

Patient Prioritization and Resource Allocation in Mass Casualty Incidents

Alex F. Mills

A dissertation submitted to the faculty of the University of North Carolina at Chapel Hill in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Statistics and Operations Research.

Chapel Hill
2012

Approved by:

Nilay Tanik Argon

Wendell G. Gilland

Kristen Hassmiller Lich

Vidyadhar G. Kulkarni

Serhan Ziya

© 2012
Alex F. Mills
ALL RIGHTS RESERVED

Abstract

**ALEX F. MILLS: Patient Prioritization and Resource Allocation
in Mass Casualty Incidents.**

(Under the direction of Nilay Tanik Argon and Serhan Ziya.)

Mass-casualty incidents, such as multi-car traffic accidents, plane crashes, and terrorist bombings, create a sudden spike in demand for emergency resources in an area. Providers of emergency medical services must act quickly to make decisions that will affect the lives of injured patients. Particularly important is triage, the process of classifying patients and prioritizing them for transportation from the scene of the incident. The most widely used standard for mass-casualty triage, START, prescribes a fixed priority ordering among the different classes of patients, without explicitly accounting for resource limitations. We develop policies to improve the resource allocation phase of START by explicitly incorporating resource limitations. Next, we develop policies for assigning resources when two or more incidents occur at the same time and demand the same set of resources. Current standards, such as START, do not prescribe how to handle such situations—these decisions are most often made in an *ad hoc* manner. Finally, we examine the problem of efficiently routing a large number of patients affected by a major disaster, such as a biological, chemical, or nuclear incident, to facilities where they can be treated. We provide insight on how resources can be used effectively to treat patients as quickly as possible. Throughout this work, we focus on policies that are analytically justified, intuitive, broadly applicable, and easy to implement. Using numerical results and simulation, we demonstrate that implementing policies based on quantitative analysis can make a meaningful impact by increasing the expected number of survivors in a mass-casualty incident.

Acknowledgments

A doctoral dissertation is never the product of just one person, and this dissertation is no exception. I am grateful to the many family members, friends, teachers, and colleagues who have influenced me, supported me, and encouraged me to excel.

Several individuals deserve special acknowledgement for specific help in the completion of this dissertation. My advisers, Nilay Tanık Argon and Serhan Ziya, have always given me excellent advice throughout the time that I have worked with them. I am extremely grateful for their support and I am proud to have been their student. James E. Winslow provided perspective from the medical point of view, which greatly improved this work. Finally, I am grateful for the love and support of all of my family, especially my wife, Alicia M. Frame, without whom I never would have made it through graduate school.

The National Science Foundation (grant number CMMI 0927607) provided financial support for this project, for which I am grateful.

Table of Contents

List of Tables	viii
List of Figures	ix
1 Introduction	1
2 Literature Review	5
2.1 Emergency Medicine Literature	5
2.2 Operations Research Literature	10
3 Patient Prioritization with Resource-Based START (ReSTART)	14
3.1 Model Formulation	15
3.2 A Simple Condition for Fixed-Priority Ordering	18
3.3 Priority Decisions for Two Classes of Patients	21
3.3.1 Characterization of the Optimal Policy	22
3.3.2 Determining the Optimal Threshold	23
3.3.3 Sensitivity of the Optimal Policy to the Number of Patients	25
3.4 A New Policy for Patient Triage: ReSTART	26
3.4.1 QuickDynamic-ReSTART (QD-ReSTART)	27
3.4.2 QuickStatic-ReSTART (QS-ReSTART)	30
3.5 Simulation Study	31
3.5.1 Experimental Setup	31
3.5.2 Comparison of QD-ReSTART and QS-ReSTART with START and InvSTART	37
3.5.3 Sensitivity Analysis on the Reward Functions	40

3.6	Adapting ReSTART to Changing Conditions on the Field	43
3.6.1	Adaptive QD-ReSTART and QS-ReSTART	44
3.6.2	Performance of Adaptive ReSTART policies	45
3.7	Discussion	47
4	Resource Allocation in a Multiple Location Mass-Casualty Incident . .	49
4.1	Model Description	50
4.2	Analysis of Two-Incident Problem	53
4.2.1	Upper and lower bounds for t^*	54
4.2.2	Upper and lower bounds for t^{**}	56
4.2.3	Application of Lower and Upper Bounds	58
4.2.4	Decentralized Resource Allocation	59
4.3	Analysis for More than Two Incidents	61
4.4	Heuristic Policies and Simulation Study	63
4.4.1	Policy Descriptions	64
4.4.2	Simulation Method	65
4.4.3	Simulation Results	67
4.5	Discussion	69
5	Dynamic Routing of Casualties in a Mass-Casualty Incident	70
5.1	Dynamic programming formulation	73
5.2	Analytical Results	77
5.3	Heuristic Policies	81
5.3.1	Marginal value approximation	81
5.3.2	Policy improvement heuristics	82
5.4	Numerical Results	86
5.5	Simulation Study	87
6	Conclusions	91

Appendix	Proofs of Analytical Results	95
Bibliography		122

List of Tables

3.1	Fitted parameters for the five survival scenarios.	34
3.2	Mean values of ϕ over 500 instances for the five survival scenarios, excluding instances for which $t^* = 0$. All 95% half-widths are less than 0.01.	37
3.3	Mean reduction in the critical mortality rate obtained using QD-ReSTART instead of START.	38
3.4	Mean reduction in the critical mortality rate obtained using QD-ReSTART instead of InvSTART.	39
3.5	Mean reduction in the critical mortality rate obtained using QS-ReSTART as opposed to START, InvSTART, or QD-ReSTART (negative numbers indicate an increase in critical mortality).	41
3.6	Mean reduction in the critical mortality rate obtained using QD-ReSTART(0.5) and QS-ReSTART(0.5) instead of START on instances with perturbed reward function parameters.	42
3.7	Mean and 95% half-width for reduction in the critical mortality rate obtained using adaptive ReSTART policies instead of START for different simulated scenarios.	48
4.1	Summary of heuristic policies for two-location and multi-location MCI response.	66
4.2	Mean decrease in critical mortality over 200 instances compared to Decentralized-Simul-S with fluid optimal resource allocation.	68
5.1	Heuristic results as a percentage of optimal value. Minimum, quartiles, maximum, and mean are calculated across 10,000 randomly generated instances.	85
5.2	Simulation results for $m = 3, n = \{3, 4, 5\}, p = 50$. (*) indicates significance at the 0.05 level.	90

List of Figures

3.1	Example of reward functions satisfying Assumption 3.1.	22
3.2	Visualizations of QD-ReSTART(0.5) and QS-ReSTART(0.5).	30
3.3	Survival probability (reward) functions, $g(t)$, and START class dis- tributions for the five scenarios.	35
4.1	Histograms for ψ^* (left) and ψ^{**} (right) from the numerical experiment. . .	64
5.1	Evacuation and decontamination of casualties in the aftermath of a mass-casualty event involving the release of a hazardous substance.	73
5.2	Example problem instance with $\mathcal{R} = \{1, 2\}$ and $\mathcal{S} = \{A, B\}$; num- bers on arcs indicate mean travel times $1/\tau_{ij}$ for $i \in \mathcal{R}, j \in \mathcal{S}$	78

Chapter 1

Introduction

In a mass-casualty incident (MCI), the number of patients is large enough to overwhelm the emergency response resources in an area, so even critically injured patients may have to wait to be transported to a hospital. Managing the emergency response in an MCI requires making many decisions. Depending on the scale of the incident and the severity of the patients' injuries, these decisions may include how many resources are needed to respond to the incident or incidents, how to classify patients, which patients should get priority for transportation to a hospital, and to which facility each patient should be sent. Because of the time-sensitive nature of emergency medicine and the chaotic environment present at the scene of an MCI, these decisions must be made quickly, and often with limited information. This dissertation analyzes some of the important prioritization and resource allocation decisions that must be made when responding to an MCI. We examine several of these decisions with the goals of developing policies that are practical and effective, and providing insights that can guide decision makers in an emergency. We consider improvements to patient prioritization that can be made within the framework of existing triage standards, we examine how to allocate resources in scenarios not considered by current protocols, and we attempt to quantify some of the observations that have been made by experts in the medical field.

One of the most important concepts involved in the response to a mass-casualty incident is triage. Triage is the process of classifying patients according to their medical conditions and injury characteristics and then determining the order in which they will be treated. According to the current practice, priority assignments in an MCI are made in a very

simple way. The triage class of a patient automatically determines the patient’s priority. For example, START, which stands for Simple Triage and Rapid Transport, is the most widely adopted mass-casualty triage protocol in the U.S. (Lerner 2008). As a triage tool, START is primarily geared toward the quick assessment and classification of patients. According to START, patients are classified into five different classes. *Minor* patients are those who are capable of walking away from the scene; *delayed* patients are those for whom treatment may be delayed by some time without risking their lives; *immediate* patients are those whose conditions will deteriorate most rapidly without care; *expectant* patients are those who are expected to die no matter what care is given; and finally *dead* patients. Expectant and dead are often considered to be the same class; we treat them as such in this work. Once the patients are classified, priorities are determined as follows: resources are allocated first to patients in the immediate class, then to those in the delayed class. Once the system is cleared of patients in these time-critical classes, resources may be used for those in the minor class. Beyond the simple static priority policy, START does not address any other aspects of the emergency response effort.

There is a clear benefit from adopting a static, predetermined prioritization policy that depends solely on patients’ triage classes: it is simple and thus easy to implement. However, there are also drawbacks to this type of simple policy. A number of recent articles in the emergency response literature have questioned whether START is in fact too simple, mainly because it completely ignores resource limitations (Garner et al. 2001, Sacco et al. 2005, Frykberg 2005, Jenkins et al. 2008). The main argument is that patient priorities should depend on the extent of resource limitation relative to demand. It might not always be sensible to give priority to patients in the immediate class over those in the delayed class. For example, there may be so many immediate patients that there is not enough time to serve the delayed patients within a time frame that will give them a good chance of survival. In Chapter 3, we study the problem of prioritizing patients during the response to an MCI, and we propose a policy, ReSTART, that builds on the START classification system by prioritizing patients based on the level of resource availability.

Policies such as START, which were developed primarily to classify patients, are also insufficient to deal with many of the other decisions that must be made when responding

to an MCI. One of these relevant issues is how to dispatch resources, such as ambulances, to care for patients during an emergency involving more than one mass-casualty incident. Multiple location mass-casualty incidents are frequently associated with terrorist attacks (e.g., the London subway bombings of 2005 (Lockey et al. 2005), where casualties were spread among four subway stations). In addition, catastrophic weather events and other natural disasters (such as tornadoes or earthquakes) can also strain emergency medical services (EMS) systems, causing mass-casualty incidents in more than one location, each of which contributes demand for the same limited emergency resources. While EMS systems are able to plan ahead in allocating resources during daily operations, resource allocation decisions in an MCI are often made in an *ad hoc* manner by emergency response coordinators. Some states maintain “strike forces” of ambulances in case additional resources are requested by a locality due to a major disaster or MCI (State of California EMS Authority 2003), but to our knowledge, analytical results do not exist to assist responders in deciding when and where to dispatch these resources. Moreover, the decision on where to allocate resources will be affected by the method used for patient prioritization. In Chapter 4, we analyze how resources should be allocated to two or more MCIs in conjunction with the prioritization of patients.

Policies such as START also do not provide guidance on how to coordinate a prolonged response effort involving primarily patients who do not have life-threatening injuries. Large scale natural or man-made disasters often result in a very large number of casualties who do not have life-threatening conditions but nonetheless need to be evacuated and receive some kind of care. Even though these patients do not need advanced care, they may crowd and overwhelm emergency rooms and other treatment facilities. For example, casualties in a biological, chemical, or nuclear disaster must be transported out of the affected area and then undergo decontamination; casualties from an earthquake, flood, or hurricane must be moved to a safe place to undergo triage or to receive treatment for injuries. Confusion and fear may add to the difficulty in performing an efficient evacuation. Emergency responders may also encounter limited resources and unexpected events, such as infrastructure damage on the evacuation routes or overcrowding at some of the facilities that are set up to receive casualties. Nonetheless, the main goal of a large scale casualty evacuation is to ensure the

safe and timely treatment of as many casualties as possible. It is therefore important for emergency responders to be able to make decisions dynamically throughout a prolonged evacuation to make it as efficient as possible.

In Chapter 5, we study how to make routing decisions dynamically so that the evacuation can be adjusted to allow for the largest possible throughput of patients. While different types of disasters have different requirements in terms of resources needed or treatment provided, a common theme is that the dynamic management of the evacuation can improve outcomes and lessen confusion. In a study of thousands of patients evacuated from an earthquake, Tanaka et al. (1998) suggested that to improve evacuation, “disaster officials must know the capabilities and capacity of each area hospital at all times to select appropriate triage and mode of transport for each victim”. Hick et al. (2011) suggest that during a nuclear incident, emergency managers should bypass hospitals that are “completely overwhelmed”. These studies give qualitative observations about evacuation of casualties in a disaster; they suggest that dynamic management of casualty evacuation could be helpful, by incorporating information about the level of congestion at each casualty collection point. We attempt to provide quantitative insights to this problem in order to assist decision makers. Specifically, we consider the problem of transporting casualties without life-threatening injuries from areas affected by a disaster to a set of casualty collection points, where they receive treatment. This problem is particularly relevant in a biological, chemical, or nuclear disaster. In these types of disasters, casualties are evacuated from a “hot zone”, where there is risk of exposure to the agent, to a “warm zone”, where they wait to undergo decontamination, and finally to a “cold zone”, where they are safe to leave the area or proceed to additional treatment for injuries or illnesses (Boardman et al. 2008).

Throughout this dissertation, we maintain a common theme: development of analytical, numerical, and simulation results that lead to fundamental insights about the problem and toward policies that are both effective and practical. In Chapter 2, we review relevant literature from both the medical community and the operations community. In Chapters 3, 4, and 5, we provide the main results on patient prioritization, resource allocation, and routing, respectively. Finally, in Chapter 6, we summarize our work and suggest areas for further research. Proofs of analytical results are provided in the Appendix.

Chapter 2

Literature Review

The literature review is divided into two parts. In Section 2.1, we review articles by physicians and emergency medical personnel that examine decision making in mass-casualty incidents and suggest modifications or improvements. This review covers the work of Sacco et al. (2005), which proposes a completely new system of triage based on linear programming, and includes a discussion as to why this proposal has been received mostly with skepticism within the medical community. In Section 2.2, we review recent contributions within the Operations Research community to problems of patient prioritization, resource allocation, and routing/evacuation.

2.1 Emergency Medicine Literature

In the emergency medicine literature, there are numerous articles on patient triage and prioritization in the context of mass-casualty incidents. Some of these articles report details on emergency response efforts and outcomes in past mass-casualty incidents and daily emergencies, some provide general discussion on patient triage practices, discuss their shortcomings and point to future research directions, and some propose new triage methods or modifications to the existing ones (Frykberg 2005, Garner et al. 2001, Bostick et al. 2008, Jenkins et al. 2008, Hick et al. 2008, Lerner 2008, Kahn et al. 2009, Lerner et al. 2010). To a large extent, this literature does not utilize mathematical modeling and analysis. The notable exceptions to this are the articles by Sacco and his co-authors (Sacco et al. 2005 and 2007, Navin et al. 2009) in which the authors describe the Sacco Triage Method (STM) and test its performance by numerical experiments.

A review of the principles of triage by Frykberg (2005) outlines some of the many issues involved in making decisions about patient care in a mass-casualty incident. Of greatest relevance to our problem, Frykberg argues that the “main factor” that distinguishes a mass-casualty incident from daily emergency response is that the supply of resources is overwhelmed by the demand. He notes that unlike in routine trauma care, the limitation of available resources in a mass-casualty incident necessitates that some of the most severely injured patients will be abandoned so that resources may be diverted to those who have a good chance of survival. Frykberg emphasizes that while this principle may be difficult for medical providers to accept, it is necessary because the goal of triage in a mass-casualty incident should be to do the “greatest good for the greatest number.” Furthermore, he notes that the most difficult decision may be to determine which patients should be denied care, because this decision depends on many factors, such as the number of patients, severity of their injuries, and specific resources that are available. Garner et al. (2001) also emphasize the importance of considering the scale of the incident and the availability of resources in a mass-casualty incident, noting that current guidelines may sometimes prioritize patients whose injuries are severe enough that they should instead be considered unsalvageable based on the scale of the mass-casualty incident: “the decision not to treat certain patients is dependent on many factors other than just physiologic or anatomic signs of severe injury.” The authors add that such decisions can only be made with “knowledge of the available resources.”

As mentioned in the introduction, the *de facto* triage standard in the United States of America is Simple Triage and Rapid Treatment (START). It is important to realize that START actually consists of two phases: *classification* of the patients into triage classes and *allocation* of the available resources in a fixed-priority manner. Although some modifications to START have been proposed in the emergency response literature, they have generally been concerned with the classification phase (Jenkins et al. 2008, Lerner 2008). That is, modifications to START have mostly focused on developing new or revised criteria for placing patients in the different triage classes, in hopes of reducing classification errors. Discrepancies in classification are generally referred to as *overtriage*, if a patient is given an inappropriately severe classification, or *undertriage*, if a patient is given an inappropriately

mild classification. The notable exception to this generalization is again the work by Sacco et al., which aims to redefine both phases of triage.

Limited work on statistical analysis of triage protocols has generally focused on determining how well emergency responders are able to accurately classify patients (Garner et al. 2001) and on how well the triage classifications correspond to the actual severity of the injury, as determined by emergency physicians (Kahn et al. 2009). One reason for the scarcity of this type of analysis is that the chaotic environment of an MCI often leads to poor or insufficient documentation. For example, a retrospective analysis of the Oklahoma City Bombing determined that “documentation of the process behind the triage decisions made in Oklahoma City was practically nonexistent” (Hogan et al. 1999). Studies such as Kahn et al. (2009) used a standardized form to collect detailed triage-related information about patients in daily emergencies and determine START classification *a posteriori*, but it is not clear that such a form would be useful for tracking triage decisions in an MCI, especially because it is impossible to predict the location and timing of an MCI, and hence the forms would not necessarily be available.

The articles cited in this section highlight the general lack of medical research on prioritization and resource allocation in an MCI. Nonetheless, despite the lack of research in this area, the need to study resource allocation in mass-casualty triage is clearly expressed by Frykberg (2005) and other authors, and we pursue this direction in this dissertation. At the intuitive level, it may not be difficult to see that taking resource limitations into account could improve patient outcomes. To the best of our knowledge, the only proposed resource-based method in the emergency response literature is the Sacco Triage Method (STM) developed by W.J. Sacco and his co-authors (Sacco et al. 2005 and 2007, Navin et al. 2009). Sacco et al. propose a method that challenges the existing standards in *both* phases of triage: classification and resource allocation. Sacco Triage Method, or STM, involves solving a linear program to allocate resources to patients in thirteen different triage classes (Sacco et al. 2005 and 2007, Navin et al. 2009). The authors suggest that the patient data could be collected quickly and transmitted by the responders in the field to an emergency response center, where the LP would be solved. They further suggest that solving the LP with hypothetical scenarios could be useful for emergency planning and training.

While the latter seems like an obvious use for STM, critics have argued that it would be impractical to implement the volume of data collection and transmission that STM would require in order for it to be useful in the immediate aftermath of a real emergency. In a response to Sacco et al. (2005), Cone and MacMillan (2005) praised STM for its ability to incorporate information about the number of patients and resources, noting that no other existing triage system has been developed and evaluated from a mathematical perspective; however, the authors devoted much of their response to critiquing operational issues that could result if STM were implemented in an actual emergency situation. Anecdotal evidence suggests that the assertions of Cone and MacMillan (2005) may be correct: for example, in a retrospective analysis of the response to a bridge collapse, Hick et al. (2008) note that STM “would not have been helpful in this setting because no data on the casualties would have been entered by the time of transport.”

Part of the resistance to STM could also be due to the fact that it uses a different scheme for classifying patients from the current practice of START. According to STM, each patient is assigned an RPM (Respiration, Pulse, and Motor response) score, which is an integer from 0 to 12, with lower scores being associated with more critical conditions. Patients are then given priority according to their RPM score as determined by the solution to the linear program. Thus, implementation of STM requires the emergency responders to drop a patient classification scheme with which they are already familiar and which has already received wide acceptance, and switch to another scheme that is largely unknown by the emergency response community. Furthermore, with this switch, they will need to use prioritization policies for which they will not have an intuitive feel, because priorities are determined by the solution to the linear program, which is like a black box for the emergency response personnel. The thirteen possible RPM scores means that there are potentially thirteen different priority classes to deal with, and the examples provided in Sacco et al. (2005) show that the priority policy suggested by STM can be quite complex and difficult to interpret. For example, in a particular setting in which only six of the thirteen RPM scores were present, the authors reported that the optimal priority ordering for the RPM scores was 6, 8, 7, 4, 3, 4, 2, 3 (note that some scores are repeated, because they can obtain the highest priority repeatedly at different times).

The work of Hick et al. (2008) highlights another potential problem of STM: the training and comfort level of emergency responders. Although the responders to the bridge collapse did not use triage tags to identify the patients, paramedics “believed uniformly that they were rapidly able to sort red [Immediate], yellow [Delayed], and green [Minor] casualties without use of a triage tool” (Hick et al. 2008). In other words, paramedics relied on their training and experience in the classification phase of triage. Although there is no uniform national standard for mass-casualty triage, START is widely accepted (Lerner 2008). For example, Mistovich and Karren (2007), a widely used training manual for emergency medical technicians (EMTs), provides training in START. Introducing a radically different classification would require significant re-training and pose communication problems across jurisdictions if not uniformly implemented. Hence, there are good reasons to believe that for an improved triage system to have a chance of acceptance and implementation, it must be compatible with or similar to the widely adopted START classification.

Finally, the perceived proprietary nature of STM has hindered its ability to be accepted. Kahn et al. (2010) notes that “every published article on the Sacco Triage Method has been funded and written by the principals of ThinkSharp, Inc., the company that sells the Sacco Triage Method.” We further note that the numerical studies showing the benefit of using STM were constructed by the authors of the method, and to our knowledge, no systematic numerical studies or simulations have been used to test the performance of STM.

Notwithstanding the objections noted above, we believe that the promising numerical results of Sacco et al. (2005 and 2007) support our assertion that more information about patients and resources may improve the outcome of triage. However, the criticism of STM highlights the fact that the current research must focus on identifying policies that are much simpler and easier to implement than STM. In this respect, we diverge from STM. Rather than attempting to develop a decision support tool that claims to solve the triage and prioritization problem in real time, we provide general characterizations of the optimal policy, and we develop simple, practical policies based on those characterizations.

One aspect of the STM formulation that distinguishes it from the models used in papers from the operations literature is the manner in which criticality is modeled. In STM, patients do not die; rather, their condition deteriorates over time. In other words, delaying

service of a patient means that the system will simply receive a smaller reward from his or her eventual service. On the other hand, in the papers from the operations literature that we will review in Section 2.2, the concept of a deadline (for a job) or lifetime (for a patient) is used. In general, such deadlines or lifetimes may be deterministic or random, and if they are random, they may be known or unknown to the decision maker at the time when he or she must make the decision. In such a model, rather than having deterioration continuously or at several points in time, the amount of reward earned depends only on whether the job receives service before or after its deadline. The deterioration model simplifies the analysis somewhat because it does not need to deal with patients leaving the system through means other than service (i.e., abandonments).

2.2 Operations Research Literature

In the operations literature, interest in patient triage and prioritization in mass-casualty incidents has been relatively recent. Most of the classical scheduling literature is not directly applicable to the patient triage problem because it deals with problems over a longer time scale or with arrivals. Among the few articles in the literature that do apply to the patient triage problem, the two that are most relevant are Argon et al. (2008) and Jacobson et al. (2012) because they also deal with identifying effective patient prioritization policies in the aftermath of mass-casualty incidents. Argon et al. (2008) and Jacobson et al. (2012) consider models in which there are a finite number of patients at different casualty levels who are in need of a service that is in limited supply. The objective is to determine the order in which these patients will be served to maximize the system’s reward. In its simplest form, the reward is the expected number of survivors. In these two articles, each patient has a random lifetime whose probability distribution depends on the class of the patient, with more critical patients having shorter lifetimes in some stochastic sense. If a patient is not served before his or her lifetime, he or she dies and a reward is not earned. An alternative way of viewing this class of model is that when the patient dies, the system incurs the “cost” of delaying response to the patient. If a patient is served before his or her lifetime, the patient either definitely survives (in Argon et al. 2008) or survives with some

probability that depends on the class of the patient (in Jacobson et al. 2012). Even though this type of formulation captures the delay cost in a very direct and realistic way, one implicit assumption in these models is that the survival probability of a patient does not change with time. Even with this simplifying assumption, the resulting models turn out to be quite difficult to analyze. Allowing survival probabilities to be time-dependent would add another level of complexity to these models and render them analytically intractable. Despite the difficulty of the analysis, both Argon et al. (2008) and Jacobson et al. (2012) showed that it is possible to find simple but effective heuristic policies for patient prioritization. In particular, a threshold policy, where priority depended on the total number of patients, performed extremely well (Jacobson et al. 2012).

In addition to Argon et al. (2008) and Jacobson et al. (2012), there are three other articles that consider models in which patients may die, but their survival probabilities are not time-dependent. The work of Glazebrook et al. (2004), which predates both Argon et al. (2008) and Jacobson et al. (2012), is not specifically interested in patient triage and prioritization, but provides one result within a generic job scheduling framework that is also partially relevant in patient triage context. Specifically, the authors establish the near-optimality of a state-independent fixed priority ordering policy when lifetimes are long, which generally is not the case for seriously injured patients. Using a model that is very similar to that considered in Argon et al. (2008), Li and Glazebrook (2010) develop a heuristic method by applying a single-step of the policy-improvement algorithm on the state-independent policy proposed by Glazebrook et al. (2004). Childers and Taaffe (2010) report the results of a numerical study of a model that is again similar to that of Argon et al. (2008).

In our work on patient prioritization, we diverge from the recent work in the operations literature by allowing survival probabilities to be time-dependent, but without explicit lifetimes. In other words, we capture the “cost” of delaying service through a reward function that depends on time.

In contrast to the relative sparsity of research in triage, techniques from operations research have been applied more extensively in the area of evacuation planning in many different contexts (Regnier 2011); however, these efforts are not necessarily relevant to MCIs.

For example, many such articles have dealt with evacuation in advance of a disaster, which is not relevant to the aftermath of an MCI, and these generally fall into one of two categories: articles dealing with evacuation of a facility, and articles dealing with vehicle allocation and/or routing. The latter type is more closely related to our research on routing patients in the aftermath of an MCI. For example, Tayfur and Taaffe (2009) and Childers and Taaffe (2010) use optimization to identify staffing and transportation schedules to minimize the cost of evacuating a hospital within a certain period of time. Bish et al. (2011) consider the similar problem of evacuating patients from a hospital to a set of receiving hospitals while minimizing risk. Facility evacuations (and specifically hospital evacuations) can benefit greatly from advance planning because they are more controlled than the scene of a mass-casualty event, and are amenable to solution by optimization methods because most of the parameters are known.

Another area related to our research is vehicle routing, which is a well-studied problem in the field of operations research (Laporte 2007). Two articles that address vehicle routing and allocation in a mass-casualty incident are Gong and Batta (2007) and Jotshi et al. (2009). Gong and Batta (2007) study the problem of determining an initial allocation of ambulances to groups of casualties in a disaster, using a deterministic model to minimize the makespan, which in this context is the total amount of time needed to evacuate all the casualties. This model is most relevant to natural disasters, where the clusters of casualties may be far apart and hence it is desirable to allocate the resources to the clusters in a predetermined way or only at certain points in time, rather than dynamically. Jotshi et al. (2009) consider both dispatching and routing emergency vehicles in a disaster. The model of Jotshi et al. (2009) is designed to use large amounts of data about medical facilities, casualties, vehicle locations, and roadways in order to suggest a solution to the vehicle routing problem. While there are advantages to models that incorporate large amounts of data, the drawback of this type of model is that the large amount of data that must be collected and entered to solve the problem in real time, which may hinder its adoption as a practical tool. Nevertheless, these articles provide a context for our research on casualty routing. Both of these articles use optimization models, such as network flows and mathematical programming, to find an initial solution to the evacuation problem or to re-solve the problem at different points in

time.

In our study of casualty routing, we consider a different component of the evacuation problem; namely, we assume that the ambulance allocation decision has been made and we study the queueing of casualties at the collection points and its effect on dynamic routing decisions. This queueing is especially relevant in situations where a large number of casualties all need similar service (such as decontamination) and where the service can be provided at one of several locations.

Both the medical literature and the operations literature highlight the need for research to support decision making in mass-casualty incidents. The results demonstrated in the remainder of this dissertation can begin to fill in some of the gaps we have identified herein.

Chapter 3

Patient Prioritization with Resource-Based START (ReSTART)

The weak reception of Sacco Triage Method by the medical community, which was discussed in detail in Section 2.1, suggests that for a resource-based prioritization method to have some use in practice, it has to be relatively simple and somewhat intuitive, and its adoption should not necessitate a complete overhaul of the current classification scheme. These objectives provide the main motivation behind this chapter. Our goal is not to propose a real-time solution method but rather to carry out mathematical and numerical analysis to generate insights that would be useful in the design of effective yet simple prioritization policies. Specifically, we aim to provide answers to some of the questions that the emergency response community may face in the process of developing resource-based prioritization policies. For example, is it possible to develop policies that are simple enough to be implemented in practice yet have substantial benefits over standard policies that do not consider resource limitations? If so, what are the main characteristics of these policies, and in what kind of mass-casualty incidents are they likely to be most beneficial?

To answer these questions, we develop a fluid model in which patients do not die and leave the system. However, their survival probability (more generally, their reward) declines with time and possibly hits zero at some point. This means that according to our formulation, all patients, even those with zero survival probability, must be transported to the hospital. Although this may sound unrealistic at first, it does not have any practical drawback as long as the objective function is chosen reasonably. Since our objective is to

maximize the expected number of survivors (more generally, the total expected reward), the optimal policy would be always such that patients with zero survival probability (dead patients) are the last patients to receive service. In other words, the optimal solution will ignore those dead patients as long as there are patients with a positive probability of survival. Thus, our formulation achieves some mathematical simplicity without sacrificing realism in any crucial way. The details of our model are provided in Section 3.1.

Before we proceed, it is important to note that the objective that we consider in this chapter, i.e., the maximization of the expected number of survivors, is consistent with the widely accepted and practiced emergency response principle of *doing the greatest good for the greatest number* (Kennedy et al. 1996, Frykberg 2005). However, triage has always been a somewhat contentious practice, because it essentially entails favoring certain individuals over others. There is a long line of discussion and research on the ethical dimensions of triage. For more on this issue, which is beyond the scope of this dissertation, we refer the reader to Winslow (1982), Baker and Strosberg (1992), and references therein.

3.1 Model Formulation

We consider a scenario where there are many injured patients who need to be transported to a hospital. In particular, we consider the case where ambulances or other transportation resources are limited in supply so that at least some of the patients will have to wait for some time before being transported. We assume that at time zero the patients have already been separated into N classes based on their injury characteristics and medical conditions, and moved to a single area of the site where they are given basic treatment and prepared for loading onto the ambulances. According to our formulation, there will be no new patient arrivals after time zero. Thus, our model is a better fit for incidents where a significant percentage of the patients are quickly accounted for and thus the response effort does not necessitate a time-consuming search and rescue activity. Nevertheless, in Section 3.6.2, we use a simulation study to consider cases where some of the patients arrive with some delay. We denote the set of classes by $\mathcal{I} = \{0, 1, \dots, N - 1\}$ and the number of patients in class i by n_i , where $n_i > 0$. We also assume that all patients need to be transported to the same hospital via the same transportation mode (e.g., via ground transportation) so that

the transportation time of a patient does not depend on the patient's class. For simplicity, we will frequently use the word "service" to refer to the process of transporting a patient to the hospital.

We approach this problem from the perspective of the emergency response coordinator, who decides the order in which patients should be transported, with the objective of maximizing the overall expected reward or gain from the system. To this end, we assume that each class i has an associated non-negative reward function $f_i(t)$, which is the expected reward earned by the system if a class i patient is served at time t . To capture the fact that no patient would benefit from a delay in service, we assume that $f_i(t)$ is monotone non-increasing in t for each $i \in \mathcal{I}$. For mathematical tractability, we further assume that the first-order derivative of $f_i(t)$ with respect to t exists for each $i \in \mathcal{I}$, and is denoted by $f'_i(t)$. The function $f_i(t)$ can be interpreted as the probability that a patient of class i ultimately survives if taken into service at time t . With this interpretation, maximizing the total expected reward is then equivalent to maximizing the expected number of survivors.

Our goal is to develop a model that captures the essential features of the patient prioritization problem but is simple enough to allow mathematical analysis and development of easy-to-implement policies that are expected to perform well in practice. Towards that end, we propose a fluid formulation where different classes of patients in the system correspond to different classes of fluid and service of those patients corresponds to a flow of the respective fluid out of the system. Without loss of generality, we assume that the service rate is one patient per unit time; therefore, when patients of only class i are flowing out of the system at time t , reward is earned at a rate of $f_i(t)$.

Define a set of decision functions $\mathbf{r}(t) \equiv \{r_i(t) : [0, \infty) \rightarrow [0, 1], i \in \mathcal{I}\}$, where $r_i(t)$ is the rate at which we choose to serve class i patients or the fraction of the total service capacity allocated to class i patients at time $t \geq 0$. We restrict ourselves to decision functions that have finitely many discontinuities, which is needed to obtain solutions that switch priorities only finitely many times and hence are applicable in practice. We now state our

optimization problem as follows:

$$\begin{aligned}
& \max_{\mathbf{r}(t), t \in [0, \infty)} \sum_{i=0}^{N-1} \int_0^{\infty} r_i(s) f_i(s) ds \\
& \text{subject to} \quad \sum_{i=0}^{N-1} r_i(t) \leq 1, \quad \forall t \in [0, \infty) \\
& \quad \int_0^{\infty} r_i(t) dt = n_i, \quad i \in \mathcal{I}.
\end{aligned} \tag{3.1}$$

We first note that, as one would expect, it is suboptimal to leave any of the available capacity unused as long as there is fluid in the system. (The proof is omitted as it immediately follows from the assumptions that the reward functions $f_i(t)$ are non-increasing in t and there are no further arrivals.) The practical implication of this result is that in the rest of this chapter we do not need to consider solutions that involve idling. Since the service rate is one patient per unit time, under non-idling policies the fluid will be cleared from the system, i.e., transportation of the patients will be complete, at time $T = \sum_{i \in \mathcal{I}} n_i$. Thus, we can restrict ourselves to the time interval $[0, T]$.

Our fluid formulation allows the total service capacity to be allocated to more than one patient class at any particular point in time. The practical interpretation of such an allocation can be problematic because transportation vehicles cannot be allocated in a continuous manner. This would especially be difficult to deal with when there are few vehicles to allocate. However, the following proposition resolves this concern. (Proofs of Proposition 3.1 and all other propositions and theorems are provided in the Appendix.)

Proposition 3.1. *There exists an optimal solution to (3.1) where only one class of patient is served at any given time.*

Proposition 3.1 implies that we can restrict the set of policies we consider to those which serve only one patient class at any point in time. For practice, this result suggests that at any point in time, there is only one highest-priority class and all transportation resources available should be allocated to that class unless the number of such patients is less than the number of resources.

Proposition 3.1 is also useful technically as it allows us to consider a formulation that

is equivalent to but easier to analyze than (3.1). Define the set-valued decision variable $\mathbf{W} = \{W(i) : i \in \mathcal{I}\}$, where $W(i)$ is the set of time points during which class i is served. Then, we can rewrite our optimization problem in the following way:

$$\begin{aligned}
& \max_{\mathbf{W}} \quad \sum_{i=0}^{N-1} \int_{W(i)} f_i(t) dt & (3.2) \\
& \text{subject to} \quad \mu(W(i)) = n_i, \quad i \in \mathcal{I}, \\
& \quad \bigcup_{i=0}^{N-1} W(i) = [0, T], \\
& \quad W(i) \cap W(j) = \emptyset, \quad \forall i \neq j,
\end{aligned}$$

where $\mu(W(i))$ is the total amount of time spent serving class i patients. In the rest of our analytical work, we will focus on this formulation. For both practical and technical reasons, we restrict ourselves to solutions \mathbf{W} such that for each $i \in \mathcal{I}$, $W(i)$ consists of finitely many intervals, each of which is closed on the left and open on the right.

3.2 A Simple Condition for Fixed-Priority Ordering

Many triage methods that are used in practice, such as START, assign a fixed priority to each class of patients. To be more precise, in a fixed-priority method, the triage class of each patient determines his or her priority level, which does not change with time throughout the response effort. Although several examples show that in general, the optimal policy to (3.2) is not a fixed-priority policy (i.e., the optimal policy is such that the priority ordering of the patient classes changes with time), due to the simplicity of fixed-priority policies, it is still important to investigate conditions under which the optimality of a such a policy is guaranteed.

It turns out that one condition that ensures an optimal fixed-priority relationship between two classes of patients is an ordering between the derivatives of their respective reward functions.

Proposition 3.2. *Suppose that there exist two classes, i and j , which have the property*

that

$$f'_j(t) \leq f'_i(t) \quad \forall t \in [0, T]. \quad (3.3)$$

Given a feasible solution to (3.2), where some class i patients are served before some class j patients, there exists another solution where

- (i) no class i patients are served before class j patients, and
- (ii) the expected total reward obtained under the new solution is at least as large as the expected total reward obtained under the existing solution.

Proposition 3.2 implies that if class j patients deteriorate at least as fast as class i patients over $[0, T]$, then there exists an optimal solution where class j has priority over class i at all times. In order to intuitively understand Proposition 3.2, it helps to think about the “opportunity cost” of delaying service to each class for a period of time. If the expected reward function of class j always decreases faster than that of class i , we will forego more expected reward by delaying the service of class j than we would by delaying the service of class i , for any arbitrary amount of time. The optimal policy is to delay the service of the class for which there is less to lose with time.

There are a few important points worth emphasizing regarding Proposition 3.2. First, the proposition does not assume an ordering between $f_i(\cdot)$ and $f_j(\cdot)$. This means that the deterioration rates, not the nominal values of the expected rewards (e.g., survival probabilities), determine dominance. For example, it is possible for class i patients to have lower survival probabilities than class j patients at all times and yet be assigned a lower priority than class j patients if their health conditions do not deteriorate as fast. Second, it is crucial to ensure that the ordering in (3.3) holds for all $t \in [0, T]$. One might expect that if $f'_j(t) \leq f'_i(t) \forall t \in [0, t_0]$ for some $t_0 < T$, then we could at least apply the result of the proposition for the time period $[0, t_0]$. Nonetheless, we can easily construct examples where this intuition does not hold. And third, it might seem at first that Condition (3.3) does not depend on the scale of the mass-casualty incident, i.e., the number of patients of various classes, transportation capacity, etc. However, recall that T is the time the response effort is over, hence T increases with the number of patients and the response time per patient.

This means that as the scale of the incident gets larger, T becomes larger and as a result it becomes increasingly less likely for (3.3) to hold.

Proposition 3.2 directly leads to the complete characterization of an optimal policy when all classes can be ordered according to Condition (3.3):

Corollary 3.1. *Suppose that*

$$f'_{N-1}(t) \leq f'_{N-2}(t) \leq \dots \leq f'_0(t) \quad \forall t \in [0, T]. \quad (3.4)$$

Then there exists an optimal policy under which there is a fixed-priority ordering $(N-1, N-2, \dots, 0)$ among the patients so that for any $i \in \{N-2, N-3, \dots, 0\}$, class $i+1$ patients have priority over class i patients.

The policy prescribed by Corollary 3.1 is very simple and practical. Since the priority ordering is fixed, there is less room for mistakes during implementation. However, the optimality of the policy is only guaranteed under a relatively strong condition. Therefore, it is important to investigate what the optimal policy would be like when Condition (3.4) does not hold.

We pursue this question in the following section where we focus on the case with only two patient classes. There are two main reasons why we consider this particular scenario. First, characterization of the optimal policy—except under Condition (3.4)—appears to be very difficult when there are more than two patient classes. However, more importantly, the two-class case fits perfectly well with the widely adopted triage classification used in START. As we described in Section 1, even though START puts patients into four classes, patients who are in the *expectant* class have almost no chance to survive and those with *minor* injuries do not carry a risk of dying from their injuries. Therefore, the success of the response effort depends almost entirely on how priority decisions for patients in the *immediate* and *delayed* classes are handled. Hence, our analysis of the two-class case in the next section will help us explore how START can be expanded to include resource limitations using our formulation.

3.3 Priority Decisions for Two Classes of Patients

Suppose that there are only two patient classes, which we name class I (*immediate*) and class D (*delayed*). By letting $g(t) \equiv f_D(t) - f_I(t)$, we can rewrite the optimization problem (3.2) as

$$\begin{aligned} & \max_{W(D)} \int_{W(D)} g(t) dt + C & (3.5) \\ \text{subject to} & \quad \mu(W(D)) = n_D \\ & \quad W(D) \subseteq [0, T], \end{aligned}$$

where $C \equiv \int_0^T f_I(t) dt$ is a constant.

From Proposition 3.2, we know that if patients of class I consistently deteriorate faster than patients of class D over $[0, T]$, then class I should have priority over class D at all times. Although this is a possibility in practice, a more likely scenario is one where patients with very critical conditions, who would be classified as *immediate* according to START, have rapidly diminishing survival probabilities, which approach zero fairly quickly, while the survival probability function for *delayed* patients would be relatively flat initially and then start to decline rapidly after some point. We formally state this scenario in Assumption 3.1 and present an example pair of reward functions, $f_I(t)$ and $f_D(t)$ satisfying Assumption 3.1 in Figure 3.1.

Assumption 3.1. *There exists $t_m \in [0, T]$ such that $f'_I(t) < f'_D(t) < 0$ for all $t < t_m$, $f'_I(t_m) = f'_D(t_m)$, and $f'_D(t) < f'_I(t) < 0$ for all $t > t_m$. Equivalently, the reward gap function $g(t)$ has a unique maximum at $t_m \in [0, T]$, is increasing for $t < t_m$, and decreasing for $t > t_m$.*

We delay a more detailed discussion on the justification of Assumption 3.1 until Section 3.5.1, where we demonstrate that estimates for survival probability functions for immediate and delayed classes, which are based on data from Sacco et al. (2005), satisfy Assumption 3.1. The remainder of this section is devoted to characterizing the solution to (3.5) under Assumption 3.1 and providing insights into the patient prioritization problem based on our




Figure 3.1: Example of reward functions satisfying Assumption 3.1.

analytical results.

3.3.1 Characterization of the Optimal Policy

In this section, we establish a useful structural result for the optimization problem (3.5).

Proposition 3.3. *There exists an optimal policy in which $W(D)$ is a single interval.*

Proposition 3.3 implies that there is an optimal policy where once the service of delayed patients starts, it is never interrupted by the service of immediate patients. Hence, the question of finding the optimal set of time points where delayed patients should be served is reduced to the question of finding the time at which we should *begin* serving delayed patients. In other words, under Assumption 3.1, we can describe the optimal policy by a single time point, or threshold, and hence the optimization problem (3.5) reduces to

$$\max_{t \in [0, n_I]} v(t), \quad (3.6)$$

where $v(t) \equiv \int_t^{t+n_D} g(x) dx$. Let t^* denote a solution to problem (3.6). Then, given t^* , an optimal policy can be described as follows: serve immediate patients over $[0, t^*)$, switch to delayed patients at t^* and serve all of them over the time interval $[t^*, t^* + n_D)$, and finally switch back to serving immediate patients and serve all the remaining immediate patients over the interval $[t^* + n_D, T)$. This general description encompasses the special case where the policy is a fixed-priority ordering policy like START. More specifically, if $t^* = n_I$, then

all immediate patients are served before all delayed patients. On the other hand, if $t^* = 0$, then all delayed patients are served before all immediate patients. Threshold t^* can be interpreted as the time at which the “cost” of delaying service to delayed patients starts outweighing the “benefit” of providing service to immediate patients.

3.3.2 Determining the Optimal Threshold

We now investigate how one can obtain a solution to problem (3.6). We first show that the solution t^* is unique and it can be determined more easily by first solving a relaxation of (3.6).

Proposition 3.4. *(i) There is a unique optimal solution, \tilde{t} , to the optimization problem*

$$\max_{t \in [0, \infty)} v(t), \text{ and } \tilde{t} \in [\max\{0, t_m - n_D\}, t_m].$$

$$(ii) \ t^* = \min\{\tilde{t}, n_I\}.$$

$$(iii) \ t_m \in [t^*, t^* + n_D].$$

Part (i) of Proposition 3.4 partially characterizes \tilde{t} . Note that the difference between t^* and \tilde{t} is that while t^* is restricted to be no greater than n_I , \tilde{t} has no such restriction. Practically, \tilde{t} is the time at which the service of delayed patients should start even if there are still immediate patients in need of service, which is very likely to be the case when there are many immediate patients initially. On the other hand, if all immediate patients are served before \tilde{t} , i.e., $n_I < \tilde{t}$, then service of delayed patients should start at n_I , because idling is suboptimal. Therefore, t^* is the minimum of \tilde{t} and the time it would take to serve all immediate patients if they were given priority at all times. This is reflected in the relationship described in part (ii) of Proposition 3.4. Part (iii) of Proposition 3.4 states that the optimal service time interval for delayed patients must contain t_m . Thus, service of delayed patients should start late enough for the service interval to contain the time point at which deterioration of delayed patients becomes faster.

The following theorem provides a complete characterization of \tilde{t} , and thus a complete characterization of t^* , which subsequently leads to an algorithm for finding t^* .

Theorem 3.1. *Exactly one of the following three statements is true:*

- (i) $g(0) > g(n_D)$, in which case $t^* = \tilde{t} = 0$.
- (ii) $g(\tilde{t}) = g(\tilde{t} + n_D)$ and $g(n_I) \leq g(T)$, in which case $t^* = n_I \leq \tilde{t}$.
- (iii) $g(\tilde{t}) = g(\tilde{t} + n_D)$ and $g(n_I) > g(T)$, in which case $t^* = \tilde{t} < n_I$.

Using Theorem 3.1, it is straightforward to show that the following algorithm determines the optimal threshold t^* :

1. If $g(0) > g(n_D)$, return $t^* = 0$.
2. Else, if $g(n_I) \leq g(T)$, return $t^* = n_I$.
3. Else, return the solution of $g(t) = g(t + n_D)$.

Note that t^* is readily available if one of the two conditions in the first two steps holds. If neither holds, then one needs to determine the unique solution to $g(t) = g(t + n_D)$.

The first step in the algorithm checks whether $g(0) > g(n_D)$. This condition may hold when delayed patients deteriorate faster than immediate patients starting at either time zero or soon after, i.e., when t_m is close to or equal to zero. In this case, Theorem 3.1 indicates that $t^* = 0$, and hence, delayed patients have priority at all times. This policy, which we call Inverted START (InvSTART), is the complete opposite of START because delayed patients have priority over immediate patients at all times. The same condition may also hold when n_D is very large, suggesting that InvSTART is optimal when there are sufficiently many delayed patients. If the condition in the first step of the algorithm is not satisfied, i.e., $g(0) \leq g(n_D)$, then Theorem 3.1 states that $g(\tilde{t}) = g(\tilde{t} + n_D)$. In this case, t^* is equal to \tilde{t} or n_I , depending on whether the inequality $g(n_I) > g(T)$ holds or not. If the algorithm stops in the second step, i.e., $t^* = n_I$, then the optimal policy is START since immediate patients are given priority over all delayed patients at all times. This case, where $g(n_I) \leq g(T)$, may occur when either t_m is sufficiently large or the total number of patients is sufficiently small. Finally, if the algorithm stops in the third step, i.e., $t^* < n_I$, then the optimal policy is either InvSTART (when $t^* = 0$) or time-dependent (i.e., priority will change at time $0 < t^* < n_I$). Under a time-dependent policy, the service of immediate patients is interrupted by the service of delayed patients during $[t^*, t^* + n_D)$.

3.3.3 Sensitivity of the Optimal Policy to the Number of Patients

From Theorem 3.1, it is clear that the optimal prioritization policy depends on the number of patients in each class, since t^* is a function of both n_I and n_D . In order to better understand the relationship between the optimal policy and patient counts, we investigate how t^* changes with n_I and n_D .

By definition, \tilde{t} does not depend on n_I . Hence, by part (ii) of Proposition 3.4, t^* increases with n_I for $n_I < \tilde{t}$ but does not change for $n_I \geq \tilde{t}$. This suggests that if there are few immediate patients and START is optimal (i.e., $t^* = n_I$), then increasing the number of immediate patients will not change the policy at first (except for the time when the service of delayed patients should start), but will eventually change the policy from START to a time-dependent one. However, when there are enough immediate patients for a time-dependent policy or InvSTART to be optimal, then having more immediate patients does not change the optimal policy.

We next present a proposition that describes how t^* depends on n_D .

Proposition 3.5. *Everything else remaining the same, t^* either decreases or stays the same as n_D increases, i.e., having more delayed patients can only decrease the time the service of delayed patients starts under the optimal policy.*

Proposition 3.5 and the discussion above yield the following conclusions:

- (i) If the optimal policy is START, having more patients, regardless of their class, may make the optimal policy time-dependent.
- (ii) If the optimal policy is time-dependent, having more immediate patients will not change the policy, whereas having more delayed patients will push the optimal policy toward InvSTART.
- (iii) If the optimal policy is InvSTART, having more patients will not change the optimal policy.

3.4 A New Policy for Patient Triage: ReSTART

In this section, building on our analytical results from previous sections, we demonstrate how one could construct a new patient prioritization policy that takes into account resource limitations, yet is simple enough for practical implementation. More specifically, we carry the simple solution from Section 3.3 to practical settings where the fluid assumptions are obviously violated. We call the new policy *Resource-based START* (ReSTART) to indicate the fact that it builds on START, which is the most widely adopted triage method in U.S. It is important to emphasize that ReSTART does not propose any new medical criteria to classify patients. ReSTART uses the START classes, but unlike START, it does not necessarily give priority to immediate patients at all times. Under ReSTART, delayed patients can get priority over immediate patients depending on the relative availability of the transportation vehicles with respect to the number of patients.

Now, let θ denote the expected transportation time for each patient, and K denote the number of available transportation vehicles. Recall that in Section 3.3, we normalized the service rate to one, which is the same as assuming that $K/\theta = 1$. Incorporating generality in service rates, i.e., allowing a general number of vehicles and general transportation times, simply requires scaling the number of patients by θ/K . The description below is based on the algorithm given in Section 3.3.2.

Resource-based START (ReSTART):

1. Classify patients according to the START classes.
2. Determine the number of patients classified as immediate (n_I) and the number of patients classified as delayed (n_D). Determine θ , the expected round-trip travel time for each transportation vehicle, and K , the number of vehicles that can be used for transporting patients to the hospital.
3. Determine priorities among the immediate and delayed patients as follows:
 - (i) If $g(0) > g(n_D\theta/K)$, transport all delayed patients first, followed by all immediate patients.

- (ii) If $g(n_I\theta/K) \leq g((n_I + n_D)\theta/K)$, transport all immediate patients first, followed by all delayed patients.
- (iii) Otherwise, determine t^* such that $g(t^*) = g(t^* + n_D\theta/K)$. Transport immediate patients until time t^* or until there are no more remaining immediate patients. Then, start transporting delayed patients and continue until there are none remaining. Finally, continue with the transportation of any remaining immediate patients.

Given the reward functions $f_I(\cdot)$ and $f_D(\cdot)$, implementation of ReSTART is mostly straightforward. While in steps 3(i) and 3(ii), one only needs to check whether the given inequalities hold, in step 3(iii) a solution to $g(t^*) = g(t^* + n_D\theta/K)$ must be found. Because the right-hand side of the equation depends on n_D , θ , and K , t^* can be computed only after the incident occurs; in other words, it cannot be computed “off-line.” The computation of t^* can be done very quickly using a line-search algorithm, since $g(\cdot)$ is a unimodal function. Nevertheless, this cannot necessarily be done by hand, which might be a cause for resistance to any potential implementation of ReSTART. Therefore, we simplify the policy further by proposing an approximation for t^* based on our analytical results. In order to distinguish this approximate version of ReSTART from the exact version described above, we call it Quick-ReSTART (Q-ReSTART). In particular, we propose two different versions of Quick-ReSTART, which we call QuickDynamic-ReSTART (QD-ReSTART) and QuickStatic-ReSTART (QS-ReSTART). In the following, we describe these two policies.

3.4.1 QuickDynamic-ReSTART (QD-ReSTART)

QD-ReSTART is essentially the same as ReSTART except that it does not use the exact value of t^* , but rather an approximation for t^* . From Proposition 3.4, with the proper scaling for the expected travel time and the number of ambulances, we know that $\tilde{t} \in [t_m - n_D\theta/K, t_m]$. Therefore, even if we cannot locate \tilde{t} exactly, it would be reasonable to expect that approximating \tilde{t} with a choice from this interval could lead to a policy that performs well. Now, we know that there exists $\tilde{\phi} \in [0, 1]$ such that $\tilde{t} = t_m - \tilde{\phi}n_D\theta/K$. Instead of determining $\tilde{\phi}$ exactly, we approximate it by some $\phi \in [0, 1]$. Then, we use

$\tau = t_m - \phi n_D \theta / K$ instead of \tilde{t} and set $t^* = \min\{n_I, \tau\}$ by part (ii) of Proposition 3.4. We call the policy that uses this approximation “QD-ReSTART(ϕ).” Thus, QD-ReSTART is in fact a family of policies, with each policy being uniquely described by a value of the parameter $\phi \in [0, 1]$. Below is a formal description of QD-ReSTART.

QD-ReSTART(ϕ):

1. Same as ReSTART.
2. Same as ReSTART.
3. Compute $\tau = t_m - \phi n_D \theta / K$ and prioritize the immediate and delayed patients as follows:
 - (i) If $\tau \leq 0$, transport all delayed patients first, followed by all immediate patients.
 - (ii) If $\tau \geq n_I \theta / K$, transport all immediate patients first, followed by all delayed patients.
 - (iii) If $0 < \tau < n_I \theta / K$, transport immediate patients until time τ or until there are no more remaining immediate patients. Then, start transporting delayed patients (if there are any) and continue until there are no remaining delayed patients. Finally, continue with the transportation of any remaining immediate patients.

As a result of using τ in place of \tilde{t} , the inequalities in steps 3(i) and 3(ii) of ReSTART simplify to conditions that are easier to check. As an added benefit, these conditions also have insightful interpretations. Here, τ can be interpreted as a measure of the availability of resources relative to the size of the event. It is larger when there are fewer delayed patients, when transportation times are shorter, and/or when more vehicles are available for transportation. Thus, lower values of τ indicate more serious resource limitations.

Although ReSTART is the optimal policy for our fluid model under Assumption 3.1 and is likely to perform better than QD-ReSTART even under realistic conditions where the fluid assumption is relaxed, QD-ReSTART is simpler and more practical. Like ReSTART, it requires estimates for the expected travel time, number of ambulances, and number of immediate and delayed patients, which should not be difficult to determine and which are

likely to be the minimal set of requirements for any policy that takes resource limitations into account. However, unlike ReSTART, QD-ReSTART requires only an arithmetic calculation at the time of implementation. It simply uses t_m and ϕ , which can be obtained off-line before the incident based on estimates for the reward functions. Furthermore, τ depends on the function $g(\cdot)$ only through its maximizer t_m , meaning that the only estimation required is for the time at which deterioration rate of the delayed patients exceeds that of the immediate patients.

To understand how QD-ReSTART works, it is useful to examine the leftmost plot in Figure 3.2, which depicts the structure of QD-ReSTART(0.5). In this plot, the horizontal axis is for $n_D\theta/K$, which is the expected time it would take to transport all delayed patients, while the vertical axis is for $n_I\theta/K$, which is the expected time it would take to transport all immediate patients using the full transportation capacity. It is immediate from the figure that for given values of expected transportation time and number of available vehicles, the priority ordering according to QD-ReSTART depends on the initial number of immediate and delayed patients. When there are few patients of both classes (in the triangular region at the bottom left), QD-ReSTART reduces to START, giving priority to immediate patients until they are all transported. When there are sufficiently many delayed patients (far right in the figure), regardless of the number of immediate patients, the priority is reversed: QD-ReSTART reduces to InvSTART and transportation of immediate patients starts after all delayed patients are transported. On the other hand, when there are sufficiently many patients but the number of delayed patients is below a certain level, priority ordering changes with time; immediate patients have priority initially, but at some specified time, priority moves to delayed patients even if there are still immediate patients waiting. Those remaining immediate patients wait until all the delayed patients are transported. Note that this structure for QD-ReSTART is consistent with the behavior of the optimal solution to the fluid model (see Section 3.3.3).

One remaining question is how to set the value for ϕ . More empirical work is needed to make a more confident decision about this, but we demonstrate in Section 3.5 that for the survival probability data that we use in our simulation experiments, setting $\phi = 0.5$ provides a good performance for the QD-ReSTART policy. For more on the justification of

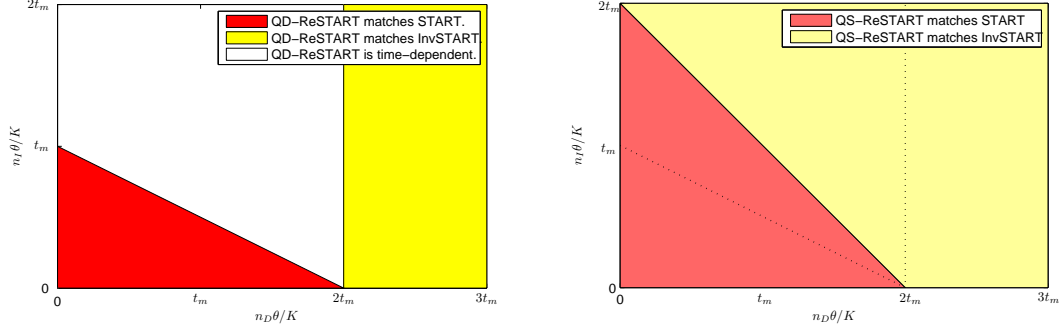


Figure 3.2: Visualizations of QD-ReSTART(0.5) and QS-ReSTART(0.5).

the use of $\phi = 0.5$, see Section 3.5.1.

3.4.2 QuickStatic-ReSTART (QS-ReSTART)

QD-ReSTART is a dynamic policy in the sense that the class that has the higher priority can change as time passes during the response effort. Although this priority switch can only happen once, one might still want to use even a simpler policy that fixes the priority levels at the beginning and does not change them later on. More precisely, one can choose either START or InvSTART given the conditions at time zero and use it until all patients are transported. We propose such a policy, which we call QS-ReSTART(ϕ), based on our analytical characterization of ReSTART, more specifically QD-ReSTART(ϕ).

We can observe from Figure 3.2 that QD-ReSTART is a time-dependent policy in only one of the three regions. To develop QS-ReSTART(ϕ), we simply divide this region into two with a line that passes through the points $(t_m/\phi, 0)$ and $(0, t_m/\phi)$, merge the left part with the already existing “START” region, and merge the right part with the already existing “InvSTART” region, thereby eliminating the time-dependent policy region completely (see the rightmost plot in Figure 3.2, where $\phi = 0.5$). The policy can then be described simply as follows: use START if $n_I + n_D \leq Kt_m/(\phi\theta)$ and use InvSTART otherwise. Note that one nice feature of this policy is that whether START or InvSTART is chosen depends on the total number of patients $n_I + n_D$, not on n_I and n_D individually. This means that once the total number of patients is determined, the policy can be determined even before triage is over. Furthermore, the policy is robust to classification errors between immediate and

delayed classes since such errors would not change the total number of critical patients.

3.5 Simulation Study

In this section, we carry out a simulation study to investigate how QD-ReSTART and QS-ReSTART perform in comparison with START and InvSTART under conditions that are more realistic than those of the fluid model. Specifically, we study a simulation model where patients are discrete entities, ambulances are discrete resources, and transportation times are stochastic. In Section 3.5.1, we provide details on our experimental setup. Our results on the comparison of QD-ReSTART and QS-ReSTART with START and InvSTART are provided in Section 3.5.2. In Section 3.5.3, we present a sensitivity analysis with respect to the reward functions that are used in our simulation study.

3.5.1 Experimental Setup

We consider a mass-casualty incident that takes place at a single location and results in a number of patients who need to be transported to a hospital. We assume that the patients have already been classified and they are ready to be transported. The initial arrival of the ambulances to the site follows a Poisson process. For each ambulance, we assume that the round-trip travel time between the incident location and the hospital has lognormal distribution with a mean of 30 minutes and a standard deviation of 12 minutes. This choice is based on an empirical study by Ingolfsson et al. (2008), which reports that a lognormal distribution with standard deviation that is equal to 40% of the mean is a good fit for ambulance travel times.

We use the *critical mortality rate*, i.e., the percent of critical patients who die, as the performance measure of interest. Frykberg (2005) states that critical mortality is the best measure of performance for mass-casualty triage because it takes into account only those patients whose conditions are serious enough to require timely treatment and who also have a non-negligible chance of survival. Note that these are the only patients for whom a priority policy can make a difference.

In our numerical experiments, we assume that patients are categorized according to

START guidelines, since START is the most widely accepted classification method in the U.S. Recall that there are four classes of patients according to START: expectant (E), immediate (I), delayed (D), and minor (M). Patients who fall into the immediate and delayed classes are considered critical patients. In order to use critical mortality rate as our performance measure, we need estimates for the survival probabilities of these critical patients as functions of time. To the best of our knowledge, the only work that attempted to estimate survival probability functions is due to Sacco and his co-authors in Sacco et al. (2005 and 2007), where the estimates are for a given initial RPM (Respiration, Pulse, and Motor response) score of a patient.

In order to obtain estimates of survival probability functions for the critical START classes (i.e., immediate and delayed classes), we utilized the RPM score-based estimates by Sacco et al. (2007). Our analysis revealed that the following three-parameter function is a good model for the reward functions $f_I(t)$ and $f_D(t)$:

$$f_i(t) = \frac{\beta_{0,i}}{\left(\frac{t}{\beta_{1,i}}\right)^{\beta_{2,i}} + 1}, \text{ for } i \in \{I, D\}, \quad (3.7)$$

where $\beta_{j,i} > 0$ for $j = 0, 1, 2$ and $i \in \{I, D\}$. Note that this function is a scaled version of the log-logistic distribution, which is commonly used in survival analysis (Cox and Oakes 1984). Furthermore, as we discuss next, this function provided a good fit to the empirical data that originated from Sacco et al. (2007).

We estimated the survival probability function for a given class $i \in \{I, D\}$ by

$$f_i(t) = \sum_{j=0}^{12} \pi_i(j) s_j(t) \text{ for } i \in \{I, D\}, \quad (3.8)$$

where $s_j(t)$ is the probability that a patient with RPM value $j \in \{0, 1, \dots, 12\}$ ultimately survives if he or she is transported at time t and $\pi_i(j)$, for $i \in \{I, D\}$ and $j \in \{0, 1, \dots, 12\}$, is the probability that a randomly selected patient who is in START class i would have an RPM score of j . (RPM can take any integer value between 0 and 12, with lower values indicating more critical conditions.)

The survival probability functions $s_j(\cdot)$, for $j \in \{0, 1, \dots, 12\}$, were estimated in Sacco

et al. (2005 and 2007) for three different types of injuries, and in this work, we use the estimates for penetrating injuries provided in Sacco et al. (2007). To obtain estimates for $\pi_i(j)$, for $j \in \{0, 1, \dots, 12\}$, we consulted Professor James E. Winslow (Winslow 2010), who informed us that his estimates for the distribution of the START class of a patient given his/her RPM score would be more reliable than those for the RPM score distribution given a patient's START class (i.e., $\pi_i(j)$, for $j \in \{0, 1, \dots, 12\}$). Hence, we expressed $\pi_i(j)$, using Bayes' Law, as follows:

$$\pi_i(j) = \frac{q_j p_j(i)}{\sum_{k=0}^{12} q_k p_k(i)}, \text{ for } i \in \{I, D\} \text{ and } j \in \{0, 1, \dots, 12\}, \quad (3.9)$$

where q_j is the probability that a randomly chosen patient has an RPM score of $j \in \{0, 1, \dots, 12\}$ and $p_j(i)$ is the probability that a patient with an RPM score of $j \in \{0, 1, \dots, 12\}$ belongs to START-class of $i \in \{E, I, D, M\}$. The probabilities $p_j(i)$ were estimated by Prof. Winslow. The only remaining estimate we need is for q_j , for $j \in \{0, 1, \dots, 12\}$. It is likely that this distribution varies depending on the type of injuries and event. Therefore, in our simulation study, we systematically considered different probability distributions for the initial RPM score of a patient (q_j). For each distribution, we determined the corresponding survival probability functions for both immediate and delayed patients using (3.8).

More specifically, we assumed that the initial RPM scores of the patients came from a discretized and rescaled version of the Beta distribution, which is typically used in the absence of data (Law 2007, Chapter 6). We considered five scenarios using five different Beta distributions with the two parameters (α_1, α_2) given by (1.5,5), (1.5,3), (1,1), (3,1.5), and (5,1.5). As we go from the first scenario to the last, the distribution changes from being right-skewed to left-skewed. Scenario 3, where $\alpha_1 = \alpha_2 = 1$, is the case where RPM scores are uniformly distributed. The corresponding empirical survival probability functions for immediate and delayed patients as well as the probability distributions for the START classes are provided in Figure 3.3.

For each of the five scenarios we constructed above, we fit the immediate and delayed reward functions ($f_I(t)$ and $f_D(t)$, respectively) to the three-parameter function given by (3.7) using Matlab's `nlinfit` function. The fitted parameters for each scenario are given

in Table 3.1. As shown in Figure 3.3, (3.7) was a very good fit for the empirical data.

Table 3.1: Fitted parameters for the five survival scenarios.


Scenario	Immediate			Delayed		
	$\beta_{0,I}$	$\beta_{1,I}$	$\beta_{2,I}$	$\beta_{0,D}$	$\beta_{1,D}$	$\beta_{2,D}$
1	0.09	17	1.01	0.57	61	2.03
2	0.15	28	1.38	0.65	86	2.11
3	0.24	47	1.30	0.76	138	2.17
4	0.40	59	1.47	0.77	140	2.29
5	0.56	91	1.58	0.81	160	2.41

The parameters β_0, β_1 , and β_2 have interpretations in the context of patient survival probabilities. Because $f_i(0) = \beta_0$, β_0 is the probability that the patient will survive if transported at time zero. By varying β_1 , the time scale of the patients' deterioration changes, i.e., increasing (decreasing) β_1 is equivalent to slowing down (speeding up) time. Finally, by varying β_2 , the shape of the curve changes. A larger value of β_2 results in a “sharper” decline in survival probability around the inflection point of the curve.

A natural question is whether functions in the form of (3.7) satisfy Assumption 1. Note that $f_i(t)$ has an inflection point at $\bar{t}_i = \beta_{1,i}[(\beta_{2,i} - 1)/(\beta_{2,i} + 1)]^{1/\beta_{2,i}}$ such that $f_i(t)$ is concave for $t > \bar{t}_i$ and convex for $t < \bar{t}_i$, where $i \in \{I, D\}$. Hence, $g(t) = f_D(t) - f_I(t)$ is concave on the interval $[\bar{t}_I, \bar{t}_D]$. This property is sufficient to ensure that if $\bar{t}_I < \bar{t}_D$, then there exists an interval $[\bar{t}_I, \bar{t}_D]$ over which the behavior of the functions $f_I(\cdot)$ and $f_D(\cdot)$ is consistent with Assumption 1. Furthermore, as we discuss in the following paragraph, Assumption 1 holds for $f_i(t)$ functions that we fitted to the data used in our experiments.

In our numerical study, we consider five possible scenarios, which mainly differ according to how “pessimistic” they are regarding the size of the mass-casualty incident and the urgency of patients. More specifically, in Scenario 1, the survival probabilities are low and most of the patients are in immediate or expectant categories; in Scenario 5, survival probabilities are much larger and there are not many immediate or expectant patients; and Scenarios 2, 3, and 4 are in the middle in terms of the severity of the event. Figure 3.3 explicitly shows the differences among these five scenarios. In the first column, we provide the empirical data on the survival probabilities, the fitted functions given by (3.7), and also

Figure 3.3: Survival probability (reward) functions, $g(t)$, and START class distributions for the five scenarios.



t_m values. We observe that as we move from Scenario 1 to Scenario 5, survival probabilities at any given time and t_m increase. In the second column, we plot the difference between the two fitted survival probability functions, i.e., $g(t)$, which shows that Assumption 1 holds over the time interval of interest. Finally, the third column shows the probability distributions for the START classes. We observe that while immediate patients constitute the highest

percentage in most of the scenarios, their frequency decreases as we move from Scenario 1 to Scenario 5, indicating a decreasing level of criticality in the patient population.

For each scenario, we generated 500 instances of 50 patients each, where each patient's START class was randomly determined using the START class probability distribution associated with that scenario. We used three different values for K , the number of ambulances available for transportation: 5, 10, and 15. The rate of initial arrival for ambulances to the event location was chosen as 10, 20, and 30 per hour when there are 5, 10, and 15 ambulances available for transportation, respectively. By keeping the number of patients and expected travel times constant but varying the number of resources, we were able to examine scenarios that ranged from resource-scarce to resource-abundant.

Within each scenario and resource level, we determined the performances of four policies, namely, START, InvSTART, QD-ReSTART(0.5), and QS-ReSTART(0.5) for each of the 500 randomly generated instances, using 200 replications for each instance. (To perform these simulation runs, we used code written in Matlab.) Under the START policy, immediate patients have priority over delayed patients at all times, while under the InvSTART policy, delayed patients have priority over immediate patients at all times. We are not aware of any implementation of InvSTART in practice, but we believe that it is useful as a point of comparison. Under QD-ReSTART, the higher-priority class possibly changes with time, although in some instances, QD-ReSTART is equivalent to START or InvSTART as can be seen in Figure 3.2. Finally, under QS-ReSTART, we have a fixed-priority policy, which is either START or InvSTART depending on the initial conditions in each scenario. In order to implement QD-ReSTART and QS-ReSTART, we needed to set an appropriate value for ϕ . From Figure 3.3, we observed that $g(t)$ is almost symmetric around t_m over $[0, 2t_m]$, which means that setting ϕ to 0.5 will yield $\tau \approx \tilde{t}$, i.e., QD-ReSTART(ϕ) will be approximately identical to ReSTART. We also calculated the value of $\tilde{\phi}$ that would make $\tau = \tilde{t}$ for each instance generated where t^* is non-zero. We then constructed 95% confidence intervals on the mean values of such $\tilde{\phi}$'s for all five scenarios and three levels of resource availability, which are shown in Table 3.2. The means of these fifteen confidence intervals ranged between 0.42 and 0.51, and the half-lengths of these intervals were all less than 0.01. Based on these observations, we set ϕ to 0.5 in all our simulation experiments that involve

Table 3.2: Mean values of ϕ over 500 instances for the five survival scenarios, excluding instances for which $t^* = 0$. All 95% half-widths are less than 0.01.

No. Ambulances	5	10	15
Scenario 1	0.43	0.43	0.42
Scenario 2	0.48	0.47	0.46
Scenario 3	0.48	0.50	0.51
Scenario 4	0.46	0.48	0.48
Scenario 5	0.46	0.47	0.48

QD-ReSTART and QS-ReSTART.

3.5.2 Comparison of QD-ReSTART and QS-ReSTART with START and InvSTART

We first compare QD-ReSTART with START and InvSTART in terms of the critical mortality rate in Tables 3.3 and 3.4, respectively. In both tables, the very last column reports the number of instances (out of 500) in which the mean performance difference is statistically significant at the 0.05 level. The numbers in the large middle column of Table 3.3 [Table 3.4] provide a statistical summary of the distribution of the mean reduction in critical mortality obtained by using QD-ReSTART as opposed to using START [InvSTART] based on 500 simulated instances: the minimum improvement (min), the first quartile (Q1), the median (med), the third quartile (Q3), and the maximum improvement (max). The sixth and seventh numbers in each row are the mean and 95% half-width for the mean (HW).

From Table 3.3, we observe that QD-ReSTART performs at least as well as START in all instances while performing better in many. The magnitude of the improvement depends on the scenario considered. We cannot claim any one scenario as being more realistic than the others; however, we can still make a few insightful observations. From Table 3.3, we can observe that going from Scenario 1 to Scenario 5, by which the survival probability estimates become more “optimistic,” performance improvement with QD-ReSTART first increases and then decreases. The most significant improvement is in Scenarios 2, 3, and 4. In Scenario 5 with a large number of ambulances, which is the most optimistic case considered in this study, the performance improvement is the smallest. This is because

Table 3.3: Mean reduction in the critical mortality rate obtained using QD-ReSTART instead of START.

									# Significant
	K	Min	Q1	Med	Q3	Max	Mean	HW	Instances
Scenario 1	5	0.0%	0.0%	0.0%	1.2%	5.8%	0.7%	0.1%	223
	10	0.0%	0.0%	0.0%	1.0%	4.8%	0.5%	0.1%	223
	15	0.0%	0.0%	0.0%	0.7%	3.5%	0.4%	0.0%	223
Scenario 2	5	0.0%	1.2%	2.3%	4.2%	8.8%	2.8%	0.2%	465
	10	0.0%	0.8%	1.6%	3.1%	6.5%	2.0%	0.1%	465
	15	0.0%	0.5%	1.0%	1.8%	4.1%	1.2%	0.1%	465
Scenario 3	5	0.0%	5.1%	6.1%	7.8%	11.4%	6.2%	0.2%	499
	10	0.0%	1.5%	2.0%	2.6%	4.5%	2.1%	0.1%	499
	15	0.0%	0.4%	0.5%	0.7%	1.5%	0.6%	0.0%	486
Scenario 4	5	2.3%	5.9%	6.8%	7.8%	10.3%	6.8%	0.1%	500
	10	0.1%	0.8%	1.1%	1.5%	3.1%	1.1%	0.0%	484
	15	0.0%	0.0%	0.0%	0.1%	0.5%	0.1%	0.0%	15
Scenario 5	5	0.0%	1.5%	2.2%	2.9%	5.7%	2.2%	0.1%	478
	10	0.0%	0.0%	0.0%	0.0%	0.3%	0.0%	0.0%	0
	15	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0

under Scenario 5, delaying transportation of the delayed patients affects them the least, so one can “afford” to use START by transporting all immediate patients first before moving on to the delayed patients. In Scenario 1, on the other hand, survival probabilities for both types of patients are so low that there is not much room for reduction in critical mortality. As a result, the difference between the performances of *any* two priority policies cannot be large. Nevertheless, even under such a pessimistic scenario, the mean improvement by QD-START is statistically significant regardless of the number of ambulances.

The number of available ambulances also has a clear effect on the performance improvement achieved by QD-ReSTART. Under all scenarios, the improvement with QD-ReSTART is larger when there are fewer ambulances. When there are many ambulances available, even when all immediate patients have priority over all delayed patients, the transportation of delayed patients will not be delayed very long. However, when resources are scarce, by switching the priority to delayed patients after a certain period of time, QD-ReSTART

saves more of these delayed patients who would otherwise have a lower chance of survival when they are transported.

Table 3.4: Mean reduction in the critical mortality rate obtained using QD-ReSTART instead of InvSTART.

									# Significant
	K	Min	Q1	Med	Q3	Max	Mean	HW	Instances
Scenario 1	5	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0
	10	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0
	15	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0.0%	0
Scenario 2	5	0.0%	0.0%	0.0%	0.0%	0.1%	0.0%	0.0%	0
	10	0.0%	0.1%	0.1%	0.1%	0.1%	0.1%	0.0%	254
	15	0.0%	0.1%	0.1%	0.2%	0.2%	0.1%	0.0%	333
Scenario 3	5	0.0%	0.2%	0.2%	0.3%	0.4%	0.2%	0.0%	392
	10	0.0%	0.5%	0.6%	0.6%	0.8%	0.6%	0.0%	499
	15	0.0%	0.6%	0.7%	0.8%	1.1%	0.7%	0.0%	498
Scenario 4	5	0.0%	0.4%	0.6%	0.8%	1.3%	0.6%	0.0%	422
	10	1.1%	1.7%	1.8%	1.9%	2.6%	1.8%	0.0%	500
	15	1.2%	2.2%	2.4%	2.6%	2.9%	2.3%	0.0%	500
Scenario 5	5	0.0%	1.2%	1.6%	1.9%	3.2%	1.5%	0.0%	494
	10	0.9%	3.2%	3.4%	3.5%	3.7%	3.3%	0.0%	500
	15	0.8%	2.6%	2.9%	3.1%	3.6%	2.8%	0.0%	500

From Table 3.4, we observe that the performance of QD-ReSTART is also at least as good as that of InvSTART in all instances considered. Perhaps not surprisingly, in contrast with START, InvSTART does better in pessimistic scenarios, with a performance matching that of QD-ReSTART. In the more optimistic scenarios, InvSTART performs poorly, especially when there are many ambulances. This is because InvSTART transports delayed patients first at the expense of delaying immediate patients, even though delayed patients can actually “afford” to wait longer.

Finally, we compare QS-ReSTART with START, InvSTART, and QD-ReSTART in terms of the critical mortality rate in Table 3.5. We can observe that the mean improvements over START and InvSTART using QS-ReSTART are smaller than those using QD-ReSTART, but not drastically. This same observation is reflected in the negative

values in the direct comparison with QD-ReSTART: QS-ReSTART performs worse than QD-ReSTART but only by a relatively small amount. This suggests that QS-ReSTART would be a reasonable alternative to QD-ReSTART if time-dependent priority levels turn out to be difficult to implement in practice. Note that under each scenario, for any given number of ambulances at least one of START or InvSTART performs very similarly to QS-ReSTART. This is expected since QS-ReSTART essentially chooses one of the two.

3.5.3 Sensitivity Analysis on the Reward Functions

To evaluate whether similar performance improvement could be achieved even without precise knowledge of the reward functions, we conducted a sensitivity analysis for the β parameters in the fitted reward functions given by (3.7). We repeated the simulations on the same instances used in the study presented in Section 3.5.2. However, this time, for each instance, we randomly perturbed the time-zero probability β_0 , the scale parameter β_1 , and the shape parameter β_2 , for both reward functions that are used as inputs to the simulation runs, but we did not perturb the reward functions while determining the operating parameters of the QD-ReSTART(0.5) and QS-ReSTART(0.5) policies (i.e., t_m). This way, we were able to test the performance of QD-ReSTART and QS-ReSTART policies when the estimates for reward functions were not accurate.

We considered two experimental settings for the perturbations. In Setting 1, each β parameter was equally likely to decrease 10%, decrease 5%, stay the same, increase 5%, or increase 10%, whereas in Setting 2, each β parameter was equally likely to decrease 20%, decrease 10%, stay the same, increase 10%, or increase 20%. In the interest of space, we here provide results only on Setting 1 and note that the main conclusions under Setting 2 are similar but the differences between the perturbed and original results are more pronounced. We also do not report our sensitivity results on the comparison of ReSTART policies with InvSTART here because the insights gained are very similar for those obtained on the comparison with START.

The results of the sensitivity analysis of comparison of ReSTART policies with START under Setting 1 are summarized in Table 3.6. This table follows the same format as Table 3.5, except that we now separate the number of instances where ReSTART policies are

Table 3.5: Mean reduction in the critical mortality rate obtained using QS-ReSTART as opposed to START, InvSTART, or QD-ReSTART (negative numbers indicate an increase in critical mortality).

		vs. START			vs. InvSTART			vs. QD-ReSTART		
		# Sig.			# Sig.			# Sig.		
	K	Mean	HW	Instances	Mean	HW	Instances	Mean	HW	Instances
Scenario 1	5	0.6%	0.1%	223	0.0%	0.0%	0	0.0%	0.0%	0
	10	0.5%	0.1%	223	0.0%	0.0%	0	0.0%	0.0%	0
	15	0.4%	0.0%	223	0.0%	0.0%	0	0.0%	0.0%	0
Scenario 2	5	2.8%	0.2%	465	0.0%	0.0%	0	0.0%	0.0%	0
	10	1.9%	0.1%	465	0.0%	0.0%	0	-0.1%	0.0%	254
	15	1.1%	0.1%	465	0.0%	0.0%	0	-0.1%	0.0%	333
Scenario 3	5	6.0%	0.2%	499	0.0%	0.0%	0	-0.2%	0.0%	392
	10	1.5%	0.1%	498	0.0%	0.0%	0	-0.6%	0.0%	499
	15	0.0%	0.0%	55	0.2%	0.0%	230	-0.5%	0.0%	485
Scenario 4	5	6.2%	0.1%	500	0.0%	0.0%	0	-0.6%	0.0%	422
	10	0.0%	0.0%	32	0.7%	0.1%	339	-1.1%	0.0%	484
	15	0.0%	0.0%	0	2.3%	0.0%	500	-0.1%	0.0%	15
Scenario 5	5	0.9%	0.1%	283	0.2%	0.1%	89	-1.3%	0.0%	473
	10	0.0%	0.0%	0	3.3%	0.0%	500	0.0%	0.0%	0
	15	0.0%	0.0%	0	2.8%	0.0%	500	0.0%	0.0%	0

Table 3.6: Mean reduction in the critical mortality rate obtained using QD-ReSTART(0.5) and QS-ReSTART(0.5) instead of START on instances with perturbed reward function parameters.

		QD-ReSTART vs. START				QS-ReSTART vs. START			
				# Better	# Worse			# Better	# Worse
	K	Mean	HW	Instances	Instances	Mean	HW	Instances	Instances
Scenario 1	5	0.6%	0.1%	223	0	0.6%	0.1%	223	0
	10	0.5%	0.1%	223	0	0.5%	0.1%	223	0
	15	0.4%	0.0%	223	0	0.4%	0.0%	223	0
Scenario 2	5	2.8%	0.2%	465	0	2.8%	0.2%	465	0
	10	2.0%	0.1%	465	0	1.9%	0.1%	465	0
	15	1.2%	0.1%	465	0	1.1%	0.1%	465	0
Scenario 3	5	6.2%	0.2%	499	0	6.0%	0.2%	499	0
	10	2.1%	0.1%	499	0	1.6%	0.1%	486	3
	15	0.6%	0.0%	461	0	0.0%	0.0%	87	48
Scenario 4	5	6.8%	0.2%	500	0	6.2%	0.2%	500	0
	10	1.2%	0.1%	455	0	0.1%	0.1%	67	58
	15	0.1%	0.0%	36	0	0.0%	0.0%	0	0
Scenario 5	5	2.2%	0.1%	441	2	0.9%	0.2%	251	94
	10	0.0%	0.0%	4	0	0.0%	0.0%	0	0
	15	0.0%	0.0%	0	0	0.0%	0.0%	0	0

statistically better from those in which they are statistically worse, both at a significance level of 0.05.

From this sensitivity analysis, we observed that while a few of the perturbed instances resulted in an increase in critical mortality, the vast majority of instances still saw improvement using QD-ReSTART. Even the first quartile had non-negative improvement over both START and InvSTART in every scenario, and QD-ReSTART performed worse than START only in 2 [82] instances (out of the 7500 simulated) at the 0.05 significance level under Setting 1 [2]. QS-ReSTART has similar trends in average performance, but START performed better than QS-ReSTART in a larger number of instances (203 [281] of the 7500 simulated under Setting 1 [2]) at the 0.05 significance level. Hence, QS-ReSTART appears to be less robust to perturbations or changes in the reward functions than QD-ReSTART. This difference can be attributed to the fact that under QD-ReSTART, a small change in

the value of t_m due to imperfect information about reward functions will only lead to a correspondingly small change in t^* , which will not affect the policy significantly. On the other hand, under QS-ReSTART, a change in the value of t_m could lead to a complete reversal of the policy (from START to InvSTART or vice versa).

3.6 Adapting ReSTART to Changing Conditions on the Field

This section relaxes some of the fundamental assumptions that we made when building our mathematical and simulation models in the earlier sections. In particular, we make three structural changes to our original setup. First, since this chapter focuses on incidents where no extensive search-and-rescue effort is needed, we originally assumed that all patients are accounted for immediately. However, even when no significant time is needed to locate and prepare patients for transportation, there could be still some delays in having some of the patients ready for transport. Hence, in this section, we will consider the possibility that not all patients are available at time zero and thus there is a delay not only in having these patients available but also in knowing the total number of patients.

Second, we originally assumed that all patients are classified into the delayed and immediate categories correctly. In reality, triage is prone to errors. More specifically, there are two types of triage errors: *undertriage*, when a patient who should have been classified as immediate is classified as delayed, and *overtriage*, when a patient who should have been classified as delayed is classified as immediate. Taking these misclassification errors into the decision making process is actually a subtle issue because it is tightly related to the estimation of the survival probabilities. Because classification errors are common in triage, survival probability functions can (and should) be estimated by taking this fact into account explicitly. Nevertheless, in practice, classification errors can be even more significant than the normally anticipated levels and thus it is of interest to investigate their effect on the performance of prioritization policies. Thus, in this section, we will also consider the possibility that patients are misclassified as a result of triage.

Third, in connection with the second issue, our original model assumed that patients are not triaged again after time zero, which may be the case in practice due to lack of resources

or poor organization. However, because many from the emergency response community emphasize the importance of *retriage* and because our new set of relaxed conditions that allow misclassification make retriage a meaningful action, we now assume that patients will go through retriage some time after the start of the response effort.

The main difficulty that arises as a result of the explicit consideration of these new features is that there is a delay of information not only on the actual number of patients from each class but also on the total number of patients who will need to be transported. As one can observe from Figure 3.2, an incorrect estimate of the number of patients can change the specific prescription by ReSTART policies, which would possibly degrade their performance. In the remainder of this section, we discuss how one can use ReSTART policies in an adaptive way and test by means of a simulation study how their performance is affected by misclassification errors, delayed availability of patients, and retriage.

3.6.1 Adaptive QD-ReSTART and QS-ReSTART

Adaptive QD-ReSTART is essentially QD-ReSTART with policy parameters updated regularly or every time an event provides new information on the number of patients. This event could be the arrival of a new patient (which could change n_I , n_D , and $n_I + n_D$) or retriage (which could change n_I and n_D but not $n_I + n_D$). More specifically, adaptive QD-ReSTART uses the most up-to-date information on the number of patients in each class and determines which class should get a priority by following the QD-ReSTART description provided in Section 3.4 and using the τ value updated with respect to the current time t , which we call $\tau(t)$. In particular, we let $\tau(t) = t_m - t - n_D(t)\theta\phi/K$, where $n_i(t)$ is the number of patients categorized as class $i \in \{I, D\}$ at time t . (Note that $\tau(0)$ corresponds to τ in the description of QD-ReSTART provided in Section 3.4.) Thus, $\tau(t)$ changes with the number of delayed patients but not with the number of immediate patients. Similarly, one could obtain an adaptive version of QS-ReSTART as follows: at time t , use START if $n_I(t) + n_D(t) \leq K(t_m - t)/(\phi\theta)$ and use InvSTART otherwise.

3.6.2 Performance of Adaptive ReSTART policies

In this section, we report the results of a numerical study on the performance of the adaptive QD-ReSTART and QS-ReSTART policies. In this study, we used the same experimental setup as in Section 3.5.2 for the total number of patients, the number of ambulances, and the survival probability functions, which gave us the same 7500 problem instances. However, this time, we also considered the possibility that there is some delay in having some of the patients ready for transportation and that the triage classification is imperfect. To understand the effects of reevaluation, we also assumed that patients would go through triage once again at some point during the response effort.

We conducted two separate simulation studies. In Study 1, we assumed that all patients are available for transportation at time zero; however, each patient can be misclassified. Retrospective studies of mass-casualty incidents found that while the overtriage rate is typically high, the undertriage rate is usually very low. For example, Frykberg (2005) reports a range of 20–80% for overtriage, while undertriage is almost non-existent. To account for the relatively wide range of possible overtriage rates, we simulated each instance twice, once with a moderate overtriage probability for each delayed patient of 0.4, and another time with a high overtriage probability of 0.6. (We let p_O denote the overtriage probability.) We set the undertriage probability for each immediate patient to 0.05. For both cases, we assumed that patients who are still waiting for transportation 40 minutes after the incident go through retriage, which then reveals patients’ true classifications. The choice of 40 minutes as the retriage time is largely arbitrary. In practice, emergency responders are urged to conduct retriage to identify cases of overtriage and undertriage (Mistovich and Karren 2007), but START and similar policies lack a standard dictating when retriage should occur (Lerner 2008).

We considered four different policies, namely the adaptive versions of START, InvSTART, QD-ReSTART(0.5), and QS-ReSTART(0.5). Note that here START and InvSTART are also adaptive in the sense that they use the updated classifications as a result of retriage. This is somewhat different from the notion of being adaptive for QD-ReSTART and QS-ReSTART, under which not only the classification of patients but also the structure

of the policy may change due to retriage. Our results on the comparison of QD-ReSTART with START are summarized in Table 3.7 under the column labeled Study 1. Numbers given are the mean percentage improvement in the critical mortality rate when using the adaptive version of QD-ReSTART as opposed to START. Comparison with Table 3.3 reveals that with misclassification and retriage, the improvements by adaptive QD-ReSTART over START are slightly lower but they are still significant. In the interest of space, we do not present here our results on the comparison of QD-ReSTART with InvSTART, which provided similar conclusions. We also omit our results on the performance of QS-ReSTART under Study 1 because the performance does not change much compared to the case without misclassification and retriage, which is reported in Table 3.5. This is an expected result because QS-ReSTART depends only on the total number of patients, and hence, the policy structure does not change with misclassification errors.

In Study 2, we considered two cases where not all patients are available at time zero. There are many reasons why some patients might be unavailable for transportation—for example, if they are trapped and specialized equipment is needed to rescue them. In particular, under the *moderate-unavailability* case, each patient has a probability 0.6 of being available at time zero, a probability 0.2 of being available at $t = 20$ minutes, and a probability 0.2 of being available at $t = 40$ minutes. For the *high-unavailability* case, each patient has a probability 0.2 of being available at time zero, has a probability 0.4 of being available at $t = 30$ minutes, and a probability 0.4 of being available at $t = 60$ minutes. Initial triage is assumed to have been done immediately after each arrival and a retriage is carried out once all patients are available. We also assume a moderate overtriage probability of 0.4 and a low undertriage probability of 0.05. The decision maker does not have advance knowledge of how many patients there are in total and the times at which new patients are going to become available. Therefore, idling to wait for a higher priority patient is not considered as long as there are patients in need of transport.

Results on the comparison of QD-ReSTART and QS-ReSTART with START under Study 2 are provided in Table 3.7. These results show that while the performance improvement with adaptive ReSTART policies is smaller than in the case where all patients are available at time zero, it is still significant, especially in Scenarios 2, 3, and 4 with a low level

of resources. As expected, the performance improvement decreases when more patients are unavailable at time zero and they are unavailable for a longer period of time.

3.7 Discussion

In this chapter, we have demonstrated that it is possible to design a prioritization policy that takes into account the three main components of the mass-casualty triage problem (i.e., the size of the event, availability of resources, and dependence of survival probabilities on time) in a very simple way, and performs better—substantially at times—than the common practice (START) that largely ignores these components. In particular, using a fluid formulation, we identified characteristics of “good” resource-based prioritization policies, which led to a simple policy that we call ReSTART and its variations. Using realistic simulations with data from emergency medicine literature, we observed that these policies have the potential to improve the critical mortality rate over START.

Qualitatively, we have observed that it is best to follow START (i.e., to give priority to immediate patients at all times) when there are *few patients* compared with the number of available resources and best to follow InvSTART (i.e., to give priority to delayed patients at all times) when there are *many delayed patients*. Otherwise, it is best to use a policy that prioritizes immediate patients initially but switches to delayed patients at *some point in time*. ReSTART gives a precise description of this structure by quantifying what *few patients*, *many delayed patients*, and *some point in time* really mean. Even if practitioners do not follow this description exactly, they can still build another policy having a structure that is similar to that of ReSTART, perhaps by coming up with new definitions for what it means to have few patients or many delayed patients. In short, our analytical characterization can provide a broad outline for the type of policy that is expected to work well in practice.

Table 3.7: Mean and 95% half-width for reduction in the critical mortality rate obtained using adaptive ReSTART policies instead of START for different simulated scenarios.

[illegible]

Chapter 4

Resource Allocation in a Multiple Location Mass-Casualty Incident

In the case of multiple-location MCIs, resource dispatching decisions must be made within the context of triage. Specifically, the number of resources dispatched to each location, and the time at which they are dispatched, may depend on the distribution of casualties at each location and on the prioritization policy. In this chapter, we adapt the most important concept from Chapter 3.1, namely that criticality of a patient is captured via a survival probability function that depends on the patient's triage class and the time at which the patient is transported, to the multiple location case, and we use this model to analyze resource allocation decisions in that case. Because resource dispatching decisions are made within the context of triage, a complete solution to the problem must tell the emergency responders where to send the resources and which triage class should have priority. When these decisions are made together, we call the resulting policy centralized, because one decision maker makes both decisions. When the decisions are made sequentially (i.e., resources are first allocated and then patients are prioritized), we call the resulting policy decentralized, because prioritization of triage classes may be decided on the scene of each incident, after resources have been assigned.

We assume that the main difference between the locations is the efficiency with which they can use the resources. That is, the same resource, if allocated to one location, will serve patients at a faster rate than if it is allocated to another location. This difference in service rate may be due to a number of factors, such as the physical characteristics of the

locations (terrain, etc.) and the relative distance of the location to its assigned hospital. For example, the same number of ambulances can serve patients at a much faster rate from a location one mile from the hospital than from a location twenty miles away because they can make more frequent trips. As in Chapter 3, we consider only transportation resources, such as ambulances and Emergency Medical Technicians, as the limiting factor, and hence we do not model the treatment of the patients at the hospital.

4.1 Model Description

In this chapter, we assume that there are two time-critical patient classes at each location; namely, the *immediate* and *delayed* classes, which we denote by I and D , respectively. Although the results do not directly depend on the use of START classes, we interpret them in this manner to provide a more intuitive understanding of the problem.

We assume that the patients are located at one of several locations and denote the set of locations by \mathcal{J} . We assume that at each location $j \in \mathcal{J}$, all patients have identically distributed travel times. This does not mean that patients from all the locations must be transported to the same hospital, but it does mean that all patients at a given location must be transported to the same hospital. Because we are interested in assigning resources to the different locations, of key interest is the efficiency with which those resources may be used. Each location $j \in \mathcal{J}$ has a resource speedup factor of s_j , which is the relative rate at which a resource can transport patients from location j . By “relative rate,” we mean a given resource can transport s_j/s_k times as many patients per unit time if that resource is assigned to location j than if it is assigned to location k . Intuitively, we may think of the mean travel time from location j to its assigned hospital as being $1/s_j$.

Letting n_{ij} be the number of class- i patients at incident j for all $i \in \{I, D\}, j \in \mathcal{J}$, the

fluid relaxation that we are trying to solve is the following:

$$\begin{aligned}
& \max_{\mathbf{r}(t), t \in [0, \infty)} \sum_{i \in \{I, D\}} \sum_{j \in \mathcal{J}} \int_0^\infty r_{ij}(x) s_j f_i(x) dx \\
& \text{subject to} \quad \sum_{i \in \{I, D\}} \sum_{j \in \mathcal{J}} r_{ij}(t) \leq 1, \quad \forall t \in [0, \infty) \\
& \quad \int_0^\infty s_j r_{ij}(x) dx = n_{ij}, \quad \forall i \in \{I, D\}, j \in \mathcal{J}.
\end{aligned} \tag{4.1}$$

In the above problem, the decision is $r_{ij}(t)$, the fraction of the available resources devoted to class- i patients at location j at time t . Without loss of generality, we have normalized the problem so that one unit of resources can transport patients at location j at rate s_j , i.e., the mean travel time from location j is $1/s_j$. It turns out that we can treat each class–incident pair as if it were a separate patient class, which has an important implication as to the structure of the optimal solution.

Proposition 4.1. *The general formulation (4.1) is equivalent to the formulation (3.1) with one incident and $2|\mathcal{J}|$ patient classes.*

Corollary 4.1. *There is an optimal solution to (4.1) where at every time t , $r_{ij}(t) = 0$ or $r_{ij}(t) = 1$ for every $i \in \{I, D\}$, $j \in \mathcal{J}$.*

If we limit ourselves only to solutions having the structure proved in Corollary 4.1, then we can re-write the problem formulation as follows. Let $W(i, j)$ be the set of time points in $[0, T]$ where class- i patients at location j are served, where T is the time horizon needed to serve all patients (i.e., $T = \sum_{i \in \{I, D\}} \sum_{j \in \mathcal{J}} n_{ij}$). Let $\mathbf{W} = \{W(i, j), i \in \{I, D\}, j \in \mathcal{J}\}$.

$$\begin{aligned}
& \max_{\mathbf{W}} \sum_{i \in \{I, D\}} \sum_{j \in \mathcal{J}} \int_{W(i, j)} s_j f_i(x) dx \\
& \text{subject to} \quad \bigcup_{i \in \{I, D\}} \bigcup_{j \in \mathcal{J}} W(i, j) = [0, T], \\
& \quad W(i, j) \cap W(k, l) = \emptyset \quad \forall (k, l) \neq (i, j) \\
& \quad \int_{W(i, j)} s_j dx = n_{ij}, \quad \forall i \in \{I, D\}, j \in \mathcal{J}.
\end{aligned} \tag{4.2}$$

Corollary 4.1 demonstrates that an optimal solution to (4.2) will also be an optimal solution to (4.1). In the remainder of the analytical results, we seek an optimal solution to (4.2), and moreover we only admit solutions that have the following structure: the set of time points where a given class of patients is transported is the union of finitely many intervals, each of which is closed on the left and open on the right. This limitation is useful both for technical reasons (so that the set of time points where we transport a given class does not include a collection of zero-measure points that are not adjacent to intervals in which we transport patients) and for practical reasons (because it would be impossible to implement a policy that changed resource allocation infinitely many times). In other words, the policy we seek maps time intervals to patient classes. The choice of closed on the left and open on the right does not affect the nature of the solution, but it provides consistency within proofs.

The type of solution that allocates all resources to a single location at any given time may not be desirable or may be deemed to be unfair because one of the locations may be “starved.” However, note that while only one location is served at any given time, this solution structure does not necessarily imply a fixed ordering between the locations. That is, in the optimal solution, it is possible that the resources are moved back and forth between the locations. On the other hand, among patients in the same class, there is a distinct relationship between the different locations, which we now establish formally.

Proposition 4.2. *For a given class of patients, it is always advantageous to prioritize patients in that class at a location with a larger speedup factor over patients in that class at a location with a lower speedup factor. That is, in an optimal solution, class i patients at incident j will all be transported before any class i patients at incident l if $s_j > s_l$.*

Proposition 4.2 confirms the intuition that if patients are identical, it is advantageous to assign resources to the location that can use them more efficiently. Proposition 4.2 also proves that the optimal policy has a relatively simple structure, because at any given point in time, many of the patient classes will be dominated by patients of the same class at a location with a faster speedup factor.

4.2 Analysis of Two-Incident Problem

The simplest case of simultaneous MCIs is when there are two incidents. The two-incident case also has the advantage of being easier to analyze than the general case, and it can provide us with some insight into the types of policies that are likely to work well with more than two incidents.

Let \underline{I} and \underline{D} denote the immediate and delayed classes at the incident with the smaller speedup factor. Let I and D denote the immediate and delayed classes at the incident with the larger speedup factor. For ease of notation we assume that the parameters of the problem are scaled so that the speedup factor at the first location is 1 and the speedup factor at the second location is $s > 1$.

We assume that the following structure is present in the reward functions: immediate patients deteriorate faster at the beginning of the response effort, and delayed patients deteriorate faster later on, after some point in time. Formally, the reward functions of classes I and D satisfy the following:

Assumption 4.1. *There exist times t_m , \underline{t}_m , and \bar{t}_m such that $0 \leq \underline{t}_m \leq t_m \leq \bar{t}_m$, and*

- $f'_I(t) < f'_D(t) < 0$ for all $t < t_m$, and $f'_D(t) < f'_I(t) < 0$ for all $t > t_m$.
- $sf'_I(t) < f'_D(t) < 0$ for all $t < \bar{t}_m$, and $f'_D(t) < sf'_I(t) < 0$ for all $t > \bar{t}_m$.
- $f'_I(t) < sf'_D(t) < 0$ for all $t < \underline{t}_m$, and $sf'_D(t) < f'_I(t) < 0$ for all $t > \underline{t}_m$.

Assumption 4.1 states that initially, the reward from transporting immediate patients decreases faster than the reward from transporting delayed patients, but after a certain amount of time, the reverse is true. One sufficient (but not necessary) condition for Assumption 4.1 would be for $f_D(t)$ to be concave decreasing and $f_I(t)$ to be convex decreasing over the time horizon. In other words, as time goes by, the rate at which delayed patients deteriorate increases, while the rate at which immediate patients deteriorate decreases. This makes intuitive sense because the delayed patients are initially stable and begin to worsen more rapidly later in the response effort, while immediate patients lose much of their ability to survive early on in the response effort.

Recall that as a result of Proposition 4.2, we know that we should give class I priority over class \underline{I} and class D priority over class \underline{D} at all times. The next proposition further characterizes the structure of the optimal policy that we seek:

Proposition 4.3. *There is an optimal solution where once the transportation of class D or \underline{D} patients begins, it does not end until all of those patients have been transported.*

In other words, once class D or \underline{D} becomes the highest priority among the remaining classes, it remains the highest priority. Therefore, we are interested in characterizing the time t^* , the time when we would begin transporting patients from class D , and time t^{**} , the time when we would begin transporting patients from class \underline{D} . In the rest of this chapter, we limit our search for optimal solutions of the type described in Proposition 4.3.

The next two sections show how to calculate an interval in which t^* must lie. First, we prove the following lemma that will be useful in this analysis.

Lemma 4.1. *Let $g(t) : \mathbb{R}^+ \rightarrow \mathbb{R}$ be a function that satisfies the following: there exists $t_m > 0$ such that $g'(t) > 0 \forall t < t_m$ and $g'(t) < 0 \forall t > t_m$. Let c_1 and c_2 be constants such that $0 < c_2 \leq c_1$. For any $t_1, t_2 \geq 0$, which satisfy*

$$g(t_1) \leq g(t_1 + c_1) \tag{4.3}$$

and

$$g(t_2) \geq g(t_2 + c_2), \tag{4.4}$$

it must be the case that $t_1 \leq t_2$.

4.2.1 Upper and lower bounds for t^*

In our patient prioritization problem, both classes of delayed patients (D and \underline{D}) will compete for resources with the immediate patients. That is, when determining the time at which to switch priority to delayed patients, we must consider two different classes of delayed patients. On the other hand, if we consider our patient prioritization problem *without* the delayed patients at the incident with the smaller speedup factor; i.e., with $n_{\underline{D}} = 0$, then only the delayed patients at the incident with the larger speedup factor (i.e., those in class

D) will compete for resources with the immediate patients. This means that the delayed patients may be able to afford to wait a bit longer before receiving priority compared to the case where $n_{\underline{D}} > 0$. This is the intuition behind the upper bound problem, which we now define formally.

Define the “upper bound problem” to be the following one-location instance of (4.2):

$$\begin{aligned}
& \max_{\mathbf{W}} \quad \sum_{i \in \{1,2\}} \int_{W(i)} f_i(x) dx & (4.5) \\
& \text{subject to} \quad \bigcup_{i \in \{1,2\}} W(i) = [0, T], \\
& \quad \quad \quad W(i) \cap W(j) = \emptyset \quad \forall j \neq i \\
& \quad \quad \quad \int_{W(i)} dx = n_i, \quad \forall i \in \{1, 2\},
\end{aligned}$$

where $f_1(t) = sf_I(t)$, $f_2(t) = sf_D(t)$, $n_1 = n_I/s + n_{\underline{I}}$, and $n_2 = n_D/s$. Note that in the above formulation, we name the two patient classes 1 and 2 to make it clear that they do not refer to the classes in the original problem, and $W(i)$ is the set of time points where we transport class i patients, $i \in \{1, 2\}$.

We can prove that in the optimal solution to the upper bound problem, the time at which we begin transporting class 2 patients is an upper bound to t^* , which leads to the following result.

Proposition 4.4. *Let W^* be any optimal solution to (4.2), and let $t^* = \inf\{W^*(D)\}$. Then $t^* \leq \min\{n_I/s + n_{\underline{I}}, t_m\}$. That is, we should begin transporting class D patients no later than time t_m or when all immediate patients have been transported.*

The argument used in the proof of Proposition 4.4 can also be used under a certain condition to give another upper bound for t^* :

Corollary 4.2. *If $\underline{t}_m \leq n_I/s$, then $t^* \leq n_I/s$.*

Because $\underline{t}_m \leq t_m$, in the case when $\underline{t}_m \leq n_I/s \leq t_m$, Corollary 4.2 will give a tighter upper bound than Proposition 4.4. As a side note, Corollary 4.2 also demonstrates that $\underline{t}_m \leq n_I/s$ is a sufficient (but not necessary) condition for class D to dominate class \underline{I} .

Now, to determine a lower bound for t^* , consider what happens if we group the patients from our problem at one location as follows: assume that all patients from class D and \underline{D} have reward $sf_D(t)$, and all the patients from class I and \underline{I} have reward $f_I(t)$. This grouping means that delayed patients are even more attractive in terms of reward than in the two-incident problem, because they are all able to take advantage of the larger reward, while all the immediate patients take the small reward. Therefore, intuitively it would be advantageous to switch to delayed patients even earlier than we would in the two-incident problem.

It turns out that this intuition is true. Define the “lower bound problem” to be (4.5) with $f_1(t) = f_I(t)$, $f_2(t) = sf_D(t)$, $n_1 = n_I/s + n_{\underline{I}}$, and $n_2 = n_D/s + n_{\underline{D}}$. Analyzing the lower bound problem yields the following lower bound for t^* .

Proposition 4.5. *We have $t^* \geq \min\{n_I/s + n_{\underline{I}}, \underline{t}_m - (n_D/s + n_{\underline{D}})\}$. That is, we should begin transporting class D patients no earlier than time $\underline{t}_m - (n_D/s + n_{\underline{D}})$, unless all immediate patients have been transported.*

Taken together, Propositions 4.4 and 4.5 and Corollary 4.2 suggest that the time to begin transportation of class D patients lies in the interval $[\underline{t}_m - (n_D/s + n_{\underline{D}}), t_m]$, or in the case where $\underline{t}_m \leq n_I/s$, the time lies in the interval $[\underline{t}_m - (n_D/s + n_{\underline{D}}), \min\{t_m, n_I/s\}]$. In either case, if all immediate patients have been transported, then we should begin transporting class D patients regardless of whether we have reached this interval.

4.2.2 Upper and lower bounds for t^{**}

In the previous section, we identified an interval in which t^* , the time at which we should begin transporting class D patients, must lie. We now approach the related question of determining an interval in which t^{**} , the time at which we should begin transporting class \underline{D} patients, must lie. We will take a similar approach of solving a single-incident version of the prioritization problem. We will develop these results under the assumption that t^* is known; conveniently, it turns out that the upper and lower bounds for t^{**} depend on t^* only through the priority relationship between class D and class \underline{D} , which we already established in Proposition 4.2. In other words, even though we assume that we have t^* in

obtaining our analytical results for t^{**} , one does not need to know t^* to determine upper and lower bounds to t^{**} as long as one understands that class D patients will always have priority over class \underline{D} patients.

In the fluid model, we finish transporting class D patients at time $t^* + n_D/s$. There will be $n_I - \min\{n_I, st^*\}$ class I patients remaining at that time, because prior to t^* , we will serve class I patients (at rate s per unit time) as long as any are present. Moreover, there will be $n_{\underline{I}} - \max\{0, t^* - n_I/s\}$ class \underline{I} patients remaining, because we will serve class \underline{I} patients only after all class I patients have been served, which takes n_I/s time units. Note that if the number of class \underline{I} patients remaining is zero, then $t^{**} = t^* + n_D/s$. In other words, the only remaining patients are in class \underline{D} and thus they should be transported at that time.

Define the upper bound problem for t^{**} to be the following:

$$\begin{aligned} \max_{\mathbf{W}} \quad & \sum_{i \in \{1,2\}} \int_{W(i)} f_i(x) dx \\ \text{subject to} \quad & W(1) \cup W(2) = [t^* + n_D/s, T], \\ & W(1) \cap W(2) = \emptyset \\ & \int_{W(i)} dx = n_i, \quad \forall i \in \{1,2\}, \end{aligned} \tag{4.6}$$

where $f_1(t) = sf_I(t)$, $f_2(t) = f_D(t)$, $n_1 = n_{\underline{I}} + n_I/s - t^*$, and $n_2 = n_{\underline{D}}$.

Proposition 4.6. *We have $t^{**} \leq \max\{\bar{t}_m, t^* + n_D/s\}$. That is, we should begin transporting class \underline{D} patients no later than time \bar{t}_m , unless class D patients still remain to be transported at that time.*

Define the lower bound problem for t^{**} to be (4.6) with $f_1(t) = f_I(t)$, $f_2(t) = f_D(t)$, $n_1 = n_{\underline{I}} + n_I/s - t^*$, and $n_2 = n_{\underline{D}}$.

Proposition 4.7. *We have $t^{**} \geq \min\{n_D/s + n_I/s + n_{\underline{I}}, t_m - n_{\underline{D}}\}$. That is, we should not begin transporting class \underline{D} patients until at least time $t_m - n_{\underline{D}}$, unless no other patients remain in the system.*

Propositions 4.6 and 4.7 imply that the optimal time to begin transporting class \underline{D} patients lies in the interval $[t_m - n_{\underline{D}}, \bar{t}_m]$, with the two exceptions that class \underline{D} patients should never be transported before class D patients (the max in Proposition 4.6) and we should never wait to transport class \underline{D} patients if it would cause idling (the min in Proposition 4.7).

4.2.3 Application of Lower and Upper Bounds

In this section, we discuss the implications of the results in Sections 4.2.1 and 4.2.2 to the two-incident mass-casualty triage problem. It is important to recognize that the most important result in these sections is structural. The optimal policy for the model begins by prioritizing immediate patients. Within the immediate class, those patients at the location with the larger speedup factor should be prioritized. After a certain amount of time (namely, at t^*), priority should switch to delayed patients at the location with the larger speedup factor. Following transportation of those patients, we should switch back to immediates until time t^{**} , when we should prioritize the delayed patients at the location with the smaller speedup factor.

Recall that earlier we scaled the prioritization problem by the inverse of the transportation rate (which is the number of transportation resources divided by the mean transportation time) so that it was equal to one per unit time. If we allow for general scaling, the analytical results state that

$$t^* \in [\underline{t}_m - \theta/K (n_D/s + n_{\underline{D}}), t_m], \quad (4.7)$$

$$t^* \in [\underline{t}_m - (n_D/s + n_{\underline{D}}), \min\{t_m, n_I\theta/(Ks)\}], \text{ if } \underline{t}_m \leq n_I\theta/(Ks), \quad (4.8)$$

and

$$t^{**} \in [t_m - n_{\underline{D}}\theta/K, \bar{t}_m], \quad (4.9)$$

where θ is the mean travel time and K is the number of resources. Because calculating t^* and t^{**} exactly requires solving a nonlinear optimization problem, it is unlikely that such exact calculations would be practical in an actual mass-casualty incident. However, calculating

(4.7), (4.8), and (4.9) is fairly simple, as it does not require any computation. Thus, the intervals can be used as guidelines for the emergency planner or incident commander. In Section 4.4 we discuss how this can be done effectively.

4.2.4 Decentralized Resource Allocation

The analysis we have conducted for the two-incident problem so far has relied on the fact that the resources can be reassigned from one incident to the other without delay. Although the structure of the optimal solution may involve switching the resources between the two locations at most twice, if the incidents are far apart or in the case of infrastructure damage, switching the resources may be costly and time consuming. Thus, it is desirable to have policies for decentralized resource allocation. We use the term decentralized to mean that prioritization decisions are not made by a central authority in concert with the resource allocation decisions. In a decentralized policy, resources are first allocated to the two incidents by the central authority, and they remain tied to their assignment regardless of the numbers or types of patients remaining at each incident over time. A manager at each incident then prioritizes patients independently of the other.

Decentralized prioritization policies might be attractive in the case of a mass-casualty incident if there are high overheads involved with using a centralized policy; for example, when there is a shortage of dispatchers or coordinators, or when it is time-consuming to move resources back and forth between the two locations. Another motivation for using a decentralized policy (specifically, one that divides the resources between the two locations) is that emergency planners may believe that it is unfair or unethical to devote all available resources to a single location at any given time, or they may believe that doing so would expose them to legal liabilities. As we mentioned in the introduction to Chapter 3, determining the ethicality of a given triage policy is beyond the scope of this dissertation. Our main goal is to show that there are decentralized policies that would be practical to implement, and in many cases can be expected to perform nearly as well as the centralized policy that is optimal for the fluid model. It remains up to the decision makers to determine whether the additional benefits of implementing a centralized policy (such as the optimal policy for our model) outweigh the costs of doing so.

The first type of decentralized policy we consider is sequential resource allocation, or in other words, fixed priority per location. From Proposition 4.2, we already know that within a given class of patients, it is desirable to prioritize the locations in order of decreasing speedup factor. A natural heuristic extension is to prioritize locations in order of decreasing speedup factor, without respect to the number and types of patients at each location. In other words, the location with the largest speedup factor would be assigned exclusive use of the resources until all of its critical patients have been transported; then, the location with the next-largest speedup factor would receive use of the resources, and so on. While a given location has use of all of the resources, patients could be prioritized using any appropriate means, such as START, ReSTART, QD-ReSTART, or QS-ReSTART.

The second type of decentralized policy is simultaneous resource allocation. In this case, the centralized emergency planner makes a single decision of how to divide the resources between the incidents at the beginning of the response. The resources then remain tied to their assigned incident for the entirety of the response effort. Again, with the resources that are provided to each location, patients could be prioritized using any appropriate means, such as START, ReSTART, QD-ReSTART, or QS-ReSTART. We can examine the simultaneous resource allocation policy using the fluid formulation for the two-incident case considered in Section 4.1. Suppose that we allocate fraction $\alpha \in (0, 1)$ of the resources to the incident with the smaller speedup factor, and fraction $1 - \alpha$ to the other incident, and that each location independently uses START or QD-ReSTART(0.5). For simplicity, we will assume that this resource allocation must be fixed in advance and cannot be changed during the response effort, even when one of the locations no longer has critical patients to serve. This is a reasonable assumption when the resources are assigned to locations; for example, they would begin serving patients classified as *minor* at their assigned location, even if critical patients remain at the other location. Then, the total reward earned by the

system under the fluid model is given by

$$\begin{aligned}
S(\alpha) = & \alpha \int_0^{\frac{n_I}{\alpha}} f_I(t) dt + \alpha \int_{\frac{n_I}{\alpha}}^{\frac{n_I+n_D}{\alpha}} f_D(t) dt \\
& + s(1-\alpha) \int_0^{\frac{n_I}{s(1-\alpha)}} f_I(t) dt + s(1-\alpha) \int_{\frac{n_I}{s(1-\alpha)}}^{\frac{n_I+n_D}{s(1-\alpha)}} f_D(t) dt,
\end{aligned} \tag{4.10}$$

in the case where START is used.

It is very difficult analytically to find the value of α maximizing $S(\alpha)$ because of the way it depends on the reward functions $f_I(t)$ and $f_D(t)$. In our simulations, we will calculate the optimal value of α numerically for the sake of comparison; however, we also need an easier way to determine α that has a reasonable expectation of working well. Therefore, we propose the *Equal Critical Service Time* (ECST) heuristic, which is to set α to

$$\alpha_0 = \frac{s(n_I + n_D)}{s(n_I + n_D) + n_I + n_D}, \tag{4.11}$$

which results in

$$\frac{n_I + n_D}{\alpha_0} = \frac{n_I + n_D}{s(1-\alpha_0)}.$$

In other words, the resources are assigned so that both incidents require the same amount of time to transport all their critical patients (i.e., excluding minor and expectant patients), which implies that the average waiting time for a randomly selected critical patient does not depend on his or her location.

4.3 Analysis for More than Two Incidents

Although the analysis becomes somewhat more complicated when there are more than two incidents because there are more decisions, we can extend some of the results of Section 4.2 to the same problem with more than two incidents. Suppose that there are k incidents, $\{1, 2, \dots, k\}$, and that for all $i \in \{1, \dots, k\}$, location i has n_{Di} delayed patients with instantaneous reward function $f_D(t)$, n_{Ii} immediate patients with instantaneous reward function $f_I(t)$, and resource speedup factor s_i such that $1 = s_1 < s_2 < \dots < s_k$. Note that as in

the two location problem, different locations do not necessarily send patients to the same hospital. The speedup factor s_i indicates the efficiency with which location i can use its resources to send patients to its assigned hospital.

Remark 4.1. By extending Proposition 4.2, it is the case that for all $i \in \{2, \dots, k\}$, immediate (delayed) patients at location i dominate immediate (delayed) patients at location $i - 1$. Hence, at any point in time, there are only two classes of patients, immediate and delayed, each at the location with the largest possible subscript (but not necessarily the same location), that could be eligible for service.

In a manner similar to Section 4.2, we can give upper and lower bounds for $t^{(j)}$, the time at which we should begin serving the delayed patients at location j , by solving an upper bound problem. For the upper bound, we must solve (4.5) with $f_1(t) = f_I(t)$, $f_2(t) = f_D(t)$, $n_1 = \sum_{i=1}^j n_{Ii}/s_i$, and $n_2 = n_{Dj}/s_j$; for the lower bound, we must solve (4.5) with $f_1(t) = f_I(t)$, $f_2(t) = s_j f_D(t)$, $n_1 = \sum_{i=1}^j n_{Ii}/s_i$, and $n_2 = \sum_{i=1}^j n_{Di}/s_i$.

Similarly to Proposition 4.2, if we define $t^{(j,1)} = \arg \max_t (s_j f_D(t) - f_I(t))$, then if $t^{(j,1)} \leq n_{Ij}/s_j$, we have $t^{(j)} \leq n_{Ij}/s_j$. We omit the proofs of these bounds, as they follow the same reasoning as the proofs of Propositions 4.4 and 4.5, and of Corollary 4.2.

Once patients remain at only $k - 1$ locations, the multi-location problem may be resolved with one fewer location. Clearly, while this method provides a centralized policy, it will become less useful as k increases.

We can also extend the ECST heuristic introduced in Section 4.2.4 to the case where there are more than two locations. In this case, we would choose to allocate fraction α_j of the resources to each location j such that

$$\frac{n_{I1} + n_{D1}}{\alpha_1} = \frac{n_{I2} + n_{D2}}{s_2 \alpha_2} = \dots = \frac{n_{Ik} + n_{Dk}}{s_k \alpha_k}$$

and $\sum_i \alpha_i = 1$. Each location can then use its own resources independently, and the average amount of time that a critical patient waits is the same at each location.

4.4 Heuristic Policies and Simulation Study

In Section 4.2, we noted that while it is not necessarily possible to quickly determine the two thresholds needed to completely determine the optimal policy to the two-location fluid model, we are able to give upper and lower bounds on these thresholds, which can be used to implement a centralized priority policy. In Section 4.2.3, we proved that t^* and t^{**} must lie in certain intervals. Could using these intervals really provide a good approximation to the optimal policy? To explore this question, we conducted a numerical study on the two-incident problem. We generated 1,000 instances, each with the number of patients uniformly drawn from $[20, 150]$, using the reward function and patient class distribution given by Scenario 3 in Figure 3.3. For each instance, we calculated t^* and t^{**} using the Matlab optimization toolbox. We then computed the intervals bounding these values according to Section 4.2. Using these data points, we calculated ψ^* , which we define as the location of t^* within the interval given in (4.7) or (4.8) as appropriate. Specifically, if we let l^* denote the lower bound for t^* and u^* denote the upper bound for t^* , then ψ^* is the unique solution in $[0, 1]$ to the equation $t^* = l^* + \psi^*(u^* - l^*)$. Similarly, we calculated ψ^{**} , the location of t^{**} within the interval (4.9). Distributions of ψ^* and ψ^{**} for this numerical study are shown in Figure 4.1. For the wide range of instances generated in the numerical study, t^* most frequently lies near the middle of its interval and t^{**} generally lies near its lower bound. While these observations will be dependent on the survival probability functions, it should be noted that in many cases the intervals themselves are not extremely large, and so any choice from within the interval may be reasonable. Moreover, emergency coordinators can “tune” their selection based on how aggressively they wish to treat delayed versus immediate patients. In our simulations, we simply use the endpoints or midpoints of the intervals, which can be determined by a hand calculation, and as we will demonstrate in section 4.4.3, this heuristic performs surprisingly well.

As we noted previously, using a centralized policy may not be desirable in all situations. Therefore, using insights from the results of Sections 4.2, we also develop two decentralized policies that may be expected to perform well (similarly to a centralized policy), especially when s is large or close to 1. We compare all the heuristic policies to a decentralized version




Figure 4.1: Histograms for ψ^* (left) and ψ^{**} (right) from the numerical experiment.

of the commonly used START protocol.

4.4.1 Policy Descriptions

We determined in Section 4.2 that a centralized policy (that is, a policy in which resources are directed to serve the locations and patient classes by a centralized decision maker) will be optimal under the fluid model. In particular, using the relationship from Proposition 4.2 and the threshold times given in Section 4.2, each point in time during the MCI response effort maps directly to a location-class pair. Hence, we expect a centralized policy to perform very well in practice, even though patients and resources are actually discrete. However, estimating the two thresholds is a non-trivial sub-problem that must be solved in order to implement such a policy. We will test two different heuristics: one where t^* and t^{**} are calculated exactly (called Centralized-Exact), and one where they are estimated by a point within the intervals given in Section 4.2 (called Centralized-Est). The Centralized-Est heuristic is parameterized by ψ , the quantile of the interval that is used for the estimate. We will use $\psi \in \{0.0, 0.5, 1.0\}$, which corresponds to lower bounds, midpoints, and upper bounds.

On the other hand, a centralized policy may not be desirable in practice. We also test decentralized policies to see which ones would be expected to work well. As we discussed before, two types of decentralized policies should be reasonable. One heuristic serves the locations sequentially, in order of decreasing speedup factor (called Decentralized-Seq). In Decentralized-Seq, we use QD-ReSTART(0.5) to prioritize patients at each location. The other possibility is to serve the locations simultaneously. In this case, the resources must be allocated in some manner between the locations. We test the heuristic that applies QD-ReSTART(0.5) with resources allocated according to the ECST heuristic (called

Decentralized-Simul-R) and the heuristic that applies START with the fluid optimal allocation (called Decentralized-Simul-S). The exact details of all the heuristic policies are given in Table 4.1.

Based on the fact that QD-ReSTART is designed to be an improvement on START, we expect Decentralized-Simul-S may perform worse than Decentralized-Simul-R; note that Decentralized-Simul-S is included for purposes of comparison because it is the type of policy that is likely to be implemented currently, since many EMS agencies use START. However, note that even Decentralized-Simul-S has an advantage over what could be used in practice because we are numerically calculating the fluid optimal allocation, which is unlikely to be possible in the aftermath of an MCI. On the other hand, based on the results of Section 4.2, we expect the Centralized-Exact policy to perform very well, because it is the optimal policy for the fluid model. Although we cannot guarantee that Centralized-Exact is the optimal policy for the simulation model, it is unlikely that any policy that is easy to implement would perform significantly better. Hence, we will also compare the other policies to Centralized-Exact, because a policy that can do as well, or nearly as well, as the Centralized-Exact policy would be an excellent candidate for implementation.

4.4.2 Simulation Method

To test the heuristic policies, we conducted a series of simulations using code written in Matlab. We used the same five survival probability scenarios given in Figure 3.3, which include survival probability functions for immediate and delayed classes, as well as distribution of the patient classes. For each scenario, we generated 200 instances, each with 100 total patients. For each instance, we generated the number of patients at the location with the larger speedup factor uniformly over $\{1, 2, \dots, 99\}$ (ensuring there would always be at least one patient at each location); the remainder of the patients were assigned to the other location. Each patient was randomly assigned one of the four START classes according to the patient class distribution. Finally, the number of ambulances was randomly selected from $\{5, 6, \dots, 15\}$. We chose not to separate the results by number of ambulances because preliminary runs showed no appreciable effect of the number of ambulances on which policy performed the best. Nevertheless, random selection of the number of ambulances ensured

Table 4.1: Summary of heuristic policies for two-location and multi-location MCI response.

Policy Type	Heuristic	Description (Two-Location)	Description (Multi-Location)
Centralized	Exact	Follow the optimal policy for the fluid model, computing thresholds t^* and t^{**} exactly via nonlinear optimization or line search.	Follow the optimal policy for the fluid model, computing thresholds exactly via nonlinear optimization or line search.
Centralized	$\text{Est}(\psi)$	Follow the optimal policy for the fluid model, but estimate t^* and t^{**} by $l^* + \psi(u^* - l^*)$.	Follow the optimal policy for the fluid model, but estimate all thresholds using the point ψ -quantile of the interval estimate.
Decentralized	Seq	Serve the location with the larger speedup factor first, using QD-ReSTART at each location.	Serve the locations in order of decreasing speedup factor, using QD-ReSTART at each location.
Decentralized	Simul-R	Use the ECST heuristic, i.e., allocate fraction α_0 of the resources to the slower location and $(1 - \alpha_0)$ to the faster location, rounding in favor of the faster location when resources are discrete. At each location, independently apply QD-ReSTART.	Use the ECST heuristic, i.e., allocate fraction α_i of the resources to location L_i , for all $i \in \{1, 2, \dots, k\}$, with any remainder (due to discreteness) going to location k . At each location, independently apply QD-ReSTART.
Decentralized	Simul-S	Use the allocation α^* of the resources to the slower location and $(1 - \alpha^*)$ to the faster location, where α^* is the maximizer of $S(\alpha)$, rounding in favor of the faster location when resources are discrete. At each location, independently apply START.	Use the ECST heuristic, i.e., allocate fraction α_i of the resources to location L_i , for all $i \in \{1, 2, \dots, k\}$, with any remainder (due to discreteness) going to location k . At each location, independently apply START.

scenarios with a variety of levels of resource availability.

4.4.3 Simulation Results

The simulation results for the heuristics defined in Table 4.1 are presented in Table 4.2. The table presents the mean decrease in critical mortality compared to Decentralized-Simul-S. Half-widths for each of these differences were 0.1% or smaller. When interpreting the table, note that our baseline policy, Decentralized-Simul-S, has the advantage that we numerically calculated the optimal resource allocation in the fluid model before running the simulation. The other simultaneous policy, Decentralized-Simul-R, does not have this advantage—rather, we use the ECST heuristic, which is more practical. In reality, it is unlikely that emergency responders would be computing the optimal resource allocation when using START; therefore, the performance of START may actually be worse in practice, which would increase the relative performance of the other policies. It is also important to note that we do not claim that Centralized-Exact is the optimal policy. However, Centralized-Exact is a useful point of comparison for two reasons: first, it is the optimal solution to the fluid model, and second, it would be the most complicated of the heuristics to implement.

Four main observations appear from the simulation results. The first observation is that using START independently at each location, even with the optimal resource allocation, performs poorly on average compared to the centralized policies for all values of s , indicating that there is potentially large room for improvement over current EMS practices.

The second observation is that with a decentralized policy, it is still possible to achieve a reasonably good improvement, either by using Decentralized-Simul-R (which appears to work better in the case of lower values of s , i.e., when the resources can be used with similar efficiency at both locations) or by using the Decentralized-Seq (which works very well, especially in the case of higher values of s).

The third observation is that in scenarios where t_m is close to zero, namely scenarios 1 and 2, which are the most pessimistic scenarios in terms of survival probability, Decentralized-Seq performs very similarly to the centralized policies. In these scenarios, the centralized policies end up prescribing the sequential policy in many instances.

Table 4.2: Mean decrease in critical mortality over 200 instances compared to Decentralized-Simul-S with fluid optimal resource allocation.

Scenario	s	Centralized				Decentralized	
		Exact	Est(0.0)	Est(1.0)	Est(0.5)	Seq	Simul-R
1	1.2	0.9%	0.9%	0.9%	0.9%	0.7%	0.7%
	1.5	1.1%	1.1%	1.1%	1.1%	1.0%	0.7%
	2.0	1.3%	1.3%	1.3%	1.3%	1.3%	0.7%
	3.0	1.8%	1.8%	1.8%	1.8%	1.8%	0.6%
	5.0	2.3%	2.3%	2.3%	2.3%	2.3%	0.4%
2	1.2	3.1%	3.1%	3.1%	3.1%	2.6%	2.7%
	1.5	3.7%	3.6%	3.6%	3.7%	3.3%	2.8%
	2.0	4.2%	4.0%	4.1%	4.1%	3.9%	2.6%
	3.0	5.1%	4.9%	5.0%	5.1%	5.0%	2.4%
	5.0	6.4%	6.2%	6.2%	6.4%	6.4%	2.2%
3	1.2	6.6%	6.3%	6.3%	6.6%	6.1%	5.7%
	1.5	7.5%	7.2%	7.3%	7.5%	7.2%	5.7%
	2.0	8.5%	8.0%	8.2%	8.4%	8.4%	5.3%
	3.0	10.2%	9.5%	9.7%	10.1%	10.2%	4.9%
	5.0	12.5%	11.6%	11.3%	12.1%	12.4%	4.9%
4	1.2	7.7%	6.9%	6.8%	7.6%	7.2%	5.9%
	1.5	10.1%	8.9%	9.3%	9.9%	9.7%	6.3%
	2.0	11.8%	10.1%	10.7%	11.5%	11.7%	5.6%
	3.0	14.4%	12.5%	12.7%	13.9%	14.4%	4.8%
	5.0	18.0%	16.1%	15.9%	16.7%	17.9%	4.2%
5	1.2	5.1%	3.8%	4.2%	5.0%	4.9%	2.6%
	1.5	7.0%	5.2%	6.1%	6.8%	6.9%	2.2%
	2.0	9.3%	7.5%	8.4%	9.0%	9.2%	1.5%
	3.0	12.9%	11.3%	12.0%	12.4%	13.0%	1.1%
	5.0	16.0%	14.8%	14.9%	15.5%	15.9%	0.1%

The final observation is that all the centralized policies perform very similarly; in particular, Centralized-Est(0.5) results in almost the same improvement as Centralized-Exact. We conclude that if it is possible to use a centralized policy, the lower and upper bounds presented in Section 4.2 are very useful, because they can be used to quickly approximate the thresholds without sacrificing much performance compared to finding the exact thresholds.

4.5 Discussion

In this chapter, we demonstrated that current triage protocols (such as START) may not be sufficient to deal with the problem of resource allocation in a two-location mass-casualty incident. However, the type of formulation used in Chapter 3 to develop the ReSTART policy also leads to insights that will be useful in making resource allocations in a multiple-location mass-casualty incident. Specifically, under the fluid formulation, the optimal policy devotes all available resources to one class-location pair at any time, switching between them according to two thresholds. Within each patient class, patients at the location with the larger speedup factor should be prioritized. While calculating the exact threshold times for beginning the transportation of the delayed patients is computationally more difficult than calculating the single threshold for ReSTART, quick approximations of these thresholds nevertheless perform very well in simulations. Moreover, if a policy with centrally dispatched resources is not feasible or not desirable, we have developed decentralized heuristics that can potentially still result in significant decreases in critical mortality compared to using START, even if START could be used with the fluid optimal resource allocation, which is unlikely to be the case.

Finally, the results in this chapter demonstrate that the idea of resource-based patient prioritization (namely, taking into account resource limitations and survival probability functions to attempt to maximize the expected number of survivors) is promising not only in a simple situation, but also in more complicated scenarios such as the multiple-location mass-casualty incident, and these results can be used when developing plans or policies to respond to such an incident.

Chapter 5

Dynamic Routing of Casualties in a Mass-Casualty Incident

In this chapter, we study the problem of dynamic decision making in a large scale casualty evacuation following an MCI. Large scale natural or man-made disasters often result in a large number of casualties who need to be evacuated in order to receive some kind of care. For example, casualties in a biological, chemical, or nuclear disaster must be transported out of the affected area and then undergo decontamination; casualties from an earthquake, flood, or hurricane must be moved to a safe place to undergo triage or to receive treatment for injuries. The main goal of the response to such an incident is to ensure the safe and timely treatment of as many casualties as possible. Emergency responders may encounter limited resources and unexpected events, such as infrastructure damage on the evacuation routes or overcrowding at some of the facilities that are set up to receive casualties. Because conditions change during the response effort, it is therefore important for emergency responders to be able to make decisions dynamically throughout the evacuation to make the evacuation as efficient as possible.

We specifically consider the problem of transporting casualties from areas affected by the disaster to a set of casualty collection points, where they receive treatment. This problem is particularly relevant in a mass-casualty event involving the release of hazardous substances (HAZMAT). Hazardous substances include biological, chemical, or nuclear agents, and they may be released deliberately (as in a terrorist attack) or accidentally (as in a meltdown at a nuclear power plant). HAZMAT events pose a particular challenge because both

casualties and responders are at risk from exposure, and the U.S. government plans and prepares for such scenarios (U.S. Department of Homeland Security 2005). In such an event, casualties must first be evacuated away from a site where they could be affected by the substance and then undergo decontamination. These two steps must be completed before the casualties can receive additional care for injuries or illnesses (Hrdina et al. 2009). To avoid cross-contamination, it would be ideal to establish decontamination as close to the scene of the incident as possible, but in a mass-casualty event, it usually is not possible to establish scene decontamination if there is a risk of continued exposure (Robenshtok et al. 2003). Moreover, health care facilities often do not rely on public safety agencies to provide decontamination services, but instead elect to provide those services prior to admission of patients, both for self-referred and ambulance arrivals (Hick et al. 2003). For some types of incidents, such as detonation of an improvised nuclear device or radiological dispersal device, it is desirable to establish radiation triage/treatment/transport sites at intermediate locations so that patients can be evaluated and decontaminated before being transferred to medical care facilities (Hrdina et al. 2009).

Three distinct zones are established during the evacuation of casualties in an event involving hazardous substances (Boardman et al. 2008). “Hot zones” are areas where there is ongoing risk of exposure to the hazardous substance. Hot zones usually include the site(s) of the incident, but may also include areas downwind of the incident. Because both casualties and rescuers have a risk of exposure, the type of treatment carried out in a hot zone is very limited. The primary goal of rescuers is to evacuate the casualties away from the hot zone. “Warm zones” are areas outside the hot zones where contaminated casualties are present. The risk of exposure in a warm zone is only due to the presence of contaminated casualties (i.e., through cross-contamination), as there is no direct exposure to the substance. Warm zones include the casualty collection points where the casualties are gathered as they wait to be decontaminated. Warm zones may be established in the field and at medical facilities that are equipped to receive contaminated casualties. Finally, the “cold zone” includes all areas not in a warm or hot zone. In order to pass from the warm zone to the cold zone, casualties must undergo thorough decontamination that may include removal of clothing, showering or flushing with water, and scrubbing with soap. In order to meet

medical and regulatory requirements, the personnel performing the decontamination must be specially trained and equipped (U.S. Occupational Safety and Health Administration 2005, Hick et al. 2003). Therefore, it is not usually practical to establish a large number of parallel servers for decontamination. Figure 5.1 presents a schematic diagram showing the movement of casualties between the zones.

In this chapter, we study how to route casualties to the collection points dynamically, so that the evacuation can be adjusted to account for congestion and avoid overwhelming or under-utilizing any of the collection points. While different types of events have different requirements in terms of the exact resources needed or treatment provided, a common theme that emerges in the literature is that the dynamic management of the evacuation can improve outcomes and lessen confusion. In a study of thousands of patients evacuated from an earthquake, Tanaka et al. (1998) suggested that to improve evacuation, “disaster officials must know the capabilities and capacity of each area hospital at all times to select appropriate triage and mode of transport for each victim”. Robenshtok et al. (2003) state that in a chemical event, “specific hospitals should be designated to receive only chemical agent casualties.” Hick et al. (2011) suggest that during a nuclear incident, emergency managers should bypass hospitals that are “completely overwhelmed”. Hrdina et al. (2009) note that a mass-casualty event involving radiation “will require a wider distribution of patients” than a typical mass-casualty event and thus “a networked system” for assigning patients to medical facilities is an essential part of any model for responding to such an incident. These studies give qualitative observations about management of casualties in a disaster; they suggest that dynamic management of casualty evacuation could be helpful, for example, by incorporating information about the capabilities and the level of congestion at each casualty collection point.

While a number of articles in the literature consider evacuation and casualty management problems, existing approaches generally use optimization models, such as network flows and mathematical programming, which are more suited for planning purposes or for determining an efficient travel path for emergency vehicles. In this chapter, we consider a different aspect of the evacuation problem that to our knowledge has not received attention; namely, the queueing of casualties at the collection points and its effect on dynamic

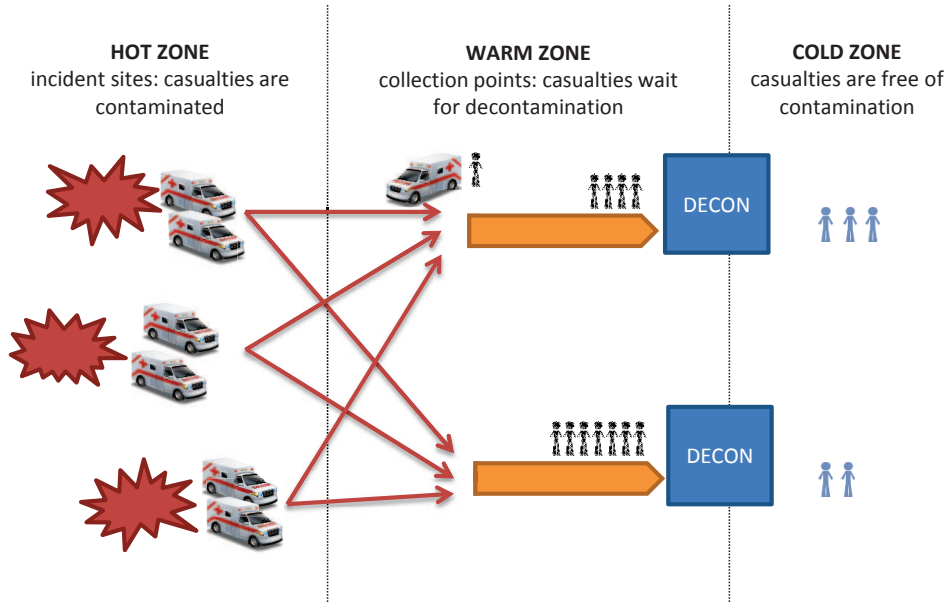


Figure 5.1: Evacuation and decontamination of casualties in the aftermath of a mass-casualty event involving the release of a hazardous substance: casualties are removed from the hot zone to the warm zone, where they wait to receive decontamination service before being released to the cold zone.

routing decisions. This queueing is especially relevant in situations such as those depicted in Figure 5.1, namely, those situations where a large number of casualties all need similar service (such as decontamination) and where the service can be provided at one of several locations.

5.1 Dynamic programming formulation

In this section, we define the main problem and the Markov decision process (MDP) formulation we will use throughout the rest of the chapter. We then briefly discuss similar models in the literature and our analytical approach.

Consider a disaster with multiple casualty clusters. Each casualty cluster is considered a hot zone, where casualties are exposed to danger. We assume that each cluster has an ample supply of casualties. Each cluster has a transportation resource or set of resources that can be used to transport the casualties to one of several casualty collection points (that we also refer to as “stations”), which may be hospitals, shelters, field triage areas, or ad hoc decontamination facilities. The rate at which casualties can be transported from

a given cluster to a given station is inversely proportional to the travel time between the cluster and the station. Once a casualty reaches the collection area, he or she waits to receive service from a single server. For medical reasons, it is always advantageous for the casualties to be located at the collection point as opposed to the scene of the incident (i.e., it is better to be in the warm zone than the hot zone), so we do not charge holding costs for casualties who are waiting at the collection points. Once completing service, the patient enters the cold zone and the system earns a reward. The reward can depend on the station where the casualty receives service, and thus the reward can be used to represent different capabilities at the different stations. In the absence of information about the capabilities of the different stations, it may be convenient to think of all rewards as being one, in which case the reward will correspond to the number of patients who have completed service (and thus the expected total discounted reward is simply the discounted throughput of the system). We seek to find the routing policy that maximizes the expected total discounted reward earned by the system.

We let \mathcal{R} denote the set of all transportation resources. Each resource corresponds to a specific cluster, although there may be more than one resource per cluster. For the remainder of this chapter, we assume a one-to-one correspondence (i.e., one resource per cluster) without loss of generality. Since clusters have an ample supply of casualties, having more than one resource at a cluster is the same as having more than one cluster. We also let \mathcal{S} denote the set of stations, each of which has a single server with a dedicated queue. A resource $i \in \mathcal{R}$ can transport casualties to station $j \in \mathcal{S}$ with a travel time that is exponentially distributed with rate $\tau_{ij} > 0$. This travel is preemptive, i.e., the resources can be re-routed while en route. Travel back to the cluster is instantaneous. The server at station j works according to an exponential distribution with rate $\mu_j > 0$. When a casualty finishes service at station j , the system receives a reward $r_j > 0$.

In order to formulate the problem as a MDP, and therefore obtain policies based on analytical results, we make several assumptions. Before we provide our MDP formulation, we discuss some of these assumptions in more detail. First, we assume that each casualty cluster has an ample supply of patients. This is clearly an approximation, even though the number of casualties in a mass-casualty event is generally large. Later, we relax this

assumption in our simulation study. Second, we assume that each collection point operates as a single server queue. Adding additional servers may be desirable; however, adding additional servers in a decontamination setting is not simple: for reasons such as available space, equipment, protective gear, and utility hookups, adding an additional server in the decontamination setting cannot be done quickly if at all (U.S. Occupational Safety and Health Administration 2005). Decontamination capacity may be further limited by the number of trained personnel available: at smaller facilities, they may number as few as one 2-person team (Hick et al. 2003). At larger facilities, it is “reasonable to expect” that ambulatory (walking) and non-ambulatory patients may be decontaminated in parallel (Hick et al. 2003); in this case, we do not model the ambulatory patients because they self-transport to the facility. Finally, we assume that vehicles have Markovian travel times and become instantaneously available after dropping off a casualty. Again, we relax these assumptions in the simulation study.

We denote the state of the system by $X(t) = (X_1(t), \dots, X_{|\mathcal{S}|}(t))$, i.e., the number of casualties at each of the stations, at time t , where $|\mathcal{S}|$ is the cardinality of \mathcal{S} . The state space, which we denote by \mathcal{Q} , is $\{(x_1, x_2, \dots, x_{|\mathcal{S}|}) : x_i \geq 0 \text{ for } i = 1, 2, \dots, |\mathcal{S}|\}$. Then the expected total discounted reward with discount rate $\alpha > 0$ is

$$E \left[\int_0^\infty e^{-\alpha t} r(X(t)) dt \right],$$

where $r(x) = \sum_{j \in \mathcal{S}} r_j \mu_j I_j(x)$, $x \in \mathcal{Q}$, and $I_j(x)$ is the indicator function that takes value 1 if $x_j > 0$ and 0 if $x_j = 0$.

We next formulate this optimization problem as a Markov Decision Process. Throughout this chapter, we use uniformization, and we define the finite uniformization constant $\beta \equiv \sum_{j \in \mathcal{S}} \mu_j + \tau$, where $\tau = \sum_{i \in \mathcal{R}} \tau_i$ and $\tau_i = \max_{j \in \mathcal{S}} \tau_{ij}$. In other words, by observing the system only at transitions, which occur according to a Poisson process with rate β , we study the discrete time MDP that is embedded in the continuous process. Without loss of generality, we let $\beta = 1$. A decision epoch occurs any time service is completed at one of the stations and any time a resource finishes transporting a patient and becomes available. If the event occurring at a decision epoch is a service completion, the state decreases by

one at the station completing the service. If the event occurring at a decision epoch is a transportation completion, the state increases by one at the station receiving the patient. At each decision epoch, the action taken is to assign a station to each resource. We only admit stationary policies; that is, under an admissible policy, each state corresponds to a single station for each resource, independent of time. We denote the set of admissible policies by \mathcal{P} .

Let $V(x)$ denote the maximum expected total discounted reward that can be obtained by an admissible policy when the system starts in state $x \in \mathcal{Q}$. We are interested in finding a policy in \mathcal{P} that yields $V(x)$. That is, we want to solve

$$\max_{\pi \in \mathcal{P}} E \left[\int_0^\infty e^{-\alpha t} r(X^\pi(t)) dt \mid X^\pi(0) = x \right],$$

where $X^\pi(t)$ is the system state at time t under policy π and $x \in \mathcal{Q}$. $V(x)$ must satisfy the optimality equation, which is

$$V(x) = \frac{1}{1 + \alpha} \left(r(x) + \sum_{j \in \mathcal{S}} \mu_j V(x - I_j(x)e_j) + \tau V(x) + \sum_{i \in \mathcal{R}} \max_{j \in \mathcal{S}} \{ \tau_{ij} M_j(x) \} \right), \quad (5.1)$$

where e_j is the vector having the j th component equal to one and all others equal to zero, and $M_j(x) = V(x + e_j) - V(x)$ is the marginal value of having an additional casualty at station j . Our objective is to find a policy that satisfies (5.1).

The queueing system considered in this chapter is closely related to a make-to-stock inventory problem—see, e.g., Veatch and Wein (1996). Namely, if we consider the transportation time to be production time and consider the service completions at different servers to be demand realizations for different products, then maximizing the discounted reward in the evacuation problem is equivalent to minimizing the discounted cost of lost sales in the make-to-stock inventory problem, provided there are multiple production resources and multiple products.

Veatch and Wein (1996) studied the make-to-stock inventory problem with a single production resource and multiple products. In that article, two different inventory regimes (lost sales and backordering) and two different cost methods (discounted and long-run average)

were considered. However, the index policies developed in the article were derived from results concerning the version of the problem with backordering and long-run average costs, and the authors note that the backordering case is somewhat easier to analyze than the lost sales case. Moreover, multiple production resources were not considered in Veatch and Wein (1996). Subsequent articles concerning make-to-stock inventory, such as Ha (1997), consider only the backordering case. Because backordering does not map to any notion in our model, it is difficult to apply any of the results from work on make-to-stock inventory systems to the casualty evacuation problem. However, the difficulty in determining the optimal policies in these articles suggests that it is not likely that we can fully characterize the optimal policy in any simple form. Nevertheless, we can provide some partial characterizations of the optimal solution, which we present in the next section.

5.2 Analytical Results

In this section, we establish several analytical results that provide some insight into the structure of the optimal routing policy, i.e., the solution to the dynamic programming problem given in (5.1).

The first proposition establishes a monotonicity relationship among any set of two resources for any two given stations. Proofs of the following proposition and all other propositions appear in Appendix . To aid in exposition of the analytical results, we need the following definition.

Definition 5.1. For a resource $k \in \mathcal{R}$, stations $j, l \in \mathcal{S}$, and state $x \in \mathcal{Q}$, we say that j is (x, k) -preferable to l if and only if $\tau_{kj}M_j(x) \geq \tau_{kl}M_l(x)$.

When station j is (x, k) -preferable to station l , this means that when the system is in state x , sending resource k to station j will yield an expected total discounted reward at least as large as sending resource k to station l , if the optimal policy is used for all other decisions. Thus, we can safely eliminate the possibility of routing resource k to station l in state x .

Proposition 5.1. *Consider any pair of resources $i, k \in \mathcal{R}$ and any pair of stations $j, l \in \mathcal{S}$,*




Figure 5.2: Example problem instance with $\mathcal{R} = \{1, 2\}$ and $\mathcal{S} = \{A, B\}$; numbers on arcs indicate mean travel times $1/\tau_{ij}$ for $i \in \mathcal{R}, j \in \mathcal{S}$.

and without loss of generality label them such that $\frac{\tau_{ij}}{\tau_{il}} \geq \frac{\tau_{kj}}{\tau_{kl}}$. For any state $x \in \mathcal{Q}$, if j is (x, k) -preferable to l , then j is (x, i) -preferable to l .

Corollary 5.1. *If $\frac{\tau_{ij}}{\tau_{il}} \geq \frac{\tau_{kj}}{\tau_{kl}}$, it is suboptimal in any state to route resource k to station j while routing resource i to station l .*

An example where the above proposition applies is presented in Figure 5.2. In the figure, casualties from cluster 1 can be transported to station A twice as fast as to station B. Casualties from cluster 2 can be transported to station A only one-and-a-half times as fast as to station B. At a given state of the system, if it is optimal for a casualty from cluster 2 to go to station A, then it is also optimal for a casualty from cluster 1 to go to station A. This conclusion is intuitive from examining the picture. Responders at cluster 1 would prefer to go to station A versus station B more strongly than those at cluster 2: at cluster 2, the difference in travel times between the two stations is not as large. Another way to state this result, using Corollary 5.1 is to say that one of the four possible actions, namely the action that simultaneously sends a casualty from cluster 1 to station B and sends a casualty from cluster 2 to station A, will be suboptimal regardless of the system state.

For the remainder of this section, we will make use of a finite-horizon version of the problem. Specifically, we will examine $V^n(x)$, the maximum expected reward earned when the system will run for n additional epochs (each epoch occurring according to a Poisson

process with rate β). We have

$$\begin{aligned}
V^{n+1}(x) &= \frac{1}{1+\alpha} \left(r(x) + \sum_{j \in \mathcal{S}} \mu_j V^n(x - I_j(x)e_j) + \tau V^n(x) + \sum_{i \in \mathcal{R}} \max_{j \in \mathcal{S}} \{ \tau_{ij} (V^n(x + e_j) - V^n(x)) \} \right), \\
\end{aligned} \tag{5.2}$$

and we assume $V^0(x) = 0$. If a result concerning $V^n(\cdot)$ holds for all values of n , then the same result will hold for $V(\cdot)$, because $V^n(\cdot)$ converges to $V(\cdot)$ uniformly as n approaches infinity as long as $V^0(x)$, which we assumed to be zero for all x , is bounded (Ross 1983, Chapter 2.3, Proposition 3.1).

The next proposition demonstrates two additional types of monotonicity present in this model. Specifically, it is always preferable to have more casualties in the warm zone (since they would otherwise be in the hot zone), and it is always preferable to have more casualties in the cold zone (having completed service) than in the warm zone.

Proposition 5.2. *The following inequalities hold for all values of $n \geq 0$, $x \in \mathcal{Q}$:*

$$V^n(x + e_j) \geq V^n(x), \quad \forall j \in \mathcal{S}. \tag{5.3}$$

$$V^n(x) + r_j \geq V^n(x + e_j), \quad \forall j \in \mathcal{S}. \tag{5.4}$$

Next, we show an analytical result that can be proved in the case where $\tau_{kj} = \tau_k$ for all $k \in \mathcal{R}$, $j \in \mathcal{S}$; that is, mean travel times depend only on the resource used and/or the location, but not on the destination station. This case approximates the scenario where the medical facilities are close to one another and/or the primary difference in transportation rates is due to a feature of the incidents, rather than the facilities (e.g., one of the incidents is in an area that is difficult to access). Specifically, the following result shows that if the the fastest station has the largest reward, then it is the optimal destination for all casualties whenever it has the shortest queue.

Proposition 5.3. *Suppose that $\tau_{kj} = \tau_k$ for all $k \in \mathcal{R}$, $j \in \mathcal{S}$, and that $r_1 \geq r_j$ and $\mu_1 \geq \mu_j$*

for all $j \in \mathcal{S}$. If $x_1 \leq x_j$ for all $j \in \mathcal{S}$, then

$$V^n(x + e_1) \geq V^n(x + e_j), \quad (5.5)$$

$$V^n(T_j x) \geq V^n(x), \quad (5.6)$$

for all $n \geq 0$ and for all $j \in \mathcal{S}$, where T_j is the transformation that swaps the 1st element of x with the j th element of x .

Note that (5.5) implies that the optimal routing is to station 1, because it implies that $\tau_{k1}V^n(x + e_1) \geq \tau_{kj}V^n(x + e_j)$ for all $j \neq 1$, which in turn implies that $\tau_{k1}(V^n(x + e_1) - V^n(x)) = \max_{j \in \mathcal{S}} \{\tau_{ij}(V^n(x + e_1) - V^n(x))\}$. Equation (5.6) does not have a practical interpretation—it is used in the proof of (5.5).

The above proposition gives an “agreeability” condition for optimality. Specifically, if a station has the largest reward, the fastest service, and the shortest queue, it should be chosen. Numerically, it appears that the result holds even if mean travel times for each resource are not the same, as long as the fastest station also has the shortest travel time (i.e., if $\tau_{k1} \geq \tau_{kj}$ for all $j \neq 1$). However, our proof of Proposition 5.3 does not work in this more general case.

Finally, we extend the results of Propositions 5.2 and 5.3 to the infinite horizon case, by making use of the fact that $V^n(\cdot)$ converges to $V(\cdot)$.

Corollary 5.2. *The following inequalities hold for all values of x .*

$$V(x + e_j) \geq V(x), \quad \forall j \in \mathcal{S}. \quad (5.7)$$

$$V(x) + r_j \geq V(x + e_j), \quad \forall j \in \mathcal{S}. \quad (5.8)$$

Corollary 5.3. *Suppose that $\tau_{kj} = \tau_k$ for all $k \in \mathcal{R}$, $j \in \mathcal{S}$, and that $r_1 \geq r_j$ and $\mu_1 \geq \mu_j$ for all $j \in \mathcal{S}$. If $x_1 \leq x_j$ for all $j \in \mathcal{S}$, then*

$$V(x + e_1) \geq V(x + e_j). \quad (5.9)$$

Therefore, it is optimal to send casualties to the fastest station if it also has the largest

reward and the shortest queue.

Although the analytical results in this section do not give a complete characterization of the optimal policy, they do confirm that the optimal policy behaves in certain ways that are somewhat intuitive. Moreover, as we develop heuristic policies in the next section, we will be able to check their consistency with the results of this section.

5.3 Heuristic Policies

In this section, we develop heuristics that can approximate the optimal policy. We use two methods: marginal value function approximation and the one-step policy improvement heuristic. Marginal value function approximation involves finding an approximate value for $M_j(x)$, and then using that approximation in (5.1). Note that the function $M_j(x)$ used in the approximation may not correspond to any admissible policy. On the other hand, the one-step policy improvement requires first setting a static policy and then applying a single step of the policy improvement algorithm.

5.3.1 Marginal value approximation

Recall from (5.1) that if the marginal value function $M_j(x)$ can be calculated, then the optimal policy is for each resource i to route the next casualty to the station with the largest index $\tau_{ij}M_j(x)$. A marginal value function approximation simply finds a way to approximate $M_j(x)$.

We choose to approximate $M_j(x)$ under the assumptions that station j is excluded as a possible decision, and that the policy does not depend on x_j . In other words, we calculate the value functions assuming that casualties already in queue for station j are processed, and that the rest of the system operates as if station j does not exist. In this case, $M_j(x) = V(x + e_j) - V(x)$ is simply the expected reward earned by the $(x_j + 1)$ st item at station j , which is r_j , discounted from the time of the $(x_j + 1)$ st service completion. The service completion time of the $(x_j + 1)$ st customer is an $\text{Erlang}(x_j + 1, \mu_j)$ random variable. Hence, the expected reward is $r_j \left(\frac{\mu_j}{\mu_j + \alpha} \right)^{x_j + 1}$.

Therefore the marginal value approximation (MVA) heuristic is for each resource $i \in \mathcal{R}$ to route casualties to the station with the largest index $\tau_{ij}r_j \left(\frac{\mu_j}{\mu_j + \alpha}\right)^{x_j + 1}$. The MVA index has the advantage of being simple, and it satisfies the “agreeability” condition that the optimal policy possesses (see Proposition 5.3 and the discussion following it), namely that if a given station has the shortest travel time, the fastest server, the largest reward, and the shortest queue, then it would be chosen by the MVA index.

5.3.2 Policy improvement heuristics

In this section, we develop heuristics based on applying a single step of the policy improvement algorithm, starting from a static policy. In other work on queueing systems, this method has been shown to provide good heuristics (see, e.g., Argon et al. (2009) and Liu et al. (2010)). A static policy means that routing decisions are state-independent. Therefore, it is usually much easier to find the optimal static policy (or at least a static policy that performs well) than to find the optimal dynamic policy, and the dynamic policy that results from applying the single step of the policy improvement algorithm is guaranteed to perform better than the static policy. In this section, we use a randomized static policy where casualties from cluster i are routed to a station j randomly with probability ρ_{ij} at each decision epoch, where $\sum_{j \in \mathcal{S}} \rho_{ij} = 1$. In this way, casualties from cluster i arrive to station j at rate $\rho_{ij}\tau_{ij}$ (i.e., each arrival that would have occurred at rate τ_{ij} only occurs with probability ρ_{ij}). Let $\lambda_j = \sum_{i \in \mathcal{R}} \rho_{ij}\tau_{ij}$ and $\rho_j = \lambda_j/\mu_j$. Then ρ_j is the traffic intensity into station j ; each station is M/M/1 with arrival rate λ_j and service rate μ_j .

Since any static policy Γ is completely determined by $\{\rho_{ij}, i \in \mathcal{R}, j \in \mathcal{S}\}$, in order to apply one step of the policy improvement algorithm, we need to determine the value function (and thus the marginal value function) associated with a generic static policy.

Proposition 5.4. *Let Γ be a static policy having probabilities $\{\rho_{ij}, i \in \mathcal{R}, j \in \mathcal{S}\}$. The value function $V^\Gamma(x)$ associated with Γ is*

$$\sum_{j \in \mathcal{S}} \left[\frac{\mu_j r_j}{\alpha} - \frac{2\mu_j r_j}{\lambda_j + \alpha - \mu_j + \eta_j} \left(\frac{\lambda_j + \mu_j + \alpha - \eta_j}{2\lambda_j} \right)^{x_j} \right], \quad (5.10)$$

where $\eta_j = \sqrt{(\mu_j + \lambda_j + \alpha)^2 - 4\lambda_j\mu_j}$.

Note that the value function in (5.10) is separable by station (because each station is a separate M/M/1 queue). By using (5.10), we can compute the marginal value function $M_j^\Gamma(x)$ associated with static policy Γ , which is

$$\frac{\mu_j r_j}{\lambda_j} \left(\frac{\mu_j - \lambda_j + \alpha - \eta_j}{\mu_j - \lambda_j - \alpha - \eta_j} \right) \left(\frac{\mu_j + \lambda_j + \alpha - \eta_j}{2\lambda_j} \right)^{x_j}. \quad (5.11)$$

$M_j^\Gamma(x)$ is the static policy marginal expected total discounted reward from having one additional customer at station j when the system state is $x \in \mathcal{Q}$ under static policy Γ . Applying one step of the policy improvement method means that we should use $M_j^\Gamma(x)$ in (5.1). Namely, the one-step policy improvement heuristic (PIH) is to send customers to the station with the largest index $\tau_{ij} M_j^\Gamma(x)$.

In the remainder of this section, we focus on three different ways to choose static policy Γ for use in the policy improvement heuristic.

Optimal static policy. The one-step policy improvement heuristic should work best if we start with the optimal randomized static policy, i.e., the policy that yields the largest expected total discounted reward from among all randomized static policies. The total expected reward starting with the empty system (i.e., $x_j = 0$ for all $j \in \mathcal{S}$) can be written in closed form as

$$\sum_{j \in \mathcal{S}} \frac{\mu_j r_j}{\alpha} - \sum_{j \in \mathcal{S}} \frac{2\mu_j r_j}{\sum_{i \in \mathcal{R}} \rho_{ij} \tau_{ij} + \alpha - \mu_j + \sqrt{(\sum_{i \in \mathcal{R}} \rho_{ij} \tau_{ij} + \mu_j + \alpha)^2 - 4\mu_j \sum_{i \in \mathcal{R}} \rho_{ij} \tau_{ij}}} \quad (5.12)$$

by evaluating (5.10) with $x_j = 0$ for all $j \in \mathcal{S}$. Finding the optimal static policy requires maximizing the above expression subject to $\sum_{j \in \mathcal{S}} \rho_{ij} = 1, \forall i \in \mathcal{R}$. Unfortunately, the objective function is nonlinear in ρ_{ij} and so the problem is difficult to solve except for cases with a very small number of clusters and stations, where it can be solved numerically. Where practical in our numerical experiments, we report the value of the policy improvement heuristic with optimal static policy (PIH-O) for comparison.

Fluid approximation. Next, we consider optimization of the fluid approximation of the randomized static policy. A fluid system is one in which the discrete entities are treated

as continuous; transportation of casualties to the stations and subsequent service thereof is equivalent to the flow of a fluid. Recall that we interpret the inverse of τ_{ij} to be the travel time to station j using resource i . Therefore, the fluid flowing from resource i to station j first arrives at time $1/\tau_{ij}$, flows at rate $\rho_{ij}\tau_{ij}$ units of fluid per time unit, and continuously earns a reward of r_j per unit of fluid. This reward is discounted to time zero at discount factor α . Thus, the total reward earned by fluid flowing from resource i to station j is $\int_{1/\tau_{ij}}^{\infty} \rho_{ij}\tau_{ij}r_j e^{-\alpha t} dt = \frac{\rho_{ij}\tau_{ij}r_j}{\alpha} e^{-\frac{\alpha}{\tau_{ij}}}$. In order to linearize the objective function without adding additional variables, we assume that fluid does not arrive to station j at a total rate faster than μ_j . Such a solution would actually be feasible, but it could not increase the expected total discounted reward because the fluid can never flow out of station j (and thus earn a reward) at a rate faster than μ_j .

The fluid approximation is thus the following linear program:

$$\begin{aligned} \max \quad & \sum_{i \in \mathcal{R}} \sum_{j \in \mathcal{S}} \tau_{ij} r_j e^{-\frac{\alpha}{\tau_{ij}}} \rho_{ij} \\ \text{s.t.} \quad & \sum_{i \in \mathcal{R}} \rho_{ij} \tau_{ij} \leq \mu_j, \quad \forall j \in \mathcal{S} \end{aligned} \tag{5.13}$$

$$\sum_{j \in \mathcal{S}} \rho_{ij} = 1, \quad \forall i \in \mathcal{R} \tag{5.14}$$

$$\rho_{ij} \geq 0, \quad \forall i \in \mathcal{R}, j \in \mathcal{S}. \tag{5.15}$$

The above linear program is an example of a fractional multiple-knapsack problem. If the problem is relaxed by removing constraints (5.13) (which would be a good approximation when stations are fast compared to transportation resources), then it can be solved optimally via greedy choice, by choosing for each $i \in \mathcal{R}$, $\rho_{ij'} = 1$ for $j' = \arg \max_{j \in \mathcal{S}} \tau_{ij} r_j e^{-\frac{\alpha}{\tau_{ij}}}$ and 0 for all others. In our numerical experiments, we will report the results of the policy improvement heuristic using both the probability assignment from the fluid LP (PIH-F) and from greedy choice (PIH-G).

Table 5.1: Heuristic results as a percentage of optimal value. Minimum, quartiles, maximum, and mean are calculated across 10,000 randomly generated instances.

α	Heuristic	Min	Q1	Q2	Q3	Max	Mean
0.1	Greedy	62%	95%	97%	99%	100%	96%
	Fluid	92%	96%	97%	99%	100%	97%
	Static-Opt	93%	97%	98%	99%	100%	98%
	MVA	45%	91%	97%	99%	100%	93%
	PIH-G	96%	99%	100%	100%	100%	99%
	PIH-F	96%	99%	100%	100%	100%	99%
	PIH-O	99%	100%	100%	100%	100%	100%
0.7	Greedy	61%	92%	95%	97%	100%	94%
	Fluid	82%	94%	95%	97%	100%	96%
	Static-Opt	88%	95%	97%	98%	100%	97%
	MVA	89%	96%	99%	100%	100%	98%
	PIH-G	93%	99%	100%	100%	100%	99%
	PIH-F	93%	99%	100%	100%	100%	99%
	PIH-O	97%	99%	100%	100%	100%	100%
2.0	Greedy	61%	91%	94%	96%	100%	93%
	Fluid	76%	92%	94%	96%	100%	94%
	Static-Opt	86%	94%	95%	97%	100%	95%
	MVA	93%	98%	99%	100%	100%	99%
	PIH-G	93%	99%	100%	100%	100%	99%
	PIH-F	93%	99%	100%	100%	100%	99%
	PIH-O	96%	99%	100%	100%	100%	100%

5.4 Numerical Results

In order to test the performance of the heuristics we developed in Section 5.3, we conducted numerical experiments with randomly generated instances. We generated 10,000 instances, each with two clusters and two locations. By using a small number of clusters and locations, we were able to calculate the value of $V(\cdot)$ numerically. Each instance had service rates and transportation rates chosen uniformly in $[1, 10]$. We rejected instances where $\sum_{i \in R} \max_{j \in S} \tau_{ij} > \sum_{j \in S} \mu_j$, a condition that indicates that the system is unstable under a static policy.

For each instance, the optimal policy was calculated using the value iteration algorithm under a truncated state space with 125 states. The expected total discounted reward of each heuristic was then calculated, assuming that the systems starts empty. The expected total discounted reward of each heuristic is expressed in Table 5.1 as a percentage of the reward under the optimal policy. For comparison, the performance of the static policies alone (prior to applying one step of the policy improvement algorithm) is also shown.

All heuristics perform quite well on at least a majority of instances, but there is a clear tradeoff between simplicity and performance. Marginal value approximation has the simplest index to calculate, but performs poorly in a small number of instances. Nonetheless, MVA still achieves an average of 98% of the value of the optimal policy. On the other hand, policy improvement performs quite well, even when the initial static policy is not the optimal one. On this set of instances, a single step of the policy improvement algorithm applied to the optimal static policy is essentially indistinguishable from the optimal dynamic policy. However, the optimal static policy itself must be calculated numerically, and this can only be accomplished for problems with a relatively small number of clusters and stations. Fortunately, most of the benefit of the policy improvement heuristic seems to come from using the heuristic rather than from the static policy itself. This can be seen in Table 5.1.

The results are also somewhat sensitive to the discount factor α . If α is very small, then the problem approaches a problem of maximizing the long-run average reward. As we would expect, for small α , static policies tend to do very well (and consequently, policy improvement does well). On the other hand, MVA begins to perform poorly, because the

estimate of the marginal value in future steps becomes relatively more important. The opposite effect is observed when α is very large. In that case, the estimate of the future marginal value is less important and MVA performs very well. However, there is a marked decrease in performance of the static policies. Notably, the performance of the policy improvement heuristic is extremely good and does not seem to depend on α . MVA is the simplest of all the policies, yet it performs best when time criticality (as measured by discounting) is greatest.

While we have shown that our heuristics produce numerical results that are very close to the optimal policy, it is obvious that numerical experiments alone are not sufficient to draw conclusions about their applicability. In particular, it is important to consider scenarios with more than two casualty clusters and two collection points, and where the distributional assumptions are violated. Adding more clusters and collection points quickly makes the MDP intractable due to the curse of dimensionality. Indeed, even the optimal randomized static policy, which must be calculated numerically, is very time consuming to obtain when there are more than two clusters and two collection points. Therefore, simulation is the tool of choice for examining the performance of these policies in a more realistic environment.

5.5 Simulation Study

In this section, we conduct a simulation study to test the performance of our policies in the case where the assumptions made in the formulation of the MDP are not accurate. In particular, as we discussed in Section 5.1, ambulance travel times are not exponentially distributed, the return time for ambulances is non-negligible, and the number of casualties will necessarily be finite.

In addition to having more realistic distributions, we also construct the simulation study to be geographically realistic. We suppose that a biological, chemical, or nuclear terrorist attack has taken place in an urban area that is 100 square miles (ten miles by ten miles). There are n incidents, randomly located within the urban area with p patients each. There are m facilities, $m - 1$ of which are randomly located within the urban area and one of which

is located between 15 and 30 miles outside the area. Each incident has two transportation resources, which can travel to the medical facilities at an average speed of 40 mph. Distance between points is Euclidean, and the travel time has a Lognormal distribution. This selection is based on the empirical observation that the Lognormal distribution best represented travel times of ambulances (Ingolfsson et al. 2008). A single server decontaminates patients at each facility before they enter the facility for treatment, and the decontamination requires Triangular(1,5,2) minutes. This selection is based on a literature review that suggests that various agencies maintain decontamination standards ranging from one minute to five minutes, although “actual showering time will be an incident-specific decision” (U.S. Occupational Safety and Health Administration 2005). Therefore, we will simulate an incident in which incident commanders have prescribed a decontamination time of two minutes, but where the actual time may be as little as one or as much as five minutes, depending on the individual.

We simulate the scenario under four different policies:

- All patients routed to the nearest facility.
- All patients routed to the facility with the shortest queue.
- Patients routed according to MVA policy.
- Patients routed according to PIH-G policy.

MVA and PIH-G were chosen as candidate policies because they strike a balance between simplicity and complexity, and thus of the policies we have developed, have the maximum potential for usability. These factors have been identified as among the most important in constructing models for medical decision making (Brandeau et al. 2009). “Nearest facility” is a policy that is used in practice in daily emergencies as well as in major disasters; for example, when sarin gas was released in 15 Tokyo subway stations, providers were unable to get information about hospital availability due to poor communication, and many patients were taken to the nearest hospital (Okumura et al. 1998). While it is not clear whether the “shortest queue” policy has been used in practice, intuitively it is a reasonable myopic

alternative for cases where providers are concerned about bypassing nearby hospitals due to congestion; see, e.g., observations in Hick et al. (2011).

The results shown in Table 5.2 show the change in several performance metrics for $m = 3, n = \{3, 4, 5\}, p = 50$. Results are given as 95% confidence intervals. Note that flow time and make span are non-discounted performance measures for which a decrease is desirable. On the other hand, an increase is desirable for the average number of busy servers. For these metrics, the discount factor was used only to determine the policy. Finally, discounted throughput is presented because it is the metric our policies intend to maximize; thus, an increase is desirable.

Table 5.2: Simulation results for $m = 3, n = \{3, 4, 5\}, p = 50$. (*) indicates significance at the 0.05 level.

Change in performance metric versus choosing the nearest hospital, 95% confidence interval

n	Policy	Disc.	Flowtime (hours)	Makespan (hours)	Avg. busy servers	Disc. Tpt.
3	MVA	0.1	[-12.4, -3.1] *	[-0.1, 0.0]	[-0.0, 0.0]	[0.2, 0.9] *
3	MVA	0.7	[-28.4, -12.7] *	[-0.1, 0.0]	[-0.0, 0.0]	[1.2, 2.6] *
3	PIH-G	0.1	[-40.0, -17.6] *	[-0.2, 0.0]	[0.0, 0.1] *	[1.3, 2.9] *
3	PIH-G	0.7	[-43.4, -19.7] *	[-0.2, 0.1]	[0.0, 0.1] *	[2.3, 4.4] *
4	MVA	0.1	[-16.9, -3.7] *	[-0.1, 0.0]	[-0.0, 0.0]	[0.2, 1.1] *
4	MVA	0.7	[-71.6, -42.3] *	[-0.3, -0.1] *	[0.0, 0.1] *	[2.5, 4.1] *
4	PIH-G	0.1	[-82.5, -48.3] *	[-0.5, -0.2] *	[0.0, 0.1] *	[3.2, 5.3] *
4	PIH-G	0.7	[-95.6, -56.9] *	[-0.3, 0.0]	[0.0, 0.1] *	[4.1, 6.3] *
5	MVA	0.1	[-38.5, -9.2] *	[-0.2, -0.0] *	[0.0, 0.0] *	[0.5, 2.1] *
5	MVA	0.7	[-165.3, -101.1] *	[-0.7, -0.3] *	[0.1, 0.1] *	[3.6, 6.0] *
5	PIH-G	0.1	[-198.2, -126.1] *	[-0.9, -0.5] *	[0.1, 0.2] *	[7.5, 11.6] *
5	PIH-G	0.7	[-210.6, -133.0] *	[-0.7, -0.2] *	[0.0, 0.1] *	[6.3, 9.4] *

Change in performance metric versus choosing the hospital with the shortest queue, 95% C.I.

n	Policy	Disc.	Flowtime (hours)	Makespan (hours)	Avg. busy servers	Disc. Tpt.
3	MVA	0.1	[-1417.1, -1210.4] *	[-20.5, -17.6] *	[0.7, 0.8] *	[50.5, 56.7] *
3	MVA	0.7	[-1430.1, -1223.0] *	[-20.5, -17.6] *	[0.7, 0.8] *	[28.3, 31.7] *
3	PIH-G	0.1	[-1438.9, -1230.7] *	[-20.6, -17.7] *	[0.7, 0.8] *	[52.0, 58.3] *
3	PIH-G	0.7	[-1442.1, -1233.0] *	[-20.5, -17.6] *	[0.7, 0.8] *	[29.6, 33.3] *
4	MVA	0.1	[-1711.0, -1453.5] *	[-19.2, -16.5] *	[0.8, 0.9] *	[61.1, 69.2] *
4	MVA	0.7	[-1757.4, -1500.4] *	[-19.4, -16.7] *	[0.8, 0.9] *	[34.8, 38.6] *
4	PIH-G	0.1	[-1764.5, -1510.1] *	[-19.5, -16.8] *	[0.8, 0.9] *	[64.9, 72.6] *
4	PIH-G	0.7	[-1775.3, -1521.1] *	[-19.4, -16.7] *	[0.8, 0.9] *	[37.0, 40.3] *
5	MVA	0.1	[-1836.2, -1524.6] *	[-17.5, -14.9] *	[0.8, 0.9] *	[63.1, 73.6] *
5	MVA	0.7	[-1942.0, -1637.5] *	[-17.9, -15.4] *	[0.8, 1.0] *	[35.3, 39.2] *
5	PIH-G	0.1	[-1967.3, -1670.1] *	[-18.0, -15.6] *	[0.9, 1.0] *	[72.0, 81.3] *
5	PIH-G	0.7	[-1976.4, -1680.3] *	[-17.8, -15.3] *	[0.8, 0.9] *	[38.7, 41.9] *

Chapter 6

Conclusions

Emergency responders to mass-casualty incidents must make many decisions in a short amount of time and in a chaotic environment in an attempt to save as many patients as possible. Because these decisions must be made quickly, and often with limited information, it is important to have quantitative results to inform these decisions by suggesting the types of policies that would likely save the most lives. Some of the most important decisions that must be made in the aftermath of a mass-casualty incident involve patient prioritization and resource allocation. The questions of how to allocate the resources that respond to an incident and how to prioritize patients for access to those resources are closely related. Moreover, while current triage standards prescribe a policy for prioritizing patients (based on fixed priorities for the classes defined in Simple Triage and Rapid Treatment, or START), such standards generally do not deal with resource allocation during multiple incidents or with routing of casualties to multiple collection points or treatment facilities. Having established the need for quantitative analysis of decision making in mass-casualty incidents, we studied problems involving prioritization, transportation, and routing decisions for patients in a mass-casualty incident.

In Chapter 3, we studied the problem of prioritizing patients at a single-location mass-casualty incident. One of the most important contributions of Chapter 3 is the way in which patient criticality was modeled. While prior work modeled criticality with lifetimes, we used decreasing survival probability functions, which are more general and can more closely represent the reality, where the survival probability of a patient with critical injuries decreases

over time. We established that a simple fixed-priority policy is unlikely to maximize the expected number of survivors. Instead, we proposed a policy called Resource-based START (ReSTART), which is based on the analytical characterization of the optimal policy for a simple model. In developing this model, our objective was not to come up with the most realistic representation of the actual system but rather to obtain a formulation that captured the “right” trade-offs and was simple enough to allow analytical characterization of the optimal policy. Using a simulation study that provided a more realistic representation of the reality, we showed that this approach works quite well, since the results demonstrate the consistently good performance of ReSTART across many scenarios. In particular, we found that ReSTART results in significantly lower critical mortality on sets of randomly generated problem instances than the standard practice of static prioritization according to START classes.

In Chapter 4, we expanded the problem under consideration to include simultaneous multi-location mass-casualty incidents. Such incidents are commonly linked to extreme weather events or terrorist attacks. The START protocol does not prescribe how to allocate resources in these types of events. We found that to maximize the expected number of survivors, the decisions about patient prioritization and resource allocation must be made together (i.e., they depend on one another). An optimal solution to our simple model required changing the resources from one location to the other at most twice. Nonetheless, computing exactly when to switch priorities was more difficult than in the single location case. To overcome this obstacle, we provided upper and lower bounds for the threshold times and demonstrated in a simulation study that these estimates performed nearly as well as calculating the threshold times exactly. One important contribution of Chapter 4 is that we established decentralized heuristics that also worked well. Decentralized prioritization policies are those where the patient prioritization decision is made after the resource allocation decision. Even the heuristic that performed the worst, which used ReSTART after pre-allocating resources according to a simple (and clearly suboptimal) plan, significantly outperformed START. This result shows that even incremental changes in prioritization policies (such as replacing START with ReSTART) will still lead to a decrease in mortality among critical patients.

In Chapter 5, we studied the related problem of routing patients to one of several collection points or treatment facilities. This problem arises most commonly in larger mass-casualty events where patients are spread across a wide geographic area. In modeling the routing of patients, we took a somewhat different approach, using a Markov decision process to model the effects of queueing at the casualty collection points in order to maximize the expected total discounted number of patients treated. As in the previous chapters, we attempted not to model the problem very realistically, but rather to capture the general tradeoffs inherent in the problem. While it turned out to be very difficult to solve the MDP in general, we were able to use the model to prove several analytical results regarding the structure of the solution. Most importantly, our numerical study clearly showed the advantage of using a dynamic control policy (that is, a policy that bases the routing on the number of patients waiting in line at each facility). Even very good static policies performed quite poorly in some instances. Moreover, obtaining a reasonable dynamic policy turns out to be easy: calculating the policy prescribed by the Policy Improvement Heuristic is possible for problems with even a large number of casualty collection points, where the MDP itself is intractable. An even more practical heuristic, Marginal Value Approximation, yields a simple index, and it performs well under heavy discounting, i.e., when time criticality is high. These heuristics suggest that dynamic routing of patients may be both beneficial and practical in the aftermath of a mass-casualty incident.

Throughout this dissertation, we used formulations that led to insights about the solution structures that are simple yet effective. Using these insights, we suggested policies that could be implemented without necessitating large amounts of data or computing power. This aspect of the work is important because requiring providers to solve an optimization problem in real-time may be an obstacle to implementing a practical policy.

There are areas of future work that would be valuable both for better understanding of prioritization and resource allocation decisions in mass-casualty incidents and for developing policies to support these decisions. The first area of future work is the estimation problem related to the reward functions (survival probability functions). In Section 3.5.1, we demonstrated one way of coming up with estimates for our simulation study. Although that was the best we could do with what is available in the literature, we refrain from mak-

ing any strong claims regarding how realistic the resulting estimated survival probability functions are. In the simplest case, a study where experts are surveyed can be carried out to estimate survival probability functions corresponding to each START class, which can then be updated as new data become available with time. Once these estimates are obtained, the policy parameters of ReSTART can be more easily determined and be ready to be used in case of a mass-casualty incident. The second area of future research is to expand the analysis of Chapter 5 and integrate it with the results of Chapters 3 and 4. While each chapter suggests insights that can be useful in responding to mass-casualty incidents, they focus on different performance measures: in Chapters 3 and 4, we focused on minimizing critical mortality (i.e., maximizing the expected number of survivors from among the critical patients); in Chapter 5, we focused on maximizing the expected total discounted number of patients who complete treatment, with the aim of maximizing throughput. The second objective is more relevant to the patients who do not have critical injuries, as within this group it is less likely that different patients will have widely differing chances of survival. In reality, emergency planners must consider both objectives: maximizing the expected number of survivors from among the critical patients, and then distributing the non-critical patients in an efficient manner so they can be treated quickly. This effort is complicated by the fact that non-critical patients may self-transport to a treatment facility, reducing the amount of control that can be exercised by the decision maker. Such self-triage and transportation (i.e., exogenous arrivals to treatment facilities) is one important subject that we did not consider in this dissertation, and it should be included in future work.

It would be naïve to claim that any of the policies proposed in this dissertation are readily available for implementation in the way we described them. Any policy that would change the adopted practice needs to be scrutinized carefully by the medical community before being formally proposed as an alternative, and such scrutiny may necessitate some adjustments. Nevertheless, we believe that the structural properties of the models in this dissertation provide insights that can be useful in efforts to develop decision support tools that can be used in practice, and ultimately can lead to more lives being saved in the aftermath of a mass-casualty incident.

Appendix

Proofs of Analytical Results

Proof of Proposition 3.1. It is sufficient to show that there exists an optimal solution where only one class is served at a time almost everywhere. That is, the set of points over which more than one point is served simultaneously will have measure zero. Because changing the solution over a set of points with measure zero does not change the value of the objective function in (3.1), we can change the solution at these points so that only one class is served.

Denote the total expected reward associated with any solution S of (3.1) by $z(S)$. Consider any solution S such that there exist two classes of patients (without loss of generality, classes 0 and 1) served simultaneously over some set of positive measure. Denote by $Y \subseteq [0, T]$ the set of all points where classes 0 and 1 are served simultaneously under S . We will construct a solution \bar{S} such that $z(\bar{S}) \geq z(S)$ but in which the set of points where classes 0 and 1 are served simultaneously has measure zero.

Y can be partitioned into $Y_0 \cup Y_1 \cup \dots$, where Y_0 is a set of points of measure zero, and $\forall j \in \{1, 2, \dots\}$, Y_j is an open interval such that where *either* $f'_1(t) \leq f'_0(t)$ or $f'_1(t) \geq f'_0(t)$ for all $t \in Y_j$, and $r_i(t)$ is continuous over $t \in Y_j$ for $i = 0, 1$. Now, take any of the open intervals, $Y_j = (a, b)$, where $0 \leq a < b \leq T$ and $f'_1(t) \geq f'_0(t)$ for $t \in (a, b)$. That is, the reward gap function $g_{1,0}(t) = f_1(t) - f_0(t)$ is non-decreasing for $t \in (a, b)$. (The case where $f'_1(t) \leq f'_0(t)$ is symmetric, and hence its proof is omitted.)

Because $r_0(t)$ and $r_1(t)$ are continuous over (a, b) , there must exist $c \in (a, b)$ such that

$$\int_a^c r_1(t)dt = \int_c^b r_0(t)dt. \quad (\text{A.1})$$

In particular, such c must exist because $\lim_{c \rightarrow a} \int_a^c r_1(t)dt = 0$, $\lim_{c \rightarrow b} \int_c^b r_0(t)dt = 0$, the left-hand side of (A.1) is non-decreasing in c (because $r_1(t) \geq 0$), and the right-hand side of (A.1) is non-increasing in c (because $r_0(t) \geq 0$).

To construct \bar{S} , change the service during (a, b) as follows: during (a, c) serve class 0 at rate $r_0(t) + r_1(t)$, and during (c, b) serve class 1 at rate $r_0(t) + r_1(t)$. Note that the total amount of service to each class within (a, b) is unchanged, hence the constraints in (3.1) are satisfied for \bar{S} . Now, we have

$$\begin{aligned} z(\bar{S}) &= z(S) + \int_a^c r_1(t)f_0(t)dt + \int_c^b r_0(t)f_1(t)dt - \int_a^c r_1(t)f_1(t)dt - \int_c^b r_0(t)f_0(t)dt \\ &= z(S) + \int_c^b r_0(t)g_{1,0}(t)dt - \int_a^c r_1(t)g_{1,0}(t)dt \\ &\geq z(S) + g_{1,0}(c) \left(\int_c^b r_0(t)dt - \int_a^c r_1(t)dt \right) = z(S). \end{aligned}$$

Here the inequality follows because $r_i(t) \geq 0$ for $t \in (a, b)$ and $i = 0, 1$, and $g_{1,0}(t)$ is non-decreasing, which implies that $g_{1,0}(c) \geq g_{1,0}(t)$ for all $t \in [c, b]$ and $g_{1,0}(c) \leq g_{1,0}(t)$ for all $t \in [a, c]$. Finally, the last equation follows by (A.1). We can then repeat this construction for the remaining intervals over which more than one class is served under S . We conclude that for any solution that serves more than one class at any given time, there is another solution that performs at least as well by serving only one class at any point in time. The result immediately follows. \square

Proof of Proposition 3.2. Consider a solution \mathbf{W} , where class i is served at least partially before class j . Let $Y \subseteq W(i)$ be the set of all points such that $\forall s \in Y, \exists t \in W(j)$ where $s < t$. That is, Y is the set of points where i is served before j . Similarly, let $Z \subseteq W(j)$ be the set of points such that $\forall s \in Z, \exists t \in W(i)$ where $s > t$. That is, Z is the set of all points where j is served after i .

Partition $Y = Y_0 \cup Y_1 \cup \dots \cup Y_{m_i}$, where $1 \leq m_i < \infty$, such that Y_0 is a set of points having measure zero; $Y_k, k = 1, 2, \dots, m_i$, are open intervals; and $p > q \implies \forall s \in Y_p, t \in Y_q : s > t$. Similarly, partition $Z = Z_0 \cup Z_1 \cup \dots \cup Z_{m_j}$, where $1 \leq m_j < \infty$, such that Z_0 is a set of points having measure zero; $Z_k, k = 1, 2, \dots, m_j$, are open intervals; and $p > q \implies \forall s \in Z_p, t \in Z_q : s > t$. The fact that m_i and m_j are finite follows from our assumption that $W(i)$ contains only finitely many intervals.

Let $Y_1 \equiv (a_i, b_i)$ and $Z_{m_j} \equiv (a_j, b_j)$, where $0 \leq a_i < b_i \leq a_j < b_j \leq T$. Let $\epsilon =$

$\min\{b_i - a_i, b_j - a_j\}$. We will define a new solution $\bar{\mathbf{W}}$ that performs at least as well as \mathbf{W} :

$$\bar{W}(i) = (W(i) \setminus (a_i, a_i + \epsilon)) \cup (b_j - \epsilon, b_j)$$

$$\bar{W}(j) = (W(j) \setminus (b_j - \epsilon, b_j)) \cup (a_i, a_i + \epsilon)$$

$$\bar{W}(k) = W(k), \quad \forall k \in \mathcal{I} \setminus \{i, j\}.$$

Since the constraint set of (3.2) is satisfied for \mathbf{W} , and the construction of $\bar{\mathbf{W}}$ does not change the measure of any of the solution sets, except by adding and subtracting sets of the same measure (ϵ), then the constraints of (3.2) are satisfied for $\bar{\mathbf{W}}$.

Now, let $g(t) = f_j(t) - f_i(t)$, for all $t \in [0, T]$, and let $z(\mathbf{W})$ be the total expected reward obtained from using solution \mathbf{W} . Then,

$$\begin{aligned} z(\bar{\mathbf{W}}) &= z(\mathbf{W}) - \int_{a_i}^{a_i + \epsilon} f_i(t) dt + \int_{b_j - \epsilon}^{b_j} f_i(t) dt - \int_{b_j - \epsilon}^{b_j} f_j(t) dt + \int_{a_i}^{a_i + \epsilon} f_j(t) dt \\ &= z(\mathbf{W}) + \int_{a_i}^{a_i + \epsilon} (f_j(t) - f_i(t)) dt - \int_{b_j - \epsilon}^{b_j} (f_j(t) - f_i(t)) dt \geq z(\mathbf{W}) + \epsilon g(a_i + \epsilon) - \epsilon g(b_j - \epsilon) \\ &\geq z(\mathbf{W}), \end{aligned}$$

where the first inequality holds because $f_j'(t) \leq f_i'(t)$, and hence $g'(t) \leq 0$ for all $t \in [0, T]$, and the second inequality holds because $a_i + \epsilon \leq b_i \leq a_j \leq b_j - \epsilon$ and $g'(t) \leq 0$ for all $t \in [0, T]$.

Note that for $\bar{\mathbf{W}}$ we guarantee that either Y or Z will have at least one fewer open interval than \mathbf{W} . Hence, we will be able to repeat this procedure at most $m_i + m_j$ times until Y and Z are of measure zero. At that point we can set the service of points in Y and Z arbitrarily, because sets of points of measure zero do not affect the expected total reward. Then, the resulting solution will have a reward at least as large as $z(\mathbf{W})$ but without any service of class i before class j .

□

Proof of Proposition 3.3. To prove the result, we show that in any optimal solution to (3.5),

$W(D)$ is a single interval, plus possibly a set of zero-measure points. Now, suppose this is not true, that is, in the optimal solution there are at least two intervals contained in $W(D)$ with non-zero measure, such that the points between these two intervals are not in $W(D)$. In other words, there exist $0 \leq a_1 < b_1 < a_2 < b_2 \leq T$ such that $(a_1, b_1) \cup (a_2, b_2) \subseteq W(D)$, but $(b_1, a_2) \not\subseteq W(D)$. We will show that such a solution cannot be optimal. We must have one of the following three cases:

Case 1 ($t_m \leq b_1$): Let $\bar{W}(D) \equiv (W(D) \setminus (a_2, b_2)) \cup (b_1, b_1 + b_2 - a_2)$. Then, if we let $z(\mathbf{W})$ be the reward obtained by using solution \mathbf{W} , we have

$$\begin{aligned} z(\bar{\mathbf{W}}) &= z(\mathbf{W}) + \int_{b_1}^{b_1+b_2-a_2} g(t) dt - \int_{a_2}^{b_2} g(t) dt = z(\mathbf{W}) + \int_{b_1}^{b_1+b_2-a_2} (g(t) - g(t+a_2-b_1)) dt \\ &> z(\mathbf{W}), \end{aligned}$$

implying that \mathbf{W} is not optimal. Here, the inequality follows from the facts that $g(t)$ is decreasing in t for all $t > t_m$, $t_m \leq b_1$, and $b_1 < a_2 < b_2$.

Case 2 ($b_1 < t_m \leq a_2$): Let $\bar{W}(D) \equiv (W(D) \setminus (a_1, a_1 + t_m - b_1)) \cup (b_1, t_m)$. Then, we have

$$\begin{aligned} z(\bar{\mathbf{W}}) &= z(\mathbf{W}) + \int_{b_1}^{t_m} g(t) dt - \int_{a_1}^{a_1+t_m-b_1} g(t) dt = z(\mathbf{W}) + \int_{a_1}^{a_1+t_m-b_1} (-g(t) + g(t+b_1-a_1)) dt \\ &> z(\mathbf{W}), \end{aligned}$$

implying that \mathbf{W} is not optimal. Here, the inequality follows from the facts that $g(t)$ is increasing in t for all $t < t_m$ and $a_1 < b_1 < t_m$.

Case 3 ($a_2 < t_m$): Let $\bar{W}(D) \equiv (W(D) \setminus (a_1, b_1)) \cup (a_1 + a_2 - b_1, a_2)$. Then, we have

$$\begin{aligned} z(\bar{\mathbf{W}}) &= z(\mathbf{W}) + \int_{a_1+a_2-b_1}^{a_2} g(t) dt - \int_{a_1}^{b_1} g(t) dt = z(\mathbf{W}) + \int_{a_1}^{b_1} (-g(t) + g(t + a_2 - b_1)) dt \\ &> z(\mathbf{W}), \end{aligned}$$

implying that \mathbf{W} is not optimal. Here, the inequality follows from the facts that $g(t)$ is increasing in t for all $t < t_m$, $a_2 < t_m$, and $a_1 < b_1 < a_2$. \square

Proof of Proposition 3.4. (i) By the Fundamental Theorem of Calculus and the fact that $g(t)$ is continuous,

$$v'(t) = g(t + n_D) - g(t), \text{ and} \quad (\text{A.2})$$

$$v''(t) = g'(t + n_D) - g'(t), \text{ for } t \geq 0, \quad (\text{A.3})$$

where $v''(\cdot)$ is the second derivative of $v(t)$. Now, by Assumption 3.1, we have $v'(t) < 0$ for $t \geq t_m$, and if $t_m \geq n_D$, then $v'(t) > 0$ for $t \leq t_m - n_D$. Hence, any global maximizer of $v(t)$ over $[0, \infty)$ must be in $[\max\{0, t_m - n_D\}, t_m]$.

We next show that \tilde{t} is unique. Note that for all $t \in (\max\{0, t_m - n_D\}, t_m)$, $g'(t) > 0$ and $g'(t + n_D) < 0$. Then from Equation (A.3), we have $v''(t) < 0$ for $t \in (\max\{0, t_m - n_D\}, t_m)$. Hence, if $t_m > 0$, then there is a unique maximizer of $v(t)$ in $(\max\{0, t_m - n_D\}, t_m)$. Otherwise, $t_m = 0$ is the unique maximizer of $v(t)$. In summary, there is a unique maximizer $\tilde{t} = \arg \max_{t \in [0, \infty)} v(t)$, and we have $v'(t) > 0$ for $t < \tilde{t}$ and $v'(t) < 0$ for $t > \tilde{t}$.

(ii) If $\tilde{t} \leq n_I$, then \tilde{t} is also the global maximizer of $v(t)$ for the domain $[0, n_I]$, i.e., $t^* = \tilde{t}$. Otherwise, because $v'(t) > 0$ for $t < \tilde{t}$, we have $v(n_I) > v(t)$ for all $t < n_I$. Hence, n_I is the global maximizer of $v(t)$ for the domain $[0, n_I]$, i.e., $t^* = n_I$.

(iii) Because $t^* \leq \tilde{t}$ by part (ii) and $\tilde{t} \leq t_m$ by part (i), we have $t^* \leq t_m$. Moreover, Assumption 3.1 states that $t_m \leq n_I + n_D$. Using this assumption and the fact that $t_m \leq \tilde{t} + n_D$ by part (i), we conclude that $t_m \leq \min\{n_I + n_D, \tilde{t} + n_D\} = n_D + \min\{n_I, \tilde{t}\} = n_D + t^*$, where the last equation is due to part (ii). \square

Proof of Theorem 3.1. From Proposition 3.4, we know that \tilde{t} exists, is unique, and is in the

interval $[\max\{0, t_m - n_D\}, t_m]$. Because \tilde{t} is a maximizer of $v(t)$ over $[0, \infty)$, and $v'(t) = g(t + n_D) - g(t)$ is defined for all $t \geq 0$, then either \tilde{t} is a stationary point (i.e., it satisfies $v'(\tilde{t}) = 0$), or $v'(0) < 0$ and hence $\tilde{t} = 0$. The latter corresponds to case (i), where $t^* = 0$ by part (ii) of Proposition 3.4.

We now show that when \tilde{t} is a stationary point of $v(t)$, then exactly one of statements (ii) or (iii) must hold. Because \tilde{t} is a stationary point, by definition $v'(\tilde{t}) = 0$, or in other words, $g(\tilde{t}) = g(\tilde{t} + n_D)$. Furthermore, from the proof of part (i) of Proposition 3.4, we know that $g(T) - g(n_I) = v'(n_I) \geq 0$ if and only if $\tilde{t} \geq n_I$. Part (ii) of Proposition 3.4 completes the proof. \square

Proof of Proposition 3.5. Let $t^*(n_D)$ and $\tilde{t}(n_D)$ denote the values of t^* and \tilde{t} , respectively, when the number of class 1 patients is n_D . We first show that $\tilde{t}(n_D)$ either stays at zero (case 1) or decreases with n_D (case 2).

Case 1 ($\tilde{t}(n_D) = 0$): We will show that for any $\bar{n}_D > n_D$, $\tilde{t}(\bar{n}_D) = 0$. From Theorem 3.1, we know that $g(0) \geq g(n_D)$, and hence by Assumption 3.1 and the assumption that $n_D > 0$, we have $t_m < n_D$, which implies that $g(n_D) > g(\bar{n}_D)$. Hence, $g(0) > g(\bar{n}_D)$, which by Theorem 3.1 yields that $\tilde{t}(\bar{n}_D) = 0$.

Case 2 ($\tilde{t}(n_D) > 0$): We will show that for any $\bar{n}_D > n_D$, $\tilde{t}(\bar{n}_D) < \tilde{t}(n_D)$. From Proposition 3.4, we know that $\tilde{t}(\bar{n}_D) \leq t_m$. Therefore, it is sufficient to show that $\tilde{t}(\bar{n}_D)$ cannot be in the interval $[\tilde{t}(n_D), t_m]$. To do this, take any $t \in [\tilde{t}(n_D), t_m]$. We will show that t cannot be the maximizer of $v(t)$ when there are \bar{n}_D class 1 patients. We have $g(t) \geq g(\tilde{t}(n_D)) \geq g(\tilde{t}(n_D) + n_D) > g(\tilde{t}(n_D) + \bar{n}_D) \geq g(t + \bar{n}_D)$, where the first inequality follows from the facts that $g'(s) > 0$ for $s < t_m$ and $t \leq t_m$; the second inequality follows from Theorem 3.1; the third inequality follows from the facts that $g'(s) < 0$ for $s > t_m$, $\tilde{t}(n_D) + n_D \geq t_m$ by Proposition 3.4, and $\bar{n}_D > n_D$; and the final inequality follows from the facts that $g'(s) < 0$ for $s > t_m$ and $t \geq \tilde{t}(n_D)$. Hence, we conclude that $g(t) > g(t + \bar{n}_D)$ for any $t \in [\tilde{t}(n_D), t_m]$. However, Theorem 3.1 implies that the unique maximizer of $v(s)$ over $s \in [0, \infty)$ is either equal to zero or is a stationary point, which satisfies $g(s) = g(s + \bar{n}_D)$ when there are \bar{n}_D class 1 patients. Since in this case, neither one of these holds for any $t \in [\tilde{t}(n_D), t_m]$, we conclude that $t \in [\tilde{t}(n_D), t_m]$ cannot be the maximizer of $v(s)$ over

$s \in [0, \infty)$ when there are \bar{n}_D class 1 patients.

We showed that $\tilde{t}(n_D)$ either stays at zero or decreases with n_D . Thus, $t^*(n_D)$ either decreases or stays the same (at n_I or zero) as n_D increases since $t^*(n_D) = \min\{\tilde{t}(n_D), n_I\}$ by part (ii) of Proposition 3.4. \square

Proof of Proposition 4.1. Let $\mathcal{K} = \{I, D\} \times \mathcal{J}$ be a set of classes, and for each $(i, j) \in \mathcal{K}$, let the number of class- k patients be n_{ij}/s_j , and let the instantaneous reward earned by transporting class- k patients at time t is $s_j f_i(t)$. Then the resulting optimization formulation would be

$$\begin{aligned} \max_{\mathbf{r}(t), t \in [0, \infty)} \quad & \sum_{(i,j) \in \mathcal{K}} \int_0^\infty r_{ij}(x) s_j f_i(x) dx \\ \text{subject to} \quad & \sum_{(i,j) \in \mathcal{K}} r_{ij}(t) \leq 1, \quad \forall t \in [0, \infty) \\ & \int_0^\infty r_{ij}(x) dx = n_{ij}/s_j, \quad \forall (i, j) \in \mathcal{K}. \end{aligned} \tag{A.4}$$

By multiplying both sides by s_j in the last constraint of (A.4), we observe that (A.4) is equivalent to (4.1). \square

Proof of Corollary 4.1. Proposition 4.1 demonstrates that we can treat each class-incident pair as a class. Proposition 3.1 demonstrated that in such a problem, there exists an optimal solution where only one class of patients is served at any time, which is sufficient to show in this problem that only one class-incident pair is served at any time. \square

Proof of Proposition 4.3. Recall that we assume that the set of time points during which a given class of patients is transported is the union of finitely many intervals, each of which is closed on the left and open on the right. We will show that any solution in which $W(D)$ contains two or more disjoint intervals cannot be an optimal solution.

Take any solution where $W(D)$ contains two or more disjoint intervals. Then there exist a, b, c, d , such that $a < b < c < d$, $[a, b) \in W(D)$, and $[c, d) \in W(D)$ but no interval contained in $[b, c)$ is in $W(D)$. If any interval contained in $[b, c)$ is in $W(\underline{D})$, then Proposition 4.2 immediately applies and \mathbf{W} is not optimal. Otherwise, there must be $[u, v) \subseteq [b, c)$,

such that $v - u \leq b - a$, $v - u \leq d - c$, and $[u, v)$ is in either $W(I)$ or $W(\underline{I})$. We complete the proof for the case where $[u, v) \in W(I)$. The proof in the other case proceeds by an identical argument.

If $t_m < u$, let $\tilde{W}(D) = (W(D) \cup [u, v)) \setminus [d - (v - u), d)$ and $\tilde{W}(I) = (W(I) \cup [d - (v - u), d)) \setminus [u, v)$. Let $\tilde{W}(\underline{D}) = W(\underline{D})$ and $\tilde{W}(\underline{I}) = W(\underline{I})$. Then,

$$\begin{aligned} v(\tilde{\mathbf{W}}) &= v(\mathbf{W}) + \int_u^v g(t)dt - \int_{d-(v-u)}^d g(t)dt \\ &= v(\mathbf{W}) + \int_u^v (g(t) - g(t + d - v))dt > v(\mathbf{W}), \end{aligned}$$

so \mathbf{W} cannot be an optimal solution. Above, the inequality follows from the fact that $t_m < u < v < d$.

If $t_m \geq v$, let $\tilde{W}(D) = (W(D) \cup [u, v)) \setminus [a, a + (v - u))$ and $\tilde{W}(I) = (W(I) \cup [a, a + (v - u)) \setminus [u, v)$. If $t_m \in [u, v)$, let $\tilde{W}(D) = (W(D) \cup [t_m, v)) \setminus [a, a + (v - t_m))$ and $\tilde{W}(I) = (W(I) \cup [a, a + (v - t_m)) \setminus [t_m, v)$. In either case, computations analogous to the above will produce the same conclusion, that \mathbf{W} cannot be an optimal solution. We conclude that if $W(D)$ contains two or more disjoint intervals, \mathbf{W} cannot be an optimal solution. Finally, we note that an analogous argument can be completed for $W(\underline{D})$. \square

Proof of Lemma 4.1. Suppose not. Then $t_2 < t_1$, and hence $t_2 + c_2 < t_1 + d_1$ because $c_2 \leq c_1$. To satisfy (4.3), we must have $t_m > t_1$, and to satisfy (4.4), we must have $t_m < t_2 + c_2$. Hence, t_m , the maximizer of g , must be in the interval $(t_1, t_2 + c_2)$. But then, by the assumption of the lemma, we must have $g(t_2) < g(t_1)$ and $g(t_2 + c_2) > g(t_1 + c_1)$. Together with (4.3) and (4.4), this implies that

$$g(t_1 + c_1) > g(t_2) \geq g(t_2 + c_2) > g_{ij}(t_1 + c_1) \geq g(t_1),$$

which is a contradiction. \square

Proof of Proposition 4.4. First, assume that $t^* > 0$. Otherwise, the result is immediate. Recall our assumption that the set of time points in which a given class of patients is transported is a union of intervals, each of which is closed on the left and open on the right.

Therefore, for sufficiently small $\epsilon > 0$, the interval $[t^* - \epsilon, t^*)$ is in either $W(I)$ or $W(\underline{I})$, depending on which class of immediate patients is transported just prior to t^* .

Consider the first case, where $[t^* - \epsilon, t^*) \in W(I)$. Let \mathbf{W} be the optimal solution to (4.2), and let $\bar{\mathbf{W}}$ be an solution constructed by shifting the transportation of the class D patients slightly earlier by an amount of time $\epsilon \geq 0$, while the class I patients who would have been transported in $[t^* - \epsilon, t^*)$ are delayed until after the transportation of class D patients. Then, letting $g(t) = f_D(t) - f_I(t)$,

$$\begin{aligned} v(\bar{\mathbf{W}}) &= v(\mathbf{W}) + \int_{t^* - \epsilon}^{t^*} sg(t)dt - \int_{t^* + n_D/s - \epsilon}^{t^* + n_D/s} sg(t)dt \\ &\leq v(\mathbf{W}), \end{aligned}$$

where the inequality follows from the optimality of \mathbf{W} . This implies that

$$\int_{t^* - \epsilon}^{t^*} sg(t)dt - \int_{t^* + n_D/s - \epsilon}^{t^* + n_D/s} sg(t)dt \leq 0.$$

Denote the left-hand side of the above inequality, which is change in total reward by performing the shift of size ϵ , by $H(\epsilon)$, and denote the right-derivative of $H(\epsilon)$ with respect to ϵ by $H^+(\epsilon)$. Note that $H^+(\epsilon) = sg(t^* - \epsilon) - sg(t^* + n_D/s - \epsilon)$. Now $H(0) = 0$, and $H(\epsilon) \leq 0$ for all $\epsilon \geq 0$ (because \mathbf{W} is optimal), so it must be the case that $H^+(0) \leq 0$, or in other words,

$$sg(t^*) \leq sg(t^* + n_D/s). \quad (\text{A.5})$$

By repeating the argument for the second case, where $[t^* - \epsilon, t^*) \in W(\underline{I})$ and denoting $\bar{g}(t) = sf_D(t) - f_I(t)$, we obtain

$$\bar{g}(t^*) \leq \bar{g}(t^* + n_D/s). \quad (\text{A.6})$$

Hence, we conclude that t^* must satisfy either (A.5) or (A.6) depending on the interval in which $[t^* - \epsilon, t^*)$ lies.

Now we turn to the upper bound problem. Denote by t^u the time at which, in the optimal solution to the upper bound problem, we would transport the class 2 patients even

if class 1 patients remain. For now, assume that $0 < t^u < n_{\underline{I}} + n_I/s$. If $t^u = n_{\underline{I}} + n_I/s$ then the result is trivial. In the special case where $t^u = 0$, then Theorem 3.1 states that $sg(t^u) \geq sg(t^u + n_D/s)$. Otherwise, By Theorem 3.1, $sg(t^u) = sg(t^u + n_D/s)$. Either way,

$$sg(t^u) \geq sg(t^u + n_D/s). \quad (\text{A.7})$$

Furthermore, it is easy to check that (A.7) implies

$$\bar{g}(t^u) \geq \bar{g}(t^u + n_D/s). \quad (\text{A.8})$$

We conclude by applying Lemma 4.1, since both $g(\cdot)$ and $\bar{g}(\cdot)$ satisfy the assumptions of the lemma: in the case where $[t^* - \epsilon, t^*) \in W(I)$, (A.5) holds. In that case, letting $c_1 = c_2 = n_D/s$, Lemma 4.1 states that (A.7) implies $t^* \leq t^u$. In the case where $[t^* - \epsilon, t^*) \in W(\underline{I})$, (A.6) holds. By the same argument, (A.8) implies $t^* \leq t^u$.

We have now established that t^u is an upper bound to t^* . Now, by Proposition 3.4, we have that for some $\phi \in [0, 1]$, $t^u = \min\{n_I/s + n_{\underline{I}}, t_m - \phi n_D/s\}$, and hence $t^u \leq \min\{n_I/s + n_{\underline{I}}, t_m\}$, because zero is the most conservative value for ϕ . \square

Proof of Corollary 4.2. Suppose not. Then in an optimal solution, transportation of class D patients does not begin until after n_I/s . Hence, for sufficiently small $\epsilon > 0$ $[t^* - \epsilon, t^*) \in W(\underline{I})$ and from the proof of Proposition 4.4, t^* must satisfy (A.6). However, (A.6) cannot hold because $\bar{g}'(t) < 0$ for all $t > \underline{t}_m$, and $\underline{t}_m \leq n_I/s \leq t^*$. \square

Proof of Proposition 4.5. If $t^* = n_I/s + n_{\underline{I}}$, then the solution to the lower bound problem is immediately a lower bound on t^* : transportation of delayed patients must begin no later than time $n_I + n_{\underline{I}}/s$, or else there would be idling in the optimal policy, which is ruled out by the fact that reward functions are non-increasing. For the same reason, transportation of class 2 patients in the lower bound problem must begin by this time, so the value t^u , which is the time at which transportation of class 2 patients begins, will be no larger than $n_I + n_{\underline{I}}/s$.

We now consider the case where $t^* < n_I/s + n_{\underline{I}}$. Recall our assumption that the set of time points in which a given class of patients is transported is a union of intervals, each of

which is closed on the left and open on the right. Therefore, for sufficiently small $\epsilon > 0$, the interval $[t^* + n_D/s, t^* + n_D/s + \epsilon)$ is in exactly one of $W(I)$, $W(\underline{I})$, or $W(\underline{D})$, depending on which class of patients is transported just prior after class D .

Take the first case, where $[t^* + n_D/s, t^* + n_D/s + \epsilon) \in W(I)$. Let \mathbf{W} be the optimal solution to the problem, and let $\bar{\mathbf{W}}$ be the solution constructed by shifting the transportation of the class D patients slightly later by an amount ϵ , while in exchange the class I patients who would have been transported during $[t^* + n_D/s, t^* + n_D/s + \epsilon)$ are instead transported just before the class D patients. Then

$$\begin{aligned} v(\bar{\mathbf{W}}) &= v(\mathbf{W}) - \int_{t^*}^{t^* + \epsilon} sg(t)dt + \int_{t^* + n_D/s}^{t^* + n_D/s + \epsilon} sg(t)dt \\ &\leq v(\mathbf{W}), \end{aligned}$$

where the inequality follows by the optimality of \mathbf{W} . This implies that

$$-\int_{t^*}^{t^* + \epsilon} sg(t)dt + \int_{t^* + n_D/s}^{t^* + n_D/s + \epsilon} sg(t)dt \leq 0$$

Denote the left-hand side of the above inequality, which is change in total reward by performing the shift of size ϵ , by $H(\epsilon)$, and denote the right-derivative of $H(\epsilon)$ with respect to ϵ by $H^+(\epsilon)$. Note that $H^+(\epsilon) = -sg(t^* + \epsilon) + sg(t^* + n_D/s + \epsilon)$. Because $H(0) = 0$ and $H(\epsilon) \leq 0$ for all $\epsilon \geq 0$, it must be the case that $H^+(0) \leq 0$, or in other words,

$$sg(t^*) \geq sg(t^* + n_D/s). \quad (\text{A.9})$$

By repeating the argument for the case where $[t^* + n_D/s, t^* + n_D/s + \epsilon) \in W(\underline{I})$, we obtain

$$\bar{g}(t^*) \geq \bar{g}(t^* + n_D/s). \quad (\text{A.10})$$

Now we examine the case where $[t^* + n_D/s, t^* + n_D/s + \epsilon) \in W(\underline{D})$. We break this case into two subcases. For sufficiently small $\delta > 0$, $[t^* + n_D/s + n_{\underline{D}}, t^* + n_D/s + n_{\underline{D}} + \delta)$ is contained in either $W(I)$ or $W(\underline{I})$; i.e., either class I or class \underline{I} is transported after completing the transportation of class \underline{D} patients.

In the first subcase, let $\bar{\mathbf{W}}$ be the solution constructed by shifting the transportation of the class D and class \underline{D} patients slightly later by an amount $\delta \geq 0$, while in exchange some of the class I patients are transported just before the class D patients. Then, denoting $\underline{g}(t) = f_D(t) - sf_I(t)$,

$$\begin{aligned} v(\bar{\mathbf{W}}) &= v(\mathbf{W}) - \int_{t^*}^{t^*+\delta} sg(t)dt + \int_{t^*+n_D/s}^{t^*+n_D/s+\delta} (sg(t) - \underline{g}(t))dt + \int_{t^*+n_{\underline{D}}+n_D/s}^{t^*+n_{\underline{D}}+n_D/s+\delta} \underline{g}(t)dt \\ &\leq v(\mathbf{W}), \end{aligned}$$

where the inequality follows by the optimality of \mathbf{W} . This implies that

$$-\int_{t^*}^{t^*+\delta} sg(t)dt + \int_{t^*+n_D/s}^{t^*+n_D/s+\delta} (sg(t) - \underline{g}(t))dt + \int_{t^*+n_{\underline{D}}+n_D/s}^{t^*+n_{\underline{D}}+n_D/s+\delta} \underline{g}(t)dt \leq 0$$

Denote the left hand side of the above inequality as $G(\delta)$, and denote its right-derivative with respect to δ as $G^+(\delta)$. $G^+(\delta) = -sg(t^* + \delta) + sg(t^* + n_D/s + \delta) - \underline{g}(t^* + n_D/s + \delta) + \underline{g}(t^* + n_{\underline{D}} + n_D/s + \delta)$. Again, it must be the case that $G^+(0) \leq 0$, or in other words,

$$sg(t^*) + \underline{g}(t^* + n_D/s) \geq sg(t^* + n_D/s) + \underline{g}(t^* + n_{\underline{D}} + n_D/s).$$

Using the definitions of $g(\cdot)$ and $\underline{g}(\cdot)$ and rearranging some terms, we can re-write the above inequality as

$$sf_D(t^*) - sf_D(t^* + n_{\underline{D}}) + f_D(t^* + n_D/s) - f_D(t^* + n_D/s + n_{\underline{D}}) \geq sf_I(t^*) - sf_I(t^* + n_D/s + n_{\underline{D}}).$$

The expression $f_D(t^* + n_D/s) - f_D(t^* + n_D/s + n_{\underline{D}})$ is non-negative because the reward function $f_D(\cdot)$ is nonincreasing, so we can multiply it by s , which is greater than one, without invalidating the inequality. Therefore,

$$sf_D(t^*) - sf_D(t^* + n_D/s) + sf_D(t^* + n_D/s) - sf_D(t^* + n_D/s + n_{\underline{D}}) \geq sf_I(t^*) - sf_I(t^* + n_D/s + n_{\underline{D}}),$$

or in other words,

$$sf_D(t^*) - sf_D(t^* + n_D/s + n_{\underline{D}}) \geq sf_I(t^*) - sf_I(t^* + n_D/s + n_{\underline{D}}).$$

By applying the definition of $g(\cdot)$, we re-write the above inequality as $g(t^*) \geq g(t^* + n_{\underline{D}} + n_D/s)$, or equivalently

$$g(t^*) \geq g(t^* + n_{\underline{D}} + n_D/s). \quad (\text{A.11})$$

By repeating the argument for the second subcase, where $[t^* + n_D/s + n_{\underline{D}}, t^* + n_D/s + n_{\underline{D}} + \delta) \in W(\underline{I})$, we obtain

$$\bar{g}(t^*) \geq \bar{g}(t^* + n_D/s + n_{\underline{D}}). \quad (\text{A.12})$$

Hence, we conclude that t^* must satisfy exactly one of (A.9), (A.10), (A.11), or (A.12).

At this point, we turn to analyzing the lower bound problem. By definition, t^l , the time at which we begin transporting class 2 patients in the lower bound problem, must be in $[0, n_I/s + n_{\underline{I}}]$. If $t^l = 0$, then the result is immediate. Otherwise, Theorem 3.1 states that

$$\bar{g}(t^l) \leq \bar{g}(t^l + n_D/s + n_{\underline{D}}), \quad (\text{A.13})$$

Furthermore, it is easy to check that (A.13) also implies

$$g(t^l) \leq g(t^l + n_D/s + n_{\underline{D}}). \quad (\text{A.14})$$

We conclude by applying Lemma 4.1: if (A.9) holds, then letting $c_2 = n_D/s$ and $c_1 = n_{\underline{D}} + n_D/s$, Lemma 4.1 states that (A.14) implies $t^l \leq t^*$; if (A.10) holds, then by the same argument, (A.13) implies $t^l \leq t^*$; if (A.12) holds, then letting $c_1 = c_2 = n_{\underline{D}} + n_D/s$, Lemma 4.1 states that (A.14) implies $t^l \leq t^*$; finally, if (A.12) holds, then by the same argument, (A.13) implies $t^l \leq t^*$.

We have now established that t^l is a lower bound to t^* . Now, by Proposition 3.4, we have that for some $\phi \in [0, 1]$, $t^l = \min\{n_I/s + n_{\underline{I}}, \underline{t}_m - \phi(n_D/s + n_{\underline{D}})\}$, and hence $t^l \geq \min\{n_I/s + n_{\underline{I}}, \underline{t}_m - (n_D/s + n_{\underline{D}})\}$, because in this case, one is the most conservative value for ϕ . \square

Proof of Proposition 4.6. First, assume that $t^{**} > t^* + n_D/s$. Otherwise, the solution to the upper bound problem for t^{**} immediately is an upper bound on t^{**} . Now, for sufficiently small $\epsilon > 0$, the interval $[t^{**} - \epsilon, t^{**})$ is in either $W(I)$ or $W(\underline{I})$, because we know that $\sup W(D) = t^* + n_D/s$. Consider the case where $[t^{**} - \epsilon, t^{**}) \in W(I)$.

Let \mathbf{W} be the optimal solution to (4.2), and let $\bar{\mathbf{W}}$ be the solution constructed by shifting the service of the class \underline{D} patients slightly earlier by an amount $\epsilon \geq 0$, while the class I patients who would otherwise have been transported during $[t^{**} - \epsilon, t^{**})$ are delayed until after the transportation of class \underline{D} patients. Then

$$\begin{aligned} v(\bar{\mathbf{W}}) &= v(\mathbf{W}) + \int_{t^{**}-\epsilon}^{t^{**}} \underline{g}(t) dt - \int_{t^{**}+n_{\underline{D}}-\epsilon}^{t^{**}+n_{\underline{D}}} \underline{g}(t) dt \\ &\leq v(\mathbf{W}), \end{aligned}$$

where the inequality follows by the optimality of \mathbf{W} . This implies that

$$\int_{t^{**}-\epsilon}^{t^{**}} \underline{g}(t) dt - \int_{t^{**}+n_{\underline{D}}-\epsilon}^{t^{**}+n_{\underline{D}}} \underline{g}(t) dt \leq 0.$$

Denote the left hand side of the above equation by $H(\epsilon)$, and denote its right-derivative with respect to ϵ by $H^+(\epsilon)$. Then $H^+(\epsilon) = \underline{g}(t^{**} - \epsilon) - \underline{g}(t^{**} + n_{\underline{D}} - \epsilon)$. Now $H(0) = 0$, and $H(\epsilon) \leq 0$ for all $\epsilon \geq 0$, so it must be the case that $H^+(0) \leq 0$, or in other words,

$$\underline{g}(t^{**}) \leq \underline{g}(t^{**} + n_{\underline{D}}). \quad (\text{A.15})$$

By repeating the argument for the other case (where $[t^{**} - \epsilon, t^{**}) \in W(\underline{I})$), we obtain

$$g(t^{**}) \leq g(t^{**} + n_{\underline{D}}). \quad (\text{A.16})$$

Hence, we conclude that t^{**} must satisfy either (A.15) or (A.16).

Now, consider the upper bound problem for t^{**} . Denote the time at which we begin transporting class 2 patients in the upper bound problem by t^u . Assume that $t^u < n_{\underline{I}} +$

$n_I/s + n_D/s$; otherwise t^u is a trivial upper bound to t^{**} . Theorem 3.1 states that

$$\underline{g}(t^u) \geq \underline{g}(t^u + n_{\underline{D}}), \quad (\text{A.17})$$

which also implies

$$g(t^u) \geq g(t^u + n_{\underline{D}}). \quad (\text{A.18})$$

We conclude by applying Lemma 4.1: if (A.15) holds, then letting $c_1 = c_2 = n_{\underline{D}}$, Lemma 4.1 states that (A.17) implies $t^{**} \leq t^u$; on the other hand, if (A.16) holds, then by the same argument, (A.18) implies $t^{**} \leq t^u$.

We have now established that t^u is an upper bound to t^{**} . Now, by Proposition 3.4, we have that for some $\phi \in [0, 1]$, $t^u = \bar{t}_m - \phi n_{\underline{D}}$, and hence $t^u \leq \bar{t}_m$, because zero is the most conservative value for ϕ .

□

Proof of Proposition 4.7. First, assume that $t^{**} < n_{\underline{I}} + n_I/s + n_D/s$. Otherwise, the solution to the lower bound problem is immediately a lower bound to t^{**} . Now, for sufficiently small $\epsilon > 0$, the interval $[t^{**} + n_{\underline{D}}, t^{**} + n_{\underline{D}} + \epsilon)$ is in either $W(I)$ or $W(\underline{I})$. Consider the case where $[t^{**} + n_{\underline{D}}, t^{**} + n_{\underline{D}} + \epsilon) \in W(I)$.

Let \mathbf{W} be the optimal solution to (4.2), and let $\bar{\mathbf{W}}$ be the solution constructed by shifting the service of the class \underline{D} patients later by an amount $\epsilon \geq 0$, while the class I patients who would have been transported in $[t^{**} + n_{\underline{D}}, t^{**} + n_{\underline{D}} + \epsilon)$ are instead transported before the class \underline{D} patients. Then

$$\begin{aligned} v(\bar{\mathbf{W}}) &= v(\mathbf{W}) - \int_{t^{**}}^{t^{**}+\epsilon} \underline{g}(t)dt + \int_{t^{**}+n_{\underline{D}}}^{t^{**}+n_{\underline{D}}+\epsilon} \underline{g}(t)dt \\ &\leq v(\mathbf{W}), \end{aligned}$$

where the inequality follows by the optimality of \mathbf{W} . This implies that where

$$-\int_{t^{**}}^{t^{**}+\epsilon} \underline{g}(t)dt + \int_{t^{**}+n_{\underline{D}}}^{t^{**}+n_{\underline{D}}+\epsilon} \underline{g}(t)dt \leq 0$$

Denote the left hand side of the above inequality by $H(\epsilon)$, and denote its right-derivative with respect to ϵ by $H^+(\epsilon)$, which is equal to $-\underline{g}(t^{**} + \epsilon) + \underline{g}(t^{**} + n_{\underline{D}} + \epsilon)$. Now $H(0) = 0$, and $H(\epsilon) \leq 0$ for all $\epsilon \geq 0$, so it must be the case that $H^+(0) \leq 0$, or in other words,

$$\underline{g}(t^{**}) \geq \underline{g}(t^{**} + n_{\underline{D}}). \quad (\text{A.19})$$

By repeating the argument for the other case (where $[t^{**} + n_{\underline{D}}, t^{**} + n_{\underline{D}} + \epsilon) \in W(\underline{I})$), we obtain

$$g(t^{**}) \geq g(t^{**} + n_{\underline{D}}). \quad (\text{A.20})$$

Hence, we conclude that t^{**} must satisfy either (A.19) or (A.20).

Now we consider the lower bound problem. Let t^l be the time at which we begin transporting class 2 patients in the lower bound problem. We will show that t^l is a lower bound to t^{**} . If $t^l = t^* + n_D/s$ then the result is trivial. Otherwise, Theorem 3.1 states that

$$g(t^l) \leq g(t^l + n_{\underline{D}}). \quad (\text{A.21})$$

Furthermore, it is easy to check that (A.21) implies

$$\underline{g}(t^l) \leq \underline{g}(t^l + n_{\underline{D}}). \quad (\text{A.22})$$

We conclude by applying Lemma 4.1: if (A.19) holds, then letting $c_1 = c_2 = n_{\underline{D}}$, Lemma 4.1 states that (A.21) implies $t^{**} \geq t^l$; on the other hand, if (A.20) holds, then by the same argument, (A.22) implies $t^{**} \geq t^l$.

We have now established that t^l is an upper bound to t^{**} . Now, by Proposition 3.4, we have that for some $\phi \in [0, 1]$, $t^l = t_m - \phi n_{\underline{D}}$, and hence $t^l \geq t_m - n_{\underline{D}}$, because one is the most conservative value for ϕ . \square

Proof of Proposition 5.1. For $x \in \mathcal{Q}$, if station j is (x, k) -preferable to station i , then

$$\tau_{kj} M_j(x) \geq \tau_{ki} M_i(x). \quad (\text{A.23})$$

By the condition that

$$\frac{\tau_{ij}}{\tau_{il}} \geq \frac{\tau_{kj}}{\tau_{kl}},$$

or equivalently,

$$\tau_{kl} \geq \frac{\tau_{kj}\tau_{il}}{\tau_{ij}},$$

we have

$$\tau_{kl}M_l(x) \geq \frac{\tau_{kj}\tau_{il}}{\tau_{ij}}M_l(x), \quad (\text{A.24})$$

because $M_l(x)$ is non-negative (a fact whose proof we postpone until Proposition 5.2). By combining (A.23) and (A.24), we conclude that

$$\tau_{ij}M_j(x) \geq \tau_{il}M_l(x),$$

which implies that in station i is (x, l) -preferable to station j . \square

Proof of Proposition 5.2. We first prove (5.3) by induction. The base case is that $V^0(x + e_j) \geq V^0(x)$ for all $x \in \mathcal{Q}$, which is trivial because $V^0(x) = 0$ for all $x \in \mathcal{Q}$. We now show that if $V^n(x + e_j) \geq V^n(x)$ for all $j \in \mathcal{S}$ and for all x then the same will hold for $n + 1$. From (5.2), we have

$$\begin{aligned} (1 + \alpha)V^{n+1}(x + e_j) &= r(x + e_j) + \sum_{k \in \mathcal{S}} \mu_k V^n(x + e_j - I_k(x + e_j)e_k) \\ &\quad + \sum_{l \in \mathcal{R}} \max_{k \in \mathcal{S}} \{ \tau_{lk} V^n(x + e_j + e_k) + (\tau_l - \tau_{lk}) V^n(x + e_j) \} \\ &\geq r(x) + \sum_{k \in \mathcal{S}} \mu_k V^n(x - I_k(x)e_k) \\ &\quad + \sum_{l \in \mathcal{R}} \max_{k \in \mathcal{S}} \{ \tau_{lk} V^n(x + e_k) + (\tau_l - \tau_{lk}) V^n(x) \} \\ &= (1 + \alpha)V^n(x), \end{aligned}$$

where the inequality follows from the fact that $r(x + e_j) \geq r(x)$ for all $j \in \mathcal{S}$ and for all $x \in \mathcal{Q}$, the fact that $\tau_l \geq \tau_{lk}$ for all $l \in \mathcal{R}$ and $k \in \mathcal{S}$, and the inductive hypothesis. Note that in the above inequality, we apply the inductive hypothesis to the terms in the first

summation—namely, $V^n(x + e_j - I_k(x + e_j)e_k)$ —for every case where $j \neq k$ or $j = k$ and $x_j > 0$. In the case where $j = k$ and $x_j = 0$, $V^n(x + e_j - I_k(x + e_j)e_k)$ is equivalent to $V^n(x + e_j - e_j)$, which is $V^n(x)$; moreover, $V^n(x - I_k(x)e_k)$ is also equivalent to $V^n(x)$. Hence, applying the inductive hypothesis to this term is not necessary.

We now prove (5.4) by induction. The base case states that $V^0(x) + r_j \geq V^0(x + e_j)$ for all $j \in \mathcal{S}$ and for all x , which is trivial because $V^0(x) = 0$ for all x and $r_j > 0$. We now show that if $V^n(x) + r_j \geq V^n(x + e_j)$ for all $j \in \mathcal{S}$ and for all x then the same will hold for $n + 1$.

$$\begin{aligned}
(1 + \alpha)(V^{n+1}(x) + r_j) &= r(x) + \sum_{k \in \mathcal{S} \setminus \{j\}} \mu_k V^n(x - I_k(x)e_k) + \mu_j V^n(x - I_j(x)e_j) + \tau V^n(x) \\
&\quad + \sum_{l \in \mathcal{R}} \max_{k \in \mathcal{S}} \{ \tau_{lk} (V^n(x + e_k) - V^n(x)) \} + (1 + \alpha)r_j \\
&\geq r(x) + \sum_{k \in \mathcal{S} \setminus \{j\}} \mu_k V^n(x - I_k(x)e_k) + \mu_j V^n(x - I_j(x)e_j) \\
&\quad + \sum_{l \in \mathcal{R}} \max_{k \in \mathcal{S}} \{ \tau_{lk} V^n(x + e_k) + (\tau_l - \tau_{lk}) V^n(x) \} + r_j,
\end{aligned}$$

since $\alpha > 0$. Using the fact that $\beta = \sum_j \mu_j + \tau = 1$, we distribute r_j to each term and obtain

$$\begin{aligned}
(1 + \alpha)(V^{n+1}(x) + r_j) &\geq r(x) + \sum_{k \in \mathcal{S} \setminus \{j\}} \mu_k (V^n(x - I_k(x)e_k) + r_j) + \mu_j (V^n(x - I_j(x)e_j) + r_j) \\
&\quad + \sum_{l \in \mathcal{R}} \max_{k \in \mathcal{S}} \{ \tau_{lk} (V^n(x + e_k) + r_j) + (\tau_l - \tau_{lk}) (V^n(x) + r_j) \}.
\end{aligned}$$

We now apply the inductive hypothesis to every term except one, as well as the fact that

$I_k(x + e_j) = I_k(x)$ for all $k \neq j$, to obtain

$$\begin{aligned}
(1 + \alpha)(V^{n+1}(x) + r_j) &\geq r(x) + \sum_{k \in \mathcal{S} \setminus \{j\}} \mu_k(V^n(x + e_j - I_k(x + e_j)e_k)) + \mu_j(V^n(x - I_j(x)e_j) + r_j) \\
&\quad + \sum_{l \in \mathcal{R}} \max_{k \in \mathcal{S}} \{ \tau_{lk}(V^n(x + e_j + e_k)) + (\tau_l - \tau_{lk})(V^n(x + e_j)) \}.
\end{aligned} \tag{A.25}$$

Now, for the third term on the right-hand side of (A.25), we apply the inductive hypothesis *only* in the case where $x_j > 0$, in which case $I_j(x + e_j) = I_j(x)$, and we obtain

$$\begin{aligned}
(1 + \alpha)(V^{n+1}(x) + r_j) &\geq r(x) + \sum_{k \in \mathcal{S} \setminus \{j\}} \mu_k(V^n(x + e_j - I_k(x + e_j)e_k)) \\
&\quad + \mu_j(V^n(x + e_j - I_j(x + e_j)e_j) + r_j(I_j(x + e_j) - I_j(x))) \\
&\quad + \sum_{l \in \mathcal{R}} \max_{k \in \mathcal{S}} \{ \tau_{lk}(V^n(x + e_j + e_k)) + (\tau_l - \tau_{lk})(V^n(x + e_j)) \}.
\end{aligned} \tag{A.26}$$

Note that in the case where $x_j = 0$, (A.26) is simply a re-writing of (A.25), since in that case, $V^n(x + e_j - I_j(x + e_j)e_j) = V^n(x - I_j(x)e_j)$ and $I_j(x + e_j) - I_j(x) = 1$.

Finally, applying the fact that $r(x + e_j) = r(x) + \mu_j r_j(I_j(x + e_j) - I_j(x))$, we end up with

$$\begin{aligned}
(1 + \alpha)(V^{n+1}(x) + r_j) &\geq r(x + e_j) + \sum_{k \in \mathcal{S}} \mu_k V^n(x + e_j - e_k) \\
&\quad + \sum_{l \in \mathcal{R}} \max_{k \in \mathcal{S}} \{ \tau_{lk} V^n(x + e_j + e_k) + (\tau_l - \tau_{lk}) V^n(x + e_j) \} \\
&= (1 + \alpha) V^{n+1}(x + e_j).
\end{aligned}$$

□

Proof of Proposition 5.3. We prove the proposition by induction. The base case requires that the inequalities hold for $V^0(\cdot)$, which is trivial because $V^0(\cdot)$ is always zero. Suppose that (5.5) and (5.6) hold for all $j \in \mathcal{S}$ for some arbitrary $n \geq 0$. We will proceed by showing

that when $x_1 \leq x_j \forall j \in \mathcal{S}$, the same inequalities hold for $V^{n+1}(\cdot)$ for all $j \in \mathcal{S}$.

Inductive step for (5.5). Case 1: $0 < x_1 < x_j$.

$$(1 + \alpha)V^{n+1}(x + e_j) = r(x + e_j) + \sum_{k \in \mathcal{S}} \mu_k V^n(x + e_j - I_k(x + e_j)e_k) + \tau V^n(x + e_j) \\ + \sum_{l \in \mathcal{R}} \tau_{l1}(V^n(x + e_j + e_1) - V^n(x + e_j)),$$

because according to the inductive hypothesis at epoch n and $\tau_{l1} \geq \tau_{lj}$ for all $j \in \mathcal{S}$, station 1 is optimal for all resources. Rearranging terms, we obtain

$$(1 + \alpha)V^{n+1}(x + e_j) = r(x + e_j) + \sum_{k \in \mathcal{S}} \mu_k V^n(x + e_j - I_k(x + e_j)e_k) \\ + \left(\tau - \sum_{l \in \mathcal{R}} \tau_{l1} \right) V^n(x + e_j) + \sum_{l \in \mathcal{R}} \tau_{l1} V^n(x + e_j + e_1). \quad (\text{A.27})$$

We use the fact that $r(x + e_1) = r(x + e_j)$ (because $x_1 > 0$ and $x_j > 0$) and then apply (5.5) to the second, third, and fourth terms on the right-hand side of (A.27) (note that for the j th summand within the second term, we are using the fact that $x_1 < x_j$ to apply (5.5) for the state $x - e_j$). For the second term, we also use the fact that $0 < x_1 < x_j$ to conclude that $I_k(x + e_1) = I_k(x + e_j)$ for all $k \in \mathcal{S}$. For the fourth term, we use the fact that $x_1 < x_j$ to apply the inductive hypothesis to state $x + e_1$. We conclude that

$$(1 + \alpha)V^{n+1}(x + e_j) \leq r(x + e_1) + \sum_{k \in \mathcal{S}} \mu_k V^n(x + e_1 - I_k(x + e_1)e_k) \\ + \left(\tau - \sum_{l \in \mathcal{R}} \tau_{l1} \right) V^n(x + e_1) + \sum_{l \in \mathcal{R}} \tau_{l1} V^n(x + 2e_1).$$

Finally, re-arranging the terms on the right hand side of the above inequality and applying (5.3) results in

$$(1 + \alpha)V^{n+1}(x + e_j) \leq r(x + e_1) + \sum_{k \in \mathcal{S}} \mu_k V^n(x + e_1 - I_k(x + e_1)e_k) + \tau V^n(x + e_1) \\ + \sum_{l \in \mathcal{R}} \tau_{l1}(V^n(x + 2e_1) - V^n(x + e_1)).$$

The right-hand side of the above inequality is less than $(1 + \alpha)V^{n+1}(x + e_1)$, because the latter has a maximum inside the last summation.

Case 2: $0 = x_1 < x_j$. This case is similar to the previous one, except that we must treat the μ_1 term separately.

$$\begin{aligned} (1 + \alpha)V^{n+1}(x + e_j) &= r(x + e_j) + \mu_1 V^n(x + e_j) + \sum_{k \in \mathcal{S} \setminus \{1\}} \mu_k V^n(x + e_j - I_k(x + e_j)e_k) \\ &\quad + \tau V^n(x + e_j) + \sum_{l \in \mathcal{R}} \tau_{l1} (V^n(x + e_j + e_1) - V^n(x + e_j)), \end{aligned}$$

because according to the inductive hypothesis at epoch n and the assumption that $\tau_{l1} \geq \tau_{lj}$, station 1 is optimal for all resources. Because $r_j \leq r_1$,

$$\begin{aligned} (1 + \alpha)V^{n+1}(x + e_j) &\leq r(x + e_j) + \mu_1 r_1 + \mu_1 (V^n(x + e_j) - r_j) \\ &\quad + \sum_{k \in \mathcal{S} \setminus \{1\}} \mu_k V^n(x + e_j - I_k(x + e_j)e_k) \\ &\quad + \tau V^n(x + e_j) + \sum_{l \in \mathcal{R}} \tau_{l1} (V^n(x + e_j + e_1) - V^n(x + e_j)), \end{aligned}$$

and by applying (5.4), the above inequality implies that

$$\begin{aligned} (1 + \alpha)V^{n+1}(x + e_j) &\leq r(x + e_j) + \mu_1 r_1 + \mu_1 V^n(x) + \sum_{k \in \mathcal{S} \setminus \{1\}} \mu_k V^n(x + e_1 - I_k(x + e_1)e_k) \\ &\quad + \left(\tau - \sum_{l \in \mathcal{R}} \tau_{l1} \right) V^n(x + e_j) + \sum_{l \in \mathcal{R}} \tau_{l1} V^n(x + e_j + e_1). \end{aligned}$$

Now, by noting that $r(x + e_1) = \mu_1 r_1 + r(x + e_j)$ (since $x_1 = 0$ and $x_j > 0$) and then applying (5.5),

$$\begin{aligned} (1 + \alpha)V^{n+1}(x + e_j) &\leq r(x + e_1) + \sum_{k \in \mathcal{S}} \mu_k V^n(x + e_1 - I_k(x + e_1)e_k) \\ &\quad + \left(\tau - \sum_{l \in \mathcal{R}} \tau_{lj} \right) V^n(x + e_1) + \sum_{l \in \mathcal{R}} \tau_{lj} V^n(x + 2e_1), \end{aligned}$$

which can be rearranged to give

$$(1 + \alpha)V^{n+1}(x + e_j) \leq r(x + e_1) + \sum_{k \in \mathcal{S}} \mu_k V^n(x + e_1 - I_k(x + e_1)e_k) + \tau V^n(x + e_1) \\ + \sum_{l \in \mathcal{R}} \pi_{lj}(V^n(x + e_j + e_1) - V^n(x + e_1)),$$

since (5.3) holds. The right hand side of the above inequality is less than $(1 + \alpha)V^{n+1}(x + e_1)$, because the latter has a maximum inside the last summation.

Case 3: $x_1 = x_j$. In this case, (5.6) implies (5.5).

Inductive step for (5.6). Case 1: $x_j > x_1 > 0$.

$$(1 + \alpha)V^{n+1}(T_j x) \geq r(T_j x) + \sum_{k \in \mathcal{S}} \mu_k V^n(T_j x - I_k(T_j x)e_k) \\ + \tau V^n(T_j x) + \sum_{l \in \mathcal{R}} \pi_{lj}(V^n(T_j x + e_j) - V^n(T_j x)), \quad (\text{A.28})$$

by observing that the maximum in V^{n+1} is greater than any particular value. Noting that $x_1 > 0$ and $x_j > 0$, we rearrange terms to obtain

$$(1 + \alpha)V^{n+1}(T_j x) \geq r(T_j x) + \sum_{k \in \mathcal{S} \setminus \{1, j\}} \mu_k V^n(T_j x - I_k(T_j x)e_k) + \mu_1 V^n(T_j x - e_1) \\ + \mu_j V^n(T_j x - e_j) + \left(\tau - \sum_{l \in \mathcal{R}} \pi_{lj} \right) V^n(T_j x) + \sum_{l \in \mathcal{R}} \pi_{lj} V^n(T_j x + e_j).$$

Now, apply the inductive hypothesis on (5.6) for n . Because $x_i < x_j$ and using the fact that $r(T_j x) = r(x)$ because both x_i and x_j are positive,

$$\begin{aligned}
(1 + \alpha)V^{n+1}(T_j x) &\geq r(x) + \sum_{k \in \mathcal{S} \setminus \{1, j\}} \mu_k V^n(x - I_k(x)e_k) + \mu_1 V^n(x - e_j) \\
&\quad + \mu_j V^n(x - e_1) + \left(\tau - \sum_{l \in \mathcal{R}} \tau_{lj} \right) V^n(x) + \sum_{l \in \mathcal{R}} \tau_{lj} V^n(x + e_1) \\
&= r(x) + \sum_{k \in \mathcal{S} \setminus \{1, j\}} \mu_k V^n(x - I_k(x)e_k) + \mu_j V^n(x - e_j) + \mu_j V^n(x - e_1) \\
&\quad + (\mu_1 - \mu_j) V^n(x - e_j) + \tau V^n(x) + \sum_{l \in \mathcal{R}} \tau_{lj} (V^n(x + e_1) - V^n(x)).
\end{aligned}$$

Next, applying the inductive hypothesis (that (5.5) holds for n for the state $x - e_j - e_1$) and the condition that $\mu_1 \geq \mu_j$, we obtain

$$\begin{aligned}
(1 + \alpha)V^{n+1}(T_j x) &\geq r(x) + \sum_{k \in \mathcal{S} \setminus \{1, j\}} \mu_k V^n(x - I_k(x)e_k) + \mu_j V^n(x - e_j) + \mu_j V^n(x - e_1) \\
&\quad + (\mu_1 - \mu_j) V^n(x - e_1) + \tau V^n(x) + \sum_{l \in \mathcal{R}} \tau_{lj} (V^n(x + e_1) - V^n(x)).
\end{aligned}$$

Finally, because we have assumed that $\tau_{lj} = \tau_{l1}$ and because the inductive hypothesis implies that $\tau_{l1}(V^n(x + e_1) - V^n(x)) = \max_{k \in \mathcal{S}} \tau_{lk}(V^n(x + e_k) - V^n(x))$, we conclude that

$$\begin{aligned}
(1 + \alpha)V^{n+1}(T_j x) &\geq r(x) + \sum_{k \in \mathcal{S}} \mu_k V^n(x - I_k(x)e_k) + \tau V^n(x) \\
&\quad + \sum_{l \in \mathcal{R}} \max_{k \in \mathcal{S}} \tau_{lk} (V^n(x + e_k) - V^n(x)) \\
&= (1 + \alpha)V^n(x).
\end{aligned}$$

Case 2: $x_j > x_1 = 0$. Re-arranging terms in (A.28) and noting that $x_1 = 0$ and $x_j > 0$

yields

$$(1 + \alpha)V^{n+1}(T_j x) \geq r(T_j x) + \sum_{k \in \mathcal{S} \setminus \{1, j\}} \mu_k V^n(T_j x - I_k(T_j x)e_k) + \mu_1 V^n(T_j x - e_1) + \mu_j V^n(T_j x) \\ + \left(\tau - \sum_{l \in \mathcal{R}} \tau_{lj} \right) V^n(T_j x) + \sum_{l \in \mathcal{R}} \tau_{lj} V^n(T_j x + e_j).$$

Now, applying the inductive hypothesis on (5.6) for n and using the fact that $r(T_j x) = r(x) + r_1 \mu_1 - r_j \mu_j$ since $x_j > x_1 = 0$,

$$(1 + \alpha)V^{n+1}(T_j x) \geq r(x) + r_1 \mu_1 - r_j \mu_j + \sum_{k \in \mathcal{S} \setminus \{1, j\}} \mu_k V^n(x - I_k(x)e_k) + \mu_1 V^n(x - e_j) \\ + \mu_j V^n(x) + \left(\tau - \sum_{l \in \mathcal{R}} \tau_{lj} \right) V^n(x) + \sum_{l \in \mathcal{R}} \tau_{lj} V^n(x + e_1) \\ = r(x) + r_1 \mu_1 - r_j \mu_j + \sum_{k \in \mathcal{S} \setminus \{1, j\}} \mu_k V^n(x - I_k(x)e_k) + \mu_j V^n(x - e_j) + \mu_j V^n(x) \\ + (\mu_1 - \mu_j) V^n(x - e_j) + \tau V^n(x) + \sum_{l \in \mathcal{R}} \tau_{lj} (V^n(x + e_1) - V^n(x)).$$

Since (5.4) holds and $\mu_1 \geq \mu_j$, we have

$$(1 + \alpha)V^{n+1}(T_j x) \geq r(x) + r_1 \mu_1 - r_j \mu_j + \sum_{k \in \mathcal{S} \setminus \{1, j\}} \mu_k V^n(x - I_k(x)e_k) + \mu_j V^n(x - e_j) \\ + \mu_j V^n(x) + (\mu_1 - \mu_j)(V^n(x) - r_j) + \tau V^n(x) + \sum_{l \in \mathcal{R}} \tau_{lj} (V^n(x + e_1) - V^n(x)),$$

and because $r_1 \geq r_j$,

$$(1 + \alpha)V^{n+1}(T_j x) \geq r(x) + \sum_{k \in \mathcal{S} \setminus \{1, j\}} \mu_k V^n(x - I_k(x)e_k) + \mu_j V^n(x - e_j) + \mu_1 V^n(x) \\ + \tau V^n(x) + \sum_{l \in \mathcal{R}} \tau_{lj} (V^n(x + e_1) - V^n(x))$$

Finally, because we have assumed that $\tau_{lj} = \tau_{l1}$ and because the inductive hypothesis

implies that $\tau_{l1}(V^n(x + e_1) - V^n(x)) = \max_{k \in S} \tau_{lk}(V^n(x + e_k) - V^n(x))$, we conclude that

$$\begin{aligned}
(1 + \alpha)V^{n+1}(T_j x) &\geq r(x) + \sum_{k \in S} \mu_k V^n(x - I_k(x)e_k) + \tau V^n(x) \\
&\quad + \sum_{l \in \mathcal{R}} \max_{k \in S} \tau_{lk}(V^n(x + e_k) - V^n(x)) \\
&= (1 + \alpha)V^{n+1}(x).
\end{aligned}$$

Case 3: $x_1 = x_j$. In this case, (5.6) holds with equality because swapping the 1st and j th elements of x results in no change to the state. \square

Proof of Proposition 5.4. Since the decisions in the static policy are random, the expected total discounted reward is separable by station. That is, we can write $V^\Gamma(x) = \sum_{j \in S} W_j^\Gamma(x_j)$. We compute $W_j^\Gamma(x_j)$, the expected total discounted reward for a specific station $j \in S$. with arrivals at rate λ_j and departures at rate μ_j , when there are x_j customers waiting. For ease of exposition, we suppress the subscript j (corresponding to the station) and the superscript Γ (corresponding to the static policy) everywhere they appear in this proof. By uniformizing the process corresponding to the station with uniformization constant $\lambda + \mu$, we can study the embedded discrete-time Markov chain by observing the queue at the station only at transitions, i.e., arrivals and service completions. Arrivals occur with probability $\lambda/(\lambda + \mu)$ and service completions occur with probability $\mu/(\lambda + \mu)$. By incorporating the discount factor α , we can define $W(x)$ according to the following recursion:

$$W(x) = \frac{\lambda + \mu}{\lambda + \mu + \alpha} \left(\frac{\lambda}{\lambda + \mu} W(x + 1) + \frac{\mu}{\lambda + \mu} (r + W(x - 1)) \right),$$

or in other words,

$$\lambda W(x) - (\lambda + \mu + \alpha)W(x - 1) + \mu W(x - 2) = \mu r, \quad (\text{A.29})$$

for $x \geq 2$. The boundary condition is

$$\lambda W(1) - (\lambda + \alpha)W(0) = 0, \quad (\text{A.30})$$

because there are no service completions when there are zero casualties at the station.

Because $\alpha > 0$, $W(x)$ converges as $x \rightarrow \infty$. We find the value to which $W(x)$ converges, which we denote by $W(\infty)$, by solving

$$\lambda W(\infty) - (\lambda + \mu + \alpha)W(\infty) + \mu W(\infty) = -\mu r,$$

which yields $W(\infty) = \frac{\mu r}{\alpha}$. We can then solve the recurrence relation (A.29) by guessing $W(x) - W(\infty) = b^x$ and writing the characteristic equation

$$\lambda b^x - (\lambda + \mu + \alpha)b^{x-1} + \mu b^{x-2} = 0 \quad (\text{A.31})$$

for $x \geq 2$. Because $b \neq 0$, this equation is equivalent to

$$\lambda b^2 - (\lambda + \mu + \alpha)b + \mu = 0,$$

which can be solved using the quadratic formula to obtain

$$b = \frac{\lambda + \mu + \alpha \pm \sqrt{(\lambda + \mu + \alpha)^2 - 4\lambda\mu}}{2\lambda}.$$

Because there are two roots to the characteristic equation, we conclude that

$$W(x) - W(\infty) = c_1 b_1^x + c_2 b_2^x, \quad (\text{A.32})$$

where $b_1 = \frac{\lambda + \mu + \alpha + \sqrt{(\lambda + \mu + \alpha)^2 - 4\lambda\mu}}{2\lambda}$ and $b_2 = \frac{\lambda + \mu + \alpha - \sqrt{(\lambda + \mu + \alpha)^2 - 4\lambda\mu}}{2\lambda}$. Observe that if $\lambda \leq \mu$, $b_1 > 1$, while $0 < b_2 < 1$, so it must be true that $c_1 = 0$; otherwise $W(x)$ would not converge to $W(\infty)$ as $x \rightarrow \infty$. Finally, we determine the coefficient c_2 using the boundary condition: (A.30) implies that

$$\lambda \left(c_2 b_2 + \frac{\mu r}{\alpha} \right) - (\lambda + \alpha) \left(c_2 + \frac{\mu}{\alpha} \right) = 0. \quad (\text{A.33})$$

Solving the above equation yields $c_2 = -\frac{\mu r}{\lambda + \alpha - \lambda b_2}$, and therefore we conclude that

$$W(x) = \frac{\mu r}{\alpha} - \frac{\mu r}{\lambda + \alpha - \lambda b_2} b_2^x, \quad (\text{A.34})$$

or by plugging in b_2 and simplifying,

$$W(x) = \frac{\mu r}{\alpha} - \frac{2\mu r}{\lambda + \alpha - \mu + \sqrt{(\lambda + \mu + \alpha)^2 - 4\lambda\mu}} \left(\frac{\lambda + \mu + \alpha - \sqrt{(\lambda + \mu + \alpha)^2 - 4\lambda\mu}}{2\lambda} \right)^x. \quad (\text{A.35})$$

By summing the above expression over all stations, we obtain (5.10). Note that we showed that (A.29), (A.30), and $\lambda \leq \mu$ together imply that $W(x)$ must be the expression given in (A.35). Moreover, using algebraic manipulation it is straightforward to show that if $W(x)$ is given by the expression in (A.35), then it satisfies (A.29) and (A.30), without needing $\lambda \leq \mu$. \square

Bibliography

- N. Argon, S. Ziya, and R. Righter. Scheduling impatient jobs in a clearing system with insights on patient triage in mass casualty incidents. *Probability in the Engineering and Informational Sciences*, 22:301–332, 2008.
- N. Argon, L. Ding, K. D. Glazebrook, and S. Ziya. Dynamic routing of customers with general delay costs in a multiserver queueing system. *Probability in the Engineering and Informational Sciences*, 23(2), 2009.
- R. Baker and M. Strosberg. Triage and equality: An historical reassessment of utilitarian analyses of triage. *Kennedy Institute of Ethics Journal*, 2(2):103–123, 1992.
- D. Bish, E. Agca, and R. Glick. Decision support for hospital evacuation and emergency response. *Annals of Operations Research*, 2011.
- C. H. Boardman, S. D. Tuorinsky, D. C. Caneva, J. D. Malone, and W. L. Jackson. Field management of chemical casualties. In S. D. Tuorinsky, editor, *Medical Aspects of Chemical Warfare*, pages 485–509. Office of the Surgeon General, United States Army, 2008.
- N. Bostick, I. Subbarao, F. Burkle, E. Hsu, J. Armstrong, and J. James. Disaster triage systems for large-scale catastrophic events. *Disaster Medicine and Public Health Preparedness*, 2(1):S35–39, Sept. 2008.
- M. L. Brandeau, J. H. McCoy, N. Hupert, J.-E. Holty, and D. M. Bravata. Recommendations for modeling disaster responses in public health and medicine: A position paper the society for medical decision making. *Medical Decision Making*, 29:438–460, 2009.
- A. Childers and K. Taaffe. Healthcare facility evacuations: Lessons learned, research activity, and the need for engineering contributions. *Journal of Healthcare Engineering*, 1(1):125–140, 2010.
- D. Cone and D. MacMillan. Mass-casualty triage systems: a hint of science. *Academic Emergency Medicine: Official Journal of the Society for Academic Emergency Medicine*, 12(8):739–741, Aug. 2005.
- D. Cox and D. Oakes. *Analysis of survival data*. CRC Press, 1984.
- E. Frykberg. Triage: principles and practice. *Scandinavian Journal of Surgery*, 94(4): 272–278, 2005.
- A. Garner, A. Lee, K. Harrison, and C. Schultz. Comparative analysis of multiple-casualty incident triage algorithms. *Annals of Emergency Medicine*, 38(5):541–548, Nov. 2001.
- K. D. Glazebrook, P. S. Ansell, R. T. Dunn, and R. R. Lumley. On the optimal allocation of service to impatient tasks. *Journal of Applied Probability*, 41(1):51–72, 2004.

- Q. Gong and R. Batta. Allocation and reallocation of ambulances to casualty clusters in a disaster relief operation. *IIE Transactions*, 39(1):27–39, 2007.
- A. Y. Ha. Optimal dynamic scheduling policy for a Make-To-Stock production system. *Operations Research*, 45(1):42–53, Jan. 1997.
- J. Hick, J. Ho, W. Heegaard, D. Brunette, A. Lapine, T. Ward, and J. Clinton. Emergency medical services response to a major freeway bridge collapse. *Disaster Medicine and Public Health Preparedness*, 2(Supplement 1):S17–S24, Sept. 2008.
- J. L. Hick, P. Penn, D. Hanfling, M. A. Lappe, D. O’Laughlin, and J. L. Burstein. Establishing and training health care facility decontamination teams. *Annals of emergency Medicine*, 42(3), 2003.
- J. L. Hick, D. M. Weinstock, C. N. Coleman, D. Hanfling, S. Cantrill, I. Redlener, J. L. Bader, P. Murrain-Hill, and A. R. Knebel. Health care system planning for and response to a nuclear detonation. *Disaster Medicine and Public Health Preparedness*, 5 (Supplement_1):S73–88, Mar. 2011.
- D. Hogan, J. Waeckerle, D. Dire, and S. Lillibridge. Emergency department impact of the oklahoma city terrorist bombing. *Annals of Emergency Medicine*, 34(2):160–167, Aug. 1999.
- C. M. Hrdina, C. N. Coleman, S. Bogucki, J. L. Bader, R. E. Hayhurst, J. D. Forsha, D. Marcozzi, K. Yeskey, and A. R. Knebel. The RTR medical response system for nuclear and radiological mass-casualty incidents: A functional TRIage-TReatment-TRansport medical response model. *Prehospital and Disaster Medicine*, 24(03):167–178, 2009.
- A. Ingolfsson, S. Budge, and E. Erkut. Optimal ambulance location with random delays and travel times. *Health Care Management Science*, 11(3):262–274, 2008.
- E. U. Jacobson, N. Argon, and S. Ziya. Priority assignment in emergency response. *Operations Research*, 60(4), 2012.
- J. Jenkins, M. McCarthy, L. Sauer, G. Green, S. Stuart, T. Thomas, and E. Hsu. Mass-casualty triage: time for an evidence-based approach. *Prehospital and Disaster Medicine*, 23(1):3–8, Feb. 2008.
- A. Jotshi, Q. Gong, and R. Batta. Dispatching and routing of emergency vehicles in disaster mitigation using data fusion. *Socio-Economic Planning Sciences*, 43(1):1–24, 2009.
- C. Kahn, C. Schultz, K. Miller, and C. Anderson. Does START triage work? an outcomes assessment after a disaster. *Annals of Emergency Medicine*, 54(3):424–430.e1, Sept. 2009.
- C. Kahn, C. Schultz, K. Miller, and C. Anderson. In reply. *Annals of Emergency Medicine*, 55(6):580–581, 2010. ISSN 01960644.
- K. Kennedy, R. Aghababian, L. Gans, and C. Lewis. Triage: Techniques and applications

- in decisionmaking. *Annals of Emergency Medicine*, 28(2):136–144, Aug. 1996.
- G. Laporte. What you should know about the vehicle routing problem. *Naval Research Logistics*, 54(8):811–819, 2007.
- A. Law. *Simulation Modeling and Analysis*. McGraw Hill, 2007.
- E. Lerner. Mass casualty triage: An evaluation of the data and development of a proposed national guideline. *Disaster Medicine and Public Health Preparedness*, 2(1):S25–34, Sept. 2008.
- E. Lerner, R. Schwartz, P. Coule, and R. Pirrallo. Use of SALT triage in a simulated Mass-Casualty incident. *Prehospital Emergency Care*, 14(1):21–25, 2010.
- D. Li and K. Glazebrook. An approximate dynamic programming approach to the development of heuristics for the scheduling of impatient jobs in a clearing system. *Naval Research Logistics*, 2010.
- N. Liu, S. Ziya, and V. G. Kulkarni. Dynamic scheduling of outpatient appointments under patient no-shows and cancellations. *Manufacturing and Service Operations Management*, 12(2), 2010.
- D. Lockey, R. MacKenzie, J. Redhead, D. Wise, T. Harris, A. Weaver, K. Hines, and G. Davies. London bombings july 2005: The immediate pre-hospital medical response. *Resuscitation*, 66(2):ix – xii, 2005.
- J. Mistovich and K. Karren. *Prehospital Emergency Care*. Prentice Hall, 8th edition, 2007.
- D. Navin, W. Sacco, and G. McGill. Application of a new Resource-Constrained triage method to Military-Age victims. *Military Medicine*, 174:1247–1255, Dec. 2009.
- T. Okumura, K. Suzuki, A. Fukuda, A. Kohama, N. Takasu, S. Ishimatsu, and S. Hinohara. The tokyo subway sarin attack: Disaster management, part 1: Community emergency response*. *Academic Emergency Medicine*, 5(6):613–617, 1998.
- E. Regnier. Evacuation planning. In *Wiley Encyclopedia of Operations Research and Management Science*. 2011.
- E. Robenshtok, A. Eisenkraft, and S. Luria. Preparedness of local and national health systems for a chemical terrorist attack. In J. Shemer and Y. Shoenfeld, editors, *Terror and Medicine*. 2003.
- S. Ross. *Introduction to Stochastic Dynamic Programming*. Academic Press, 1983.
- W. Sacco, D. Navin, K. Fiedler, R. W. II, W. Long, and R. B. Jr. Precise formulation and evidence-based application of resource-constrained triage. *Academic Emergency Medicine*, 12(8):759–770, 2005.
- W. Sacco, D. Navin, R. Waddell, K. Fiedler, W. Long, and R. Buckman. A new Resource-Constrained triage method applied to victims of penetrating injury. *The Journal of*

Trauma: Injury, Infection, and Critical Care, 63(2):316–325, 2007.

- State of California EMS Authority. Ambulance strike team/medical task forces (AST) guidelines. Technical Report EMSA#215, July 2003. URL <http://www.emsa.ca.gov/pubs/pdf/emsa215.pdf>.
- H. Tanaka, A. Iwai, J. Oda, Y. Kuwagata, T. Matsuoka, T. Shimazu, and T. Yoshioka. Overview of evacuation and transport of patients following the 1995 Hanshin-Awaji earthquake. *The Journal of Emergency Medicine*, 16(3):439–444, 1998.
- E. Tayfur and K. Taaffe. A model for allocating resources during hospital evacuations. *Computers & Industrial Engineering*, 57(4):1313–1323, Nov. 2009.
- U.S. Department of Homeland Security. National planning scenarios. Technical report, 2005. URL <http://media.washingtonpost.com/wp-srv/nation/nationalsecurity/earlywarning/NationalPlanningScenariosApril2005.pdf>.
- U.S. Occupational Safety and Health Administration. *OSHA Best Practices for Hospital-Based First Receivers of Victims from Mass Casualty Incidents Involving the Release of Hazardous Substances*. Government Printing Office, 2005.
- M. H. Veatch and L. M. Wein. Scheduling a Make-To-Stock queue: Index policies and hedging points. *Operations Research*, 44(4):634–647, July 1996.
- G. R. Winslow. *Triage and Justice*. University of California Press, 1982.
- J. Winslow. Personal communication, 2010.