Health Care Manag Sci (2014) 17:77–87
DOI 10.1007/s10729-013-9236-0

Prioritization strategies for patient evacuations

Ashley Kay Childers · Maria E. Mayorga · Kevin M. Taaffe

Received: 17 May 2012 / Accepted: 7 April 2013 / Published online: 11 May 2013

Abstract Evacuation from a health care facility is considered last resort, and in the event of a complete evacuation, a standard planning assumption is that all patients will be evacuated. A literature review of the suggested prioritization strategies for evacuation planning—as well as the transportation priorities used in actual facility evacuations—shows a lack of consensus about whether critical or non-critical care patients should be transferred first. In addition, it is implied that these policies are "greedy" in that one patient group is given priority, and patients from that group are chosen to be completely evacuated before any patients are evacuated from the other group. The purpose of this paper is to present a dynamic programming model for emergency patient evacuations and show that a greedy, "all-or-nothing" policy is not always optimal as well as discuss insights of the resulting optimal prioritization strategies for unit- or floor-level evacuations.

Keywords Hospitals · Evacuation · Prioritization · Policy-making

1 Background and related literature

There is limited literature related to health care facility evacuations. Most facilities are designed with a variety of system redundancies in place to protect against the threat of evacuation. For example, in the event of a fire, patients in the most immediate danger would likely be evacuated "horizontally" through fire doors to a safe area of the facility while the fire was contained. There are, however, a variety of potential threats that could require a complete evacuation such as loss of power, flooding, exposure to hazardous materials, or a bomb threat. Such events either pose a direct risk to patients or could damage the facility's ability to provide services. When holding patients in the facility greatly reduces the quality of care, facility administrators are faced with decisions about whether to evacuate or shelter in place. Complete facility evacuations are expensive and can introduce additional risks to patients, and the questions associated with patient prioritization are highly ethical. When planning for such an evacuation, it is often assumed that all patients will be evacuated.

Evacuations necessitated by internal, facility-specific emergencies present a variety of complex problems: is an evacuation necessary? When should the evacuation begin? Which patients can be discharged early? Which patients would benefit from evacuation, and which patients should be sheltered in place? In what order should patients be chosen for evacuation? These same questions relate to external, regional threats, but during such events, there are additional questions associated with how to best utilize shared, community resources or how to handle new, incoming patients. In the event of a community-wide disaster, a health care facility is expected to be a resource to the affected population, and therefore the facility may see a sudden, increase in patient demands. There are a variety of papers that address surge capacity and triage for incoming patients (see, e.g., [1–6]), but the literature pertaining to health care facility evacuations is limited (see, e.g., [7–9]). Recently, however, Bish et al. [10] presented an integer programming model

A. K. Childers (✉) · M. E. Mayorga · K. M. Taaffe
Department of Industrial Engineering, Clemson University, 110 Freeman Hall, Clemson, SC 29634-0920, USA
e-mail: childer@clemson.edu

M. E. Mayorga
e-mail: mayorga@clemson.edu

K. M. Taaffe
e-mail: taaffe@clemson.edu

to determine patient allocation strategies for transporting patients to receiving facilities. Their "hospital evacuation transportation model" minimizes the risks (a function of the emergency threat and patient transportation) to determine the number of patients to be transported to each receiving hospital based on the availability of transportation resources. While Bish et al. [10] examine how to allocate patients for transport with the assumption that there are patients available to satisfy the resulting transportation plan, we examine the problem by assuming that the transportation resources are available to satisfy the patient prioritization plan.

A review of the suggested patient prioritization strategies for evacuation planning as well as the transport priorities used in actual health care facility evacuations shows there is a lack of consensus about whether critical care or non-critical care patients should be transferred away from a facility first in the event of a complete evacuation. In addition, it is implied that these policies are "greedy" policies in that one of these groups is chosen to be completely evacuated before any patients are evacuated from the other group. In these papers, described later in this section, patients are typically categorized as critical and non-critical care patients.

The purpose of this paper is to examine a model for unit- or floor-level patient prioritization during emergency evacuations and discuss implications for selection guidelines when patients can be classified into one of two groups: critical care and non-critical care. We consider the following factors as relevant to evacuation decisions: how quickly patients can be evacuated (inclusive of preparation and transport times); how quickly patients die while waiting to be evacuated either according to the death rates associated with normal operating conditions or because of how the actual emergency threatens patients; and the likelihood that a patient could survive transportation. We assign a reward for an evacuated patient as well as a penalty for a death.

Once evacuation orders are given, it is highly likely that the number of patients in the facility would be quickly reduced by as many early discharges as possible. We also assume that a complete health care facility evacuation would be carried out by multiple unit- or department-level evacuations occurring simultaneously and that the clinicians that typically staff that unit or department would be the ones actually carrying out the evacuation. Finally, this paper does not address how to handle incoming patients due to a regional disaster.

During an actual emergency evacuation, transportation decisions would likely depend on more than the number and classification of the patients in the system. The availability of resources and beds at receiving facilities are likely a large determinant of how patients are chosen for evacuation. Particularly during a regional disaster such as the threat of a hurricane, it is likely that evacuation prioritization decisions will be very highly dependent on the transportation resources that can be made available. Since multiple facilities often have contracts with the same transportation services, they may therefore be competing for access to ambulances, buses, and other resources during a regional disaster. Reducing the prioritization problem to a decision based on the number of patients in the system, as presented in this paper, is a simplified representation of the system in practice, but it is the first step in understanding the problem. During an evacuation, clinicians and administrators would be required to make a variety of decisions under uncertainty, continually changing conditions, and incomplete or limited information. Having understood and practiced the insights from this research during evacuation planning and tabletop simulations may enable decision makers to be more comfortable with the patient prioritization discussion.

For a more detailed discussion of the literature related to evacuation prioritization, see [11, 12]. An extensive review of the literature pertaining to evacuation guidelines and actual patient transfers shows that there is limited attention to patient prioritization. The papers listed in Table 1 represent the only papers that, to our knowledge, discuss the order in which patients should be—or actually were—chosen for evacuation. Please see [12] for a more detailed description of these papers. It is important to note that there are no clearly defined recommendations for whether critical or non-critical care patients should be given priority and as such, the policies chosen in actual facility evacuations illustrate this lack of consensus. As these policies are assumed to be greedy policies, in the case of a limited evacuation window, such decisions would likely only allow for patient transfers from that patient group without giving any opportunity for patient transfers from the other group.

Because there are no consistent prioritization policies, we aim to address the following questions in this paper: is the optimal policy in fact greedy? When should a single patient group be given priority? What is the effect of choosing a non-optimal policy? While there are no papers that address these problems directly, there are a variety of papers that clearly indicate that methods for ethically prioritizing patients for evacuation would be a beneficial contribution:

– "Familiarity with and utilization of a framework for ethical decision-making may facilitate health care professionals in maneuvering through disaster-instigated ethical dilemmas" from "Disaster Ethics, Healthcare and Nursing: A Model Case Study to Facilitate the Decision Making Process" in the *Online Journal of Health Ethics* [24].
– "The goals of triage in different environments and contexts can lead to divergent perspectives of what

Table 1 Summary of prioritization strategies in the literature

	Critical care patients first	Non-critical care patients first
Evacuation plans and guidelines	Center for bioterrorism preparedness and planning [13] Gray and Hebert [15] - summary of Hurricane Katrina response Protocol at Lindy Boggs Meddical Center [15] Memorial Medical Center after Katrina [18]	AORNs guidance statement [14] New York center for terrorism preparedness and Planning [16] Johnson simulation of evacuation risks [17] Moskop and Iserson discussion of triage [19] Lach et al. geriatric patient evacuation [20] Texas Medicals policy until 2005 [21]
Actual Facility Evacuations	Facilities not in immediate danger after the Northridge, California earthquake [9] TexasMedical Center after Tropical Storm Allison [22] Charity Hospital after Hurricane Katrina [15] Texas Medical during Hurricane Katrina [21]	Facilities in immediate danger after the Northridge, California earthquake [9] Lindy Boggs Medical Center after Katrina [15] Memorial Medical Center after Katrina [23]

constitutes ethically sound decision making," and "a system that allows real-time classification of risks and benefits...would be a great advantage in ethical decision making" from "Lifeboat Ethics: Considerations in the Discharge of Inpatients for the Creation of Hospital Surge Capacity" in Disaster Medicine and Public Health Preparedness [25].

– "Advance agreement is needed among key parties about which patients will be evacuated first. Several disputes developed over priorities in the days after Katrina. There was disagreement, for example, over whether the sickest patients or those more likely to survive should be evacuated first" from "Hospitals in Hurricane Katrina: Challenges Facing Custodial Institutions in a Disaster" by The Urban Institute [15].
– "Medical and social needs must be considered in trigging evacuees. The traditional medical model for triage in the U.S. is to treat the most critically injured first; in an overwhelming disaster situations, health care providers may shift to battlefield triage practices in which those with the highest probability of survival are treated first. Little is known about lay clinicians' abilities to shift paradigms during response" from AHRQ's "Recommendations for a National Mass Patient Evacuee Movement, Regulating, and Tracking System" [26].
– "The ethical decisions inherent in triage decisions should not be first considered during a real event. Rather, they should be rehearsed and discussed long before they are needed" from "Terrorism and the Ethics of Emergency Medical Care" in the *Annals of Emergency Medicine* [27].

The organization of this paper is as follows. We first present a dynamic programming model that we use to determine the optimal selection strategy for two patient groups. The behavior and properties of this model are explored, and then policy determination and evaluation are discussed. The paper concludes with a discussion of one possible model extension and a summary of the research conclusions and future work.

2 Single server evacuation model

During an actual emergency evacuation event, transportation decisions would likely depend on a variety of factors such as the availability of transportation resources, the availably of beds at receiving facilities, or whether the emergency is regional or facility-specific. For now, we reduce the prioritization problem to a decision based on the number of patients in the system to be evacuated. In this section, a single team evacuation model is introduced as a tool for patient selection decisions when patients are split into two patient classifications.

2.1 Model description

In the event of a planned evacuation (some advanced notice of the emergency is given), a facility would likely discharge as many patients as possible to reduce the number of patient transfers. Of those patients that remain under the care of the staff, patients will be removed from the system if they 1) are successfully transported and transferred to another facility, 2) die during the evacuation transportation process, or 3) die while they are waiting to be selected for evacuation. It is assumed that the input parameters are stationary. In reality, the rates at which patients can be evacuated, the rates at which patients die while waiting to be evacuated, and the probability of a successful evacuation are likely to change over time as the evacuation window decreases. We use the following parameters to model the system:

- i patient type, $i = 1, 2$.
- N_i the number of type i patients at the beginning of the time horizon to be evacuated,
- $x_i(t)$ the number of type i patients remaining in the system at time t,
- λ_i the rate at which type i patients can be evacuated,
- p_i the probability that a type i patient will survive transport,
- α_i the rate at which type i patients die while waiting to be evacuated,
- l_i^e the reward associated with the completed evacuation of a type i patient, and
- l_i^d the penalty incurred when a type i patient dies (either waiting or during transport).

Let policy π represent the sequence of decisions, $\pi = a_1, a_2, ...$, that represents the choice for how the evacuation team should be allocated based on the current state where a_k is the vector of actions taken at each decision epoch. Let Π be the set of all such policies where $\left(x_1^k, x_2^k\right)$ represents the state of the system at time n_k, and a_k denotes the actions taken at time k under policy π. The resulting n-stage expected reward, with initial state, $\left(x_1^0, x_2^0\right) = (x_1, x_2)$, is given by

$$v_n(x_1, x_2) = E^{\pi}_{(x_1^k, x_2^k)}\left[\sum_{k=0}^{n-1} R\left(x_1^k, x_2^k, a_k\left(x_1^k, x_2^k\right)\right)\right], \quad (1)$$

where $R()$ denotes the reward achieved at stage or time k. This leads to the optimal n-stage expected reward of

$$v_n^*(x_1, x_2) = \sup_{\pi \in \Pi} v_n^{\pi}(x_1, x_2). \quad (2)$$

During an evacuation, the teams would choose from which patient class to evacuate the next patient. This leads to the following decision at any epoch:

$$a_k = \begin{cases} (\lambda_1, 0) & \text{evacuate type 1 next - Policy 1} \\ (0, \lambda_2) & \text{evacuate type 2 next - Policy 2} \end{cases}. \quad (3)$$

In order to convert this continuous time problem to a discrete time problem, the transition rates are uniformized by scaling the maximum rate of transition, γ. In this case, the maximum rate of transition is $\gamma = \lambda_1 + \lambda_2 + N_1\alpha_1 + N_2\alpha_2$. Therefore, the optimality equation for the single server finite horizon evacuation decision problem is then shown below in Eq. 4, where $\nu(x_1, x_2)$ represents the long run average reward.

$$\begin{aligned} \nu(x_1, x_2) &= x_1\alpha_1\left[\nu(x_1-1, x_2) - l_1^d\right] + x_2\alpha_2\left[\nu(x_1, x_2-1) - l_2^d\right] \\ &+ \max\begin{cases} \lambda_1 p_1\left[\nu(x_1-1, x_2) + l_1^e\right] + \lambda_1(1-p_1)\left[\nu(x_1-1, x_2) - l_1^d\right] \\ + \left[(N_1 - x_1)\alpha_1 + (N_2 - x_2)\alpha_2 + \lambda_2\right]\nu(x_1, x_2), \\ \lambda_2 p_2\left[\nu(x_1, x_2-1) + l_2^e\right] + \lambda_2(1-p_2)\left[\nu(x_1, x_2-1) - l_2^d\right] \\ + \left[(N_1 - x_1)\alpha_1 + (N_2 - x_2)\alpha_2 + \lambda_1\right]\nu(x_1, x_2) \end{cases} \end{aligned} \quad (4)$$

The first two terms in Eq. 4 represent patient deaths while waiting for evacuation. Based on the number of patients remaining in the system, a patient death—unrelated to the evacuation decision but dependent on the normal operations of the facility or the emergency type—may occur. When a patient death occurs, the number of patients waiting in the system is reduced and a penalty is assigned. The final term represents the choice: either evacuate a Type 1 patient or a Type 2 patient. With both choices, there is a chance of an unsuccessful evacuation, so the number of patients in the system will be reduced, and either a reward or penalty will be assigned based on the probability of a successful evacuation. In addition, each decision is associated with a fictitious transition rate as a result of uniformization.

Assume that patients are categorized as critical or non-critical care patients (let Type 1 represent critical care patients and Type 2 represent non-critical care patients). We assume

- non-critical patients can be evacuated more quickly than critical patients ($\lambda_1 < \lambda_2$),
- non-critical patients have a higher probability of successful evacuation ($p_1 < p_2$), and
- critical patients die while waiting at a quicker rate than non-critical patients ($\alpha_1 > \alpha_2$).

For the discussion in this paper, it is assumed that the reward associated with a completed, successful patient evacuation is the same for both patient classes, and in addition, the value is equal in magnitude to the cost of a lost life from either of the patient classes (either during transportation or while waiting on evacuation). That is, $l_1^e = l_2^e = l_1^d = l_2^d = 1$. This assumption follows the utilitarian logic: every patient counts equally towards the greater good such that no patient is considered more valuable than another [28, 29].

2.2 Dynamic programming results

The dynamic program was solved using value iteration and returns the optimal policy for a given set of parameters. Therefore, the optimal action (which patient type should be evacuated as a function of the number of each type of patient remaining in the system) is returned, and this set of policies can be classified as one of three types:

- a greedy policy that chooses to evacuate all critical patients (Type 1) first,
- a greedy policy that chooses to evacuate all non-critical patients (Type 2) first, or
- a "switching policy" that, at the beginning of the evacuation window, gives priority first to non-critical care patients (Type 2) but switches priority to critical care patients (Type 1) while there are still non-critical care patients waiting to be evacuated (see Fig. 1).

Policy diagrams are used to represent which patient type—Type 1 or Type 2—should be chosen at any combination of (x_1, x_2) patients remaining in the system. The example shown in Fig. 1 represents a switch from non-critical care priority (when the number of patients in both patient groups is high) to critical care priority (when the number of patients in both groups is low). The initial state space is represented in the upper right corner of the figure; for this case, there are 40 patients in the facility awaiting evacuation. When there are 20 critical care patients and 20 non-critical care patients waiting to be evacuated ($x_1 = 20$ and $x_2 = 20$), the optimal action is to evacuate a Type 2 (non-critical) patient. The only actions that affect the number of patients remaining, and therefore position on the figure, are deaths and evacuations. Above the policy switching curve (shown as the dashed line between Type 2 to Type 1 in Fig. 1), evacuations decrease the number of non-critical care patients remaining (move down in the figure),

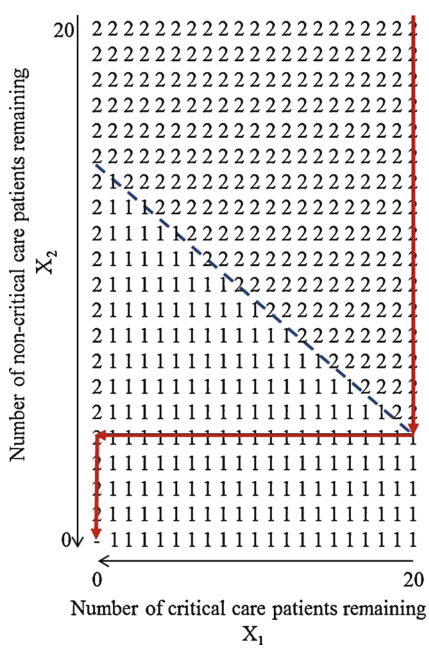

Fig. 1 Sample policy diagram (path depicting a sample set of optimal actions)

and deaths can either move the decisions down (when a non-critical care patient dies) or to the left (when a critical care patient dies). Below the policy switching curve, evacuations decrease the number of critical care patients remaining (move left in the figure), and deaths can move the decisions as previously described.

Let's assume that there are 20 critical care patients and 20 non-critical care patients in the system to be evacuated. At this point, the optimal decision is to evacuate a Type 2 or non-critical care patient (represented by the top, right point on the Fig. 1). The actual policy realization (or "path" through the diagram) will depend on patient deaths, which are random events. However, as an example, assume that no patients die while waiting to be evacuated. The optimal policy decisions are represented by the continuous arrows on the diagram. Type 2 patients will be evacuated until there are only 4 Type 2 patients remaining. At that time, the optimal policy is to switch priority to Type 1 patients. After all Type 1 patients are evacuated, the final four remaining Type 2 patients will be evacuated.

3 Characterization of the optimal policy

Analytical characterization of the optimal policy was not possible through functional structures that facilitate proving that switching curve is optimal (e.g., super-modularity), but we were able to show concavity and super-modularity computationally. Based on thousands of trials, we observed that the optimal policy has at most one switching curve and that if a switching curve is present the optimal policy is to begin by assigning priority to non-critical patients and then switch to critical care patients, and there is at most one switch from the time the queue of patients is full until the end of the evacuation period. To implement a switching policy, the evacuation team or teams select non-critical patients for evacuation first. At some point, depending on the patient classification rates, a switch should be made so all remaining critical care patients are selected for evacuation. Once there are no critical care patients remaining, any remaining non-critical care patients should be evacuated.

In order to investigate how the input parameters affect the optimal policy, a sensitivity analysis was performed to examine which policy was optimal: a greedy policy in favor of critical care patients, a greedy policy in favor of non-critical care patients, or a switching policy. In the test cases where a switching policy was optimal, we were interested in the "location" of the switching curve that therefore defines the size of the prioritization region for each patient group. A full factorial design of experiments (DOE) was created to study the effects of the input parameters. In order to decrease the number of runs, the ratios of the following input parameters were selected as the

Table 2 DOE parameters

	Type 2	Type 1		
		High	Medium	Low
λ_i	4	3.5	3	2
p_i	1.0	0.95	0.75	0.5
α_i	0.001	0.075	0.050	0.025

Table 3 Model inputs for batch testing

λ_1	λ_2	α_1	α_2	p_1	p_2	r
1.5	2.5	0.05	0.01	0.5	0.85	1.76
2	3	0.06	0.02	0.75	0.9	1.67
2.5	3.5	0.07	0.03	0.9	0.95	1.58
3	4	0.08	0.04	1	1	1.5

adjustable factors: λ_1/λ_2, p_1/p_2, and α_1/α_2. The parameters for the non-critical care patients were held constant while the parameters for the critical care patients were varied at a low, medium, and high setting. Varying these six factors at three levels each resulted in 27 possible scenarios for the sensitivity analysis. The values chosen to remain constant for the non-critical care patients and the values chosen to be varied for the critical care patients are shown in Table 2. In the context of critical and non-critical care patients, a downward shift in the location of the switching curve results in a greedy policy in favor of the non-critical care patients. A shift upwards results in a greedy policy in favor of the critical care patients. This implies that a switching policy should begin with the evacuation team—or teams—selecting non-critical care patients first for evacuation. At some point, depending on the patient classification rates, a switch should be made so that all remaining critical care patients are selected for evacuation. Once all critical care patients have been removed from the system, all remaining non-critical care patients should be transported away from the facility.

The results show that increasing the value for any one of the critical care patient parameters increases the region of critical care patient priorities until the policy becomes a greedy policy in favor of the critical care patients. This result is intuitive: if the rate at which critical care patients can be evacuated increases so that patients are moved out of the system more quickly, eventually critical care patients should get priority. The same applies for the probability with which they can be successfully evacuated to another facility. Similarly, as the rate at which critical care patients die while waiting to be evacuated increases, it seems obvious that these patients should be the first to be taken out of harm's way.

4 Policy determination

In the event that a patient group can be evacuated more quickly and successfully as well as die more quickly while waiting to be evacuated, it is obvious that this patient group should be given priority. Such a scenario could occur based on the patients' relative location to the hazard. It is more likely, however, that patients will be categorized as critical or non-critical care patients. This assumes that non-critical care patients can be evacuated more quickly, that non-critical care patients have a higher probability of a successful evacuation, and that critical care patients die while waiting to be evacuated more quickly than non-critical care patients $\alpha_1 > \alpha_2$. In these cases, it was determined that a greedy policy may be predicted by examining a ratio, r, consisting of the Type 1 and Type 2 parameter values. Similar to the $c\mu$ rule [30, 31], which gives priority to jobs with the largest value of $c_i\mu_i$ (where c_i is the reward if a job is completed and μ_i is the service rate for jobs $i = 1, 2, ..., I$), the optimal policy for the evacuation systems can be predicted by considering the following relationship.

$$r = \frac{\lambda_1 p_1 \alpha_1}{\lambda_2 p_2 \alpha_2} \tag{5}$$

In order to further examine this ratio, r, we conducted a full DOE for 6 different initial evacuation populations with the parameters shown in Table 3. This resulted in 4,096 unique tests at each of the initial evacuation populations.

There is a value, $U_{(x1,x2)}$, of this ratio—increasing with an increasing initial population size— beyond which the optimal policy is always a Type 1 policy. Though we have not yet been able to mathematically determine this bound, Table 4 shows the values based on the DOE.

It makes sense that the patient group i with the higher value of $\lambda_i p_i \alpha_i$ will more quickly contribute to the reward or incur a penalty; therefore, these patients should be evacuated first. It should be noted, however, that none of these ratios predicts that a switching policy will be optimal. Based on extensive testing, if the ratio is between 1 and some upper bound, $U_{(x1,x2)}$, the optimal policy could be either one of the three possible policies, and the model should be run to

Table 4 Ratio value to predict type 1 policies in evacuations

Initial state	$U_{(x1,x2)}$
(5,5)	1.127
(10,10)	1.339
(20,20)	1.587
(40,40)	2.116
(60,60)	2.679
(100,100)	3.214

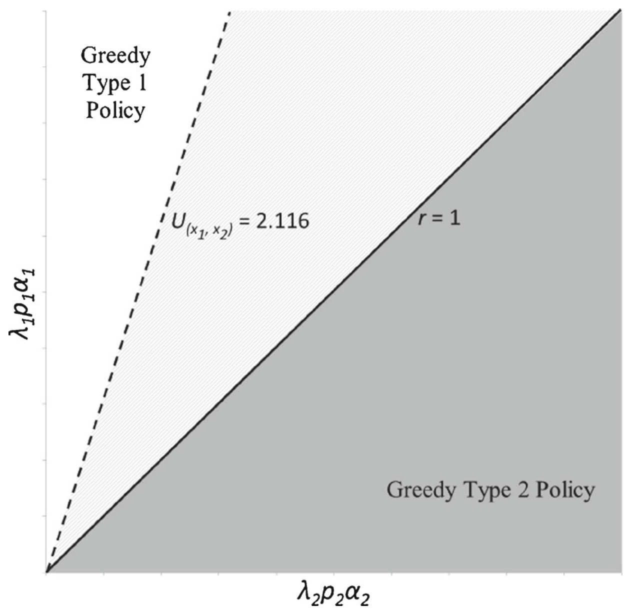

Fig. 2 Sample of ratio values for predicting the optimal policy

determine which is optimal. In summary, given a particular state space i, if

– $r < 1$, a Greedy Type 2 policy is the optimal policy,
– $r > 1$ and $r < U_{(x1,x2)}$, a switching policy may occur, or
– $r > U_{(x1,x2)}$, a Greedy Type 1 policy is the optimal policy.

Figure 2 illustrates an example of the r and $U_{(x1,x2)}$ values for a sample case of 40 of each patient type in the system waiting for evacuation. The dark grey and white regions demonstrate cases where the optimal policy is known by using these ratios. For smaller initial patient populations, the value of $U_{(x1,x2)}$ is smaller, and therefore the region where the optimal policy can be determined using these values increases. As the initial patient population to be evacuated increases, so does the number of decisions to be made, and therefore the percentage of cases for which r can be used to determine the optimal policy decreases.

We noted before that any switching policy begins with the evacuation of non-critical care patients (Type 2) and then switches to the evacuation of critical care patients (Type 1) as the evacuation continues and the number of patients in the system decreases. Figure 3 shows the frequency of each of the three possible policy types, and with larger initial patient populations, there is more opportunity for a switching policy to be optimal. Since switching policies switch only from non-critical to critical patients, if the optimal policy is to evacuate the non-critical care patient when there is only one of each patient type in the system—the optimal policy is 2 at (1,1)—then the optimal policy will always be to give priority to the non-critical care patients. This explains why the percentage of trials in which the optimal policy is a Greedy Type 2 policy remains constant in Fig. 3. Based on the input parameters we chose for the 4,096 trials, 26.2% resulted in a Greedy Type 2 policy. These results lead us to our next discussion.

5 Policy evaluation

Now that we have shown that a greedy policy is not always optimal, it is important that we consider the benefits of choosing the best policy with respect to the losses associated with choosing any other policy. Table 5 examines the effects of choosing either greedy policy for all evacuations by examining the difference in the optimal policy value and value resulting from the chosen policy. As shown in Table 1, published strategies for evacuation are greedy policies. In

Fig. 3 Classification of the optimal policy for 4,096 trials

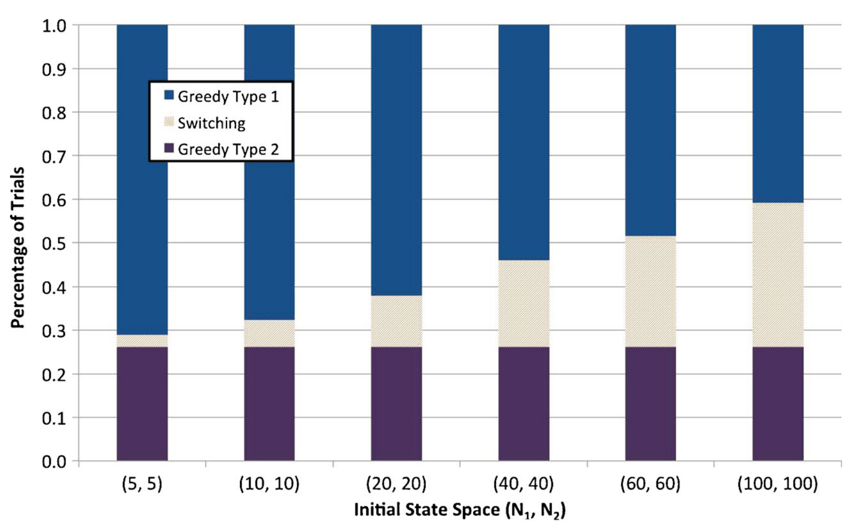

addition, it may be difficult for clinicians to implement the optimal switching policy; thus we compare the optimal policy to the performance of always choosing the "wrong" policy. Based on the parameters chosen for the trials, when the initial patient population is low, a Greedy Type 1 policy is most often optimal. With an increasing state space, however, the probability that the optimal policy is a switching policy increases.

The tables below show the average objective function loss or objective function solution gap (the value sacrificed by using a suboptimal policy), coefficient of variation, and the average objective function loss per patient. The average objective function loss is calculated by the absolute value of the difference in the value functions of the optimal and evaluated policies at the start of the evacuation. The average objective function loss per patient is the average objective function loss divided by the total number of patients in the system to be evacuated (the initial state space). Since the units are in number of lives, the objective function loss can be interpreted as the number of lives lost by not selecting the optimal policy. The loss per patient can be thought of as the increased chance of death per patient as compared to the optimal policy.

Based on the values chosen for our tests, a Greedy Type 1 policy is most often the optimal policy, and the average objective function loss and the average lost value per patient per patient is less than when a Greedy Type 2 policy is always chosen. However, note that as the state space increases, the lost value per patient increases for a Greedy Type 1 policy and remains relatively stable for a Greedy Type 2 policy. In addition, the coefficient of variation is greater for each initial patient population when a Greedy Type 1 policy is always chosen. Table 6 examines only the subset of tests in which the optimal policy is a switching policy. Since optimal switching policies would be difficult for health care facilities to determine and carry out, we examine the effects of always choosing either greedy policy when the optimal policy is actually a switching policy.

In these cases, when the optimal policy is a switching policy, the average objective function loss is less when a Greedy Type 2 policy is chosen, and in this case, the coefficient of variation is greater for each initial patient population when a Greedy Type 2 policy is always chosen. However, we have determined that we can narrow the need for mathematically determining the optimal policy by considering the ratio r. The optimal policy is only unknown when $1 < r < U_{(x1,x2)}$. Within this range, the optimal policy may be any one of the three possibilities. When r is closer to 1, if the optimal policy is not a switching policy, it will be a Greedy Type 2 policy, and when r is closer to $U_{(x1,x2)}$, if the optimal policy is not to switch, it will be to choose all Type 1 patients first. This ratio, that we will call r_s, reflects the highest value within this unknown range at which the optimal policy could be a Greedy Type 2 policy (otherwise it is a switching policy). Beyond r_s, if the optimal policy is not a switching policy, it is a Greedy Type 1 policy. The exact r_s where these switches occur has not been proven mathematically but was determined to be 1.0125 for all initial patient populations for the combination of 4,096 trials we tested. In summary, when

- $r < 1$, the optimal policy is a Greedy Type 2 policy,
- $1 < r < 1.0215$, the optimal policy is either to switch or to choose a Greedy Type 2 policy,
- $1.0125 \leq r \leq U_{(x1,x2)}$, the optimal policy is either to switch or to choose a Greedy Type 1 policy, or
- $r > U_{(x1,x2)}$, the optimal policy is a Greedy Type 1 policy.

6 Extension: two evacuation teams

There will likely be multiple teams available to evacuate patients during a unit-, floor-, or department-level evacuations. In this section, a model for assigning two servers–or evacuation teams–is examined with dynamic programming.

Table 5 The effects of always choosing a greedy policy

	Initial Patient Population					
	(5, 5)	(10, 10)	(20, 20)	(40, 40)	(60, 60)	(100, 100)
Always choose a greedy Type 1 policy						
Obj function solution gap	0.06	0.22	0.78	2.48	4.44	8.90
CV	2.17	2.08	1.89	1.69	1.56	1.41
Lost value/patient	0.01	0.01	0.02	0.03	0.04	0.05
Always choose a greedy Type 2 policy						
Obj function solution gap	0.25	0.78	2.16	4.88	7.32	10.40
CV	0.99	1.03	1.10	1.19	1.26	1.35
Value/patient	0.03	0.04	0.05	0.06	0.06	0.05

Table 6 The effects of always choosing a greedy policy when it is optimal to switch

	Initial Patient Population					
	(5, 5)	(10, 10)	(20, 20)	(40, 40)	(60, 60)	(100, 100)
Always choose a greedy Type 1 policy						
Obj Function Solution Gap	0.01	0.07	0.35	1.34	2.67	5.69
CV	0.91	0.99	0.97	0.94	0.92	0.89
Lost Value/Patient	0.001	0.003	0.009	0.017	0.022	0.028
Always choose a greedy Type 2 policy						
Obj Function solution gap	0.003	0.018	0.093	0.41	0.83	1.88
CV	1.33	1.33	1.39	1.30	1.25	1.21
Lost value/patient	0.003	0.009	0.002	0.005	0.006	0.009

Assume that that there are two evacuation teams available to move patients, and that the two teams can be allocated to patient evacuations according to any one of the following three policies: either both teams can evacuate Type 1 patients, the teams can be split so that one team is dedicated to moving Type 1 and one team is dedicated to moving Type 2 patients, or both teams can evacuate Type 2 patients. This leads to the following decision at any epoch:

$$\pi = \begin{cases} (2\lambda_1, 0) & \text{both teams evacuate Type 1 patients - Policy 1} \\ (\lambda_1, \lambda_2) & \text{teams split between both patient types - Policy 2} \\ (0, 2\lambda_2) & \text{both teams evacuate Type 2 patients - Policy 3} \end{cases}, \tag{6}$$

The state description as well as the n-stage expected reward remains the same as in the previously discussed models. The uniformization rate, based on the maximum rate of transition, used for this model is $\gamma = 2\lambda_1 + 2\lambda_2 + N_1\alpha_1 + N_2\alpha_2$. The fictitious transition rate in this case is $(N_1 - X_1)\alpha_1 + (N_2 - X_2)\alpha_2 + 2\lambda_2$ when Policy 1 is chosen, $(N_1 - X_1)\alpha_1 + (N_2 - X_2)\alpha_2 + \lambda_1 + \lambda_2$ when Policy 2 is chosen, and $(N_1 - X_1)\alpha_1 + (N_2 - X_2)\alpha_2 + 2\lambda_1$ when Policy 3 is chosen. The optimality equation used to determine how the two evacuation teams should be allocated and therefore prioritize the patients for evacuation is shown in Eq. 7 below.

$$\begin{aligned} v(X_1, X_2) &= X_1\alpha_1\left[v(X_1-1, X_2) - l_1^d\right] + X_2\alpha_2\left[v(X_1, X_2-1) - l_2^d\right] \\ &+ \max\begin{cases} 2\lambda_1 p_1\left[v(X_1-1, X_2)+l_1^e\right] + 2\lambda_1(1-p_1)\left[v(X_1-1, X_2)-l_1^d\right] \\ +\left[(N_1-X_1)\alpha_1+(N_2-X_2)\alpha_2+2\lambda_2\right]v(X_1, X_2), \\ \lambda_1 p_1\left[v(X_1-1, X_2)+l_1^e\right] + \lambda_1(1-p_1)\left[v(X_1-1, X_2)-l_1^d\right] \\ +\lambda_2 p_2\left[v(X_1, X_2-1)+l_2^e\right] + \lambda_2(1-p_2)\left[v(X_1, X_2-1)-l_2^d\right] \\ +\left[(N_1-X_1)\alpha_1+(N_2-X_2)\alpha_2+\lambda_1+\lambda_2\right]v(X_1, X_2), \\ 2\lambda_2 p_2\left[v(X_1, X_2-1)+l_2^e\right] + 2\lambda_2(1-p_2)\left[v(X_1, X_2-1)-l_2^d\right] \\ +\left[(N_1-X_1)\alpha_1+(N_2-X_2)\alpha_2+2\lambda_1\right]v(X_1, X_2) \end{cases} \end{aligned} \tag{7}$$

As before, the first two terms of the optimality equation represent patient deaths while waiting for evacuation, and the final term represents the choice between the three allocation options, and each includes a fictitious transition rate.

A number of tests showed that the optimal policy for this model was either to assign both evacuation teams to work together on a critical care patient (Policy 1) or to assign both evacuation teams to work together on the non-critical care patients (Policy 2). It was never optimal to split the teams between the two patient groups (Policy 2). Though the structural properties of the optimality equations are difficult to prove, it is relatively easy to show that the evacuation teams should be allocated to the same patient group (see [32]). In summary, when there are two evacuation teams available to move patients, the optimal policy is either to assign both teams to transport Type 1 patients or to assign both teams to evacuate Type 2 patients; it is never optimal to split the evacuation teams between the two patient groups. However, a switching policy may still be the optimal choice, and in the case of critical and non-critical care patients, a switching policy should begin with non-critical care patient evacuations followed by a switch to critical care patient evacuations.

7 Conclusions

The purpose of this research is to provide insights into the problem of patient prioritization during complete evacuations from health care facilities. To date, there is limited research related to this problem, most likely due to the fact that most facilities have system redundancies in place as well as the highly ethical nature of the prioritization discussion. In the few cases where patient prioritization strategies are suggested or explained in the literature, there is a lack of consensus about (1) which patients *should* be selected first and (2) which patients *were* selected first during actual emergency evacuations. These policies are all greedy; either critical care or non-critical care patients are given priority. We have shown, however, that a greedy policy is not always optimal. The optimal policies from the evacuation dynamic programming model presented in this paper can be characterized as one of three different policy types: either a greedy, Type 1 policy; a switching policy; or a greedy, Type 2 policy.

Based on certain conditions, the optimal policy can be determined without the use of the dynamic programming model. First, certain emergency types could result in a situation where one patient group can be evacuated more quickly and more successfully as well as die more quickly while waiting to be evacuated. Therefore it is obvious that this patient group should be given priority. Next, if the optimal policy is to choose a non-critical care patient when there are only two patients in the system (one critical and one non-critical), then the optimal policy will always be to give priority to non-critical care patients, no matter how large the initial evacuation population is. This is because there is at most one switch in a switching policy, and any switching policy begins by giving priority to non-critical care patients and then switching to evacuate all remaining critical care patients from the facility. Finally, if the ratio $r = \lambda_1 p_1 \alpha_1 / \lambda_2 p_2 \alpha_2 < 1$, priority should always be given to non-critical care patients. There is an upper value on this ratio, $U_{(x1,x2)}$, that is increasing with an increasing initial evacuation population, for which the optimal policy for any set of parameters with a ratio above $U_{(x1,x2)}$ is to give priority to critical care patients. Therefore, it is only when $1 < r < U_{(x1,x2)}$ that the optimal policy is unknown and may be any one of the three possible policies. Within this window, if $r < 1.0125$, the optimal policy is either to choose a switching policy or a Greedy Type 2 policy. If $r \geq 1.0125$, the optimal policy is either to choose a switching policy or a Greedy Type 1 policy.

By considering the ratio of the patient classification rates, the ratio r can be used to determine whether priority should be given to critical care patients, whether priority should be given to non-critical care patients, or whether the optimal policy is unknown. This provides some insight on patient prioritization that was previously unknown as demonstrated by the lack of consensus on patient prioritization strategies. These and other insights from this paper could be tested and discussed in table-top simulation exercises. For example, in the event that there are 2 evacuation teams, it is never optimal to split two evacuation teams between the two patient groups so that one team evacuates Type 1 patients and one team evacuates Type 2 patients. Instead, the optimal policy is either to assign both teams to evacuate critical care patients or both teams to evacuate non-critical care patients. While the optimal policy may still be a switching policy, both evacuation teams should focus the same patient group.

There are certainly opportunities for a variety of modeling extensions. Though not discussed in this paper, we consider the costs associated with choosing to keep a patient in the facility rather than evacuate (holding costs). Such costs could include the monetary costs associated with caring for a particular patient or the opportunity costs associated with selecting the "wrong" patient for evacuation. While it is not as easy to predict the structure of the optimal policy when holding costs are included, any optimal policy still has at most one switch when these additional costs are included. The models are also most sensitive to holding costs and can even affect the model such that choosing a particular patient is "not worth it." In addition, the current model looks only at the prioritization decision; a more realistic model would consider the availability of transportation resources as well as the availability of beds at receiving facilities. It may be beneficial to extend the number of patient classifications. Along those same lines, the rewards for the patient classifications could be expanded rather than solely than counting the number of saved lives. For example, the number of life years after the evacuation, quality-adjusted life years, or some other measure would allow for further prioritization within the groups. In addition, the effects of incorrectly classifying a patient as a critical or non-critical care patient—or incorrectly assigning the input parameters—need consideration. Regardless of how the models are altered, a better understanding of the actual values for the input parameters would improve the modeling efforts. While the rates of evacuation chosen in the previous discussion were based on observations at mock evacuations, there are no data available for estimating the holding costs. Knowing the input values would allow a better analysis of the location and movement of the switching curve. We have also developed a simulation model and in future work hope to examine the effect of evacuation policies (number of lives saved and lost, time to complete an evacuation, frequency that all patients can be evacuated in a given time window). Once the dynamic program is used to determine the optimal policy, we can code this policy into the simulation model. Switching policies, however, cannot yet be coded quickly and tested on a large scale as we can do with the dynamic programming trials.

Finally, there needs to be continued discussion among health care workers about the ethical dilemmas associated with making evacuation decisions as well as other scarce resource allocation decisions. The insights from this research should be used to encourage such a discussion and could be used in tabletop simulation exercises for evacuation planning.

References

1. Argon NT, Ziya S, Righter R (2008) Scheduling impatient jobs in a clearing system with insights on patient triage in mass casualty incidents. Probab Eng Inform Sc 22(3):301–332
2. Cookson R, Dolan P (2000) Principles of justice in health care rationing. J Med Ethics 26(5):323–329
3. Hick JL, Barbera JA, Kelen GD (2009) Refining surge capacity: conventional, contingency, and crisis capacity. Disaster Med Public Health Prep 3(2 Suppl):S59–S67
4. Jacobson EU, Argon NT, Ziya S (2012) Priority assignment in emergency response. Oper Res 60(4):813–832
5. Schultz CH, Stratton S (2007) Improving hospital surge capacity: a new concept for emergency credentialing of volunteers. Ann Emerg Med 49(5):602–609
6. Sinuff T, Kahnamoui K, Cook DJ, Luce JM, Levy MM (2004) Rationing critical care beds: a systematic review. Crit Care Med 32(7):1588–1597
7. Bovender J, Carey B (2006) A week we don't want to forget: Lessons learned from Tulane. Front Health Serv Manage 23(1): 3–12
8. McGlown KJ (1999) Determinants of the evacuation of health care facilities. PhD thesis. University of Alabama at Birmingham
9. Schultz CH, Koenig KL, Lewis RJ (2003) Implications of hospital evacuation after the Northridge, California, earthquake. N Engl J Med 348(14):1349–1355
10. Bish DR, Acga E, Glick R (2011) Decision support for hospital evacuation and emergency response. Ann Oper Res. doi:10.1007/s10479-011-0943-y
11. Childers AK, Visagamurthy G, Taaffe KM (2009) Prioritizing patients for evacuation from a healthcare facility. TRR: J TRB 2137:38–45
12. Childers A K, Taaffe K M (2010) Healthcare facility evacuations. Lessons learned, research activity, and the need for engineering contributions. J Healthc Eng 1(1):125–140
13. Bioterrorism Hospital Preparedness Program (2006) Hospital evacuation protocol. Technical report. New York City department of health and mental hygiene
14. Guidance Statement AORN (2007) Mass casualty, triage, and evacuation. AORN J 85(4):792–800
15. Gray BH, Hebert K (2007) Hospitals in Hurricane Katrina: challenges facing custodial institutions in a disaster. J Health Care Poor U 18(2):283–298
16. New York Centers for Terrorism Preparedness and Planning (2006) Hospital evacuation protocol. Technical report. Bioterrorism hospital preparedness program
17. Johnson C (2006) Using computer simulations to support a risk-based approach for hospital evacuation. Technical report. Glasgow accident analysis group. University of Glasgow
18. Hodge JG, Garcia AM, Anderson ED, Kaufman T (2009) Emergency legal preparedness for hospitals and health care personnel. Disaster Med Public Health Prep 3(2 Suppl):S37–S44
19. Moskop JC, Iserson KV (2007) Triage in medicine, Part II. Underlying values and principles. Ann Emerg Med 49(3):282–287
20. Lach HW, Langan JC, James DC (2005) Disaster planning: are gerontological nurses prepared? J Gerontol Nurs 31(11): 21–27
21. Sexton K, Alperin L, Stobo J (2007) Lessons from Hurricane Rita: The University of Texas Medical Branch Hospital's evacuation. Acad Med 82(8):792–796
22. Bowers P, Maguire M, Silva P, Kitchen R (2006). Everybody out!!, Nursing 36(5 Suppl E D):12–13
23. Okie S (2008) Dr. Pou the hurricane: implications for patient care during disaster. N Engl J Med 358(1):1–5
24. Powell-Young Y, Baker J, Jogan J (2006) Disaster ethics, health care and nursing: a model case study to facilitate the decision making process. Online J Health Ethics 3(2)
25. Kraus CK, Levy F, Kelen GD (2007) Lifeboat ethics: considerations in the discharge of inpatients for the creation of hospital surge capacity. Disaster Med Public Health Prep 1(1):51–56
26. Agency for Healthcare Research and Quality (2009) Recommendations for a national mass patient and evacuee movement, regulating, and tracking system. Technical report. AHRQ
27. Pesik N, Keim ME, Iserson KV (2001) Terrorism and the ethics of emergency medical care. Ann Emerg Med 37(6):642–646
28. Goodin RE (1995) Utilitarianism as a public policy. Cambridge University Press, New York, pp 3–23
29. Rachels J, Rachels S (2007) The elements of moral philosophy, 5th edn. McGraw-Hill, New York, pp 100–116
30. Iravani SMR, Kolfal B (2004) When does the $c\mu$ rule apply to finite-population queueing system? Oper Res Lett 33(3):301–304
31. Van Miegham J (1995) Dynamic scheduling with convex delay costs: the generalized $c\mu$ rule. Ann Appl Prob 5(3):809–833
32. Childers AK (2010) Prioritizing patients for emergency evacuation from a healthcare facility. PhD thesis. Clemson University