



Operations Research

Publication details, including instructions for authors and subscription information:
<http://pubsonline.informs.org>

Dynamic Distribution of Patients to Medical Facilities in the Aftermath of a Disaster

Alex F. Mills, Nilay Tanık Argon, Serhan Ziya

To cite this article:

Alex F. Mills, Nilay Tanık Argon, Serhan Ziya (2018) Dynamic Distribution of Patients to Medical Facilities in the Aftermath of a Disaster. *Operations Research* 66(3):716-732. <https://doi.org/10.1287/opre.2017.1695>

Full terms and conditions of use: <http://pubsonline.informs.org/page/terms-and-conditions>

This article may be used only for the purposes of research, teaching, and/or private study. Commercial use or systematic downloading (by robots or other automatic processes) is prohibited without explicit Publisher approval, unless otherwise noted. For more information, contact permissions@informs.org.

The Publisher does not warrant or guarantee the article's accuracy, completeness, merchantability, fitness for a particular purpose, or non-infringement. Descriptions of, or references to, products or publications, or inclusion of an advertisement in this article, neither constitutes nor implies a guarantee, endorsement, or support of claims made of that product, publication, or service.

Copyright © 2018, INFORMS

Please scroll down for article—it is on subsequent pages

INFORMS is the largest professional society in the world for professionals in the fields of operations research, management science, and analytics.

For more information on INFORMS, its publications, membership, or meetings visit <http://www.informs.org>

Dynamic Distribution of Patients to Medical Facilities in the Aftermath of a Disaster

Alex F. Mills,^a Nilay Tanik Argon,^b Serhan Ziya^b

^a Department of Operations and Decision Technologies, Kelley School of Business, Indiana University, Bloomington, Indiana 47405;

^b Department of Statistics and Operations Research, University of North Carolina at Chapel Hill, Chapel Hill, North Carolina 27599

Contact: millsaf@indiana.edu,  <http://orcid.org/0000-0003-4087-0441> (AFM); nilay@unc.edu (NTA); ziya@unc.edu,

 <http://orcid.org/0000-0003-1558-6051> (SZ)

Received: December 10, 2013

Revised: March 16, 2015; April 18, 2016;
March 7, 2017; July 18, 2017

Accepted: August 3, 2017

Published Online in Articles in Advance:
May 23, 2018

Subject Classifications: dynamic programming/optimal control: applications; health care: ambulance service

Area of Review: Policy Modeling and Public Sector OR

<https://doi.org/10.1287/opre.2017.1695>

Copyright: © 2018 INFORMS

Abstract. In the aftermath of a disaster, emergency responders must transport a large number of patients to medical facilities, using limited transportation resources (such as ambulances). Decisions about where to send the patients are typically made in an ad hoc manner by responders on the scene. Using a Markov decision process formulation, we develop two heuristic policies that use limited information such as mean travel times and congestion levels to determine (a) how to allocate ambulances to patient locations and (b) which medical facility should be the destination for those ambulances. In a simulation study, we incorporate patient survival rates and service times for different types of traumatic injuries, and show that the proposed heuristics can provide substantial improvement in the expected number of survivors compared to the common practice of transporting to the nearest facility, even when the decision maker has only limited up-to-date information about the system state. In particular, a myopic approach that considers only what is best for the next patient to be transported increases the expected number of survivors in almost all scenarios considered. Using a more sophisticated one-step policy improvement approach provides further improvement when the event involves patients who do not deteriorate rapidly, especially when the transportation is not the bottleneck and the casualties are spread over many locations. We demonstrate the effectiveness of the proposed heuristics on a case study of a hypothetical earthquake, where casualty data is generated using computer software developed by the U.S. government.

Funding: This work was partially supported by the National Science Foundation [Grants CMMI-0927607, CMMI-1234212, and CMMI-1635574].

Supplemental Material: The e-companion is available at <https://doi.org/10.1287/opre.2017.1695>.

Keywords: mass-casualty incident • disaster response • transportation • routing • queueing network

1. Introduction

Disasters or mass-casualty events can affect hundreds of people and place a huge burden on emergency medical services. In the aftermath of such events, the goal of emergency response management is to ensure the safety of the affected people and the timely treatment of as many patients as possible. This requires several complex decisions and periodically re-evaluating these decisions based on changing system conditions such as crowding at the treatment facilities. In this article, we study *how to distribute patients from areas affected by a disaster to medical facilities that participate in the response effort*. This dynamic decision-making problem, which is commonly called the *patient-distribution problem* (CDC 2010), arises in the aftermath of large-scale emergency events affecting multiple regions that are geographically separated but close enough to be served by the same group of treatment facilities (e.g., in the same urban area).

In most civilian settings, patient-distribution decisions are made by the Emergency Medical Services (EMS) transport officer, who acts under the authority of the incident commander. This officer considers each

hospital's medical capabilities, capacity or congestion level, and distance from the scene when deciding where to transport patients. Available information is synthesized in an ad hoc manner to quickly clear patients from the scene without overwhelming any particular hospital to the extent possible. From our personal communication with emergency response participants and planners in Indiana and North Carolina, we learned that these ad hoc decisions vary widely depending on the local protocol and the individual incident. It appears that there is no national or regional standard in the United States for selecting destination hospitals for patients. While extensive guidelines for field triage have been established and well-studied, the authors of these guidelines have stated that "research is needed that includes multiple sites, multiple EMS agencies, trauma and nontrauma hospitals" and that "[the] impact of geography on triage, [and] issues regarding proximity to trauma centers... [are] poorly understood" (Sasser et al. 2012, p. 15).

The importance of patient-distribution decisions has also been noted in published literature. Larson et al.

(2006) review the operational response to five emergencies from 1989 to 2004, including two hurricanes, the Oklahoma City bombing, an airline crash, and a chemical attack on the Tokyo subway system. Among these, the incident with the worst operational patient distribution was the subway attack. In the aftermath of this incident, due to the geographical dispersion of patients, poor communication, and the lack of a patient-distribution policy, some ambulances traveled from hospital to hospital looking for available beds, while others dropped victims off at nearby hospitals with no consideration of bed availability, and still others made the (sometimes incorrect) decision of traveling to distant hospitals. Larson et al. (2006, p. 494) conclude that cities must establish procedures to “assign victims to hospitals, taking into account the victims’ locations and conditions and the hospitals’ capacities.”

In the aftermath of a disaster, non-ambulatory patients with serious injuries require two types of resources: (i) emergency vehicles such as ambulances or helicopters for *transportation* and (ii) hospital resources such as beds, medical personnel, and equipment for *treatment*. The salient feature of the patient-distribution problem is that there are enough patients requiring transportation to hospitals at the same time that the available treatment capacity at a single facility is not enough to provide all of them with timely treatment. As a result, the patient-distribution problem following a disaster features two types of endogeneity not typically found together in existing models of emergency medical services. First, because casualties arrive all at once and transportation resources are limited, the decision of where to send a particular patient affects the demand at future points in time (i.e., the number of patients remaining on scene), unlike in existing EMS models where future demand arrives independently of the past. Second, because treatment resources are limited, the decision of where to send a particular patient affects the available treatment capacity at future points in time, which is not explicitly modeled in studies of EMS operations.

Our goal is to develop dynamic decision rules for distribution of non-ambulatory patients to treatment facilities, accounting for congestion levels at the facilities in addition to travel times and hospital capabilities. We consider a general formulation of the patient-distribution problem where the number of non-ambulatory patients is sufficiently large that transportation resources (such as emergency vehicles) and hospital treatment resources (such as beds, personnel, and medical equipment) are limited, and hence the location and numbers of patients and treatment resources must be taken into account. We build a mathematical model that explicitly accounts for the limited availability of these two types of resources, and assumes that transportation and treatment times, and hence the amount of time spent waiting for these

resources, are subject to uncertainty. In our model, patients are characterized by their geographic locations, which determine the distribution of travel times to receiving facilities.

After reviewing the literature in Section 2, we formulate the patient-distribution problem as a Markov decision process (MDP) in Section 3. We first obtain some analytical characterizations for the optimal solution to this MDP. For example, we provide conditions under which patients should be routed to the least congested facility. Based on the structure of our MDP formulation, we develop two policies that use travel times, service rates, and state information to make dynamic patient-distribution decisions (see Section 4). Specifically, we use two strategies in developing these policies, i.e., (i) a myopic strategy that optimizes the decision for the next patient to be transported and (ii) a policy improvement approach that has worked well in the literature. We use a discrete-event simulator to conduct a randomized simulation study that gives insights into when each patient-distribution policy would be expected to work well (see Section 5). Finally, we demonstrate the effectiveness of our proposed patient-distribution policies on a case study using data from a disaster simulator (see Section 6). Proofs of all analytical results and additional technical material are provided in an electronic companion.

2. Literature Review

Emergency vehicle allocation and routing problems in the aftermath of a mass-casualty event have been studied in several recent articles. Gong and Batta (2007) consider the problem of determining an initial static allocation of ambulances to groups of patients in a disaster and use deterministic optimization methods to suggest an allocation that minimizes the total time needed to evacuate all patients. Jotshi et al. (2009) consider the problem of both allocating and routing emergency vehicles in a disaster, using parameters such as travel times and number of patients, which are assigned weights by the user. In another relevant work, Zayas-Cabán et al. (2013) develop policies to mobilize ambulances from surrounding areas in response to a mass-casualty incident and to allocate this additional capacity among affected regions. The objective is to clear the system as quickly as possible while keeping the cost of moving vehicles under a certain budget. Dean and Nair (2014) formulate a mixed-integer program (MIP) for optimal patient transportation from a single location. Mills et al. (2018) demonstrate the benefit of sharing information about emergency department and inpatient beds when assigning patients to hospitals while also accounting for the probabilistic need for inpatient resources. Our work differs substantially from these four articles. Neither Gong and Batta (2007) nor Zayas-Cabán et al. (2013) consider

treatment capacity at the hospitals. Jotshi et al. (2009) consider congestion at the receiving facilities, but they assume that this (and other system state information) is exogenous and independent of the decision on where to send the patients. Mills et al. (2018) consider capacity at the hospitals when making the patient-distribution decision, but they assume the capacity is static and they do not consider transportation resources to be limited. By contrast, we consider two types of resources (transportation and hospital) and model the effects of patient-distribution decisions on future resource availability. Dean and Nair (2014) consider both transportation and treatment capacity as endogenous to the patient-distribution decisions, however they do not consider any uncertainty in travel or treatment times. Instead, they make all patient-distribution decisions at once by solving an integer program. While such a global optimization approach can be used to study different scenarios offline, it may not be a practical way to make decisions at the scene of a disaster. This approach requires substantial data entry at the time of the event, and solving an integer program may yield a distribution policy that is not easy to implement.

There is also a vast literature on the daily operations of EMS, but with a focus on ambulance location and redeployment problems to optimize coverage for randomly occurring calls; see, e.g., Ingolfsson et al. (2008), Erkut et al. (2008), Maxwell et al. (2010), McLay and Mayorga (2013), and Yue et al. (2012). In all of these articles, ambulance requests arrive one at a time. For example, in both Maxwell et al. (2010) and McLay and Mayorga (2013), patients arrive according to a Poisson process and are assigned to an ambulance, which is thereafter busy for a random amount of time. Thus, assignment decisions affect future resource availability, but not future demand (which is independent). These assumptions are appropriate for daily EMS operations, where demand arises one patient at a time. Furthermore, since these papers are concerned with coverage, none of them account for hospital congestion. However, in the aftermath of a disaster when many patients arrive at once, both future demand (i.e., the number of patients remaining at each location) and available hospital capacity depend on the decision.

To our knowledge, there is no work in the literature that studies the patient-distribution problem in daily emergencies. The closest problem studied in this literature is the ambulance diversion problem; see Deo and Gurvich (2011) and Allon et al. (2013). Ambulance diversion occurs when hospitals turn away ambulances due to high patient loads. Both of these articles use queueing systems to model congestion levels at hospitals, but neither considers the scarcity of ambulances, which is a more salient feature in disasters than in daily emergencies. They also do not consider travel times to hospitals, assuming they would be roughly equal,

although Deo and Gurvich (2011) suggest inclusion of differing travel times as a future research direction.

Finally, we note that dynamic routing of customers or jobs to servers is a well-studied problem in the queueing literature. In the case where servers are homogeneous, routing customers to the shortest queue is optimal in many situations; see Hordijk and Koole (1992) for one such result, and Whitt (1986) for several examples and counterexamples. A number of articles have developed index policies for routing customers in more complex systems, such as those with heterogeneous servers, impatient customers, or general delay costs; see, e.g., van Mieghem (1995), Mandelbaum and Stolyar (2004), Armony (2005), Glazebrook et al. (2009), and Argon et al. (2009). These articles assume that customers are routed to a server instantaneously, either upon their arrival (after which they may wait in a queue at their assigned server), or just before service (after waiting in a single queue). In our problem, this assumption would correspond to ambulances traveling instantaneously to medical facilities and would eliminate a feature that is essential for a realistic representation of the actual system. Hence, in our formulation, transporting a patient to a medical facility takes time and requires the allocation of a limited resource.

We conclude this section by returning to the patient-distribution problem in the aftermath of a disaster. If transportation resources are limited but hospital treatment resources are not, prior models of mass-casualty incident response could be appropriate. On the other hand, if transportation is instantaneous but treatment takes time and hospital resources are limited, classical queue-routing models could be applied. However, to our knowledge, none of the existing literature provides a clear and practical way to solve the patient-distribution problem when the number of patients is large enough that both types of resources are limited.

3. Markov Decision Process Formulation and Analysis

In this section, we formulate the problem under consideration as an MDP. To focus on the fundamental trade-off inherent to this problem and develop solution methods that can lead to heuristic policies (see Section 4), we purposely abstract away from certain features of the problem. We later conduct simulation experiments to test the proposed methods under much more general conditions where the simplifying assumptions are relaxed (see Sections 5 and 6).

In our formulation, we consider a disaster with multiple casualty locations that are geographically separated, and multiple receiving facilities such as major hospitals. Denote the set of locations by \mathcal{L} and the set of facilities by \mathcal{F} . Patients must be transported using a transportation resource such as an ambulance. Ambulances may be dedicated to one of the casualty locations or they may be flexible (i.e., they can operate

between any location and any facility). Let n_i denote the number of ambulances dedicated to location $i \in \mathcal{L}$ and n^f denote the number of flexible ambulances. A resource assigned to location $i \in \mathcal{L}$ can transport patients to facility $j \in \mathcal{F}$ with a travel time that is exponentially distributed with rate τ_{ij} . Primarily for tractability, we assume that the return travel is instantaneous. Once a patient reaches the facility, he or she receives treatment from one of b_j servers (which may represent beds or medical teams), each of which provides service that takes an exponentially distributed time with rate μ_j . If there are more patients than servers at the receiving facility, the remaining patients wait in a queue. After completion of treatment of a patient, the system earns a finite reward r_j , which possibly depends on the facility $j \in \mathcal{F}$ where the patient receives service. Thus, the rewards can be used to represent differences in quality or capability among the facilities. The performance measure of interest is the expected total discounted reward earned by the system. (In the simplest case, all rewards can be set to one, in which case the objective function reduces to the expected discounted throughput.) By discounting the rewards, we capture the diminishing benefits of medical treatments provided to the patients as the patients' conditions deteriorate with time. For example, if r_j represents the survival probability of a patient treated at facility j , then the objective function corresponds to the expected number of survivors. We seek an optimal dynamic policy that will specify, for each state, the destination of each dedicated resource, and the assigned location and destination of each flexible resource.

Denote the system state by $(\mathbf{W}(t), \mathbf{X}(t)) = (W_1(t), W_2(t), \dots, W_{|\mathcal{L}|}(t), X_1(t), \dots, X_{|\mathcal{F}|}(t))$, where $W_i(t)$ is the number of patients at location $i \in \mathcal{L}$ at time t , $X_j(t)$ is the number of patients at facility $j \in \mathcal{F}$ at time $t \geq 0$, and $|\cdot|$ denotes the cardinality of a set. The state space, which we denote by \mathcal{S} , is equivalent to $\mathbb{N}^{|\mathcal{L}|} \times \mathbb{N}^{|\mathcal{F}|}$, where \mathbb{N} is the set of nonnegative integers. The expected total discounted reward with discount rate $\alpha > 0$ is given by $E[\int_0^\infty e^{-\alpha t} r(\mathbf{X}(t)) dt]$, where $r(\mathbf{x}) \equiv \sum_{j \in \mathcal{F}} r_j (b_j \wedge x_j) \mu_j$ and \wedge is the minimum operator.

We formulate this optimization problem as an MDP, using uniformization (as in Lippmann 1975) with uniformization constant $\beta \equiv \sum_{j \in \mathcal{F}} \mu_j b_j + \sum_{i \in \mathcal{L}} \sum_{j \in \mathcal{F}} \tau_{ij} \cdot (n_i + n^f)$, and study the embedded discrete-time MDP. For this problem it is sufficient to consider only stationary deterministic policies (see, e.g., theorem 6.2.10 in Puterman 1994). We denote the set of such policies by \mathcal{P} .

Let $V(\mathbf{w}, \mathbf{x})$ denote the maximum expected total discounted reward that can be obtained by a policy in \mathcal{P} when the system starts in state $(\mathbf{w}, \mathbf{x}) \in \mathcal{S}$. We are interested in finding a policy in \mathcal{P} that yields $V(\mathbf{w}, \mathbf{x})$. In other words, for any initial state $(\mathbf{w}, \mathbf{x}) \in \mathcal{S}$, we want to solve

$$\max_{\pi \in \mathcal{P}} E \left[\int_0^\infty e^{-\alpha t} r(\mathbf{X}^\pi(t)) dt \mid \mathbf{W}^\pi(0) = \mathbf{w}, \mathbf{X}^\pi(0) = \mathbf{x} \right], \quad (1)$$

where $\{\mathbf{W}^\pi(t), t \geq 0\}$ and $\{\mathbf{X}^\pi(t), t \geq 0\}$ are multidimensional stochastic processes representing the number of patients at each location and facility, respectively, at time $t \geq 0$ under policy $\pi \in \mathcal{P}$.

The decision variables are $\mathbf{d} = \{d_{ij}, i \in \mathcal{L}, j \in \mathcal{F}\}$, where d_{ij} is the number of dedicated ambulances of location i to be sent to facility j , and $\mathbf{f} = \{f_{ij}, i \in \mathcal{L}, j \in \mathcal{F}\}$, where f_{ij} is the number of flexible ambulances to be assigned to location i with destination j . These variables must satisfy the constraints

$$\sum_{j \in \mathcal{F}} (d_{ij} + f_{ij}) \leq w_i \quad \forall i \in \mathcal{L}, \quad (2)$$

$$\sum_{j \in \mathcal{F}} d_{ij} \leq n_i \quad \forall i \in \mathcal{L}, \quad (3)$$

$$\sum_{i \in \mathcal{L}} \sum_{j \in \mathcal{F}} f_{ij} \leq n^f. \quad (4)$$

Define $\mathcal{A}(\mathbf{w})$ to be the set of actions (\mathbf{d}, \mathbf{f}) that satisfy (2)–(4). Then for $(\mathbf{w}, \mathbf{x}) \in \mathcal{S}$, the optimality equations are

$$\begin{aligned} V(\mathbf{w}, \mathbf{x}) &= \frac{1}{\alpha + \beta} \left[\sum_{j \in \mathcal{F}} \mu_j (b_j \wedge x_j) [r_j + V(\mathbf{w}, \mathbf{x} - \mathbf{e}_j) - V(\mathbf{w}, \mathbf{x})] \right. \\ &\quad + \beta V(\mathbf{w}, \mathbf{x}) + \max_{(\mathbf{d}, \mathbf{f}) \in \mathcal{A}(\mathbf{w})} \left\{ \sum_{i \in \mathcal{L}} \sum_{j \in \mathcal{F}} \tau_{ij} (d_{ij} + f_{ij}) \right. \\ &\quad \left. \left. \cdot [V(\mathbf{w} - \mathbf{e}_i, \mathbf{x} + \mathbf{e}_j) - V(\mathbf{w}, \mathbf{x})] \right\} \right]. \quad (5) \end{aligned}$$

Throughout, \mathbf{e}_k is a vector with k th component equal to one and all others equal to zero (the size of \mathbf{e}_k will be clear from the context). In (5), the first summation (involving μ_j) computes the marginal expected reward earned due to a service completion at the next epoch, the term $\beta V(\mathbf{w}, \mathbf{x})$ is the uniformization term, and the term inside the max operator computes the marginal expected future reward due to a patient transfer from location i to facility j at the next epoch.

We begin our analysis by considering the case where the value function $V(\cdot)$ is known, in which case we can completely characterize the optimal decisions. Define $m_{ij}^{\mathbf{w}, \mathbf{x}} = \tau_{ij} (V(\mathbf{w} - \mathbf{e}_i, \mathbf{x} + \mathbf{e}_j) - V(\mathbf{w}, \mathbf{x}))$, for $(\mathbf{w}, \mathbf{x}) \in \mathcal{S}$, $i \in \mathcal{L}$, and $j \in \mathcal{F}$. From (5), solving the following integer program (with C a constant) yields the optimal action:

$$\begin{aligned} &\text{maximize} \quad \sum_{i \in \mathcal{L}} \sum_{j \in \mathcal{F}} m_{ij}^{\mathbf{w}, \mathbf{x}} (d_{ij} + f_{ij}) + C \\ &\text{s.t.} \quad (2), (3), \text{ and } (4). \quad (6) \end{aligned}$$

Proposition 1. *Algorithm 1 returns an optimal solution to (6) for given $V(\cdot)$.*

Algorithm 1 (Algorithm for obtaining the optimal policy, given values of $V(\cdot)$)

- 1: **function** FINDPOLICY($\{m_{ij}^{\mathbf{w}, \mathbf{x}}\}, \mathbf{w}, \{n_i\}, n^f$)
- 2: **for all** $i \in \mathcal{L}, j \in \mathcal{F}$ **do** $d_{ij} \leftarrow 0, f_{ij} \leftarrow 0$

```


3: list ← {(i, j), i ∈ ℒ, j ∈ ℱ}
4: SortDescending(list, mijw,x)
   ▷ Sorts list in descending order of mijw,x.
5: for k = 1 to Length(list) do
6:   (i, j) ← list[k]
7:   while ∑s ∈ ℱ (dis + fis) < wi do
8:     if ∑s ∈ ℱ dis < ni then dij ← dij + 1
9:     else if ∑r ∈ ℒ ∑s ∈ ℱ frs < nf then fij ← fij + 1
10:    else break
    
```

In Algorithm 1, ambulances are systematically assigned to routes (each route comprises one location and one facility) in order of decreasing marginal reward $m_{ij}^{w,x}$, with priority given to assigning dedicated ambulances. Given that this greedy-like algorithm can be used to assign ambulances, we use this solution structure to design near optimal policies (in Section 4).

We continue our analysis by considering cases where some simple policy, such as a “shortest-wait”, is optimal, as in many queue-routing problems (see, e.g., Whitt 1986 and Hordijk and Koole 1992). Incorporation of transportation times makes our problem different from many well-studied problems. To demonstrate, suppose that for two facilities $j, l \in \mathcal{F}$, we have $r_j \geq r_l$, $\mu_j \geq \mu_l$, $b_j \geq b_l$, and $\tau_{kj} \geq \tau_{kl}$ for all $k \in \mathcal{L}$. Now, suppose that for the current state, we have $x_j \leq x_l$. In this case, facility j appears to be a better choice than facility l with respect to every single factor: It has a higher reward, faster service, more servers, fewer patients waiting, and a shorter travel time from every location. Nonetheless, an optimal policy may send the next eligible patient from a location to facility l . Example 1 demonstrates this phenomenon.

Example 1. Consider the example shown in Figure 1, where $\mu_A = 8$ per hour, $\mu_B = 6.5$ per hour, $b_A = b_B = 2$, $r_A = r_B = 1$, $n_1 = n_2 = 1$, $n^f = 1$, and $\alpha = 0.7$. (Note that facility A is a “better” choice for both locations because it is closer and faster than facility B.) We obtain the optimal policy using the value iteration algorithm and truncating the state space to 50 patients on each dimension; see Figure 2. We observe that ambulances may be routed to facility B even when facility A has the shorter queue (center plot of Figure 2(a))—indeed, we

Figure 1. An Example with Two Locations (1 and 2) and Two Facilities (A and B)



Note. Numbers on the arcs indicate the mean travel times (in minutes) between the locations and treatment facilities.

may route *all* ambulances to facility B, even when facility A has the shorter queue (all plots of Figure 2(b)). This is because routing ambulances to facility A too frequently would overwhelm it while leaving facility B at risk of becoming idle. One can fill the queue at facility A relatively “quickly” with the dedicated ambulance from location 1, which has many patients available. Therefore, when facility A has a sufficient number of patients in its queue to keep it busy for a while, one can afford to devote some of the transportation resources to taking patients to facility B. Nevertheless, Figure 2(c) shows that when the total number of patients at the scene is smaller, then the intuition to send patients to the better facility holds.

Unlike in Example 1, however, if any two facilities are equivalent in terms of travel times, then as we next state in Proposition 2, it is always optimal to choose the one that is superior with respect to the other criteria, which aligns more closely with our intuition on queue routing problems. A situation in which no facility is better than the other in terms of travel times may arise, for example, when all the patient locations are in a remote area and facilities are in the same nearby city (in this case, travel times may be roughly the same regardless of which facility is chosen).

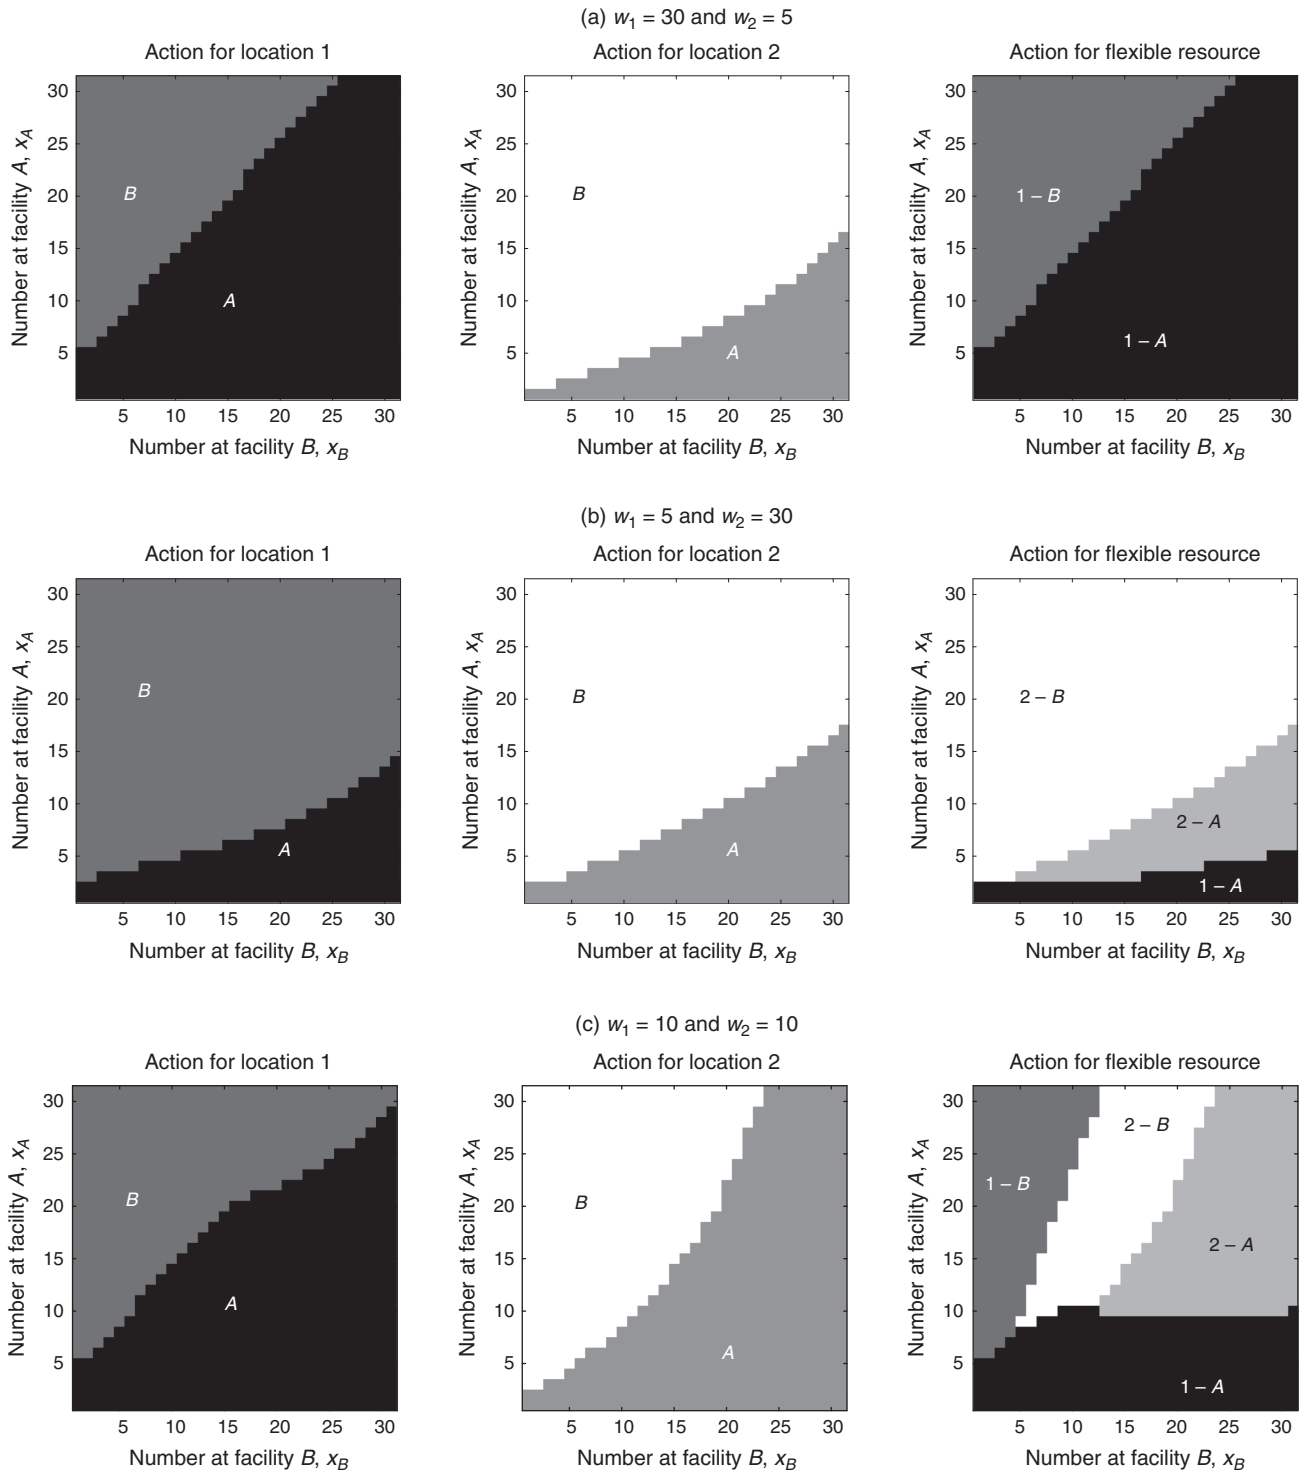
Proposition 2. *Suppose that for two facilities $j, l \in \mathcal{F}$, we have $r_j \geq r_l$, $\mu_j \geq \mu_l$, $b_j \geq b_l$, and $\tau_{kj} = \tau_{kl}$ for all $k \in \mathcal{L}$. If $x_j \leq x_l$, then $V(\mathbf{w}, \mathbf{x} + \mathbf{e}_j) \geq V(\mathbf{w}, \mathbf{x} + \mathbf{e}_l)$, for all $(\mathbf{w}, \mathbf{x}) \in \mathcal{S}$.*

In our model, each patient’s total delay consists of scene waiting time, transportation time, hospital waiting time, and treatment time. We now turn our analysis to optimal policies that result from simplifying some aspect of this structure. In particular, we study the case wherein there is no waiting at the hospitals, which is equivalent to assuming that the number of servers b_j is sufficiently large that $b_j \wedge x_j = x_j$ for all possible values of x_j , $j \in \mathcal{F}$. In this case, the expected marginal reward $m_{ij}^{w,x}$ has the following structure:

Proposition 3. *Suppose that b_j is sufficiently large that $b_j \wedge x_j = x_j$ for all possible values of x_j and all $j \in \mathcal{F}$. Then $m_{ij}^{w,x} = \tau_{ij}(r_j \mu_j / (\alpha + \mu_j) - \kappa_i(\mathbf{w}))$, where $\kappa_i(\mathbf{w}) \geq 0$ does not depend on \mathbf{x} or j . Furthermore, for all $i \in \mathcal{L}$, $\kappa_i(\mathbf{w}) < \max_{j \in \mathcal{F}} \{r_j \mu_j / (\mu_j + \alpha)\} [(\max_{l \in \mathcal{F}} \tau_{il}(n^f + n_i)) / (\max_{l \in \mathcal{F}} \tau_{il}(n^f + n_i) + \alpha)]^{w_i}$.*

Proposition 3 shows that in the absence of hospital waiting, the index policy used in Algorithm 1 chooses location—facility pairs according to $m_{ij}^{w,x}$, which is inversely proportional to the mean travel time τ_{ij}^{-1} , directly proportional to the reward r_j , and directly proportional to $\mu_j / (\mu_j + \alpha)$, which accounts for differences in service rate. This index is modulated by $\kappa_i(\mathbf{w})$, which depends on the number of patients remaining to be transported and diminishes exponentially in the number of patients. Therefore, when the number of patients

Figure 2. Optimal Policy for Example 1 When There Are (a) 30 Patients at Location 1 and 5 Patients at Location 2, (b) 5 Patients at Location 1 and 30 Patients at Location 2; and (c) 10 Patients at Each Location



Notes. The leftmost (center) plots show the optimal destination facility for dedicated ambulances from location 1 (2) as a function of the number of patients at each facility. The rightmost plots show the optimal origin-destination (OD) pair for the flexible ambulance.

at a location is large, we can approximate $m_{ij}^{w,x}$ by $\tau_{ij} r_j \mu_j / (\alpha + \mu_j)$, which leads to a policy that is similar to simply choosing the nearest facility when treatment facilities are not too different in terms of service rate

and quality. On the other hand, facilities with large $\kappa_i(\mathbf{w})$ (which can happen when the number of patients is smaller) are less likely to be selected by the flexible ambulances. This result is consistent with what we

observe in Figure 2, even though the number of beds in Example 1 is small.

4. Heuristic Policies

In Proposition 3, we showed that when there is no hospital waiting, the optimal solution is an index policy independent of \mathbf{x} , although it depends on \mathbf{w} . With such a policy, the decision maker will need information only on one component of the state. Policies that are \mathbf{w} -independent but are \mathbf{x} -dependent are in fact even more desirable because in the aftermath of a disaster obtaining information from the scene could be difficult but the medical facilities are expected to provide information to coordinators at the scene of a mass-casualty incident. Therefore, in this section, we seek heuristic policies that are \mathbf{w} -independent but may depend on \mathbf{x} . We take two approaches to develop such heuristic policies: a *Myopic* approach (Section 4.1) and a *Policy Improvement* approach (Section 4.2). Because of possible problems in communication infrastructure following a disaster, information on the congestion level at the hospitals may not be perfect. We address this issue in Section 4.3 by providing a modification for our heuristics in case of uncertain state information.

4.1. Myopic Approach

Consider the Myopic policy, which maximizes the expected reward of the patient reaching a medical facility in the next decision epoch (if any). More specifically, assume without loss of generality that the uniformization constant $\beta = 1$. Then, if a patient from location i is sent to facility j , the patient reaches facility j in the next epoch with probability τ_{ij} , in which case she becomes the $(x_j + 1)$ st patient. If $x_j + 1 \leq b_j$, the patient begins treatment immediately. Otherwise, she joins the queue and her hospital waiting time is an Erlang random variable with $x_j + 1 - b_j$ phases, each having mean $1/(b_j\mu_j)$, after which her service time is exponentially distributed with mean $1/\mu_j$. Hence, the expected discounted reward for the current patient is $r_j[1/(1+\alpha)] \cdot [\mu_j/(\mu_j + \alpha)][b_j\mu_j/(b_j\mu_j + \alpha)]^{(x_j+1-b_j)^+}$ if she reaches facility j (which occurs with probability τ_{ij}), and zero otherwise. Because $(1+\alpha)^{-1}$ is positive and does not depend on i or j , we use the value

$$\tilde{m}_{ij}^x = \tau_{ij}r_j \left(\frac{\mu_j}{\mu_j + \alpha} \right) \left(\frac{b_j\mu_j}{b_j\mu_j + \alpha} \right)^{(x_j+1-b_j)^+} \quad (7)$$

as an approximation for $m_{ij}^{\mathbf{w},\mathbf{x}}$ in Algorithm 1.

The resulting policy, which we call the Myopic policy, will yield the same set of actions for all values of \mathbf{w} ; it can also be written as a set of switching curves that are linear in the number of waiting patients. From location $i \in \mathcal{L}$, facility $j \in \mathcal{F}$ is chosen by the Myopic policy if and only if

$$(x_j + 1 - b_j)^+ \leq a_{kj}(x_k + 1 - b_k)^+ + c_{ijk}, \quad (8)$$

for all $k \in \mathcal{F} \setminus \{j\}$, where

$$a_{kj} = \frac{\log(b_k\mu_k/(b_k\mu_k + \alpha))}{\log(b_j\mu_j/(b_j\mu_j + \alpha))}$$

$$\text{and } c_{ijk} = \frac{\log(\tau_{ik}r_k\mu_k(\mu_j + \alpha)/(\tau_{ij}r_j\mu_j(\mu_k + \alpha)))}{\log(b_j\mu_j/(b_j\mu_j + \alpha))}.$$

Expressing the heuristic in this linear form yields structural insights into the Myopic policy. In particular, for a given $i \in \mathcal{L}$ and $k \in \mathcal{F}$, $\min_{j \in \mathcal{F}} c_{ijk}$ is a lower bound on the incident size necessary for facility k to be chosen for patients from location i . To demonstrate, in Example 1, we have $c_{1AB} = 16.6$ and $c_{2AB} = 7.2$, implying that patients from location 1 would be sent to facility B only if the incident has at least 17 patients in total; the same can be said of location 2 for incidents of 8 or more total patients. Hence, in this scenario, the Myopic policy recommends a response effort involving *only* facility A , when there are fewer than 8 total patients.

Finally, we can show that the Myopic policy is consistent with Proposition 2 in that it chooses the optimal action when the conditions of Proposition 2 are satisfied. It is also consistent with Proposition 3 as it reduces to the index $\tau_{ij}r_j\mu_j/(\mu_j + \alpha)$ when $w_i, i \in \mathcal{L}$ and $b_j, j \in \mathcal{F}$ are large.

4.2. Policy Improvement Approach

The Myopic policy is concerned with only the very next patient. In this section, we take the opposite approach and develop a \mathbf{w} -independent heuristic based on the case where there are infinitely many patients. For tractability, we also assume that servers are pooled at each facility, i.e., facility j has a single server and service times are exponentially distributed with rate $b_j\mu_j$. Define $V_\infty(\mathbf{x})$ to be the expected total discounted reward for any state \mathbf{x} for this limiting case, which satisfies the following optimality equation:

$$V_\infty(\mathbf{x}) = \frac{1}{\alpha + \beta} \left[\sum_{j \in \mathcal{F}} \mu_j b_j [r_j + V_\infty(\mathbf{x} - \mathbf{e}_j) - V_\infty(\mathbf{x})] \right. \\ \left. + \beta V_\infty(\mathbf{x}) + \max_{(\mathbf{d}, \mathbf{f}) \in \mathcal{B}} \left\{ \sum_{i \in \mathcal{L}} \sum_{j \in \mathcal{F}} \tau_{ij} (d_{ij} + f_{ij}) \right. \right. \\ \left. \left. \cdot [V_\infty(\mathbf{x} + \mathbf{e}_j) - V_\infty(\mathbf{x})] \right\} \right], \quad (9)$$

where \mathcal{B} is the set of (\mathbf{d}, \mathbf{f}) that satisfy (3) and (4). We will use a policy improvement approach for (9) to obtain a \mathbf{w} -independent heuristic. Using a policy improvement approach requires (i) calculating the value function for a state-independent (static) policy in closed form, and (ii) applying a single step of the policy improvement algorithm to (9) to obtain an index policy. Several studies from the queueing-control literature successfully implemented this approach to develop near-optimal policies (see, e.g., Krishnan 1990, Ansell et al. 2003, and Argon et al. 2009).

We use *Bernoulli splitting* to obtain an initial static policy, where dedicated ambulances from location $i \in \mathcal{L}$ are sent to facility $j \in \mathcal{F}$ independently with probability $\rho_{ij} \geq 0$ at each decision epoch, and flexible ambulances are sent from location $i \in \mathcal{L}$ to facility $j \in \mathcal{F}$ with probability $\theta_{ij} \geq 0$.

Proposition 4. *Let γ be a Bernoulli splitting policy having probabilities $\{\rho_{ij}, \theta_{ij}, i \in \mathcal{L}, j \in \mathcal{F}\}$ and denote by $V_\infty^\gamma(\mathbf{x})$ the value function associated with γ . Then we have*

$$V_\infty^\gamma(\mathbf{x} + \mathbf{e}_j) - V_\infty^\gamma(\mathbf{x}) = \frac{b_j \mu_j r_j}{\lambda_j} \cdot \frac{b_j \mu_j - \lambda_j + \alpha - \eta_j}{b_j \mu_j - \lambda_j - \alpha - \eta_j} \cdot \left(\frac{b_j \mu_j + \lambda_j + \alpha - \eta_j}{2\lambda_j} \right)^{x_j}, \quad (10)$$

where $\lambda_j \equiv \sum_{i \in \mathcal{L}} \tau_{ij}(n_i \rho_{ij} + n^f \theta_{ij})$ and $\eta_j \equiv [(b_j \mu_j + \lambda_j + \alpha)^2 - 4\lambda_j b_j \mu_j]^{1/2}$.

Applying one step of the policy improvement algorithm corresponds to using (10) in place of $V_\infty(\mathbf{x} + \mathbf{e}_j) - V_\infty(\mathbf{x})$ inside the maximum operator in (9). We call the resulting policy the *Policy Improvement Heuristic (PIH)*, which is implemented using $\hat{m}_{ij}^x = \tau_{ij}(V_\infty^\gamma(\mathbf{x} + \mathbf{e}_j) - V_\infty^\gamma(\mathbf{x}))$ in place of $m_{ij}^{w,x}$ in Algorithm 1. Like the Myopic policy, the PIH will yield the same set of actions for all values of \mathbf{w} ; it can also be written as a set of linear switching curves. Unlike the Myopic policy, PIH depends on an initial static policy γ . In Section EC.2 of the electronic companion, we describe three such initial static policies; the simplest of the three (which we use in our simulation experiments of Section 5) uses a greedy algorithm to assign routing probabilities ρ_{ij} and θ_{ij} in decreasing order of index $\tau_{ij} r_j \exp(-\alpha/\tau_{ij})$, which is based on the objective function of a reward-maximizing fluid approximation.

4.3. Modified Heuristics for Uncertain State Information

Until now, we assumed that state information about congestion at each facility (i.e., \mathbf{x}) is known at every decision epoch for implementation of the proposed heuristics. While this information is simple to collect at each facility, it may not be easy for responders on the scene to obtain this information continuously, or the information that is transmitted might not be accurate. We now demonstrate how the Myopic policy and PIH can be modified to incorporate uncertain state information.

Let Λ_j be a random variable denoting the decision maker's belief about the number of patients at facility $j \in \mathcal{F}$. Then, we modify the Myopic and PIH policies to use the indices

$$\tau_{ij} r_j \left(\frac{\mu_j}{\mu_j + \alpha} \right) \mathbb{E} \left[\left(\frac{b_j \mu_j}{b_j \mu_j + \alpha} \right)^{(\Lambda_j + 1 - b_j)^+} \right] \quad (11)$$

and

$$\frac{\tau_{ij} b_j \mu_j r_j}{\lambda_j} \left(\frac{b_j \mu_j - \lambda_j + \alpha - \eta_j}{b_j \mu_j - \lambda_j - \alpha - \eta_j} \right) \cdot \mathbb{E} \left[\left(\frac{b_j \mu_j + \lambda_j + \alpha - \eta_j}{2\lambda_j} \right)^{\Lambda_j} \right], \quad (12)$$

respectively. Whether (11) and (12) can be further simplified depends on the distribution of Λ_j . In the specific case of intermittent status updates (i.e., state information is correct but it is obtained only sporadically), we develop an approximation that can be used as long as the decision maker keeps track of the following: (i) the state of each facility at the last status update $\{x_j, j \in \mathcal{F}\}$, (ii) the number of patients sent to each facility since the last status update, which we denote by $\{y_j, j \in \mathcal{F}\}$, and (iii) the amount of time since the last status update, which we denote by $s \geq 0$. The decision maker then can approximate the current state of each facility as $\Lambda_j = (x_j + y_j - \psi_j(s))^+$, where $\psi_j(s)$ is a Poisson random variable with mean $b_j \mu_j s$. Note that $\psi_j(s)$ is the number of *potential* departures at facility $j \in \mathcal{F}$ during a time period of length s if all beds are occupied, and is thus an approximation for the number of departures from facility j . In Section EC.4 of the electronic companion, we demonstrate an application of this approximation.

5. Simulation Study

Solving our original MDP given by (5) numerically is very time consuming even for a small number of locations and facilities; thus, it is difficult to compare the heuristics to the optimal solution (5). Therefore, we first conducted a numerical study using the infinite-patient approximation, i.e., under the modeling assumptions of Section 4.2, and found that PIH closely approximates the optimal policy, while the performance of the Myopic policy depends heavily on α (results of this study are presented in Section EC.3 of the electronic companion). Although this is an encouraging conclusion, we are more interested in the performance of the heuristics under conditions that are more realistic than those assumed by our mathematical models given in Sections 3 and 4. Hence, we dedicated the remainder of our numerical analysis to a simulation study, where we relax the assumptions that ambulance travel times and treatment times are exponentially distributed and assignments are preemptive, and we explicitly model both the forward and return travel of ambulances. We also use a reward function based on patient survivability. In this section, we use a randomized study to quantify the difference between the heuristics and certain baseline policies as the characteristics of problem instances vary, and we gain insights into situations where particular heuristics are most valuable. Later, in Section 6, we use our simulation model to conduct a realistic case study with data corresponding to a hypothetical disaster.

5.1. Problem Instances

Each generated problem instance corresponds to a trauma-related mass-casualty event in an area that is 10 miles by 10 miles, resulting in three, four, or five patient locations. Initial patient counts are generated from a discrete Uniform distribution over $\{5, 6, \dots, 75\}$. Each location has one dedicated ambulance, which can travel at an average speed of 40 mph, and there are two flexible ambulances with the same speed. There are three Level I or Level II trauma centers (TCs), two of which are in the urban area and one of which is 15 to 30 miles outside the area. The presence of the outside TC is important because it helps us observe the behavior of our heuristics in determining how congested a nearby facility must be before it should be bypassed. Distances are calculated using the L^1 -norm (Manhattan distance); travel time from a location to a hospital (and vice versa) has a lognormal distribution (Ingolfsson et al. 2008). By varying the number of locations, number of patients at each location, and location of the facilities, we induce variability in the level of congestion on the system in two dimensions: More patients mean more congestion for both the ambulances and the treatment facilities, while increased travel distance means more congestion for ambulances. Therefore, some problem instances naturally experience more scene waiting time while others experience more hospital waiting time, allowing us to examine the relative performance of each heuristic under these different settings.

We considered two types of traumatic injuries that result from a natural or man-made disaster, i.e., *blunt* trauma (e.g., from falling or being struck by an object) and *penetrating* trauma (e.g., from a firearm or explosive device). We fit emergency department (ED) treatment times by injury type and TC level (I or II) using over 380,000 observations from the 2012 National Trauma Data Bank (NTDB) (American College of Surgeons 2012). Using SAS software, we fit distributions commonly used to model the time taken to perform a task—exponential, Weibull, Gamma, and lognormal (Law 2007). The Lognormal distribution was the best fit for blunt trauma, while the Exponential distribution was the best fit for penetrating trauma. In this simulation study, we set all rewards to one; hence the performance measure of interest is the discounted throughput. To determine discount rates, we examined data for survival probability of blunt trauma and penetrating trauma patients from Sacco et al. (2005, 2007). We averaged the survival probability functions of the critical patients and observed that an exponentially decreasing function (of the form $e^{-\alpha t}$) provided a good fit (root mean squared error (RMSE) < 0.01). Using MATLAB, we determined the best-fit α for each trauma type. As expected, penetrating injuries had larger α , indicating a higher level of urgency. We provide all the parameter settings of our simulation setup in Table EC.5 of the electronic companion.

5.2. Heuristics and Baseline Policies

Because of the practical difficulty in obtaining and implementing an optimal policy that satisfies (5), our goal is to compare the state-dependent heuristics of Section 4 to some baseline policies. For both Myopic policy and PIH, we can determine the action without solving an optimization problem, and both are w -independent, which is practically advantageous. Because we relax the assumption of instantaneous return time for ambulances in our base model of Section 3, we modify the dynamic heuristics so that when a flexible ambulance drops off a patient at facility $j \in \mathcal{F}$, it next travels to location $i \in \mathcal{L}$ with the largest index $\tau_{ij} \max_{k \in \mathcal{F}} \tilde{m}_{ik}^x$ (for Myopic) or $\tau_{ij} \max_{k \in \mathcal{F}} \hat{m}_{ik}^x$ (for PIH). That is, at the time of dropoff, for *each* location, we compute the maximum index over all facilities (i.e., over $k \in \mathcal{F}$ in the maximum operators), and then we scale each location's index by τ_{ij} , reflecting the travel time back to location i from the current facility j . Then, when an ambulance (dedicated or flexible) reaches a location i , the destination facility is determined by the index $\max_{j \in \mathcal{F}} \tilde{m}_{ij}^x$ (for Myopic) or $\max_{j \in \mathcal{F}} \hat{m}_{ij}^x$ (for PIH). Another modification we have incorporated is that in the Bernoulli splitting policy used to initialize PIH, we approximate λ_j , the arrival rate of patients to facility j , by $\sum_{i \in \mathcal{L}} [n_i \rho_{ij} \tau_{ij} / 2 + n^f \theta_{ij} \sum_{k \in \mathcal{F}} \sum_{l \in \mathcal{L}} \theta_{lk} (1/\tau_{ij} + 1/\tau_{lk})^{-1}]$. Comparing this expression to the one in Proposition 4, it can be seen that we doubled the total travel time for dedicated ambulances and averaged the return travel time for flexible ambulances over all locations to which a flexible ambulance may be assigned.

As mentioned in Section 1, there is no specific national or regional standard for selecting patients' destinations, and the results may vary from one locality to another and from one incident to another. As a result of the lack of any standard policy with which to compare our heuristics, we consider two baseline policies: a baseline static (BLS) policy, where all dedicated ambulances take patients to the *nearest facility* (i.e., with the largest τ_{ij} among all $j \in \mathcal{F}$) and all flexible ambulances are assigned to the fastest route (i.e., with the largest τ_{ij} among all $i \in \mathcal{L}, j \in \mathcal{F}$); and a baseline dynamic (BLD) policy, where all patients are sent to the facility that has the *smallest time until service completion* at the time they begin transportation (i.e., with the smallest $\tau_{ij}^{-1} + (b_j \mu_j)^{-1} (x_j + 1 - b_j)^+ + \mu_j^{-1}$ over $j \in \mathcal{F}$ for dedicated ambulances and over $i \in \mathcal{L}, j \in \mathcal{F}$ for flexible ambulances). Similar to the proposed heuristics, BLD sends a flexible ambulance completing transportation at facility $j \in \mathcal{F}$ to location $i \in \mathcal{L}$ with the largest index $\tau_{ij} / \min_{k \in \mathcal{F}} \{ \tau_{ik}^{-1} + (b_k \mu_k)^{-1} (x_k + 1 - b_k)^+ + \mu_k^{-1} \}$. Note that the BLD policy uses similar logic to the Myopic heuristic. The BLS policy would maximize the discounted throughput, if service at the facilities were instantaneous (since it would minimize the time to service completion for the current patient, which would result

in the ambulance being available for future patients as soon as possible). On the other hand, the BLD policy would be equivalent to sending a patient to the facility with the least congestion if transportation were instantaneous (because in such a case, $\tau_{ij}^{-1} = 0$), a policy that is optimal in many queue routing problems in the literature (see, e.g., Whitt 1986). Thus, each baseline policy represents a policy that would make sense if the model were simplified by assuming that transportation or treatment has infinite capacity and is infinitely fast, and hence the baseline policies serve as useful points of comparison for the heuristics.

5.3. Randomized Study Results

We varied the number of locations (three, four, and five) and trauma type (blunt and penetrating), and then randomly generated 300 problem instances for each combination of number of locations and trauma type as described in Section 5.1. In the randomized study, we generated locations of the incidents and the facilities within the urban area by a two-dimensional Uniform random variate and the distance to the outlying facility by a Uniform random variate. Generated trauma centers are either Level I (with probability 0.58), with median 39 ED beds or Level II (with probability 0.42) with median 25 ED beds. These probabilities for Level I and II trauma centers are obtained from a national survey by MacKenzie and Hoyt (2003), while the numbers of beds are taken from Rivara et al. (2006). We added variability to the number of beds for each trauma center according to a Uniform random variate with range 10, centered about the median. Utilization of ED beds at a trauma center is very high during normal operations (MacKenzie and Hoyt 2003), and beds for patients from a disaster would become available over a period of time as patients in the ED would be moved to create surge capacity (Hick et al. 2004, Peleg and Kellermann 2009). In our simulation, 10% of beds were available initially, and an additional 30% of them

became available over a two-hour period as existing patients were discharged or moved to other facilities.

All simulations were written in C++ and run using 100 replications and common random numbers across policies. In Table 1, we report the minimum (“Min”), first quartile (“Q1”), median (“Med”), third quartile (“Q3”), and maximum (“Max”) percentage improvement in discounted number of patients treated by the Myopic and PIH policies versus the two baseline policies across the 300 scenarios. We also report the number of instances (out of 300) where each proposed heuristic is better at the 0.05 statistical significance level (“# Sig”).


Table 1 shows that the heuristics resulting from our analysis were valuable. Both heuristics were substantially better than BLS in almost all instances, since choosing the nearest hospital completely ignores congestion. Furthermore, both heuristics provided a statistically significant improvement over BLD in most instances, albeit by a smaller amount. The BLD policy uses state information to dynamically assign patients to hospitals, incorporating both travel times and congestion at the hospital, and hence is expected to be competitive. In Table 1, PIH and Myopic policies appear to behave similarly, except under certain performance measures and injury types. When patients have blunt injuries, PIH provides a statistically better improvement over the baseline policies in a larger number of instances than Myopic. On the other hand, Myopic has better worst-case performance under penetrating injuries, especially for a small number of locations. We next examine the simulation results more closely to determine whether certain characteristics of problem instances predict the observed performance improvement and the difference between PIH and Myopic. We present results for the penetrating injury type here; results for the blunt type are presented in Section EC.5 of the electronic companion.

All dynamic heuristics considered (Myopic, PIH, and BLD) improve patient distribution as compared

Table 1. Simulation Results: Randomized Study

Trauma type	No. of locations	Policy	vs. Baseline static						vs. Baseline dynamic					
			Min (%)	Q1 (%)	Med (%)	Q3 (%)	Max (%)	# Sig	Min (%)	Q1 (%)	Med (%)	Q3 (%)	Max (%)	# Sig
Blunt	3	PIH	-11	37	48	67	246	298	-23	2	4	7	28	259
		Myopic	0	31	42	66	231	300	-12	-1	1	3	23	189
	4	PIH	0	38	48	63	241	299	-9	3	4	6	18	292
		Myopic	1	33	43	58	235	300	-7	0	2	3	12	230
	5	PIH	0	17	45	56	242	300	0	3	4	6	11	300
		Myopic	0	16	41	51	226	299	-9	1	2	4	8	261
Penetrating	3	PIH	-60	6	28	51	201	239	-54	0	5	12	76	215
		Myopic	0	20	35	56	192	300	-13	3	11	26	131	261
	4	PIH	-35	19	39	55	223	284	-20	3	6	11	53	259
		Myopic	2	23	36	54	213	300	-7	3	6	12	90	268
	5	PIH	-14	17	38	51	215	297	-13	4	7	10	40	279
		Myopic	1	19	35	50	200	300	-7	3	5	10	52	281

Figure 3. (Color online) Boxplots of Improvement over the BLS (Nearest Facility) Policy by Number of Facilities Used Under the Nearest Facility Policy (Penetrating Injury)



to BLS (nearest facility) by striking a balance between transporting patients to a nearby facility (which uses transportation resources efficiently) and transporting patients to a facility with a short wait (which uses treatment resources efficiently). This improvement compared to the BLS is highly dependent on the geographic distribution of the locations and facilities, specifically, by how many facilities are ever used under BLS. We found that when all locations share a single nearest facility, dynamic policies provide over 100% improvement on the median problem instance, while when all three facilities are used by the nearest facility policy, the performance improvement is usually much smaller though still significant, with the median around 10% (see Figure 3, which shows the difference between the dynamic heuristics and the nearest facility policy in the randomized study). Controlling for this effect, we see that when the location distribution is geographically unbalanced (i.e., not all facilities are used in the nearest facility policy), the performance improvement is also strongly affected by the distance to the facilities. Dynamic policies perform better compared to BLS when the average distance to the farthest facility is shorter (see Figure 4). For penetrating injuries, one extra mile distance is associated with a reduction of approximately 3.0 (2.1) percentage points in performance improvement for PIH over BLS when 1 (2) facilities are used under the nearest facility policy. The distance effect appears to be insignificant when all three facilities are used by the nearest facility policy. We also observe that BLD is the most affected by the mean distance among all three dynamic policies.


Examining Table 1 and Figures 3 and 4, we observe that Myopic and PIH perform somewhat similarly in general. However, in individual instances, the relative

performance of these two policies may vary widely. In Figure 5, we see that PIH performs better than Myopic when the average distance to the nearest facility is small, but Myopic performs better when the average distance to the nearest facility is large. This effect becomes more pronounced with fewer locations (in fact, it is statistically insignificant for blunt injuries when there are five locations). As the mean distance to the nearest facility becomes larger, scene waiting time takes a greater proportion of the total delay as compared to hospital waiting time. Moreover, because the number of locations is a proxy for the number of patients (since the number of patients is drawn from the same distribution for each location), when there are fewer locations (and thus fewer patients), hospital waiting time makes up a smaller proportion of the total delay. A similar conclusion can be made when PIH is compared with BLD (see Figure 6). This observation makes sense because PIH uses a more sophisticated model of queue waiting for future patients, while Myopic and BLD consider the waiting time of only the current patient. This effect is much stronger for the penetrating injury type than the blunt injury type. An intuitive explanation would be that when patients deteriorate rapidly (as in the case of penetrating injuries) and the travel times are longer, discounted rewards from future patients are smaller, and hence, a policy that acts myopically with respect to the current patient would be expected to perform well.

5.4. Robustness Checks

We conducted four robustness checks to verify that the main insights of the simulation study hold over a variety of different scenarios. One of the benefits of our proposed heuristics is that they can be applied in

Figure 4. (Color online) Improvement over BLS (Nearest Facility) Policy by Distance to Farthest Facility (Penetrating Injury)




Notes. Lines labeled as “Fit” are computed using least-squares regression. Fitted line slope is statistically significant at the 0.05 level for all cases with one or two unique nearest facilities.

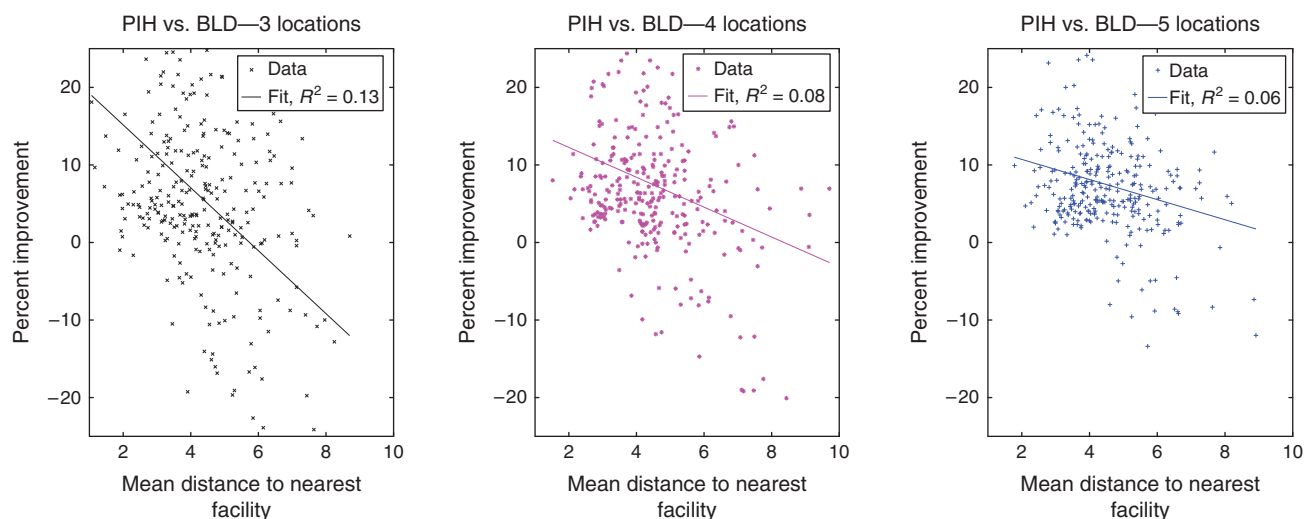
scenarios with limited state updates, as in Section 4.3. As part of our robustness checks, we repeated the entire randomized study with state updates according to a Poisson process, with four, two, or zero updates per hour, with all dynamic policies using the same state information. Table 2 shows the results when there are no state updates after time zero (details of this

experiment and additional results for intermittent state updates are given in Section EC.4 in the electronic companion). We observed that the Myopic heuristic and PIH are both robust to limited state updates. Specifically, the performance of the two proposed policies against the BLS policy is slightly worse compared to Table 1, but the performance against the baseline

Figure 5. (Color online) Percent Improvement Using PIH vs. Myopic by Distance to Nearest Facility (Penetrating Injury)



Notes. Lines labeled as “Fit” are computed using least-squares regression. Fitted line slope is statistically significant at the 0.05 level for all cases.

Figure 6. (Color online) Percent Improvement Using PIH vs. BLD by Distance to Nearest Facility (Penetrating Injury)

Notes. Lines labeled as “Fit” are computed using least-squares regression. Fitted line slope is statistically significant at the 0.05 level for all cases.

dynamic policy is better, which implies that BLD is not as robust to limited state updates. In particular, the median improvement of PIH over BLD is two to three times as high in the absence of state updates than with perfect information, indicating a degradation in the performance of the BLD policy when accurate state information is not available.

We also repeated the randomized experiments by doubling the number of flexible and dedicated ambulances and observed that both Myopic and PIH provide even slightly better performance than in the case with the original setup (see Table EC.7 in the electronic companion). One difference was that with more ambulances, PIH provided consistently better performance than Myopic in terms of the number of instances that had significant improvement (see Figure EC.5 in the electronic companion). This result is similar to our previous observation that PIH performed better than

Myopic when travel times were relatively short, and suggests that overall transportation capacity (depending on the number of ambulances and travel time distributions) is a factor in choice of a preferred heuristic. To study the effect of travel time variability, we repeated the entire randomized study, reducing the coefficient of variation of travel times by half, which did not change any of the results in a meaningful way (see Table EC.8 in the electronic companion). We finally repeated the same suite of simulation experiments presented in Section 5.3 with a small number of patients, uniformly drawn over $\{5, 6, \dots, 15\}$ per location. With a smaller number of patients, there is less room for improvement because the system is not so congested (see Table EC.9 in the electronic companion). However, we still observed substantial improvement except in the case of blunt injury, where Myopic did not result in significant improvement over BLD in the majority

Table 2. Simulation Results for Modified Heuristics with No Information Updates After Time Zero

Trauma type	No. of locations	Policy	vs. Baseline static						vs. Baseline dynamic					
			Min (%)	Q1 (%)	Med (%)	Q3 (%)	Max (%)	# Sig	Min (%)	Q1 (%)	Med (%)	Q3 (%)	Max (%)	# Sig
Blunt	3	PIH	-8	32	41	67	214	290	-8	4	12	21	58	249
		Myopic	-5	16	31	52	195	290	-27	-5	1	8	47	161
	4	PIH	-6	32	41	55	233	278	-9	2	10	19	57	238
		Myopic	-7	13	29	48	209	289	-26	-4	2	8	43	171
	5	PIH	-8	10	39	50	228	253	-8	1	9	18	49	226
		Myopic	-8	8	20	42	172	282	-23	-5	2	7	36	180
Penetrating	3	PIH	-61	15	26	46	186	272	-52	3	11	18	60	240
		Myopic	-5	11	24	43	173	291	-20	0	7	16	61	217
	4	PIH	-7	21	32	51	211	283	-10	6	14	23	61	266
		Myopic	-8	10	23	44	182	294	-18	-1	7	17	47	211
	5	PIH	-12	13	33	46	206	265	-17	7	16	26	62	248
		Myopic	-9	7	17	36	150	285	-28	-1	6	15	47	211

of instances. This suggests that when the number of patients is small, careful design of a dynamic policy (such as PIH) is required to achieve improvement over a simple baseline policy.

6. Case Study: Earthquake in San Francisco, CA

To demonstrate the effectiveness of our heuristic policies and to verify the insights gained from our randomized study in Section 5, we developed a case study using HAZUS-MH software, which is published by the U.S. government (Federal Emergency Management Agency 2016). HAZUS-MH uses publicly available data along with geographic information to estimate direct and indirect losses from hazards including earthquakes, hurricanes, and floods, and maps the results on a geographic information system (GIS) platform. For example, the map shown in Figure 7 was generated by the HAZUS-MH software when running our case study problem instance.


We chose San Francisco, California, for our case study for several reasons: (i) San Francisco is prone to earthquakes, which can be modeled with HAZUS-MH; (ii) San Francisco has 10 hospitals with EDs that are capable of treating earthquake patients (making the decision problem highly complex); and (iii) high-quality data about ED capacities and capabilities is available from the State of California (California Office of Statewide Health Planning and Development 2016).

We used HAZUS-MH to generate casualty estimates for 194 census tracts in San Francisco, assuming an

earthquake of magnitude 7.0 occurred 60 miles south of the city (at the approximate location of a major earthquake that occurred in 1989) on a weekday at 2 P.M. The scenario generated 488 critically injured patients who would require transportation by ambulance. To generate travel times, we assumed the patients are clustered at the geographic center of each census tract, as identified in the GIS software. We used the Google Maps application programming interface (API) to obtain best-case, typical, and worst-case estimates of the travel time from each cluster to each of the ten hospitals on a weekday at 2 P.M. Since a more detailed travel time distribution was not available, the travel times generated from Google Maps were interpreted as parameterizing a triangular distribution in accordance with common practice in simulation studies (Law 2007, chapter 6). We determined the hospital ED bed capacities using the CA-OSHPD data. Finally, we determined that the San Francisco Fire Department has 56 ambulances suitable for patient transport (California Emergency Medical Services Agency 2018). In the first simulation, we assumed all 56 ambulances were flexible; in subsequent simulations, we dedicated 10, 20, and 30 ambulances to randomly selected locations to assess the effect of ambulance flexibility. Other parameters not discussed here were handled as in Section 5 with the blunt injury mechanism since blunt injuries are common in earthquakes (Gautschi et al. 2008), and each simulation was repeated for 100 replications.

Table 3 shows that the proposed heuristics (PIH and Myopic) lead to an increased expected number of survivors over the baselines in a large-scale incident

Figure 7. (Color online) Distribution of Patients at Residential Buildings Under an Earthquake Scenario, City of San Francisco



Source. HAZUS-MH software (Federal Emergency Management Agency 2016).

Notes. Darker colors indicate census tracts with greater numbers of patients. Black dots indicate hospitals with EDs suitable for receiving patients.

Table 3. Comparison of Myopic Heuristic and PIH to the BLS and BLD Policies for the San Francisco Earthquake Scenario

No. of dedicated ambulances	Expected number of survivors (out of 488)				% improvement by			
					Myopic vs.		PIH vs.	
	Myopic	PIH	BLS	BLD	BLS	BLD	BLS	BLD
0	358	362	336	244	6	47	8	48
10	343	346	323	229	6	50	7	51
20	319	323	303	205	5	55	7	57
30	287	291	276	173	4	66	6	69

Note. All differences between policies are statistically significant at the 0.05 level.

with many locations and facilities. One interesting result from the case study is that the improvement is greater over the BLD policy than over the BLS policy, a contrast to the randomized study. The case study has two features that suggest that BLS, which sends patients to the nearest facility, may not perform poorly. First, there are a large number of facilities that are geographically proximate to the highly populated areas of the city. Second, an earthquake is the type of disaster that generates moderate numbers of patients over a widespread area, rather than large number of patients concentrated at a small number of locations (which might occur, for example, in a coordinated terrorist event). This means that the patients themselves are geographically dispersed. Thus, sending each patient to the nearest facility is unlikely to overwhelm one facility while leaving another idle. This observation is consistent with Figure 3, where we found more limited improvement over BLS when many facilities are used. Nonetheless, despite the fact that BLS would be expected to work well in such a scenario, the PIH still provides an average of 5%–8% improvement, equivalent to 15–26 lives, over BLS depending on the number of dedicated ambulances, which is both statistically significant and meaningful.

As expected, fixing the location of additional ambulances always results in a decrease in performance of all policies, but the effects are different when measured against different baselines. In general, we see that the dynamic policies (proposed and baseline) leverage flexibility more than BLS, but the two proposed heuristics continue their robust performance when more ambulances are dedicated, while BLD is less robust. This result is similar to our observation in Section 5 that BLD is less robust to intermittent state updates than PIH and Myopic.

7. Conclusion

In the aftermath of a disaster, distributing patients to medical facilities where they can receive care is a major operational challenge. In this paper, we formulated the patient-distribution problem as an MDP, which led to

several analytical and numerical results that could be used in designing emergency response plans.

We provide strong evidence that dynamic policies considering both the scene waiting time and the hospital waiting time can substantially improve the outcome of a response effort compared with simply distributing patients to the geographically nearest facility. However, this information must be used carefully. Our work suggests that it is especially important to use such dynamic policies that take into account changing hospital congestion levels *when responding to disasters with multiple incident locations that share one or two nearest facilities but where none of the hospitals involved in the response effort are too far away.*

The dynamic policies proposed in this study (PIH and Myopic) perform very well in scenarios considered in our simulation study. On the basis of observations from the tested scenarios, we make the following recommendations:

1. To increase the expected number of patients saved, we recommend at the very least using the Myopic policy, which is straightforward to implement and which performed better than the nearest facility rule in almost all scenarios we tested.

2. For events where patients do not deteriorate very rapidly, such as events resulting in blunt trauma, further improvement may be achieved via the more sophisticated PIH, especially when transportation is not the bottleneck (because the round-trip travel time to hospitals is not too long or because there are many ambulances) and the patients are spread over many incident locations. EMS providers are best positioned to use their region-specific knowledge to determine whether a particular incident meets these criteria. In the paper, we provide one such exemplary and realistic scenario that involves a hypothetical earthquake that takes place in San Francisco, California.

The heuristics presented in this paper could be incorporated into training exercises to help emergency responders better understand their capability to respond to specific disaster scenarios. The observations and results of this paper could be also used to develop rules of thumb as part of an emergency response plan.

For example, the heuristics can be used to determine the types of events for which distant hospitals should be included in the response. Finally, we believe that Myopic and PIH are simple enough to be implemented in a decision support tool that could assist in real-time response, provided at least some information on the hospital congestion levels can be obtained. Although a simple calculation would need to be performed, given the ubiquity of portable technology such as mobile phones, this could easily be automated. Moreover, such devices would not necessarily require pervasive network connectivity as we observed that our heuristic policies are effective even with limited state updates. Our policies are tailored to the situation with many patients arriving at once. However, similar ideas could be used in future work to develop patient-distribution policies for daily emergencies with patients arriving over a period of time in small numbers.

Although we demonstrated our policies on a realistic simulated case study, further experimentation with retrospective disaster data would increase understanding of the impact of changes in the way patients are distributed. Therefore, this article also highlights the need for improved communication and reporting in the aftermath of a disaster.

Acknowledgments

The authors thank J. E. Winslow, MD, L. M. Smith, MD, and B. Hiestand, MD, for providing valuable advice from a medical perspective, and B. Courtney, MPH, T. Stephens, D. Fletcher, J. Mugele, MD, C. Rochelle, and J. S. Sargent, who shared their expert knowledge about emergency response, planning, and coordination, during interviews with the first author. The information gained from those interviews improved this article. K. Mayo provided assistance with setting up the HAZUS software. The authors also thank the associate editor and anonymous referees for their comments on an earlier version of this article, which helped to significantly improve it.

References

Allon G, Deo S, Lin W (2013) The impact of size and occupancy of hospital on the extent of ambulance diversion: Theory and evidence. *Oper. Res.* 61(3):544–562.

American College of Surgeons (2012) National Trauma Data Bank. Accessed April 30, 2018, <https://www.facs.org/quality-programs/trauma/ntdb>.

Ansell PS, Glazebrook KD, Kirkbride C (2003) Generalised “join the shortest queue” policies for the dynamic routing of jobs to multi-class queues. *J. Oper. Res. Soc.* 54(4):379–389.

Argon NT, Ding L, Glazebrook KD, Ziya S (2009) Dynamic routing of customers with general delay costs in a multiserver queueing system. *Probab. Engrg. Inform. Sci.* 23(2):175–203.

Armony M (2005) Dynamic routing in large-scale service systems with heterogeneous servers. *Queueing Systems* 51(3–4):287–329.

California Emergency Medical Services Agency (2018) EMS Systems Division—Transportation. Accessed April 30, 2018, <https://emsa.ca.gov/transportation/>.

California Office of Statewide Health Planning and Development (2016) Office of Statewide Health Planning and Development. Accessed November 28, 2016, <http://oshpd.ca.gov/>.

Centers for Disease Control and Prevention (2010) Interim planning guidance for preparedness and response to a mass casualty event resulting from terrorist use of explosives. Accessed December 15, 2016, <https://stacks.cdc.gov/view/cdc/5705/>.

Dean MD, Nair SK (2014) Mass-casualty triage: Distribution of victims to multiple hospitals using the SAVE model. *Eur. J. Oper. Res.* 238(1):363–373.

Deo S, Gurvich I (2011) Centralized vs. decentralized ambulance diversion: A network perspective. *Management Sci.* 57(7):1300–1319.

Erkut E, Ingolfsson A, Erdoğan G (2008) Ambulance location for maximum survival. *Naval Res. Logist.* 55(1):42–58.

Federal Emergency Management Agency (2016) HAZUS software. Accessed November 28, 2016, <https://www.fema.gov/hazus-software>.

Gautschi OP, Cadosch D, Rajan G, Zellweger R (2008) Earthquakes and trauma: Review of triage and injury-specific, immediate care. *Prehospital Disaster Medicine* 23(2):195–201.

Glazebrook KD, Kirkbride C, Ouenniche J (2009) Index policies for the admission control and routing of impatient customers to heterogeneous service stations. *Oper. Res.* 57(4):975–989.

Gong Q, Batta R (2007) Allocation and reallocation of ambulances to casualty clusters in a disaster relief operation. *IIE Trans.* 39(1):27–39.

Hick JL, Hanfling D, Burstein JL, DeAtley C, Barbisch D, Bogdan GM, Cantrill S (2004) Health care facility and community strategies for patient care surge capacity. *Ann. Emergency Medicine* 44(3):253–261.

Hordijk A, Koole G (1992) On the assignment of customers to parallel queues. *Probab. Engrg. Inform. Sci.* 6(4):495–511.

Ingolfsson A, Budge S, Erkut E (2008) Optimal ambulance location with random delays and travel times. *Health Care Management Sci.* 11(3):262–274.

Jotshi A, Gong Q, Batta R (2009) Dispatching and routing of emergency vehicles in disaster mitigation using data fusion. *Socio-Econom. Planning Sci.* 43(1):1–24.

Krishnan KR (1990) Joining the right queue: A state-dependent decision rule. *IEEE Trans. Automatic Control* 35(1):104–108.

Larson RC, Metzger MD, Cahn MF (2006) Responding to emergencies: Lessons learned and the need for analysis. *Interfaces* 36(6):486–501.

Law AM (2007) *Simulation Modeling and Analysis* (McGraw Hill, New York).

Lippmann S (1975) Applying a new device in the optimization of exponential queueing systems. *Oper. Res.* 23(4):687–710.

MacKenzie KJ, Hoyt DB (2003) National inventory of hospital trauma centers. *J. Amer. Medical Assoc.* 289(12):1515–1522.

Mandelbaum A, Stolyar AL (2004) Scheduling flexible servers with convex delay costs: Heavy-traffic optimality of the generalized $c\mu$ -rule. *Oper. Res.* 52(6):836–855.

Maxwell MS, Restrepo M, Henderson SG, Topaloglu H (2010) Approximate dynamic programming for ambulance redeployment. *INFORMS J. Comput.* 22(2):266–281.

McLay LA, Mayorga ME (2013) A dispatching model for server-to-customer systems that balances efficiency and equity. *Manufacturing Service Oper. Management* 15(2):205–220.

Mills AF, Helm JE, Jola-Sanchez AF, Tatikonda MV, Courtney BA (2018) Coordination of autonomous healthcare entities: Emergency response to multiple casualty incidents. *Production Oper. Management* 27(1):184–205.

Peleg K, Kellermann AL (2009) Enhancing hospital surge capacity for mass casualty events. *J. Amer. Medical Assoc.* 302(5):565–567.

Puterman ML (1994) *Markov Decision Processes: Discrete Stochastic Dynamic Programming* (John Wiley & Sons, Hoboken, NJ).

Rivara FP, Nathens AB, Jurkovich GJ, Maier RV (2006) Do trauma centers have the capacity to respond to disasters? *J. Trauma* 61(4):949–953.

Sacco WJ, Navin DM, Fiedler KE, Waddell RK II, Long WB, Buckman RF Jr (2005) Precise formulation and evidence-based application of resource-constrained triage. *Acad. Emergency Medicine* 12(8):759–770.

- Sacco WJ, Navin DM, Waddell RK, Fiedler KE, Long WB, Buckman RF (2007) A new resource-constrained triage method applied to victims of penetrating injury. *J. Trauma* 63(2):316–325.
- Sasser SM, Hunt RC, Faul M, Sugerman D, Pearson WS, Dulski T, Wald MM, et al. (2012) Guidelines for field triage of injured patients: Recommendations of the National Expert Panel on Field Triage, 2011. *Morbidity Mortality Weekly Report* 61(1):1–20.
- van Mieghem JA (1995) Dynamic scheduling with convex delay costs: The generalized c|mu rule. *Ann. Appl. Probab.* 5(3):809–833.
- Whitt W (1986) Deciding which queue to join: Some counterexamples. *Oper. Res.* 34(1):55–62.
- Yue Y, Marla L, Krishnan R (2012) An efficient simulation-based approach to ambulance fleet allocation and dynamic redeployment. Hoffmann J, Selman B, eds. *Proc. 26th AAAI Conf. Artificial Intelligence* (AAAI Press, Palo Alto, CA), 398–405.
- Zayas-Cabán G, Lewis ME, Olson M, Schmitz S (2013) Emergency medical service allocation in response to large scale events. *IIE Trans. Healthcare Systems Engrg.* 3(1):57–68.

Alex F. Mills is assistant professor of operations and decision technologies in the Kelley School of Business at Indiana University. His research interests are in service systems and healthcare operations, with a focus on healthcare system response to disruptions, demand management in healthcare services, and management of healthcare providers.

Nilay Tanik Argon is associate professor of statistics and operations research at the University of North Carolina. Her research interests are in stochastic modeling of manufacturing and service systems, queueing systems, healthcare operations, and statistical output analysis for computer simulation.

Serhan Ziya is professor of statistics and operations research at the University of North Carolina. His research interests are in service operations, with a focus on healthcare operations, queueing systems, revenue management, and pricing under inventory considerations.