THE IMPACT OF DEMAND UNCERTAINTY ON STOCKPILE AND DISTRIBUTION
DECISIONS DURING INFLUENZA PANDEMIC

by

ANDREW M. WALDMAN

A THESIS

submitted in partial fulfillment of the requirements for the degree

MASTER OF SCIENCE

Department of Industrial and Manufacturing Systems Engineering

College of Engineering

KANSAS STATE UNIVERSITY

Manhattan, Kansas

2014

Approved by:

Major Professor
Dr. Jessica L. Heier Stamm

Copyright

ANDREW WALDMAN
2014

Abstract

The main goal of public health emergency preparedness efforts is to mitigate the impact of events on the health of the population. However, decision-makers must also remain conscientious of the costs associated with these efforts. Planning is further complicated by uncertainty about the location and volume of demand that will need to be met in an emergency, the speed with which demand must be met, and the potential scarcity of needed items once an emergency occurs. To address these challenges, public health emergency planners often keep inventory stockpiles that are distributed when an event happens. Managing these stockpiles is a difficult task, and inefficient stockpile location and equipment distribution strategies can be costly both in terms of cost and public health impact.

This research is motivated by challenges faced by state public health departments in creating stockpile location and equipment distribution strategies. The primary emphasis is on facemasks and respirators used by health workers during an influenza pandemic, but the approach is generalizable to other scenarios. The model proposed here uses a two-stage approach to generate a holistic solution to the problem. The first stage uses a pull distribution strategy to make stockpile location decisions. Additionally, it determines how counties should be assigned to stockpiles to minimize both storage and distribution costs. The second stage adopts a push distribution strategy to determine optimal delivery routes based on the county assignments made in stage one. This stage offers guidance for public health planners who have made location-allocation decisions but who then face a different distribution scenario than what was anticipated in the original planning phase. Recourse methods for managing demand uncertainty are also proposed.

A case study of the state of Kansas is conducted using the methods introduced in the thesis. The computational results yield several significant insights into the tradeoffs and costs of various facility location-allocation and vehicle routing decisions:

- For the tested range of storage and distribution cost parameters, multiple stockpile locations are preferred over a single location.
- In a pull distribution system, storage costs play a greater role in location-allocation decisions than distribution costs.

- In the push distribution system, finding an optimal vehicle routing plan is computationally intensive for stockpiles with a large number of assigned counties.
- Efficient heuristics perform well to design recourse routing plans when realized demand is greater than expected.
- In the event that planners wish to specify routes well in advance, the results of this research suggest adopting a robust routing plan based on higher-than-expected demand levels.

This thesis makes three important contributions. The first is an optimization approach that considers multiple distribution strategies. This is especially relevant when stockpiling for an influenza pandemic where stockpiles need to be located significantly before the material is needed, during which time the distribution strategy may change. Second, the case study demonstrates that the proposed methods are applicable to a large-scale problem arising in practice. Finally, this research illustrates for decision-makers the tradeoffs between different stockpile management strategies and between optimal and heuristic methods.

Table of Contents

List of Figures

List of Tables

Acknowledgements

I would like to take this opportunity to thank my major advisor, Dr. Jessica Heier Stamm, whose guidance, patience, and encouragement have allowed me to apply my degree beyond the classroom. The wisdom and industrial engineering skillset I learned from her will carry with me well beyond the completion of my degree.

I would also like to thank the members of my committee, Dr. David Ben-Arieh and Dr. Kerry Priest. Their backgrounds in this area of research will provide diverse perspectives and help broaden the applicability of this thesis. I am grateful to have had the opportunity to work with them on this project. Additionally, thank you to Zac Graves, Michael McNulty, and Emily Farley at KDHE. Their collaboration and input has made this research possible.

Finally, I would like to thank my family and the students, faculty, and staff of the IMSE Department at Kansas State University. Through these relationships, I have been challenged to pursue my passions both inside the classroom and out. I am thankful to have had such positive role models to shape my experience.

Chapter 1 - Introduction

During an influenza pandemic, the virus spreads quickly and without warning through a large population. A large-scale emergency such as this can lead to a surge in the demand for protective equipment. However, since these events are rare, it is difficult for public health planners to develop efficient preparedness strategies. The work in this thesis is motivated by the need to locate, allocate, and distribute stockpiles of protective equipment when demand is unknown. Although this research focuses specifically on facemasks and N95 respirators used in an influenza pandemic, the approach is generalizable to other scenarios.

1.1 Research Motivation

This research is motivated by a partnership with the Kansas Department of Health and Environment (KDHE) and considers stockpile location-allocation and vehicle routing decisions made at Kansas' state level. Currently, equipment stockpiles are maintained locally by hospitals, public health departments, and other health care organizations. There is limited visibility at the state level of the total quantity and location of protective equipment. KDHE aims to establish state-level stockpiles, also known as medical materiel caches, to improve pandemic preparedness.

KDHE seeks to understand the advantages and disadvantages of maintaining a single centralized materiel cache versus multiple cache locations throughout the state. These decisions are impacted by the cost of storing and distributing the stockpiled equipment. Different distribution strategies are also considered. In conversations with Zac Graves, Medical Countermeasures Program Manager at KDHE, he mentioned that if a central materiel cache were used, KDHE would distribute supplies to every county in the state. This "push" distribution strategy may also be adopted in an urgent situation to get initial equipment inventory into local communities. On the other hand, if a stockpile were opened in each of the seven Kansas health regions, each county could potentially travel to the stockpile to pick up materials. This is termed a "pull" distribution strategy. This research considers the implications of opening a combination of any of the seven stockpiles and seeks to quantify the costs and tradeoffs between scenarios where realized demand is greater than expected.

Influenza pandemics can pose a number of challenges for public health responders due to demand uncertainty. As a result, there is a need for tools to help planners make more informed and justifiable decisions. Operations research methods are particularly suited to help guide these decisions. While many models in the literature address related problems, to the author's knowledge, this optimization approach is unique in that it uses multiple distribution strategies to address problems that may arise when stockpiles need to be located significantly before the material is needed, during which time the distribution strategy may change. Additionally, this model uses optimization uncertainty with recourse to quantify tradeoffs between optimal and heuristic distribution strategies.

1.2 Research Goals

The aim of this thesis is to study the effect of demand uncertainty on facility location-allocation and vehicle routing decisions. More informed decisions in these applications could save innumerable lives by stockpiling the greatest amount of resources at the lowest possible cost, thus increasing availability of protective equipment. The first goal of this work is to determine the number of stockpiles to open and to which counties they should be assigned to minimize total storage and distribution costs. The second goal of the research, which depends directly on the results of the stockpile location-allocation decisions made in the first stage, is to provide recommendations about the routes that vehicles should take to minimize cumulative distance traveled while accounting for demand uncertainty.

1.3 Contributions

This thesis makes three important contributions to operations research applications in emergency preparedness. The first is an optimization approach that uses a two-stage, holistic approach to the problem to consider multiple distribution strategies. This is especially relevant when stockpiling for an influenza pandemic where stockpiles need to be located significantly before the material is needed, in which time the distribution strategy may change. The first stage determines where stockpiles should be located based on the assumption that a pull distribution strategy will be employed. The second stage then looks at county assignments and solves the vehicle routing problem assuming a push distribution strategy. Optimization with recourse is used to display how vehicle routes can be modified to account for realized demand that is greater

than expected. This knowledge will allow public health officials to implement routes that will perform well, regardless of realized demand.

Second, the case study demonstrates that the proposed methods are applicable to a large-scale problem arising in practice. Working closely with KDHE, the research is implemented in the state of Kansas to quantify the effects of various facility location-allocation and vehicle routing decisions. By accounting for scenarios with uncertain demand and incorporating optimization with recourse, this study shows how the model can be used to mitigate the cost of corrective actions that may be needed when realized demand is greater than expected.

Finally, this research illustrates for decision-makers the tradeoffs between different stockpile management strategies and between optimal and heuristic decision-making methods. This is advantageous for public health officials because it provides a method for determining vehicle routes quickly, or when someone is not readily available to modify the input parameters used in the optimization model.

1.4 Outline

This thesis begins with a brief overview of the problem being addressed by this research and the current challenges faced by public health decision-makers. Additionally, Chapter 1 summarizes the goals of this research and how the results provide relevant contributions to emergency preparedness. Chapter 2 reviews the literature relevant to past influenza pandemics and their effects on public health planning. It then explores deterministic, stochastic, and robust models of facility location and vehicle routing problems. Finally, the review looks at a few ways these operations research techniques can be applied in the public health sector. Chapter 3 outlines the general methods used to solve the models introduced in this research. These include a two-stage optimization model, as well as optimization under uncertainty with recourse. Chapter 4 discusses the specific assumptions and parameters used in a computational study with KDHE. This chapter also assesses the results of this study and provides a recommendation by which public health planners can make informed decisions. Finally, Chapter 5 summarizes the findings of this paper and generalizes recommendations based on trends in the findings. To conclude, the chapter indicates relevant areas of future work.

Chapter 2 - Literature Review

This chapter introduces literature relevant to this research. The first section discusses the background of pandemics on both a national and international level, and how these events influence public policy and emergency preparedness decisions. The next section looks specifically at facility location and vehicle routing methods, and how these types of problems are modified to account for different types of uncertainty. These two problem classes are most directly related to this research and are further developed in Chapters 3 and 4. The third section introduces various applications of facility location and vehicle routing models in public health scenarios. Finally, the last section outlines the differentiating elements of this thesis and their contributions to operations research applications in emergency preparedness.

2.1 Pandemic and Emergency Preparedness

This research focuses on infectious disease and pandemics because they are the easiest for which to stockpile materials. Unlike a natural disaster, they do not necessarily depend on the number of people affected by the event, but rather by the number of emergency responders that may be needed. Emergency preparedness planning is especially relevant for an influenza pandemic because its onset is rapid and unpredictable. Demand often overwhelms the supply of commercially-available products during the course of the outbreak. In these situations, it is important that responders have access to an adequate supply of protective equipment such as facemasks and N95 respirators. Another benefit to stockpiling these materials is that they are universal and not dependent on the strain of the influenza virus.

An important differentiation must be made between the levels of severity with which an infectious disease can affect a population. More specifically, severity refers to how widespread the disease is. From an emergency preparedness perspective, this influences the number of people who may require protection. In general, an influenza virus can present itself as an outbreak, an epidemic, or a pandemic [50].

Outbreaks occur when a disease infects a greater number of people than expected. This typically confines itself to a community or a region, although it can appear in several countries. The unique element, though, is that outbreaks can last for years, appearing in different places of the world at different times, and manifesting in new communities or in communities where the

disease has been absent for a long period of time. Epidemics are different in that they spread much more rapidly than outbreaks, although the terms are often used interchangeably. When an outbreak or an epidemic appears in several countries at the same time, it may develop into a pandemic. Pandemic refers to a disease outbreak that covers a much wider geographical area, generally caused by a virus strain that humans have no or very little immunity against. A key characteristic of this new strain is that it is easily transmissible between humans or from animal to human [36].

The influenza virus exists in two forms, known as Influenza A and Influenza B. Influenza B viruses circulate only among humans, while Influenza A viruses also are found in animals such as chicken and pigs. When an Influenza A virus develops a subtype that can infect humans – which occurred with the avian flu virus in 1997 and 1999 – it is likely to then develop to the point where it can spread from human to human. When this happens, there is typically no vaccine available, resulting in a pandemic that rapidly spreads to numerous other countries [48].

While seasonal influenza outbreaks are common, influenza pandemics are more rare, although more likely to occur as people become highly internationally mobile. According to [37], there have only been four influenza pandemics since 1900. In 1918, the Spanish flu affected 20 – 40% of the population worldwide. This pandemic was unique in that there were high mortality rates among healthy adults. When it ended in 1919, an estimated 50 million people died from the virus, nearly 700,000 of which were from the United States. However, unlike other pandemics, many victims did not die from the flu itself, but often from pneumonia and other complications caused by bacteria. Because the virus was not quickly identified, there was little that could be done to prepare for its onset [37].

In 1957, an influenza virus was identified in Asia to which most people under the age of 65 had little immunity. Predicting a pandemic, health officials monitored outbreaks and produced a limited supply of vaccine. During the pandemic, a panel of experts from the World Health Organization (WHO) found that spread in some countries followed public gatherings, such as conferences and festivals, and broke out first in camps, army units, and schools. They advised that people avoid crowding in public places in the hope that this would reduce the severity of the pandemic [39].

However, when the virus came to the United States in the summer of 1957, it was not immediately detected. Children were returning to school and had the opportunity to spread the

influenza virus before it was discovered. The disease spread rapidly in classrooms and peaked among school children and young adults, who then brought it home to their families. By December, officials believed the virus was gone. However, another wave appeared in the early months of 1958. In the end, approximately two million people died during this pandemic, about 70,000 of which were from the United States [37].

A "second wave" is not uncommon during a pandemic, and must be accounted for when stockpiling emergency equipment that will be needed. Not only does it influence the number of people who will require treatment, but it also lengthens the timespan for which equipment must be stockpiled. This concept is not incorporated in the research because, as was the case in the 1957-58 pandemic, the second wave usually occurs months after the first, allowing sufficient time to replenish the stockpile. Chapter 3 will discuss how the length of the pandemic is incorporated in the model and how it can be easily manipulated to account for a longer time frame, which may be necessary if considering the effects of a second wave.

In 1968, a new subtype of the influenza virus emerged in Hong Kong. It was the mildest flu pandemic to date, with a death toll in the United States of around 34,000. There are a number of potential reasons why this particular pandemic affected fewer people. In some ways, it was similar to the 1957 pandemic flu virus, so many Americans may have developed an immunity to the strain. Additionally, it did not reach the United States until December when students were away from school; thus, they did not have the chance to spread the virus to one another. Finally, there was improved medical care and vaccine availability for infected patients [37].

The most recent influenza pandemic occurred in 2009. The virus was nicknamed "swine flu" because of its resemblance to a strain found in pigs. The Centers for Disease Control and Prevention (CDC) estimates that 43 – 89 million people had the H1N1 virus, but estimates only 10,000 – 20,000 deaths [37]. There are many reasons this pandemic resulted in fewer infections and deaths compared to earlier pandemics, the primary reason being that the virus caused less severe illness despite being easily transmissible [24]. Additionally, organizations like WHO and the CDC increased their level of preparedness, as well as the urgency with which they were able to create a vaccine for the virus strain [37].

According to MacKellar [34], "Policies to respond to pandemic influenza fall into three time frames-measures that can be taken before the emergence of a new virus, measures that can be undertaken in the immediate aftermath of its emergence, and measures that can be taken once

6

the pandemic has been established." The third time frame would include research done to develop and administer a new vaccine to those infected. The research in this thesis, though, is primarily concerned with the first time frame; specifically, the steps that can be taken to ensure health care employees have adequate access to emergency equipment during an influenza pandemic [38]. In 2005, the U.S. Department of Health and Human Services (HHS) created the National Pandemic Preparedness and Response Task Group. This organization, which works with the CDC and other international agencies, works to prepare for potential pandemics. In addition to vaccine development and production, they invest heavily in research related to the distribution of vaccines and emergency equipment. Other strategic decisions include stockpiling of antiviral medications, risk and communications, and international partnership opportunities [21].

Planning for a crisis such as an influenza pandemic raises ethical concerns that must be addressed by public health officials. In situations where health needs greatly exceed available human and material resources, difficult decisions must be made regarding where and to whom materials are made available. Although scientific evidence and applications can help make these decisions, they often fail to consider the overall effect on people; they cannot determine, for example, whether a decision is just. Despite the fact that ethics have little to no contribution to the understanding of an influenza virus, they do contribute to debates such as who is most willing to be inconvenienced by a situation, and how should the burdens of negative outcomes be distributed across the population. It is important for decision makers to balance the ethical component with the urgency of logistical and scientific needs. Failing to do so could result in loss of the public's trust. Using an ethical framework to guide decision-making can help mitigate some of the unavoidable backlash from an influenza pandemic [47].

The concept of ethical decision making in emergency preparedness is relevant to this research because the decisions made can affect the lives of thousands of people. For example, it is not uncommon to intentionally stockpile in counties with a higher population density, despite the fact that it may greatly increase cost. The expectation here is that more supplies will be readily available for a greater number of people. However, in addition to potentially increasing cost, using this population-weighted approach may move resources further away from less populated counties.

There are a wide range of decisions that must be made in emergency preparedness settings, such as where to stockpile resources and how to distribute those resources to the affected population. Operations research provides a method by which researchers can incorporate parameters to make the best, most informed decisions.

2.2 Facility Location and Vehicle Routing Methods

There are many ways in which operations research can be used to influence decisions and policies made during an influenza pandemic. This research is most concerned with facility location and vehicle routing applications of operations research. One aspect of these types of problems is demand uncertainty. For this reason, nearly all of the models in the literature incorporate stochastic or robust optimization. This section will introduce various methods used to solve general problems of these types, and it will also discuss how the models can be modified to account for risk and uncertainty.

Decision models can be divided into three categories: deterministic, stochastic, and robust. In situations where it is reasonable to assume that parameters are known at the time that models are solved, deterministic methods are appropriate. However, because this assumption is not reasonable when preparing for an influenza pandemic, models that explicitly account for uncertainty are needed. According to Rosenhead et al. [42], in some situations, parameters are uncertain but their values follow a probability distribution that is known to the decision maker. This information is used to develop stochastic optimization models where the objective is generally to optimize some expected value. In other situations, these probability distributions are not known, and robust optimization models are used to optimize the worst-case performance of a system. In both models, the purpose is to find a solution that will perform well regardless of the random parameter's realized value [46].

In general, facility location problems are *NP*-hard. The main purpose is to determine where to open p facilities from n possible locations, and p is either specified in advance or determined by the optimization model. In a deterministic situation where input values are known, the model also decides which demands are assigned to each open location. A simple example of this is the Weber problem, in which a single facility is placed to minimize the demand-weighted sum of distances from a given set of sites [9]. The p-median problem, introduced by Hakimi [25], expanded the Weber problem to consider the locations of p unique facilities, still with the

objective of minimizing some measure of transportation cost such as distance. The Hakimi property states that there is an optimal solution where facilities are located on the nodes of the network rather than along the edges. Because parameters are known, the problem is still deterministic [46].

The Hakimi property was further developed by Louveaux [33], who presented the capacitated p-median problem with uncertain demands and costs. In this type of stochastic facility location problem, the most common objective is to optimize the mean expected outcome. The goal of this model is to choose facility locations and decide which customers they will serve. Since demands are random and facilities are capacitated, a facility may not be able to satisfy the demands of all the counties assigned to it, so a penalty for unmet demand is introduced.

Another section of literature discusses models that use a probabilistic approach to make facility location decisions. The goal of these models is to maximize the probability that the outcome is good, or constraining the probability that it is bad. The exact definitions of "good" and "bad" are solution-specific. Frank [19] was first to introduce "max-probability" centers, which are "points that maximize the probability that the maximum weighted distance from the point is within a given limit." This research is furthered by Berman et al. [4], who consider the effect of uniformly distributed demand instead of normally distributed. Carbone [6] introduced chance-constrained programming as it applies to the facility location problem. These models utilize constraints to limit the degree to which an objective value may be undesirable to the decision maker.

Hodgson [26] performed sensitivity analysis on both deterministic and stochastic solutions. Using simulation to estimate regret, the research found that optimal p-median solutions are relatively insensitive to errors in travel distances. Interestingly, he also found that solutions are especially insensitive to errors in demand. Cooper [10] performs sensitivity analysis on the problem itself. Expanding on the Weber problem, he considers how the solution is affected when the locations of demand points are not known with certainty, but rather lie within "uncertainty circles." Drezner [16] generalizes the Weber problem on a sphere with random customer locations. His research concludes that as the number of demand points approaches infinity, the difference between minimum and maximum costs approaches zero.

When probability distributions are unknown, it is necessary to develop robust optimization models to decide facility locations. In these models, uncertain parameters are

captured using a scenario approach, where deterministic values represent potential realizations of these parameters. The two most common measures of robustness are cost and regret, the latter of which refers to the difference between the cost of a particular solution in a given scenario and the optimal cost that could be achieved in that scenario with perfect foresight. Some models seek to optimize the minimax cost or regret, meaning they find the solution that minimizes the maximum cost or regret across all scenarios. This type of optimization finds the best worse-case scenario.

Labbé et al. [32] consider the facility location problem on a network with uncertain node weights; these weights are only estimated. Given an optimal solution for the deterministic problem, the model seeks to find the maximum regret where weights differ from their estimates by no more than a predetermined tolerance gap. Carrizosa [7] elaborates on this research by redefining the robustness measure used by Labbé et al. Instead, the model seeks to find the minimum value of the tolerance gap required to violate a cost limit. The more the tolerance is able to change without exceeding the limit, the more robust the solution.

Gupta and Rosenhead [23] introduce the robustness measures when decisions are made over time. Schilling [43] expands upon this by using a scenario-based approach. A set-covering model is used to maximize the number of facilities in common among all scenarios while satisfying all demands. In the first model, a fixed number of facilities are opened in each scenario, while future models do not make this specification. This identifies a tradeoff between the total number of facilities opened and the number of facilities in common. However, Daskin et al. [12] showed that Schilling's methods produce the worst possible results when transportation costs are taken into consideration. Finally, Kouvelis et al. [30] developed the concept of *p*-robustness, although the term was coined by Snyder and Daskin [45]. This method introduces a constraint on the regret in any scenario to *p*.

The second portion of this research is concerned with the vehicle routing problem, which is also *NP*-hard. The deterministic version of this problem consists of a set of customers, each with a known location and known demand, which are to be supplied by a set of vehicles with known capacity. The objective is typically to minimize some cost, often distribution cost. All vehicle routing models share a series of constraints necessary to find a feasible solution:

- demands of all customers must be met,
- vehicles cannot exceed capacity,
- the route must not contain subtours.

The vehicle routing problem was first considered by Dantzig and Ramser [13]. They introduced a heuristic that used linear programming (LP) models to solve for a single-depot problem. A key assumption of their model is that a customer's demand is met whenever it is serviced. This implies that it is only necessary (and optimal) to stop at each customer exactly one time. Garvin [20] relaxes this assumption by using a more complicated mixed integer LP, and Balinski and Quandt [3] introduce a binary integer LP formulation. Golden et al. [22] modified these applications to incorporate multiple depots. In this type of model, it is important to clarify whether the vehicle must end its tour at the same depot at which it started.

2.3 Facility Location and Vehicle Routing Applications in Emergency Preparedness

There are two types of uncertainty that manifest themselves in this research: time uncertainty and demand uncertainty. Time is twofold in that it refers to both when the next pandemic will occur and for how long it will affect a population. Demand refers to the number of people that will be affected by the disease. Once again, from the perspective of emergency preparedness, this can have multiple interpretations. If stockpiling antiviral medications, for example, demand is the predicted number of infected people that will require the medication. In other situations, it may be more appropriate to stockpile emergency equipment for public health responders, as is the case with this research. In the latter strategy, demand is generally less random. This section will look at how these two types of uncertainty have been considered specifically in the public health setting, and how researchers have used this information to make facility location and vehicle routing decisions.

As mentioned earlier, influenza pandemics are rare. There have only been four since 1900, although there have been many more threats and perceived pandemics in that time [37]. When a pandemic does occur, most recently with the H1N1 virus in 2009-10, the situation is further complicated because planners cannot base their decisions solely on experience from earlier pandemics. For the most part, they can only incorporate this information into a mathematical model that can then be used to project possible scenarios and influence control strategies. However, even this approach is reactive in nature and does not help guide decisions that must be made before a pandemic is perceived [18].

To determine how resources can best be allocated for a pandemic influenza, some research considers the effect of stockpile location and

Alternatively, some researchers use an economic approach. Drake, Chalabi, and Coker [14] note that all probabilistic models to date assume that pandemics occur according to a Poisson stochastic process, sometimes expressed as a constant annual probability of a pandemic. However, timing is the second largest source of pandemic parameter uncertainty after mortality rate [44].

Drake et al. evaluate uncertainty using a case study of influenza antiviral stockpiling in Cambodia. They compare this to having no stockpile, and evaluate based on cost and disability adjusted life years (DALYs). Essentially, they use a present value cost analysis to assess the benefits of a given stockpiling strategy. Whereas the stockpile investment occurs right now, health gains are realized in the future and must be discounted back to present value. They found that the time-to-pandemic parameter has elasticity slightly greater than two. This means that as the time-to-pandemic value changes, it has a greater than proportional impact on the incremental cost-effectiveness ratio (ICER) [15].

One difficulty in using an economic model is that countries have unique characteristics that need to be taken into consideration. For example, a stockpiling strategy in one country may not work in another due to accessibility of the stockpile, which may result in a misallocation of limited resources. It is important to keep in mind the landscape in which the health scenario may take place [8].

Carrasco et al. [8] use an epidemic-economic model to study how mortality and cost are affected by stockpile sizes for Brazil, China, Guatemala, India, Indonesia, New Zealand, Singapore, the United Kingdom, the United States, and Zimbabwe. Not surprisingly, they noticed that antiviral stockpiling reduced mortality considerably, and had greater cost-avoidance potential in countries with more resources, such as the United States. However, they also estimated that, based on antiviral pricing at the time, stockpiling is not cost-effective for two-thirds of the world's population [8]. When the scope of the problem is narrowed to encompass a more uniform population such as a single state, these particular economic factors have less influence on stockpiling decisions. It is then easier to directly see the effects of these decisions on cost.

Many facility location and vehicle routing applications in health care use a similar cost minimization approach. Daskin [53] was one of the first to consider how reliability of the location model influenced facility location decisions. He uses a maximum covering integer programming (IP) model to incorporate a facility's ability or inability to serve a demand based on the probability that it is busy. Drezner [55] also considers reliability, but he introduces the fact that if a customer cannot be served by a given facility, it will have to be served by an alternative facility.

Connecting both facility location and inventory, Balcik and Beamon [54] develops a maximal covering model that determines the number and location of the centers, as well as the amount of inventory that should be held at each. They use a set of scenarios and the probability of each scenario happening to incorporate uncertainty into the model. Huang et al. [52] use a stochastic optimization model to estimate the inventory levels of ventilators throughout Texas. The model quantifies risk as either the expected number of patients not receiving required material, or the probability that at least one patient in the state will not receive material Campbell and Jones [5] were the first to consider inventory levels in a facility location model without using scenarios. They choose to do this because scenarios are generally based on significant historical data, which may not always be available in disaster response scenarios. Instead, they introduce a risk parameter that is the probability that any inventory stored at a facility will be destroyed or inaccessible. This approach relies on historical natural disaster data, which is more accessible and accurate.

Finally, Rawles and Turnquist [40] look at emergency preparedness with service quality constraints. They developed a model to minimize expected costs while incorporating penalties for unmet demand. Using stochastic optimization, the two-stage model seeks to minimize the total expected cost while constraining the quality of the response to the uncertain parameters. To expand, they want to ensure that the distance stockpiled supplies have to be moved as a result of a given scenario is within a specified limit. Similarly, the research done by Murali et al. [35] considers how facilities must be located in response to a bio-terror attack when there is possible damage to the transportation network. In some scenarios, it may be impossible to predict whether a customer is willing and able to travel to an assigned facility. These papers help demonstrate how transportation costs can be influenced by facility location decisions.

Following Hurricanes Katrina, Wilma, and Rita, which caused over $100 billion in damage, it became a priority for public health officials to look at the entire emergency preparedness network as a whole [41]. In addition to the research outlined above where facility locations were decided with the assumption that the transportation network would be altered following a natural disaster, researchers such as Rawls and Turnquist [41] took a more holistic approach. They developed a model that combined facility location, demand uncertainty, post-disaster distribution of supplies, survival of stockpiled materials, and the condition of the transportation network. This research is advantageous in that it most accurately represents the unanticipated behavior and outcomes of a hurricane or other natural disaster.

The research done by Özdamar et al. [51] is specifically concerned with vehicle routing problems in natural disasters. The model addresses transportation problems that must be solved at given time intervals during aid delivery. This is because the model considers ongoing aid where plans are regularly periodically updated to include new requests and modifications to the transportation network. In the context of this particular emergency preparedness plan, a customer in one period may become a depot in another due to a surplus of supplies. This is important because it implies that vehicles do not have to return to the node in which they started or even to depots at all. As a result, there are no closed-loop tours. Additionally, this model relaxes the constraint that a customer's demand must be satisfied by only one vehicle and instead allows for split delivery.

Finally, Créupt et al. [11] represent an emergency problem as a dynamic vehicle routing problem with time windows (DVRPTW). The setting of this problem, although still in the health care field, is unique. In this scenario, patients request doctors when needed, and an algorithm determines which doctor is assigned to the patient based on criteria such as distance, reaction time, and ability to respond to the patient's need. This application of the vehicle routing problem introduces unique opportunities for further research in similar settings. For example, rather than a stationary customer, an algorithm could assign vehicles to counties based on the spread of the pandemic throughout the state.

The model introduced in this research is different than those discussed in this chapter because it incorporates both facility location-allocation decisions as well as vehicle routing decisions. It incorporates a pull distribution strategy in the former and a push distribution strategy in the latter. This approach is especially justified when stockpiling for an influenza

pandemic because stockpiles need to be located significantly before the emergency equipment is needed, during which time the distribution strategy may change. Furthermore, this research applies multiple methods of optimization under uncertainty with recourse to display costs and tradeoffs with various solutions. This is of particular relevance to public health planners because they may not have the technical capabilities to re-optimize routes with computer software, so the model allows them to make decisions based on the recourse values of a heuristic.

Chapter 3 - Methods

During a pandemic, it is important for public officials to respond in a way that protects the most people within a reasonable amount of time. Advance preparation is key in achieving this goal. This chapter introduces the pandemic preparedness setting that motivates this research, as well as the considerations and assumptions that must be made in order to make the best decisions. Additionally, this chapter introduces the two-stage optimization approach used to solve the problem, and concludes with an analysis of the optimization uncertainty with recourse methods used to capture the demand uncertainty present in public health systems.

3.1 Problem Definition

It is widely accepted practice for health care facilities to carry a stockpile of equipment that may be needed in an emergency. Depending on the situation, this equipment may never be used even though it may have cost thousands or millions of dollars to develop the emergency preparedness plan, purchase necessary materials, and distribute the materials accordingly.

In an influenza pandemic, the disease is typically spread through the airborne transmission of particles or large droplets. This can be avoided by the simple use of OSHA-approved facemasks and respirators. The scope of this research is the management of facemask and respirator stockpiles for pandemic preparedness. To most efficiently respond to the situation, stockpiles may be maintained by individual facilities, groups of facilities, or state health officials. It is the responsibility of these organizations to provide emergency equipment to employees who will be most exposed to the infection in their daily work. Typically, these include public health department employees, hospital personnel, and EMS attendants. Based on their level of exposure, though, these employees do not use the same type of equipment, nor do they consume these materials at the same rate. Decision-support models are needed to account for these system characteristics in preparedness planning.

Decisions made in this type of emergency preparedness planning typically include where stockpiles should be opened, how customers should be assigned to those stockpiles, and how much material should be held in inventory. Unlike the models outlined in Chapter 2, this research expands beyond the stockpile location-allocation problem and incorporates a vehicle routing problem to generate a cost-effective distribution strategy. Decisions in this stage of the

model include which counties each vehicle will serve, and in what order the vehicles will deliver equipment. The methods by which these decisions are made are outlined in the following sections.

3.2 Modeling Approach

This section introduces the two-stage optimization approach used to model and solve this problem. The first stage determines which stockpiles should be open and which counties assigned to minimize both storage and distribution costs. This model uses a pull distribution strategy in which vehicles make round trips between the stockpile location and a single county. According to KDHE, this is likely to occur in events where counties pick up their supplies from the stockpile.

The second stage builds on the first stage decisions. However, it adopts a push distribution strategy, in which delivery vehicles are routed from the stockpile and visit multiple counties before returning to the stockpile location. This vehicle routing model minimizes the cumulative distance that vehicles must travel to satisfy the demand of every county. Thus, the second stage offers guidance for public health planners who have made location-allocation decisions for pandemic preparedness but who then face a different distribution scenario than what was anticipated in the original planning phase.

The second stage is dependent upon the first in that vehicles can only deliver to the given county assignments made in the first stage. Although the two stages are interdependent, they are modeled separately in this research for two reasons. First, in practice, it is likely that location-allocation decisions will be made well in advance of routing decisions, and some routing parameters or constraints may not be available at that time. Second, the two-stage approach is more computationally tractable for large problems.

Finally, the models are extended to assess the impact of demand uncertainty on the routing decisions made in the second stage. Using optimization under uncertainty with recourse, county demands are multiplied by a factor α to assess what will happen if the realized demand is greater than the estimated demand. Recourse is calculated based on the resulting increase in cumulative travel distance. A robust rather than stochastic approach is used because there is no historical data to justify an underlying probability distribution to the scenarios used in this model.

3.3 Stockpile Location-Allocation Model

This section describes the first stage of the two-stage approach: the stockpile location-allocation model. The result of this model determines which of the predetermined stockpiles are to be opened, as well as which counties are to be serviced by the given stockpile to minimize total storage and distribution costs. The model notation, parameters, and mathematical formulation are provided below.

3.3.1 Notation and Parameters

Most of the parameters in the first stage of the model are used to calculate equipment demand in each county. Also included is the distance matrix that quantifies how far counties are from each of the stockpile locations.

C	= set of county locations
S	= set of stockpile locations
H	= set of types of health equipment
D	= set of types of health department
P_d	= percentage of employees in department $d \in D$ that will require health equipment
F_s	= storage cost per square foot at stockpile facility $s \in S$
m	= distribution cost per mile
t	= duration of pandemic (in days)
w	= total number of cubic feet in a trailer
Q_h	= quantity of equipment $h \in H$ in one box
B_h	= number of boxes of equipment h that can be stacked in a 1' x 1' x 8' space
E_{cd}	= number of employees from county $c \in C$ in department d
N_{hd}	= number of equipment type h required by employee in department d
L_{cs}	= distance in miles from county c to stockpile s
λ_s	= county in which stockpile s is located

3.3.2 Decision Variables

The decisions made in the first stage of the model include which stockpiles are to be open and to which counties they will deliver.

Z_s = 1 if stockpile s is open;

       0 if otherwise

A_{cs} = 1 if county c is assigned to stockpile s;

       0 if otherwise

3.3.3 Mathematical Formulation

The complete formulation of the stockpile location-allocation model is as follows:

Minimize

$$\sum_{c=1}^{|C|}\sum_{d=1}^{|D|}\sum_{s=1}^{|S|}\sum_{h=1}^{|H|} \frac{(E_{cd}P_d\ N_{dh}\ t)}{(Q_h\ B_h)}\ F_s\ A_{cs}\ +$$

$$\sum_{c=1}^{|C|}\sum_{d=1}^{|D|}\sum_{s=1}^{|S|}\sum_{h=1}^{|H|} \left\lceil \frac{(E_{cd}P_d\ N_{dh}\ t)}{(Q_h\ B_h)}\ /\left(\frac{w}{8}\right)\right\rceil 2\ m\ L_{cs}\ A_{cs}$$

(**3.1**)

subject to

$$\sum_{s=1}^{|S|} A_{cs} = 1 \qquad \forall\, c \in C$$

(**3.2**)

$$A_{cs} \leq Z_s \qquad \forall\, c \in C, s \in S$$

(**3.3**)

$$Z_s \leq \sum_{c=1}^{|C|} A_{cs} \qquad \forall\, s \in S$$

(**3.4**)

$$A_{\lambda_s s} = Z_s \qquad \forall\, s \in S$$

(**3.5**)

$$A_{cs}, Z_s \in \{0,1\} \qquad \forall\, s \in S, c \in C$$

(**3.6**)

The components are described in more detail below.

OBJECTIVE: The main objective of the first stage of the model is to minimize the total cost accrued by storing emergency equipment at given stockpiles and then distributing the respirators to assigned counties throughout the state. One key assumption, made in conjunction with KDHE collaborators, is that there are no fixed costs associated with opening and maintaining a stockpile. This is because some of the stockpile locations are existing hospitals, thus it is not necessary for KDHE to open an entirely new building. For this reason, the objective function only takes storage and distribution costs into consideration when deciding stockpile locations and county assignments.

The first summation calculates the total storage cost of a given solution. Because storage cost is influenced by the stockpile's cost per square foot, it is necessary to first determine how many square feet of storage each county needs to occupy at the stockpile facility. This is represented by the fraction shared by both cost summations. In this term, the numerator corresponds to the total demand of a county for a pandemic that lasts t days. Demand is directly related to the number of employees in the county, E_{cd}, and the percentage of employees in each department that require emergency response equipment, P_d. The denominator corresponds to the number of pieces of equipment that can fit in 1' x 1' x 8' space. This number is found by calculating the maximum number of boxes that can fit in one cubic foot and then multiplying it by eight.

The second summation in Equation (3.1) considers the way in which vehicles will deliver the respirators to the counties. In this model, a multiplier of two is used, assuming that one truck will make exactly one round trip to each county. By doing so, it accounts for the worst case scenario, or upper bound, of total distance traveled. Although this assumption does not solve for optimal stockpile assignments, it is adequate in that it is feasible for all scenarios.

In order to calculate these distribution costs, it is necessary to calculate the number of trips that are required for one vehicle to satisfy the demand of a given county. This cost calculation can be seen in the second summation of Equation (3.1), where the model uses an arbitrarily chosen cost per mile, m, to see how counties are assigned to stockpiles when this variable increases or decreases. As is anticipated, cost per mile is directly correlated with the total distribution cost; thus it is also directly correlated with the number of open stockpiles.

The second summation calculates the total space requirement of the allotted equipment. However, in Equation (3.1), the number of trips that a vehicle must take is calculated by dividing the space requirement (in cubic feet) by $w / 8$. This is because the space requirement calculation already takes into consideration the fact that stacks are eight feet high, so the truck can only fit, at most, $w / 8$ stacks of equipment.

As mentioned earlier in this section, the objective function must account for both the storage and distribution costs. Therefore, it is calculated by simply summing the two cost equations. There are certain tradeoffs between the two independent variables (storage cost per square foot and distribution cost per mile) in this stage of the model. For example, it is intuitive that as storage cost per square foot increases, it is more desirable to store a county's equipment at a stockpile that is further away, thus resulting in fewer open stockpiles. Conversely, as distribution cost increases, it is more desirable to open more stockpiles so vehicles do not have to travel as far to make deliveries. These tradeoffs are further discussed in Chapter 4.

CONSTRAINTS: The constraints in the first stage of the model govern the feasible assignment of counties to stockpiles.

Constraint (3.2) states that every county must be assigned to exactly one stockpile, regardless of whether or not it is available. Constraint (3.3) specifies that the stockpile must be open in order for counties to be assigned to it. Together, these two constraints ensure that counties can only be assigned to exactly one open stockpile. Constraint (3.4) ensures that if a stockpile is open, it serves at least one county. Without the clarification of this constraint, it would be possible for the stockpile to be labeled as "open" even though it does not serve any counties.

Finally, Constraint (3.5) states that if a stockpile s is open, it must deliver to the county in which it is located, λ_s. This constraint is necessary because it is possible that, if storage costs are high and transportation costs are relatively low, it may be more desirable for one stockpile to be open and serve nearby counties, although receive its equipment from another stockpile. This type of assignment is undesirable from KDHE's perspective.

3.4 Vehicle Routing Model

Once the first-stage model determines which counties are assigned to which stockpile, the second stage consisted of a vehicle routing optimization model. The purpose of this model is to determine the route that a given number of vehicles must travel in order to minimize the cumulative distance traveled by those vehicles.

This approach demonstrates a push distribution strategy, where the delivery is not necessarily based on observed need, but rather as a means of ensuring local availability in anticipation of a sudden large demand. In a pull strategy, customers would order equipment as the pandemic began to infect more people. However, it is more appropriate for two reasons to assume the former in this situation. First, pandemics are generally unpredictable, random, and fast-spreading. Demand cannot necessarily be determined based on historical data or a forecasting model, as is generally the case with manufactured products. Instead, public health officials tend to take the more conservative approach of prepositioning stockpiles throughout the state. Secondly, a push system allows for more efficient system-wide planning on behalf of KDHE.

The following sections outline the notation and constraints used to develop the model and describe cutting-plane techniques that improve computational efficiency.

3.4.1 Notation and Parameters

The parameters of a vehicle routing problem include the demand of each county, the vehicle capacity, and the county-to-county distance matrix. Notations for these parameters are outlined below.

C $\quad$ = set of county locations

S $\quad$ = set of stockpile locations

V $\quad$ = set of vehicles

u $\quad$ = vehicle capacity

R_i $\quad$ = demand of county $i \in C$

G_{ij} $\quad$ = distance from county i to county j

The vehicle routing model is solved for each stockpile location and its assigned counties. Demand R_i of county i is assumed to be 25% of its total demand, or the supplies required for 30 days. The routing plan is designed to push out supply quickly, with the same routes to then be repeated later in the pandemic response.

3.4.2 Decision Variables

The major decision made in this stage of the model is the route taken by each vehicle. This is captured by the binary decision variable, X_{ijv}, which indicates whether vehicle v traverses the road between county i and county j. Additionally, the binary variable, Y_{iv}, indicates that county i is served by vehicle v.

X_{ijv} $\quad = 1$ if vehicle $v \in V$ traverses edge (i, j);

$\quad\quad$ 0 if otherwise

Y_{iv} $\quad = 1$ if county i is serviced by vehicle v;

$\quad\quad$ 0 if otherwise

3.4.3 Mathematical Formulation

The complete formulation of the vehicle routing model can be seen below. It is important to note that this model will only solve for a single stockpile location and its assigned counties.

$$\text{Minimize} \quad\quad \sum_{i=1}^{|C|}\sum_{j=1}^{|C|}\sum_{v=1}^{|V|} X_{ijv}\ G_{ij} \quad\quad\quad (3.7)$$

$$\text{subject to} \quad\quad \sum_{v=1}^{|V|} X_{iiv} = 0 \quad\quad \forall\ i \in C \quad\quad\quad (3.8)$$

$$\sum_{v=1}^{|V|} Y_{iv} = 1 \quad\quad \forall\ i \in C: i \neq S \quad\quad\quad (3.9)$$

23

$$\sum_{i=1}^{|C|}(R_i \, Y_{iv}) \leq u \qquad \forall \, v \in V \tag{3.10}$$

$$\sum_{v=1}^{|V|} Y_{iv} = |V| \tag{3.11}$$

$$\sum_{j=1}^{|C|} X_{jiv} = Y_{iv} \qquad \forall \, i \in C, v \in V \tag{3.12}$$

$$\sum_{j=1}^{|C|} X_{ijv} = \sum_{j=1}^{|C|} X_{jlv} \qquad \forall \, i \in C, l \in C, v \in V \tag{3.13}$$

$$X_{i,i+1} + X_{i+1,i+2} + \ldots + X_{n-1,n} \leq (n-1) \tag{3.14}$$

$$X_{ijv}, Y_{iv} \in \{0,1\} \qquad \forall \, i \in C, j \in C, v \in V \tag{3.15}$$

OBJECTIVE: In the second stage, the objective is to determine what routes the vehicles should take in order to minimize the cumulative distance traveled while still satisfying the demand of each county. The objective function simply sums the total distance traveled by every vehicle

CONSTRAINTS: This section will focus on the traditional constraints used to find a feasible solution to a vehicle routing problem. It is important to note that this model will only solve for a single stockpile location and its assigned counties.

Constraint (3.8) simply states that a vehicle cannot traverse edge (*i, i*). This constraint seems relatively intuitive and redundant, especially considering the fact that the distance matrix is designed so that the distance from county *i* to county *i* is zero. By traversing this edge, the

objective function is not affected, but no additional demand is satisfied either. However, when the model was first run, this issue arose, so Constraint (3.8) was introduced.

Constraint (3.9) states that each county in the set C must be visited exactly one time by one unique vehicle, with the exception of the stockpile county. By doing so, it is implied that the county's demand must be satisfied in full by that single delivery. In effect, this eliminates the possibility of dividing the demand amongst multiple vehicles. When considered in addition to truck capacity, it is easy to see how partitioning the demand may ultimately decrease the total miles traveled by the vehicles; however, from an implementation standpoint, it would not be desirable for multiple vehicles to visit the same county. In the worst case scenario, it is feasible and potentially optimal for one county to be visited by every vehicle. Although unlikely, the health care environment must be kept in consideration.

Constraint (3.10) ensures that a vehicle is not assigned to more counties than it can satisfy with a given capacity u.

Constraint (3.11) specifies that the stockpile – county i – must be served by exactly $|V|$ vehicles. This means that every vehicle originates at the county in which the stockpile is located.

The final two constraints are trivial in nature, but critical to any vehicle routing problem. Constraint (3.12) simply states that if a county i is assigned to a vehicle v, then that vehicle's route must include exactly one edge from some other county j into county i. Conversely, a county cannot be on a vehicle's route if it was not assigned to that vehicle. Finally, constraint (3.13) maintains the continuity of the route. In reality, if a vehicle travels from county i to county j, the next sequence in the route must be from county j to some other county l.

Constraint (3.14) is the general formulations of the subtour elimination constraints. These are discussed in detail in the next section.

3.4.4 Subtour Elimination Constraints

Subtours in the vehicle routing problem include any route that does not begin and end at the county in which the stockpile is located. Such solutions are not allowed. This section illustrates the concept of subtours and describes the computational approach to finding solutions without them.

A complete model formulation includes all subtour elimination constraints. In the implementation, these constraints are added one at a time because there are exponentially many.

Consider Figure 1 below, where the route is discontinuous. If subtour elimination constraints are not included, this tour is feasible because all of the constraints are satisfied: the truck goes to county z twice, all counties are visited, and it is assumed that vehicle capacity is not exceeded.
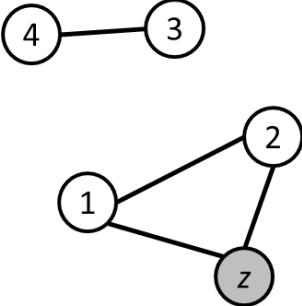

Figure 1. Subtour in which the route is discontinuous

It is clear that this particular route is not possible because it shows that one vehicle must travel route $z – 1 – 2 – z$ as well as $3 – 4 – 3$. However, the vehicle does not stop at every county while also satisfying the constraint that it must begin and end at the stockpile z. To ensure the optimal solution does not include these subtours, subtour elimination constraints are incorporated into the model. These constraints require that every county assigned to a vehicle must be included in the route, and that the route, in its entirety, must begin and end at the stockpile location. This is implemented using a script that identifies subtours.

The script first looks at the given optimal solution. If it notices that a vehicle has a subtour as defined above, it writes a constraint to eliminate that unique subtour from the feasible region. Finally, it resolves to find a new optimal solution, and the process is repeated until there are no subtours in any vehicle's route. In the example introduced in Figure 1, the subtour elimination constraint takes the form of a cover cut:

$$X_{Z1} + X_{12} + X_{Z2} \leq 2 \qquad\qquad (\textbf{3.16})$$

This cut states that only two of the three arcs that comprise the subtour can exist in the final solution. However, as mentioned above, this cut only eliminates a single subtour. Further iterations may introduce more subtours, as shown in Figure 2. The constraint in Equation (3.16) is still satisfied, yet the solution is still not feasible.
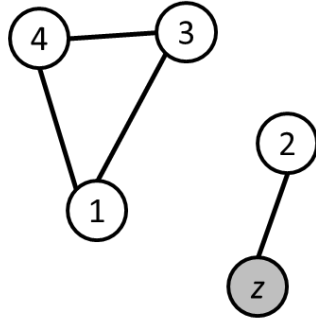
Figure 2. Another solution containing a subtour

For a vehicle that must travel to *C* counties, subtours may range from size 2 to size $|C| - 2$ counties. In general, the summation of edges in a subtour consisting of *n* nodes (where $n < |C|$) must be less than or equal to $n - 1$, as shown in Equation (3.14).

It is easy to see why run time is so strongly correlated with the number of counties being serviced by a stockpile. In the example with five nodes, every solution containing subtours has one subtour of size three and one subtour of size two. Therefore, the model may need to complete ten iterations (five choose two) to eliminate all possible subtours. As $|C|$ increases, so too does the number of iterations. The optimal solution to the above example would be the route shown in Figure 3.
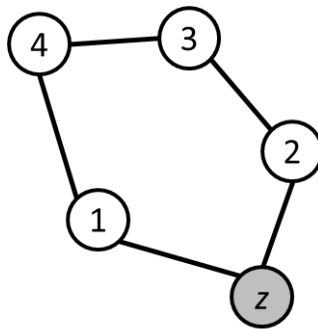

Figure 3. The optimal route with no subtours

3.4.5 Symmetry Constraints

Even with the iterative addition of subtour elimination constraints, the vehicle routing problem is still computationally challenging. To improve the computational efficiency of the model, it was necessary to introduce symmetry cuts. In a vehicle routing problem, it is possible

27

that a vehicle in one iteration may have the exact same route as another vehicle in a different iteration. This is redundant and ultimately increases the solving time and number of cuts that must be implemented to solve the model to optimality. Symmetry cuts are a quick and easy way to avoid the unnecessary run time.

For example, consider a county whose demand is equal to the vehicle capacity, *m*. Without these constraints, the model is still feasible regardless of which vehicle delivers from the stockpile to the given county. Similarly, if the model has found an optimal route, it does not matter which vehicle travels this route, but by nature of the model, it will explore this same route for every vehicle in an attempt to find a better solution. However, optimality will not change. Thus, the following symmetry constraints are introduced to eliminate this redundancy.

$$Y_{v,v} + Y_{v,v-1} + \ldots + X_{v,1} = 1 \qquad\qquad (\textbf{3.17})$$

3.5 Optimization under Uncertainty with Recourse

The final section of this chapter introduces an optimization approach to support stockpile distribution routing under demand uncertainty. After stockpile locations are determined in the first stage of the model, the push distribution approach used in the second stage is much more susceptible to discrepancies in demand because optimal routes are constrained by vehicle capacity. When realized demand is greater than the estimated demand used to create the vehicle routes, it is likely that the vehicle will no longer be able to deliver to every county on the route without exceeding capacity. In this case, there are several possible recourse actions, including: (a) solve a new vehicle routing problem using the model from Section 3.4 and the new demand parameters; (b) use the previously-determined routes, but re-solve part of the routing problem to account for the increased demand; or (c) use the previously-determined routes and a heuristic update rule to account for the increased demand. Each of these options may result in increased distances and may require additional vehicles.

This thesis proposes a scenario-based approach to evaluate the tradeoffs between the three recourse strategies described above. There is little historical data to justify exactly how much emergency equipment a county will need during an influenza pandemic. Instead, scenarios are used to predict county demand at varying stages of pandemic severity. These scenarios are

determined with the introduction of a severity factor, α, which is multiplied by county demand. For example, if it is estimated in stage one that a county has a demand of 1,000 respirators, a severity factor of two (α = 2) would result in a realized demand of 2,000 respirators.

This parameter is incorporated into Equation (3.10), which constrains the amount of equipment in a vehicle to a given capacity. The resulting constraint is shown in Equation (3.18). Here, parameter R_i represents the equipment demand in county i, and Y_{iv} is the binary decision variable used to determine if county i is served by vehicle v.

$$\sum_{i=1}^{|C|} (R_i \, \alpha \, Y_{iv}) \leq u \qquad \forall \, v \in V \qquad\qquad (\textbf{3.18})$$

The following step-by-step process is used to determine the three recourse strategies. Strategy 1 (step 1) gives the optimal solution for the vehicle routing problem with the new parameter values. Heuristic recourse solutions are calculated in two different ways (strategies 2 and 3) to help decision-makers visualize how the routes may be affected if realized demand is greater than estimated demand. These methods are differentiated below in steps 5 and 6.

1. Using computer software, find the optimal vehicle routing plan, denoted here as T_α^*, for all desired values of α. The cumulative distance traveled on each of these routes is denoted as $F(T_\alpha^*)$.
2. For all values of α ≠ 1, begin scheduling vehicles using the optimal routes of α = 1 until vehicle capacity is exceeded.
3. When capacity is exceeded, put the vehicle's remaining counties in a set, M.
4. Re-optimize set M using computer software to determine additional vehicles needed to satisfy demand. This new vehicle routing plan is denoted as T_α^C.
5. Re-optimize set M using a simple sweep heuristic [56] to determine additional vehicles needed to satisfy demand. This new vehicle routing plan is denoted as T_α^H.
6. Calculate two recourse values using $F(T_\alpha^C) - F(T_\alpha^*)$ and $F(T_\alpha^H) - F(T_\alpha^*)$ for all values of α.

By considering recourse with this scenario-based approach, decision-makers can determine vehicle routes based on tradeoffs between the number of vehicles required to implement a plan in a given scenario, the cumulative miles traveled in routing plan, and the number of recourse miles resulting when realized demand is greater than estimated demand. While re-optimizing the remaining set of counties, *M,* using computer software is obviously better to decrease distance traveled, it is only useful if someone is readily available who can effectively manipulate the data in the software. Otherwise, public health officials may implement a sweep heuristic because it is less technical, but still performs well compared to the optimal solution.

Chapter 4 introduces a computational study performed with KDHE, where optimization under uncertainty with recourse is used to evaluate and decide which routing strategy should be implemented in an influenza pandemic. This approach is used to identify tradeoffs between recourse strategies. The metrics considered include robustness, or how well a solution performs under a variety of severity factor levels, and computational requirements.

Chapter 4 - Computational Study

The models described in Chapter 3 were implemented using data from the state of Kansas. This chapter outlines the process for gathering the necessary information and estimated model parameters. Then it discusses trends identified by modifying certain parameters of the model. Finally, it introduces optimization with recourse to address the issue of demand uncertainty and its effect on the model's outcome.

4.1 Motivating Scenario

In early 2013, the Medical Countermeasure Program of KDHE wanted to reevaluate the way in which they stockpile and distribute emergency respirators throughout the state of Kansas. The question that sparked this research was whether there should be a stockpile in every health region or a central materiel cache containing enough respirators to satisfy the demand of every county in the state. There are pros and cons to both decisions, but this prompted additional opportunities that KDHE was not originally considering.

KDHE currently divides the state into seven health regions to allow for more efficient decision making, dissemination of information, and patient care. Each of these regions has a candidate stockpile location, represented by a black triangle in Figure 4. In general, the counties in which the stockpiles are located cannot be changed due to preexisting agreements. For this reason, the model does not look at where the stockpiles should be located, but rather which stockpiles should be opened based on these given locations.
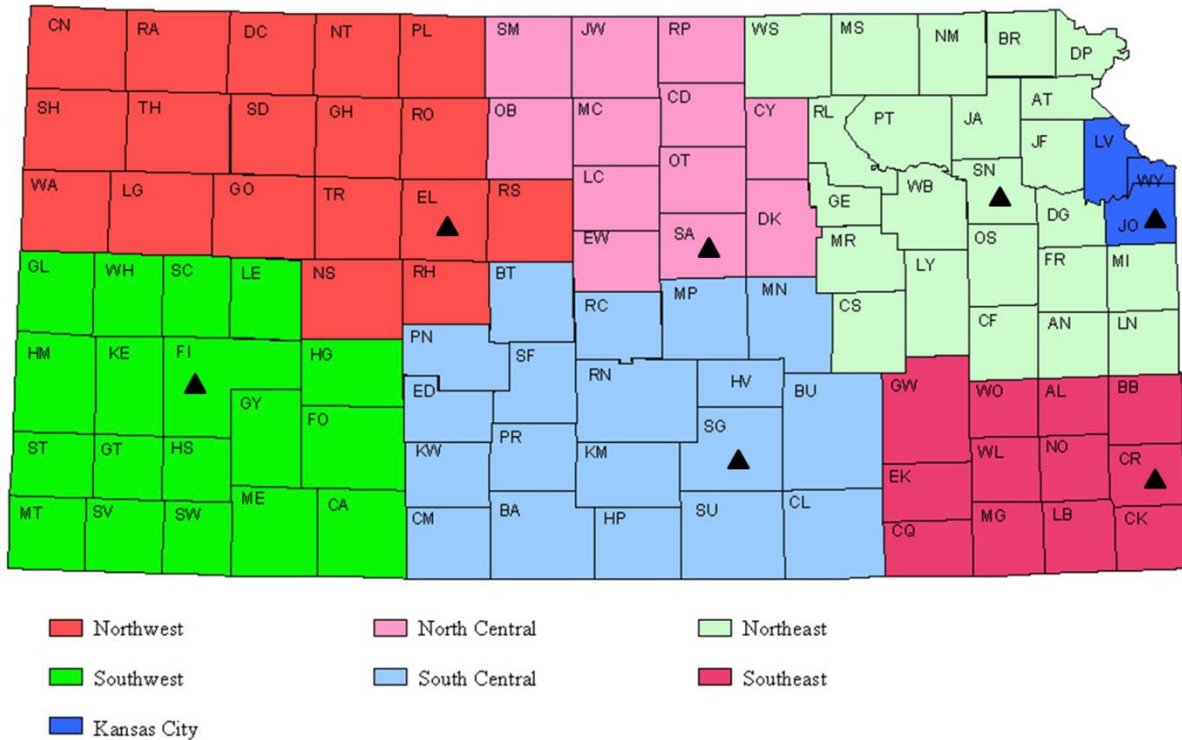
Figure 4. The seven health regions of Kansas and their stockpile locations [29]

Something important to note at this point is that this model generates proposals of various opportunities and their respective outcomes. Regardless of the costs and benefits, counties retain the right to choose whether or not to participate in the stockpile and distribution strategy. The results presented here are obtained assuming all 105 counties will choose to participate. If a county opts to receive medical equipment in another way, the results, especially in the vehicle routing problem, may change.

4.2 Type of Demand

KDHE posits that there are three primary service providers during an influenza pandemic, and each of these providers is likely exposed to the virus at varying levels. This information greatly influences the equipment demand at each county.

The scope of service providers considered in this research is emergency medical services (EMS) attendants, hospital personnel such as doctors and nurses, and public health employees. EMS attendants and hospital personnel have a high exposure risk because they are often in direct

contact with known or suspected pandemic patients. These employees, by OSHA standards, are required to wear approved N95 respirators while caring for these patients. In its most basic form, an N95 respirator is a filtering facepiece that reduces the user's exposure to small inhalable particles. It is meant to form a tight seal to the face, and it generally requires a "fit test" to determine what size will provide adequate protection to the user [15]. Despite its inconvenience, it is important to note that this opportunity cost was not taken into consideration in the final model.

On the other hand, public health employees have only a medium exposure risk because their jobs require close contact exposures to known or suspected sources, such as the general public and outpatients, although they may not necessarily come into direct contact with affected patients. OSHA standards only require employees at this level to wear FDA-cleared facemasks rather than N95 respirators. Facemasks are less protective than respirators, and generally only protect against exposure to splashes of large droplets rather than inhalable particles. However, they are more convenient in that they are cheaper and do not require a fit test.

An important element that was discussed when drafting the model was the type of disposal cost that would need to be taken into consideration. Both facemasks and N95 respirators are designed to be disposed once worn in the presence of an infectious person. Upon doing so, the equipment is considered potentially contaminated and contact should be avoided. However, our research did not reveal that there were any special procedures or additional equipment needed in the disposal process. This information is necessary because it allows the model to assume that respirators and facemasks can be disposed on-site rather than having to be removed by KDHE. There are opportunities for employers to provide reusable respirators that can be cleaned and repaired when necessary, but they are significantly more costly; thus we assume that these products will not be part of the stockpile.

4.3 Demand Estimation

For a deterministic model, one of the most critical components is accurate data. However, in the public health environment, it is nearly impossible to forecast demand with certainty, especially during a pandemic. In this research, parameters were derived based on data provided by partners and on surveys administered to Kansas hospitals and local health departments.

Additional sensitivity analysis on population and demand is performed in the implementation of optimization with recourse.

4.3.1 EMS Data Set

The Kansas Board of Emergency Medical Services (BEMS) provided a data set containing the exact number of EMS attendants in each county. For the most part, every region contained about 600 attendants, with the exceptions being those three regions that contained Sedgwick, Shawnee, and Johnson counties. Respectively, those regions contained about 3100, 2400, and 2000 attendants, reflecting the higher population densities of the state's largest cities.

4.3.2 Hospital Survey

Hospital personnel, such as physicians, nurses, and respiratory specialists, are primary responders in a pandemic. They have a high exposure to infected patients, so it is imperative they have an adequate supply of respirators.

Working with the Preparedness Project Director at the Kansas Hospital Education and Research Foundation, a brief survey was administered to 128 hospitals throughout the state of Kansas. The survey, which can be found in Appendix A, seeks to gain a better understanding of the number of pandemic response personnel at hospitals, the daily respirator usage by hospital personnel during a pandemic, and the frequency with which wearers change respirators. Of the 128 hospitals that received the survey, 36 (28%) responded from 32 of the 105 counties in Kansas. These counties represent both rural and urban communities where populations vary greatly, which helps justify the conclusions drawn from survey responses.

Based on the information received, it was concluded that there is approximately one hospital employee per thousand people in the county in which the hospital is located. While daily usage and the frequency with which wearers change respirators varied, an OSHA-drafted document was later received that allows the model to generalize this parameter. The relevant information from this document is introduced in Section 4.3.4.

4.3.3 KALHD Survey

The Kansas Association of Local Health Departments (KALHD) is a nonprofit organization that seeks to improve the health and protection of Kansas residents by strengthening local health departments. Its members are from 99 of the 100 health departments in Kansas [28].

The support of KALHD is important to this research because its members, many of whom are the directors or administrators of local health departments, are valuable constituents of KDHE. Implementation of any project, especially one that changes the way in which they operate and function as a health department, requires their buy-in. For this reason, it was beneficial to maintain a level of transparency and clear communication with KALHD throughout this research.

Working with contacts at the state level, a short survey was distributed to KALHD members to help determine how many public health employees were in each county. This survey, which can be found in Appendix A, was used to estimate the number of public health employees in each county. In total, responses were received from 14 of the 99 public health departments (14%). However, these counties represented 38.2% of the total population of Kansas, implying that many of the responses were from more densely populated counties. Those counties that did respond were from six of the seven Kansas regions (all except for the southeast) and comprised both urban and rural communities where population varies greatly. Furthermore, the responses were relatively consistent. Similar to hospital personnel, it was concluded that there is approximately one public health employee per thousand people in the county.

4.3.4 Census Information

The Population Division of the United States Census Bureau performs a national census every 10 years. This data, which was most recently collected in 2010, performs regression analysis to predict population numbers of future years. In this research, we use the predicted population estimations for 2012. Once this information is collected, it is broken down by county; this information was used to ultimately predict respirator and facemask demand in the model.

These demand numbers are used to predict how many service providers will require equipment in an influenza pandemic. The proposed OSHA guidance [15] indicates that, based on their level of exposure to infectious patients, employees will be affected according to the predictions in Table 1 below.

Table 1. Equipment requirement for each service provider [15]

Occupational Setting	Proportion of At-risk Employees	N95 Respirators per Employee per Shift	Facemasks per Employee per Shift
EMS Attendants	100%	8	0
Hospital Personnel	33%	4	0
Public Health Employees	90%	0	2

In order to determine exactly how many employees served in each occupational setting, survey were distributed to the appropriate health department, as outlined in the previous sections. The results of these surveys were combined with the information gathered in the census to derive the number of employees each county had in the various occupational settings. Finally, this data was used to calculate demand information by incorporating the length of the pandemic and the number of facepieces used per day, as outlined in Table 1.

Additionally, using the information gathered from the most recent census, the Census Bureau found the population-weighted centroid of each county in the state of Kansas. Although there may not necessarily be a health facility at this exact location, these coordinates are valid in that they provide an optimistic location where emergency equipment is truly available to the greatest number of citizens. These longitude and latitude coordinates were used in the Google Distance Matrix API to create the county-to-county distance matrix used in the model. A critical element of this service is that it calculates the shortest driving distance, which is significantly more valuable and more realistic than straight-line distances between county population centroids.

4.4 Implementation

This section outlines the steps taken and state-specific assumptions made to implement the deterministic model. Implementation assumptions were made based on conversations with contacts at KDHE.

4.4.1 Stockpile Location-Allocation

The purpose of this stage of the model is to decide which of the seven candidate stockpile locations should be open and which counties should be assigned to each open stockpile to minimize cost. Assumptions about stockpile operations are made based on traditional location-allocation applications and conversations with partners at KDHE. First, it is necessary to clarify that any combination of the seven stockpiles may be open; however, if a stockpile is open, it must provide supplies to the county in which it is located. Second, because most of the stockpile locations are existing hospitals, it is not necessary to account for a fixed cost of opening a stockpile. Third, each county is assigned to exactly one stockpile location. For stockpile operations, equipment will be stored in 8' vertical space, and capacity is assumed to be unlimited at each location.

Finally, this model assumes a pull distribution strategy in which each county will send someone to its designated stockpile to pick up the supplies. This implies that at least two trips must be made between each county and the stockpile. If the county demand exceeds vehicle capacity, more trips will have to be made to and from the stockpile until demand is met. In this study, we assume supplies will be delivered using a 6' x 12' cargo trailer with a total capacity of 396 ft^3.

Because the candidate stockpiles are located in counties throughout the state, it is assumed that each has a unique storage cost per square foot. The Saline County Economic Development Office provided an estimate of hospital storage costs in Saline County, and the others are estimated relative to this based on their locations in Kansas. For example, counties in western Kansas are generally rural and less populated, while counties to the east are urban and more densely populated. These estimates, found in Table 2, greatly influenced stockpile location-allocation decisions.

Table 2. Storage cost per square foot for each candidate stockpile location

Stockpile County	Storage Cost per Square Ft.
Crawford	$90
Ellis	$95
Finney	$85
Johnson	$120
Saline	$100
Sedgwick	$110
Shawnee	$105

The decisions in this stage of the model are also influenced by distribution cost per mile. This model examines multiple values to demonstrate the tradeoff between this cost parameter and the number of open stockpiles. This model was used to determine the mileage cost breakpoints, or the costs at which it is more cost-effective to open another stockpile. The results are discussed further in Section 4.5.

4.4.2 Vehicle Routing

The vehicle routing model of Section 3.4 is used to find optimal routes for distributing a fraction of the demand of each county using a push delivery strategy. On average, an influenza pandemic will last about 120 days [15], so to preserve storage space and cost, it may not be necessary to supply the county's entire demand in a single delivery. Instead, it may be advantageous for KDHE to implement a schedule that specifies what percentage, *p*, of emergency equipment is supplied in a one delivery. Then, based on the rate at which counties use the equipment, KDHE can determine when the succeeding deliveries occur so that all counties have at least enough equipment to meet the needs of health care employees. In this model, we assume that each delivery will provide enough equipment to satisfy 25%, or 30 days' worth, of the counties' demands. This number is generally small enough that facilities will not be overwhelmed with the inventory, yet large enough to maintain adequate care of the population while awaiting another delivery from the stockpile facility.

This stage of the model utilizes a push distribution strategy where each of a given number of vehicles must deliver to all of the counties on its route without exceeding vehicle capacity.

The goal of this model is to minimize the cumulative distance traveled by the vehicles. To provide a more implementable distribution plan for KDHE, the model operates using the facility location-allocation solution from stage one that resulted in five stockpile locations. In cases where a county's demand exceeds vehicle capacity, the model is not solvable because the capacity constraint is violated. To account for this, it is assumed that vehicles will deliver equipment to the county at capacity until its remaining demand is less than vehicle capacity. The remaining demand is then used to solve the vehicle routing problem.

For stockpiles to which there is a large number of assigned counties, finding an optimal vehicle routing plan is computationally intensive. In these situations, a heuristic is introduced to determine the optimal routes. This heuristic solves the traveling salesperson problem (TSP) to find the optimal route that an uncapacitated vehicle can deliver to each facility. Then vehicles are assigned to counties in the order of the TSP solution until capacity is exceeded. It is important to note that this heuristic still maintains the constraint that counties may only be served by exactly one vehicle and split deliveries are not feasible. Therefore, if the next county on the TSP route cannot be served by a vehicle without exceeding capacity, the vehicle returns to the stockpile county.

4.4.3 Optimization under Uncertainty with Recourse

For influenza pandemic planning, it is important to explore the impact on stockpile planning decisions under a range of demands. This research uses scenarios based on severity factor values of 1.5, 2, 2.5, and 3. These values were chosen to examine the impact of a wide range of possible demands but to limit the number of individual cases that were considered. An inherent assumption in the choice of scenarios is that a county's realized demand would never be more than three times that of its estimated demand.

This research compares three types of recourse, as described in Section 3.5, and analyzes the tradeoffs associated with each of them. The results of these recourse models are outlined in the following section.

4.5 Results

The results of this two-stage model can be used to influence stockpile location-allocation and vehicle routing decisions made by public health officials, but they can also be used as a decision support tool to guide conversations regarding public policy in the state of Kansas.

Table 3 displays the results of the stockpile location-allocation model and the associated costs. The first column represents the mileage cost breakpoints, and the second column shows the stockpiles that should be open to minimize total cost at the respective cost per mile. The third column represents the total storage and distribution costs incurred from the given scenario, and the last column shows what percentage of the total cost is due to storage.

Table 3. Quantitative results of the stockpile location-allocation model

Cost/Mile	Open Stockpiles	Total Cost	Storage Cost as % of Total Cost
< $ 0.32	Finney	$ 6.88 million	96%
$ 0.32	Finney, Crawford	$ 6.89 million	96%
$ 1.67	Finney, Crawford, Ellis	$ 7.65 million	89%
$ 1.77	Finney, Crawford, Ellis, Saline	$ 7.70 million	89%
$ 2.11	Finney, Crawford, Ellis, Saline, Shawnee	$ 7.85 million	90%

Many conclusions can be drawn from these results, and the information can also be used to influence decisions made by KDHE. For the tested range of storage and distribution cost parameters, and given that transportation reimbursement in Kansas is $.56 per mile [17], multiple stockpile locations are preferred over a single location. It is also clear that storage cost accounts for a much greater percentage of total cost than do distribution cost; and, as a result, stockpiles are opened in order of increasing storage cost per square foot.

From the perspective of public safety, it is generally desired by KDHE to have more stockpiles open to decrease the time needed to deliver respirators in the case of an influenza pandemic. For this reason, stage two of the model was performed with the assumption that five stockpiles are open. These the five stockpiles and their assigned counties can be seen in Figure 5 below.
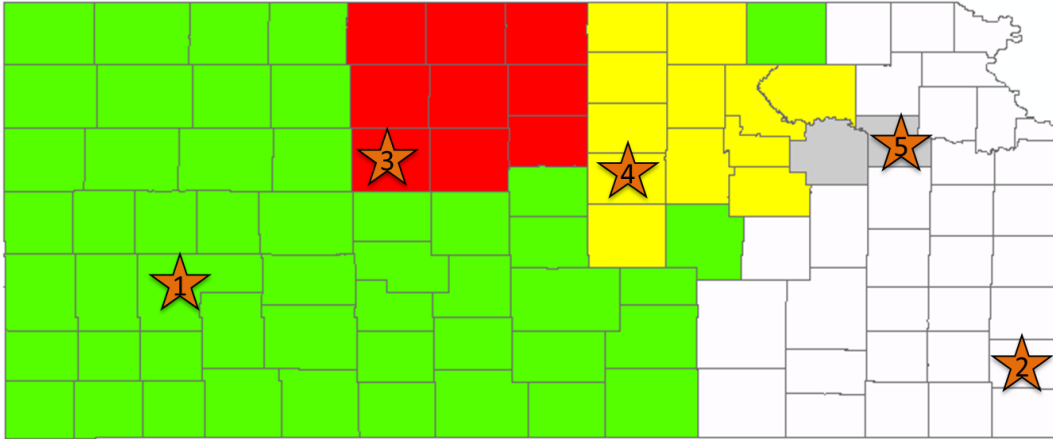
Figure 5. Stockpile location-allocation decisions made in stage one

(1: Finney, 2: Crawford, 3: Ellis, 4: Saline, 5: Shawnee)

Using Saline County (stockpile 4) as an example, Figure 6 below shows the optimal distribution routes that two vehicles should take to minimize cumulative distance traveled. For the sake of simplicity, the figure displays straight-line distances, but the distances actually used in the model are those of county-to-county road networks.
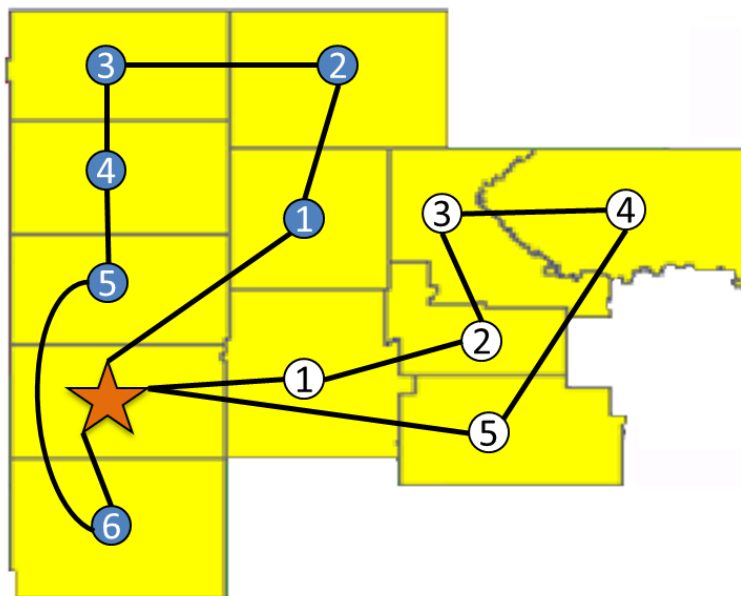

Figure 6. Optimal vehicle routing plan for a stockpile in Saline County

Optimization with recourse was incorporated into these results to demonstrate how demand uncertainty influenced vehicle routes. Figure 7 shows the results of the second type of recourse, where computer software is used to re-optimize distribution routes to those counties whose demand will no longer fit on the vehicles using the optimal routes when $\alpha = 1$. The results from the third recourse strategy are shown in Figure 8. This method uses a simple sweep heuristic to re-optimize routes rather than computer software.

It can be observed that some routes have identical solutions. In the Saline County example below, the re-optimized routing plan found using computer software at $\alpha = 2$ is identical to both re-optimized solutions at $\alpha = 2.5$. These results are merely a coincidence and cannot necessarily infer any type of relationship between the two severity factors, but they do indicate a degree of robustness in the routing decisions and show decision-makers how one route can be implemented to feasibly satisfy two different levels of realized demand.
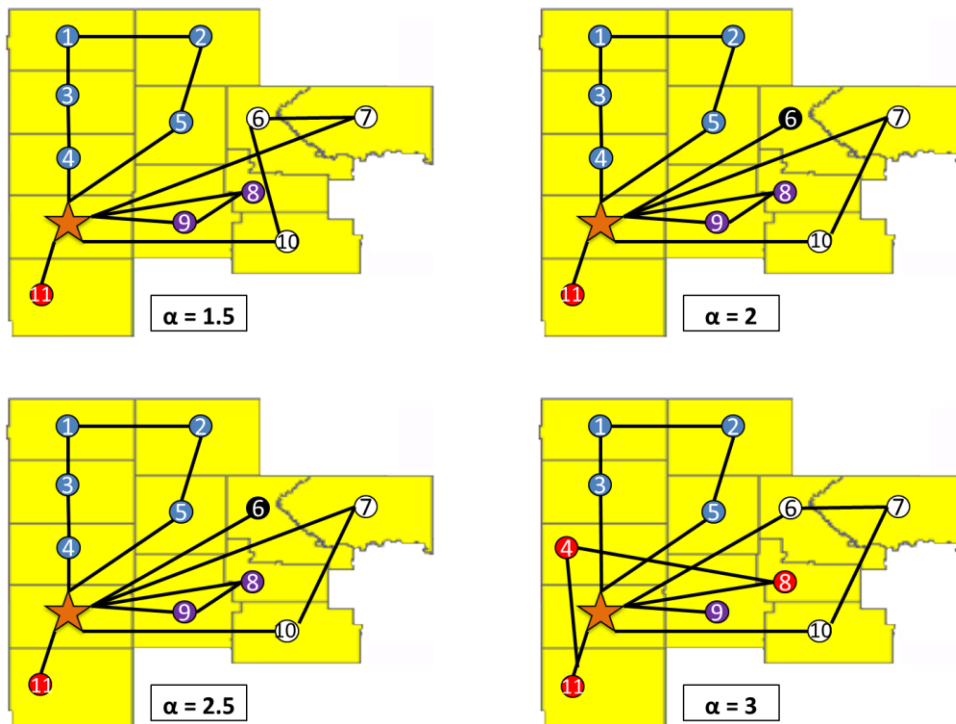

Figure 7. Vehicle routes when counties are re-optimized using computer software
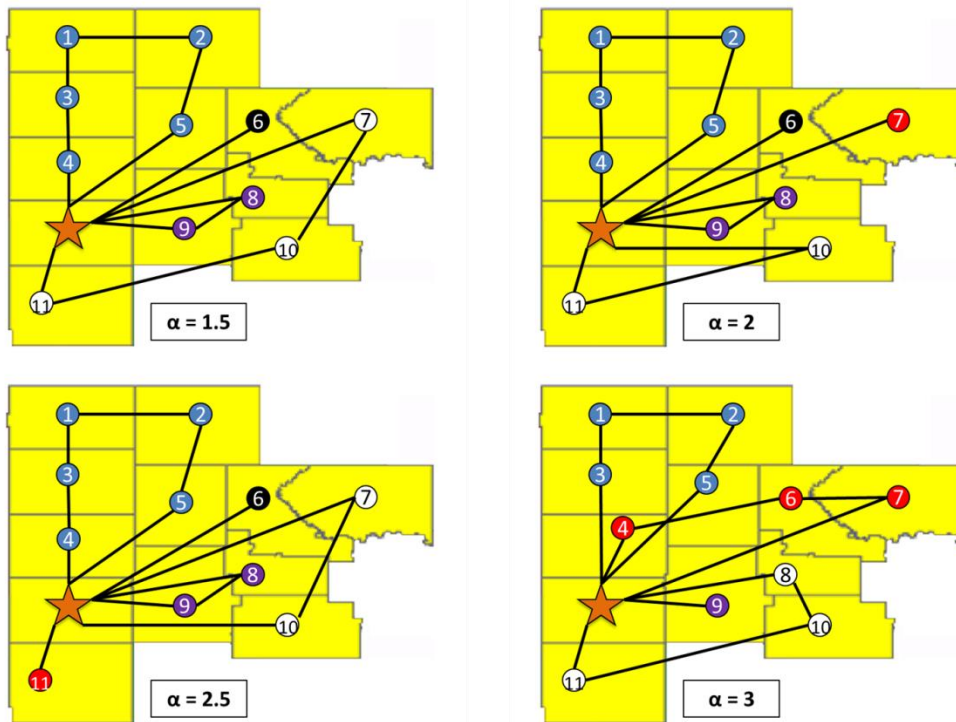
Figure 8. Vehicle routes when counties are re-optimized using a simple sweep heuristic

Table 4 outlines the quantitative results of these solutions. There are many tradeoffs that must be considered when decision-makers analyze these results. In addition to the mileage increase associated with new solutions, it may be of significant importance to also consider the number of vehicles required to fulfill the routing plan. For example, if the cost to obtain a vehicle is sufficiently high, it may be more desirable to employ a routing solution that uses fewer vehicles, and vice versa.

Table 4. Quantitative results of optimization with recourse for a stockpile in Saline County

Severity Factor, α	Optimal		Recourse w/ Computer Software		Recourse w/ Sweep Heuristic	
	# vehicles	Distance	# vehicles	Distance	# vehicles	Distance
1	2	510.3	--	--	--	--
1.5	3	580	4	616.8	4	717.4
2	4	672.3	5	740.7	5	810.9
2.5	5	737.9	5	740.7	5	740.7
3	6	888	6	888	6	893

From these results, the following observations can be made:

- The maximum number of additional vehicles required in comparison to complete re-optimization at each value of α is one, regardless of whether the routes are re-optimized with computer software or a simple sweep heuristic.

- The additional distance required by recourse with computer software re-optimization, in comparison to the optimal solution, is generally small (at most, 10.2% at α = 2).

- The additional distance required by recourse with a sweep heuristic can be significant, in comparison both to complete re-optimization and recourse with computer software re-optimization.

- The greatest additional distance required by recourse, in comparison to the optimal solution, is 23.7% using the sweep heuristic at α = 1.5.

4.6 Recommendations

This section analyzes the results of the model and outlines a potential solution. However, based on the results in Section 4.5, it is not trivial which vehicle routing plan should be adopted. Rather than serving as a decision-making tool, this computational study is intended to guide discussions that planners at KDHE can use to devise and implement the best emergency preparedness plan.

Assuming the stockpile location-allocation solution at Saline County found in stage one of the model, it is recommended that decision-makers choose routes based on the optimal solution when α = 2. This was chosen for two main reasons. First, the optimal routing plan

requires four vehicles, which is midway between the fewest and greatest number of vehicles to distribute supplies in any scenario. Second, both recourse re-optimization plans require the use of a fifth vehicle. However, if the severity factor is ever greater than two, the route will require at least five vehicles. Therefore, it is more beneficial to plan for the route that requires four vehicles in hopes that the severity factor is two or less.

Planners at KDHE must also look at the negative implications of such a plan, though. If this solution is chosen but demand is realized to have a severity factor of three, the worst case scenario using recourse with a sweep heuristic has a distance of 893 miles. This is a difference of 220.7 miles and two vehicles.

These recommendations are relevant only to the example used in this computational study, in which the stockpile at Saline County is assigned to 11 nearby counties. The results are summarized and used to draw more general recommendations in Chapter 5.

Chapter 5 - Conclusions and Future Work

Since 1900, only four influenza pandemics have occurred, but they have resulted in the deaths of nearly 50 million people worldwide. As an airborne virus, influenza is easily prevented if individuals have access to the proper protective equipment, such as facemasks and N95 respirators. This thesis quantifies the costs associated with stockpiling, allocating, and distributing emergency respirators. It demonstrates how operations research techniques can be used to support the development of emergency preparedness strategies and policies in the public health sector. However, in situations regarding public health and safety, it is important for decision-makers to consider not only the costs associated with these decisions, but also their effects on the population being served.

Many applications of operations research in emergency preparedness efforts fail to consider both stockpile and distribution strategies or to account for parameter uncertainty. This research introduces a two-stage optimization approach that quantifies tradeoffs in stockpile storage and distribution costs and accounts for different distribution strategies. Additionally, the second stage of this model uses optimization under uncertainty with recourse to account for situations in which realized county demand is much greater than was estimated.

Working with KDHE, a computational study was conducted using the methods introduced in the thesis. The results yield several informative insights into the tradeoffs and costs of various facility location-allocation and vehicle routing problems. The following recommendations are supported by the results of the computational study:

1. Under current operational assumptions, it is beneficial to open multiple stockpiles rather than a single, centralized materiel cache.
2. Storage costs play an important role in location-allocation decisions, and thus should be carefully estimated.
3. In the push distribution system, finding an optimal vehicle routing plan is computationally intensive for stockpiles with a large number of assigned counties. Thus, decision-makers should consider using a modified TSP heuristic to find an efficient vehicle routing plan.
4. The combination of using a pull distribution strategy to make stockpile location-allocation decisions and a push distribution strategy to make vehicle routing decisions

offers flexibility and computational advantages while still examining stockpile management holistically, and thus should be considered an important planning tool.

5. If it is desirable to specify routes well in advance, planners should adopt a robust routing plan based on higher-than-expected demand levels.

6. These results should be used to inform discussions about stockpile location, allocation, and distribution strategies, but do not necessarily reflect the optimal solutions for every population.

This thesis makes important advances in the modeling and analysis of public health preparedness stockpile planning. Furthermore, it opens the door for future work in several related areas.

First, future research could 7explore the potential benefits of integrating location, allocation, and routing decisions in a single model. By doing so, the model could make more informed stockpile location-allocation decisions based on the resulting optimal vehicle routing plan. It is likely that new heuristics will be needed to solve an integrated problem efficiently so that the model is useful in practice. Second, further research can be done to examine the impact of non-uniform, stochastic consumption rates on optimal location, allocation, and routing decisions. This model assumes emergency equipment is consumed at a uniform rate and does not consider implications of the spread of the virus throughout the population over time. Third, there are opportunities to include logistical parameters in the model based on realistic applications. For example, the model could incorporate economies of scale to see how facility location-allocation decisions change if storage is discounted based on the amount of equipment held in inventory at a given stockpile. Finally, it would be worthwhile to incorporate a stochastic element into this model. Based on the characteristics of influenza pandemics, it is reasonable to conclude that they appear and spread more rapidly in densely populated areas. Probability distributions could be introduced to represent the likelihood that a pandemic will affect a given county, which preparedness planners could use in stochastic optimization models to make more strategic facility location decisions.

References

[1]   3M 1870 Surgical Mask N95 3M 1870 Surgical Mask N95, Box of 20. *Amazon.com*. Web. <http://www.amazon.com/3M-1870-Surgical-Mask-N95/dp/B000BZ1OFO>.

[2]   3M Particulate Respirator 8612F Mask N95 FDA Approved, 2-Count. *Amazon.com*. Web. <http://www.amazon.com/Particulate-Respirator-8612F-Approved-2-Count/dp/B001FA0XMO>.

[3]   Balinski, M. L. and Quandt, R. E., "On an Integer Program for a Delivery Problem," *Operations Research* 12. (1964): 300-304.

[4]   Berman, O. Wang, J. Drezner, Z. and Wesolowsky, G. O., "The Minimax and Maximin Location Problems on a Network with Uniform Distributed Weights," *11E Transactions* 35(11) (2003): 1017-1025.

[5]   Campbell, Ann Melissa, and Philip C. Jones. "Prepositioning Supplies in Preparation for Disasters," *European Journal of Operational Research* 209.2 (2011): 156-65. Web.

[6]   Carbone, R. "Public Facilities Location Under Stochastic Demands," *INFOR* 12(3) (1974): 261-270.

[7]   Carrizosa, E. and Nickel, S. "Robust Facility Location," *Mathematical Methods in Operations Research.* 58 (2003): 331-349.

[8]   Carrasco, L. R., V. J. Lee, M. I. Chen, D. B. Matchar, J. P. Thompson, and A. R. Cook. "Strategies for Antiviral Stockpiling for Future Influenza Pandemics: A Global Epidemic-economic Perspective," *Journal of The Royal Society Interface* 8.62 (2011): 1307-313.

[9]   Chen, Pey-Chun, Pierre Hansen, Brigitte Jaumard, and Hoang Tuy. "Weber's Problem with Attraction and Repulsion," *Journal of Regional Science* 32.4 (1992): 467-86. Web.

[10] Cooper, L. "Bounds on the Weber Problem Solution Under Conditions of Uncertainty," *Journal of Regional Science.* 18(1) (1978): 87-93.

[11] Créput, Jean-Charles, Hajjam, Amir, Abderrafiaa Koukam, and Oliver Kuhn. "Dynamic Vehicle Routing Problem for Medical Emergency Management," *Self Organizing Maps - Applications and Novel Algorithm Design*. Rijeka, Croatia: InTech, 2011. 234-50. Print.

[12] Daskin, M. S., Hopp W. J. and Medina, B. "Forecast Horizons and Dynamic Facility Location Planning," *Annals of Operations Research.* 40 (1992): 125-151.

[13] Dantzig, G. B., and J. H. Ramser. "The Truck Dispatching Problem," *Management Science* 6(1) (1959): 80-91.

< http://andresjaquep.files.wordpress.com/2008/10/2627477-clasico-dantzig.pdf>

[14] Drake, Tom, Zaid Chalabi, and Richard Coker. "Buy Now, Saved Later? The Critical Impact of Time-to-pandemic Uncertainty on Pandemic Cost-effectiveness Analyses," *U.S. National Library of Medicine*. National Institutes of Health, 14 Jan. 2014. Web. <http://www.ncbi.nlm.nih.gov/pubmed/24369125>.

[15] "Proposed Guidance on Workplace Stockpiling of Respirators and Facemasks for Pandemic Influenza." OSHA, 2008. Unpublished document. <https://www.osha.gov/dsg/guidance/proposedGuidanceStockpilingRespirator .pdf>.

[16] Drezner, Z. "Stochastic Analysis of the Weber Problem on the Sphere," *Journal of the Operational Research Society.* 40(12) (1989): 1137-1144.

[17] "Employee Travel Expense Reimbursement Handbook," Office of Financial Management. Kansas Department of Administration, 1 July 2013. Web. June 2014. <http://www.da.ks.gov/ar/employee/travel/travbkSMART14.pdf>.

[18] Fisman, David. "Modelling an Influenza Pandemic: A Guide for the Perplexed," *Canadian Medical Association Journal* 181.3-4 (2009): 171-73. Web.

[19] Frank, H. "Optimum Locations on a Graph with Probabilistic Demands," *Operations Research* 14(3) (1966): 409-421.

[20] Garvin, W., Grandall, H., John, J. and Spellman, R., "Applications of Linear Programming in the Oil Industry," *Management Science 3*. (1957): 407-430.

[21] Gellin, Bruce G. "Testimony," *Pandemic Influenza Preparedness*. U.S. Department of Health and Human Services, 30 June 2005. Web.

[22] Golden, B. L., Magnati, T.L. and Nguyen H.Q. "Implementing Vehicle Routing Algorithm," *Networks 7*. (1977): 113-148.

[23] Gupta, S. K. and Rosenhead, J. "Robustness in Sequential Investment Decisions," *Management Science.* 15(2) (1968): B18-B29.

[24] "H1N1 Flu," *CDC*. Centers for Disease Control and Prevention, 16 June 2010. Web. <http://www.cdc.gov/h1n1flu/cdcresponse.htm>.

[25] Hakimi S. L. "Optimal Location of Switching Centers and the Absolute Centers and Medians of a Graph," *Operations Research* 12 (1964): 450-459. Web.

[26] Hodgson, M. J. "Stability of Solutions to the *p*-median Problem Under Induced Data Error," *INFOR* 29(2), (1991): 167-183.

[27] "Hospital Bioterrorism Regions," *Kansas Department of Health and Environment*. Web. Mar. 2013. <http://kic.kdhe.state.ks.us/kic/OHA/ksmap.html#hosp>.

[28] *Kansas Association of Local Health Departments*. 2012. Web. Oct. 2013. <www.kalhd.org>.

[29] "Kansas Pandemic Influenza Preparedness and Response Plan," Version 2.2.3. Kansas Department of Health and Environment, Jan. 2014. Web. <http://www.kdheks.gov/cphp/download/KS_PF_Plan.pdf>.

[30] Kouvelis, P. Kurawarwa`la, A. A. and Gutiérrez, G. J. "Algorithms for Robust Single and Multiple Period Layout Planning for Manufacturing Systems." *European Journal of Operational Research.* 63 (1992): 287-303.

[31] Kulkarni, R.V., and P.R. Bhave. "Integer Programming Formulations of Vehicle Routing Problems," *European Journal of Operational Research* 20.1 (1985): 58-67.

[32] Labbé, M., Thisse, J.-F and Wendell, R. E. "Sensitivity Analysis in Minisum Facility Locations Problems," *Operations Research* 39(6) (1991): 961-969.

[33] Louveaux, F. V. "Discrete Stochastic Location Models," *Annals of Operations Research 6* (1986): 23-34.

[34] MacKellar, Landis. "Pandemic Influenza: A Review," *Population and Development Review* 33.3 (2007): 429-51. Web.

[35] Murali, Pavankumar, Fernando Ordóñez, and Maged M. Dessouky. "Facility Location Under Demand Uncertainty: Response to a Large-scale Bio-terror Attack," *Socio-Economic Planning Sciences* 46.1 (2012): 78-87. Web.

[36] Nordqvist, Christian. "What Is a Pandemic? What Is an Epidemic?" *Medical News Today*. MediLexicon International, 5 May 2009. Web. <http://www.medicalnewstoday.com/articles/148945.php>.

[37] "Pandemic Flu History," *Flu.gov*. U.S. Department of Health and Human Services. Web. Nov. 2013. <http://www.flu.gov/pandemic/history/>.

[38] "Pandemic Influenza," *Emergency Management*. Douglas County Sheriff, Aug. 2012. Web. < http://www.dcsheriff.net/emergencymanagement/pandemic-influenza/>.

[39] Pike, John. "1957 Asian Flu Pandemic," *Pandemic Influenza, National and Community Measures*. Global Security, 13 July 2011. Web.
< http://www.globalsecurity.org/security/ops/hsc-scen-3_pandemic-1957.htm>.

[40] Rawles, Carmen G., and Mark A. Turnquist. "Pre-Positioning Planning for Emergency Response with Service Quality Constraints," *OR Spectrum* 33 (2011): 481-98. Web.

[41] Rawls, Carmen G., and Mark A. Turnquist. "Pre-positioning of Emergency Supplies for Disaster Response," *Transportation Research Part B: Methodological* 44.4 (2009): 521-34.

[42] Rosenhead, J., Elton, M. and Gupta, S.K. "Robustness and Optimality as Criteria for Strategic Decisions," *Operational Research Quarterly* 23(4) (1972): 413-413.

[43] Schilling, D. A. "Strategic Facility Planning: The Analysis of Options," *Decision Sciences.* 13 (1982): 1-14.

[44] Siddiqui, Rudy M., and Edmunds, W. John. "Cost-effectiveness of Antiviral Stockpiling and Near-Patient Testing for Potential Influenza Pandemic," *Emerging Infectious Diseases* 14(2) (Feb. 2008): 267-274.

[45] Snyder, Lawrence V., and Daskin, Mark S. "Stochastic-robust Location Problems," *IIE Transactions* 38.11 (2006): 971-85.

[46] Snyder, Lawrence V. "Facility Location under Uncertainty: A Review," *IIE Transactions* 38.7 (2006): 547-64.

[47] Thompson, Alison K., Karen Faith, Jennifer L. Gibson, and Ross E.G. Upshur. "Pandemic Influenza Preparedness: An Ethical Framework to Guide Decision-Making," *BMC Medical Ethics* 7.12 (2006).

[48] "Transmission of Influenza Viruses from Animals to People," *CDC*. Centers for Disease Control and Prevention, 22 Mar. 2012. Web.
<http://www.cdc.gov/flu/about/viruses/transmission.htm>.

[49] United States. United States Census Bureau. U.S. Department of Commerce. *Kansas QuickLinks*. 2010. Web. <http://quickfacts.census.gov/qfd/states/20000lk.html>.

[50] "What Are Epidemics, Pandemics, and Outbreaks?" *Cold, Flu, & Cough Health Center*. WebMD. Web. Nov. 2013. <http://www.webmd.com/cold-and-flu/what-are-epidemics-pandemics-outbreaks>.

[51] Özdamar, Linet, Ediz Ekinci, and Beste Küçükyazici. "Emergency Logistics Planning in Natural Disasters," *Annals of Operations Research* 129.1-4 (2004): 217-45.

[52] Huang, Hsin-Chan, Ozgur Araz, David Morton, Paul Damien, and Lauren Meyers. "Stockpiling Ventilators for Pandemic Influenza. University of Texas at Austin. Working Paper.

[53] Daskin, M.S. "A Maximum Expected Covering Location Model: Formulation, Properties, and Heuristic Solution," *Transportation Science* 17(1) (1983): 48-80

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

[55] Drezner, Z., ed. *Facility Location: A Survey of Applications and Methods,* Springer-Verlag, New York. 1995.

[56] Laporte, G., Gendreau, M., Potvin, J.-Y. and Semet, F. *Classical and Modern Heuristics for the Vehicle Routing Problem.* International Transactions in Operational Research, 7 (2000): 285–300.

Appendix A - Public Health Surveys

Survey Distributed to Kansas Hospitals

Question 1. With which hospital are you affiliated?

Question 2. In which county of Kansas is your facility located?

Question 3. On a given day during an influenza pandemic response, how many personnel (nurses, support, volunteers, etc.) would require an N95 respirator?

Question 4. How many N95 respirators, on average, does each person use per day?

Question 5. Do wearers change masks every time they change patients? If no, please explain further.

Question 6. Please provide any additional information you think would be valuable in estimating the number of N95 respirators needed.

Survey Distributed to KALHD Members

Question 1. On a given day during a pandemic response, how many personnel (nurses, support, etc.) would require an N95 respirator?

Question 2. Approximately how long do responders wear respirators, or how many, on average, does each person use per day?

Question 3. Do wearers change masks every time they change patients?