\}$. Note that $l_j \in L \ \forall j \in J$. By Lemma 1 and the assumption that $g_j(d(\cdot, \cdot), d^j)$ is non-decreasing in $d(\cdot, \cdot) \ \forall j \in J$, for $i_t \in I$ and $d_t \in D$,

$$\begin{aligned}
V_T(i_T, d_T) &= G(i_T, d_T) \\
&= \sum_{j \in J} g_j(d(i_T, l_j), d_T^j) \\
&\geq \sum_{j \in J} g_j(d(proj_L(i_T), l_j), d_T^j) \\
&= G(proj_L(i_T), d_T) \\
&= V_T(proj_L(i_T), d_T)
\end{aligned}$$

Assume $V_{t+1}(i_{t+1}, d_{t+1}) \geq V_{t+1}(proj_L(i_{t+1}), d_{t+1}) \ \forall\ i_{t+1} \in I$ and $d_{t+1} \in D$. Consider period t. For $i_t \in I \setminus \hat{I}$ and $d_t \in D$, $Z_t(i_t, d_t, i_{t+1}^\pi(i_t, d_t)) = Z_t(i_t, d_t, i_{t+1}^{\pi_p}(i_t, d_t))$. By Lemma 1 and the assumption that f is non-decreasing, for $i_t \in \hat{I}$ and $d_t \in D$,

$$\begin{aligned}
Z_t(i_t, d_t, i_{t+1}^\pi(i_t, d_t)) &= f(d(i_t, i_{t+1}^\pi(i_t, d_t))) + \sum_{d_{t+1} \in D} p_t(d_t, d_{t+1}) V_{t+1}(i_{t+1}^\pi(i_t, d_t), d_{t+1}) \\
&\geq f(d(i_t, proj_L(i_{t+1}^\pi(i_t, d_t)))) \\
&\qquad + \sum_{d_{t+1} \in D} p_t(d_t, d_{t+1}) V_{t+1}(proj_L(i_{t+1}^\pi(i_t, d_t)), d_{t+1}) \\
&= Z_t(i_t, d_t, i_{t+1}^{\pi_p}(i_t, d_t))
\end{aligned}$$

Thus, $i_{t+1}^{\pi_p}$ is also an optimal policy function for period t. By Lemmas 1 and 2 and the assumption that $f(d(\cdot,\cdot))$ and $g_j(d(\cdot,\cdot),\cdot)$ are non-decreasing in $d(\cdot,\cdot)$ $\forall j \in J$, for $i_t \in I$ and $d_t \in D$,

$$\begin{aligned}
V_t(i_t, d_t) &= G(i_t, d_t) + f(d(i_t, i_{t+1}^{\pi}(i_t, d_t))) + \sum_{d_{t+1}\in D} p_t(d_t, d_{t+1}) V_{t+1}(i_{t+1}^{\pi}(i_t, d_t), d_{t+1}) \\
&\geq G(proj_L(i_t), d_t) + f(d(proj_L(i_t), proj_L(i_{t+1}^{\pi}(i_t, d_t)))) \\
&\qquad + \sum_{d_{t+1}\in D} p_t(d_t, d_{t+1}) V_{t+1}(proj_L(i_{t+1}^{\pi}(i_t, d_t)), d_{t+1}) \\
&= G(proj_L(i_t), d_t) + Z_t(proj_L(i_t), d_t, proj_L(i_{t+1}^{\pi}(i_t, d_t))) \\
&\geq G(proj_L(i_t), d_t) + Z_t(proj_L(i_t), d_t, i_{t+1}^{\pi}(proj_L(i_t), d_t)) \qquad (3.23) \\
&= V_t(proj_L(i_t), d_t)
\end{aligned}$$

where line 3.23 follows from the fact that

$$i_{t+1}^{\pi}(proj_L(i_t), d_t) = \underset{i_{t+1}\in I}{\arg\min}\, Z_t(proj_L(i_t), d_t, i_{t+1})$$

By mathematical induction, $i_{t+1}^{\pi_p}$ is an optimal policy function for $t \in \{0, ..., T-1\}$ and $V_t(i_t, d_t) \geq V_t(proj_L(i_t), d_t)$ $\forall$ $i_t \in I$, $d_t \in D$, and $t \in \{0, ..., T\}$. Thus,

$$\begin{aligned}
V_0(i_0^{\pi}, d_0) &\geq V_0(proj_L(i_0^{\pi}), d_0) \\
&= V_0(i_0^{\pi_p}, d_0)
\end{aligned}$$

and $\pi_p = \{i_0^{\pi_p}, ..., i_T^{\pi_p}\}$ is an optimal policy satisfying the conditions of the theorem. □

3.5.8 Proof of Theorem 6

Proof. Let $V_t(i_t, d_t)$, $t \in \{0, ..., T\}$, be as in equations 3.1 for the system with movement cost f and let $\hat{V}_t(i_t, d_t)$, $t \in \{0, ..., T\}$, be the corresponding cost-to-go functions for the same system with movement cost $\hat{f}$. Then,

$$V_T(i_T, d_T) = G(i_T, d_T) = \hat{V}_T(i_T, d_T) \quad \forall i_T \in I, d_T \in D$$

Thus, $V_T(i_T, d_T) \leq \hat{V}_T(i_T, d_T)$ $\forall i_T \in I, d_T \in D$. Assume $V_t(i_t, d_t) \leq \hat{V}_t(i_t, d_t)$ $\forall i_t \in I, d_t \in D$. Then,

$$\begin{aligned}
&f(d(i_{t-1}, i_t)) + \sum_{d_t\in D} p_{t-1}(d_{t-1}, d_t) V_t(i_t, d_t) \\
&\qquad \leq \hat{f}(d(i_{t-1}, i_t)) + \sum_{d_t\in D} p_{t-1}(d_{t-1}, d_t) \hat{V}_t(i_t, d_t) \quad \forall i_{t-1}, i_t \in I, d_{t-1} \in D
\end{aligned}$$

and

$$\begin{aligned}
V_{t-1}(i_{t-1}, d_{t-1}) &= G(i_{t-1}, d_{t-1}) + \min_{i_t \in I} \left\{ f(d(i_{t-1}, i_t)) + \sum_{d_t \in D} p_{t-1}(d_{t-1}, d_t) V_t(i_t, d_t) \right\} \\
&\leq G(i_{t-1}, d_{t-1}) + \min_{i_t \in I} \left\{ \hat{f}(d(i_{t-1}, i_t)) + \sum_{d_t \in D} p_{t-1}(d_{t-1}, d_t) \hat{V}_t(i_t, d_t) \right\} \\
&= \hat{V}_{t-1}(i_{t-1}, d_{t-1}) \quad \forall i_{t-1}, i_t \in I, d_{t-1} \in D
\end{aligned}$$

Thus, by induction,

$$\begin{aligned}
V_0(i_0, d_0) &\leq \hat{V}_0(i_0, d_0) \quad \forall i_0 \in I, d_0 \in D \\
\Rightarrow \quad \min_{i_0 \in I} V_0(i_0, d_0) &\leq \min_{i_0 \in I} \hat{V}_0(i_0, d_0) \quad \forall d_0 \in D
\end{aligned}$$

Thus,

$$\min_{\bar{i} \in I_s} \bar{V}_0(\bar{i}, d_0) - \min_{i_0 \in I} V_0(i_0, d_0) \geq \min_{\bar{i} \in I_s} \bar{V}_0(\bar{i}, d_0) - \min_{i_0 \in I} \hat{V}_0(i_0, d_0) \quad \forall d_0 \in D \qquad \square$$

3.5.9 Proof of Theorem 7

Proof. Let $\bar{V}_0(\bar{i}, d_0)$ be as in equation 3.4 for the stationary system with cost to serve G and let $\hat{\bar{V}}_0(\bar{i}, d_0)$ be the corresponding cost of the same system with cost to serve $\hat{G}$. Let $\bar{i}^* = \arg\min_{\bar{i} \in I_s} \bar{V}_0(\bar{i}, d_0)$. Then,

$$\begin{aligned}
\bar{V}_0(\bar{i}^*, d_0) &= G(\bar{i}^*, d_0) + \sum_{d \in D} \left[G(\bar{i}^*, d) \sum_{t=1}^{T} p^t(d_0, d) \right] \\
&\leq G(\bar{i}, d_0) + \sum_{d \in D} \left[G(\bar{i}, d) \sum_{t=1}^{T} p^t(d_0, d) \right] \quad \forall \bar{i} \in I_s
\end{aligned}$$

Multiplying by a and adding $(T+1)b$ to each side,

$$\begin{aligned}
aG(\bar{i}^*, d_0) + b + \sum_{d \in D} &\left[(aG(\bar{i}^*, d) + b) \sum_{t=1}^{T} p^t(d_0, d) \right] \\
&\leq aG(\bar{i}, d_0) + b + \sum_{d \in D} \left[(aG(\bar{i}, d) + b) \sum_{t=1}^{T} p^t(d_0, d) \right] \quad \forall \bar{i} \in I_s \\
&\equiv \hat{\bar{V}}_0(\bar{i}^*, d_0) \leq \hat{\bar{V}}_0(\bar{i}, d_0) \quad \forall \bar{i} \in I_s
\end{aligned}$$

Thus, $\bar{i}^*$ is also optimal for the stationary system with cost to serve $\hat{G}$ and the difference in cost between the two stationary systems is

$$\hat{\bar{V}}_0(\bar{i}^*, d_0) - \bar{V}_0(\bar{i}^*, d_0) = (a-1)G(\bar{i}^*, d_0) + b + \sum_{d \in D} \left[((a-1)G(\bar{i}^*, d) + b) \sum_{t=1}^{T} p^t(d_0, d) \right] \tag{3.24}$$

We now introduce some additional notation. Let $\boldsymbol{d} = (d_1, ..., d_T) \in D \times .... \times D = \boldsymbol{D}$ so that $\boldsymbol{d}$ is a string of demand realizations for time periods 1 through T, $\boldsymbol{D}$ be the set of all possible strings of demand realizations $\boldsymbol{d}$, the probability of seeing the string of demand realizations $\boldsymbol{d} \in \boldsymbol{D}$ be

$$p(\boldsymbol{d}) = \prod_{t=0}^{T-1} p_t(d_t, d_{t+1}),$$

and the location of the mobile inventory at time $t \in \{1, ..., T-1\}$ under policy $\pi \in \Pi$ and string of demand realizations $\boldsymbol{d} \in \boldsymbol{D}$ be represented using the more compact notation:

$$\begin{aligned} i^{\pi}_{1,\boldsymbol{d}} &= i^{\pi}_1(i^{\pi}_0, d_0) \\ i^{\pi}_{t+1,\boldsymbol{d}} &= i^{\pi}_{t+1}(i^{\pi}_{t,\boldsymbol{d}}, d_t) \end{aligned}$$

The cost of responding to disasters over time using the mobile inventory system under policy $\pi \in \Pi$ can be written as

$$V^{\pi}_0 = G(i^{\pi}_0, d_0) + f(d(i^{\pi}_0, i^{\pi}_{1,\boldsymbol{d}'})) + \sum_{\boldsymbol{d} \in \boldsymbol{D}} p(\boldsymbol{d}) \sum_{t=1}^{T-1} \left[G(i^{\pi}_{t,\boldsymbol{d}}, d_t) + f(d(i^{\pi}_{t,\boldsymbol{d}}, i^{\pi}_{t+1,\boldsymbol{d}}))\right] + G(i^{\pi}_{T,\boldsymbol{d}}, d_T)$$

where $\boldsymbol{d}' \in \boldsymbol{D}$. This formulation calculates the probability weighted average cost of all potential strings of demand realizations. Equation 3.24 can be rewritten as

$$\hat{\bar{V}}_0(\bar{i}^*, d_0) - \bar{V}_0(\bar{i}^*, d_0) = (a-1)G(\bar{i}^*, d_0) + b + \sum_{\boldsymbol{d} \in \boldsymbol{D}} p(\boldsymbol{d}) \sum_{t=1}^{T} [(a-1)G(\bar{i}^*, d_t) + b]$$

Let V^{π}_0 be the cost of the mobile inventory system with cost to serve G under policy π and $\hat{V}^{\pi}_0$ be the corresponding cost of the same system with cost to serve $\hat{G}$. Let $\pi^* = \min_{\pi \in \Pi} V^{\pi}_0$. By Theorem 2,

$$V^{\pi^*}_0 \leq \bar{V}_0(\bar{i}^*, d_0)$$

Using this and the assumption $f(y) \geq 0 \; \forall y \geq 0$,

$$G(i^{\pi^*}_0, d_0) + \sum_{\boldsymbol{d} \in \boldsymbol{D}} p(\boldsymbol{d}) \sum_{t=1}^{T} G(i^{\pi^*}_{t,\boldsymbol{d}}, d_t) \leq \bar{V}_0(\bar{i}^*, d_0) = G(\bar{i}^*, d_0) + \sum_{\boldsymbol{d} \in \boldsymbol{D}} p(\boldsymbol{d}) \sum_{t=1}^{T} G(\bar{i}^*, d_t)$$

Multiplying by $(a-1)$ and adding $(T+1)b$ to each side,

$$\begin{aligned} &(a-1)G(i^{\pi^*}_0, d_0) + b + \sum_{\boldsymbol{d} \in \boldsymbol{D}} p(\boldsymbol{d}) \sum_{t=1}^{T} \left[(a-1)G(i^{\pi^*}_{t,\boldsymbol{d}}, d_t) + b\right] \\ &\qquad \leq (a-1)G(\bar{i}^*, d_0) + b + \sum_{\boldsymbol{d} \in \boldsymbol{D}} p(\boldsymbol{d}) \sum_{t=1}^{T} [(a-1)G(\bar{i}^*, d_t) + b] \end{aligned}$$

$$
\begin{aligned}
\equiv \hat{V}_0^{\pi^*} - V_0^{\pi^*} &\leq \hat{\bar{V}}_0(\bar{i}^*, d_0) - \bar{V}_0(\bar{i}^*, d_0) \\
&= \min_{\bar{i}' \in I_s} \hat{\bar{V}}_0(\bar{i}', d_0) - \min_{\bar{i} \in I_s} \bar{V}_0(\bar{i}, d_0)
\end{aligned}
$$

Furthermore, π^* is feasible, though not necessarily optimal, for the mobile inventory system with cost to serve $\hat{G}$ and thus, recalling that $V_0^{\pi^*} = \min_{\pi \in \Pi} V_0^{\pi}$,

$$
\begin{aligned}
& \min_{\pi' \in \Pi} \hat{V}_0^{\pi'} - \min_{\pi \in \Pi} V_0^{\pi} \leq \min_{\bar{i}' \in I_s} \hat{\bar{V}}_0(\bar{i}', d_0) - \min_{\bar{i} \in I_s} \bar{V}_0(\bar{i}, d_0) \\
\equiv \quad & \min_{\bar{i} \in I_s} \bar{V}_0(\bar{i}, d_0) - \min_{\pi \in \Pi} V_0^{\pi} \leq \min_{\bar{i}' \in I_s} \hat{\bar{V}}_0(\bar{i}', d_0) - \min_{\pi' \in \Pi} \hat{V}_0^{\pi'}
\end{aligned}
$$

□

3.5.10 Proof of Theorem 10

Proof. Since f is non-decreasing in the distance moved, the optimal path will be to move directly from i_0 to i_T. It is sufficient to show that there exists an optimal solution in which the inventory is moved the same distance in each period. Suppose there is no such optimal solution. Then in an optimal solution there is some period t when the inventory moves a distance b and some period t' when the inventory moves a distance $a < b$. The slope of the line through a point $f(x)$ and $f(\frac{a+b}{2})$ is

$$
\frac{f(\frac{a+b}{2}) - f(x)}{\frac{a+b}{2} - x}
$$

which is monotonically non-decreasing in x given the convexity of f. Thus,

$$
\begin{aligned}
& \frac{f\left(\frac{a+b}{2}\right) - f(a)}{\frac{a+b}{2} - a} \leq \frac{f\left(\frac{a+b}{2}\right) - f(b)}{\frac{a+b}{2} - b} \\
\equiv \quad & \frac{f\left(\frac{a+b}{2}\right) - f(a)}{\frac{b-a}{2}} \leq \frac{f(b) - f\left(\frac{a+b}{2}\right)}{\frac{b-a}{2}} \\
\equiv \quad & f\left(\frac{a+b}{2}\right) + f\left(\frac{a+b}{2}\right) \leq f(a) + f(b)
\end{aligned}
$$

Thus, we can decrease the cost (or keep the same cost) by instead moving the same distance $f\left(\frac{a+b}{2}\right)$ in periods t and t'. This argument holds for all differences in movement distances across periods, contradicting our assumption that there is no optimal solution in which the inventory is moved the same distance in each period. □

3.5.11 Proof of Lemma 3

Proof. By assumption, $f(\cdot)$ is convex.

$$\begin{aligned}
h(\lambda\tau_1 + (1-\lambda)\tau_2) &= (\lambda\tau_1 + (1-\lambda)\tau_2) f\left(\frac{\bar{d}}{\lambda\tau_1 + (1-\lambda)\tau_2}\right) \\
&= (\lambda\tau_1 + (1-\lambda)\tau_2) f\left(\frac{\lambda\tau_1}{\lambda\tau_1 + (1-\lambda)\tau_2}\frac{\bar{d}}{\tau_1} + \frac{(1-\lambda)\tau_2}{\lambda\tau_1 + (1-\lambda)\tau_2}\frac{\bar{d}}{\tau_2}\right) \\
&\leq (\lambda\tau_1 + (1-\lambda)\tau_2)\left(\frac{\lambda\tau_1}{\lambda\tau_1 + (1-\lambda)\tau_2} f\left(\frac{\bar{d}}{\tau_1}\right) + \frac{(1-\lambda)\tau_2}{\lambda\tau_1 + (1-\lambda)\tau_2} f\left(\frac{\bar{d}}{\tau_2}\right)\right) \\
&= \lambda\tau_1 f\left(\frac{\bar{d}}{\tau_1}\right) + (1-\lambda)\tau_2 f\left(\frac{\bar{d}}{\tau_2}\right) \\
&= \lambda h(\tau_1) + (1-\lambda) h(\tau_2)
\end{aligned}$$

□

3.5.12 Proof of Theorem 11

Proof. By Theorem 10, the optimal path will be to move directly from i_0 to i_T and there exists an optimal solution in which the distance moved in all periods in which there is non-zero movement will be the same. Since the cost to move function is the same in all periods, the problem becomes determining the number of periods τ in which the inventory should experience non-zero movement. The cost to move as a function of τ is

$$C_f(\tau) := \tau\bar{f} + \tau f\left(\frac{d(i_0, i_T)}{\tau}\right) + (T-\tau)f(0)$$

$C_f(\tau)$ is convex in τ since $\tau f\left(\frac{d(i_0,i_T)}{\tau}\right)$ is convex by Lemma 3 and sums of convex functions are convex. Thus, the integer τ^* that minimizes $C_f(\tau)$ is the smallest $\tau \in \mathbb{Z}^+$ such that $C_f(\tau) \leq C_f(\tau+1)$. □

3.5.13 Proof of Theorem 12

Proof. Since $f_c(r)$ is non-decreasing in r, the optimal path will be to move directly from i_0 to i_T. Thus, this theorem considers the problem of moving the inventory from some starting location on a line, say $x(0)$, to some ending location on the line, $x(T)$, where $x(T) \geq x(0)$ without loss of generality and the distance between $x(0)$ and $x(T)$ is $d(i_0, i_T)$. Using continuous-time optimal control theory, the problem can be written as

$$\min_{r(t)} \int_0^T f_c(r(t))dt$$

where $\dot{x}(t) = r(t)$ and $x(0)$ and $x(T)$ are given. The Hamiltonian is

$$H = f_c(r) + pr$$

The adjoint equation is $\dot{p}(t) = -\bigtriangledown_x H = 0 \Rightarrow p(t)$ is constant $\forall t \in [0,t]$. By the Pontryagin Minimum Principle,

$$r^*(t) = \arg\min_r H = \arg\min_r f_c(r) + p(t)r$$

Since $p(t)$ is constant, f_c is convex and non-decreasing, and $r(t) \geq 0 \ \forall t \in [0,T]$, an optimal solution $r^*(t)$ exists in which $r^*(t)$ is constant $\forall t \in [0,T]$. Since $x(0)$ and $x(T)$ are fixed, $r^*(t) = d(i_0, i_T)/T \ \forall t \in [0,T]$. □

3.5.14 Proof of Theorem 13

Proof. Since $f_c(r)$ is non-decreasing in r, the optimal path will be to move directly from i_0 to i_T. By Theorem 12, there exists an optimal solution in which the rate of movement at all times in which there is non-zero movement will be constant. Thus, the problem becomes determining the amount of time t out of T to spend moving the inventory. The cost to move as a function of t is

$$C_f(t) := t\bar{f}_c + tf_c\left(\frac{d(i_0, i_T)}{t}\right) + (T-t)f_c(0)$$

$C_{f_c}(t)$ is convex in t since $tf_c\left(\frac{d(i_0,i_T)}{t}\right)$ is convex by Lemma 3 and sums of convex functions are convex. Thus, there exists a minimum t^* of $C_{f_c}(t)$ over $[0,T]$. □

3.5.15 Proof of Theorem 16

Proof. Assume the demands are temporally independent. Let

$$p_t(d_{t+1}) = p_t(d_t, d_{t+1}) \quad \forall\ d_t \in D, t \in \{0,1,...,T-1\}$$

It is sufficient to show that

$$\hat{V}_t(i_t, d_t) = V_t(i_t, d_t) \quad \forall\ (i_t, d_t) \in I \times D, t \in \{0,1,...,T\}$$

By definition,

$$\hat{V}_T(i_T, d_T) = V_T(i_T, d_T) \quad \forall\ (i_T, d_T) \in I \times D$$

Assume

$$\hat{V}_{t+1}(i_{t+1}, d_{t+1}) = V_{t+1}(i_{t+1}, d_{t+1}) \quad \forall\ (i_{t+1}, d_{t+1}) \in I \times D$$

Then, for all Base States (i_t, d_t), $\hat{V}_t(i_t, d_t) = V_t(i_t, d_t)$. Let $(i_t, d_t'^*) =$ $\arg\max_{(i_t, d_t') \in S_{(i_t,d_t)}} \nu(i_t, d_t')$ where

$$\begin{aligned}\nu(i_t, d_t') &= G(i_t, d_t') + \min_{i_{t+1}\in I}\left\{f(d(i_t, i_{t+1})) + \sum_{d_{t+1}\in D} p_t(d_t', d_{t+1})\hat{V}_{t+1}(i_{t+1}, d_{t+1})\right\} \\ &= G(i_t, d_t') + \min_{i_{t+1}\in I}\left\{f(d(i_t, i_{t+1})) + \sum_{d_{t+1}\in D} p_t(d_{t+1})V_{t+1}(i_{t+1}, d_{t+1})\right\}\end{aligned}$$

For all non-Base States (i_t, d_t),

$$\begin{aligned}
\hat{V}_t(i_t, d_t) &= \Psi(S_{(i_t,d_t)}) \\
&= G(i_t, d_t) + \max_{(i_t,d'_t)\in S_{(i_t,d_t)}} \nu(i_t, d'_t) - G(i_t, d'^*_t) \\
&= G(i_t, d_t) + \min_{i_{t+1}\in I} \left\{ f(d(i_t, i_{t+1})) + \sum_{d_{t+1}\in D} p_t(d_{t+1}) V_{t+1}(i_{t+1}, d_{t+1}) \right\} \\
&= V_t(i_t, d_t)
\end{aligned}$$

Thus, by induction, $\hat{V}_t(i_t, d_t) = V_t(i_t, d_t) \ \forall \ (i_t, d_t) \in I \times D, t \in \{0, 1, ..., T\}$. $\square$

4 Joint Dynamic Facility Relocation and Inventory Management

An important consideration which has been disregarded thus far is inventory management. As discussed in Section 1, it is suboptimal to consider relocation decisions and inventory management decisions separately. Thus, in this section, we consider making dynamic relocation and inventory management decisions simultaneously. This section is organized as follows. Section 4.1 details the model and analytical results. In Section 4.1.1, we develop intuition on optimal relocation and inventory management policies by considering a special case of the problem. In Section 4.1.2, we return to the general dynamic relocation and inventory management problem and develop results regarding optimal relocation and inventory management policies. The proofs of the results in this section can be found in Section 4.2.

4.1 Model and Results

In this section we describe our model and analytical results for relocating a mobile inventory and managing inventory ordering decisions to respond to disasters over time. This problem is called the dynamic relocation and inventory management problem. It is a sequential decision making problem and we model it using DP. The timing of the problem is as follows: at the beginning of each period, the location of the inventory and the initial inventory level are known and a decision is made on where to move the inventory and how much, if any, new inventory to order from the new location. Relocation and order lead time is considered instantaneous and demand is realized throughout the rest of the time period. See Figure 9 for a visualization of the timing of the problem. We assume a finite horizon with T time periods as it is unrealistic to forecast disasters infinitely into the future and that our objective is to minimize cost.

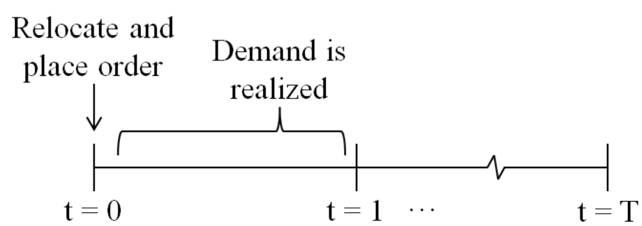

Figure 9: The decision making timeline

For reasons of tractability, we assume that at most one potential disaster site may experience demand in each period. This assumption allows us to disregard allocation issues when demand exceeds supply, simplifying the problem. This is a reasonable assumption as, in most applications, the mobile inventory will serve a particular geographic region which will likely not experience more than one disaster in a given period. Unlike most retail inventory management models, we assume that it is not possible to backlog demand as demand satisfaction is urgent in the disaster relief setting. Instead, unsatisfied demand is considered lost or is served from an outside supplier through a rush order and incurs a corresponding shortage cost penalty. Let $d(\cdot,\cdot)$ be the Euclidean distance metric. Our notation is as follows:

I :	set of potential inventory locations
J :	set of potential disaster sites
S :	set of suppliers
l_j :	location of potential disaster site $j \in J$
l_s :	location of supplier $s \in S$
p_{jt} :	probability potential disaster site $j \in J$ will experience demand in period $t \in \{0,...,T-1\}$ where $\sum_{j \in J} p_{jt} = 1 \; \forall t \in \{0,...,T-1\}$
d_{jt} :	demand level of potential disaster site $j \in J$ if it experiences demand in period $t \in \{0,...,T-1\}$
b_j :	unit shortage cost for potential disaster site $j \in J$
h :	unit holding cost per period
i_0^0 :	initial inventory location, $i_0^0 \in I$
x_0 :	initial inventory on hand, $x_0 \geq 0$
$f(d(i,i'))$:	cost to move the inventory from $i \in I$ to $i' \in I$ in one period
$g_j(d(i,l_j))$:	unit cost to serve demand at potential disaster site $j \in J$ from $i \in I$
$K\left(\min_{s \in S} d(i,l_s)\right)$:	fixed cost to order from $i \in I$
$c\left(\min_{s \in S} d(i,l_s)\right)$:	unit purchase cost from $i \in I$

I may be continuous or discrete (though a few of our results specify specific forms). Assume that J contains a dummy site j' with $d_{j't} = 0 \; \forall t$, $b_{j'} = 0$, and $g_{j'}(\cdot) = 0$; then $p_{j't}$ is the probability that no potential disaster sites experience demand in period t. Note that f may include maintenance or operational costs. Also note that a stationary warehouse owned by the organization could be considered a supplier in S; in this case, the penalty costs b_j could be the cost to serve potential disaster site j from the warehouse rather than the mobile inventory. We assume all costs are greater than or equal to zero. For notational simplicity, we will write $f_{ii'} = f(d(i,i'))$, $g_{ij} = g_j(d(i,l_j))$, $K_i = K(\min_{s \in S} d(i,l_s))$ and $c_i = c(\min_{s \in S} d(i,l_s))$. We can find the minimum cost-to-go when the inventory is located at site $i_t^0 \in I$ with inventory level $x_t \geq 0$ at the beginning of period $t \in \{0,...,T-1\}$ using the following DP equations; these DP equations describe the **dynamic relocation and inventory management (DRIM)** problem:

$$V_t(i_t^0, x_t) = \min_{i_t \in I, y_t \geq x_t} \{f_{i_t^0 i_t} + K_{i_t}\delta(y_t - x_t) + G_t(i_t, y_t) - c_{i_t} x_t\} \tag{4.1}$$

where

$$G_t(i,y) = c_i y + \sum_{j\in J} p_{jt}\big[g_{ij}\min(y,d_{jt}) + h\max(0,y-d_{jt})$$
$$+ b_j\max(0,d_{jt}-y) + V_{t+1}(i,\max(0,y-d_{jt}))\big]$$

and $V_T(i_T^0, x_T) = 0$. The cost-to-go is the cost to relocate and possibly order new inventory plus the expected cost to serve demand and hold remaining inventory or pay for a shortage of inventory plus the expected future period cost. Note that $min(y,d) = d - max(0, d - y)\ \forall y, d \in \mathbb{R}$ and thus we can re-write $G_t(i,y)$ as

$$G_t(i,y) = c_i y + \sum_{j\in J} p_{jt}\Bigg[h\max(0,y-d_{jt}) + g_{ij}d_{jt}$$
$$+ (b_j - g_{ij})\max(0,d_{jt}-y) + V_{t+1}(i,\max(0,y-d_{jt}))\Bigg]$$

We can also write the DRIM problem in terms of policies; this notation will help us describe our results below. Let $\pi = \{\pi_i, \pi_y\} \in \Pi$, where $\pi_i = \{i_0^\pi, ..., i_{T-1}^\pi\}$ and $\pi_y = \{y_0^\pi, ..., y_{T-1}^\pi\}$, be a policy consisting of sequences of functions i_t^π and y_t^π that map states (i_t^0, x_t) into decisions of where to move the inventory, $i_t^\pi(i_t^0, x_t)$, and the order-up-to level, $y_t^\pi(i_t^0, x_t)$, for period t. Let Π be the set of all feasible policies for the DRIM problem described by the DP equations 4.1 and the cost of the mobile inventory system under policy π be represented by $V_0^\pi(i_0^0, x_0)$ where $V_T^\pi(i_T^0, x_T) = 0$,

$$V_t^\pi(i_t^0, x_t) = f_{i_t^0 i_t^\pi(i_t^0,x_t)} + K_{i_t^\pi(i_t^0,x_t)}\delta(y_t^\pi(i_t^0,x_t) - x_t) + G_t^\pi(i_t^\pi(i_t^0,x_t), y_t^\pi(i_t^0,x_t)) - c_{i_t^\pi(i_t^0,x_t)}x_t$$

and

$$G_t^\pi(i,y) = c_i y + \sum_{j\in J} p_{jt}\Bigg[h\max(0,y-d_{jt}) + g_{ij}d_{jt}$$
$$+ (b_j - g_{ij})\max(0,d_{jt}-y) + V_{t+1}^\pi(i,\max(0,y-d_{jt}))\Bigg]$$

The DRIM problem can then be written $min_{\pi\in\Pi} V_0^\pi(i_0^0, x_0)$.

The need to make both dynamic relocation and inventory management decisions makes the DRIM problem very complex. We first gain some intuition by examining a special case of the problem in Section 4.1.1. In Section 4.1.2, we return to the general DRIM problem and develop results characterizing the optimal relocation and inventory management policies with the help of the intuition and results developed in Section 4.1.1.

4.1.1 Special Case

In this section, we develop intuition on optimal relocation and inventory management policies by considering a special case of the DRIM problem. We consider a one-dimensional setting in which we have one supplier and one potential disaster site located L units apart, as illustrated in Figure 10. The set of potential inventory locations is all locations on the line between the supplier and the potential disaster site, i.e. $I = [0, L]$. As there is only one potential disaster site, throughout this section, we will drop the subscript j on p_{jt}, d_{jt}, b_j, and $g_j(\cdot)$. This special case is more tractable than the general DRIM problem and elucidates an important trade-off which is difficult to tease out in the general two-dimensional setting: when is it advantageous to move toward a supplier at the expense of moving farther from potential disaster sites and vice versa.

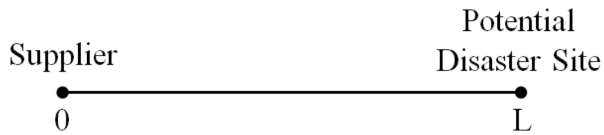

Figure 10: Network of the special case

The DP equations can be written as:

$$\begin{aligned} V_t(i_t^0, x_t) = \min_{i_t \in I, y_t \geq x_t} \Big\{ & f(|i_t - i_t^0|) + K(i_t)\delta(y_t - x_t) + c(i_t)(y_t - x_t) \\ & + p_t\big[g(L - i_t)\min(y_t, d_t) + h\max(0, y_t - d_t) \\ & \qquad + b\max(0, d_t - y_t) + V_{t+1}(i_t, \max(0, y_t - d_t))\big] \\ & + (1 - p_t)\big[hy_t + V_{t+1}(i_t, y_t)\big] \Big\} \end{aligned}$$

where $V_T(i_T^0, x_T) = 0$.

If we assume that $f(\cdot)$ is concave, a reasonable assumption given that greater fuel efficiencies are generally reached at higher speeds, we can prove that there exists an optimal policy such that we will only move toward the supplier in a period in which we also place an order. We use the following Lemma in the proofs of Theorems 17 and 22:

Lemma 4. *If f is a concave function, then $f(y_1) + f(y_2) \geq f(y_1 + y_2)$ for y_1, $y_2 \in \mathbb{R}$*

Theorem 17. *Assume $g(\cdot)$ is non-decreasing, $f(\cdot)$ is concave and non-decreasing over $(0, \infty)$, and $f(0) \leq \lim_{d \to 0^+} f(d)$. $\exists$ an optimal policy $\pi^* = \{\pi_i, \pi_y\} \in \Pi$, where $\pi_i = \{i_0^\pi, ..., i_{T-1}^\pi\}$ and $\pi_y = \{y_0^\pi, ..., y_{T-1}^\pi\}$, such that if $y_t^\pi(i_t^0, x_t) = x_t$ then $i_t^\pi(i_t^0, x_t) \in [i_t^0, L]$ $\forall$ $i_t^0 \in I$, $x_t \geq 0$, and $t \in \{0, ..., T-1\}$.*

By admitting a discontinuity in $f(\cdot)$ at 0, the conditions of Theorem 17 allow for a fixed cost to move. Theorem 17 assures that there will not be a period in which we move the inventory towards the supplier (and away form the potential disaster site) without also placing an order. If f is not concave, this result does not necessarily hold as cost savings may be achieved by moving towards a better ordering position over several periods rather than all at once when placing an order. If we also assume that $K(\cdot)$, $c(\cdot)$, and $g(\cdot)$ are concave, a reasonable assumption due to ordering and shipping economies of scale, we can also prove the following:

Theorem 18. *Assume $K(\cdot)$, $c(\cdot)$, and $g(\cdot)$ are concave and non-decreasing, $f(\cdot)$ is concave and non-decreasing over $(0,\infty)$, and $f(0) \leq \lim_{d\to 0^+} f(d)$. $\exists$ an optimal policy $\pi^* = \{\pi_i, \pi_y\} \in \Pi$, where $\pi_i = \{i_0^\pi, ..., i_{T-1}^\pi\}$ and $\pi_y = \{y_0^\pi, ..., y_{T-1}^\pi\}$, such that $\forall\ i_t^0 \in I$, $x_t \geq 0$, and $t \in \{0, ..., T-1\}$:*

1. *$i_t^\pi(i_t^0, x_t) \in \{0, i_t^0, L\}$*
2. *if $i_t^\pi(i_t^0, x_t) = L$ and $x_t \geq d_t$ then $y_t^\pi(i_t^0, x_t) = x_t$*
3. *if $y_t^\pi(i_t^0, x_t) = x_t$ and*
$$p_t \left(g(L - i_t^0) - g(0)\right) \min(x_t, d_t) \geq f(L - i_t^0) + f(L) \tag{4.2}$$
then $i_t^\pi(i_t^0, x_t) = L$
4. *if $y_t^\pi(i_t^0, x_t) > x_t$ and*
$$\begin{aligned} K(i) - K(0) + (c(i) - c(0)) &\left(y_t^\pi(i_t^0, x_t) - x_t\right) \\ &\geq p_t \left(g(L) - g(L - i)\right) \min\left(y_t^\pi(i_t^0, x_t), d_t\right) + f(i_t^0) + f(L) \end{aligned} \tag{4.3}$$
for $i \in \{i_t^0, L\}$, then $i_t^\pi(i_t^0, x_t) = 0$

Theorem 18.1 assures that there exists on optimal policy such that the inventory is always located at 0, i_0^0, or L; even more, once the inventory moves to 0 or L, it will in all periods afterward only be located at either 0 or L. Due to the concavity of the cost functions, there is no additional value in considering any other inventory location; for any ordering decision, the cost of a decision to move the inventory in one direction can always be improved upon by moving the inventory completely to the endpoint of $I = [0, L]$ in that direction. Corollary 4 follows directly from Theorems 17 and 18.1:

Corollary 4. *Assume $K(\cdot)$, $c(\cdot)$, and $g(\cdot)$ are concave and non-decreasing, $f(\cdot)$ is concave and non-decreasing over $(0,\infty)$, and $f(0) \leq \lim_{d\to 0^+} f(d)$. $\exists$ an optimal policy $\pi^* = \{\pi_i, \pi_y\} \in \Pi$, where $\pi_i = \{i_0^\pi, ..., i_{T-1}^\pi\}$ and $\pi_y = \{y_0^\pi, ..., y_{T-1}^\pi\}$, such that if $y_t^\pi(i_t^0, x_t) = x_t$ then $i_t^\pi(i_t^0, x_t) \in \{i_t^0, L\}$ $\forall\ i_t^0 \in I$, $x_t \geq 0$, and $t \in \{0, ..., T-1\}$.*

By Corollary 4, if we do not place an order in a given period, we will either keep the inventory in the same location or move it to L, the location of the potential disaster site; if the initial inventory location is L and we do not place an order, we will not move the inventory.

Theorem 18.2 assures that in period t we will only possibly order from position L if our initial inventory level is less than d_t; that is, we will not place an order from position L if our initial inventory level is greater than or equal to d_t. This makes intuitive sense; if the inventory is located at L, it is as far as possible from the supplier and ordering costs will only possible decrease due to changes in location in the next period. Any inventory ordered in excess of d_t in that period is not needed for that period and can be ordered in the next period instead for potentially a lower cost.

Theorem 18.3 gives a condition under which we move the inventory to L when we do not place an order. It assures that in a given period if we do not place an order and 4.2 holds, we will move the inventory to L. Condition 4.2 holds when the expected difference in cost between serving demand from the starting location i_t^0 and serving demand from L is greater than or equal to the cost to move from i_t^0 to L plus the maximum cost to move from L to another location in the network. That is, if the expected savings in serving demand from L exceeds the cost to move to L plus the cost to move to potentially another location in the next period, we will move the inventory to L.

Theorem 18.4 further characterizes optimal relocation and inventory management policies for the special case by giving a condition under which we move the inventory to 0, the location of the supplier, when placing an order. It assures that in a given period if we place an order and 4.3 holds for $i \in \{i_t^0, L\}$, we will move the inventory to 0. For $i \in \{i_t^0, L\}$, condition 4.3 holds when the difference in cost between ordering from i and ordering from 0 is greater than or equal to the expected increase in cost between serving demand from 0 and serving demand from i plus the cost to move from i_t^0 to 0 plus the maximum cost to move from 0 to another location in the network. That is, if the savings from ordering from 0 exceeds the expected increase in cost from serving demand from 0 plus the cost to move to 0 plus the cost to move to potentially another location in the next period, we will move the inventory to 0.

Theorems 17 and 18.1 through 18.4 and Corollary 4 help to characterize optimal inventory relocation and ordering policies in this special case. In Section 4.1.2, we extend a few of these results to the general case of the DRIM problem and develop other results regarding optimal relocation and inventory management policies.

4.1.2 Dynamic Relocation and Inventory Management Problem

In this section, we return to the general DRIM problem described by the DP equations 4.1 and develop results regarding optimal relocation and inventory management policies.

We first examine characteristics of optimal inventory management decisions. Consider the case where, for staffing or other planning reasons, we are restricted to movement policies that do not depend on the inventory level. In this case, our movement policy consists of a sequence of locations $\{i_0, ..., i_{T-1}\}$ detailing the path for the mobile inventory to follow throughout the horizon. With the following result, we show that a multiperiod (s,S) policy is optimal in this special case. A **multiperiod** (s, S) **policy** is a policy consisting of scalars s_t and S_t $\forall t \in \{0, ..., T-1\}$; in each period t if the initial inventory on hand is less than or equal to s_t, we place an order to bring the inventory up to S_t, and do not place an order otherwise. Define the following:

Assumption 2. *Assume we are restricted to movement policies that do not depend on the inventory level. Additionally, assume* $b_j - g_{ij} > c_{i'}$ $\forall i, i' \in I$ *and* $j \in J$, $K_i = K$ $\forall i \in I$ *for some* $K \geq 0$, *and* $c_i + h > 0$ $\forall i \in I$.

Assumption 2 requires that our movement policies do not depend on the inventory level, the effective penalty cost is greater than the purchase cost for all locations, the fixed order cost is the same for all locations, and the purchase and holding cost is greater than 0 for all locations. We can prove the following:

Theorem 19. *Under Assumption 2, a multiperiod (s, S) policy is optimal.*

The proof of Theorem 19 requires showing certain convexity properties of the cost-to-go functions (specifically K-convexity; see, for example, Snyder and Shen (2011) for the definition and some basic properties of K-convex functions) so that a simple (s, S) policy is optimal. When considering the general case of the DRIM problem, represented by DP equations 4.1, however, these convexity properties no longer hold. In an (s, S) policy, in each period t, the inventory management policy always remains the same no matter what the initial inventory level: if the inventory level is less than s_t, order up to S_t, and do not order otherwise. However, in the DRIM problem, different initial inventory levels affect our movement decisions and different movement decisions correspond to different ordering costs and thus different ordering policies. Example 2 illustrates this.

Example 2. *Consider an instance of the DRIM problem in which there are two potential inventory locations ($I = \{1, 2\}$), two potential disaster sites ($J = \{1, 2\}$), and two periods in the planning horizon ($T = 2$). Assume $i_0^0 = 1$, $f_{1,2} = f_{2,1} = 16$, $f_{1,1} = f_{2,2} = 0$, $K_1 = K_2 = 0$, $c_1 = 1.5$ and $c_2 = 0.5$, $h = 1$, $b_1 = b_2 = 4$, and the remaining parameters are as follows:*

g_{ij}	1	2
1	1	2.5
2	2	0.5

p_{jt}	0	1
1	0.5	0.5
2	0.5	0.5

d_{jt}	0	1
1	25	5
2	5	20

Figure 11 shows the optimal amount to order and minimum cost by location decision and the overall optimal inventory policy for varying inventory levels for periods $t \in \{0, 1\}$. The first column shows the figures relevant to starting in period 0 in location 1. From the first graph, we see that if we remain in location 1, the optimal ordering policy is to order up to 5 if the inventory is less than 5, not place an order if the inventory is between 5 and 5.5, order up to 10 if the inventory is between 5.5 and 10, and not place an order otherwise. If we move to location 2, the optimal ordering policy is to order up to 25 if the inventory is less than 25 and to not order otherwise. From the second graph, we see that our optimal movement policy is to move to location 2 if the initial inventory level is less than or equal to 14.1 and to remain in location 1 otherwise. The third graph gives the resultant optimal inventory management policy: order up to 25 if the inventory is less than or equal to 14.1 and do not place an order otherwise.

The second and third columns show the same graphs for period $t = 1$ for the two possible starting locations. From the first graph for starting in location 1, we see that if we remain in location 1, the optimal ordering policy is to order up to 5 if the inventory is less than 5 and to not order otherwise. If we move to location 2, the optimal ordering policy is to order up to 20 if the inventory is less than 20 and to not order otherwise. From the second graph, we see that our optimal movement policy is to remain in location 1 if the initial inventory *level is between* 2.8 *and* 14 *and to move to location* 2 *otherwise. The third graph gives the resultant optimal inventory management policy: order up to* 20 *if the inventory is less than or equal to* 2.8, *order up to* 5 *if the inventory is between* 2.8 *and* 5, *do not place an order if the inventory is between* 5 *and* 14, *order up to* 20 *if the inventory is between* 14 *and* 20, *and do not place an order otherwise. Thus, an* (s, S) *policy is not optimal in period* 1 *if we start in location* 1.

Example 2 also illustrates a counter-intuitive managerial insight. One may expect that the lower our initial inventory level, the more we will want to order. However, the dynamic location decision makes this not true. Consider starting in period 1 in location 1 with starting inventory level x and reference again the optimal ordering policy displayed in the third graph in column 2 of Figure 11. We see that if $2.8 \leq x < 5$ we order $5 - x$, if $5 \leq x < 14$ we order 0, and if $14 \leq x < 20$ we order $20 - x$. This is because if $2.8 \leq x < 14$, we remain in location 1 and we move to location 2 otherwise and each location has it's own ordering costs and expected future period costs resulting in different optimal ordering decisions. In this case, these differences in optimal ordering decisions create a counterintuitive optimal ordering policy in which we do not necessarily order less with a greater initial inventory level.

While an (s, S) policy is not necessarily optimal for the general case of the DRIM problem as illustrated by Example 2, we can prove a few results regarding the structure of optimal inventory management policies. Theorem 20 assures that there exists an optimal policy such that, in any state, we either order nothing or order up to some sum of potential future period demands. Corollary 5 follows directly from Theorem 20 and assures that if our initial inventory level is weakly greater than the maximum possible demand for the rest of the horizon, we do not place an order.

Theorem 20. *$\exists$ an optimal policy $\pi^* = \{\pi_i, \pi_y\} \in \Pi$, where $\pi_y = \{y_0^\pi, ..., y_{T-1}^\pi\}$, such that $y_t^\pi(i_t^0, x_t) = x_t$ or $y_t^\pi(i_t^0, x_t) = \sum_{\tau=t}^{T-1} d_{j_\tau \tau}$, for some $j_\tau \in J$ $\forall \tau \in \{t, ..., T-1\}$, $\forall i_t^0 \in I$, $x_t \geq 0$, and $t \in \{0, ..., T-1\}$.*

Corollary 5. *Let $d_t^m = \max_j d_{jt}$. $\exists$ an optimal policy $\pi^* = \{\pi_i, \pi_y\} \in \Pi$, where $\pi_y = \{y_0^\pi, ..., y_{T-1}^\pi\}$, such that $y_t^\pi(i_t^0, x_t) = x_t$ $\forall$ $x_t \geq \sum_{\tau=t}^{T-1} d_\tau^m$, $i_t^0 \in I$, and $t \in \{0, ..., T-1\}$.*

Theorem 20 and Corollary 5 allow us to consider a smaller set of possible ordering decisions when determining the optimal policy which may reduce the computation time required to solve the problem. That is, for each period t and each starting inventory location i_t and level x_t, we only need to consider ordering nothing and ordering up to the $|J|^{T-t}$ possible sums of potential future period demands instead of considering the infinite number of possible decisions $y_t \geq x_t$.

We can also develop a result analogous to Theorem 5 in Section 3.1 for the dynamic relocation problem that helps to characterize the optimal relocation policy for the DRIM problem. Under a few basic assumptions, this result states that it is sufficient to consider a smaller feasible set of inventory locations. This allows us to reduce the size of the problem and thus the time it takes to solve the problem which is worthwhile as realistic instances will have a large state space and thus will take a long time to solve. Let $conv(C)$ and $Int(C)$ denote the convex hull and interior of a set C, respectively, and define the following:

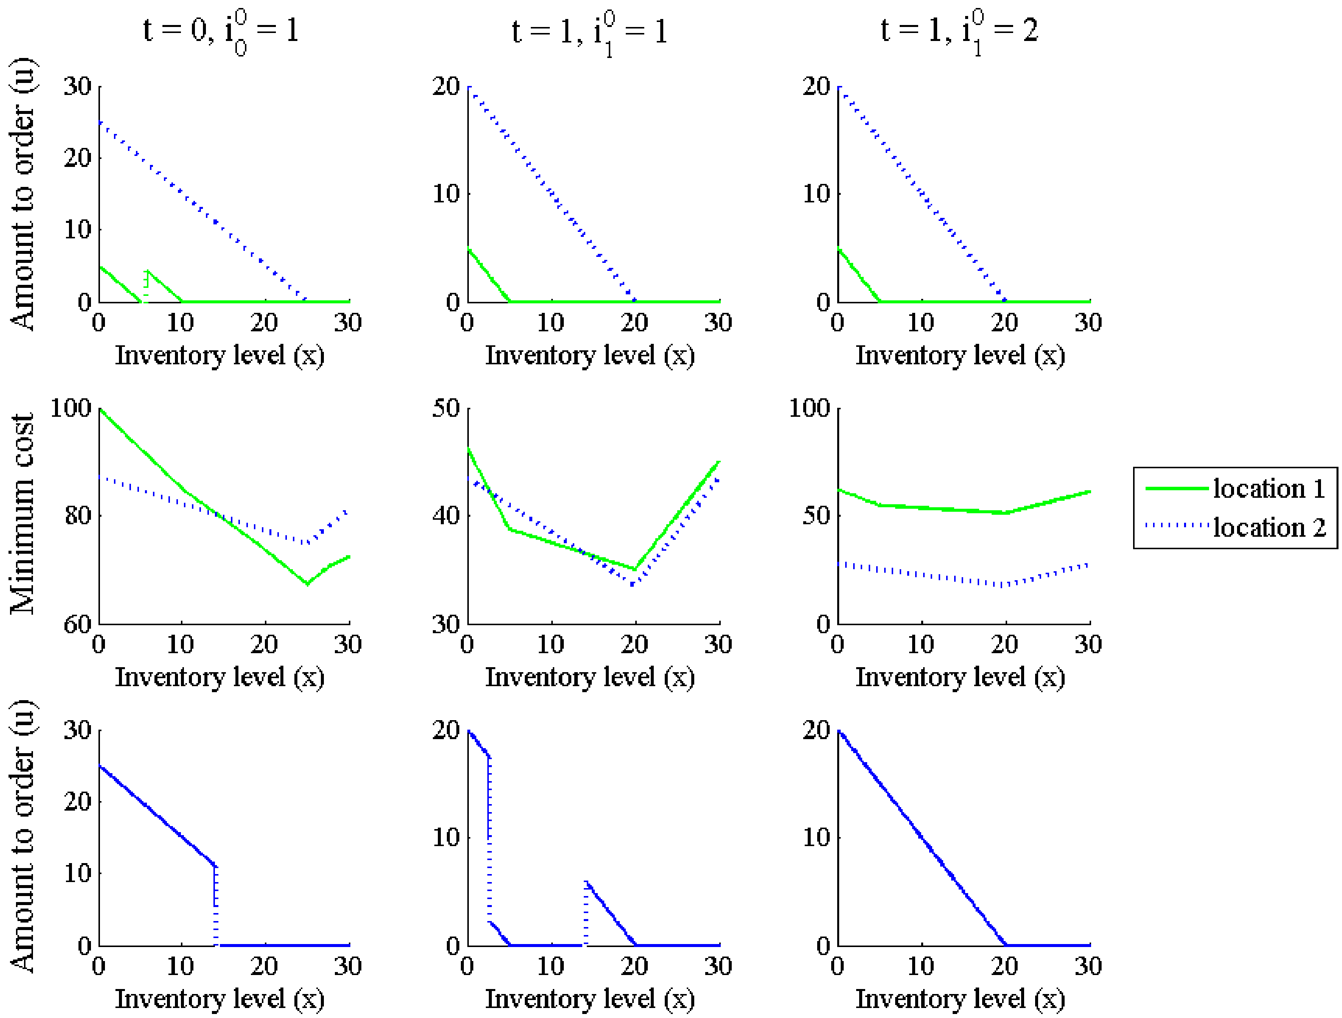

Figure 11: Example 2 amount to order and minimum cost by location and overall optimal inventory policy for varying starting inventory levels for periods $t \in \{0, 1\}$. From the graph of the overall optimal inventory policy (third row), we see that an (s, S) policy is not optimal in period 1 if we start in location 1.

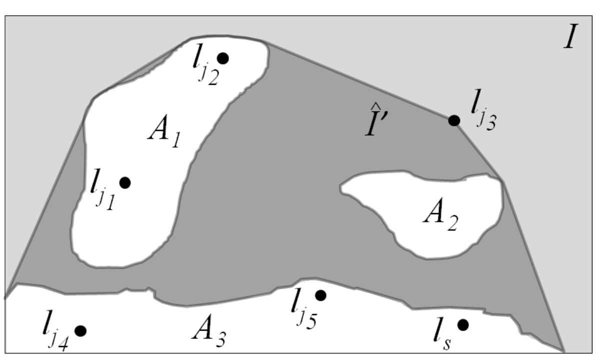

Figure 12: Example network satisfying Assumption 3. I consists of the shaded regions, the white regions A_1, A_2, and A_3 are infeasible regions, $\hat{B}_s$ is $A_1 \cup A_2 \cup A_3$, and $\hat{I}'$ is the dark gray shaded region. Under Assumption 3, Theorem 21 assures that there exists an optimal solution to the dynamic relocation-inventory problem such that the inventory is always located within $\hat{I}'$.

Assumption 3. *Assume $I \subseteq \mathbb{R}^2$ is connected and closed and $f(\cdot)$, $g_j(\cdot)$, $K(\cdot)$, and $c(\cdot)$ are non-decreasing $\forall j \in J$. Let $B = \{A_1, ..., A_l\}$ be a set of finitely many mutually disjoint, connected, closed, and bounded subsets of $\mathbb{R}^2$ such that $i \notin I \ \forall \ i \in Int(A_n)$ and $A_n \in B$, $conv(I \cup \{l_j | j \in J \cup S\} \cup i_0^0) \subseteq I \cup B$, and $I \cup B$ is convex.*

Furthermore, let $\hat{B}_s = \{A_n \in B | \{A_n \cap conv(\{l_j | j \in J \cup S\} \cup i_0^0) \neq \emptyset\} \vee \{A_n \cap conv(\hat{B}_s \backslash A_n) \neq \emptyset\}\}$ and $\hat{I}' = I \cap conv(\hat{B}_s \cup \{l_j | j \in J \cup S\} \cup i_0^0)$. See Figure 12 for an illustration of a network satisfying Assumption 3. Assumption 3 defines the set B of all infeasible regions for the mobile inventory, or areas where the inventory cannot be located (e.g. land in the case of inventory on a ship). Theorem 21 assures that it is sufficient for the DRIM problem to consider only inventory locations within $\hat{I}'$, the feasible inventory locations within the convex hull of the potential disaster sites, the supplier locations, and the initial inventory location extended to include any overlapping infeasible regions (see the dark gray shaded region in Figure 12 for an example):

Theorem 21. *Under Assumption 3, $\exists$ an optimal policy $\pi^* = \{\pi_i, \pi_y\} \in \Pi$, where $\pi_i = \{i_0^\pi, ..., i_{T-1}^\pi\}$, such that $i_t^\pi(i_t^0, x_t) \in \hat{I}' \ \forall \ i_t^0 \in \hat{I}'$, $x_t \geq 0$, and $t \in \{0, ..., T-1\}$.*

Theorems 19 and 20 and Corollary 5 help to characterize the optimal ordering policy and Theorem 21 helps to characterize the optimal relocation policy. It is also worthwhile to understand how the optimal relocation and ordering policies relate to each other. To this end, we extend a few of the results of Section 4.1.1 to the general DRIM problem. Specifically, we extend Theorems 17 and 18.2 to Theorems 22 and 23, respectively, below.

Theorem 17 assures that in the special case there exists an optimal policy such that we will not move the mobile inventory toward the supplier, and away from the potential disaster site, in a period in which we do not place an order. We can extend this result to the general case of the DRIM problem. The analogous result states that we will not move the mobile inventory toward the supplier at the expense of moving it farther from all potential disaster sites in a period in which we do not place an order. Figure 13 illustrates this result.

Theorem 22. *Assume $g(\cdot)$ is non-decreasing, $f(\cdot)$ is concave and non-decreasing over $(0, \infty)$, and $f(0) \leq \lim_{d \to 0^+} f(d)$. $\exists$ an optimal policy $\pi^* = \{\pi_i, \pi_y\} \in \Pi$, where $\pi_i = \{i_0^\pi, ..., i_{T-1}^\pi\}$ and $\pi_y = \{y_0^\pi, ..., y_{T-1}^\pi\}$, such that if $y_t^\pi(i_t^0, x_t) = x_t$ then $\exists j \in J$ such that $d(i_t^\pi(i_t^0, x_t), l_j) \leq d(i_t^0, l_j) \ \forall \ i_t^0 \in I$, $x_t \geq 0$, and $t \in \{0, ..., T-1\}$.*

Recall that Theorem 20 and Corollary 5 allow us to consider a smaller set of possible ordering actions which may reduce the computation time required to solve the problem; we can further reduce the set of possible ordering decisions by extending Theorem 18.2 from the special case to the general case of the DRIM problem. In the special case, this result assures that we will not place an order from position L if our initial inventory is greater than or equal to d_t; this is because we do not need any more inventory to satisfy the current period's demand and our ordering costs will only possibly decrease in the next period as we are currently located as far as possible from the supplier. To extend this result to the general case of the DRIM problem, we need to classify which potential inventory locations

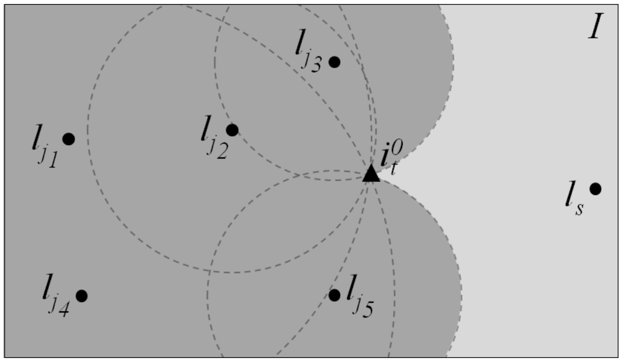

Figure 13: Example network illustrating Theorem 22. The triangle represents the initial mobile inventory location i_t^0 and the black dots represent the potential disaster sites and the supplier location. Each potential disaster site has a dark gray circle around it indicating all locations at least as close to it as i_t^0. Theorem 22 assures that in this example network there exists an optimal solution such that we will not move anywhere within the light gray area if we do not place an order.

are farthest from a supplier. Let $I_m = \{i | i \in I, \min_{s\in S} d(i, l_s) = \max_{i'\in I}(\min_{s\in S} d(i', l_s))\}$. I_m is the set of all potential inventory locations farthest from a supplier. Theorem 23 assures that there exists an optimal policy such that we will not place an order from position $i \in I_m$ if our current inventory level is weakly greater than the maximum possible current period demand.

Theorem 23. *Assume $K(\cdot)$ and $c(\cdot)$ are non-decreasing. $\exists$ an optimal policy $\pi^* = \{\pi_i, \pi_y\} \in \Pi$, where $\pi_i = \{i_0^\pi, ..., i_{T-1}^\pi\}$ and $\pi_y = \{y_0^\pi, ..., y_{T-1}^\pi\}$, such that if $i_t^\pi(i_t^0, x_t) \in I_m$ then $y_t^\pi(i_t^0, x_t) = x_t \ \forall \ x_t \geq \max_{j\in J} d_{jt}$, $i_t^0 \in I$, and $t \in \{0, ..., T-1\}$.*

Thus, from a position $i \in I_m$, we will only possibly place an order if the current inventory level is not enough to satisfy the current period's potential demand. This makes intuitive sense; if the inventory is located in I_m, it is as far as possible from a supplier and ordering costs will only possibly decrease due to changes in location in subsequent periods. Any order placed when the initial inventory is greater than $\max_{j\in J} d_{jt}$ is not needed for that period and can be postponed resulting in potential holding cost and ordering cost savings.

4.2 Proofs

4.2.1 Relevant Notation and Concepts

We will use aspects of the following in the proofs of Theorems 17, 18, and 20 through 23.

$$Z_t(i_t^0, x_t, i_t, y_t) = f_{i_t^0 i_t} + K_{i_t}\delta(y_t - x_t) + G_t(i_t, y_t) - c_{i_t} x_t$$

Then,

$$V_t(i_t^0, x_t) = \min_{i_t \in I, y_t \geq x_t} Z_t(i_t^0, x_t, i_t, y_t)$$

and relocation and order policies i_t^π and y_t^π are optimal for period t if they satisfy

$$\left(i_t^\pi(i_t^0, x_t), y_t^\pi(i_t^0, x_t)\right) = \operatorname*{arg\,min}_{i_t \in I, y_t \geq x_t} Z_t(i_t^0, x_t, i_t, y_t) \ \forall i_t^0 \in I, x_t \geq 0$$

Let

$$\boldsymbol{D}^t = \left\{\boldsymbol{d} = (d_{j_t t}, ..., d_{j_{T-1} T-1}) \mid j_\tau \in J \ \forall \tau \in \{t, ..., T-1\}\right\}$$

be the set of all possible strings of future demand realizations starting in period t. Let the probability of seeing the vector of demand realizations $\boldsymbol{d} \in \boldsymbol{D}^t$ be

$$p(\boldsymbol{d}) = \prod_{\tau=t}^{T-1} p_{j_\tau \tau}$$

Let i_t^* be some predefined location for the inventory in period t. Let the inventory level at the beginning of the period, the location of the mobile inventory throughout the period, and the order up to level of the period under policy $\pi \in \Pi$ and vector of demand realizations $\boldsymbol{d} \in \boldsymbol{D}^t$ for periods $\tau \in \{t+1, ..., T-1\}$ be denoted $x_{\tau \boldsymbol{d}}^\pi$, $i_{\tau \boldsymbol{d}}^\pi$, and $y_{\tau \boldsymbol{d}}^\pi$, respectively, where

$$\begin{aligned}
&i_{t\boldsymbol{d}}^\pi = i_t^* \\
&x_{t+1\boldsymbol{d}}^\pi = \max(0, y_t^* - d_{jt}) \\
&i_{\tau\boldsymbol{d}}^\pi = i_\tau^\pi(i_{\tau-1\boldsymbol{d}}^\pi, x_{\tau\boldsymbol{d}}^\pi) \\
&y_{\tau\boldsymbol{d}}^\pi = y_\tau^\pi(i_{\tau-1\boldsymbol{d}}^\pi, x_{\tau\boldsymbol{d}}^\pi) \\
&x_{\tau\boldsymbol{d}}^\pi = \max(y_{\tau-1\boldsymbol{d}}^\pi - d_{j_{\tau-1}\tau-1}, 0)
\end{aligned}$$

4.2.2 Proof of Lemma 4

Proof. Since f is concave,

$$\begin{aligned}
\frac{y_2}{y_1+y_2} f(0) + \frac{y_1}{y_1+y_2} f(y_1+y_2) &\leq f\left(\frac{y_2}{y_1+y_2}(0) + \frac{y_1}{y_1+y_2}(y_1+y_2)\right) \\
\Rightarrow \frac{y_1}{y_1+y_2} f(y_1+y_2) &\leq f(y_1)
\end{aligned} \tag{4.4}$$

Similarly,

$$\frac{y_2}{y_1+y_2} f(y_1+y_2) \leq f(y_2) \tag{4.5}$$

By 4.4 and 4.5,

$$\begin{aligned}
\frac{y_1}{y_1+y_2} f(y_1+y_2) + \frac{y_2}{y_1+y_2} f(y_1+y_2) &\leq f(y_1) + f(y_2) \\
\equiv f(y_1+y_2) &\leq f(y_1) + f(y_2)
\end{aligned}$$

□

4.2.3 Proof of Theorem 17

Proof. Consider period $T-1$. For $i_{T-1}^0 \in I$ and $x_{T-1} \geq 0$, let

$$(i_{T-1}^*, y_{T-1}^*) = \underset{i_{T-1}\in I, y_{T-1}\geq x_{T-1}}{\arg\min} Z_{T-1}(i_{T-1}^0, x_{T-1}, i_{T-1}, y_{T-1})$$

Set $y_{T-1}^\pi(i_{T-1}^0, x_{T-1}) = y_{T-1}^*$. If $y_{T-1}^\pi(i_{T-1}^0, x_{T-1}) > x_{T-1}$ or $y_{T-1}^\pi(i_{T-1}^0, x_{T-1}) = x_{T-1}$ and $i_{T-1}^* \geq i_{T-1}^0$, set $i_{T-1}^\pi(i_{T-1}^0, x_{T-1}) = i_{T-1}^*$. Otherwise, $y_{T-1}^\pi(i_{T-1}^0, x_{T-1}) = y_{T-1}^* = x_{T-1}$ and $i_{T-1}^* < i_{T-1}^0$; for $i \leq i_{T-1}^0$, let

$$\begin{aligned} F(i) =& f(i_{T-1}^0 - i) + p_{T-1}[g(L-i)\min(y_{T-1}^*, d_{T-1}) + h\max(0, y_{T-1}^* - d_{T-1}) \\ & + b\max(0, d_{T-1} - y_{T-1}^*)] + (1-p_{T-1})hy_{T-1}^* \end{aligned}$$

Note that $F(i_{T-1}^*) = Z_{T-1}(i_{T-1}^0, x_{T-1}, i_{T-1}^*, y_{T-1}^*)$ and $F(i)$ is weakly decreasing over $[0, i_{T-1}^0]$. Thus,

$$Z_{T-1}(i_{T-1}^0, x_{T-1}, i_{T-1}^*, y_{T-1}^*) = F(i_{T-1}^*) \geq F(i_{T-1}^0) = Z_{T-1}(i_{T-1}^0, x_{T-1}, i_{T-1}^0, y_{T-1}^*)$$

Set $i_{T-1}^\pi(i_{T-1}^0, x_{T-1}) = i_{T-1}^0$. Defined in this way, i_{T-1}^π and y_{T-1}^π are optimal policies for period $T-1$ such that if $y_{T-1}^\pi(i_{T-1}^0, x_{T-1}) = x_{T-1}$, then $i_{T-1}^\pi(i_{T-1}^0, x_{T-1}) \in [i_{T-1}^0, L]$ $\forall$ $i_{T-1}^0 \in I$ and $x_{T-1} \geq 0$.

Assume i_τ^π and y_τ^π are optimal policies for periods $\tau \in \{t+1, ..., T-1\}$ such that if $y_\tau^\pi(i_\tau^0, x_\tau) = x_\tau$ then $i_\tau^\pi(i_\tau^0, x_\tau) \in [i_\tau^0, L]$ $\forall$ $i_\tau^0 \in I$, $x_\tau \geq 0$, and $\tau \in \{t+1, ..., T-1\}$.

Consider period t. For $i_t^0 \in I$ and $x_t \geq 0$, let

$$(i_t^*, y_t^*) = \underset{i_t\in I, y_t\geq x_t}{\arg\min} Z_t(i_t^0, x_t, i_t, y_t)$$

Set $y_t^\pi(i_t^0, x_t) = y_t^*$. If $y_t^\pi(i_t^0, x_t) > x_t$ or $y_t^\pi(i_t^0, x_t) = x_t$ and $i_t^* \geq i_t^0$, set $i_t^\pi(i_t^0, x_t) = i_t^*$. Otherwise, $y_t^\pi(i_t^0, x_t) = y_t^* = x_t$ and $i_t^* < i_t^0$; let

$$\begin{aligned} \bar{i}_{t+1,1}^\pi &= i_{t+1}^\pi(i_t^*, \max(0, y_t^* - d_t)) \\ \bar{i}_{t+1,0}^\pi &= i_{t+1}^\pi(i_t^*, y_t^*) \\ \bar{y}_{t+1,1}^\pi &= y_{t+1}^\pi(i_t^*, \max(0, y_t^* - d_t)) \\ \bar{y}_{t+1,0}^\pi &= y_{t+1}^\pi(i_t^*, y_t^*) \end{aligned}$$

and, for $i \leq i_t^0$,

$$\begin{aligned} F(i) =& f(i_t^0 - i) + p_t[g(L-i)\min(y_t^*, d_t) + h\max(0, y_t^* - d_t) + b\max(0, d_t - y_t^*) \\ & + f(|\bar{i}_{t+1,1}^\pi - i|) + K(\bar{i}_{t+1,1}^\pi)\delta(\bar{y}_{t+1,1}^\pi - \max(0, y_t^* - d_t)) \\ & + G_{t+1}(\bar{i}_{t+1,1}^\pi, \bar{y}_{t+1,1}^\pi) - c(\bar{i}_{t+1,1}^\pi)\max(0, y_t^* - d_t)] \\ & + (1-p_t)[hy_t^* + f(|\bar{i}_{t+1,0}^\pi - i|) + K(\bar{i}_{t+1,0}^\pi)\delta(\bar{y}_{t+1,0}^\pi - y_t^*) + G_{t+1}(\bar{i}_{t+1,0}^\pi, \bar{y}_{t+1,0}^\pi) \\ & - c(\bar{i}_{t+1,0}^\pi)y_t^*] \end{aligned}$$

Note that $F(i) \geq Z_t(i_t^0, x_t, i, y_t^*)$ and $F(i_t^*) = Z_t(i_t^0, x_t, i_t^*, y_t^*)$ and

$$\begin{aligned} F(i_t^*) - F(i_t^0) = & p_t[f(i_t^0 - i_t^*) - f(0) + (g(L - i_t^*) - g(L - i_t^0)) \min(y_t^*, d_t) \\ & + f(|\bar{i}^\pi_{t+1,1} - i_t^*|) - f(|\bar{i}^\pi_{t+1,1} - i_t^0|)] \\ & + (1 - p_t)[f(i_t^0 - i_t^*) - f(0) + f(|\bar{i}^\pi_{t+1,0} - i_t^*|) - f(|\bar{i}^\pi_{t+1,0} - i_t^0|)] \end{aligned}$$

We will show that

$$f(i_t^0 - i_t^*) + f(|\bar{i}^\pi_{t+1,1} - i_t^*|) \geq f(0) + f(|\bar{i}^\pi_{t+1,1} - i_t^0|)$$

It will follow similarly that

$$f(i_t^0 - i_t^*) + f(|\bar{i}^\pi_{t+1,0} - i_t^*|) \geq f(0) + f(|\bar{i}^\pi_{t+1,0} - i_t^0|)$$

and together with the assumption that $g(\cdot)$ is non-decreasing, it will follow that $F(i_t^0) \leq F(i_t^*)$.

Let $f^0 = \lim_{d \to 0^+} f(d)$ and $f^+(d) = f(d) - f^0$ for $d \geq 0$. Note that $i_t^* \neq i_t^0$. If $i_t^* = \bar{i}^\pi_{t+1,1}$,

$$f(i_t^0 - i_t^*) + f(|\bar{i}^\pi_{t+1,1} - i_t^*|) = f(0) + f(|\bar{i}^\pi_{t+1,1} - i_t^0|)$$

If $i_t^0 = \bar{i}^\pi_{t+1,1}$,

$$\begin{aligned} f(i_t^0 - i_t^*) + f(|\bar{i}^\pi_{t+1,1} - i_t^*|) & \geq 2f(0) \\ & = f(0) + f(|\bar{i}^\pi_{t+1,1} - i_t^0|) \end{aligned}$$

Assume $i_t^* \neq \bar{i}^\pi_{t+1,1}$ and $i_t^0 \neq \bar{i}^\pi_{t+1,1}$. Note that $f^+(\cdot)$ is non-decreasing and, since sums of concave functions are concave, $f^+(\cdot)$ is concave over $(0, \infty)$. Then,

$$\begin{aligned} f(i_t^0 - i_t^*) + f(|\bar{i}^\pi_{t+1,1} - i_t^*|) & = 2f^0 + f^+(|i_t^* - i_t^0|) + f^+(|\bar{i}^\pi_{t+1,1} - i_t^*|) \\ & \geq 2f^0 + f^+(|i_t^* - i_t^0| + |\bar{i}^\pi_{t+1,1} - i_t^*|) && (4.6) \\ & \geq 2f^0 + f^+(|\bar{i}^\pi_{t+1,1} - i_t^0|) && (4.7) \\ & = f^0 + f(|\bar{i}^\pi_{t+1,1} - i_t^0|) \\ & \geq f(0) + f(|\bar{i}^\pi_{t+1,1} - i_t^0|) && (4.8) \end{aligned}$$

Where 4.6 follows from Lemma 4, 4.7 follows from the triangle inequality, and 4.8 follows from the assumption $f(0) \leq \lim_{d \to 0^+} f(d) = f^0$. Thus, $F(i_t^0) \leq F(i_t^*)$

$$\Rightarrow Z_t(i_t^0, x_t, i_t^0, y_t^*) \leq F(i_t^0) \leq F(i_t^*) = Z_t(i_t^0, x_t, i_t^*, y_t^*)$$

Set $i_t^\pi(i_t^0, x_t) = i_t^0$. Defined in this way, i_t^π and y_t^π are optimal policies for period t such that if $y_t^\pi(i_t^0, x_t) = x_t$ then $i_t^\pi(i_t^0, x_t) \in [i_t^0, L]$ $\forall$ $i_t^0 \in I$ and $x_t \geq 0$.

By induction, i_t^π and y_t^π are optimal policy functions for $t \in \{0, ..., T-1\}$ and $\pi^* = \{\pi_i, \pi_y\}$, where $\pi_i = \{i_0^\pi, ..., i_{T-1}^\pi\}$ and $\pi_y = \{y_0^\pi, ..., y_{T-1}^\pi\}$, is an optimal policy satisfying the conditions of the theorem. □

4.2.4 Proof of Theorem 18

Proof. Consider period $T-1$. For $i^0_{T-1} \in I$ and $x_{T-1} \geq 0$, let

$$(i^*_{T-1}, y^*_{T-1}) = \underset{i_{T-1}\in I, y_{T-1}\geq x_{T-1}}{\arg\min} Z_{T-1}(i^0_{T-1}, x_{T-1}, i_{T-1}, y_{T-1})$$

Let

$$\begin{aligned} F(i) =& f(|i - i^0_{T-1}|) + K(i)\delta(y^*_{T-1} - x_{T-1}) + c(i)(y^*_{T-1} - x_{T-1}) \\ &+ p_{T-1}[g(L-i)\min(y^*_{T-1}, d_{T-1}) + h\max(0, y^*_{T-1} - d_{T-1}) + b\max(0, d_{T-1} - y^*_{T-1})] \\ &+ (1 - p_{T-1})hy^*_{T-1} \end{aligned}$$

and note that $F(i) = Z_{T-1}(i^0_{T-1}, x_{T-1}, i, y^*_{T-1})$

1. If $i^*_{T-1} \notin \{0, i^0_{T-1}, L\}$: because a concave, monotonic transformation of a concave function is concave function is concave and a positive linear combination of concave functions is concave, $F(i)$ is concave over $[0, i^0_{T-1})$ and $(i^0_{T-1}, L]$. Furthermore, $F(i^0_{T-1}) \leq \lim_{i\to i^{0-}_{T-1}} F(i)$ and $F(i^0_{T-1}) \leq \lim_{i\to i^{0+}_{T-1}} F(i)$. Thus, $F(i^{**}_{T-1}) \leq F(i^*_{T-1})$ where $i^{**}_{T-1} = \arg\min_{i\in\{0, i^0_{T-1}, L\}} F(i)$. Then,

$$Z_{T-1}(i^0_{T-1}, x_{T-1}, i^{**}_{T-1}, y^*_{T-1}) = F(i^{**}_{T-1}) \leq F(i^*_{T-1}) = Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, y^*_{T-1})$$

   Set $i^*_{T-1} = i^{**}_{T-1}$.

2. If $i^*_{T-1} = L$, $x_{T-1} \geq d_{T-1}$, and $y^*_{T-1} > x_{T-1}$: For $y \geq x_{T-1} \geq d_{T-1}$, let

$$\begin{aligned} Q(y) =& f(|i^*_{T-1} - i^0_{T-1}|) + K(i^*_{T-1})\delta(y - x_{T-1}) + c(i^*_{T-1})(y - x_{T-1}) \\ & p_{T-1}[g(L - i^*_{T-1})d_{T-1} + h(y - d_{T-1})] + (1 - p_{T-1})hy \end{aligned}$$

   Note that $Q(y) = Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, y)$. $Q(y)$ is increasing over $[x_{T-1}, \infty)$. Thus,

$$Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, x_{T-1}) = Q(x_{T-1}) \leq Q(y^*_{T-1}) = Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, y^*_{T-1})$$

   Set $y^*_{T-1} = x_{T-1}$.

3. If $y^*_{T-1} = x_{T-1}$,

$$p_{T-1}(g(L - i^0_{T-1}) - g(0))\min(x_{T-1}, d_{T-1}) \geq f(L - i^0_{T-1}) + f(L) \tag{4.9}$$

   and $i^*_{T-1} \neq L$: $i^*_{T-1} \neq L \Rightarrow i^*_{T-1} \in \{0, i^0_{T-1}\}$. Thus, by 4.9 and the assumption that $g(\cdot)$ is non-decreasing,

$$\begin{aligned} p_{T-1}(g(L - i^*_{T-1}) - g(0))\min(x_{T-1}, d_{T-1}) &\geq p_{T-1}(g(L - i^0_{T-1}) - g(0))\min(x_{T-1}, d_{T-1}) \\ &\geq f(L - i^0_{T-1}) + f(L) \end{aligned}$$

Thus,

$$\begin{aligned}F(i^*_{T-1}) - F(L) &= f(|i^*_{T-1} - i^0_{T-1}|) - f(L - i^0_{T-1})\\ &\quad + p_{T-1}(g(L - i^*_{T-1}) - g(0))\min(x_{T-1}, d_{T-1})\\ &\geq 0\end{aligned}$$

$$\Rightarrow Z_{T-1}(i^0_{T-1}, x_{T-1}, L, y^*_{T-1}) = F(L) \leq F(i^*_{T-1}) = Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, y^*_{T-1})$$

Set $i^*_{T-1} = L$.

4. If $y^*_{T-1} > x_{T-1}$,

$$\begin{aligned}&K(i) - K(0) + (c(i) - c(0))\left(y^*_{T-1} - x_{T-1}\right)\\ &\qquad \geq p_{T-1}\left(g(L) - g(L-i)\right)\min\left(y^*_{T-1}, d_{T-1}\right) + f(i^0_{T-1}) + f(L)\end{aligned}$$

for $i \in \{i^0_{T-1}, L\}$, and $i^*_{T-1} \neq 0$: $i^*_{T-1} \neq 0 \Rightarrow i^*_{T-1} \in \{i^0_{T-1}, L\}$. Thus,

$$\begin{aligned}F(i^*_{T-1}) - F(0) &= f(|i^*_{T-1} - i^0_{T-1}|) - f(|0 - i^0_{T-1}|) + K(i^*_{T-1}) - K(0)\\ &\quad + (c(i^*_{T-1}) - c(0))(y^*_{T-1} - x_{T-1})\\ &\quad + p_{T-1}(g(L - i^*_{T-1}) - g(L))\min(y^*_{T-1}, d_{T-1})\\ &\geq K(i^*_{T-1}) - K(0) + (c(i^*_{T-1}) - c(0))(y^*_{T-1} - x_{T-1})\\ &\quad - p_{T-1}(g(L) - g(L - i^*_{T-1}))\min(y^*_{T-1}, d_{T-1}) - f(i^0_{T-1})\\ &\geq 0\end{aligned}$$

$$\Rightarrow Z_{T-1}(i^0_{T-1}, x_{T-1}, 0, y^*_{T-1}) = F(0) \leq F(i^*_{T-1}) = Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, y^*_{T-1})$$

Set $i^*_{T-1} = 0$.

Set $i^\pi_{T-1}(i^0_{T-1}, x_{T-1}) = i^*_{T-1}$ and $y^\pi_{T-1}(i^0_{T-1}, x_{T-1}) = y^*_{T-1}$. Defined in this way, i^π_{T-1} and y^π_{T-1} are optimal policies for period $T-1$ satisfying the conditions of the Theorem.

Consider period t. Assume i^π_τ and y^π_τ are optimal policies for periods $\tau \in \{t+1, ..., T-1\}$ satisfying the conditions of the Theorem. For $i^0_t \in I$ and $x_t \geq 0$, let

$$(i^*_t, y^*_t) = \underset{i_t \in I, y_t \geq x_t}{\arg\min}\, Z_t(i^0_t, x_t, i_t, y_t)$$

1. If $i^*_t \notin \{0, i^0_t, L\}$: Let

$$\begin{aligned}\hat{D}^* &= \{(\mathbf{d}, \tau) \mid \mathbf{d} \in \mathbf{D}^t, t \in \{t+1, ..., T-1\}, i^\pi_{\iota\mathbf{d}} = i^*_t\ \forall \iota \in \{t+1, ..., \tau\}\},\\ \hat{D}^{*+} &= \{(\mathbf{d}, \tau) \mid \mathbf{d} \in \mathbf{D}^t, t \in \{t+1, ..., T-1\}, i^\pi_{\iota\mathbf{d}} = i^*_t\ \forall \iota \in \{t+1, ..., \tau-1\}, i^\pi_{\tau\mathbf{d}} \neq i^*_t\}\\ \hat{D}^*_C &= \{(\mathbf{d}, \tau) \mid \mathbf{d} \in \mathbf{D}^t, t \in \{t+1, ..., T-1\}\} \setminus (\hat{D}^* \cup \hat{D}^{*+})\end{aligned}$$

Furthermore, for $i \in I$, let

$$\begin{aligned}
F(i) =& f(|i - i_t^0|) + K(i)\delta(y_t^* - x_t) + c(i)(y_t^* - x_t) \\
&+ p_t[g(L-i)\min(y_t^*, d_t) + h\max(0, y_t^* - d_t) + b\max(0, d_t - y_t^*)] + (1-p_t)hy_t^* \\
&+ \sum_{(\mathbf{d},\tau)\in\hat{D}^*} p(\mathbf{d})[K(i)\delta(y_{\tau\mathbf{d}}^\pi - x_{\tau\mathbf{d}}^\pi) + c(i)(y_{\tau\mathbf{d}}^\pi - x_{\tau\mathbf{d}}^\pi) + g(L-i)\min(y_{\tau\mathbf{d}}^\pi, d_{j_\tau\tau}) \\
&\qquad\qquad + h\max(0, y_{\tau\mathbf{d}}^\pi - d_{j_\tau\tau}) + b\max(0, d_{j_\tau\tau} - y_{\tau\mathbf{d}}^\pi)] \\
&+ \sum_{(\mathbf{d},\tau)\in\hat{D}^{*+}} p(\mathbf{d}) f(|i_{\tau\mathbf{d}}^\pi - i|) \ + \sum_{(\mathbf{d},\tau)\in\hat{D}_C^*} p(\mathbf{d}) f(|i_{\tau\mathbf{d}}^\pi - i_{\tau-1\mathbf{d}}^\pi|) \\
&+ \sum_{(\mathbf{d},\tau)\in\hat{D}^{*+}\cup\hat{D}_C^*} p(\mathbf{d})[K(i_{\tau\mathbf{d}}^\pi)\delta(y_{\tau\mathbf{d}}^\pi - x_{\tau\mathbf{d}}^\pi) + c(i_{\tau\mathbf{d}}^\pi)(y_{\tau\mathbf{d}}^\pi - x_{\tau\mathbf{d}}^\pi) \\
&\qquad\qquad + g(L - i_{\tau\mathbf{d}}^\pi)\min(y_{\tau\mathbf{d}}^\pi, d_{j_\tau\tau}) + h\max(0, y_{\tau\mathbf{d}}^\pi - d_{j_\tau\tau}) \\
&\qquad\qquad + b\max(0, d_{j_\tau\tau} - y_{\tau\mathbf{d}}^\pi)]
\end{aligned}$$

Note that $F(i) \geq Z_t(i_t^0, x_t, i, y_t^*)$ and $F(i_t^*) = Z_t(i_t^0, x_t, i_t^*, y_t^*)$. Because a concave, monotonic transformation of a concave function is concave and a positive linear combination of concave functions is concave, $F(i)$ is concave over $(0, i_t^0)$ and (i_t^0, L). Furthermore, $F(0) \leq \lim_{i\to 0^+} F(i)$, $F(i_t^0) \leq \lim_{i\to i_t^{0-}} F(i)$, $F(i_t^0) \leq \lim_{i\to i_t^{0+}} F(i)$, and $F(L) \leq \lim_{i\to L^-} F(i)$. Thus, $F(i_t^{**}) \leq F(i_t^*)$ where $i_t^{**} = \arg\min_{i\in\{0, i_t^0, L\}} F(i)$. Then,

$$Z_t(i_t^0, x_t, i_t^{**}, y_t^*) \leq F(i_t^{**}) \leq F(i_t^*) = Z_t(i_t^0, x_t, i_t^*, y_t^*)$$

Set $i_t^* = i_t^{**}$.

2. If $i_t^* = L$, $x_t \geq d_t$, and $y_t^* > x_t$: Let

$$\begin{aligned}
\bar{i}_{t+1,1}^\pi &= i_{t+1}^\pi(i_t^*, y_t^* - d_t) \\
\bar{i}_{t+1,0}^\pi &= i_{t+1}^\pi(i_t^*, y_t^*) \\
\bar{y}_{t+1,1}^\pi &= y_{t+1}^\pi(i_t^*, y_t^* - d_t) \\
\bar{y}_{t+1,0}^\pi &= y_{t+1}^\pi(i_t^*, y_t^*)
\end{aligned}$$

and, for $y \geq x_t \geq d_t$,

$$\begin{aligned}
Q(y) =& f(|i_t^* - i_t^0|) + K(i_t^*)\delta(y - x_t) + c(i_t^*)(y - x_t) \\
&+ p_t[g(L - i_t^*)d_t + h(y - d_t) + f(|\bar{i}_{t+1,1}^\pi - i_t^*|) + K(\bar{i}_{t+1,1}^\pi)\delta(\bar{y}_{t+1,1}^\pi - (y - d_t)) \\
&\qquad + G_{t+1}(\bar{i}_{t+1,1}^\pi, \bar{y}_{t+1,1}^\pi) - c(\bar{i}_{t+1,1}^\pi)(y - d_t)] \\
&+ (1 - p_t)[hy + f(|\bar{i}_{t+1,0}^\pi - i_t^*|) + K(\bar{i}_{t+1,0}^\pi)\delta(\bar{y}_{t+1,0}^\pi - y) + G_{t+1}(\bar{i}_{t+1,0}^\pi, \bar{y}_{t+1,0}^\pi) \\
&\qquad - c(\bar{i}_{t+1,0}^\pi)y]
\end{aligned}$$

Note that $Q(y) \geq Z_t(i_t^0, x_t, i_t^*, y)$ and $Q(y_t^*) = Z_t(i_t^0, x_t, i_t^*, y_t^*)$. By the assumption that $K(\cdot)$ and $c(\cdot)$ are non-decreasing, $c(i_t^*) \geq c(i)\ \forall i \in I$; thus,

$$
\begin{aligned}
Q(y_t^*) - Q(x_t) = {} & K(i_t^*) + c(i_t^*)(y_t^* - x_t) \\
& + p_t[h(y_t^* - x_t) + K(\bar{i}_{t+1,1}^\pi)[\delta(\bar{y}_{t+1,1}^\pi - (y_t^* - d_t)) - \delta(\bar{y}_{t+1,1}^\pi - (x_t - d_t))] \\
& \qquad - c(\bar{i}_{t+1,1}^\pi)(y_t^* - x_t)] \\
& + (1 - p_t)[h(y_t^* - x_t) + K(\bar{i}_{t+1,0}^\pi)[\delta(\bar{y}_{t+1,0}^\pi - y^*) - \delta(\bar{y}_{t+1,0}^\pi - x_t)] \\
& \qquad \quad - c(\bar{i}_{t+1,0}^\pi)(y_t^* - x_t)] \\
\geq {} & K(i_t^*) + c(i_t^*)(y_t^* - x_t) + p_t[-K(\bar{i}_{t+1,1}^\pi) - c(\bar{i}_{t+1,1}^\pi)(y_t^* - x_t)] \\
& + (1 - p_t)[-K(\bar{i}_{t+1,0}^\pi) - c(\bar{i}_{t+1,0}^\pi)(y_t^* - x_t)] \\
= {} & K(i_t^*) - p_t K(\bar{i}_{t+1,1}^\pi) - (1 - p_t) K(\bar{i}_{t+1,0}^\pi) \\
& + [c(i_t^*) - p_t c(\bar{i}_{t+1,1}^\pi) - (1 - p_t) c(\bar{i}_{t+1,0}^\pi)](y_t^* - x_t) \\
\geq {} & 0
\end{aligned}
$$

Thus,

$$Z_t(i_t^0, x_t, i_t^*, x_t) \leq Q(x_t) \leq Q(y_t^*) = Z_t(i_t^0, x_t, i_t^*, y_t^*)$$

Set $y_t^* = x_t$.

Let

$$
\begin{aligned}
\bar{i}_{t+1,1}^\pi &= i_{t+1}^\pi(i_t^*, \max(0, y_t^* - d_t)) \\
\bar{i}_{t+1,0}^\pi &= i_{t+1}^\pi(i_t^*, y_t^*) \\
\bar{y}_{t+1,1}^\pi &= y_{t+1}^\pi(i_t^*, \max(0, y_t^* - d_t)) \\
\bar{y}_{t+1,0}^\pi &= y_{t+1}^\pi(i_t^*, y_t^*)
\end{aligned}
$$

and, for $i \in I$,

$$
\begin{aligned}
F(i) = {} & f(|i - i_t^0|) + K(i)\delta(y_t^* - x_t) + c(i)(y_t^* - x_t) \\
& + p_t[g(L - i)\min(y_t^*, d_t) + h\max(0, y_t^* - d_t) + b\max(0, d_t - y_t^*) + f(|\bar{i}_{t+1,1}^\pi - i|) \\
& \qquad + K(\bar{i}_{t+1,1}^\pi)\delta(\bar{y}_{t+1,1}^\pi - \max(0, y_t^* - d_t)) + G_{t+1}(\bar{i}_{t+1,1}^\pi, \bar{y}_{t+1,1}^\pi) \\
& \qquad - c(\bar{i}_{t+1,1}^\pi)\max(0, y_t^* - d_t)] \\
& + (1 - p_t)[hy_t^* + f(|\bar{i}_{t+1,0}^\pi - i|) + K(\bar{i}_{t+1,0}^\pi)\delta(\bar{y}_{t+1,0}^\pi - y_t^*) + G_{t+1}(\bar{i}_{t+1,0}^\pi, \bar{y}_{t+1,0}^\pi) \\
& \qquad - c(\bar{i}_{t+1,0}^\pi)y_t^*]
\end{aligned}
$$

Note that $F(i) \geq Z_t(i_t^0, x_t, i, y_t^*)$.

3. If $y_t^* = x_t$,

$$p_t(g(L - i_t^0) - g(0))\min(x_t, d_t) \geq f(L - i_t^0) + f(L) \tag{4.10}$$

and $i_t^* \neq L$: $i_t^* \neq L \Rightarrow i_t^* \in \{0, i_t^0\}$. Thus, by 4.10 and the assumption that $g(\cdot)$ is non-decreasing,

$$\begin{aligned} p_t(g(L - i_t^*) - g(0)) \min(x_t, d_t) &\geq p_t(g(L - i_t^0) - g(0)) \min(x_t, d_t) \\ &\geq f(L - i_t^0) + f(L) \end{aligned}$$

Thus,

$$\begin{aligned} F(i_t^*) - F(L) = &f(|i_t^* - i_t^0|) - f(|L - i_t^0|) \\ &+ p_t[(g(L - i_t^*) - g(0)) \min(x_t, d_t) + f(|\bar{i}_{t+1,1}^{\pi} - i_t^*|) - f(|\bar{i}_{t+1,1}^{\pi} - L|)] \\ &+ (1 - p_t)[f(|\bar{i}_{t+1,0}^{\pi} - i_t^*|) - f(|\bar{i}_{t+1,0}^{\pi} - L|)] \\ \geq &\, p_t(g(L - i_t^*) - g(0)) \min(x_t, d_t) + f(L - i_t^0) - f(L) \\ \geq &\, 0 \end{aligned}$$

$$\Rightarrow Z_t(i_t^0, x_t, L, y_t^*) \leq F(L) \leq F(i_t^*) = Z_t(i_t^0, x_t, i_t^*, y_t^*)$$

Set $i_t^* = L$.

4. If $y_t^* > x_t$,

$$\begin{aligned} K(i) - K(0) + (c(i) - c(0))\,(y_t^* - x_t)& \\ \geq p_t\,(g(L) - g(L - i)) \min\,(y_t^*, d_t) + f(i_t^0) + f(L)& \end{aligned}$$

for $i \in \{i_t^0, L\}$, and $i_t^* \neq 0$: $i_t^* \neq 0 \Rightarrow i_t^* \in \{i_t^0, L\}$. Thus,

$$\begin{aligned} F(i_t^*) - F(0) = &f(|i_t^* - i_t^0|) - f(|0 - i_t^0|) + K(i_t^*) - K(0) + (c(i_t^*) - c(0))(y_t^* - x_t) \\ &+ p_t[(g(L - i_t^*) - g(L)) \min(y_t^*, d_t) + f(|\bar{i}_{t+1,1}^{\pi} - i_t^*|) - f(|\bar{i}_{t+1,1}^{\pi} - 0|)] \\ &+ (1 - p_t)[f(|\bar{i}_{t+1,0}^{\pi} - i_t^*|) - f(|\bar{i}_{t+1,0}^{\pi} - 0|)] \\ \geq &\, K(i_t^*) - K(0) + (c(i_t^*) - c(0))(y_t^* - x_t) \\ &- p_t(g(L) - g(L - i_t^*)) \min(y_t^*, d_t) - f(i_t^0) - f(L) \\ \geq &\, 0 \end{aligned}$$

$$\Rightarrow Z_t(i_t^0, x_t, 0, y_t^*) \leq F(0) \leq F(i_t^*) = Z_t(i_t^0, x_t, i_t^*, y_t^*)$$

Set $i_t^* = 0$.

Set $i_t^{\pi}(i_t^0, x_t) = i_t^*$ and $y_t^{\pi}(i_t^0, x_t) = y_t^*$. Defined in this way, i_t^{π} and y_t^{π} are optimal policies for period t satisfying the conditions of the Theorem.

By induction, i_t^{π} and y_t^{π} are optimal policy functions for $t \in \{0, ..., T-1\}$ and $\pi^* = \{\pi_i, \pi_y\}$, where $\pi_i = \{i_0^{\pi}, ..., i_{T-1}^{\pi}\}$ and $\pi_y = \{y_0^{\pi}, ..., y_{T-1}^{\pi}\}$, is an optimal policy satisfying the conditions of the theorem. □

4.2.5 Proof of Theorem 19

Proof. Let $i_{-1} = i_0^0$ and $\Pi_i = \{\{i_0, ..., i_{T-1}\} | i_t \in I \ \forall t\}$. The problem can be written $\min_{\pi_i \in \Pi_i} \bar{V}_0^{\pi_i}(x_0)$ where

$$\bar{V}_t^{\pi_i}(x_t) = \min_{y_t \geq x_t} \left\{ f_{i_{t-1}i_t} + K\delta(y_t - x_t) + \bar{G}_t(y_t) - c_{i_t}x_t \right\}$$

$\bar{V}_T^{\pi_i}(x_T) = 0$ and

$$\begin{aligned} \bar{G}_t(y) =& c_{i_t}y + \sum_{j \in J} p_{jt} \Big[h\max(0, y - d_{jt}) + g_{i_tj}d_{jt} + (b_j - g_{i_tj})\max(0, d_{jt} - y) \\ & + \bar{V}_{t+1}^{\pi_i}(\max(0, y - d_{jt})) \Big] \end{aligned}$$

Consider a specific movement policy $\pi_i = \{i_0, ..., i_{T-1}\} \in \Pi_i$ and let us optimize over the inventory decisions. We will show that under π_i, a multiperiod (s, S) policy is optimal. Thus, taking the minimization over all $\pi_i \in \Pi_i$, a multiperiod (s, S) policy is optimal for the problem and the result will be proved.

Under π_i, we can disregard the fixed cost terms $f_{i_{t-1}i_t}$ and $g_{i_tj}d_{jt}$ in optimizing over the inventory decisions:

$$\bar{V}_t^{\pi_i}(x_t) = \min_{y_t \geq x_t} \left\{ K\delta(y_t - x_t) + \bar{G}_t(y_t) - c_{i_t}x_t \right\} \tag{4.11}$$

$$\bar{G}_t(y) = c_{i_t}y + \sum_{j \in J} p_{jt} \Big[h\max(0, y - d_{jt}) + (b_j - g_{i_tj})\max(0, d_{jt} - y) + \bar{V}_{t+1}^{\pi_i}(\max(0, y - d_{jt})) \Big]$$

Suppose for $t \in \{0, ..., T-1\}$ $\bar{G}_t(y)$ is K-convex and has unconstrained minimum $S_t = \min_y \bar{G}_t(y)$. Let s_t be the largest y such that $\bar{G}_t(y) = K + \bar{G}_t(S_t)$. In the case where $K = 0$, $s_t = S_t \ \forall \ t \in \{0, ..., T-1\}$. By well known properties of K-convex functions, the optimal y_t in 4.11 is (for details, see Snyder and Shen (2011) Lemma 4.4 (b)):

$$y_t^* = \begin{cases} S_t & \text{if } x_t \leq s_t \\ x_t & \text{if } x_t > s_t \end{cases}$$

which is equivalent to an (s, S) policy. Thus, the result will be proved if we can show the functions $\bar{G}_t(y)$ are K-convex and continuous and $\lim_{|y| \to \infty} \bar{G}_t(y) = \infty$, so that the minimizing scalar S_t exists, $\forall t \in \{0, ..., T-1\}$. Similarly to Shreve (2005), let

$$H_{tj}(y) = (b_j - g_{i_tj})\max(0, -y) + \bar{V}_{t+1}^{\pi}(\max(0, y))$$

Note that $\max(0, y-d)$ and $\max(0, d-y)$ are convex $\forall d \in \mathbb{R}$ and that $b_j - g_{ij} > c_{i'} \ \forall i, i' \in I$ and $j \in J$ by assumption $\Rightarrow b_j - g_{i_tj} > 0 \ \forall j \in J$ since $c_i \geq 0 \ \forall i \in I$. If $H_{tj}(y)$ is K-convex and continuous $\forall j \in J$, then $\bar{G}_t(y)$ is K-convex and continuous since non-negative multiples of convex functions are convex, $\alpha_1 f_1(x) + \alpha_2 f_2(x)$ is $(\alpha_1 K_1 + \alpha_2 K_2)$-convex for $f_1(x)$ and $f_2(x)$ K_1- and K_2-convex functions, respectively, and $\alpha_1, \alpha_2 \geq 0$, and sums of continuous functions are continuous. Furthermore,

$$\lim_{|y|\to\infty} c_{i_t} y + \sum_{j\in J} p_{jt}[h \max(0, y - d_{jt}) + (b_j - g_{i_t j}) \max(0, d_{jt} - y)] = \infty \tag{4.12}$$

since its derivative tends to $c_{i_t} + h$ as $y \to \infty$ which is positive by assumption and $c_{i_t} - \sum_{j\in J} p_{jt}(b_j - g_{i_t j})$ as $y \to -\infty$ which is negative by the assumption $b_j - g_{ij} > c_{i'}$ $\forall i, i' \in I$ and $j \in J$. Thus, if

$$\lim_{|y|\to\infty} \bar{V}_{t+1}^{\pi_i}(\max(0, y - d)) \geq 0 \quad \forall d \in \mathbb{R}$$

then $\lim_{|y|\to\infty} \bar{G}_t(y) = \infty$. Accordingly, the result will be proved if we can show that $H_{tj}(y)$ is K-convex and continuous and

$$\lim_{|y|\to\infty} \bar{V}_{t+1}^{\pi_i}(\max(0, y - d)) \geq 0 \quad \forall t \in \{0, ..., T-1\},\ j \in J,\ d \in \mathbb{R}$$

$\bar{V}_T^{\pi_i}(\max(0, y - d)) = 0$ $\forall d \in \mathbb{R}$ and is thus K-convex and

$$\lim_{|y|\to\infty} \bar{V}_T^{\pi_i}(\max(0, y - d)) = 0 \geq 0 \quad \forall d \in \mathbb{R}$$

Additionally, $H_{T-1j}(y)$ is convex and thus is K-convex and continuous $\forall j \in J$. Assume

$$\lim_{|y|\to\infty} \bar{V}_{t+2}^{\pi_i}(\max(0, y - d)) \geq 0 \quad \forall d \in \mathbb{R}$$

and $H_{t+1j}(y)$ is K-convex and continuous $\forall j \in J$. Then, an (s, S) policy is optimal in period $t+1$ and $\bar{V}_{t+1}^{\pi_i}(y) = Q_{t+1}(y) - c_{i_{t+1}} y$ where

$$Q_{t+1}(y) = \begin{cases} K + \bar{G}_{t+1}(S_{t+1}) & \text{if } y < s_{t+1} \\ \bar{G}_{t+1}(y) & \text{if } y \geq s_{t+1} \end{cases}$$

From this equation together with 4.12 and the assumption

$$\lim_{|y|\to\infty} \bar{V}_{t+2}^{\pi_i}(\max(0, y - d)) \geq 0 \quad \forall d \in \mathbb{R}$$

we see that

$$\lim_{|y|\to\infty} \bar{V}_{t+1}^{\pi_i}(\max(0, y - d)) \geq 0 \quad \forall d \in \mathbb{R}$$

Each piece of $\bar{V}_{t+1}^{\pi_i}(y)$ is continuous and at the break point $y = s_{t+1}$, $\bar{G}_{t+1}(y) = K + \bar{G}_{t+1}(S_{t+1})$ by definition of s_{t+1}. Thus, $\bar{V}_{t+1}^{\pi_i}(y)$ is continuous and, since sums of continuous functions are continuous, $H_{tj}(y)$ is continuous $\forall j \in J$. Since $H_{t+1j}(y)$ is K-convex $\forall j \in J$, $\bar{G}_{t+1}(y)$ is K-convex and thus $\bar{V}_{t+1}^{\pi_i}(y)$ is K-convex (see Snyder and Shen (2011) Lemma 4.4 (c) for details). All that is left to prove is that $H_{tj}(y)$ is K-convex $\forall j \in J$ and the result will follow from induction. By Shreve (2005) together with the assumption $b_j - g_{ij} > c_{i'}$ $\forall i, i' \in I$ and $j \in J$ and K-convexity of $\bar{V}_{t+1}^{\pi_i}(y)$, $H_{tj}(y)$ is K-convex $\forall j \in J$. $\square$

4.2.6 Proof of Theorem 20

Proof. Recall that J contains a dummy site J with $d_{jt} = 0\ \forall t$. Consider period $T-1$, the last decision making period. For $i^0_{T-1} \in I$ and $x_{T-1} \geq 0$, let

$$(i^*_{T-1}, y^*_{T-1}) = \underset{i_{T-1}\in I, y_{T-1}\geq x_{T-1}}{\arg\min} Z_{T-1}(i^0_{T-1}, x_{T-1}, i_{T-1}, y_{T-1})$$

Set $i^\pi_{T-1}(i^0_{T-1}, x_{T-1}) = i^*_{T-1}$. If $y^*_{T-1} = x_{T-1}$ or $y^*_{T-1} = d_{jT-1}$ for some $j \in J$, set $y^\pi_{T-1}(i^0_{T-1}, x_{T-1}) = y^*_{T-1}$. Otherwise, let $J^*_{T-1} = \{j | j \in J, d_{jT-1} < y^*_{T-1}\}$,

$$l = \begin{cases} \max\left(\max_{j\in J^*_{T-1}} d_{jT-1}, x_{T-1}\right), & \text{if } J^*_{T-1} \neq \emptyset \\ x_{T-1}, & \text{otherwise} \end{cases}$$

and

$$u = \begin{cases} \min_{j\in J\setminus J^*_{T-1}} d_{jT-1}, & \text{if } J\setminus J^*_{T-1} \neq \emptyset \\ \infty, & \text{otherwise} \end{cases}$$

Keeping i^0_{T-1}, x_{T-1}, and i^*_{T-1} constant, we can rewrite $Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, y)$ over $y \in [l, u]$ if $u \neq \infty$ and $y \in [l, u)$ if $u = \infty$ as

$$\begin{aligned} Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, y) =& f_{i^0_{T-1} i^*_{T-1}} + K_{i^*_{T-1}}\delta(y - x_{T-1}) + c_{i^*_{T-1}}(y - x_{T-1}) \\ &+ \sum_{j\in J^*_{T-1}} p_{jT-1}\left[g_{i^*_{T-1}j} d_{jT-1} + h(y - d_{jT-1})\right] \\ &+ \sum_{j\in J\setminus J^*_{T-1}} p_{jT-1}\left[g_{i^*_{T-1}j}\, y + b_j(d_{jT-1} - y)\right] \end{aligned}$$

Suppose $u \neq \infty$. If $l = x_{T-1}$, then $Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, y)$ is linear in y over $(l, u]$, holding all else constant, and

$$\begin{aligned} Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, l) &= \lim_{y\to l^+} Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, y) - K_{i^*_{T-1}} \\ \Rightarrow Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, l) &\leq \lim_{y\to l^+} Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, y) \end{aligned}$$

If $l \neq x_{T-1}$, then $Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, y)$ is linear in y over $[l, u]$, holding all else constant. Thus, in either case,

$$Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, l) \leq Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, y^*_{T-1})$$

or

$$Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, u) \leq Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, y^*_{T-1})$$

Set

$$y^{\pi}_{T-1}(i^0_{T-1}, x_{T-1}) = \underset{y \in \{l,u\}}{\arg\min}\, Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, y)$$

Suppose $u = \infty$. Then

$$Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, l) \leq Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, y^*_{T-1})$$

since $c_{i^*_{T-1}}$ and h are non-negative; set $y^{\pi}_{T-1}(i^0_{T-1}, x_{T-1}) = l$. Defined in this way, i^{π}_{T-1} and y^{π}_{T-1} are optimal policies for period $T-1$ and y^{π}_{T-1} is such that $y^{\pi}_{T-1}(i^0_{T-1}, x_{T-1}) = x_{T-1}$ or $y^{\pi}_{T-1}(i^0_{T-1}, x_{T-1}) = d_{jT-1}$ for some $j \in J$ $\forall i^0_{T-1} \in I$ and $x_{T-1} \geq 0$.

Assume i^{π}_{τ} and y^{π}_{τ} are optimal policies for periods $\tau \in \{t+1, ..., T-1\}$ and y^{π}_{τ} is such that $y^{\pi}_{\tau}(i^0_{\tau}, x_{\tau}) = x_{\tau}$ or $y^{\pi}_{\tau}(i^0_{\tau}, x_{\tau}) = \sum_{\iota=\tau}^{T-1} d_{j_\iota \iota}$, for some $j_\iota \in J$ $\forall \iota \in \{\tau, ..., T-1\}$, $\forall i^0_{\tau} \in I$, $x_{\tau} \geq 0$, and $\tau \in \{t+1, ..., T-1\}$. Consider period t. We prove this induction step similarly to the period $T-1$ base case, however, this step is more complicated as changing an ordering decision in period t may affect ordering decisions in future periods. For $i^0_t \in I$ and $x_t \geq 0$, let

$$(i^*_t, y^*_t) = \underset{i_t \in I, y_t \geq x_t}{\arg\min}\, Z_t(i^0_t, x_t, i_t, y_t)$$

Set $i^{\pi}_t(i^0_t, x_t) = i^*_t$. If $y^*_t = x_t$ or $y^*_t = \sum_{\tau=t}^{T-1} d_{j_\tau \tau}$, for some $j_\tau \in J$ $\forall \tau \in \{t, ..., T-1\}$, set $y^{\pi}_t(i^0_t, x_t) = y^*_t$. Otherwise, let $J^*_t = \{j \mid j \in J, d_{jt} < y^*_t\}$ and

$$\begin{aligned}
\bar{i}^{\pi}_{t+1j} &= i^{\pi}_{t+1}(i^*_t, \max(0, y^*_t - d_{jt})) \quad \forall j \in J \\
\bar{y}^{\pi}_{t+1j} &= y^{\pi}_{t+1}(i^*_t, \max(0, y^*_t - d_{jt})) \quad \forall j \in J
\end{aligned}$$

Note that

$$\begin{aligned}
Z_t(i^0_t, x_t, i^*_t, y^*_t) =& f_{i^0_t i^*_t} + K_{i^*_t}\delta(y^*_t - x_t) + c_{i^*_t}(y^*_t - x_t) \\
&+ \sum_{j \in J^*_t} p_{jt}\Big[g_{i^*_t j} d_{jt} + h(y^*_t - d_{jt}) + f_{i^*_t \bar{i}^{\pi}_{t+1j}} + K_{\bar{i}^{\pi}_{t+1j}}\delta(\bar{y}^{\pi}_{t+1j} - (y^*_t - d_{jt})) \\
&\qquad + G_{t+1}(\bar{i}^{\pi}_{t+1j}, \bar{y}^{\pi}_{t+1j}) - c_{\bar{i}^{\pi}_{t+1j}}(y^*_t - d_{jt})\Big] \\
&+ \sum_{j \in J \setminus J^*_t} p_{jt}\Big[g_{i^*_t j} y^*_t + b_j(d_{jt} - y^*_t) + f_{i^*_t \bar{i}^{\pi}_{t+1j}} + K_{\bar{i}^{\pi}_{t+1j}}\delta(\bar{y}^{\pi}_{t+1j}) \\
&\qquad + G_{t+1}(\bar{i}^{\pi}_{t+1j}, \bar{y}^{\pi}_{t+1j})\Big]
\end{aligned}$$

Lets consider changing the ordering decision y^*_t while keeping i^*_t constant and the decisions in future periods as constant as possible. Let

$$\hat{J}^*_t = \{j \mid j \in J^*_t, \bar{y}^{\pi}_{t+1j} = y^*_t - d_{jt}\}$$

Then $\{d_{jt} \mid j \in \hat{J}_t^*\}$ is the collection of period t demand realizations for which the associated period $t+1$ ordering decision is to set the order up to level to the initial inventory level or, in other words, to not order anything. Thus, keeping these ordering decisions unchanged, if we decrease y_t^* slightly we introduce a new fixed ordering cost and if we increase y_t^* slightly these ordering decisions become infeasible. In changing y_t^*, lets also change the ordering decisions $\{\bar{y}_{t+1j}^{\pi} \mid j \in \hat{J}_t^*\}$, and the future period ordering decisions which are similarly affected, in the corresponding way. We introduce some new notation. Let

$$\boldsymbol{D}^{t*} = \left\{\boldsymbol{d} = (d_{j_t t}, ..., d_{j_{T-1} T-1}) \mid \boldsymbol{d} \in \boldsymbol{D}^t, j_t \in \hat{J}_t^*\right\}$$

be the set of all possible strings of future demand realizations starting with $d_{j_t t}$ for $j_t \in \hat{J}_t^*$. Let $\Sigma_{\boldsymbol{d}\tau} = \sum_{\iota=t}^{\tau} d_{j_\iota \iota}$ and

$$\begin{aligned}
\hat{D}^* &= \left\{(\boldsymbol{d}, \tau) \mid \boldsymbol{d} \in \boldsymbol{D}^t, \tau \in \{t+1, ..., T-1\}, y_{\tau\boldsymbol{d}}^{\pi} = y_t^* - \Sigma_{\boldsymbol{d}\tau-1}\right\} \\
\hat{D}_c^* &= \left\{(\boldsymbol{d}, \tau) \mid \boldsymbol{d} \in \boldsymbol{D}^t, \tau \in \{t+1, ..., T-1\}, y_{\tau\boldsymbol{d}}^{\pi} \neq y_t^* - \Sigma_{\boldsymbol{d}\tau-1}\right\} \\
D_\Sigma &= \left\{\sum_{\tau=t}^{T-1} d_{j_\tau \tau} \mid j_\tau \in J \ \forall \tau \in \{t, ..., T-1\}\right\}
\end{aligned}$$

Let

$$\begin{aligned}
l &= \max\left(\{z \mid z \in D_\Sigma, z < y_t^*\}, x_t\right) \\
u &= \min\left(\{z \mid z \in D_\Sigma, z > y_t^*\}, \infty\right)
\end{aligned}$$

and, for $y \in [l, u]$ ($[l, u)$ if $u = \infty$),

$$\begin{aligned}
F(y) = &f_{i_t^0 i_t^*} + K_{i_t^*}\delta(y - x_t) + c_{i_t^*}(y - x_t) \\
&+ \sum_{j \in J_t^* \setminus \hat{J}_t^*} p_{jt}\Big[g_{i_t^* j} d_{jt} + h(y - d_{jt}) + f_{i_t^* \bar{i}_{t+1j}^{\pi}} + K_{\bar{i}_{t+1j}^{\pi}}\delta(\bar{y}_{t+1j}^{\pi} - (y - d_{jt})) \\
&\qquad\qquad + G_{t+1}(\bar{i}_{t+1j}^{\pi}, \bar{y}_{t+1j}^{\pi}) - c_{\bar{i}_{t+1j}^{\pi}}(y - d_{jt})\Big] \\
&+ \sum_{j \in J \setminus J_t^*} p_{jt}\Big[g_{i_t^* j} y + b_j(d_{jt} - y) + f_{i_t^* \bar{i}_{t+1j}^{\pi}} + K_{\bar{i}_{t+1j}^{\pi}}\delta(\bar{y}_{t+1j}^{\pi}) \\
&\qquad\qquad + G_{t+1}(\bar{i}_{t+1j}^{\pi}, \bar{y}_{t+1j}^{\pi})\Big] \\
&+ \sum_{j \in \hat{J}_t^*} p_{jt}[g_{i_t^* j} d_{jt} + h(y - d_{jt})] \\
&+ \sum_{(\boldsymbol{d}, \tau) \in \hat{D}^*} p(\boldsymbol{d})\Big[f_{i_{\tau-1\boldsymbol{d}}^{\pi} i_{\tau\boldsymbol{d}}^{\pi}} + g_{i_{\tau\boldsymbol{d}}^{\pi} j_\tau}\min(y - \Sigma_{\boldsymbol{d}\tau-1}, d_{j_\tau \tau}) + h\max(0, y - \Sigma_{\boldsymbol{d}\tau-1} - d_{j_\tau \tau}) \\
&\qquad\qquad + b_{j_\tau}\max(0, d_{j_\tau \tau} - (y - \Sigma_{\boldsymbol{d}\tau-1}))\Big] \\
&+ \sum_{(\boldsymbol{d}, \tau) \in \hat{D}_c^*} p(\boldsymbol{d})\Big[f_{i_{\tau-1\boldsymbol{d}}^{\pi} i_{\tau\boldsymbol{d}}^{\pi}} + K_{i_{\tau\boldsymbol{d}}^{\pi}}\delta(y_{\tau\boldsymbol{d}}^{\pi} - x_{\tau\boldsymbol{d}}^{\pi}) + c_{i_{\tau\boldsymbol{d}}^{\pi}}(y_{\tau\boldsymbol{d}}^{\pi} - x_{\tau\boldsymbol{d}}^{\pi}) + g_{i_{\tau\boldsymbol{d}}^{\pi} j_\tau}\min(y_{\tau\boldsymbol{d}}^{\pi}, d_{j_\tau \tau})
\end{aligned}$$

$$+ h\max(0, y^{\pi}_{\tau\boldsymbol{d}} - d_{j_\tau\tau}) + b_{j_\tau}\max(0, d_{j_\tau\tau} - y^{\pi}_{\tau\boldsymbol{d}})\Big]$$

Note that $F(y_t^*) = Z_t(i_t^0, x_t, i_t^*, y_t^*)$. $F(y)$ is linear over (l,u) and $f(l) \le \lim_{y\to l^+} F(y)$. Suppose $u \neq \infty$. Then $F(u) \le \lim_{y\to u^-} F(y)$. Thus, $\min_{y\in\{l,u\}} F(y) \le F(y_t^*)$. Set $y_t^{**} = \arg\min_{y\in\{l,u\}} F(y)$. Alternatively, suppose $u = \infty$. Then $F(l) \le F(y_t^*)$ since $c_{i_t^*}$ and h are non-negative. Set $y_t^{**} = l$. Then, in both cases,

$$Z_t(i_t^0, x_t, i_t, y_t^*) = F(y_t^*) \ge F(y_t^{**}) \ge Z_t(i_t^0, x_t, i_t, y_t^{**})$$

Set $y_t^\pi(i_t^0, x_t) = y_t^{**}$. Defined in this way, i_t^π and y_t^π are optimal policies for period t such that $y_t^\pi(i_t^0, x_t) = x_t$ or $y_t^\pi(i_t^0, x_t) = \sum_{\tau=t}^{T-1} d_{j_\tau\tau}$, for some $j_\tau \in J$ $\forall\tau \in \{t, ..., T-1\}$, $\forall i_t^0 \in I$ and $x_t \ge 0$.

By induction, i_t^π and y_t^π are optimal policy functions for $t \in \{0, ..., T-1\}$ and $\pi^* = \{\pi_i, \pi_y\}$, where $\pi_i = \{i_0^\pi, ..., i_{T-1}^\pi\}$ and $\pi_y = \{y_0^\pi, ..., y_{T-1}^\pi\}$, is an optimal policy satisfying the conditions of the theorem. □

4.2.7 Proof of Theorem 21

Proof. Note that $\hat{B}_s \cup \{l_j | j \in J\} \cup S \cup i_0^0$ is a closed and bounded subset of $\mathbb{R}^2$ and thus is compact by the Heine-Borel Theorem. Thus, $conv(\hat{B}_s \cup \{l_j|j \in J\} \cup S \cup i_0^0)$ is closed since the convex hull of a compact set in $\mathbb{R}^2$ is compact and every compact set is closed. Similarly, since I is closed by assumption and the intersection of closed sets is closed, $\hat{I}' = I \cap conv(\hat{B}_s \cup \{l_j|j \in J\} \cup S \cup i_0^0)$ and $conv(\hat{I}')$ are closed. Furthermore, by construction, the boundary of $conv(\hat{B}_s \cup \{l_j|j \in J\} \cup S \cup i_0^0)$ is contained in I and thus is also contained in $\hat{I}'$ and $L = conv(\hat{I}') = conv(\hat{B}_s \cup \{l_j|j \in J\} \cup S \cup i_0^0)$. Thus, L is convex and closed, its boundary is contained in $\hat{I}'$, $proj_L(i) \in \hat{I}'$ $\forall i \in \mathbb{R}^2$, and $\hat{I}' \subseteq L$.

Let $\pi \in \Pi$ be an optimal policy such that $i_t^\pi(i_t^0, x_t) \notin \hat{I}'$ for some $i_t^0 \in \hat{I}'$, $x_t \ge 0$, and $t \in \{0, ..., T-1\}$. Let

$$Z_t(i_t^0, x_t, i_t, y_t) = f_{i_t^0 i_t} + K_{i_t}\delta(y_t - x_t) + G_t(i_t, y_t) - c_{i_t}x_t$$

Then

$$(i_t^\pi(i_t^0, x_t), y_t^\pi(i_t^0, x_t)) = \underset{i_t\in I, y_t\ge x_t}{\arg\min}\, Z_t(i_t^0, x_t, i_t, y_t)$$

Consider the policy $\pi_p = \{i_0^{\pi_p}, ..., i_{T-1}^{\pi_p}\}$ such that

$$\begin{aligned}
i_t^{\pi_p}(i_t^0, x_t) &= proj_L(i_t^\pi(i_t^0, x_t)) && \forall\ i_t^0 \in \hat{I}',\ x_t \ge 0,\ t \in \{0, ..., T-1\} \\
i_t^{\pi_p}(i_t^0, x_t) &= i_t^\pi(i_t^0, x_t) && \forall\ i_t^0 \in I \setminus \hat{I}',\ x_t \ge 0,\ t \in \{0, ..., T-1\} \\
y_t^{\pi_p}(i_t^0, x_t) &= y_t^\pi(i_t^0, x_t) && \forall\ i_t^0 \in I,\ x_t \ge 0,\ t \in \{0, ..., T-1\}
\end{aligned}$$

π_p is a feasible policy such that

$$i_t^{\pi_p}(i_t^0, x_t) \in \hat{I}' \quad \forall i_t^0 \in \hat{I}',\ x_t \ge 0,\ t \in \{0, ..., T-1\}$$

Note that $l_j \in L \; \forall j \in J \cup S$. By Lemma 1,

$$d(i,i') \geq d(i, proj_L(i')) \quad \forall i \in \hat{I}',\; i' \in \mathbb{R}^2$$

and

$$d(i,l_j) \geq d(proj_L(i), l_j) \quad \forall i \in \mathbb{R}^2,\; j \in J \cup S$$

By Lemma 2,

$$d(i,i') \geq d(proj_L(i), proj_L(i')) \quad \forall i, i' \in \mathbb{R}^2$$

Thus,

$$\min_{s \in S} d(i,l_s) \geq \min_{s \in S} d(proj_L(i), l_s)$$

Since $f(\cdot)$, $g_j(\cdot)$, $K(\cdot)$, and $c(\cdot)$ are non-decreasing $\forall j \in J$,

$$\begin{array}{lll}
f(d(i,i')) \geq f(d(i, proj_L(i'))) & \equiv f_{ii'} \geq f_{iproj_L(i')} & \forall i \in \hat{I}', i' \in \mathbb{R}^2 \\
f(d(i,i')) \geq f(d(proj_L(i), proj_L(i'))) & \equiv f_{ii'} \geq f_{proj_L(i)proj_L(i')} & \forall i, i' \in \mathbb{R}^2 \\
g_j(d(i,l_j)) \geq g_j(d(proj_L(i), l_j)) & \equiv g_{ij} \geq g_{proj_L(i)j} & \forall i \in \mathbb{R}^2, j \in J \\
K\left(\min_{s \in S} d(i,l_s)\right) \geq K\left(\min_{s \in S} d(proj_L(i), l_s)\right) & \equiv K_i \geq K_{proj_L(i)} & \forall i \in \mathbb{R}^2 \\
c\left(\min_{s \in S} d(i,l_s)\right) \geq c\left(\min_{s \in S} d(proj_L(i), l_s)\right) & \equiv c_i \geq c_{proj_L(i)} & \forall i \in \mathbb{R}^2
\end{array}$$

We will use these relations throughout the rest of the proof.

Note that

$$V_T(i_T^0, x_T) = 0 = V_T(proj_L(i_T^0), x_T) \quad \forall i_T^0 \in I,\; x_T \geq 0$$

Assume

$$V_{t+1}(i_{t+1}^0, x_{t+1}) \geq V_{t+1}(proj_L(i_{t+1}^0), x_{t+1}) \quad \forall i_{t+1}^0 \in I,\; x_{t+1} \geq 0$$

Consider period t. For $i_t^0 \in I \setminus \hat{I}'$ and $x_t \geq 0$,

$$Z_t(i_t^0, x_t, i_t^\pi(i_t^0, x_t), y_t^\pi(i_t^0, x_t)) = Z_t(i_t^0, x_t, i_t^{\pi_p}(i_t^0, x_t), y_t^{\pi_p}(i_t^0, x_t))$$

For $i_t^0 \in \hat{I}'$ and $x_t \geq 0$,

$$\begin{aligned}
G_t(i_t^\pi(i_t^0, x_t), y_t^\pi(i_t^0, x_t)) - c_{i_t^\pi(i_t^0, x_t)} &\geq G_t(proj_L(i_t^\pi(i_t^0, x_t)), y_t^\pi(i_t^0, x_t)) - c_{proj_L(i_t^\pi(i_t^0, x_t))} \\
&= G_t(i_t^{\pi_p}(i_t^0, x_t), y_t^{\pi_p}(i_t^0, x_t)) - c_{i_t^{\pi_p}(i_t^0, x_t)}
\end{aligned}$$

and thus

$$\begin{aligned}Z_t(i_t^0, x_t, i_t^\pi(i_t^0, x_t), y_t^\pi(i_t^0, x_t)) &\geq Z_t(i_t^0, x_t, proj_L(i_t^\pi(i_t^0, x_t)), y_t^\pi(i_t^0, x_t)) \\ &= Z_t(i_t^0, x_t, i_t^{\pi_p}(i_t^0, x_t), y_t^{\pi_p}(i_t^0, x_t))\end{aligned}$$

Thus, $i_t^{\pi_p}$ and $y_t^{\pi_p}$ must also be optimal policy functions for period t. Furthermore, for $i_t^0 \in I$ and $x_t \geq 0$,

$$\begin{aligned}V_t(i_t^0, x_t) &= Z_t(i_t^0, x_t, i_t^\pi(i_t^0, x_t), y_t^\pi(i_t^0, x_t)) \\ &\geq Z_t(proj_L(i_t^0), x_t, proj_L(i_t^\pi(i_t^0, x_t)), y_t^\pi(i_t^0, x_t)) \\ &\geq Z_t(proj_L(i_t^0), x_t, i_t^\pi(proj_L(i_t^0), x_t), y_t^\pi(proj_L(i_t^0), x_t)) \qquad (4.13) \\ &= V_t(proj_L(i_t^0), x_t)\end{aligned}$$

where 4.13 follows from the fact that

$$\left(i_t^\pi(proj_L(i_t^0), x_t), y_t^\pi(proj_L(i_t^0), x_t)\right) = \underset{i_t \in I, y_t \geq x_t}{\arg\min}\ Z_t(proj_L(i_t^0), x_t, i_t, y_t)$$

By induction, $i_t^{\pi_p}$ and $y_t^{\pi_p}$ are optimal policy functions for $t \in \{0, ..., T-1\}$ and

$$V_t(i_t^0, x_t) \geq V_t(proj_L(i_t^0), x_t) \quad \forall i_t^0 \in I,\ x_t \geq 0,\ t \in \{0, ..., T\}$$

Thus, $\pi^* = \{\pi_i^p, \pi_y^p\}$, where $\pi_i^p = \{i_0^{\pi_p}, ..., i_{T-1}^{\pi_p}\}$ and $\pi_y^p = \{y_0^{\pi_p}, ..., y_{T-1}^{\pi_p}\}$, is an optimal policy satisfying the conditions of the theorem. □

4.2.8 Proof of Theorem 22

Proof. Consider period $T-1$. For $i_{T-1}^0 \in I$ and $x_{T-1} \geq 0$, let

$$(i_{T-1}^*, y_{T-1}^*) = \underset{i_{T-1} \in I, y_{T-1} \geq x_{T-1}}{\arg\min}\ Z_{T-1}(i_{T-1}^0, x_{T-1}, i_{T-1}, y_{T-1})$$

Set $y_{T-1}^\pi(i_{T-1}^0, x_{T-1}) = y_{T-1}^*$. If $y_{T-1}^\pi(i_{T-1}^0, x_{T-1}) > x_{T-1}$ or $y_{T-1}^\pi(i_{T-1}^0, x_{T-1}) = x_{T-1}$ and $\exists j \in J$ such that $d(i_{T-1}^*, l_j) \leq d(i_{T-1}^0, l_j)$, set $i_{T-1}^\pi(i_{T-1}^0, x_{T-1}) = i_{T-1}^*$. Otherwise, $y_{T-1}^\pi(i_{T-1}^0, x_{T-1}) = y_{T-1}^* = x_{T-1}$ and $d(i_{T-1}^*, l_j) > d(i_{T-1}^0, l_j)$ $\forall j \in J$; for $i \in I$, let

$$\begin{aligned}F(i) = &f_{i_{T-1}^0 i} + \sum_{j \in J} p_{jT-1}[g_{ij} \min(y_{T-1}^*, d_{jT-1}) + h \max(0, y_{T-1}^* - d_{jT-1}) \\ &+ b_j \max(0, d_{jT-1} - y_{T-1}^*)]\end{aligned}$$

Note that $F(i) = Z_{T-1}(i_{T-1}^0, x_{T-1}, i, y_{T-1}^*)$,

$$f_{i_{T-1}^0 i_{T-1}^0} = f(0) \leq f(d(i_{T-1}^0, i_{T-1}^*)) = f_{i_{T-1}^0 i_{T-1}^*}$$

and

$$g_{i_{T-1}^0 j} = g_j(d(i_{T-1}^0, l_j)) \leq g_j(d(i_{T-1}^*, l_j)) = g_{i_{T-1}^* j} \quad \forall j \in J$$

Thus,

$$Z_{T-1}(i^0_{T-1}, x_{T-1}, i^0_{T-1}, y^*_{T-1}) = F(i^0_{T-1}) \leq F(i^*_{T-1}) = Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, y^*_{T-1})$$

Set $i^\pi_{T-1}(i^0_{T-1}, x_{T-1}) = i^0_{T-1}$. Defined in this way, i^π_{T-1} and y^π_{T-1} are optimal policies for period $T-1$ such that if $y^\pi_{T-1}(i^0_{T-1}, x_{T-1}) = x_{T-1}$, then $\exists j \in J$ such that $d(i^\pi_{T-1}(i^0_{T-1}, x_{T-1}), l_j) \leq d(i^0_{T-1}, l_j)$ $\forall$ $i^0_{T-1} \in I$ and $x_{T-1} \geq 0$.

Assume i^π_τ and y^π_τ are optimal policies for periods $\tau \in \{t+1, ..., T-1\}$ such that if $y^\pi_\tau(i^0_\tau, x_\tau) = x_\tau$ then $\exists j \in J$ such that $d(i^\pi_\tau(i^0_\tau, x_\tau), l_j) \leq d(i^0_\tau, l_j)$ $\forall$ $i^0_\tau \in I$, $x_\tau \geq 0$, and $\tau \in \{t+1, ..., T-1\}$.

Consider period t. For $i^0_t \in I$ and $x_t \geq 0$, let

$$(i^*_t, y^*_t) = \underset{i_t \in I, y_t \geq x_t}{\arg\min}\, Z_t(i^0_t, x_t, i_t, y_t)$$

Set $y^\pi_t(i^0_t, x_t) = y^*_t$. If $y^\pi_t(i^0_t, x_t) > x_t$ or $y^\pi_t(i^0_t, x_t) = x_t$ and $\exists j \in J$ such that $d(i^*_t, l_j) \leq d(i^0_t, l_j)$, set $i^\pi_t(i^0_t, x_t) = i^*_t$. Otherwise, $y^\pi_t(i^0_t, x_t) = y^*_t = x_t$ and $d(i^*_t, l_j) > d(i^0_t, l_j)$ $\forall j \in J$; let

$$\bar{i}^\pi_{t+1j} = i^\pi_{t+1}(i^*_t, \max(0, y^*_t - d_{jt}))$$
$$\bar{y}^\pi_{t+1j} = y^\pi_{t+1}(i^*_t, \max(0, y^*_t - d_{jt}))$$

and, for $i \in I$,

$$\begin{aligned} F(i) = & f_{i^0_t i} + \sum_{j \in J} p_{jt}[g_{ij} \min(y^*_t, d_{jt}) + h \max(0, y^*_t - d_{jt}) + b_j \max(0, d_{jt} - y^*_t) \\ & \qquad + f_{i\bar{i}^\pi_{t+1j}} + K_{\bar{i}^\pi_{t+1j}} \delta(\bar{y}^\pi_{t+1j} - \max(0, y^*_t - d_{jt})) \\ & \qquad + G_{t+1}(\bar{i}^\pi_{t+1j}, \bar{y}^\pi_{t+1j}) - c_{\bar{i}^\pi_{t+1j}} \max(0, y^*_t - d_{jt})] \end{aligned}$$

Note that $F(i) \geq Z_t(i^0_t, x_t, i, y^*_t)$ and $F(i^*_t) = Z_t(i^0_t, x_t, i^*_t, y^*_t)$ and

$$\begin{aligned} F(i^*_t) - F(i^0_t) = & \sum_{j \in J} p_{jt}[f_{i^0_t i^*_t} - f_{i^0_t i^0_t} + (g_{i^*_t j} - g_{i^0_t j}) \min(y^*_t, d_{jt}) \\ & \qquad + f_{i^*_t \bar{i}^\pi_{t+1j}} - f_{i^0_t \bar{i}^\pi_{t+1j}}] \end{aligned}$$

Let $f^0 = \lim_{d \to 0^+} f(d)$ and $f^+(d) = f(d) - f^0$ for $d \geq 0$. Note that $i^*_t \neq i^0_t$. If $i^*_t = \bar{i}^\pi_{t+1j}$,

$$f_{i^0_t i^*_t} + f_{i^*_t \bar{i}^\pi_{t+1j}} = f_{i^0_t \bar{i}^\pi_{t+1j}} + f(0) = f_{i^0_t i^0_t} + f_{i^0_t \bar{i}^\pi_{t+1j}}$$

If $i^0_t = \bar{i}^\pi_{t+1j}$, then $i^*_t \neq i^0_t \Rightarrow i^*_t \neq \bar{i}^\pi_{t+1j}$

$$\Rightarrow f_{i^0_t i^*_t} + f_{i^*_t \bar{i}^\pi_{t+1j}} \geq 2f(0) = f_{i^0_t i^0_t} + f_{i^0_t \bar{i}^\pi_{t+1j}}$$

Assume $i_t^* \neq \bar{i}_{t+1j}^\pi$ and $i_t^0 \neq \bar{i}_{t+1j}^\pi$. Note that $f^+(\cdot)$ is non-decreasing and, since sums of concave functions are concave, $f^+(\cdot)$ is concave over $(0, \infty)$. Then,

$$
\begin{aligned}
f_{i_t^0 i_t^*} + f_{i_t^* \bar{i}_{t+1j}^\pi} &= f(d(i_t^0, i_t^*)) + f(d(i_t^*, \bar{i}_{t+1j}^\pi)) \\
&= 2f^0 + f^+(d(i_t^0, i_t^*)) + f^+(d(i_t^*, \bar{i}_{t+1j}^\pi)) \\
&\geq 2f^0 + f^+(d(i_t^0, i_t^*) + d(i_t^*, \bar{i}_{t+1j}^\pi)) \qquad (4.14) \\
&\geq 2f^0 + f^+(d(i_t^0, \bar{i}_{t+1j}^\pi)) \qquad (4.15) \\
&= f^0 + f(d(i_t^0, \bar{i}_{t+1j}^\pi)) \\
&\geq f(0) + f(d(i_t^0, \bar{i}_{t+1j}^\pi)) \qquad (4.16) \\
&= f_{i_t^0 i_t^0} + f_{i_t^0 \bar{i}_{t+1j}^\pi}
\end{aligned}
$$

Where 4.14 follows from Lemma 4, 4.15 follows from the triangle inequality, and 4.16 follows from the assumption $f(0) \leq \lim_{d \to 0^+} f(d) = f^0$. Furthermore, note that

$$
g_{i_t^* j} - g_{i_t^0 j} = g_j(d(i_t^*, l_j)) - g_j(d(i_t^0, l_j)) \geq 0 \quad \forall j \in J
$$

by the assumption that $g(\cdot)$ is non-decreasing and $d(i_t^*, l_j) > d(i_t^0, l_j)$ $\forall j \in J$. Thus,

$$
\begin{aligned}
& F(i_t^*) - F(i_t^0) \geq 0 \\
\Rightarrow\ & Z_t(i_t^0, x_t, i_t^0, y_t^*) \leq F(i_t^0) \leq F(i_t^*) = Z_t(i_t^0, x_t, i_t^*, y_t^*)
\end{aligned}
$$

Set $i_t^\pi(i_t^0, x_t) = i_t^0$. Defined in this way, i_t^π and y_t^π are optimal policies for period t such that if $y_t^\pi(i_t^0, x_t) = x_t$ then $\exists j \in J$ such that $d(i_t^\pi(i_t^0, x_t), l_j) \leq d(i_t^0, l_j)$ $\forall$ $i_t^0 \in I$ and $x_t \geq 0$.

By induction, i_t^π and y_t^π are optimal policy functions for $t \in \{0, ..., T-1\}$ and $\pi^* = \{\pi_i, \pi_y\}$, where $\pi_i = \{i_0^\pi, ..., i_{T-1}^\pi\}$ and $\pi_y = \{y_0^\pi, ..., y_{T-1}^\pi\}$, is an optimal policy satisfying the conditions of the theorem. □

4.2.9 Proof of Theorem 23

Proof. Consider period $T-1$. For $i_{T-1}^0 \in I$ and $x_{T-1} \geq 0$, let

$$
(i_{T-1}^*, y_{T-1}^*) = \underset{i_{T-1} \in I, y_{T-1} \geq x_{T-1}}{\arg\min} Z_{T-1}(i_{T-1}^0, x_{T-1}, i_{T-1}, y_{T-1})
$$

Set $i_{T-1}^\pi(i_{T-1}^0, x_{T-1}) = i_{T-1}^*$. If $i_{T-1}^* \notin I_m$ or $i_{T-1}^* \in I_m$ and $y_{T-1}^* = x_{T-1}$ or $x_{T-1} < \max_{j \in J} d_{jT-1}$, set $y_{T-1}^\pi(i_{T-1}^0, x_{T-1}) = y_{T-1}^*$. Otherwise, $i_{T-1}^* \in I_m$, $y_{T-1}^* > x_{T-1}$, and $x_{T-1} \geq \max_{j \in J} d_{jT-1}$; for $y \geq x_{T-1} \geq \max_{j \in J} d_{jT-1}$, let

$$
\begin{aligned}
F(y) = & f_{i_{T-1}^0 i_{T-1}^*} + K_{i_{T-1}^*} \delta(y - x_{T-1}) + c_{i_{T-1}^*}(y - x_{T-1}) \\
& + \sum_{j \in J} p_{jT-1}[g_{i_{T-1}^* j} d_{jT-1} + h(y - d_{jT-1})]
\end{aligned}
$$

Note that $F(y^*_{T-1}) = Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, y^*_{T-1})$. $F(y)$ is increasing in y over $[x_{T-1}, \infty)$. Thus,

$$Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, x_{T-1}) = F(x_{T-1}) \leq F(y^*_{T-1}) = Z_{T-1}(i^0_{T-1}, x_{T-1}, i^*_{T-1}, y^*_{T-1})$$

Set $y^\pi_{T-1}(i^0_{T-1}, x_{T-1}) = x_{T-1}$. Defined in this way, i^π_{T-1} and y^π_{T-1} are optimal policies for period $T-1$ such that if $i^\pi_{T-1}(i^0_{T-1}, x_{T-1}) \in I_m$ and $x_{T-1} \geq \max_{j\in J} d_{jT-1}$ then $y^\pi_{T-1}(i^0_{T-1}, x_{T-1}) = x_{T-1}$ $\forall$ $i^0_{T-1} \in I$, $x_{T-1} \geq 0$.

Assume i^π_τ and y^π_τ are optimal policies for periods $\tau \in \{t+1, ..., T-1\}$ such that if $i^\pi_\tau(i^0_\tau, x_\tau) \in I_m$ and $x_\tau \geq \max_{j\in J} d_{j\tau}$ then $y^\pi_\tau(i^0_\tau, x_\tau) = x_\tau$ $\forall$ $i^0_\tau \in I$, $x_\tau \geq 0$, and $\tau \in \{t+1, ..., T-1\}$.

Consider period t. For $i^0_t \in I$ and $x_t \geq 0$, let

$$(i^*_t, y^*_t) = \underset{i_t \in I, y_t \geq x_t}{\arg\min}\, Z_t(i^0_t, x_t, i_t, y_t)$$

Set $i^\pi_t(i^0_t, x_t) = i^*_t$. If $i^*_t \notin I_m$ or $i^*_t \in I_m$ and $y^*_t = x_t$ or $x_t < \max_{j\in J} d_{jt}$, set $y^\pi_t(i^0_t, x_t) = y^*_t$. Otherwise, $i^*_t \in I_m$, $y^*_t > x_t$, and $x_t \geq \max_{j\in J} d_{jt}$; let

$$\begin{aligned}\bar{i}^\pi_{t+1,j} &= i^\pi_{t+1}(i^*_t, y^*_t - d_{jt})\\ \bar{y}^\pi_{t+1,j} &= y^\pi_{t+1}(i^*_t, y^*_t - d_{jt})\end{aligned}$$

and, for $y \geq x_t \geq \max_{j\in J} d_{jt}$,

$$\begin{aligned}F(y) =& f_{i^0_t i^*_t} + K_{i^*_t}\delta(y - x_t) + c_{i^*_t}(y - x_t)\\ &+ \sum_{j\in J} p_{jt}[g_{i^*_t j} d_{jt} + h(y - d_{jt}) + f_{i^*_t \bar{i}^\pi_{t+1,j}} + K_{\bar{i}^\pi_{t+1,j}}\delta(\bar{y}^\pi_{t+1,j} - (y - d_{jt}))\\ &\qquad + G_{t+1}(\bar{i}^\pi_{t+1,j}, \bar{y}^\pi_{t+1,j}) - c_{\bar{i}^\pi_{t+1,j}}(y - d_{jt})]\end{aligned}$$

Note that $F(y^*_t) = Z_t(i^0_t, x_t, i^*_t, y^*_t)$. By the assumption that $K(\cdot)$ and $c(\cdot)$ are non-decreasing, $i^*_t \in I_m \Rightarrow K_{i^*_t} \geq K_i$ and $c_{i^*_t} \geq c_i$ $\forall i \in I$; thus,

$$\begin{aligned}F(y^*_t) - F(x_t) &= K_{i^*_t} + c_{i^*_t}(y^*_t - x_t)\\ &\quad + \sum_{j\in J} p_{jt}[h(y^*_t - x_t) + K_{\bar{i}^\pi_{t+1,j}}[\delta(\bar{y}^\pi_{t+1,j} - (y^*_t - d_t)) - \delta(\bar{y}^\pi_{t+1,j} - (x_t - d_{jt}))]\\ &\qquad - c_{\bar{i}^\pi_{t+1,j}}(y^*_t - x_t)]\\ &\geq K_{i^*_t} + c_{i^*_t}(y^*_t - x_t) + \sum_{j\in J} p_{jt}[-K_{\bar{i}^\pi_{t+1,j}} - c_{\bar{i}^\pi_{t+1,j}}(y^*_t - x_t)]\\ &= K_{i^*_t} - \sum_{j\in J} p_{jt} K_{\bar{i}^\pi_{t+1,j}} + (c_{i^*_t} - \sum_{j\in J} p_{jt} c_{\bar{i}^\pi_{t+1,j}})(y^*_t - x_t)\\ &\geq 0\end{aligned}$$

Thus,

$$Z_t(i_t^0, x_t, i_t^*, x_t) \leq F(x_t) \leq F(y_t^*) = Z_t(i_t^0, x_t, i_t^*, y_t^*)$$

Set $y_t^\pi(i_t^0, x_t) = x_t$. Defined in this way, i_t^π and y_t^π are optimal policies for period t such that if $i_t^\pi(i_t^0, x_t) \in I_m$ and $x_t \geq \max_{j \in J} d_{jt}$ then $y_t^\pi(i_t^0, x_t) = x_t \; \forall \; i_t^0 \in I$, $x_t \geq 0$.

By induction, i_t^π and y_t^π are optimal policy functions for $t \in \{0, ..., T-1\}$ and $\pi^* = \{\pi_i, \pi_y\}$, where $\pi_i = \{i_0^\pi, ..., i_{T-1}^\pi\}$ and $\pi_y = \{y_0^\pi, ..., y_{T-1}^\pi\}$, is an optimal policy satisfying the conditions of the theorem. □

5 Conclusion and Areas for Future Research

In this dissertation, we examined how to make optimal relocation and inventory management decisions for a mobile inventory in response to changing demand forecasts.

We first analyzed the dynamic relocation problem. We modeled it using dynamic programing. As large instances of the problem take a long time to solve to optimality, we presented conditions under which it is sufficient for both the stationary and dynamic relocation problems to consider a subset of the potential inventory locations and thereby reduce the size of the problems. On the operational level, we showed that under certain conditions, it is optimal in relocating the mobile inventory to move the inventory as slowly as possible along the direct path between the two locations. We proved that a similar result holds in the case of a fixed cost but that the fixed cost shortens the amount of time over which it is optimal to move the inventory at a constant, slow rate. We proved each of these results in both discrete and continuous time.

Additionally, we examined the value of inventory mobility over traditional warehouse prepositioning in the context of the dynamic relocation problem. We proved analytically that the value is weakly greater in systems which have lower movement costs and in systems which have costs to serve that are greater in a positive linear transformation sense. Our numerical results suggest that the value is greater in systems 1) that have less concentrated risk, meaning that the potential disaster sites are farther from each other, 2) in which potential disaster sites that are relatively far from each other have demands that are less correlated, 3) that have demands which are relatively rare and of high magnitude rather than frequent and of low magnitude, and 4) that serve sites which have demand characteristics more similar to tropical cyclones rather than earthquakes. The experiment that suggests this last result shows that certain potential disaster site demand characteristics lead to a greater value of inventory mobility, specifically, site demands that are easier to predict (e.g. cyclone-related demand with weather forecasts) and sites that have greater probabilities of positive demand and greater potential demand levels. Furthermore, as large problem instances take a long time to solve to optimality, we developed the Base State Heuristic (BSH) which is optimal when the demands are temporally independent. The BSH solved the cases we tested in our numerical experiments within 0.5% of optimality in a small fraction of the time required by an exact algorithm.

We also examined the joint dynamic facility relocation and inventory management (DRIM) problem. As we did for the dynamic relocation problem, we modeled the DRIM problem using dynamic programming. We first examined a special case of the problem to develop intuition on optimal relocation and inventory management policies and specifically on the trade-off of when it is advantageous to move toward a supplier at the expense of moving away from potential disaster sites and vice versa. We then returned to the general DRIM problem and developed a number of analytical results characterizing the optimal relocation and inventory management policies. Research analyzing dynamic inventory management problems often prove that multiperiod (s, S) policies are optimal; similarly, we proved that such a policy is optimal when we are restricted to movement policies that do not depend on the starting inventory level. However, we also show that because optimal relocation decisions vary with starting inventory levels and different relocation decisions affect ordering prices, a multiperiod (s, S) policy is not necessarily optimal for the problem when our movement policy can depend on the starting inventory level. Despite this, we proved that, in an optimal policy, we will either order nothing or order up to some sum of potential future period demands. A corollary to this result is that, in an optimal policy, we will not place an order if our initial inventory level is weakly greater than the maximum possible demand for the rest of the horizon.

In the context of the optimal relocation decisions, similar to the result we proved for the dynamic relocation problem, we proved that it is sufficient for optimality to consider a smaller feasible set of inventory locations and thereby reduce the size of the problem. We were also interested in developing results which help us understand how the optimal dynamic relocation and inventory management policies relate to each other. Along these lines, we proved that, in an optimal policy, we will not move the mobile inventory toward the supplier at the expense of moving farther from all potential disaster sites in a period in which we do not place an order. We also proved that, in an optimal policy, we will not place an order if we are in a location farthest from the supplier and our current inventory level is weakly greater than the maximum possible current period demand.

To our knowledge, we are the first in the literature to 1) consider the dynamic alteration of a disaster relief supply chain in response to changing demand patterns over time, 2) examine mobile inventory systems for disaster relief, 3) study dynamic facility location with demand which evolves according to a non-stationary DTMC, and 4) consider both dynamic relocation and inventory management decisions with stochastic demand. We hope there will be future work in these areas considering some of the many important extensions of this research.

For the dynamic relocation problem, as large problem instances take a long time to solve to optimality, we developed the BSH which is able to solve the cases we tested to near optimality in a small fraction of the time required by an exact algorithm. As the joint dynamic relocation and inventory management problem is even more complex than the dynamic relocation problem, it will be important for future work to develop specialized heuristics capable of solving the problem more quickly than exact solution methods.

Another area for further research is the analysis of the value of inventory mobility in disaster relief. Theorems 8 and 9 provide bounds on the value of inventory mobility; it may be worthwhile to find tighter bounds as well as simpler bounds which can be computed without solving optimization problems and to quantify when the bounds are tight. It also will be important to consider the other ways in which inventory mobility may provide value, such as decreased response times, lower total inventory levels and inventory costs, less dependency on surviving infrastructure, and increased inventory security, and to develop a more general definition of the value of inventory mobility in disaster relief.

In this research, we examined problems in which the length of the planning horizon is pre-determined and known. An area for future research is to analyze how to optimally set the length of the planning horizon. Alternatively, it may be interesting to explore a model in which the end of the horizon is unknown and determined by the stochastic occurrence of a disaster.

Note that forecasts on potential locations and severity of disasters can be continually updated since worldwide seismic, political, weather, and other data is constantly updated and available. Another interesting extension would thus be to include forecast updating. For this extension it would be useful to find some daily managerial insights such as: When is it necessary to re-run the model and when is it appropriate to stick with the current plan? Can some kind of confidence level for disaster updates be found for this trade-off? Perhaps there is a cost to change the plan due to the need to re-coordinate entities such as the ship operator, decision makers at headquarters, and relief operations the ship is already supporting elsewhere. Can an appropriate trade-off be found in this case?

All models discussed in this dissertation are limited to considering the optimal movement of a single mobile inventory. However, in a real disaster relief system, organizations may opt for using multiple supply holding ships, hurricane relief supply holding containers, or mobile health clinics in their operations rather than just a single mobile inventory or facility. Thus, the models should be extended to consider the optimal management of a system of mobile inventories or facilities. However, this would require us to greatly expand our state space to include the locations of each of the ships and the dimension of the control, making the problem much more complicated and difficult to analyze and solve. For this reason, it may be most important in this case to focus on finding heuristics which find close to optimal solutions more quickly than exact solution algorithms.

It may also be important for the optimal management of a mobile inventory system for disaster relief to consider refueling policies. The state and control spaces would need to be further expanded to include fuel levels and decisions on when to refuel. A further extension that would make the models more realistic would be to consider the fact that fuel levels may affect ship performance and thus movement costs.

A ship's movement is subject to both the control of the ship operator as well as ocean currents. Thus, the models discussed in this dissertation could be extended to allow for probabilistic movement of the ship with the currents. In this case, the state transition function will no longer be a deterministic function of the control and the DP equations will need to be updated accordingly. There may also be additional control options. As in the models presented, it will be possible to choose to have the ship remain in the same location; however, there may also be either the option of turning off the engines and letting the ship be subject to the currents at little to no cost and thus allowing the next period location of the ship to be stochastic or the option of expending fuel to keep the ship in the exact current location. If the network is large enough, however, considering probabilistic movement may not be worthwhile.

It is also of interest to consider other possible objective functions and relevant costs. For example, disaster relief organizations have severely limited budgets. Thus, it may be appropriate to set the objective to maximizing the expected demand covered or minimizing the expected losses/shortfall subject to a set budget constraint. Furthermore, the current models minimize expected costs; it may be more appropriate to minimize expected cost plus a risk measure like CVaR (see, e.g., Noyan (2012)) or some kind of worst-case-type measure.

As suggested in Section 4.1, we may interpret the shortage cost as the cost to serve a potential disaster site from an outside supplier or a stationary warehouse rather than the mobile inventory. However, if we are unable to meet demand in a reasonable amount of time through these means, it is very difficult to interpret or quantify a "shortage cost" as the "customers" in the disaster relief context are people who may suffer or lose their lives if they do not receive the supplies. Thus, it would be useful for future work to consider how to calibrate the shortage cost appropriately in the joint dynamic relocation and inventory management model.

With the joint dynamic relocation and inventory management model, it may be worthwhile for future work to allow for more than one disaster to occur in a given time period and to consider more realistic ways of modeling demand. Furthermore, it will be important for future work to remove the assumption that relocation and order lead times are instantaneous.

Other decisions which will be important to consider include modes and speeds of transportation. For example, ships can move at different speeds at different costs. Additionally, the cost to provide relief to a disaster at a certain location i is currently governed according to the deterministic cost function g_{ij}. However, this cost may be stochastic or there may be several options. For example, the relief organization may have the choice between using a small ship to transport the supplies to the disaster site or using a helicopter, at a significant cost increase, to get the supplies to the disaster site much more quickly.

This brings up a further need for extension as the models described in this dissertation do not directly model the lead time of getting supplies from the mobile inventory to the disaster sites and the effects of the length of this lead time. The length of this lead time has a direct impact on an organization's ability to serve affected communities and on the survival, health, and recovery of these communities and its consideration will be an important extension area for future research.

Bibliography

Akkihal, A., 2006. Inventory Pre-positioning for Humanitarian Operations. Master's thesis, Massachusetts Institute of Technology, Cambridge, MA.

Amouzegar, M., McGarvey, R., Tripp, R., Luangkesorn, L., Lang, T., and Roll, C., 2006. *Evaluation of Options for Overseas Combat Support Basing*. RAND Corporation, Santa Monica, California.

Arabani, A., and Farahani, R., 2012. Facility Location Dynamics: An Overview of Classifications and Applications. *Computers & Industrial Engineering*, 62: 408-420.

Balcik, B., and Beamon, B.M., 2008. Facility Location in Humanitarian Relief. *International Journal of Logistics: Research and Applications*, 11 (2): 101-121.

Ballou, R., 1968. Dynamic Warehouse Location Analysis. *Journal of Marketing Research*, 5 (3): 271-276.

Barzinpour, F., and Esmaeili, V., 2014. A Multi-Objective Relief Chain Location Distribution Model for Urban Disaster Management. *The International Journal of Advanced Manufacturing Technology*, 70 (5-8): 1291-1302.

Beamon, B.M., and Kotleba, S.A., 2006. Inventory Modeling for Complex Emergencies in Humanitarian Relief Operations. *International Journal of Logistics: Research and Applications*, 9 (1): 1-18.

Beamon, B.M., and Kotleba, S.A., 2006. Inventory Management Support Systems for Emergency Humanitarian Relief Operations in South Sudan. *International Journal of Logistics Management*, 17 (2): 187-212.

Berman, O., and Odoni, A., 1982. Locating Mobile Servers on a Network with Markovian Properties. *Networks*, 12: 73-86.

Bozkurt, M., and Duran, S., 2012. Effects of Natural Disaster Trends: A Case Study for Expanding the Pre-Positioning Network of CARE International. *International Journal of Environmental Research and Public Health*, 9: 2863-2874.

Campbell, A.M., and Jones, P.C., 2011. Prepositioning Supplies in Preparation for Disasters. *European Journal of Operations Research*, 209 (2): 156-165.

Campbell, J., 1990. Locating Transportation Terminals To Serve an Expanding Demand. *Transportation Research Part B: Methodological*, 24 (3): 173-192.

Canel, C., Khumawala, B., Law, J., and Loh, A., 2001. An Algorithm for the Capacitated, Multi-Commodity Multi-Period Facility Location Problem. *Computers & Operations Research*, 28: 411-427.

Chand, S., 1988. Decision/Forecast Horizon Results for a Single Facility Dynamic Location/Relocation Problem. *Operations Research Letters*, 7 (5): 247-251.

Chow, J., and Regan, A., 2011. Resource Location and Relocation Models with Rolling Horizon Forecasting for Wildland Fire Planning. *INFOR: Information Systems and Operational Research*, 49 (1): 31-43.

Current, J., Ratick, S., and ReVelle, C., 1997. Dynamic Facility Location when the Total Number of Facilities is Uncertain: A Decision Analysis Approach. *European Journal of Operational Research*, 110: 597-609.

Dias, J., Captivo, M.E., and Clímaco, J., 2007. Efficient Primal-Dual Heuristic for a Dynamic Location Problem. *Computers & Operations Research*, 34: 1800-1823.

Döyen, A., Aras, N., and Barbarosoğlu, G., 2012. A Two-Echelon Stochastic Facility Location Model for Humanitarian Relief Logistics. *Optimization Letters*, 6 (6): 1123-1145.

Drezner, Z., and Wesolowsky, G.O., 1991. Facility Location when Demand is Time Dependent. *Naval Research Logistics*, 38: 763-777.

Duran, S., Gutierrez, M.A., and Keskinocak, P., 2011. Pre-Positioning of Emergency Items for CARE International. *Interfaces*, 41 (3): 223-237.

Erlenkotter, D., 1981. A Comparative Study of Approaches to Dynamic Location Problems. *European Journal of Operational Research*, 6: 133-143.

Farahani, R., Drezner, Z., and Asgari, N., 2009. Single Facility Location and Relocation Problem with Time Dependent Weights and Discrete Planning Horizon. *Annals of Operations Research*, 167: 353-368.

Farahani, R., Abedian, M., and Sharahi, S., 2009. Dynamic Facility Location Problem. In Farahani, R., and Hekmatfar, M., (Eds.), *Facility Location: Concepts, Models, Algorithms and Case Studies. Contributions to Management Science* (pp. 347-372). Springer-Verlag.

da Gama, S.F., and Captivo, M., 1998. A Heuristic Approach for the Discrete Dynamic Location Problem. *Location Science*, 6: 211-223.

Ghaderi, A., and Jabalameli, M., 2013. Modeling the Budget-Constrained Dynamic Uncapacitated Facility LocationNetwork Design Problem and Solving it via Two Efficient Heuristics: A Case Study of Health Care. *Mathematical and Computer Modelling*, 57: 382-400.

Gue, K., 2003. A Dynamic Distribution Model for Combat Logistics. *Computers & Operations Research*, 30: 367-381.

Guha-Sapir, D., Below, R., and Hoyois., Ph., 2014. EM-DAT: International Disaster Database. Université Catholique de Louvain, Brussels, Belgium, http://www.emdat.be.

Guha-Sapir, D., Hoyois, Ph., and Below., R., 2015. *Annual Disaster Statistical Review 2014: The Numbers and Trends*. CRED, Brussels, Belgium.

Guha-Sapir, D., Hoyois, Ph., and Below, R., 2014. *Annual Disaster Statistical Review 2013: The Numbers and Trends*. CRED, Brussels, Belgium.

Haghani, A., and Oh, S.-C., 1996. Formulation and Solution of a Multi-commodity, Multimodal Network Flow Model for Disaster Relief Operations. *Transportation Research Part A: Policy and Practice*, 30 (3): 231-250.

Halper, R., and Raghavan, S., 2011. The Mobile Facility Routing Problem. *Transportation Science*, 45 (3): 413-434.

Hinojosa, Y., Kalcsics, J., Nickel, S., Puerto, J., and Velten, S., 2008. Dynamic Supply Chain Design with Inventory. *Computers & Operations Research*, 35: 373-391.

Houming, F., Tong, Z., Xiaoyan, Z., Mingbao, J., and Guosong, D., 2008. Research on Emergency Relief Goods Distribution after Regional Natural Disaster Occurring. *International Conference on Information Management, Innovation Management and Industrial Engineering, 2008*, 3: 156-161.

Jena, S.D., Cordeau, J., and Gendron, B., 2015. Dynamic Facility Location with Generalized Modular Capacities. *Transportation Science*, Articles in Advance: 1-16.

Jornsten, K., and Bjorndal, M., 1994. Dynamic Location Under Uncertainty. *Studies in Regional & Urban Planning*, 3: 163-184.

McCall, V., 2006. Designing and Prepositioning Humanitarian Assistance Pack-up Kits (HA PUKs) to Support Pacific Fleet Emergency Relief Operations. Master's thesis, Naval Postgraduate School, Monterey, CA.

McGarvey, R., Tripp, R., Rue, R., Lang, T., Sollinger, J., Conner, W., and Luangkesorn, L., 2010. *Global Combat Support Basing: Robust Prepositioning Strategies for Air Force War Reserve Materiel.* RAND Corporation, Santa Monica, California.

Melachrinoudis, E., Min, H., and Wu, X., 1995. A Multiobjective Model for the Dynamic Location of Landfills. *Location Science*, 3 (3): 143-166.

Melachrinoudis, E., and Min, H., 2000. The Dynamic Relocation and Phase-out of a Hybrid, Two-Echelon Plant/Warehousing Facility: A Multiple Objective Approach. *European Journal of Operational Research*, 123: 1-15.

Melo, M.T., Nickel, S., and da Gama, S. F., 2005. Dynamic Multi-Commodity Capacitated Facility Location: A Mathematical Modeling Framework for Strategic Supply Chain Planning. *Computers & Operations Research*, 33: 181-208.

National Hurricane Center, 2015. Forecast Accuracy. Florida International University, Miami, Florida, http://www.nhc.noaa.gov/.

Notteboom, T. and Carriou, P., 2009. Fuel surcharge practices of container shipping lines: Is it about cost recovery or revenue making? Proceedings of the 2009 International Association of Maritime Economists (IAME) Conference, June, Copenhagen, Denmark.

Noyan, N., 2012. Risk-Averse Two-Stage Stochastic Programming with an Application to Disaster Management. *Computers & Operations Research*, 39: 541-559.

Owen, S., and Daskin, M., 1998. Strategic Facility Location: A Review. *European Journal of Operational Research*, 111: 423-447.

Ozbay, K., and Ozguven, E.E., 2007. Stochastic Humanitarian Inventory Control Model for Disaster Planning. *Transportation Research Record: Journal of the Transportation Research Board*, 2022: 63-75.

Ozbay, K., and Ozguven, E.E., 2013. A Secure and Efficient Inventory Management System for Disasters. *Transportation Research Part C: Emerging Technologies*, 29: 171-196.

Ozdamar, L., Ekinci, E., and Kucukyazici, B., 2004. Emergency Logistics Planning in Natural Disasters. *Annals of Operations Research*, 129 (1-4): 328-245.

Powell, W., 2007. *Approximate Dynamic Programming.* John Wiley & Sons, Inc., Hoboken, New Jersey.

Qiu, J., and Sharkey, T., 2013. Integrated Dynamic Single-Facility Location and Inventory Planning Problems. *IIE Transactions*, 45: 883-895.

Rawls, C., and Turnquist, M., 2010. Pre-Positioning of Emergency Supplies for Disaster Response. *Transportation Research Part B: Methodological*, 44 (4): 521-534.

Rawls, C., and Turnquist, M., 2011. Pre-Positioning Planning for Emergency Response with Service Quality Constraints. *OR Spectrum*, 33: 481-498.

Rawls, C. and Turnquist, M., 2012. Pre-Positioning and Dynamic Delivery Planning for Short-Term Response Following a Natural Disaster. *Socio-Economic Planning Sciences*, 46: 46-54.

Richter, A., and Shen, Z.-J.M., and Shanthikumar, J.G., 2015. Dynamic Facility Relocation and the Value of Inventory Mobility in Disaster Relief. Manuscript submitted for publication.

Romauch, M., and Hartl., R.F., 2005. Dynamic Facility Location with Stochastic Demands. In Lupanov, O.B., Kasim-Zade, O.M., Chaskin, A.V., and Steinhöfel, K., (eds.), *Stochastic Algorithms: Foundations and Applications* (pp. 180-189). Springer-Verlag, Heidelberg, Germany.

Rosenthal, R., White, J., and Young, D., 1978. Stochastic Dynamic Location Analysis. *Management Science*, 24 (6): 645-653.

Rottkemper, B., Fischer, K., Blecken, A., and Danne, C., 2011. Inventory Relocation for Overlapping Disaster Settings in Humanitarian Operations. *OR Spectrum*, 33: 721-749.

Shen, Z.-J.M., Coullard, C., and Daskin, M., 2003. A Joint Location-Inventory Model. *Transportation Science*, 37 (1): 40-55.

Sheu, J., 2007. An Emergency Logistics Distribution Approach for Quick Response to Urgent Relief Demand in Disasters. *Transportation Research Part E*, 43: 687-709.

Shreve, S.E., 2005. Abbreviated Proof [In the Lost Sales Case]. In Bertsekas, D.P., *Dynamic Programming and Optimal Control, Vol. 1, 3rd ed.* (pp. 204-206). Athena Scientific, Nashua, New Hampshire.

Shu, J., 2010. An Efficient Greedy Heuristic for Warehouse-Retailer Network Design Optimization. *Transportation Science*, 44 (2): 183-192.

Shulman, A., 1991. An Algorithm for Solving Dynamic Capacitated Plant Location Problems with Discrete Expansion Sizes. *Operations Research*, 39 (3): 423-436.

Snyder, L.V., 2006. Facility Location Under Uncertainty. *IIE Transactions*, 38: 537-554.

Snyder, L.V., Daskin, M., and Teo, C.P., 2007. The Stochastic Location Model with Risk Pooling. *European Journal of Operational Research*, 179: 1221-1238.

Snyder, L.V., and Shen, Z.-J.M., 2011. *Fundamentals of Supply Chain Theory.* John Wiley & Sons, Inc., Hoboken, New Jersey.

Sweeney, D., and Tatham, R., 1976. An Improved Long-Run Model for Multiple Warehouse Location. *Management Science*, 22 (7): 748-758.

Taskin, S., and Lodree, E.J., 2010. Inventory Decisions for Emergency Supplies Based on Hurricane Count Predictions. *International Journal of Production Economics*, 126 (1): 66-75.

Thanh, P.N., Bostel, N. and Péton, O., 2008. A Dynamic Model for Facility Location in the Design of Complex Supply Chains. *International Journal of Production Economics*, 113: 678-693.

U.S. Senate, Committee on Homeland Security and Governmental Affairs, 109th Congress, 2nd Session, 2006. *Hurricane Katrina: A Nation Still Unprepared. Special Report.* U.S. G.P.O., Washington, DC.

Ukkusuri, S., and Yushimito, W., 2008. Location Routing Approach for the Humanitarian

Prepositioning Problem. *Transportation Research Record: Journal of the Transportation Research Board*, 2089: 18-25.

Verma, A., and Gaukler, G., 2011. A Stochastic Optimization Model for Positioning Disaster Response Facilities for Large Scale Emergencies. *Network Optimization*, 6701: 547-552.

Wesolowsky, G., 1973. Dynamic Facility Location. *Management Science*, 19 (11): 1241-1248.

Wesolowsky, G., and Truscott, W., 1975. The Multiperiod Location-Allocation Problem with Relocation of Facilities. *Management Science*, 22 (1): 57-65.