3-2021

The Impact of Triage Classification Errors on Military Medical Evacuation System Performance

Emily S. Graves

THE IMPACT OF TRIAGE CLASSIFICATION ERRORS ON MILITARY MEDICAL EVACUATION SYSTEM PERFORMANCE

THESIS

Emily S. Graves, 2d Lt, USAF

AFIT-ENS-MS-21-M-163

DEPARTMENT OF THE AIR FORCE
AIR UNIVERSITY

AIR FORCE INSTITUTE OF TECHNOLOGY

Wright-Patterson Air Force Base, Ohio

AFIT-ENS-MS-21-M-163

THE IMPACT OF TRIAGE CLASSIFICATION ERRORS ON MILITARY MEDICAL EVACUATION SYSTEM PERFORMANCE

THESIS

Presented to the Faculty

Department of Operational Sciences

Graduate School of Engineering and Management

Air Force Institute of Technology

Air University

Air Education and Training Command

in Partial Fulfillment of the Requirements for the

Degree of Master of Science in Operations Research

Emily S. Graves, B.S

2d Lt, USAF

25 March 2021

AFIT-ENS-MS-21-M-163

THE IMPACT OF TRIAGE CLASSIFICATION ERRORS ON MILITARY MEDICAL EVACUATION SYSTEM PERFORMANCE

THESIS

Emily S. Graves, B.S
2d Lt, USAF

Committee Membership:

Capt Phillip. R. Jenkins, PhD
Chair

Dr. Matthew J. Robbins
Member

AFIT-ENS-MS-21-M-163

Abstract

In a deployed environment, evacuation requests of injured personnel are serviced by multiple forms of evacuation including medical evacuation (MEDEVAC) and casualty evacuation (CASEVAC). This thesis focuses on the optimal dispatching policy for MEDEVAC units when triage classification errors and blood transfusion kits are considered. A discounted, infinite-horizon Markov decision process (MDP) model is formulated to analyze the MEDEVAC dispatching problem and determine the optimal policy based on the status of the MEDEVAC units in the system, the priority level of incoming requests, and the locations from which requests originate. A notional, representational scenario based in Azerbaijan is utilized to compare the optimal policy against the currently practiced policy of always dispatching the nearest available MEDEVAC unit. Multiple excursions are analyzed to understand the impact of altering problem parameters, including the misclassification rate, number of aircraft equipped with blood transfusion kits, arrival rate of incoming service requests, aircraft speed, and types of triage classification errors. Results reveal that with the application of the optimal policy found by the MDP model the performance of the MEDEVAC dispatching system improves, wherein performance is measured in terms of casualty survivability. Additionally, the inclusion of blood transfusion kits on board aircraft increase MEDEVAC system performance. This analysis is of interest to the military medical planning community and may inform the development of tactics, techniques, and procedures of future dispatching policies for MEDEVAC systems.

Acknowledgements

I first would like to express my deep gratitude to Capt Phillip R. Jenkins, PhD. for working with me as my advisor and his unending patience in assisting me with the formulation and analysis in this thesis. I am thankful for the opportunity to have him for my thesis advisor. Additionally, I would like to thank Dr. Matthew J.D. Robbins. for putting forth extra effort in being my reader and for the invaluable inputs he gave.

Emily S. Graves

Table of Contents

Page

Page

List of Figures

List of Tables

THE IMPACT OF TRIAGE CLASSIFICATION ERRORS ON MILITARY MEDICAL EVACUATION SYSTEM PERFORMANCE

I. Introduction

In a deployed environment, military emergency medical service (EMS) response system personnel seek to effectively and efficiently evacuate casualties from the battlefield to a medical treatment facility (MTF). There are two primary resources available to accomplish this task: medical evacuation (MEDEVAC) and casualty evacuation (CASEVAC). MEDEVAC platforms have dedicated medical personnel on board to treat casualties en route to an MTF, whereas CASEVAC platforms do not (Department of the Army, 2019). As such, military medical planners rely on MEDEVAC to serve as the primary link among the roles of medical care across combat operations. Moreover, CASEVAC is typically utilized only when MEDEVAC platforms are limited and/or overburdened.

Whereas a variety of platforms can be leveraged when performing an evacuation (e.g., ground ambulances, air ambulances, and sea ambulances), this thesis focuses specifically on rotary wing air assets (e.g., helicopters) in regards to evacuating casualties via MEDEVAC. Helicopters were first utilized to evacuate casualties during the Korean War and continue to be employed as the primary MEDEVAC platform. The United States (U.S.) Army employs HH-60M Black Hawk helicopters for MEDEVAC missions, which are capable of air crash rescue support; expeditious delivery of whole blood and medical supplies to meet critical requirements; rapid movement of medical personnel and accompanying equipment to address changes in battlefield requirements; and movement of patients between hospital. (Department of the Army,

2019).

It is imperative that a MEDEVAC system is effective and efficient, not only to increase survivability and decrease the time between injury and medical care, but also to retain confidence among military personnel conducting combat operations on the battlefield. More specifically, an effective and efficient system demonstrates to battlefield personnel that rapid and quality care is available upon request. Many decisions impact the effectiveness and efficiency of a MEDEVAC system, including the location, allocation, relocation, dispatching, and redeployment of assets.

This thesis focuses on the MEDEVAC dispatching problem, which seeks to determine whether to dispatch an asset and, if so, which asset to dispatch in response to a request. The MEDEVAC dispatching problem is formulated as a discounted, continuous-time Markov decision process (MDP) model over an infinite horizon. Similar to previous research, this study assumes MEDEVAC asset locations are predetermined and that redeployment does not occur. As an augment to the previous research in this area (e.g., Keneally *et al.*, 2016; Jenkins, 2017; Jenkins *et al.*, 2018), this work accounts for the possibility of triage classification errors and explicitly models blood transfusion kits on board select MEDEVAC aircraft to improve realism and explore how these characteristics might impact system performance. For example, a MEDEVAC request that is called in as a life-threatening request may actually be a non-life threatening request. Aircraft equipped with blood transfusion kits give on board medical professionals the ability to begin the necessary life saving medical procedures prior to arriving at an MTF.

Each MEDEVAC request submitted to the system is categorized as one of three priority levels: urgent, priority, or routine (Department of the Army, 2019). Urgent requests (i.e., Priority I) correspond to emergency cases that should be evacuated as soon as possible and within a maximum of one hour to maximize survivability and

minimize long-term disabilities (e.g., loss of limb or eyesight). Priority requests (i.e., Priority II) correspond to sick and wounded personnel requiring prompt medical care. Casualties are categorized as Priority II when they should be evacuated within four hours or when their medical condition could deteriorate to such a degree that they will become a Priority I request. Routine requests (i.e., Priority III) correspond to sick and wounded personnel requiring evacuation but whose condition is not expected to deteriorate significantly. Even so, routine request casualties should still should be evacuated within 24 hours to prevent further deterioration in health. The MDP model utilizes the casualty priority level to determine whether or not to dispatch a MEDEVAC unit and which MEDEVAC unit to dispatch. For example, if there is one idle MEDEVAC unit and a service request is received by the dispatching authority, the MDP model will take into account the priority level of the request and determine whether or not to dispatch. The dispatching authority may be more likely to dispatch the MEDEVAC unit for an urgent request as opposed to a routine request.

Previous MEDEVAC dispatching models assume accurate triage classification (e.g., Rettke *et al.*, 2016; Robbins *et al.*, 2020; Jenkins *et al.*, 2021b), but this is an unrealistic assumption for practical scenarios. In a deployed environment, assessing injuries can be difficult and the service member calling in a MEDEVAC request may incorrectly report the true priority level. The priority level of a casualty is not truly known to MEDEVAC staff until the injury is assessed by trained medical personnel at the casualty collection point (CCP). For example, one case of incorrect reporting could be that a Priority II request is reported as a Priority I request. If MEDEVAC staff assume correct reporting, an optimal policy might immediately deploy the closest-available MEDEVAC asset, whereas a different decision may be made if the true priority level is known.

The remainder of this thesis is organized as follows: Chapter II provides a review of

research relating to MEDEVAC systems. Chapter III describes the MDP formulation developed to determine an optimal MEDEVAC dispatch policy. Chapter IV covers an application of the formulated MDP based on a representative scenario in Azerbaijan. Chapter V concludes the thesis and proposes several directions for future research.

II. Literature Review

Related research leading up to this thesis can be divided into two subsections: civilian EMS response system research and MEDEVAC research. Throughout the research on medical response units, a variety of problem factors and decision variables have been analyzed to include the location, allocation, and relocation of response vehicles (e.g., Berman, 1981; Kolesar & Walker, 1974; Chaiken & Larson, 1972), the dispatch policy (e.g., Ignall *et al.*, 1982; Swersey, 1982; Green & Kolesar, 1984; Jenkins, 2019), and the distribution of service zones (e.g., Daskin & Stern, 1981; Jarvis, 1985). EMS response research lends itself well to common operations research techniques to include stochastic modeling, queuing, discrete optimization, and simulation modeling (Green & Kolesar, 2004).

For both military MEDEVAC dispatching and EMS response, previous research efforts formulate MDP models and leverage approximate dynamic programming (ADP) techniques to compare myopic policies against optimal or near optimal policies. A common myopic policy utilized in practice is to task the nearest available ambulance to respond to a request for service, but this policy is often not optimal in practice. Bandara *et al.* (2012) found that the optimal policy when dispatching ambulances to different zones is to reserve an ambulance to serve the designated zone with the highest call rate. In addition to analyzing dispatching policies, researchers have examined relocation procedures after a call has been serviced. Relocation creates a more dynamic model with fewer operational constraints when compared to a model with fixed ambulance locations. Jagtenberg *et al.* (2017) develop a model that determines relocation decisions with heuristics. The use of heuristics allows computations to be made in real time.

Augmenting a standard dispatching model to optimize a specific objective and then applying heuristics or ADP to lower computational time is common throughout

the field of EMS research. For example, Nasrollahzadeh *et al.* (2018) utilize an ADP technique to overcome the curse of dimensionality and produce high-quality policies that seek to simultaneously minimize both response time and fraction of high-priority late calls.

Similar to this thesis, McLay & Mayorga (2013) model classification errors for an ambulance dispatching model. The authors seek to minimize average response times when there is a positive probability of triage classification errors. The results show that improvement can be made over a myopic policy for response times, but there are diminished benefits for higher rates of classification errors. Although military MEDEVAC research mirrors EMS response research, key differences exist that warrant further exploration (Jenkins *et al.*, 2020a,b, 2021c).

Whereas military MEDEVAC dispatching is similar to EMS dispatching, the presence of possible threats, the unfamiliar environment, and the nature of combat related injuries allow for distinctions between the two. Similar to augmenting the standard EMS dispatching problem, the military MEDEVAC dispatching problem has been adjusted to allow for different scenarios to be modeled and analyzed. Keneally *et al.* (2016) utilize value iteration with dynamic programming to determine the optimal dispatching policy for military MEDEVAC for a small scale scenario in Afghanistan. Included in the model was the possibility of a high threat situation and the addition of an armed military escort.

One augmentation made by Jenkins *et al.* (2018) was to allow for admission control and queuing in the MEDEVAC dispatching model. An original assumption made by Keneally *et al.* (2016) is that when a MEDEVAC request is submitted, if the system has an available MEDEVAC unit, then a MEDEVAC unit must respond. Admission control allows the dispatching authority to turn the request away to other forms of evacuation (e.g., ground ambulances). This thesis includes admission control to allow

the dispatching authority to turn away MEDEVAC requests to reserve MEDEVAC units for more urgent requests that are expected to occur in the near future.

In addition to utilizing MDP to determine the optimal dispatching policy for military MEDEVAC, the research has moved to use ADP techniques. The MEDEVAC dispatching problem cannot be solved in a tractable amount of time utilizing exact dynamic programming techniques when the number of zones is high or the state space gets large with additional augmentations. ADP techniques seek to find high-quality solutions in an efficient manner. Multiple ADP techniques have been used to solve MEDEVAC related problems. For example, Rettke *et al.* (2016) leverage least-squares temporal differences (LSTD) within an approximate policy iteration (API) framework to solve the MEDEVAC dispatching problem. Their ADP policies improve over the myopic dispatching policy by nearly 31% in regards to a life-saving performance metric. Robbins *et al.* (2020) utilize hierarchical aggregation technique to solve the dispatching problem. Jenkins *et al.* (2021a) implement two separate ADP techniques, the first being LSTD and the second being neural network learning. Results reveal neural network learning outperforms LSTD. Additionally, Jenkins *et al.* (2021b) examine the dispatching, preemption-rerouting, redeployment (DPR) problem. This research improves the combat casualty survivability rate utilizing support vector regression within an API algorithmic framework. The ADP methods mentioned above allow researchers to explore larger-scale research than allowed by MDP methods.

In addition to accounting for classification errors, this thesis accounts for blood transfusion kits on MEDEVAC aircraft. Malsby III *et al.* (2013) examine the feasibility of performing blood transfusions on-board military MEDEVAC vehicles and found that the vehicle and altitude do not have an adverse effect on blood transfusions. Analyzing the effectiveness of a timely blood transfusion, Kotwal *et al.* (2018)

show that an early blood transfusion is associated with higher chances of battlefield survival.

This thesis contributes to existing MEDEVAC dispatching literature by explicitly accounting for classification errors and blood transfusion kits on-board MEDEVAC aircraft. Classification errors occur when the reported priority level differs from the true priority level. Blood transfusion kits on board military MEDEVAC aircraft allow for the patient to be treated more extensively and in a more timely manner prior to arriving at an MTF. These additional problem features create a more accurate depiction of the military MEDEVAC system.

III. Methodology

In this chapter, the first section provides a brief description of the MEDEVAC dispatching problem and subsequently how it will be modeled utilizing a MDP formulation. The next section describes the mathematical model formulation, detailing the parameter definitions, decision epochs, state space, action space, transition probabilities, rewards, and objective function.

3.1 Problem Description

When a service member requests a MEDEVAC unit to service injured individuals, the MEDEVAC dispatching authority analyzes the incoming information and subsequently decides whether to dispatch a MEDEVAC unit, and if they do dispatch a unit, which unit to dispatch. The information in the service request includes but is not limited to, location, number, and priority level(s) of the injured individual(s). If the dispatching authority assigns a unit to the service request, the unit assigned then departs the staging area, arrives at the CCP, loads the injured individual(s) onto the aircraft, departs the the CCP, arrives at the MTF, unloads the injured individual(s), departs the MTF, and travels back to the staging area to await the next service request from the dispatching authority. Figure 1 depicts the timeline described above.

The incorporation of admission control allows the MEDEVAC dispatching authority to turn away requests to other forms of evacuation. In this case, the service request is sent to a secondary evacuation service such as CASEVAC. If the request is not sent to another evacuation system, it is assigned a MEDEVAC unit to service it. Once assigned, the MEDEVAC crew prepares for the mission, and the unit is dispatched to complete the request. If every MEDEVAC unit is busy and a service request is submitted to the system, the service request is relayed to other forms of evacuation, and

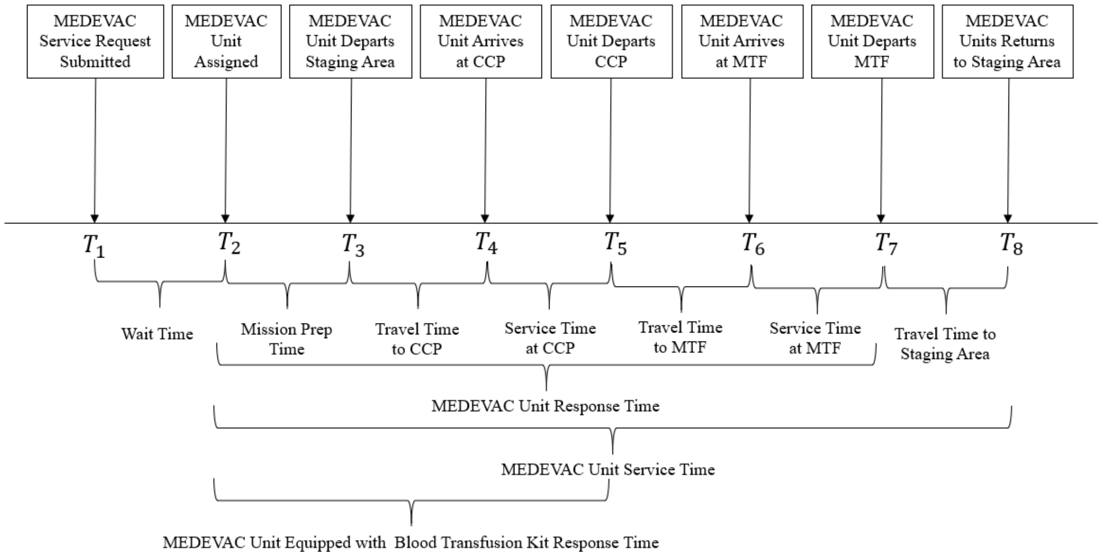

Figure 1. MEDEVAC Mission Timeline

a queue for MEDEVAC unit service is not formed. Because there are multiple forms of evacuation in a deployed environment, it is more beneficial for a service request to be relayed to a secondary evacuation service as opposed to waiting for a MEDEVAC unit. In flight from the CCP to the MTF, medical treatment is administered to the injured individual(s), but medical personnel are limited by what is available on board. If the MEDEVAC aircraft is equipped with a blood transfusion kit, medical personnel are able to start blood transfusions as soon as the injured individual(s) is loaded. Kotwal *et al.* (2018) show that the timing of a blood transfusion is critical to survival and the earlier blood is given the better chance of survival. Because of the benefits of an on board transfusion kit, the injured individual will be considered serviced when loaded onto MEDEVAC aircraft and the model will not include travel time from the CCP to the MTF or from the MTF to the staging area when calculating casualty service time. Figure 1 shows the representative service time for MEDEVAC aircraft equipped with Blood Transfusion Kits.

3.2 Problem Formulation

Because of the stochastic nature of this problem we are able to model it as a continuous-time MDP. This thesis includes classification errors and the presence of blood transfusion kits on some MEDEVAC units. This section defines the sets, parameters, and components of the MDP model.

3.2.1 Sets and Parameters

- Let $\mathcal{M} = \{1, 2, ..., |\mathcal{M}|\}$ represent the set of aircraft that are utilized to service MEDEVAC requests.
- Let $\mathcal{Z} = \{1, 2, ..., |\mathcal{Z}|\}$ represent the set of zones from which MEDEVAC requests can originate.
- Let $\mathcal{K} = \{1, 2, ..., |\mathcal{K}|\}$ represent the set of possible priority levels describing each MEDEVAC request.
- Let ϕ_{kc} denote the probability the reported priority level k of a request is in reality the true priority level c of a request. For example, $\phi_{12} = 0.9$ indicates that 90% of reported Priority 1 casualties are, in reality, Priority 2 casualties. Additionally, $\sum_{c \in \mathcal{K}} \phi_{kc} = 1, \forall\, k \in \mathcal{K}$.
- Requests for MEDEVAC service arrive to the dispatching authority according to a Poisson Process with a rate parameter of $\lambda = \sum_{z \in \mathcal{Z}} \sum_{k \in \mathcal{K}} \sum_{c \in \mathcal{K}} \lambda_{zkc}$ wherein λ_{zkc} represents the MEDEVAC request arrival rate in minutes for zone z, reported priority level k and true priority level c.
- Let μ_{mz} represent the expected service rate in minutes for MEDEVAC m servicing a zone z request. Assume service times are exponentially distributed

- Let ζ_{mz} represent the expected response rate in minutes for MEDEVAC aircraft m servicing a zone z request.
- Let $\mathcal{I}(S_t) = \{m : m \in \mathcal{M}, M_{tm} = 0\}$ represent the set of idle MEDEVAC units available to be dispatched when the state of the system is S_t at decision epoch t.
- Let $\mathcal{B}(S_t) = \{m : m \in \mathcal{M}, M_{tm} \neq 0\}$ represent the set of busy MEDEVAC units when the state of the system is S_t at decision epoch t. Of note, $\mathcal{M} = \mathcal{I}(S_t) \cup \mathcal{B}(S_t), \forall\, S_t \in \mathcal{S}$.

3.2.2 Decision Epochs

Decisions occur when a request is transmitted or when a MEDEVAC unit completes a service request (i.e., arrives back to its staging facility). The set of decision epochs is given by

$$\mathcal{T} = \{1, 2, ...\}.$$

3.2.3 State Space

Let $S_t \in \mathcal{S}$ represent the state of the system at decision epoch $t \in \mathcal{T}$. The state space is comprised of two components. The first component is the status of the MEDEVAC units and the second component is the request status of the system. Let $S_t = (M_t, R_t)$ represent the status of the MEDEVAC system at epoch t, where M_t represents the MEDEVAC status tuple, and R_t represents the request status tuple. The MEDEVAC status tuple is defined as

$$M_t = (M_{tm})_{m \in \mathcal{M}},$$

where $\mathcal{M} = \{1, 2, ..., |\mathcal{M}|\}$ represents the set of MEDEVAC units in the system and M_{tm} contains information pertaining to MEDEVAC units, m at epoch t. The state of

MEDEVAC unit $m \in \mathcal{M}$ is represented by the state variable $M_{tm} = \{0\} \cup \mathcal{Z}$, which represents the location of aircraft m at epoch t. When $M_{tm} = 0$, the unit m is idle, and when $M_{tm} = \mathcal{Z}$, unit m is busy servicing a request from zone $z \in \mathcal{Z}$.

The request status tuple R_t provides information on the current request awaiting an admission control decision at epoch t. Specifically, it provides the zone, the reported priority level, and true priority level of the request arrival given there is one at epoch t. The request status tuple is $R_t = (0, 0, 0)$ when there is not a request awaiting decision in the system. Otherwise, the request status tuple is given by

$$R_t = (Z_t, K_t, C_t)_{Z_t \in \mathcal{Z}, K_t \in \mathcal{K}, C_t \in \mathcal{K}},$$

where Z_t represents the zone from which the request originated, K_t represents the reported priority level of the request (i.e., routine, priority, or urgent), and C_t represents the true priority level of the request. If $C_t \neq K_t$, a classification error has occurred.

3.2.4 Action Space

When a request is submitted to the system, either a MEDEVAC unit is dispatched to service the request, or the request is rejected from entering the system and is serviced through other avenues. The dispatcher can only take action when a request is in the system and at least one MEDEVAC unit is available to service the request. There are two situations and sets of actions that can be taken. The first situation is if a request arrives and there are MEDEVAC units available. The dispatching authority can either pick a MEDEVAC unit to dispatch or choose not to dispatch at all (i.e., reject the request from entering the system). The second situation is if a request arrives and there are no MEDEVAC units available. The only action to be taken is to reject the request from entering the system to be serviced elsewhere. It is important

to note that if a MEDEVAC unit is dispatched, it is considered unavailable until it completes the request and returns back to the staging location.

Let $a_t^{reject} \in \{\Delta, 0, 1\}$ denote the admission control decision at epoch t. When $a_t^{reject} = 1$ the service request is rejected from the system; when $a_t^{reject} = 0$ the service request is accepted into the MEDEVAC system; and when $a_t^{reject} = \Delta$ the request status tuple is empty (i.e., $R_t = (0, 0, 0)$), indicating there is no request in the system, and the system will transition without any input from a_t^{reject}. Let $a_t^d = (a_{tm}^d)_{m \in \mathcal{I}(S_t)}$ represent the arrival request dispatch decision variable at epoch t. If $a_{tm}^d = 1$, then MEDEVAC unit $m \in \mathcal{I}(S_t)$ is dispatched to service the current request at epoch t, and 0 otherwise. Finally, let $a_t = (a_t^{reject}, a_t^d)$ represent the decision variable tuple at epoch t. The action space is constrained by

$$\sum_{m \in \mathcal{I}(S_t)} a_{tm}^d \leq \mathbb{I}_{\{R_t \neq (0,0,0)\}} \mathbb{I}_{\{a_t^{reject}=0\}},$$

where $\mathbb{I}_{\{R_t \neq (0,0,0)\}}$ is an indicator function that equals 1 if an incoming request has arrived to the system. Additionally, $\mathbb{I}_{\{a_t^{reject}=0\}}$ is an indicator function that equals 1 if the incoming request is admitted to the system. Let $\mathcal{A}(S_t)$ denote the actions available to the MEDEVAC dispatcher when the system is in state S_t at decision epoch t. The action space is below:

$$\mathcal{A}(S_t) = \begin{cases} (\Delta, \{0\}^{|\mathcal{I}(S_t)|}) & R_t = (0,0,0) \\ (1, \{0\}^{|\mathcal{I}(S_t)|}) & R_t \neq (0,0,0), \mathcal{I}(S_t) = \varnothing \\ (\{0,1\}, \{0,1\}^{|\mathcal{I}(S_t)|}) & R_t \neq (0,0,0), \mathcal{I}(S_t) \neq \varnothing. \end{cases} \tag{1}$$

The first case represents the scenario in which there are no requests in the system, and the only action available is to do nothing. The second case represents the scenario in which there is a request in the system, but there are no MEDEVAC units available to be dispatched, and the only action is to reject the request from entering the system. The third case represents the scenario in which there is a request in the system and

at least one MEDEVAC unit is available, the actions available are to not dispatch a unit if the request is rejected from entering the system, or to dispatch any available unit if the request is admitted into the system.

3.2.5 Transition Probabilities

This system transitions when an event occurs. The first event type is when a request is submitted and the MEDEVAC dispatching authority either accepts the request and dispatches an aircraft or rejects the request to a secondary means of evacuation. The system remains in a post-decision state until either another request is submitted or a dispatched unit completes service. The second event type is when a MEDEVAC unit completes service, and the unit deterministically transitions into an idle state.

When the MEDEVAC system is in state S_t at epoch t and action a_t is taken, the system immediately transitions to a post decision state denoted by S_t^a. The transition time from this post decision state to the next pre-decision state S_{t+1} is exponentially distributed with parameter $\beta(S_t, a_t)$ (i.e., the state-action sojourn time). Let $\beta(S_t, a_t)$ be defined as,

$$\beta(S_t, a_t) = \lambda + \sum_{m \in \mathcal{B}(S_t)} \mu_{m,M_{tm}} + \sum_{m \in \mathcal{I}(S_t)} a_{tm}^d \mu_{m,Z_t}. \tag{2}$$

If $\mathcal{B}(S_t) = \varnothing$ and $a_t^d = \{0\}^{\mathcal{I}(S_t)}$, indicating all MEDEVAC units are idle and there are no requests in the system, then $\beta(S_t, a_t)$ represents the state-action pair sojourn time wherein the next decision epoch occurs upon the arrival of a new MEDEVAC service request. Otherwise, $\mathcal{B}(S_t) \neq \varnothing$ and/or a unit is tasked to service an incoming request (e.g., $a_{tm}^d = 1$ for some $m \in \mathcal{I}(S_t)$) then at least one MEDEVAC unit is servicing a request. In this case, $\beta(S_t, a_t)$ represents the state-action pair sojourn time wherein the next decision epoch occurs after the arrival of a new request or a MEDEVAC

unit returns from servicing a request. The probabilistic behavior of the MEDEVAC system can be summarized using an infinitesimal generator as follows,

$$G(S_{t+1}|S_t, a_t) = \begin{cases} -[1 - p(S_t^a|S_t, a_t)]\beta(S_t, a_t), \text{if } S_{t+1} = S_t^a \\ p(S_{t+1}|S_t, a_t)\beta(S_t, a_t), \text{if } S_{t+1} \neq S_t^a. \end{cases} \tag{3}$$

Given that the system is in state S_t and action a_t is taken, the probability that the system transitions to state S_{t+1} is denoted by,

$$p(S_{t+1}|S_t, a_t) = \begin{cases} \frac{\lambda_{zkc}}{\beta(S_t,a_t)}, \text{if } R_{t+1} = (z, k, c), z \in \mathcal{Z}, k \in \mathcal{K}, c \in \mathcal{K} \\ \frac{\mu_{mz}}{\beta(S_t,a_t)}, \text{if } R_{t+1} = (0, 0, 0), M_{m,t+1} = 0, M_{tm}^a = z, m \in \mathcal{M}, z \in \mathcal{Z} \\ 0, \text{ otherwise.} \end{cases} \tag{4}$$

Note that $M_{tm}^a \in \{0\} \cup \mathcal{Z}$ denotes the post-decision state variable that contains the information pertaining to MEDEVAC unit m when action a_t is taken at epoch t.

Using uniformization as explained by Puterman (2005), we transform the continuous-time MDP into an equivalent discrete-time MDP to ease subsequent analysis. The maximum rate of the system is calculated by,

$$\nu = \lambda + \sum_{m \in M} \tau_m, \tag{5}$$

wherein

$$\tau_m = \max_{z \in Z} \mu_{mz}, \forall m \in \mathcal{M}. \tag{6}$$

When uniformization is applied, it allows the system to have self-transitions. Applying uniformization gives the following transition probabilities,

$$
\tilde{p}(S_{t+1}|S_t, a_t) = \begin{cases} 1 - \frac{[1-p(S_t^a|S_t,a_t]\beta(S_t,a_t)}{\nu}, \text{ if } S_{t+1} = S_t^a \\ \frac{[p(S_{t+1}|S_t,a_t]\beta(S_t,a_t)}{\nu}, \text{ if } S_{t+1} \neq S_t^a \\ 0, \text{otherwise.} \end{cases} \tag{7}
$$

3.2.6 Rewards

As the MEDEVAC system services requests, the model will reward itself based off of the reward function. The immediate expected reward received is a function of the request zone $z \in \mathcal{Z}$, the MEDEVAC unit $m \in \mathcal{M}$, the reported priority level $k \in \mathcal{K}$, and the true priority level $c \in \mathcal{K}$. Let $r(S_t, a_t) = \psi_{mzkc}$ represent the immediate expected reward reward given when the system is in state S_t and action a_t is taken. It is given by

$$
\psi_{mzk} = \delta\phi_{k1}e^{\frac{-\zeta mz}{60}} + \phi_{k2}e^{\frac{-\zeta mz}{240}} \tag{8}
$$

where $\delta \geq 1$ denotes the trade-off parameter utilized to alter the urgent to priority request immediate expected reward ratio. For this thesis, we set $\delta = 10$. If no MEDEVAC unit is dispatched $\psi_{mzkc} = 0$. Because of the continuous nature of the problem, uniformization is applied to transform the reward function to an equivalent discrete-time form as follows

$$
\tilde{r}(S_t, a_t) = r(S_t, a_t)\frac{\alpha + \beta(S_t, a_t)}{\alpha + \nu}, \tag{9}
$$

where $\alpha > 0$ represents the continuous-time discounting rate.

3.2.7 Objective Function

Let $A^\pi(S_t) \in \mathcal{A}(S_t)$ represent the decision function that maps the state space to the action space. This function indicates action a_t to be taken given the system is

in state S_t according to policy π. The MDP model looks to determine the optimal policy π^* from all available policies, $(A^\pi(S_t))_{\pi \in \Pi}$. The optimal policy maximizes the expected total discounted reward (ETDR). The objective is given by

$$\max_{\pi \in \Pi} \mathbb{E}^\pi \left[\sum_{t=1}^{\infty} \gamma^{t-1} \tilde{r}(S_t, A^\pi(S_t)) \right], \tag{10}$$

where $\gamma = \frac{\nu}{\nu + \alpha}$ is the uniformized discount factor. The optimal policy is found via the Bellman equation,

$$V(S_t) = \max_{a_t \in \mathcal{A}(S_t)} \left(\tilde{r}(S_t, a_t) + \gamma \sum_{S_{t+1} \in \mathcal{S}} \tilde{p}(S_{t+1}|S_t, a_t) V(S_{t+1}) \right). \tag{11}$$

The policy iteration algorithm is implemented in MATLAB to solve the Bellman equations and determine the optimal dispatching policy π^* exactly.

IV. Testing, Results, & Analysis

This chapter examines a notional scenario of a military MEDEVAC planning instance. This scenario is utilized to showcase the ability of the MDP model and allow analysis on the optimal decision policy created. This chapter examines a myriad of excursions in which model parameters of the baseline scenario are altered for sensitivity analysis. The parameters include the misclassification rate, number of MEDEVAC units equipped with blood transfusion kits, expected arrival rate of service requests, speed of the MEDEVAC aircraft, and expected proportion of each service request priority level.

4.1 Baseline Scenario

The baseline scenario represents a notional military MEDEVAC planning instance in Azerbaijan. The bases correspond to military MEDEVAC locations. In this scenario there are four bases, two of which are adjacent to MTFs. The four bases used for the scenario are located in Agdzhabedi, Karachala, Goradiz, and Salyany. Casualty cluster centers are created using Monte Carlo simulation based off of projected enemy locations. The bases, stations, and casualty cluster centers are depicted below in Figure 2.

The baseline scenario consists of six zones, three priority classes (e.g., urgent, priority, and routine), and four available MEDEVAC units. A total misclassification rate of 0.3 indicates that we expect 30% of the casualty service requests to be reported as the incorrect priority level. We assume service request priority levels can only be overestimated, so ϕ_{kc} is directly calculated using the total misclassification rate. Table 1 tabulates ϕ_{kc} for all combinations of $k, c \in \mathcal{K}$ for the baseline scenario. The subsequent analysis of this research focuses on the total misclassification rate as a

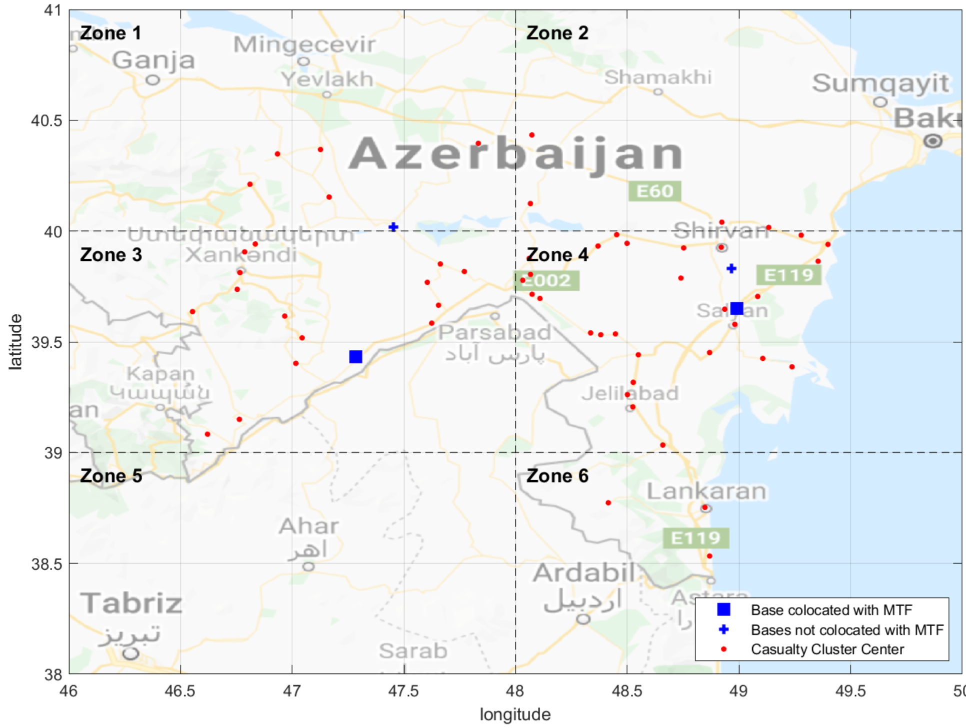

Figure 2. 6 Zone Azerbaijan Baseline Scenario

proxy for ϕ_{kc}

Table 1. Total Misclassification Rate(0.3) to ϕ_{kc}

$k\backslash c$	1	2	3
1	0.7	0.15	0.15
2	0	0.7	0.3
3	0	0	1

The arrival rate, $\lambda = 1/60$, indicates that we expect one service request per hour on average to arrive in the MEDEVAC system. The speed of the MEDEVAC aircraft is set to 150 knots, which corresponds with the current average airspeed of the HH-60M Black Hawk aircraft. Using notional data to simulate casualty points, we calculate the expected proportion of casualties originating from each zone, the expected response time by MEDEVAC unit and zone, and the expected service time by MEDEVAC unit and zone. The proportions of casualty events occurring in each

zone are listed in Table 2. The expected response times are listed in Table 3, and the expected service times are listed in Table 4.

Table 2. Proportion of Casualty Events Coming From Each Zone

Zone	1	2	3	4	5	6
Proportion	0.096	0.081	0.276	0.486	0.001	0.059

Table 3. Expected Response Time in Minutes of MEDEVAC unit to Zone

Zone\MEDEVAC	1	2	3	4
1	71.58	84.10	61.05	86.12
2	74.24	57.05	65.88	60.27
3	51.42	74.02	53.93	74.25
4	65.53	50.28	63.61	49.70
5	61.37	91.74	73.76	90.58
6	83.16	79.20	90.27	75.25

Table 4. Expected Service Time in Minutes of MEDEVAC unit to Zone

Zone\MEDEVAC	1	2	3	4
1	56.76	89.53	90.48	110.14
2	99.06	114.77	77.78	74.67
3	48.48	70.33	81.55	102.5
4	103.78	105.95	60.64	48.24
5	92.07	93.31	121.91	140.74
6	116.72	111.68	76.55	64.20

When a casualty event is misclassified, we assume the event priority has been overestimated. For example, if a MEDEVAC service request is submitted as urgent, the request may actually be urgent or it may have been misclassified and is truly a priority (or routine) request.

No blood transfusion kits are included on board any MEDEVAC aircraft for the baseline scenario, but subsequent analysis will allow for the addition of kits on a number of aircraft that are co-located with MTFs. Table 5 lists and describes the parameters of the baseline scenario.

Table 5. Baseline Parameters

Parameter	Description	Value
$\lvert\mathcal{M}\rvert$	number of MEDEVAC units	4
$\lvert\mathcal{Z}\rvert$	number of zones	6
$\lvert\mathcal{K}\rvert$	priority levels	3
λ	arrival rate	1/60
ϕ_{12}	Reported urgent, but truthfully priority request misclassification rate	0.15
ϕ_{13}	Reported urgent, but truthfully routine request misclassification rate	0.15
ϕ_{23}	Reported priority, but truthfully routine request misclassification rate	0.3

4.1.1 Myopic Policy

The MEDEVAC dispatching system currently practices a myopic policy, which immediately dispatches the nearest available MEDEVAC unit when a service request enters the system. When the myopic policy is implemented, the dispatching authority does not take system factors (e.g., priority level or number of units available) into consideration. The goal of this thesis is to analyze the optimal policy calculated from the MDP in comparison to the myopic policy.

4.1.2 Baseline Results

The optimal policy for the baseline scenario is determined in approximately 32 minutes using MATLAB 2019a on a Intel Zeon E5-2687W workstation with 64GB RAM, 10 cores, and MATLAB Parallel Computing Toolbox. The optimal policy outperforms the myopic by 1.71% in regards to ETDR. Because of the life or death nature of the MEDEVAC system, any improvement over the current policy is significant and may increase casualty survivability rates. The optimal policy differed from the myopic in 18,945 of the 132,055 states. That is, 14.3% of the optimal policy decisions are different than the myopic policy. Of the 18,945 differences, 15,969 (or 85.5%) of the differences resulted from the myopic policy recommending to send a unit to service an incoming request and the optimal policy recommending to reject the MEDEVAC request to an alternate form of evacuation. The remaining 15.5%

are differences regarding which of the available units to dispatch. The subsequent comparison and analysis of the optimal and myopic policies focuses on the expected amount of time the MEDEVAC system spends in each state depending on the policy implemented as opposed to the specific differences between the two policies.

4.1.3 Baseline Analysis of Specific System States

The optimal policy dictates which action to take given the current state of the system. For example, if all four MEDEVAC are idle and an urgent request arrives from Zone 4, the dispatching authority can examine the optimal policy for the optimal decision and see the recommendation to dispatch MEDEVAC unit 3 to service the request. Table 6 outlines five system states and displays the myopic and optimal decisions.

Scenario 1 represents a MEDEVAC system state wherein 3 of the 4 MEDEVAC units are not available and a request is in the system for Zone 5 with a reported and true priority of routine. The myopic policy recommends dispatching the last remaining MEDEVAC unit, and the optimal policy recommends reserving MEDEVAC Unit 1 and rerouting the request to other forms of evacuation. The optimal policy is reserving the remaining MEDEVAC for potential high priority requests that may arrive before one of the other three MEDEVAC units arrive back to their staging locations. Because the rerouted request is of low priority, CASEVAC or other forms of evacuation can effectively and safely service the request while avoiding the possi-

Table 6. Example MEDEVAC Unit Dispatching Policies

Scenario	S_t	π^{myopic}	π^*
1	$(0,1,1,1),(5,3,3)$	Dispatch Unit 1	Reject
2	$(4,0,0,1),(5,3,3)$	Dispatch Unit 2	Reject
3	$(5,6,0,5),(1,1,3)$	Dispatch Unit 3	Dispatch Unit 3
4	$(3,1,0,0),(3,2,3)$	Dispatch Unit 3	Dispatch Unit 4

bility of a high priority request being rerouted to other evacuation services. Scenario 2 represents a similar situation, but there are two MEDEVAC units available when the request arrives. The optimal policy recommends to not dispatch either of the units and the myopic recommends dispatching MEDEVAC Unit 2. This scenario emphasizes the optimal policy's tendency to reserve units when compared to the myopic policy of always dispatching a unit.

Scenario 3 represents a similar situation to Scenario 1 in which there is one MEDEVAC available and a request arrives in the system. The request arrives from Zone 1 and is a reported urgent request, but the true priority is routine. Because the request priority level is overestimated, the system falsely expects a bigger benefit to dispatching a MEDEVAC unit to service the request. Had the true priority level of routine been presented, the optimal policy may be to reserve the final MEDEVAC unit for a true urgent request. This scenario represents the detriment to the system that triage classification errors cause. Finally, Scenario 4 represents the difference between the myopic and optimal policies regarding which unit to send to service an incoming request when two or more units are available. For this scenario, MEDEVAC Unit 3 and MEDEVAC Unit 4 are available and a request to service Zone 3 arrives in the system. Because the myopic policy recommends dispatching MEDEVAC Unit 3, we know it is the closest. The optimal policy recommends reserving MEDEVAC Unit 3 and instead dispatching Unit 4. Because the optimal policy takes into account possible future states, dispatching the nearest MEDEVAC unit does not always lead to the largest ETDR.

These scenarios represent some of the differences and similarities between the myopic and optimal policies created in this thesis. The 1.71% improvement of the optimal policy over the myopic is a result of the 18,945 differences in system state actions as revealed in Table 6.

4.1.4 MEDEVAC Unit Availability Rates

In the MEDEVAC system, prompt service leads to better outcomes for service requests of all three priority levels, but prompt service is imperative for urgent priority casualties due to the nature of their classification. If an urgent request arrives in the system and there is at least one unit available, the urgent request will be serviced and the MEDEVAC system can avoid rejecting the request and sending it to another form of evacuation. Let the availability rate represent the expected proportion of time a specific number of MEDEVAC units are available. If the availability rate for four units is 0.44, we expect that 44% of the time there are four units available to service an incoming request in the MEDEVAC system. Figure 3 shows the proportion of time the system expects to have any number of units available for the baseline scenario.

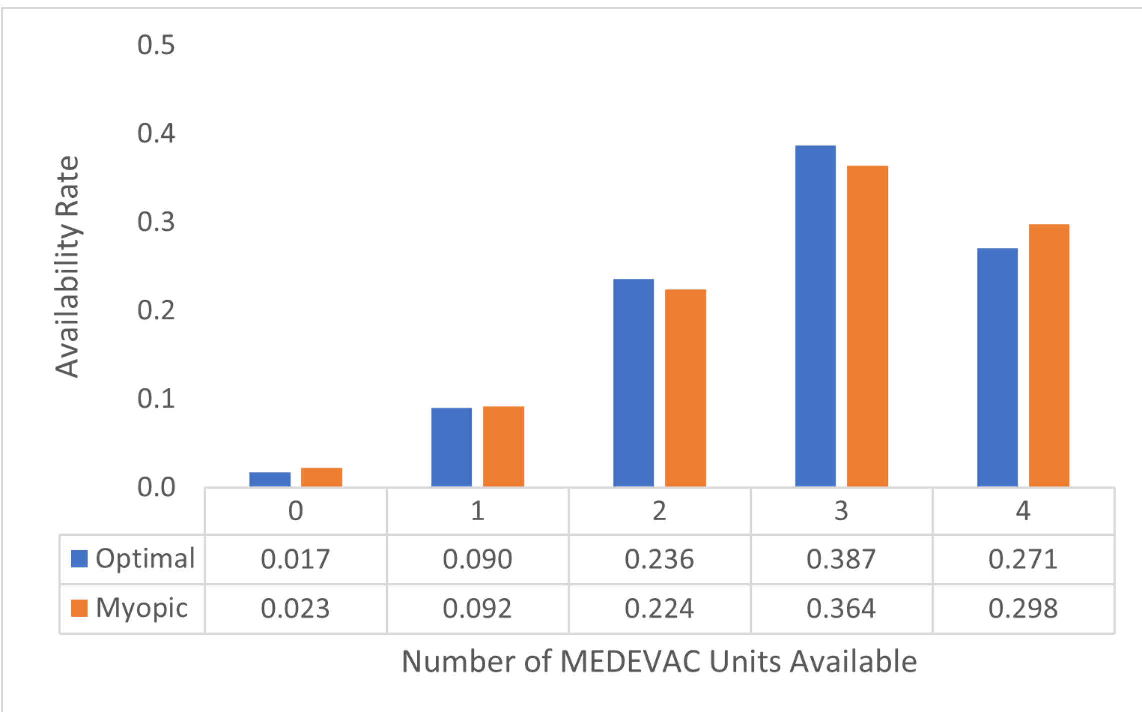

Figure 3. MEDEVAC Unit Availability Rates - Baseline Scenario

Almost 90% of the time (89.32%) the MEDEVAC system will have at least two units available. A difference between the optimal and myopic policies is that the op-

timal policy reserves a MEDEVAC unit for possible incoming urgent requests. When implementing the optimal policy, the system has zero MEDEVAC units available 1.7% of the time. This indicates if a urgent request is reported, only 1.7% of the time would we expect the MEDEVAC system to reject the request and redirect it to another form of evacuation. The myopic policy's availability rate for zero units is 2.3%. For both the myopic and optimal policy, three MEDEVAC units are available the largest amount of time. An availability rate graph with a right skew would represent a system that struggles to meet the demand in incoming service requests, whereas a graph with a significant left skew would represent a MEDEVAC system that does not efficiently use MEDEVAC units (i.e., they are sitting idle the majority of the time). Figure 3 shows a good balance between the two extremes for both the myopic and optimal policies.

4.1.5 MEDEVAC Unit Utilization Rates

Four MEDEVAC units are considered in the baseline scenario. Let the utilization rate for each MEDEVAC unit represent the expected proportion of time a MEDEVAC unit spends actively servicing requests. For example, a utilization rate of 0.4 for MEDEVAC Unit 3 indicates that we expect MEDEVAC Unit 3 to be servicing requests 40% of the time in the long run. A high utilization rate for MEDEVAC units may lead to increased wear and tear on equipment and overworking of MEDEVAC personnel. A low utilization rate may indicate that there are possibly too many MEDEVAC units assigned. Figure 4 provides the utilization rate of each unit for both the myopic and optimal policies.

Comparing the myopic and the optimal policies, the myopic policy has an unbalanced use of the MEDEVAC units. When following the myopic policy, we expect MEDEVAC Unit 1 to be used at least 5% more than the other three units. Whereas

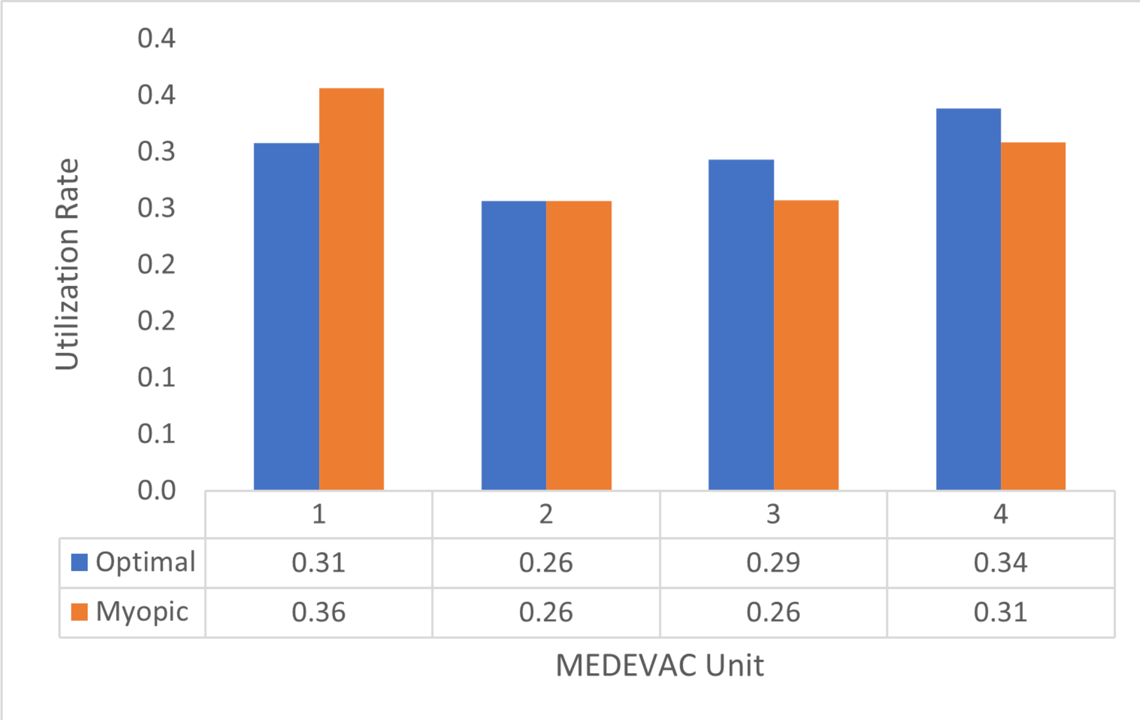

Figure 4. MEDEVAC Unit Utilization Rates - Baseline Scenario

MEDEVAC Unit 4 has the highest utilization rate when the optimal policy is implemented and we expect it to be used at least 3% more than the other three units in the optimal policy. We expect each of the four MEDEVAC units to be be busy less than 40% of the time for the both myopic and optimal policies. Because the utilization rate for MEDEVAC Unit 1 is disproportionately higher than the other three units for the myopic policy, we know that more casualties occur near MEDEVAC Unit 1 than any other unit. The optimal policy balances the burden of responding to MEDEVAC service requests more evenly across the four available units.

4.1.6 MEDEVAC Unit Zone Allocation

Because of the nature of our notional data, we expect an uneven distribution of casualty events across the 6 zones. Additionally, because dispatching policies vary on their recommendations of which MEDEVAC unit to dispatch to service a request,

the percentage of time each MEDEVAC unit spends servicing each zone varies across policies. Given that the MEDEVAC is servicing a request, Figure 5 illustrates the percentage of time each unit spends in each zone.

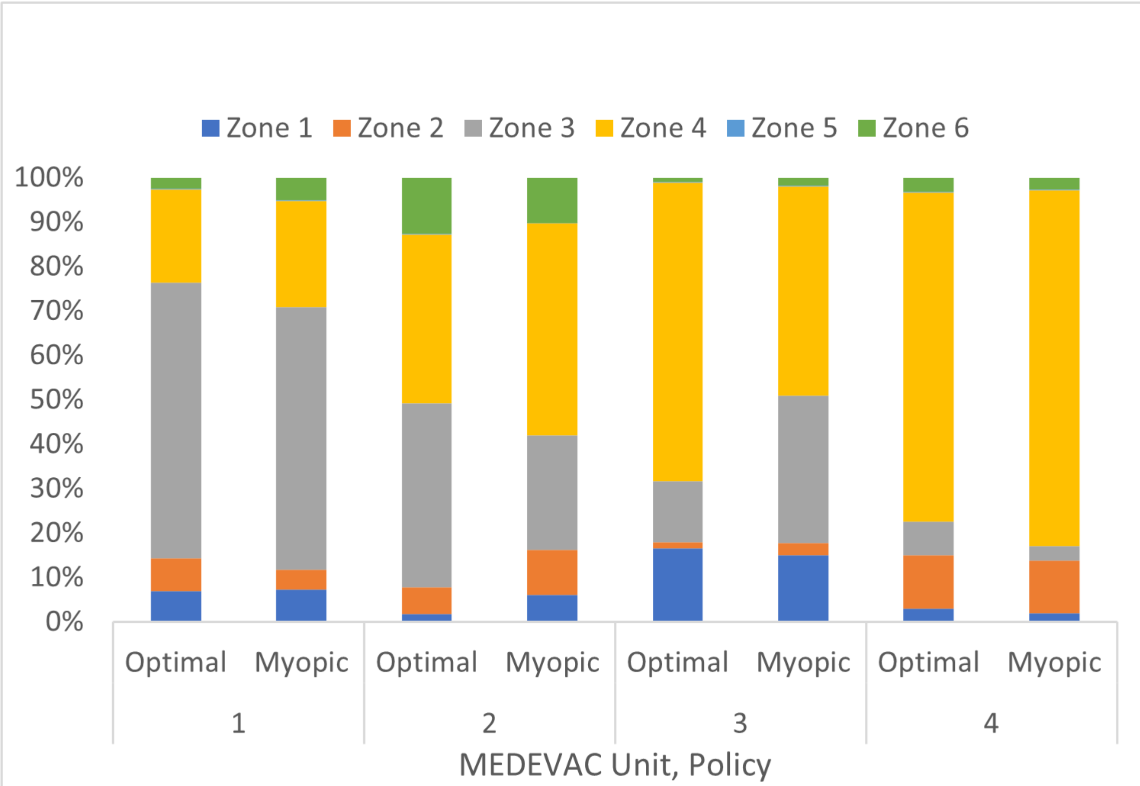

Figure 5. MEDEVAC Unit Zone Allocation

The chart illustrates that all four MEDEVAC units spend the majority of their time servicing requests in Zones 3 and 4, which is where the majority of CCPs are located. When the myopic policy is implemented, the proportion of time MEDEVAC Units 1, 2, and 4 spend servicing Zone 3 increases and the proportion of time spent servicing Zone 4 decreases. The least visited zone by all four MEDEVAC units is Zone 5. Whereas there are distinct differences between the optimal and myopic policies for zone allocation, implementing the optimal policy over the myopic policy would not drastically change the day to day operations of the MEDEVAC units. Only small adjustments need to be made for where each unit spends their time proportionally

according to the policy implemented.

4.1.7 Optimal vs Myopic Policy Analysis

Comparison of the myopic policy and optimal policy is important to understand the reason why the optimal policy has a higher ETDR in the long run. Table 7 tabulates the optimal policy, the myopic policy, and the percentage of time we can expect the policies to be the same or different. For example, given both policies are evaluating the same S_t, we expect 0.13% of the time the myopic policy to recommend dispatching MEDEVAC Unit 2 and the optimal policy to recommend not dispatching an unit.

The large majority of time, both policies recommend not dispatching a MEDEVAC unit. This is the proportion of time no request is in the system. If a request was in the system, the myopic policy would recommend dispatching the nearest available vehicle. We expect that 96.8% of the time the optimal policy and the myopic policy recommend the same course of action. The most common difference between the two is when the myopic policy recommends dispatching MEDEVAC Unit 1 and the optimal recommends dispatching MEDEVAC Unit 3, 2, or 4. This event occurs an expected 1.66% of the time. While the myopic policy always dispatches a MEDEVAC unit when a service request arrives and a MEDEVAC unit is available, the optimal policy can reserve a MEDEVAC unit for future high priority service requests. The percentage of time we expect the optimal policy to reserve units that the myopic policy would recommend to dispatch is 0.4%. The misclassification rate affects the expected amount of time the optimal policy recommends to reserve units.

The optimal policy recommends which unit to dispatch based on expected future states, but the misclassification rate introduces more uncertainty into the model and decreases the optimal policy's ability to make an effective decision. As the misclassi-

Table 7. Optimal vs Myopic Policy

π^{myopic}	π^*	%
Reject	Reject	85.8
Dispatch Unit 4	Dispatch Unit 4	3.06
Dispatch Unit 3	Dispatch Unit 3	2.71
Dispatch Unit 2	Dispatch Unit 2	2.69
Dispatch Unit 1	Dispatch Unit 1	2.53
Dispatch Unit 1	Dispatch Unit 3	0.93
Dispatch Unit 1	Dispatch Unit 2	0.56
Dispatch Unit 2	Dispatch Unit 3	0.27
Dispatch Unit 4	Dispatch Unit 1	0.27
Dispatch Unit 1	Dispatch Unit 4	0.17
Dispatch Unit 2	Dispatch Unit 4	0.17
Dispatch Unit 3	Dispatch Unit 2	0.17
Dispatch Unit 2	Dispatch Unit 1	0.15
Dispatch Unit 2	Reject	0.13
Dispatch Unit 1	Reject	0.1
Dispatch Unit 3	Reject	0.1
Dispatch Unit 3	Reject	0.09
Dispatch Unit 4	Reject	0.08
Dispatch Unit 3	Dispatch Unit 1	0.01
Dispatch Unit 4	Dispatch Unit 2	0
Dispatch Unit 4	Dispatch Unit 3	0
Reject	Dispatch Unit 1	0
Reject	Dispatch Unit 2	0
Reject	Dispatch Unit 3	0
Reject	Dispatch Unit 4	0

fication rate increases, the proportion of requests that appear high priority increase. The system responds based on the reported priority, and if the majority of requests received are urgent or routine, the system is unable to prioritize true urgent or priority casualties and is less likely to reserve units as it would with a lower misclassification rate.

4.2 Excursions

Sensitivity analysis for this thesis focuses on altering parameters likely to change in a deployed environment. We focus on the misclassification rate, number of blood

transfusion kits available on specific MEDEVAC aircraft, arrival rate of MEDEVAC service requests, MEDEVAC aircraft speed, and proportion of urgent, priority, and routine requests received by the MEDEVAC system. For each excursion we vary one parameter and keep the other parameters of the baseline scenario constant. To compare the effects of altering each parameter, we compare the optimality gap between the myopic and optimal policies for each parameter change, availability rate of MEDEVAC units, and utilization rate of MEDEVAC units.

4.2.1 Excursion 1: Misclassification Rate

The misclassification rate introduces uncertainty into the model. As the priority level of an incoming casualty event becomes less certain, the MDP model adapts taking into consideration the additional uncertainty. This excursion investigates total misclassification rates 0, 0.3, 0.5, and 0.8. We expect as the total misclassification rate increases, the optimal policy will perform more similarly to the myopic policy due to the incorrect information received. Figure 6 depicts the ETDR for both the myopic and optimal solutions for each scenario. Additionally, the right axis depicts the percent difference between the two ETDRs. A higher percent difference between the two policies indicates a larger performance gap.

As the total misclassification rate increases, the ETDR decreases slightly, and the percent difference between the myopic and optimal policies for each total misclassification rate decreases. A higher total misclassification rate means the system has to make a less informed recommendation. When the system cannot differentiate between service request priority levels accurately, it acts similarly to the myopic policy and the percent difference between the two policies decreases. The ETDR does not decrease significantly for the increased total misclassification rate because the system is not under high stress. The priority level of the incoming request is weighted heavier

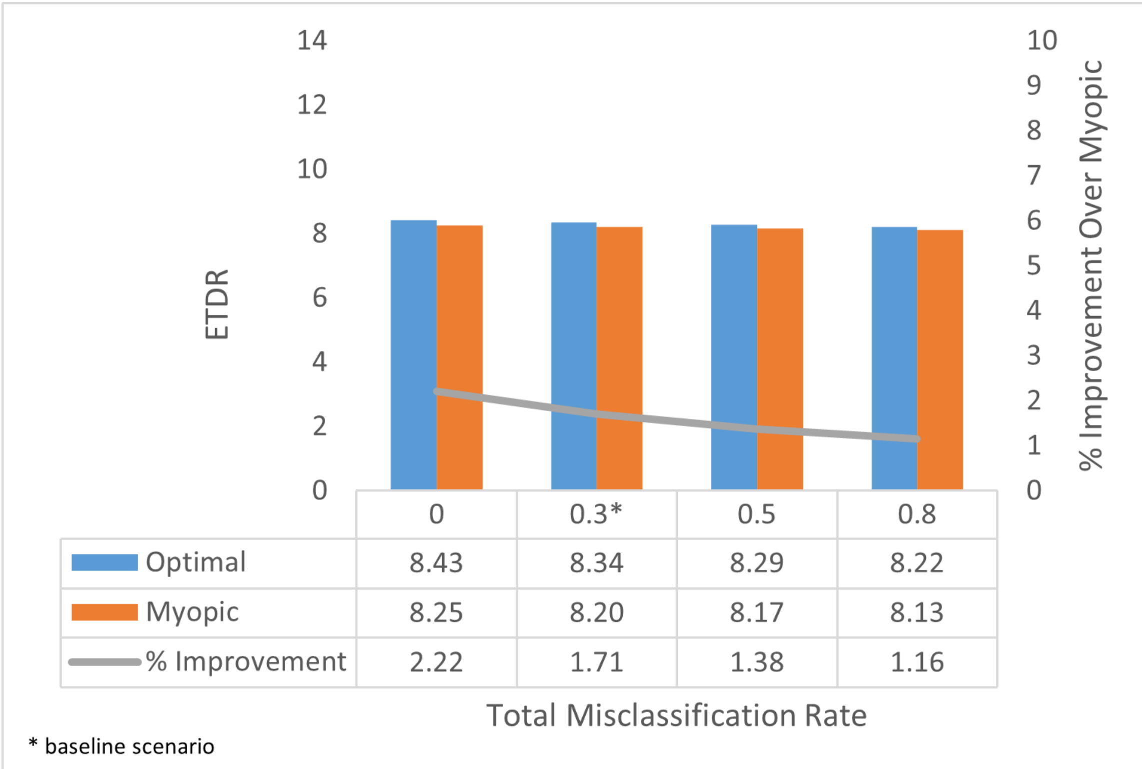

Figure 6. ETDR - Misclassification Rate

when there is only one or two MEDEVAC units available. If there are more units available, the optimal policy is more likely to recommend dispatching a MEDEVAC unit, and it will perform more similarly to the myopic policy.

Figure 7 illustrates how the MEDEVAC unit availability rate changes with each total misclassification rate. For all four total misclassification rates, we expect there to be 3 MEDEVAC units available the majority of the time. When the total misclassification rate is 0, the percentage of time we expect there to be 0 units available is lower than the other three total misclassification rates by less than 0.01. While having more MEDEVAC units available more often reduces the risk that a urgent or priority request will be rejected from the system, the total misclassification rate does not seem to have a practical impact on the availability rate for the MEDEVAC units.

As the information received by the MEDEVAC dispatching system becomes more

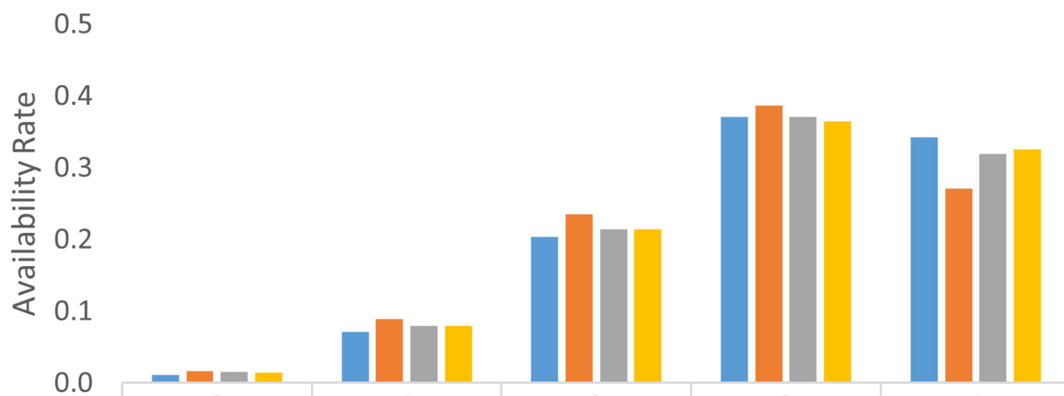

	0	1	2	3	4
0	0.011	0.071	0.204	0.371	0.343
0.3*	0.017	0.090	0.236	0.387	0.271
0.5	0.016	0.080	0.214	0.371	0.319
0.8	0.015	0.080	0.214	0.365	0.326

Number of MEDEVAC Units Available

* baseline scenario

Figure 7. MEDEVAC Unit Availability Rates - Misclassification Rate

inaccurate with increasing total misclassification rates, the practical value of the priority level in determining whether or not to send a MEDEVAC unit and which unit to send decreases. This will in turn alter the utilization rates of each MEDEVAC unit illustrated in Figure 8.

Altering the total misclassification rate affects the utilization rate of each MEDEVAC unit, but not in a specific pattern. Figure 8 illustrates how the small adjustments and changes the MDP model makes when calculating the optimal policy. MEDEVAC Unit 4 and 2 are more commonly used when the total misclassification rate is 0.8 and used less when it is 0. This could be because as the priority classification became more uncertain, the optimal policy does not attempt to reserve either unit for incoming higher priority, but these shifts could be for many different reasons that cumulatively alter the utilization rates of each unit.

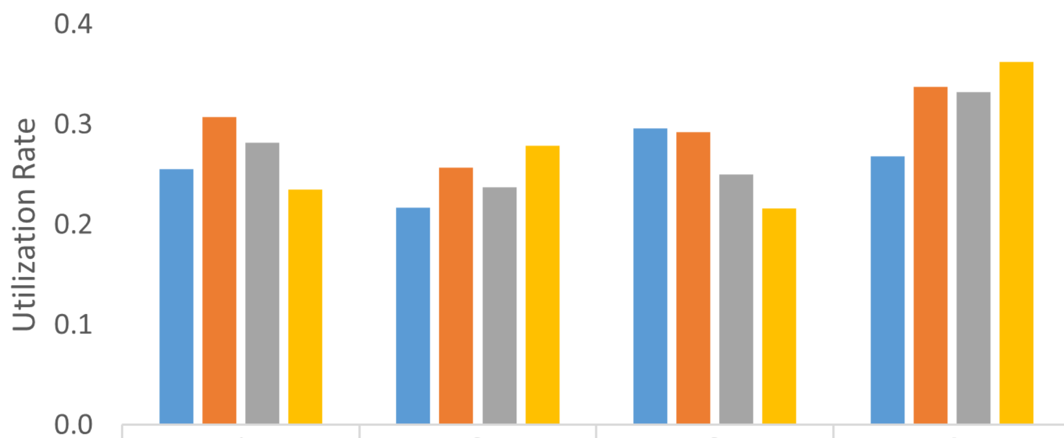

	1	2	3	4
0	0.26	0.22	0.30	0.27
0.3*	0.31	0.26	0.29	0.34
0.5	0.28	0.24	0.25	0.33
0.8	0.24	0.28	0.22	0.36

MEDEVAC Unit

* baseline scenario

Figure 8. MEDEVAC Unit Utilization Rates - Misclassification Rate

4.2.2 Excursion 2: Presence of Blood Transfusion Kits on MEDEVAC Aircraft

The presence of blood transfusion kits on board MEDEVAC units decreases the time until an injured individual receives necessary medical care. As more MEDEVAC units are equipped with blood transfusion kits, more individuals receive critical care earlier, and the overall efficiency of the MEDEVAC system increases. Adding blood transfusion kits decreases the response time of the MEDEVAC units equipped. We expect this to increase the ETDR and make the system more efficient. MEDEVAC Units 1 and 4 are co-located with MTFs for the Azerbaijan scenario indicating that they can be equipped with blood transfusion kits and deliver life saving care earlier. This excursion includes four scenarios, a scenario where neither MEDEVAC units are equipped with blood transfusion kits (i.e., the baseline scenario), only MEDEVAC

Unit 1 is equipped, only MEDEVAC Unit 4 is equipped, and both MEDEVAC Units 1 and 4 are equipped with blood transfusion kits. Figure 9 depicts the ETDR for each scenario implementing both the optimal and myopic policies and the percent difference between the two.

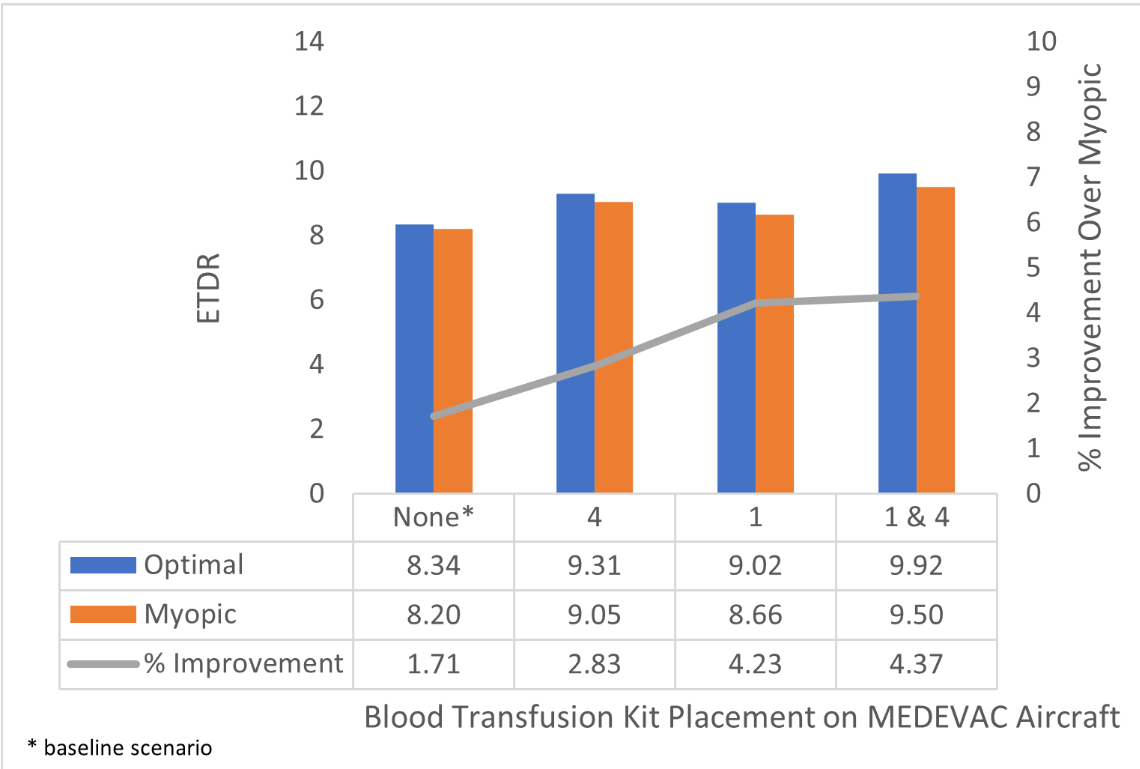

Figure 9. ETDR - Blood Transfusion Kits

Figure 9 shows that when a blood transfusion kit is included on board both MEDEVAC Units 1 and 4, the model receives the highest reward and the highest percent improvement in ETDR. When MEDEVAC Unit 4 is equipped with a blood transfusion kit, the MEDEVAC system receives a higher reward in total, and the optimal policy sees a higher improvement over the myopic policy when compared to equipping MEDEVAC Unit 1 with a blood transfusion kit. MEDEVAC Unit 4 may be better placed to service higher priority requests when equipped with a blood transfusion kit. If only one kit were to be placed, we recommend placing it on MEDEVAC Unit 4.

Figure 10 depicts the availability rate for each number of units. Although if both units could be equipped with kits, we would see a substantial increase in overall efficiency of the system.

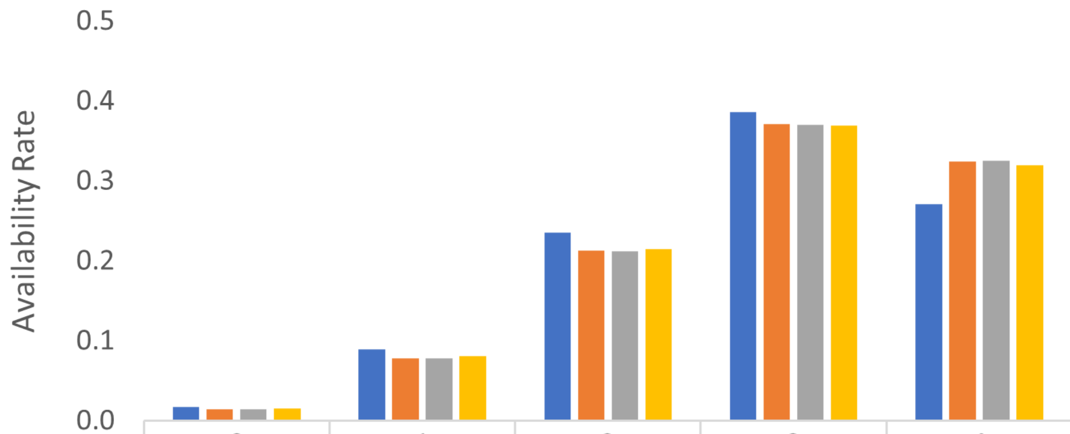

	0	1	2	3	4
None*	0.017	0.090	0.236	0.387	0.271
4	0.015	0.078	0.213	0.371	0.324
1	0.015	0.079	0.212	0.370	0.325
1 & 4	0.015	0.081	0.215	0.369	0.320

Number of MEDEVAC Units Available

* baseline scenario

Figure 10. MEDEVAC Unit Availability Rates - Blood Transfusion Kits

Figure 10 depicts that the addition of blood transfusion kits increases the expected number of units available at any point in time. Going from 0 kits to at least 1 kit decreases the expected amount of time 0, 1, 2, or 3 MEDEVAC units are available and increases the expected amount of time 4 MEDEVAC units are available. This may be because the system is more likely to reserve aircraft to care for urgent requests, when the reward is significantly increased when blood transfusion kits are included on board MEDEVAC aircraft.

Figure 11 displays the utilization graph for the blood transfusion kits, indicates a decrease in utilization for the MEDEVAC unit when it is equipped with a blood

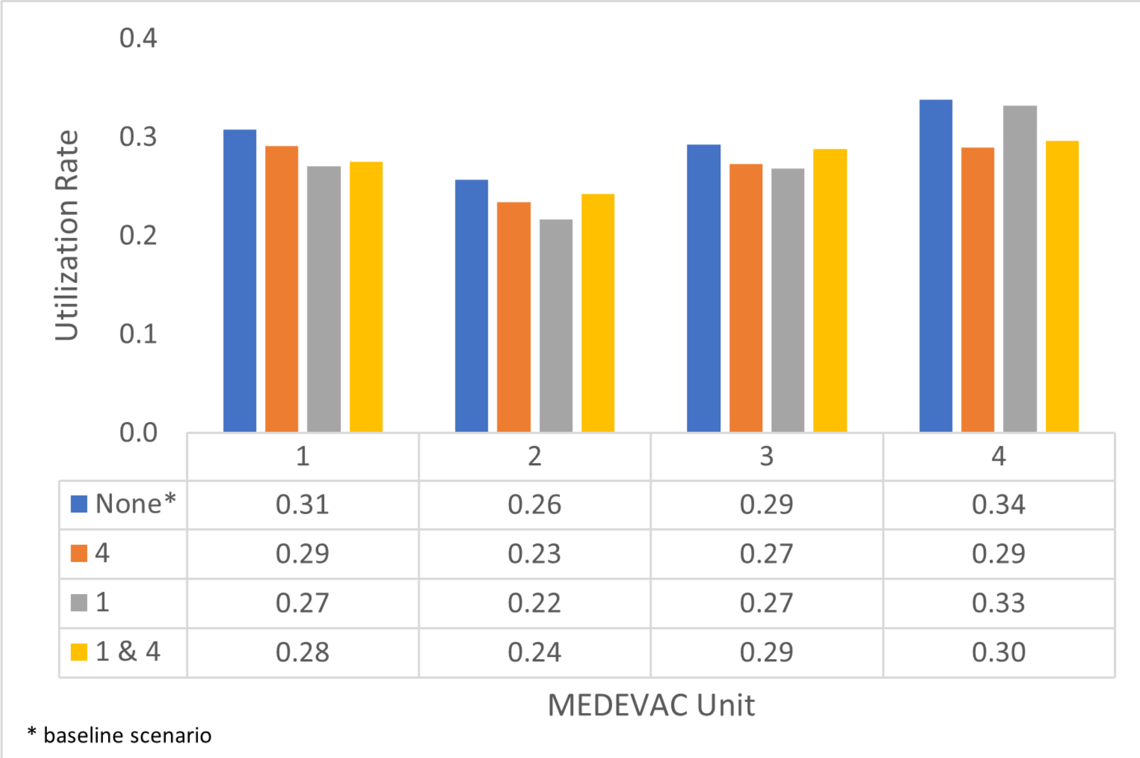

Figure 11. MEDEVAC Unit Utilization Rates - Blood Transfusion Kits

transfusion kit. The optimal policy is more likely to reserve the MEDEVAC units equipped with blood transfusion kits lowering the overall utilization and increasing the number of routine requests that are rejected from the system.

4.2.3 Excursion 3: Arrival Rate of Casualty Events

The arrival rate λ represents the frequency that the system receives MEDEVAC requests. A system that is receiving requests every few minutes will respond differently than a system receiving requests every few hours. This excursion investigates arrival rates of 1/30, 1/60, 1/90, and 1/120. An increased arrival rate introduces more stress in the system, and the dispatching policy has increased importance when compared to a system in which casualties are arriving over a longer period of time. The arrival rate is not a parameter the MEDEVAC system authorities can control, rather they have to react to the arrival rate and ensure fast and efficient care is being given to

injured individuals. We expect that as the arrival rate increases, the percent difference between the optimal and myopic policies will increase. Figure 12 shows the ETDR graph for all four arrival rates.

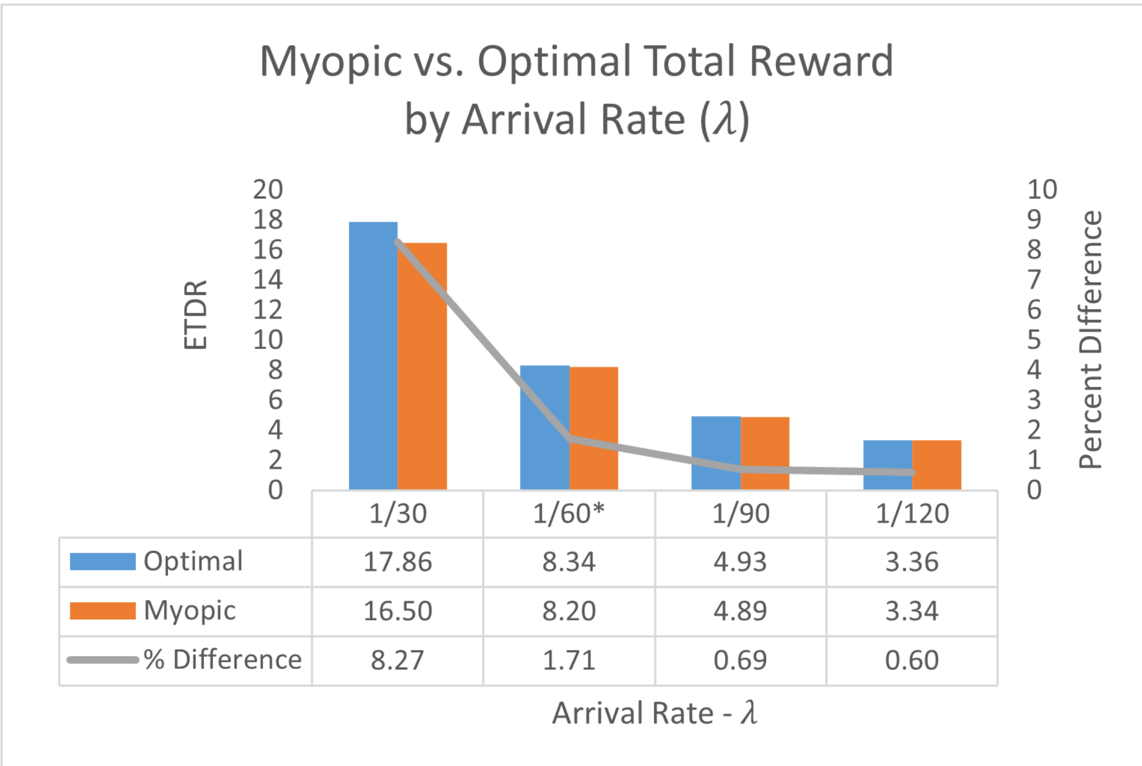

Figure 12. ETDR - Arrival Rate

Figure 12 shows the differences in reward between myopic and optimal dispatching policies for all four arrival rates. As the arrival rate decreases, we expect fewer casualty requests to arrive in the system and the expected reward to decrease due to fewer opportunities to service requests. Additionally, the percent improvement decreases. If one request arrives every two hours with $\lambda = 1/120$ there is less of a reason to reserve a MEDEVAC unit and reject the request from the system because there is a lower likelihood that another request will arrive in the system and need the MEDEVAC unit. As the arrival rate increases, whether or not the system dispatches a MEDEVAC matters more because there is a higher likelihood that an urgent request will arrive

and need a MEDEVAC unit. The availability rate chart in Figure 13 illustrates the stress on the system.

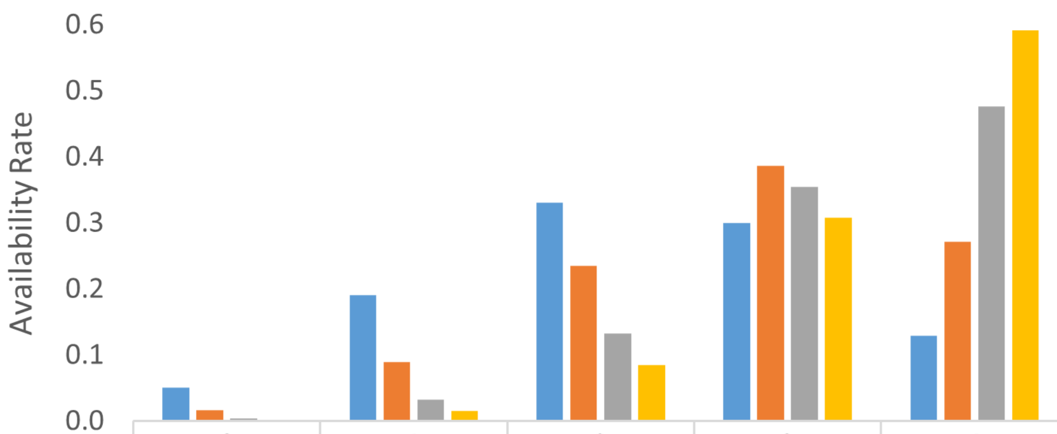

	0	1	2	3	4
1/30	0.051	0.190	0.330	0.300	0.129
1/60*	0.017	0.090	0.236	0.387	0.271
1/90	0.004	0.032	0.132	0.355	0.477
1/120	0.002	0.015	0.085	0.308	0.591

Number of MEDEVAC Units Available

* baseline scenario

Figure 13. MEDEVAC Unit Availability Rates - Arrival Rate

Looking at the availability rates for each arrival rate, as the arrival rate increases, the number of units available shifts to the left so more often fewer units are available to be dispatched. As the rate is decreased from 1/30 to 1/60 we see the biggest decrease in the expected amount of time there are zero units available. As mentioned before, this is the expected percentage of time that if a request were to arrive, the MEDEVAC dispatching authority has to reject the request from the system and relay it to other evacuation authorities. Comparing the four arrival rates, the smaller the arrival rate, the further left skewed the availability rate chart becomes. Although a left skewed chart may seem desirable, it raises the question whether the system needs as many MEDEVAC units as it currently has in place.

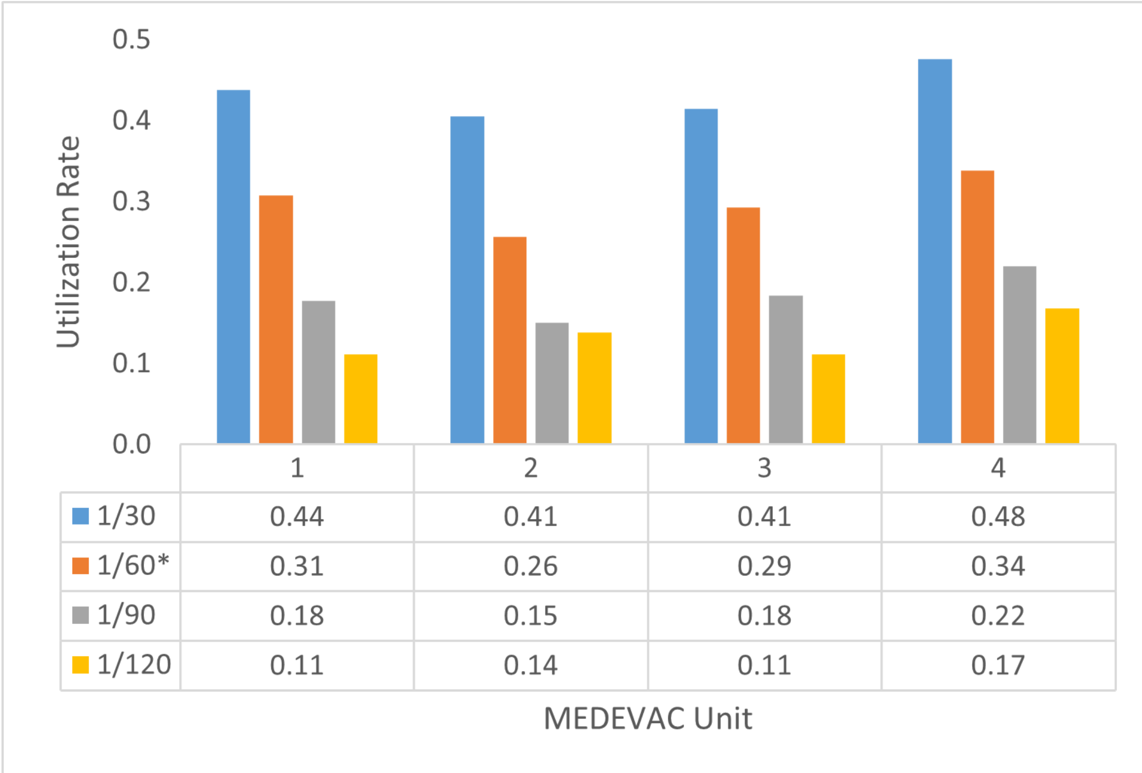

Figure 14. MEDEVAC Unit Utilization Rates - Arrival Rate

The utilization rate of each MEDEVAC unit by arrival rate show the same result in that the smaller the arrival rate, the lower we can expect the utilization rate to be. Even when the arrival rate is $\lambda = 1/30$ we expect the MEDEVAC units to be servicing requests less than half of the time. When the rate is decreased to $\lambda = 1/120$ we expect the MEDEVAC units to be busy less than 20% of the time. The utilization rates of the MEDEVAC units are highly affected by the arrival rate of incoming service requests. Moreover, accurately reporting the parameter leads to more efficient analysis.

4.2.4 Excursion 4: MEDEVAC Aircraft Speed

Similar to the presence of blood transfusion kits, increased speed of MEDEVAC aircraft leads to increased efficiency in the MEDEVAC system. The Bell V-280 Valor is a possible replacement for the current HH-60M aircraft and has a cruise speed

of 280 knots as opposed to the HH-60M cruise speed of 150 knots. A faster aircraft suggests that not only will injured individuals receive initial care from the MEDEVAC unit faster, but they will also be delivered to an MTF faster. Figure 15 depicts the ETDR for both aircraft types and the percent difference between the two.

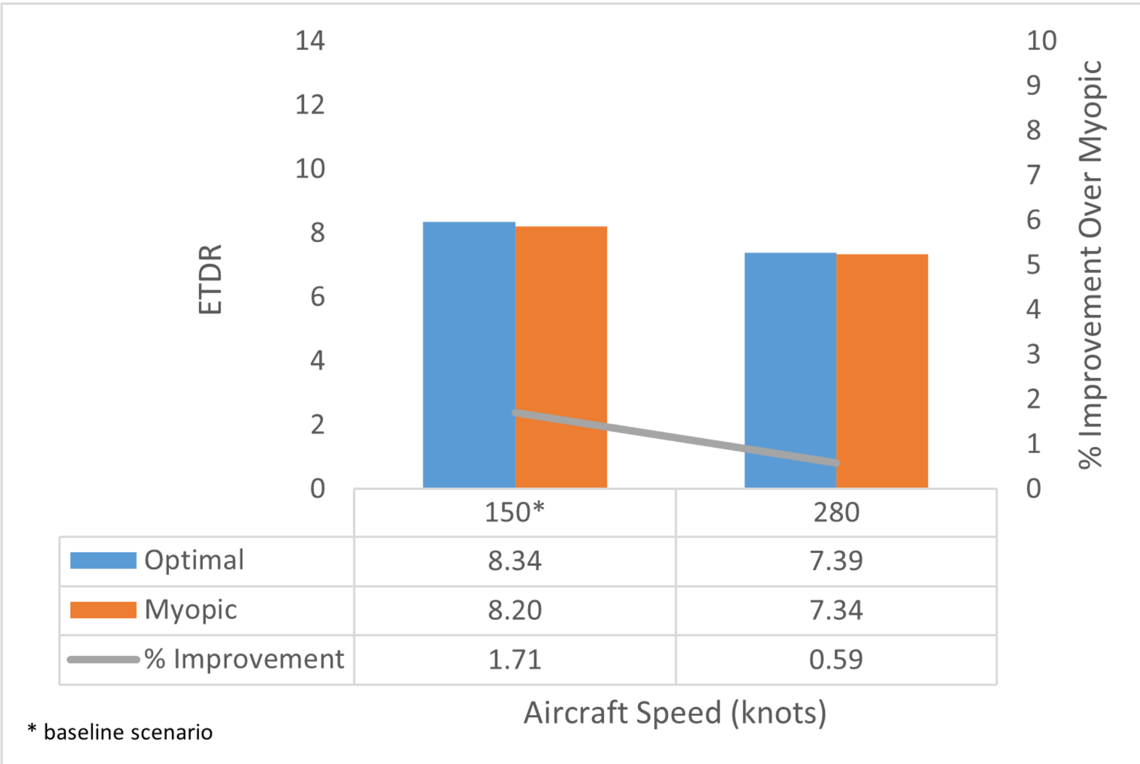

Figure 15. ETDR - Aircraft Speed

The ETDR shows that as the aircraft speed increases, the percent difference between the optimal and myopic policies decreases. Because the aircraft take less time to service a request, there are more MEDEVAC units able to be dispatched at any one time.

Because the Valor aircraft is faster, we expect to be able to complete more service requests in a faster amount of time. Therefore, more units will be available a larger proportion of time. The aircraft speed has a significant impact on the percentage of time four units are available. The availability chart depicted in Figure 16 is heavily

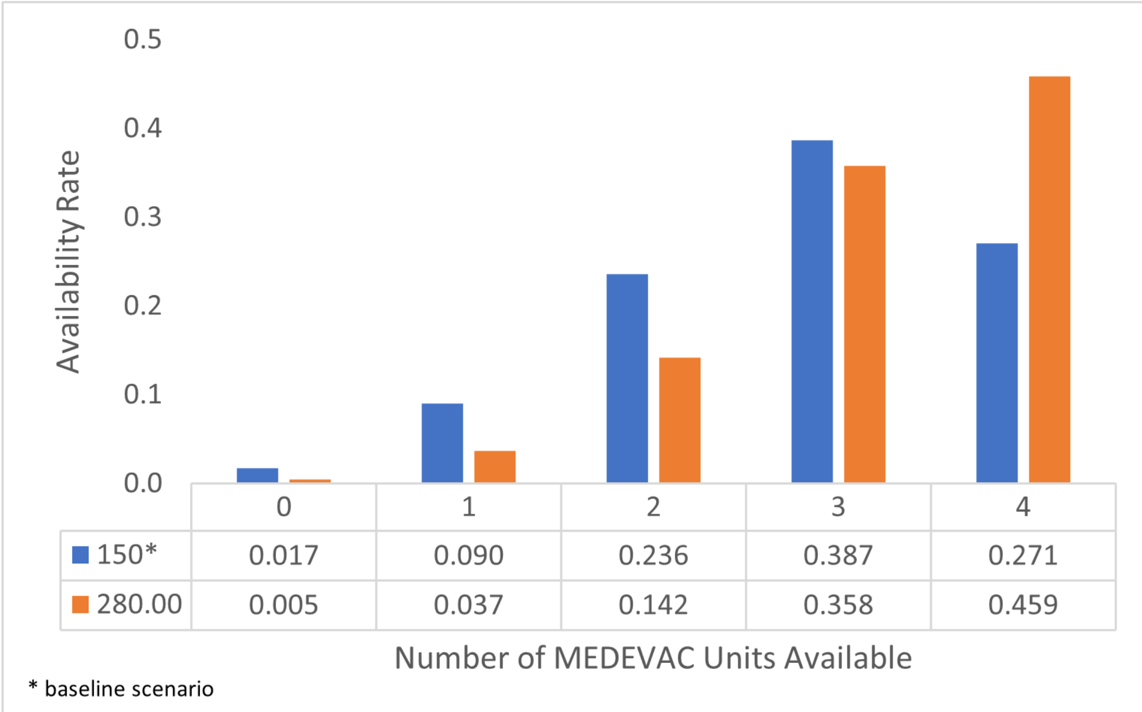

Figure 16. MEDEVAC Unit Availability Rates - Aircraft Speed

left skewed. Additional analysis should be performed on the possible removal of a MEDEVAC unit if there were three other MEDEVAC aircraft available that have an average speed of 280 knots.

Similar to the availability chart, the increase in aircraft speed significantly decreases the expected utilization of MEDEVAC unit as indicated in Figure 17. Each unit spends less time servicing requests and more time waiting for new requests to arrive. With the implementation of the faster aircraft, we expect each MEDEVAC unit to be busy for less than 25% of the time as compared to the average utilization for the HH-60M of 30%.

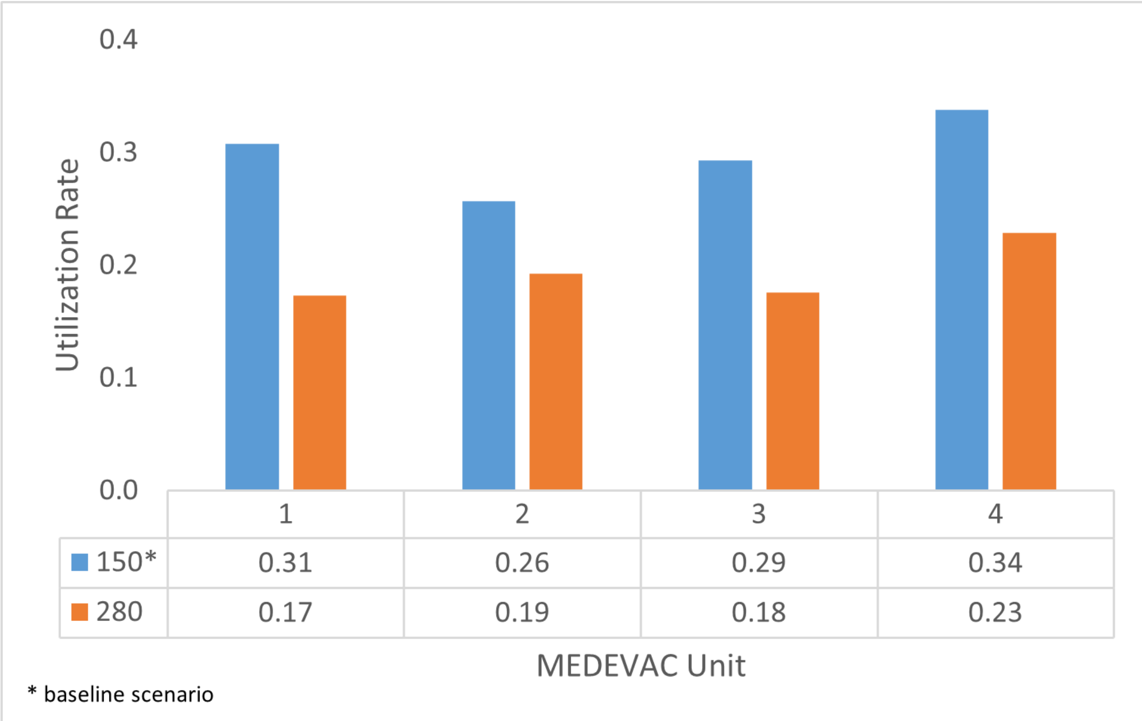

Figure 17. MEDEVAC Unit Utilization Rates - Aircraft Speed

4.2.5 Excursion 5: Proportion of Urgent, Priority and Routine Requests

The baseline scenario assumes that the MEDEVAC system receives an approximately equal proportion of urgent, priority, and routine requests. The MEDEVAC system would need to adjust if a disproportionate number of one of the priority level requests was arriving. A higher proportion of urgent requests increases stress on the system because there is less flexibility in choosing to reject a request from the system. The MEDEVAC system is the fastest form of evacuation when compared to other evacuation services (i.e., CASEVAC) and when an urgent request arrives, it is imperative it is serviced quickly. Unless there are no units available, the MEDEVAC system will not reject an urgent request. If there is a high proportion of urgent requests, there is a higher likelihood of having to reject an urgent request due to unit unavailability. Moreover, a high proportion of routine requests will increase the flex-

ibility of the dispatching authority due to the fact that should a routine MEDEVAC request be turned away by the dispatching authority, CASEVAC is capable of providing the necessary care. When there is a high number of urgent or routine requests, the dispatching authority accepts more risk when turning away an urgent request. The MEDEVAC system is much more likely to reject a routine request from the system when compared to the other priority levels. For this excursion, we create three scenarios where each priority level has an increased proportion of service requests as shown in Table 8.

Table 8. Priority Proportion Excursion Scenarios

Scenario	Urgent	Priority	Routine
Baseline	1/3	1/3	1/3
1	1/2	1/4	1/4
2	1/4	1/2	1/4
3	1/4	1/4	1/2

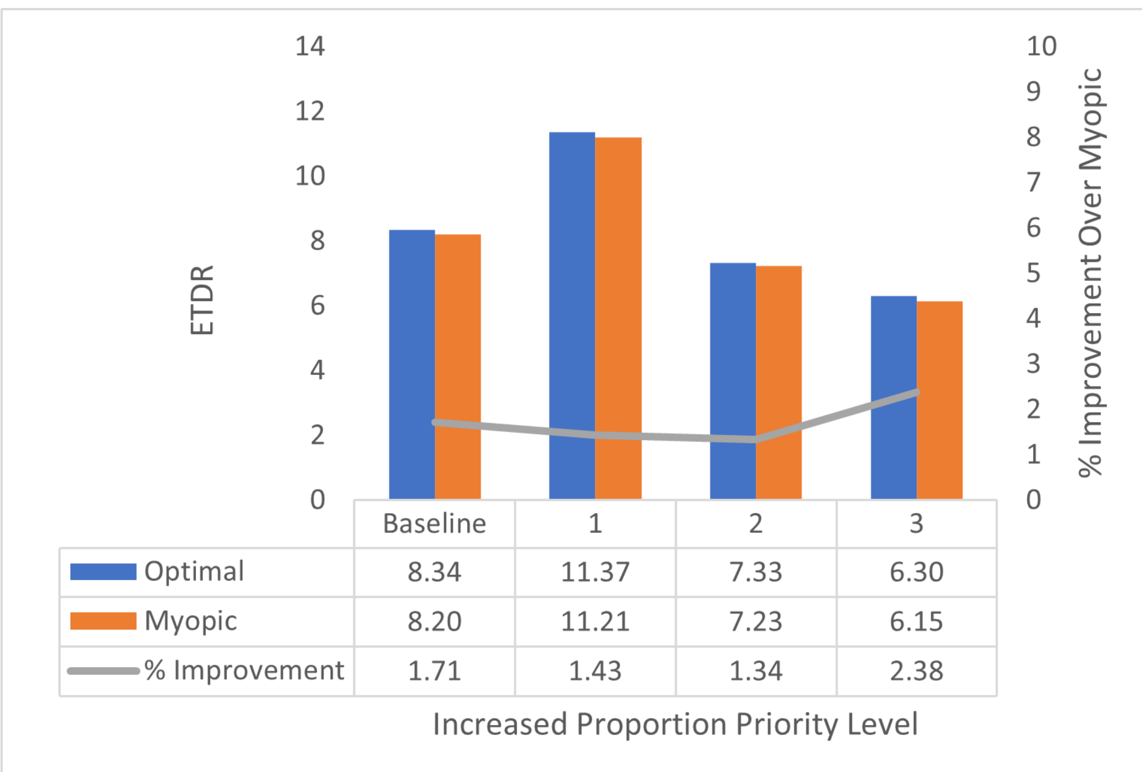

Figure 18. ETDR - Priority Proportion

The highest total ETDR for this excursion occurs when there is a higher proportion of urgent requests. This is because the system earns more reward for servicing urgent requests and there is a better opportunity to service these requests when there are more coming in. Even though there is a higher reward for priority requests than routine requests we do not see a large gap in ETDR between Scenario 2 and Scenario 3 as compared to the difference between Scenarios 1 and 2. This can be attributed to the decreased proportion of urgent requests. The optimal policy implementation has the largest improvement over the myopic when implemented on Scenario 3 with a 2.38% improvement. This could be due to the fact that the optimal policy is more likely to reject incoming requests from the system if they are routine requests. The total misclassification rate has less of an effect in this scenario as well because we assume routine requests are not being misclassified as often as higher priority requests. This gives the model more reliable information about the incoming requests and allows it to be more confident in reserving MEDEVAC units for future requests.

The availability rate chart in Figure 19 shows that, the higher proportion of lower priority requests, the farther the availability chart skews to the left. The baseline availability is the farthest right skewed because it has the lowest predictive power. The equal probabilities between all three priority probabilities give less information than the other three scenarios, which have a more likely priority level. The optimal solution for the three scenarios is able to react to the higher proportion priority and offsets the detriment to the model that the total misclassification rate brings.

Looking at the MEDEVAC unit utilization chart in Figure 20, as the proportion of lower priority casualty increases, the utilization of MEDEVAC Units 1, 2, and 4 decreases. MEDEVAC 3 sees the opposite effect, and its utilization increases as the proportion of low priority requests increases. The baseline scenario has the highest utilization for all four MEDEVAC units, and this can be explained by it's lack of

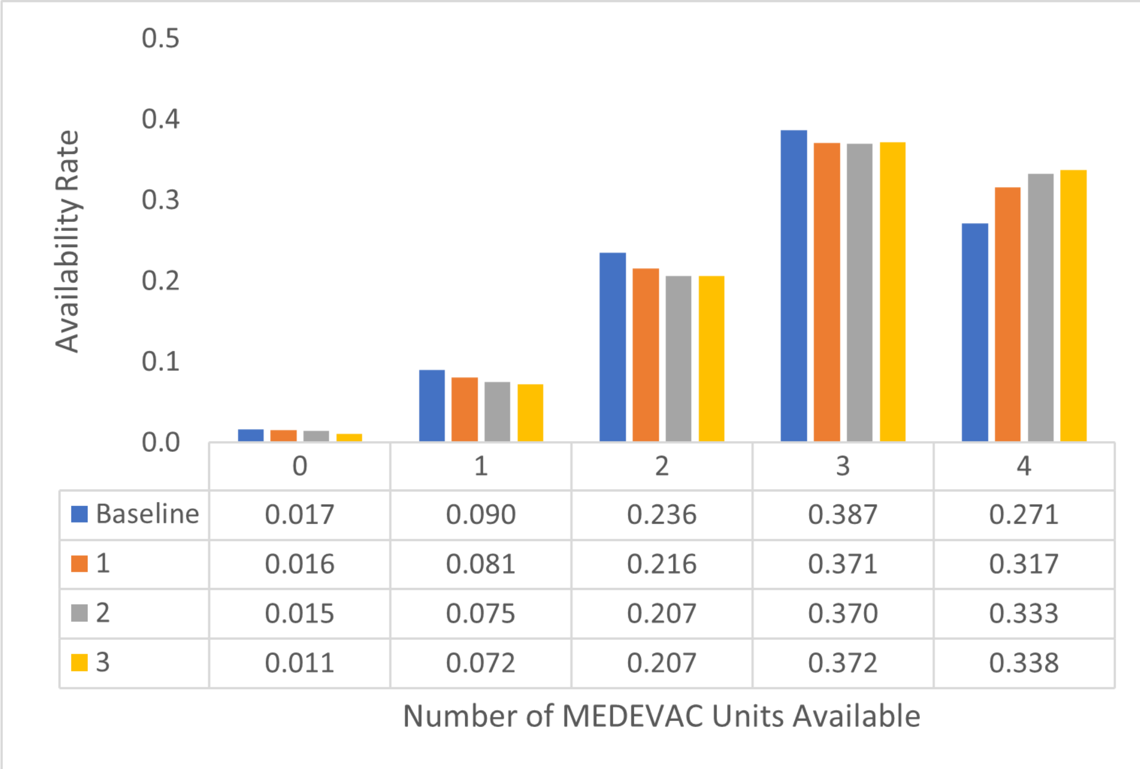

	0	1	2	3	4
Baseline	0.017	0.090	0.236	0.387	0.271
1	0.016	0.081	0.216	0.371	0.317
2	0.015	0.075	0.207	0.370	0.333
3	0.011	0.072	0.207	0.372	0.338

Figure 19. MEDEVAC Unit Availability Rates - Priority Proportion

predictive power that the other three scenarios have because of uneven priority level proportions.

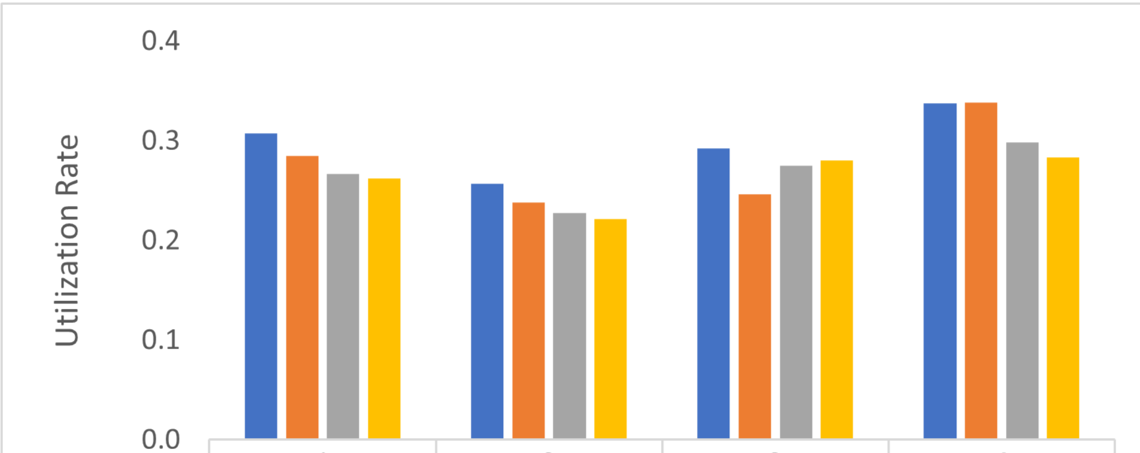

	1	2	3	4
Baseline	0.31	0.26	0.29	0.34
1	0.29	0.24	0.25	0.34
2	0.27	0.23	0.28	0.30
3	0.26	0.22	0.28	0.28

MEDEVAC Unit

Figure 20. MEDEVAC Unit Utilization Rates - Priority Proportion

V. Conclusions & Recommendations

The objective of this thesis is to determine the optimal dispatching policy of medical evacuation (MEDEVAC) units to complete MEDEVAC service requests by examining the MEDEVAC dispatching problem. Improving the performance of the MEDEVAC system leads to higher efficiency and ultimately improve battlefield survivability rates. We develop a discounted, infinite-horizon Markov decision process (MDP) to examine military medical planning scenarios. As an augmentation to previous research, this thesis incorporates the possibility of triage classification errors and the placement of blood transfusion kits on board select MEDEVAC aircraft. Triage classification errors occur when the reported priority level of a casualty event is different from the true priority level, which is assessed when the MEDEVAC unit arrives at the casualty site. Blood transfusion kits allow for life saving care to be administered as soon as an injured individual is loaded onto the MEDEVAC aircraft as opposed to waiting until the MEDEVAC unit arrives back at the medical treatment facility (MTF) and the injured individual(s) are offloaded. When a request arrives to the MEDEVAC system, the dispatching authority can accept or reject the request, and if it accepts the request, the dispatcher decides which of the available MEDEVAC units to dispatch. The dispatching authority considers priority level (e.g., urgent, priority, and routine) and zone of the incoming request to determine the dispatching decision. Requests that are rejected from the system are serviced by other means of evacuation such as casualty evacuation (CASEVAC) services. This thesis measures system performance based off of the true priority level of the incoming request and the response time of the MEDEVAC unit servicing the request. To explore the MDP model a notional scenario in Azerbaijan is created utilizing simulation of historical data. Sensitivity analyses on parameters of interest to MEDEVAC system leadership are investigated.

When the system receives a service request the MDP model considers the reported priority and the probability that the priority classification level was reported incorrectly. The MEDEVAC system is not informed of the true priority level until a MEDEVAC unit is servicing a request and the immediate expected reward is calculated accordingly. The decisions from the MDP model are made when a service request arrives in the system or when a MEDEVAC unit completes a service request. The entire state of the system is considered when a decision is made by the dispatching authority.

The current policy in place is a myopic policy that always dispatches the nearest available MEDEVAC unit without considering priority level classification. Results indicate that this myopic policy is not optimal. The optimal policy, which considers the entire state of the MEDEVAC system (i.e., the MEDEVAC units' status, location of incoming request, priority level of incoming request, and probability of a triage classification error), increases the expected total discounted reward (ETDR). The baseline scenario is 1.71% more efficient when the optimal policy is implemented when compared to the myopic policy. The improvement gaps between the myopic and optimal policies range from 0.59% to 8.27% via the excursions examined. In the long run, the optimal policies will substantially increase the survivability rates of battlefield casualties and should be considered by MEDEVAC system planners.

As the total misclassification rate of the MEDEVAC system increased, the difference between the myopic and optimal decreased. Additionally, the ETDR decreased but only slightly. The ETDR for MEDEVAC system when the total misclassification rate was zero was 8.43 and with the increase of the total misclassification rate to 0.8, we only saw a decrease of 0.19 to where the ETDR was 8.22. Further research should investigate the effect of the total misclassification rate on the MEDEVAC system with a larger arrival rate. Including blood transfusion kits on board MEDEVAC aircraft

1 and 4 saw an increase in ETDR from the baseline scenario. If only one aircraft were to be equipped with blood transfusion kits, we recommend MEDEVAC aircraft 4, which leads to the greatest increase in ETDR. The best case scenario is to equip both aircraft with blood transfusion kits for an ETDR of 9.92.

Replacing the current HH-60M aircraft with the Bell V-280 Valor leads to increased availability of more units and lower utilization of each MEDEVAC unit. The average utilization of MEDEVAC units is 30% for the baseline model and less than 25% when the increased average aircraft speed is implemented. Through our investigation of priority levels, the largest improvement gap of 2.38% is found when there is a larger proportion of routine priority requests. This is because of the increased flexibility that lower priority service requests introduce into the system. As the expected arrival rate increases, the stress on the system increases and the benefit gained from implementing the optimal policy over the myopic policy increases. The optimality gap for an arrival rate is 8.27% when the arrival rate is 1/30 but only 0.60% when the arrival rate is 1/120.

This research is limited by the computational constraints of the Markov decision process model. Whereas the 6-zone case for the baseline scenario has 132,055 states and is tractable, any additional MEDEVAC aircraft or zones included in the model increase the computational complexity, and the MDP model may not be able to solve to completion in a tractable amount of time. The application of approximate dynamic programming (ADP), which has been implemented in similar research (e.g., Jenkins *et al.* (2020c)), will lead to near optimal solutions and solve the issue of computational complexity. Further research should focus on applying ADP to the MEDEVAC dispatching model with triage classification errors and blood transfusion kits.

Bibliography

Bandara, Damitha, Mayorga, Maria E, & McLay, Laura A. 2012. Optimal dispatching strategies for emergency vehicles to increase patient survivability. *International Journal of Operational Research*, **15**(2), 195–214.

Berman, Oded. 1981. Repositioning of distinguishable urban service units on networks. *Computers & Operations Research*, **8**(2), 105–118.

Chaiken, Jan M, & Larson, Richard C. 1972. Methods for allocating urban emergency units: a survey. *Management Science*, **19**(4), 110–130.

Daskin, Mark S, & Stern, Edmund H. 1981. A hierarchical objective set covering model for emergency medical service vehicle deployment. *Transportation Science*, **15**(2), 137–152.

Department of the Army, United States. 2019 (July). *Army Techniques Publication 4-02.2 Medical Evacuation*.

Green, Linda, & Kolesar, Peter. 1984. A comparison of the multiple dispatch and M/M/c priority queueing models of police patrol. *Management Science*, **30**(6), 665–670.

Green, Linda V, & Kolesar, Peter J. 2004. Anniversary article: Improving emergency responsiveness with management science. *Management Science*, **50**(8), 1001–1014.

Ignall, E, Carter, G, & Rider, K. 1982. An algorithm for the initial dispatch of fire companies. *Management Science*, **28**(4), 366–378.

Jagtenberg, CJ, Bhulai, Sandjai, & van der Mei, RD. 2017. Optimal ambulance dispatching. *Markov Decision Processes in Practice*, 269–291.

Jarvis, James P. 1985. Approximating the equilibrium behavior of multi-server loss systems. *Management Science*, **31**(2), 235–239.

Jenkins, Phillip R. 2017. Using Markov decision processes with heterogeneous queueing systems to examine military medevac dispatching policies. *Theses and Dissertations*, **797**, 1–112.

Jenkins, Phillip R. 2019. Strategic location and dispatch management of assets in a military medical evacuation enterprise. *Theses and Dissertations*, **2292**, 1–155.

Jenkins, Phillip R, Robbins, Matthew J, & Lunday, Brian J. 2018. Examining military medical evacuation dispatching policies utilizing a Markov decision process model of a controlled queueing system. *Annals of Operations Research*, **271**(2), 641–678.

Jenkins, Phillip R, Lunday, Brian J, & Robbins, Matthew J. 2020a. Aerial MEDEVAC Operations. *Phalanx*, **53**(1), 63–66.

Jenkins, Phillip R., Lunday, Brian J., & Robbins, Matthew J. 2020b. Artificial Intelligence for Medical Evacuation in Great-Power Conflict. *War on the Rocks*.

Jenkins, Phillip R, Lunday, Brian J, & Robbins, Matthew J. 2020c. Robust, Multi-Objective Optimization for the Military Medical Evacuation Location-Allocation Problem. *Omega*, **97**(2020), 1–12.

Jenkins, Phillip R, Robbins, Matthew J, & Lunday, Brian J. 2021a. Approximate Dynamic Programming for Military Medical Evacuation Dispatching Policies. *INFORMS Journal on Computing*, **33**(1), 2–26.

Jenkins, Phillip R, Robbins, Matthew J, & Lunday, Brian J. 2021b. Approximate dynamic programming for the military aeromedical evacuation dispatching, preemption-rerouting, and redeployment problem. *European Journal of Operational Research*, **290**(1), 132–143.

Jenkins, Phillip R., Robbins, Matthew J., & Lunday, Brian J. 2021c. Optimising Aerial Military Medical Evacuation Dispatching Decisions via Operations Research Techniques. *BMJ Military Health*, 1–3.

Keneally, Sean K, Robbins, Matthew J, & Lunday, Brian J. 2016. A Markov decision process model for the optimal dispatch of military medical evacuation assets. *Health Care Management Science*, **19**(2), 111–129.

Kolesar, Peter, & Walker, Warren E. 1974. An algorithm for the dynamic relocation of fire companies. *Operations Research*, **22**(2), 249–274.

Kotwal, Russ S, Scott, Laura LF, Janak, Jud C, Tarpey, Bruce W, Howard, Jeffrey T, Mazuchowski, Edward L, Butler, Frank K, Shackelford, Stacy A, Gurney, Jennifer M, & Stockinger, Zsolt T. 2018. The effect of prehospital transport time, injury severity, and blood transfusion on survival of US military casualties in Iraq. *Journal of trauma and acute care surgery*, **85**(1S), S112–S121.

Malsby III, Robert F, Quesada, Jose, Powell-Dunford, Nicole, Kinoshita, Ren, Kurtz, John, Gehlen, William, Adams, Colleen, Martin, Dustin, & Shackelford, Stacy. 2013. Prehospital blood product transfusion by US Army MEDEVAC during combat operations in Afghanistan: a process improvement initiative. *Military Medicine*, **178**(7), 785–791.

McLay, Laura A, & Mayorga, Maria E. 2013. A dispatching model for server-to-customer systems that balances efficiency and equity. *Manufacturing & Service Operations Management*, **15**(2), 205–220.

Nasrollahzadeh, Amir Ali, Khademi, Amin, & Mayorga, Maria E. 2018. Real-Time Ambulance Dispatching and Relocation. *Manufacturing & Service Operations Management*, **20**(3), 467–480.

Puterman, Martin. 2005. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. 2nd edn. New Jersey: Wiley-Interscience.

Rettke, Aaron J, Robbins, Matthew J, & Lunday, Brian J. 2016. Approximate dynamic programming for the dispatch of military medical evacuation assets. *European Journal of Operational Research*, **254**(3), 824–839.

Robbins, Matthew J, Jenkins, Phillip R, Bastian, Nathaniel D, & Lunday, Brian J. 2020. Approximate dynamic programming for the aeromedical evacuation dispatching problem: Value function approximation utilizing multiple level aggregation. *Omega*, **91**, 102020.

Swersey, Arthur J. 1982. A Markovian decision model for deciding how many fire companies to dispatch. *Management Science*, **28**(4), 352–365.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704–0188

1. REPORT DATE *(DD–MM–YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From — To)*
25-03-2021	Master's Thesis	Oct 2019 — Mar 2021

4. TITLE AND SUBTITLE

The Impact of Triage Classification Errors on Military Medical Evacuation System Performance

5a. CONTRACT NUMBER

5b. GRANT NUMBER

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)

Graves, Emily S, 2d LT

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

Air Force Institute of Technology
Graduate School of Engineering and Management (AFIT/EN)
2950 Hobson Way
WPAFB OH 45433-7765

8. PERFORMING ORGANIZATION REPORT NUMBER

AFIT-ENS-MS–21-M-163

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)

Joint Artificial Intelligence Center
Rebecca E. Lee, Product Manager, JAIC
122 S. Clark Street
Crystal City, VA 22202
rebecca.e.lee.20.civ@mail.mil

10. SPONSOR/MONITOR'S ACRONYM(S)

JAIC

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

DISTRIBUTION STATEMENT A:
APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED.

13. SUPPLEMENTARY NOTES

This work is declared a work of the U.S. Government and is not subject to copyright protection in the United States.

14. ABSTRACT

In a deployed environment, evacuation requests of injured personnel are serviced by multiple forms of evacuation including medical evacuation (MEDEVAC) and casualty evacuation (CASEVAC). This thesis focuses on the optimal dispatching policy for MEDEVAC units when triage classification errors and blood transfusion kits are considered. A discounted, infinite-horizon Markov decision process (MDP) model is formulated to analyze the MEDEVAC dispatching problem and determine the optimal policy based on the status of the MEDEVAC units in the system, the priority level of incoming requests, and the locations from which requests originate. A notional, representational scenario based in Azerbaijan is utilized to compare the optimal policy against the currently practiced policy of always dispatching the nearest available MEDEVAC unit. Multiple excursions are analyzed to understand the impact of altering problem parameters, including the misclassification rate, number of aircraft equipped with blood transfusion kits, arrival rate of incoming service requests, aircraft speed, and types of triage classification errors. Results reveal that with the application of the optimal policy found by the MDP model the performance of the MEDEVAC dispatching system improves, wherein performance is measured in terms of casualty survivability. Additionally, the inclusion of blood transfusion kits on board aircraft increase MEDEVAC system performance. This analysis is of interest to the military medical planning community and may inform the development of tactics, techniques, and procedures of future dispatching policies for MEDEVAC systems.

15. SUBJECT TERMS

Markov Decision Process, MEDEVAC, Dispatching, Priority Levels

16. SECURITY CLASSIFICATION OF: a. REPORT	b. ABSTRACT	c. THIS PAGE	17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON / 19b. TELEPHONE NUMBER *(include area code)*
U	U	U	U	64	Capt Phillip R. Jenkins, PhD, AFIT/ENS; (312) 785-3636 x4727; phillip.jenkins@afit.edu

Standard Form 298 (Rev. 8–98)
Prescribed by ANSI Std. Z39.18