Omega xxx (xxxx) xxx

Contents lists available at ScienceDirect

Omega

journal homepage: www.elsevier.com/locate/omega

Assigning spontaneous volunteers to relief efforts under uncertainty in task demand and volunteer availability☆

Kyle E. Paret[a], Maria E. Mayorga[a,*], Emmett J. Lodree[b]

[a] Department of Industrial and Systems Engineering, North Carolina State University, Campus Box 7906, Raleigh, NC 27695, USA
[b] Department of Information Systems, Statistics and Management Science, Culverhouse College of Business, University of Alabama, Tuscaloosa, AL 35487, USA

ARTICLE INFO

Article history:
Received 28 February 2019
Accepted 15 February 2020

Keywords:
Volunteer management
Disaster operations management
Spontaneous volunteer
Markov decision process
Discrete event simulation

ABSTRACT

In the wake of a disaster, people from nearby areas often converge to assist the affected community. Spontaneous volunteers are not affiliated with relief agencies but are in a unique position to provide invaluable aid at a crucial point in the disaster cycle. Often, these volunteers are ineffectively used or refused altogether. Volunteer Reception Centers (VRCs) can benefit from improved strategies to integrate the influx of spontaneous volunteers. In this paper, a multi-server queuing model is formulated to represent the dynamics of assigning spontaneous volunteers to tasks in a post-disaster setting. In particular, we consider the case of stochastic arrival of demand for service and stochastic arrival of volunteers, whose time in service is also stochastic. These assumptions mimic disaster relief tasks such as distribution of relief items, where both beneficiaries and volunteers arrive randomly. An optimal policy for assigning volunteers to tasks is generated using a Markov Decision Process. We then use simulation to compare the optimal policy against several heuristic policies and discuss real world implications.

1. Introduction

During and after major disasters, large numbers of people unaffiliated with traditional emergency response organizations converge on the scene to offer assistance [1]. Motivated by the desire to do something for those in need, these individuals respond on impulse immediately following disaster events and are referred to as *spontaneous volunteers* [2]. Spontaneous offers to help during and following disasters are well documented, and are strongly influenced by the amount of media coverage an event receives [3]. For example, Lowe and Fothergill [2] report that 15,000 volunteers helped during the two-and-a-half week period following the September 11 attacks in New York City. The response to Hurricane Katrina in 2005 attracted 60,000 volunteers to New Orleans [4].

Research shows that spontaneous volunteers are capable of positively contributing to relief efforts during the aftermath of disasters by performing a variety of services, including search and rescue, distribution of relief items, and the assessment of community needs [5]. While spontaneous volunteers can be a valuable resource, they are often ineffectively used and can potentially hinder emergency operations by creating health, safety, and security concerns. Furthermore, spontaneous volunteers require supervision, which can distract professional responders from their duties that directly serve disaster survivors [1]. Oftentimes, the services of spontaneous volunteers are refused purely because volunteer organizations are ill-equipped to manage them. A survey of non-governmental voluntary organizations (NVOs) found that the use of spontaneous volunteers is widespread, but NVOs are not necessarily structured to effectively engage them [44]. Improved strategies for incorporating spontaneous volunteers into organized relief efforts are needed in order to achieve safe and responsive disaster management [1].

In this paper, we consider the problem of assigning spontaneous volunteers arriving at a Volunteer Reception Center (VRC) to multiple disaster relief tasks. Such volunteers are commonly utilized to sort donations and distribute relief supplies (e.g. food, clothing, clean-up kits). One characteristic that distinguishes spontaneous volunteer assignment from all other forms of labor scheduling is that spontaneous volunteers randomly join and abandon the operation [6]. As such, we represent the spontaneous volunteer assignment problem as a multi-server parallel queuing system where the servers (spontaneous volunteers) randomly arrive and depart the system. Within this framework, optimal policies for assigning spontaneous volunteers to tasks are derived from a continuous time Markov Decision Process (MDP) model. In addition, we evaluate a variety of experimental cases in a discrete-event simulation model in order to examine the performance

☆ This manuscript was processed by Associate Editor Prof. Xiuli He.

* Corresponding author.
E-mail addresses: keparet@ncsu.edu (K.E. Paret), memayorg@ncsu.edu (M.E. Mayorga), ejlodree@cba.ua.edu (E.J. Lodree).

https://doi.org/10.1016/j.omega.2020.102228
0305-0483/

of heuristic assignment policies relative to the optimal policy. We identify an effective assignment policy that is also easy to implement, and provide insights for individuals who may be in charge of managing spontaneous volunteers in real-world disaster response environments.

The remainder of this paper is organized as follows. In Section 2, we review academic literature related to disaster operations management, volunteer scheduling, and server assignment in queuing theory. In Section 3, we formulate the spontaneous volunteer assignment problem as a Markov Decision Process (MDP), and discuss sufficient conditions for the existence of steady-state solutions. In Section 4, we introduce and describe practical heuristic policies to be tested against the MDP policy. Section 5 provides a brief overview of the simulation model used to compare the performance of the MDP and heuristic policies. Section 6 details the experimental analysis generated for the simulation study and the results. Concluding remarks and implications for volunteer managers are provided in Section 7.

2. Literature review

This study is related to three areas: (1) disaster operations management, (2) volunteer scheduling, and (3) server assignment in queuing systems.

2.1. Disaster operations management

The disaster operations management (DOM) literature has experienced rapid growth since the year 2001. For example, Gupta et al. [7] surveyed 268 papers from among 25 Operations Research / Management Science (OR/MS) journals and reported an increase in the annual publication rate of 2.67 in 2001 to 33.67 in 2014 (these results are based on a three-year moving average). It is safe to say that DOM is a mainstream application area of OR/MS based on the numerous survey papers dedicated to the subject (e.g. [8–10]). Moreover, multiple focus areas with critical mass have emerged from DOM literature, and review papers dedicated to these subjects have also been published. The most prevalent DOM topics include: inventory management [11], facility location [12], and relief distribution [13]. Manpower planning, however, is another important topic within the context of DOM that has received very limited attention from an academic perspective [14,15]. Fritz and Mathewson [16] first defined the mass movement of people and supplies to the affected area as volunteer and material convergence respectively. Volunteer Management focuses on managing these volunteer resources through the entire process. Volunteer management plans cover the engagement, recruitment, placement, orientation, training, supervision, recognition, and evaluation of volunteers. Many software solutions exist (e.g. volgistics, EveryAction, MobileServe, etc.) to help volunteer managers track and schedule registered volunteers. These software solutions are also useful after disasters due to their reporting and database capabilities. However, they do not necessarily assist with the immediate assignment of incoming spontaneous volunteers following a natural disaster. Many of the current management practices related to SV, as highlighted by FEMA [17] and Red Cross [18], focus on strategic guidelines. Where operational guidelines do exist, they fail to elaborate on key points. For example, one of FEMA's concepts of operation suggests volunteer managers should "Refer unaffiliated volunteers to appropriate response agencies after initial screening" [17]. However, the referral is left to the best judgement of the volunteer manager on duty at the volunteer reception center (VRC) or volunteer site. The manager must interpret the meaning of *appropriate* as it relates to the situation, using their best judgement and experience. Prioritizing the tasks of SVs is considered to be one of the major challenges for volunteer managers during disaster response [19]. This paper addresses a known gap in the DOM literature by considering the problem of assignment/placement for spontaneous volunteers after disaster events from an operational perspective.

2.2. Volunteer scheduling

While traditional workforce scheduling problems have been studied extensively (e.g. [20]), there has been much less work that focuses on labor assignment from a DOM perspective. One of the most impactful distinctions between traditional and DOM manpower planning is the role of volunteer labor. In particular, volunteers complete a significant portion of the tasks performed during the early phases of disaster response [5]. The labor in classical personnel scheduling is supported by a paid workforce, which tends to be more stable and predictable compared to volunteer labor [21]. Falasca and Zobel [22] identify several operational characteristics that are unique to volunteer planning and scheduling for disaster relief purposes that ought to be considered from an OR/MS modeling perspective. One such characteristic is the importance of satisfying volunteers' preferences, including the type of work, the times they want to work, how long they are willing to work, and with whom they prefer to work (volunteers often contribute to relief efforts in groups, e.g., churches, sports teams, families). The issue of prioritizing volunteer preferences applies to volunteer scheduling in general (e.g. [21]), not just volunteer scheduling for disaster relief. However, manpower planning for DOM may also require assigning volunteers across geographically dispersed locations, which for the most part, is irrelevant when it comes to labor scheduling in non-DOM contexts. Falasca and Zobel [22] incorporate the above features into a bi-objective optimization model that seeks to balance the conflicting objectives of minimizing unmet task demands and maximizing volunteer preferences. Lassiter et al. [23] also take unmet task demands and volunteer preferences into account within the context of humanitarian relief, but extend Falasca and Zobel's deterministic model in an important way: they consider task uncertainty and propose a robust optimization model to handle this uncertainty.

In both Falasca and Zobel [22] and Lassiter et al. [23], there is no uncertainty associated with labor capacity or capabilities. However, empirical studies have demonstrated that spontaneous volunteerism is characterized by various forms of uncertainty, such as the times between volunteer arrivals, the amount of time volunteers contribute the relief efforts on a given day, and the sizes of volunteer groups who arrive and depart the relief effort together [6]. From this perspective, we view the studies by Falasca and Zobel [22] and Lassiter et al. [23] as appropriate for scheduling affiliated volunteers where an uncertain labor pool is significantly less of an issue. However, assignment decisions that pertain to spontaneous volunteers, the volunteer type we focus on in this paper, require a different approach. Mayorga et al. [24], which is the study that is closest to ours, proposes a framework that captures the unique characteristics of spontaneous volunteers, namely uncertainty in volunteer arrival and departure times. Specifically, they model the spontaneous volunteer assignment problem as a parallel queuing system with random server (i.e., volunteer) arrivals and abandonments, in which a deterministic amount of work known *a priori* is to be completed. The control problem is formulated as a MDP, and a policy iteration algorithm is used to generate optimal policies for problem instances. Mayorga et al. [24] also conduct a computational experiment in which the performance of several practical heuristic policies are examined through simulation, and they find that simply assigning volunteers to the queue with fewest volunteers generally works well as an assignment policy. This paper generalizes the deterministic demand model by considering stochastic demand streams, allowing for the representation of more complex disaster relief tasks. By doing so, we now have to contend with deriving an appropriate stability condition for the queuing system (done in Section 3.1), which was not necessary for the deterministic demand case described above.

In summary, we have identified only three papers in the academic DOM literature that address manpower planning for disaster relief from an operational perspective: [22–24]. Furthermore, only one of the three considers random volunteer arrival and abandonment processes [24], and is therefore relevant to the spontaneous volunteer assignment problem. We conclude our discussion here by noting that volunteer scheduling literature has occurred in contexts other than DOM. A general framework for volunteer scheduling is laid out by Sampson [21], and he applies that framework to the problem of scheduling reviewer assignments for an academic conference. Other settings in which volunteer labor assignment research has been applied are an annual music festival [25] and a bike sharing program [26]. However, these applications do not consider labor uncertainty.

2.3. Queuing theory

Optimal control of queues via server assignment is a widely studied class of problems in queuing theory with numerous variations. As such, a comprehensive review of this literature is beyond the scope of our discussion here. Instead, we review a few representative studies that focus on the control of parallel queuing systems, and we also highlight related areas of the queuing literature. A basic framework for server assignment in parallel queuing systems involves multiple customer classes, where each queue is dedicated to serving a specific class. The control problem entails dynamically allocating a fixed pool of heterogeneous servers among the queues with the goal of minimizing waiting time. Squillante et al. [27], for example, consider this problem under the assumption of Poisson arrivals and exponential service times, and propose threshold-based priority policies related to the well-known $c\mu$ rule.

Variations of the $c\mu$ policy have since been applied to many different queuing systems. For example, a generalized version of the $c\mu$ rule is optimal when holding costs are nonlinear, but convex [28]. More recently, Cao et al. [29] derived conditions under which the $c\mu$ policy is optimal for a single server queuing network with two customer classes, where customers from the lower class are upgraded to the higher class after a random amount of time. There are also some cases where the optimality of $c\mu$ is preserved for parallel queuing systems with multiple servers. Consider, for instance, the N-network that consists of two types of customers and servers; a dedicated server that can serve one customer type, and a fully flexible server that is capable of serving both customer types. Bell et al. [30] and Saghafian et al. [31] prove that the $c\mu$ rule is optimal under certain conditions in this setting, while Down et al. [32] extend the result to an N-Network with multiple servers of each type (dedicated and fully flexible) and customer upgrades.

All of the above-mentioned studies involve the assignment of a fixed pool of servers where, as is the case in most papers, service rate is the only stochastic characteristic of the servers. However, the present study considers server assignment policies where not only service rates, but also server availability, are stochastic. Besides Mayorga et al. [24], only a limited number of papers deal with optimal control of queues in the presence of unreliable servers; and to our knowledge, only two do so within the context of assigning servers to parallel queues: [31,33]. Both of these papers derive optimal policies for assigning unreliable servers to queues involving multiple customer classes, where servers are unreliable in the sense that they fail at random points in time and remain unavailable for random periods until they are repaired. Saghafian et al. [31] establish conditions under which a version of the $c\mu$ rule is optimal for a generalization of the N-network that

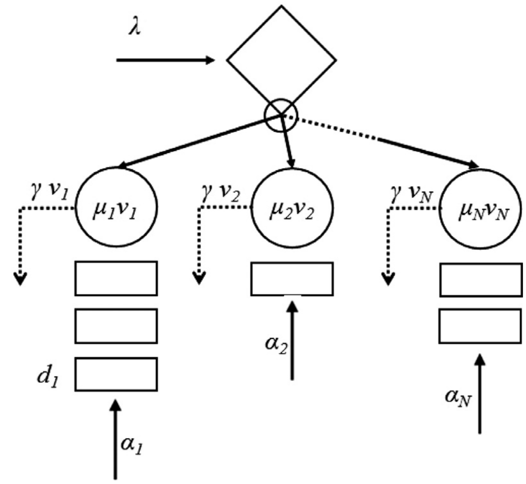

Fig. 1. Pictorial representation of the system.

has three servers (instead of two): one dedicated, one fully flexible, and one partially flexible (the authors refer to this extension as the W-network). Andradóttir et al. [33], on the other hand, investigate the prospect of using flexible servers to compensate for server unreliability under the long-run average throughput objective. Andradóttir et al. [34] also consider the throughput objective, but within the context of a tandem (not parallel) queuing system. Wu et al. [35] and Wu et al. [36] also examine the assignment of unreliable servers in a tandem system, but with the objective of minimizing longrun average holding costs. The remaining papers that address the control of queues with unreliable servers do so by routing customers [37,38] or designing server repair policies [39,40] as opposed to server assignment policies. The present work differs from these in very fundamental ways in terms of the how server unreliability is represented. First, we consider servers who arrive randomly over time from an infinite population of spontaneous volunteers. This is unlike the queuing systems mentioned above where the pool of servers is fixed and finite. Additionally, after remaining in the system for random amounts of time, all volunteers eventually abandon the system forever. They do not return after random periods of inactivity due to repair as in previous studies.

3. Model development

To model this problem, we consider a queuing system where each queue i, $(i = 1, \ldots, N)$, represents a different job or task that volunteers may be assigned to work on, as shown in Fig. 1. Spontaneous volunteers arrive to the system according to a Poisson process with rate λ. As identified in the literature, it has been shown that there is typically an abundance of volunteers following a disaster (e.g. [2–4]), effectively representing an infinite population of potential servers. It is assumed that the volunteer manager is immediately able to assign volunteers to available queues according to a specified control policy. Available queues are all queues not currently at maximum volunteer capacity. Volunteer capacity refers to the number of volunteers that can actively contribute to the relief efforts, which could be due to resource requirements (e.g. carts to transport items), limited space (e.g. a gymnasium being used as a distribution center), or safety concerns. Volunteers' times in system until abandonment are exponentially distributed with mean $1/\gamma$. Work, or demand for services (e.g. request for relief items, or donations to be sorted) arrives to each queue separately, according to a Poisson process with rate α_i.

The number of volunteers working in queue i can be represented as v_i, with the maximum volunteer capacity represented as V_i. Work is completed at each queue according to an exponential additive service rate of $v_i\mu_i$; that is, volunteers are assumed to work on a single task collaboratively. The work or demand for services at each queue i is represented by d_i. The holding cost per unit time for a job of type i is represented as h_i, which for the purposes of this study does not necessarily reflect the inventory holding cost to the relief organization. Alternatively, h_i can be thought of as a representation of the relative importance of specific goods or services during disaster response. For example, consider two tasks related to material convergence: (1) offloading truckloads of donations and (2) sorting/organizing donations. We expect volunteer managers to assign importance such that $h_1 > > h_2$ in most cases. It is assumed that if all queues reach the maximum allowable capacity, volunteers will be turned away from the reception center or directed to another relief organization. Demand may continue to accumulate without being bounded. If demand reaches zero at a queue, volunteers remain idle or participate in ancillary tasks, such as cleaning up the work area, while they wait for additional work to arrive. Immediate reassignment of spontaneous volunteers is not considered.

We can define the feasible state space of the system described above as $\mathcal{S} = \{\mathbf{v}, \mathbf{d} \mid v_i \in (0, \cdots, V_i),\ d_i \in (0, \cdots, \infty),\ i = 1 \ldots N\}$. The state of the system at time t can be defined by $(\mathbf{v}(t), \mathbf{d}(t))$, where $\mathbf{v}(t) = \{v_1(t), v_2(t), \ldots, v_N(t)\}$ and $\mathbf{d}(t) = \{d_1(t), d_2(t), \ldots, d_N(t)\}$ where $v_i(t)$ and $d_i(t)$ refer to the number of volunteers working and the number of unfinished units of work remaining at queue i at time t respectively. Feasible volunteer allocations include all queues where $v_i < V_i$. As described above, the state transitions are all Markovian and therefore this system can be formulated as a continuous time Markov decision process (MDP). The transition rates of the process are state and action dependent. Conditions to ensure stability of this system are provided in the following section.

3.1. Stability

In this section we derive conditions to guarantee that an assignment policy exists which maintains system stability for a set of input parameters, i.e. that a steady-state solution can be found ($\mathbf{E}[d_i] < \infty$ as $t \to \infty$). While the overall system is complex, we can break the model down into subsystems to more easily derive stability conditions. First, consider a random routing policy that thins the volunteer arrival rate into N independent arrival streams with probability $p_i = 1/N$, where $\sum_i p_i = 1$. The result is N subsystems, each with a volunteer arrival rate λp_i. The demand at each of the N queues can be represented as an M/M/1 queue with Markov modulated service rates. The Markov modulated service rates are determined by an underlying Birth Death (BD) process representing the number of volunteers. Births occur according to the rate volunteers arrive to the queue, λp_i, and deaths occur according to volunteer abandonment, $v_i\gamma$. The expected number of volunteers in each queue is independent of the amount of work in that queue for this policy.

Taking a closer look at the demand process, it has been shown by Queija [41] that for M/M/1 queues with Markov Modulated service rates that follow a BD process, the system is stable if and only if:

$$\alpha_i < \sum_{i=0}^{V_i} \pi_i v_i \mu_i \tag{1}$$

where π_i in Eq. (1) is the steady state probability of being in state i of the BD system, which represents the number of volunteers in queue i. Eq. (1) can be rewritten as $\alpha_i < \mu_i \mathbb{E}[v_i]$. That is, the arrival rate must be less than the average service rate to maintain stability. Given that the volunteer process in queue i is a finite BD process, we can solve for $\mathbb{E}[v_i]$ using known BD stability equations. The queue intensity ρ_i in Eq. (2) is represented as $\rho_i = \frac{\lambda p_i}{\gamma}$ for volunteers in subsystem i.

$$\mathbb{E}[v_i] = \sum_{n=0}^{V_i} \pi_i n = \sum_{n=0}^{V_i} \frac{\rho_i^n n}{n!} \frac{1}{\sum_{n=0}^{V_i} \frac{\rho_i^n}{n!}} = \rho_i \left(1 - \frac{\frac{\rho_i^{V_i}}{V_i!}}{\sum_{n=0}^{V_i} \frac{\rho_i^n}{n!}} \right) \tag{2}$$

For the whole system to maintain stability in steady state, it follows that each subsystem must satisfy $\alpha_i < \mu_i \mathbb{E}[v_i]$. Given a set of input parameters (λ, $\boldsymbol{\alpha}$, $\boldsymbol{\mu}$, $\boldsymbol{\gamma}$, $\boldsymbol{V}$) and a random thinning policy such as the one discussed above, we meet sufficient conditions for system stability if:

$$\alpha_i < \mu_i \rho_i \left(1 - \frac{\frac{\rho_i^{V_i}}{V_i!}}{\sum_{n=0}^{V_i} \frac{\rho_i^n}{n!}} \right) \text{for } i = 1, \ldots, N \tag{3}$$

Essentially, the expected service rate for each relief task must be fast enough to keep up with task arrival rates for a given volunteer assignment policy. If the sufficient conditions for stability are met, we have identified an assignment policy that maintains system stability. Thus, an optimal assignment policy is also guaranteed to result in a stable system. We discuss the process of developing an optimal assignment policy for a set of input parameters in the next section.

3.2. Optimal policy

Given a stable system with a feasible assignment policy, we next find the optimal assignment solution. The objective is to minimize the long-run average holding cost of the system. We note that solutions to steady state problems tend to be simpler and easier to implement than finite horizon or non-stationary solutions, which is beneficial in the case of relief efforts. Additionally, note that even though the true arrival rate (λ) is likely to be time varying, analyzing steady state models can indeed be helpful and tend to provide conservative results [42]. To find the optimal policy, we solve an equivalent discrete time problem by employing uniformization, as detailed in [43]. We define the maximum transition rate as: $\Gamma = \lambda + \sum_{i=1}^{N} (\alpha_i + V_i u_i + V_i \gamma_i)$.

Next, we define the recursive optimality equation (Bellman equation) for the discrete-time equivalent finite horizon problem, $J_k(\mathbf{v}, \mathbf{d})$, as the value of being in state $(\mathbf{v}, \mathbf{d})$ with $k < K$ periods left-to-go out of K, and J_0 represents the terminal cost.

$$\begin{aligned} J_k(\mathbf{v}, \mathbf{d}) = \frac{1}{\Gamma} \Bigg\{ & \sum_{i=1}^{N} h_i d_i + \lambda \min(R_i J_{k-1}(\mathbf{v}, \mathbf{d})) + \sum_{i=1}^{N} \alpha_i J_{k-1}(\mathbf{v}, \mathbf{d} + \mathbf{e}_i) \\ & + \sum_{i=1}^{N} \mu_i v_i C_i J_{k-1}(\mathbf{v}, \mathbf{d}) + \sum_{i=1}^{N} \gamma_i v_i J_{k-1}(\mathbf{v} - \mathbf{e}_i, \mathbf{d}) \\ & + \left[\Gamma - \lambda - \sum_{i=1}^{N} (\alpha_i + \mu_i v_i + \gamma v_i) \right] J_{k-1}(\mathbf{v}, \mathbf{d}) \Bigg\} \end{aligned} \tag{4}$$

The first summation in Eq. (4) corresponds to the total holding cost incurred at each queue. The second term denotes the allocation decision for the next arrival given that we are in state (**v,d**). The third term accounts for the arrival of work to each queue. The fourth term accounts for the completion of units of work, where the completion rate is proportional to the number of volunteers working at that queue. The fifth term accounts for the abandonment of volunteers proportional to the number working at each

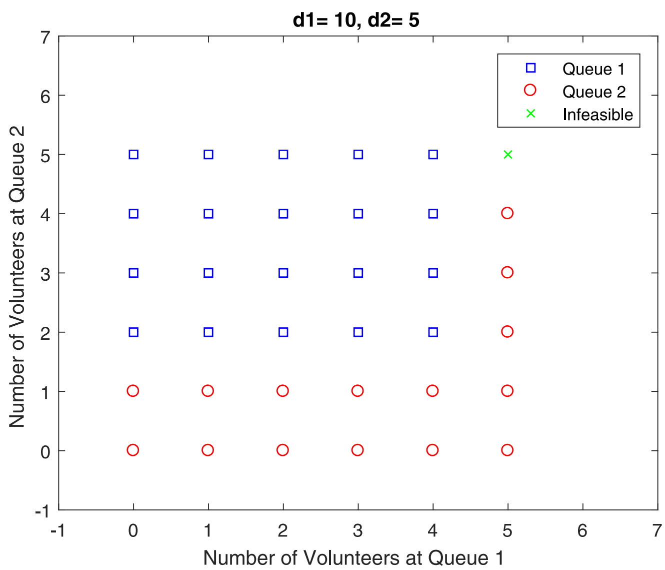

Fig. 2. Optimal Policy Assignments for a two queue system with parameters $\lambda = 3, \gamma = 0.5, V_1 = V_2 = 5, \mu_1 = 2, \mu_2 = 4, \alpha_1 = 6, \alpha_2 = 8, h_1 = h_2 = 20, d_1 = 10, d_2 = 5$.

queue. Finally the last term ensures that all transition rates add up to Γ, which is required for uniformization. Transformation operators R_i and C_i were included to simplify Eq. (4). These operators ensure that the correct value function is chosen based on the feasibility conditions. Specifically, R_i ensures that upon arrival, the volunteer may be sent to queue i if that queue has not reached its maximum capacity; C_i ensures that volunteers complete work with rate μ_i as long as there is work to complete. The transformation operators are defined below:

$$R_iJ_k(\mathbf{v},\mathbf{d}) = \begin{cases} J_k(\mathbf{v}+\mathbf{e}_i,\mathbf{d}) & \text{if } v_i < V_i \\ J_k(\mathbf{v},\mathbf{d}) & \text{otherwise} \end{cases}$$

$$C_iJ_k(\mathbf{v},\mathbf{d}) = \begin{cases} J_k(\mathbf{v},\mathbf{d}-\mathbf{e}_i) & \text{if } d_i > 0 \\ J_k(\mathbf{v},\mathbf{d}) & \text{otherwise} \end{cases}$$

Next, we initialize the value function such that J_k(**v**,**d**)=0 for all (**v**,**d**) $\in \mathcal{S}$, then using the recursive optimality Eq. (4) we apply the value iteration algorithm until $\max(J_k(\mathbf{v},\mathbf{d}) - J_{k-1}(\mathbf{v},\mathbf{d}))$ - $\min(J_k(\mathbf{v},\mathbf{d}) - J_{k-1}(\mathbf{v},\mathbf{d})) \le \epsilon$. For tractability of the value iteration algorithm we truncate the state space by limiting $d_i \le D_i$, where D_i is chosen such that the probability of demand being turned away due to truncation in steady state is less than a small percentage (e.g. 3%). The value iteration algorithm was written and run in Matlab, on a 3.60 GHz Intel(R) Core(TM) i7-4790 CPU machine with 16.00 GB of RAM. The run times to convergence were approximately 3 min for cases with 100,000 states. The MDP policy is specified by the volunteer assignment to queue i, which minimizes the right hand side of Eq. (4). We focus on minimizing the holding cost in the MDP equation as opposed to maximizing reward for completion to avoid scenarios in which a queue becomes completely inundated with demand. This would represent a severe material convergence issue, and would not be considered “optimal” in real world scenarios. An example of the MDP policy can be seen in Fig. 2. When there are two volunteers at queue 1 and three volunteers at queue 2 ($v_1 = 2, v_2 = 3$), the next volunteer should be assigned to queue 1. The MDP policy reports an infeasible assignment when $v_1 = V_1$ and $v_2 = V_2$. In this case, the volunteer would not enter the system.

The MDP policy generated using the recursive optimality equation is state dependent and thus is hard to characterize based solely upon the input parameters of the system. Additionally, implementing such a complex policy during disaster response is difficult. Small to mid-size VRCs are often not equipped with technology capable of assessing the current system and producing complex optimal policies in real time. Instead, we propose the use of heuristic policies, which are easier to implement in practice and test their performance relative to the MDP policy.

4. Heuristic policies

In this section we define four policies to compare to the MDP policy developed in the previous section. These heuristic policies act as alternatives to the MDP policy and are more easily implemented in practice. These policies come from both common sense assignment practices and existing literature on queuing models, with the caveat that they need to implementable at a reception center during disaster response. Accordingly, we define the following heuristic policies:

Fewest Volunteers (FV): An arriving volunteer is assigned to the queue with the fewest number of volunteers. In the event of a tie between multiple feasible queues, the volunteer is assigned to one of the queues randomly with equal probability. This policy is very attractive for use in disaster relief because it is extremely easy for a volunteer manager to implement and only requires basic system knowledge. The FV policy is analogous to the “join the shortest queue policy” commonly used to minimize the total number of customers in the system at any time.

Largest Weighted Demand (LWD): An arriving volunteer is assigned to the queue with the largest weighted demand from the subset of feasible queues which meet the condition $v_i < V_i$ and regardless of any other system parameters. The weighted demand is calculated as d_ih_i. This policy, is similar to a largest demand (LD) policy and should also be easy to implement in practice. The benefit of LWD over LD is that the volunteer coordinator can impart the relative importance of each task during volunteer assignment. This policy can be thought of simply as trying to “put out the biggest fire first”.

Largest Queue Clearing Time (LQCT): An arriving volunteer is assigned to the queue with the largest weighted queue clearing time at the time of assignment. Note that this is not the actual queue clearing time for the queue as the number of volunteers is dynamic and service rates are stochastic. This policy is similar to the D/V policy evaluated by Mayorga et al. [24]. These types of policies are a mix of FV and LWD in that they ensure each task is staffed (when possible) before balancing for workload. The expected queue clearing time is calculated as:

$$\frac{(h_id_i)}{(\mu_iv_i+\epsilon)} \tag{5}$$

Best Random (BR): In this case, an arriving server is assigned to queue i with probability p_i, such that $\sum_{i=1}^{N} p_i = 1$. This heuristic is included only as a reference for comparison and is not intended to be implemented in actual disaster response. BR is provided in place of a true random policy ($p_i = 1/N$) and serves as a tighter performance bound. The optimal thinning probabilities $\mathbf{p_i} = (p_1, \dots, p_N)$, in terms of resulting long run-average holding costs, can be found by solving the mathematical program found in Eq. (6).

$$\min \quad \sum_{i=1}^{N} \frac{\phi_i^2}{1-\phi_i} h_i \tag{6a}$$

$$\text{s.t.} \quad \alpha_i < \mu_i\rho_i \left(1 - \frac{\frac{\rho_i^{V_i}}{V_i!}}{\sum_{n=0}^{V_i} \frac{\rho_i^n}{n!}} \right) \text{for } i = 1, \dots, N \tag{6b}$$

Please cite this article as: K.E. Paret, M.E. Mayorga and E.J. Lodree, Assigning spontaneous volunteers to relief efforts under uncertainty in task demand and volunteer availability, Omega, https://doi.org/10.1016/j.omega.2020.102228

$$\sum_{i=1}^{N} p_i = 1 \tag{6c}$$

$$p_i > 0 \text{ for } i = 1, \dots, N \tag{6d}$$

where $\rho_i = \frac{\lambda p_i}{\gamma}$, as defined previously. The queue intensity of each M/M/1 queue with Markov modulated service rates is represented here as $\phi_i = \frac{\alpha_i}{E[v_i]\mu_i}$. The objective function, shown in Eq. (6a) minimizes the expected weighted holding costs in steady state. Constraint (6b) represents the sufficient condition for stability as discussed in Section 3. The final constraints ensure that each queue has a positive weighted probability of volunteer assignment and that the total weight sums to one.

5. Simulation model

In order to compare the heuristic policies defined above to the MDP policy, a simulation model was developed in Matlab R2017a. The model simulates a system with two unique queues, but the framework can easily be generalized to $n > 2$ queues. The code is separated into two distinct parts: (1) initialization and (2) main simulation. Model Initiation creates and assigns a variety of variables for the case being simulated. These variables include case specific system parameters, number of replications, warm-up period, simulation length, and simulation end time. The MDP optimal policy is stored in the Matlab Workspace or can also be read in from a separate file (for example using the xlsread function). Each run begins in state $v = d = 0$, i.e. zero volunteers and zero demand at all queues. The arrival times for the first volunteer and demand are sampled from the appropriate distribution.

The main simulation tracks the time until next action over all possible actions and continuously compares it to the system time, t. Possible actions include: volunteer arrival, work arrival, volunteer departure, and work completion. The holding cost is updated at each discrete time jump when a new action is triggered. When a volunteer arrival action is triggered, the simulation must decide how the incoming volunteer should be assigned. The heuristic policies are hard coded into the simulation and the MDP policy is referenced based on the current system state. Feasibility of the volunteer assignment decision is considered for all heuristics and no volunteer will be sent to a queue that is at maximum capacity. As discussed previously, if all queues are at capacity, the volunteer does not enter the system. When volunteer arrival or departure occurs, the time to next work completion must be updated preemptively, due to the additive service rate. The demand arrival, demand completion, and volunteer departure actions are modeled similarly. When an arrival or departure occurs, the respective count is updated and a new arrival or departure rate is sampled from the appropriate distribution.

6. Computational analysis and discussion

To evaluate the heuristic policies defined in Section 4, a set of computational experiments was developed. The goal is to provide insight into which volunteer assignment policies perform well over a robust set of scenarios. Experiments were designed for a system with two queues to more easily compare the relative performance of each heuristic policy to the MDP policy. A two queue system can be fully described by ten different system parameters. We set the arrival and departure of volunteers, λ and γ respectively, based on real world data collected by Lodree and Davis [6]. The remaining system parameters are varied to create a robust set of system configurations that cover a broad range of potential real world scenarios (high intensity, low server capacity, varied service rates, etc).

Table 1
Summary of fixed and 2^k factorial design parameters.

V_1	μ_1	h_1	λ	γ	
5	2.4	20	3.066	0.517	
Case	V_2/V_1	μ_2/μ_1	h_2/h_1	α_1	α_2
Low	0.6	0.5	0.5	3	1
High	2	1.5	3	5	2

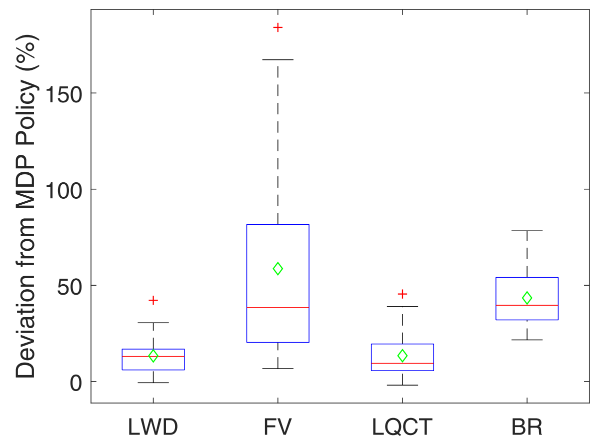

Mean	13.2%	58.5%	13.3%	43.2%

Fig. 3. Box Plot of Average Deviation from MDP policy for each of the four Heuristic Policies. Mean values are represented as diamonds.

However, care must be taken when choosing these system parameters. We must ensure that an assignment policy exists which maintains system stability, as discussed in Section 3.

We fix most parameters of queue 1, as shown in Table 1, to avoid duplicating cases. Next we assign high and low value for remaining parameters, following a 2^k factorial design ($k = 5$), resulting in 32 unique cases (referred to as the baseline cases). A summary of the fixed and varied parameters is shown in Table 1. All full list of the experimental design parameters can be found in Table A1.

Each case and policy was run for 1000 replications in the simulation model. Holding cost was chosen as the primary metric for evaluation in the simulation. Holding cost can be thought of as a proxy for unmet need, with higher costs assigned to items/tasks of greater value to beneficiaries. It represents the cost of having demand sitting at the volunteer site, and not yet given to the beneficiaries. We choose to evaluate average holding cost (AHC) specifically to allow equivalent comparison against the MDP policy, generated using value iteration. We find that the simulation reaches steady state performance in 2 model days, with a full run time of 24 days (representing 2–3 weeks following a disaster). For ease of comparison, we present percent deviation from the MDP policy. The percent deviation is defined as, $\Delta = \frac{\overline{AHC}_{Policy} - \overline{AHC}_{MDP}}{\overline{AHC}_{MDP}}$. The full results of the experimental cases, including: (1) the mean AHC, (2) percent deviation, and (3) statistical significance can be found in the Appendix. A summary of the baseline results is shown in Fig. 3 below. The box plot represents the variation in percent deviation from the MDP policy for all policies across the baseline cases.

As expected, the MDP policy performs at least 13% better than all other policies on average. To further analyze the differences, we conduct a two sample t-test to compare the sample mean AHC of each policy to the MDP policy across all cases. Using a p-value of 0.05, the MDP policy performs statistically significantly better

Table 2
Summary of sensitivity analysis conducted (min:max:increment). For example, (2.85:3.95:0.1) in S1 represent that λ was in discrete increments (2.85, 2.95, ..., 3.85, 3.95).

Case	λ	γ	V_1	V_2	μ_1	μ_2	α_1	α_2	h_1	h_2
S1	2.85 : 3.95 : 0.1	0.5	5	5	2	4	6	8	20	20
S2	2.95	0.5	5	3 : 8 : 1	2	4	6	8	20	20
S3	2.95	0.5	5	5	2 : 8 : 1	4	6	8	20	20
S4	3.95	0.5	5	5	2	4	6	2 : 16 : 2	20	20
S5	2.95	0.5	5	5	2	4	6	8	20	10 : 80 : 10

than the FV and BR policy in all 32 cases. There was no statistical difference in sample mean AHC between the MDP policy and the LWD or LQCT polices in cases 4 and 7. We can also see from Fig. 3 that the LWD and LQCT policies are preferred over FV and BR given a system with stochastic demand tasks. However, it is difficult to identify a preferred policy just by evaluating the box plot. The LQCT policy has a slightly lower median value, but higher overall variation. A similar statistical analysis was conducted on the difference in sample mean AHC for LWD - LQCT. A summary of the numerical results are included in Table A2, found in the Appendix. In 9 cases, the LQCT policy performs significantly better than the LWD policy. Out of the remaining 23 policies, LWD out performs LQCT in 5 of them. In the following sections, we conduct additional analysis to fully evaluate the performance differences between LWD and LQCT.

Interestingly, the FV policy performs worse than both the LWD and LQCT policies, suggesting that the FV policy is not appropriate for cases with stochastic demand tasks. This warrants additional discussion as it is in contrast to the recommendation of Mayorga et al. [24]. There are two significant differences between this work and Mayorga et al. [24] which contribute to the conflicting results. First and foremost is the underlying differences between deterministic and stochastic demand tasks. When a queue was cleared in the deterministic case, all volunteers would leave the system, as neither paper allows for reassignment. Given stochastic demand, it does not make sense to have volunteers leave immediately when the demand is zero, because additional demand may arrive. The FV policy likely performed well in the deterministic case because it minimized the volunteers lost after task completion. The second major difference between the two papers is the model objective. Mayorga et al. [24] compared “time to completion” as the metric when determining optimality. This metric is inappropriate in the case with stochastic demand as there is no clear completion time for stochastic tasks. Additionally, differences in relative holding cost were not considered in the previous paper, as tasks were assumed to be similar or have the same level of importance. Holding costs are a significant factor in the assignment decision of the LWD and LQCT policies.

6.1. Sensitivity analysis

We develop five different sensitivity analysis cases to further investigate the performance differences between the LWD and LQCT policies. Case S1 varies the arrival rate of spontaneous volunteers to determine the impact of varied volunteer arrival rates on AHC. Cases S2-S5 vary system parameters to create different types of queue imbalances. Cases S2 and S3 vary parameters on the volunteer side, maximum volunteer capacity and service rate respectively. The last two cases, S4 and S5, vary demand side parameters, namely arrival rate and relative holding cost, to determine the impact on AHC within the system. The arrival rate in case S4 is increased to allow for feasible solutions with higher demand arrival rates.

A summary of the fixed and varied parameters (min:max:increment) in each sensitivity case (S1–S5) can be found in Table 2. Results are provided as graphs of the sample

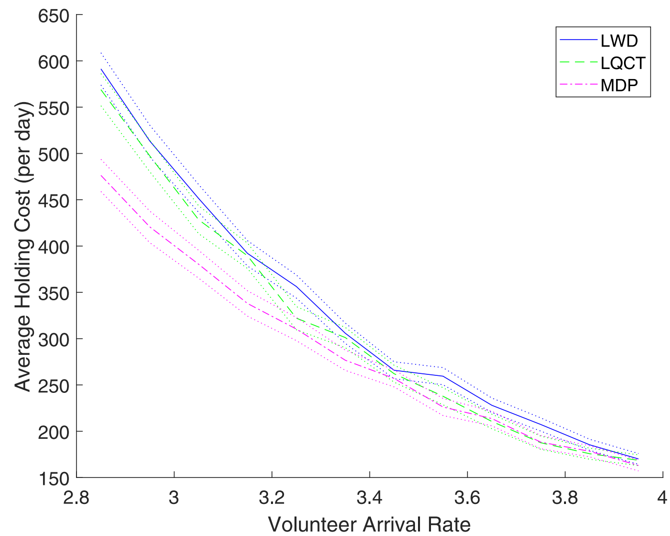

Fig. 4. AHC vs Volunteer Arrival for Case S1.

mean AHC and a 95% half width over the 1000 replications for each of policies. Full results for the sensitivity analysis, can be found in Table A3.

6.1.1. Variations in spontaneous volunteer arrival

It is assumed that in the immediate aftermath of a disaster, queues will experience extremely heavy demand and high spontaneous volunteer participation. We adjust the overall arrival rate of volunteers, λ, in Case S1 from moderate to very high and compare LWD and LQCT to the MDP solution. Fig. 4 indicates that LWD and LQCT policies perform well over the majority of the cases within S1. When volunteer arrival rates are very high, LWD and LQCT perform nearly as well as the MDP solution. Table A3 shows that for cases S1.11 and S1.12, there is no statistically significant difference between MDP and LWD or LQCT policies. The LQCT policy does perform significantly better than LWD for cases S1.8–S1.11, indicating that it may be more appropriate for use in cases with high volunteer arrival rates. At rates beyond λ=3.95 (S1.12), there is not a statistically significant difference between LWD, LQCT, or the MDP policy. If arrival rates are pushed to an extreme (e.g. $\lambda > 20$), SV positions will “always” be full, and so there is no difference between any policy (even random assignment).

As the volunteer arrival rates decrease, the average deviation grows for all policies. At moderate levels of spontaneous volunteer arrivals (S1.1-S1.4), all policies perform statistically significantly worse than the MDP policy. In these situations, volunteer assignment decisions become increasingly important to system performance. There is no significant difference between the LWD and LQCT policies in cases with moderate volunteer arrival rates. Overall, the LQCT policy performed significantly better than LWD in 5 of the 12 sub-cases tested.

6.1.2. Queue imbalance caused by spontaneous volunteers

Next we compare cases with queue imbalances in terms of spontaneous volunteers (S2 and S3). We first consider variations in

Please cite this article as: K.E. Paret, M.E. Mayorga and E.J. Lodree, Assigning spontaneous volunteers to relief efforts under uncertainty in task demand and volunteer availability, Omega, https://doi.org/10.1016/j.omega.2020.102228

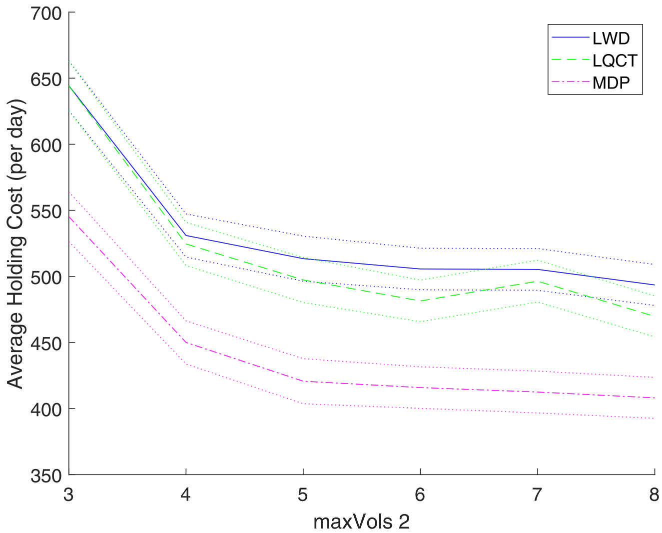

Fig. 5. AHC vs Max Volunteer V_2 for Case S2.

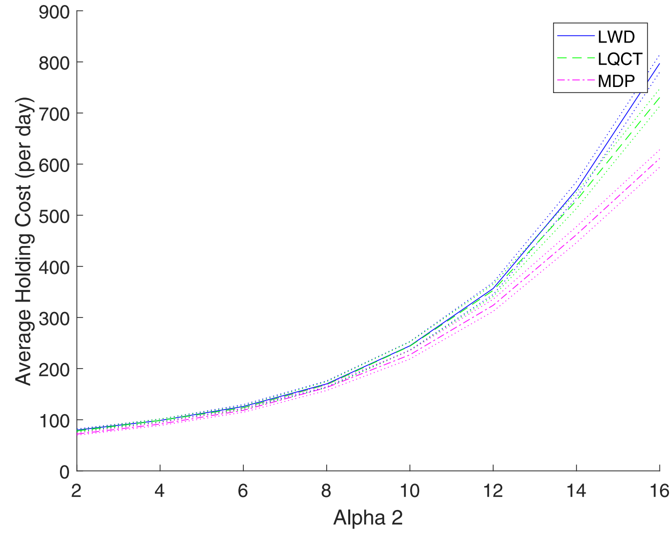

Fig. 7. AHC vs Demand Arrival Rate α_2 for Case S4.

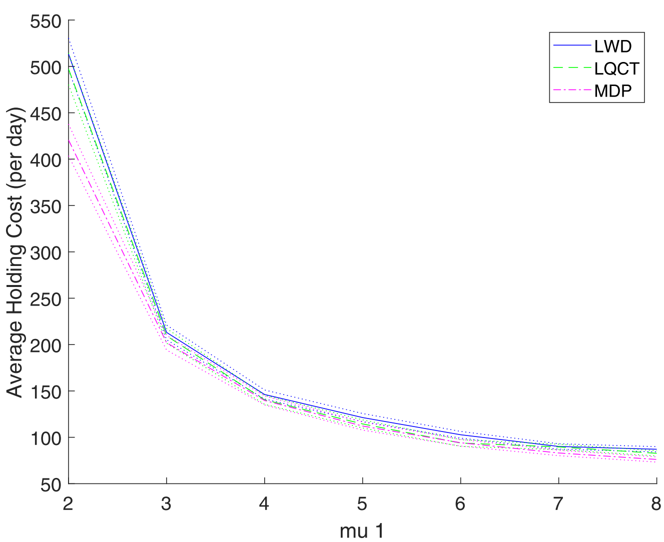

Fig. 6. AHC vs Volunteer Work Rate μ_1 for Case S3.

maximum volunteer capacity (V_i) of one task. Within Case S2, all policies perform statistically worse than the MDP policy over all parameters tested. There is no significant difference in mean AHC between the LWD and LQCT in any of the sub-cases. We can see from Fig. 5 that there exists some value of V_2 at which additional work space does not translate to improved output. Clearly, the assignment policy alone is not always enough to take advantage of the increased floor space. Understanding when this bottleneck occurs is of practical importance for volunteer managers who need to make decisions on planned facility layout with limited floor space. Planners should consider volunteer participation levels when developing floor plans for VRCs. Other options to improve space utilization would be improved volunteer recruitment or initiatives to reduce departure rate of volunteers.

Queues with imbalance in volunteer service rates (α_i) are compared in Case S3. Here, we consider demand tasks with equal relative importance ($h_1 = h_2$) but large differences in service rates. Fig. 6 shows LWD and LQCT performing well across the tested range, indicating they are fairly robust to changes in service rate. The LWD policy does not take into account the service rate differences directly, and therefore it is expected that the LQCT policy performs better. The LQCT performs significantly better than the LWD policy in 2 of the 7 cases. Overall, the LQCT policy is the best performing heuristic policy, matching MDP performance in 4 of the 7 sub-cases.

6.1.3. Queue imbalances caused by demand tasks

Next we move to cases with a queue imbalance on the demand side (S4 and S5). Case S4 varies the work arrival rate in queue 2 (α_2) producing a queue imbalance that also indirectly affects system utilization. As the arrival rate of work in queue 2 increases, we see a nonlinear increase in holding costs across all policies in Fig. 7. At higher demand arrival rates, all policies perform poorly when compared to the MDP policy. In fact, all policies perform significantly worse than the MDP policy in all but one sub-case (S4.4). It is important to note that the LWD policy is the worst performing heuristic for cases with extremely imbalanced demand rates. In cases with extremely imbalanced queues, volunteer managers must exercise caution to ensure that volunteers are not all assigned to one task. For example, at high demand arrival rates, the LWD policy reacts by sending all volunteers to queue 2, effectively ignoring queue 1 for periods of time. The LQCT policy is less affected by this phenomenon due to the volunteer count component within the heuristic. There is no significant difference in the LWD and LQCT policy in 7 of the 8 sub-cases. However, the LQCT policy is more robust to large differences in demand arrival rates.

Finally, in case S5 we consider the scenario in which the relative importance of one task is very different in to another. For example, consider a case in which a relief organization is providing food/water to beneficiaries and also sorting incoming donations at a warehouse. The immediate importance of getting the food/water to beneficiaries likely outweighs the need to sort donations. For cases where the importance is relatively similar (1.2x) LWD and LQCT perform similarly. Fig. 8 identifies a critical point around $h_2 = 1.5h_1$ where the trends change. The performance of LWD and LQCT diverges significantly to the right of this point. LQCT begins to approach the performance level of FV in cases with a large imbalances in relative importance. This is the only sensitivity case in which LWD performed statistically significantly better than LQCT (5 of 8 sub-cases). When extreme differences exist in task importance, LWD is the preferred heuristic assignment strategy.

6.2. Policy recommendation

From the 32 original cases, there is little discernible difference between the LWD and LQCT policies. The LWD policy has a slightly

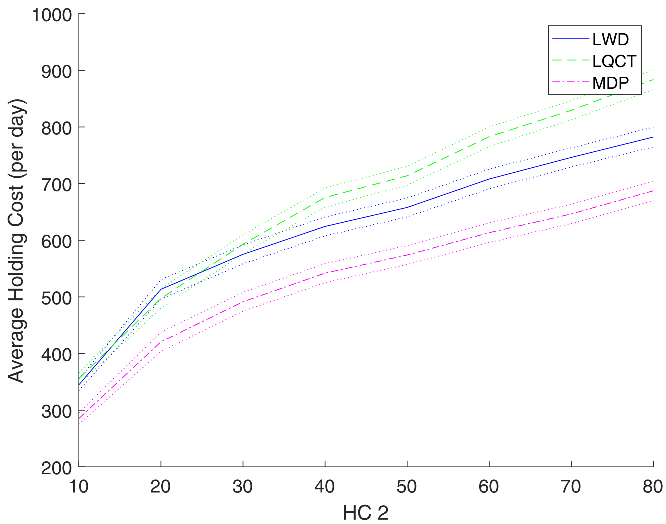

Fig. 8. AHC vs Holding cost h_2 for Case S5.

lower mean average deviation (13.2 vs 13.3) and tighter variance when compared to LQCT. However, the LQCT policy is more robust than the LWD policy, as shown by out performing LWD in 3 of the 5 sensitivity cases (S1, S3, S4). However, in cases with high spontaneous volunteer arrival rates (S1.12), there is no significant difference between the LWD and LQCT policy. One major drawback to the LQCT policy is that it does not handle large differences in holding costs between demand tasks. In cases with a high ratio of h_2/h_1, LWD is clearly the best performing policy.

We must also consider difficulty of implementation when determining a policy recommendation. The FV policy is most attractive in volunteer management due to the ease of implementation and requiring limited system information. The LWD policy, an extension of the Largest Demand policy, requires the volunteer coordinators to keep track of demand and place value on the relative importance of each task. The LQCT policy requires complete system knowledge, which may be unknown during early disaster response efforts. In comparison to all other heuristic policies tested, FV requires the least amount of system knowledge. However, the FV policy is not robust, and performs poorly in a variety of system configurations. As shown in Fig. 3, the FV policy is the worst performing on average and has an extremely large variance in AHC across all cases.

Although LQCT performs slightly better than LWD over the majority of cases tested, it has two major drawbacks: (1) implementation difficulty and (2) performance in systems with a large holding cost differences. It is expected that volunteer coordinators may not have full system knowledge immediately following a disaster making LQCT difficult to implement. Throughout the disaster response, it is expected that tasks with largely different priorities require volunteers. For cases with a large difference between priority ratings, LWD should be used over LQCT to minimize AHC. Therefore we recommend the use of the LWD volunteer assignment policy for disaster response tasks with stochastic demand.

7. Conclusions

In this section we summarize the results of the paper and discuss future work related to spontaneous volunteer assignment.

7.1. Summary

This paper modeled spontaneous volunteer assignment at a volunteer reception center (VRC) in a post disaster setting. The problem was formulated as a queuing model with stochastic arrival of demand, stochastic arrival and departure of volunteers, and stochastic and additive service rates. This is an expansion upon the previous work by Mayorga et al. [24] and allows for stochastic demand. The queuing system was formulated as a continuous time MDP and transformed to a discrete time MDP using uniformization. Solving the MDP using value iteration provided an optimal control policy. It was shown that the optimal policy is extremely complex and not suited for use in a disaster relief operation. Three different heuristic policies were introduced as implementable alternatives to the optimal MDP policy: (1) the Fewest Volunteer policy, (2) the Largest Weighted Demand policy, (3) the Largest Queue Clearing Time policy. A fourth policy, Best Random was included as a reference.

In order to test the performance of these heuristics compared to the MDP policy, a discrete event simulation model was developed in Matlab. A set of thirty two experiments were designed to test the robustness of each heuristic policy. Although the FV policy is easy to implement, it was found to not perform well in cases with stochastic demand. This is in contrast to the results found by Mayorga et al. [24], which recommended FV in cases with deterministic demand. A short discussion of the underlying differences between the two papers was included for completeness. On average the LWD and LQCT policies performed well in relation to the MDP policy when minimizing average holding cost.

Additional sensitivity analysis was conducted to further evaluate the effectiveness of LWD and LQCT. Again, both policies perform well, with LQCT performing better than LWD in the majority of cases tested. The exception to this was cases with large imbalances in relative holding costs, where LWD performed statistically significantly better than LQCT. Ease of implementation and usability was also discussed. The policy recommendation for volunteer assignment given tasks that exhibit stochastic demand is to use the LWD policy. While there are challenges with implementing any routing policy in a post disaster response, the results show that value can be gained from planning for spontaneous volunteers and choosing appropriate assignment stategies. Volunteer organizations that are able to better manage spontaneous volunteers will be better able to serve the affected population.

7.2. Limitations and future work

There are a few important limitations due to the formulation approach and model assumptions that should be discussed. With the Markov assumption, we cannot track volunteers day to day and therefore are unable to account for any individual task learning that may improve the service rate. We assume that tasks change frequently, dependent on the needs of the organizations, and therefore the benefit of learning is minimal. Similarly, we ignore volunteer training because we cannot track training across separate visits under the Markov assumptions. It is possible to include a training element within the model, through the development of additional states, but retraining would occur for each volunteer.

Future work includes efforts to add more realism to the spontaneous volunteer assignment model. We would like to incorporate volunteer preference in the model in two ways. (1) Allow for multiple classes of volunteers who are only able to work or prefer a subset of available tasks. (2) Incorporate variable volunteer departure rates based on task preference or number of other volunteers throughout the system. Finally, we would like to further relax the model to allow for reassignment of volunteers.

Table A1
Input parameters for each of the 2^k baseline cases.

Case	λ	γ	V_1	V_2	μ_1	μ_2	α_1	α_2	h_1	h_2
1	3.066	0.517	5	3	2.4	1.2	3	1	20	10
2	3.066	0.517	5	3	2.4	1.2	3	1	20	60
3	3.066	0.517	5	3	2.4	1.2	3	2	20	10
4	3.066	0.517	5	3	2.4	1.2	3	2	20	60
5	3.066	0.517	5	3	2.4	1.2	5	1	20	10
6	3.066	0.517	5	3	2.4	1.2	5	1	20	60
7	3.066	0.517	5	3	2.4	1.2	5	2	20	10
8	3.066	0.517	5	3	2.4	1.2	5	2	20	60
9	3.066	0.517	5	3	2.4	3.6	3	1	20	10
10	3.066	0.517	5	3	2.4	3.6	3	1	20	60
11	3.066	0.517	5	3	2.4	3.6	3	2	20	10
12	3.066	0.517	5	3	2.4	3.6	3	2	20	60
13	3.066	0.517	5	3	2.4	3.6	5	1	20	10
14	3.066	0.517	5	3	2.4	3.6	5	1	20	60
15	3.066	0.517	5	3	2.4	3.6	5	2	20	10
16	3.066	0.517	5	3	2.4	3.6	5	2	20	60
17	3.066	0.517	5	8	2.4	1.2	3	1	20	10
18	3.066	0.517	5	8	2.4	1.2	3	1	20	60
19	3.066	0.517	5	8	2.4	1.2	3	2	20	10
20	3.066	0.517	5	8	2.4	1.2	3	2	20	60
21	3.066	0.517	5	8	2.4	1.2	5	1	20	10
22	3.066	0.517	5	8	2.4	1.2	5	1	20	60
23	3.066	0.517	5	8	2.4	1.2	5	2	20	10
24	3.066	0.517	5	8	2.4	1.2	5	2	20	60
25	3.066	0.517	5	8	2.4	3.6	3	1	20	10
26	3.066	0.517	5	8	2.4	3.6	3	1	20	60
27	3.066	0.517	5	8	2.4	3.6	3	2	20	10
28	3.066	0.517	5	8	2.4	3.6	3	2	20	60
29	3.066	0.517	5	8	2.4	3.6	5	1	20	10
30	3.066	0.517	5	8	2.4	3.6	5	1	20	60
31	3.066	0.517	5	8	2.4	3.6	5	2	20	10
32	3.066	0.517	5	8	2.4	3.6	5	2	20	60

Table A2
Computational results of the 2^k factorial base cases.

		Mean Holding Cost per Day						Average Deviation from MDP			
	Case	MDP	LWD	FV	LQCT	BR	LWD-LQCT	LWD	FV	LQCT	BR
Baseline	1	24.4	26.1	31.5	26.2	31.6	−0.1	**6.8%**	**28.7%**	**7.4%**	**29.3%**
	2	60.5	61.4	65.1	61.4	80.7	0.0	1.5%	**7.5%**	1.5%	**33.4%**
	3	52.8	55.1	58.5	52.8	76.6	2.3	**4.3%**	**10.8%**	0.0%	**45.1%**
	4	181.8	178.5	222.9	189.9	317.0	**−11.4**	−1.8%	**22.6%**	4.4%	**74.3%**
	5	52.7	55.4	110.3	61.6	68.3	**−6.2**	**5.2%**	**109.3%**	**16.9%**	**29.7%**
	6	103.7	107.3	144.9	109.7	137.4	−2.4	3.4%	**39.7%**	**5.7%**	**32.5%**
	7	103.1	110.9	134.9	109.4	140.4	1.5	**7.6%**	**30.9%**	**6.1%**	**36.2%**
	8	270.9	275.8	295.5	264.4	420.8	11.5	1.8%	**9.1%**	−2.4%	**55.3%**
	9	17.0	19.7	27.5	20.0	20.7	−0.3	**15.7%**	**61.7%**	**17.4%**	**21.6%**
	10	29.5	31.9	36.4	30.8	41.4	**1.1**	**7.9%**	**23.2%**	**4.1%**	**40.0%**
	11	22.6	23.5	29.3	24.4	28.9	−0.9	**4.3%**	**30.0%**	**8.1%**	**28.3%**
	12	46.9	48.6	51.0	47.6	68.8	1.0	**3.6%**	**8.9%**	1.6%	**46.9%**
	13	39.8	45.6	104.9	47.5	49.5	−2.0	**14.5%**	**163.6%**	**19.4%**	**24.3%**
	14	59.5	63.9	113.9	63.9	78.2	0.0	**7.5%**	**91.6%**	**7.5%**	**31.5%**
	15	48.9	54.1	108.3	55.1	61.0	−0.9	**10.8%**	**121.6%**	**12.7%**	**24.8%**
	16	85.5	91.6	130.2	87.7	121.8	**3.9**	**7.1%**	**52.3%**	2.5%	**42.4%**
	17	22.9	28.6	30.8	29.0	34.7	−0.4	**25.0%**	**34.4%**	**26.7%**	**51.5%**
	18	50.7	57.8	57.6	54.8	76.8	**3.0**	**13.9%**	**13.5%**	**8.0%**	**51.4%**
	19	38.7	44.7	46.6	43.6	67.2	1.1	**15.5%**	**20.5%**	**12.7%**	**73.9%**
	20	108.4	116.6	152.5	111.2	171.2	**5.4**	**7.6%**	**40.7%**	2.6%	**57.9%**
	21	47.9	57.3	122.4	64.4	73.9	**−7.1**	**19.6%**	**155.6%**	**34.5%**	**54.4%**
	22	91.6	105.1	149.6	103.0	142.4	2.1	**14.7%**	**63.2%**	**12.4%**	**55.3%**
	23	80.3	89.4	126.1	95.1	141.1	**−5.7**	**11.3%**	**57.0%**	**18.3%**	**75.6%**
	24	194.7	218.1	225.2	206.1	307.3	**12.0**	**12.0%**	**15.7%**	**5.8%**	**57.9%**
	25	16.5	23.5	27.1	24.0	20.7	−0.5	**42.2%**	**64.0%**	**45.5%**	**25.1%**
	26	27.4	33.2	34.5	32.0	39.4	**1.3**	**21.4%**	**26.0%**	**16.8%**	**44.0%**
	27	21.0	27.4	28.8	27.0	28.7	0.4	**30.5%**	**37.4%**	**28.5%**	**36.6%**
	28	41.8	47.2	46.3	45.0	63.7	**2.2**	**12.8%**	**10.8%**	**7.5%**	**52.3%**
	29	38.8	49.9	110.3	53.9	49.0	**−4.1**	**28.5%**	**184.1%**	**38.9%**	**26.3%**
	30	55.3	68.0	117.6	65.3	74.8	**2.7**	**23.0%**	**112.4%**	**18.0%**	**35.1%**
	31	47.0	57.8	117.6	57.9	62.8	−0.1	**23.0%**	**150.0%**	**23.1%**	**33.4%**
	32	77.3	92.6	135.0	86.2	120.4	**6.4**	**19.7%**	**74.6%**	**11.5%**	**55.7%**

Please cite this article as: K.E. Paret, M.E. Mayorga and E.J. Lodree, Assigning spontaneous volunteers to relief efforts under uncertainty in task demand and volunteer availability, Omega, https://doi.org/10.1016/j.omega.2020.102228

Table A3
Computational results of the Sensitivity Cases.

		Mean Holding Cost per Day						Average Deviation from MDP			
	Case	MDP	LWD	FV	LQCT	BR	LWD-LQCT	LWD	FV	LQCT	BR
S1: Vary Arrival Rate, λ	S1.1	476.3	591.3	636.8	568.9	619.6	22.4	**24.1%**	**33.7%**	**19.4%**	**30.1%**
	S1.2	420.7	513.4	595.7	497.3	547.3	16.1	**22.0%**	**41.6%**	**18.2%**	**30.1%**
	S1.3	380.3	451.4	534.2	428.4	496.7	23.0	**18.7%**	**40.5%**	**12.6%**	**30.6%**
	S1.4	338.0	392.1	489.2	389.5	463.1	2.6	**16.0%**	**44.7%**	**15.2%**	**37.0%**
	S1.5	310.3	356.3	462.5	322.1	412.3	**34.2**	**14.8%**	**49.0%**	3.8%	**32.8%**
	S1.6	276.7	305.8	411.7	301.2	376.3	4.6	**10.5%**	**48.8%**	**8.9%**	**36.0%**
	S1.7	257.0	265.8	354.7	262.1	344.7	3.7	3.4%	**38.0%**	2.0%	**34.1%**
	S1.8	226.2	259.5	333.4	237.8	310.6	**21.7**	**14.7%**	**47.4%**	**5.1%**	**37.3%**
	S1.9	213.7	228.0	298.1	210.5	286.4	**17.5**	**6.7%**	**39.5%**	−1.5%	**34.0%**
	S1.10	188.3	207.3	271.3	187.6	258.5	**19.7**	**10.1%**	**44.1%**	−0.3%	**37.3%**
	S1.11	178.3	185.4	250.2	175.7	234.3	**9.6**	3.9%	**40.3%**	−1.5%	**31.3%**
	S1.12	163.0	170.1	225.3	168.7	210.0	1.4	4.3%	**38.2%**	3.5%	**28.8%**
S2: Vary V_2	S2.1	545.2	644.3	611.5	644.2	715.1	0.1	**18.2%**	**12.2%**	**18.2%**	**31.1%**
	S2.2	450.0	531.0	585.7	524.6	572.6	6.4	**18.0%**	**30.1%**	**16.6%**	**27.2%**
	S2.3	420.7	513.4	595.7	497.3	554.8	16.1	**22.0%**	**41.6%**	**18.2%**	**31.9%**
	S2.4	415.9	505.6	577.8	481.4	550.5	24.2	**21.6%**	**38.9%**	**15.8%**	**32.4%**
	S2.5	412.5	505.3	560.7	496.4	566.1	8.9	**22.5%**	**35.9%**	**20.3%**	**37.2%**
	S2.6	408.1	493.5	562.6	469.7	565.5	23.8	**20.9%**	**37.9%**	**15.1%**	**38.6%**
S3: Vary μ_1	S3.1	420.7	513.4	595.7	497.3	547.3	16.1	**22.0%**	**41.6%**	**18.2%**	**30.1%**
	S3.2	202.0	213.1	224.5	209.4	298.1	3.7	**5.5%**	**11.1%**	3.7%	**47.6%**
	S3.3	139.7	146.2	167.7	140.4	209.7	5.8	4.6%	**20.0%**	0.5%	**50.1%**
	S3.4	112.1	121.3	145.4	114.4	179.0	**6.9**	**8.2%**	**29.6%**	2.0%	**59.6%**
	S3.5	94.1	102.7	145.1	93.8	148.8	**8.9**	**9.1%**	**54.1%**	−0.3%	**58.1%**
	S3.6	83.1	90.1	136.5	89.7	133.7	0.5	**8.5%**	**64.3%**	**7.9%**	**60.9%**
	S3.7	76.1	87.1	132.9	82.8	124.3	4.4	**14.5%**	**74.6%**	**8.7%**	**63.4%**
S4: Vary α_2	S4.1	71.9	79.8	184.1	78.0	91.6	1.8	**11.0%**	**156.1%**	**8.6%**	**27.4%**
	S4.2	91.5	98.6	187.8	98.6	112.1	0.0	**7.8%**	**105.2%**	**7.7%**	**22.4%**
	S4.3	118.4	125.9	208.0	123.6	156.2	2.3	**6.4%**	**75.8%**	**4.4%**	**32.0%**
	S4.4	163.0	170.1	225.3	168.7	208.2	1.4	4.3%	**38.2%**	3.5%	**27.7%**
	S4.5	227.0	244.6	276.6	244.0	297.5	0.7	**7.8%**	**21.9%**	**7.5%**	**31.1%**
	S4.6	323.8	356.8	398.6	353.0	415.1	3.8	**10.2%**	**23.1%**	**9.0%**	**28.2%**
	S4.7	462.0	550.0	583.9	528.9	592.5	21.1	**19.0%**	**26.4%**	**14.5%**	**28.3%**
	S4.8	611.1	797.1	764.0	730.4	774.5	**66.8**	**30.4%**	**25.0%**	**19.5%**	**26.7%**
S5: Vary h_2	S5.1	285.2	344.8	534.3	355.8	427.7	−11.1	**20.9%**	**87.3%**	**24.7%**	**49.9%**
	S5.2	420.7	513.4	595.7	497.3	554.8	16.1	**22.0%**	**41.6%**	**18.2%**	**31.9%**
	S5.3	491.4	575.1	657.1	593.0	680.7	−17.9	**17.0%**	**33.7%**	**20.7%**	**38.5%**
	S5.4	542.0	624.4	718.5	675.5	791.9	**−51.1**	**15.2%**	**32.6%**	**24.6%**	**46.1%**
	S5.5	573.9	657.9	779.9	713.9	908.6	**−56.0**	**14.6%**	**35.9%**	**24.4%**	**58.3%**
	S5.6	613.4	708.1	841.3	782.8	1028.7	**−74.7**	**15.4%**	**37.1%**	**27.6%**	**67.7%**
	S5.7	646.5	746.2	902.6	829.4	1133.6	**−83.1**	**15.4%**	**39.6%**	**28.3%**	**75.3%**
	S5.8	687.4	782.2	964.0	883.8	1299.7	**−101.6**	**13.8%**	**40.2%**	**28.6%**	**89.1%**

Appendix A. Full computational results

The Tables in Appendix A provide the full results of the computational experiments. Results are shown for all policies (e.g. Fewest Volunteer, Largest Weighted Demand, Largest Queue Clearing Time, Markov Decision Process, and Best Random) over all cases. Table A1 summarizes the system parameters for each of the baseline cases. Tables A2 and A3 include the sample mean AHC from 1000 simulated replications for each policy. The difference in mean AHC between LWD and LQCT and percent deviation from the MDP policy are also included for ease of comparison. Bolded values indicate a statistically significant difference ($p < 0.05$) in sample mean AHC between MDP and all other policies. Similarly, the difference between LWD and LQCT is bolded if the difference is statistically significant.

Supplementary material

Supplementary material associated with this article can be found, in the online version, at doi:10.1016/j.omega.2020.102228.

References

[1] Fernandez LS, Barbera JA, van Dorp JR. Spontaneous volunteer response to disasters: the benefits and consequences of good intentions. J Emerg Manage 2006;5(4):57–68.
[2] Lowe S, Fothergill A. A need to help: emergent volunteer behavior after september 11th. In: Beyond September 11th: An account of post-disaster research; 2003. p. 293–314.
[3] Cottrell A. A survey of spontaneous volunteers. Australian Red Cross Research Report; 2012.
[4] Townsend FF, et al. The federal response to hurricane katrina: lessons learned. Washington, DC: The White House 2006.
[5] Whittaker J, Mclennan B, Handmer J. A review of informal volunteerism in emergencies and disasters: definition, opportunities and challenges. Int J Disaster Risk Reduct 2015;13:358–68. doi:10.1016/j.ijdrr.2015.07.010.
[6] Lodree EJ, Davis LB. Empirical analysis of volunteer convergence following the 2011 tornado disaster in tuscaloosa, alabama. Nat Hazards 2016;84(2):1109–35.
[7] Gupta S, Starr MK, Farahani RZ, Matinrad N. Disaster management from a POM perspective: mapping a new domain. Prod Oper Manage 2016;25(10):1611–37.
[8] Altay N, Green WG. OR/MS Research in disaster operations management. Eur J Oper Res 2006;175(1):475–93.
[9] Galindo G, Batta R. Review of recent developments in OR/MS research in disaster operations management. Eur J Oper Res 2013;230(2):201–11.
[10] Habib MS, Lee YH, Memon MS. Mathematical models in humanitarian supply chain management: A Systematic literature review. Math Probl Eng 2016;2016.
[11] Ortuño M, Cristóbal P, Ferrer J, Martín-Campo F, Muñoz S, Tirado G, et al. Decision aid models and systems for humanitarian logistics. a survey. In: Decision aid models for disaster management and emergencies. Springer; 2013. p. 17–44.
[12] Leiras A, de Brito Jr I, Queiroz Peres E, Rejane Bertazzo T, Tsugunobu Yoshida Yoshizaki H. Literature review of humanitarian logistics research: trends and challenges. J Human Logist Supp Chain Manage 2014;4(1):95–130.
[13] Anaya-Arenas AM, Renaud J, Ruiz A. Relief distribution networks: a systematic review. Ann Oper Res 2014;223(1):53–79.

[14] Simpson N, Hancock P. Fifty years of operational research and emergency response. J Oper Res Soc 2009;60(1):S126–39.
[15] Caunhye AM, Nie X, Pokharel S. Optimization models in emergency logistics: aliterature review. Socioecon Plann Sci 2012;46(1):4–13.
[16] Fritz C, Mathewson J. Convergent behavior: a disaster control problem. Special report for the committee on Disaster Studies Disaster Study; 1957.
[17] Foundation P. Managing spontaneous volunteers in times of disaster: the synergy of structure and good intentions. Points Light; 2004.
[18] Cross AR. Spontaneous volunteer management resource kit: helping to manage spontaneous volunteers in emergencies. Canberra: Commonwealth of Australia, Australian Red Cross; 2010.
[19] Harris M, Shaw D, Scully J, Smith CM, Hieke G. The involvement/exclusion paradox of spontaneous volunteering: new lessons and theory from winter flood episodes in england. Nonprofit Volunt Sect Q 2017. doi:10.1177/0899764016654222.
[20] Van den Bergh J, Beliën J, De Bruecker P, Demeulemeester E, De Boeck L. Personnel scheduling: a literature review. Eur J Oper Res 2013;226(3):367–85.
[21] Sampson SE. Optimization of volunteer labor assignments. J Oper Manage 2006;24(4):363–77.
[22] Falasca M, Zobel C. An optimization model for volunteer assignments in humanitarian organizations. Socioecon Plann Sci 2012;46(4):250–60.
[23] Lassiter K, Khademi A, Taaffe KM. A robust optimization approach to volunteer management in humanitarian crises. Int J Prod Econ 2015;163:97–111.
[24] Mayorga ME, Lodree EJ, Wolczynski J. The optimal assignment of spontaneous volunteers. J Oper Res Soc 2017;68(9):1106–16.
[25] Gordon L, Erkut E. Improving volunteer scheduling for the edmonton folk festival. Interfaces (Providence) 2004;34(5):367–76.
[26] Kaspari M. Optimal volunteer assignment with an application to the Denver B-cycle bike sharing program. University of Colorado Denver; 2010. Ph.D. thesis.
[27] Squillante MS, Xia CH, Yao DD, Zhang L. Threshold-based priority policies for parallel-server systems with affinity scheduling. In: American Control Conference, 2001. Proceedings of the 2001, 4. IEEE; 2001. p. 2992–9.
[28] Van Mieghem JA. Dynamic scheduling with convex delay costs: the generalized c| mu rule. Ann Appl Probab 1995:809–33.
[29] Cao P, Xie J. Optimal control of a multiclass queueing system when customers can change types. Queueing Syst 2016;82(3–4):285–313.
[30] Bell SL, Williams RJ, et al. Dynamic scheduling of a system with two parallel servers in heavy traffic with resource pooling: asymptotic optimality of a threshold policy. Ann Appl Probab 2001;11(3):608–49.
[31] Saghafian S, Van Oyen MP, Kolfal B. The âæwâg network and the dynamic control of unreliable flexible servers. IIE Trans 2011;43(12):893–907.
[32] Down DG, Lewis ME. The n-network model with upgrades. Probab Eng Inf Sci 2010;24(2):171–200.
[33] Andradóttir S, Ayhan H, Down DG. Compensating for failures with flexible servers. Oper Res 2007;55(4):753–68.
[34] Andradóttir S, Ayhan H, Down DG. Maximizing the throughput of tandem lines with flexible failure-prone servers and finite buffers. Probab Eng Inf Sci 2008;22(2):191–211.
[35] Wu C-H, Lewis ME, Veatch M. Dynamic allocation of reconfigurable resources ina two-stage tandem queueing system with reliability considerations. IEEE Trans Automat Contr 2006;51(2):309–14.
[36] Wu C-H, Down DG, Lewis ME. Heuristics for allocation of reconfigurable resources in a serial line with reliability considerations. IIE Trans 2008;40(6):595–611.
[37] Efrosinin D. Queueing model of a hybrid channel with faster link subject to partial and complete failures. Ann Oper Res 2013;202(1):75–102.
[38] Özkan E, Kharoufeh JP. Optimal control of a two-server queueing system with failures. Probab Eng Inf Sci 2014;28(4):489–527.
[39] Wu C-H, Lee W-C, Ke J-C, Liu T-H. Optimization analysis of an unreliable multi-server queue with a controllable repair policy. Comput Oper Res 2014;49:83–96.
[40] Peschansky A, Kovalenko A. On a strategy for the maintenance of an unreliable channel of a one-server loss queue. Autom Control Comput Sci 2016;50(6):397–407.
[41] Queija N. Steady-state analysis of a queue with varying service rate1997.
[42] Green LV. Capacity planning and management in hospitals. Boston, MA: Springer US; 2004. p. 15–41. doi:10.1007/1-4020-8066-2_2.
[43] Bertsekas DP. Dynamic Programming and Optimal Control. 3rd Edition. Approximate Dynamic Programming; 2011. Volume II Chapter 6.
[44] Sauer L M, et al. The utility of and risks associated with the use of spontaneous volunteers in disaster response: a survey. Disaster Medicine and Public Health Preparedness 2014;8.

Please cite this article as: K.E. Paret, M.E. Mayorga and E.J. Lodree, Assigning spontaneous volunteers to relief efforts under uncertainty in task demand and volunteer availability, Omega, https://doi.org/10.1016/j.omega.2020.102228