

Optimal Control of Parallel Queues for Managing Volunteer Convergence

Gabriel Zayas-Cabán

Department of Industrial and Systems Engineering, University of Wisconsin–Madison, Madison, Wisconsin 53706, USA,
zayasocaban@wisc.edu

Emmett J. Lodree*

Culverhouse College of Business, University of Alabama–Tuscaloosa, Tuscaloosa, Alabama 35487, USA, ejlodree@cba.ua.edu

David L. Kaufman 

College of Business, University of Michigan–Dearborn, Dearborn, Michigan 14853, USA, davidlk@umich.edu

Volunteer convergence refers to the influx of volunteers to affected areas after large-scale disasters. There are not only many benefits to volunteer convergence, but it also creates significant logistical challenges that can impede relief efforts. This study examines policies for admitting volunteers into organized relief operations, and for assigning admitted volunteers to relief tasks. We represent this problem as a queueing system where, in addition to customer arrivals and departures, random server arrivals and abandonments are also present. Then, using a Markov decision process framework, we analyze server admission and assignment policies that seek to minimize relief tasks holding costs as well as volunteer holding and rejection costs. We show that the classic $c\mu$ rule, a server allocation policy that determines where to put servers based on relief tasks holding costs and processing requirements, is optimal under both collaborative and non-collaborative service regimes and when batch server arrivals are allowed. Additionally, we find that the optimal server admission policy is a complex state-dependent policy. As a result, we propose a class of admission heuristics that depend on the number of workers in the system and the remaining system workload. In a numerical study, we show that our heuristic policies perform well with respect to long-run average costs, waiting times, number of volunteers in the system, and number of volunteers idling in the system over a range of parameter values and distributions that are based on real data from a case study. As such, they promise volunteer coordinators an effective and simple way to manage disaster volunteers.

Key words: humanitarian logistics; volunteer scheduling; Markov decision process; queueing; simulation

History: Received: January 2018; Accepted: April 2020 by Martin Starr and Sushil Gupta, after 3 revisions.

1. Introduction

In September 2001, the Red Cross registered 15,000 volunteers during the two-week period that followed the World Trade Center attacks in New York City (Lowe and Fothergill 2003). In August 2005, approximately 60,000 volunteers contributed to relief efforts in New Orleans during the aftermath of Hurricane Katrina, the most destructive natural disaster in United States history (Townsend 2006). In 1985, approximately 2 million individuals engaged in some form of volunteer activity in response to a Mexico City earthquake (Perry et al. 2001). These are all examples of a common post-disaster phenomenon known as *volunteer convergence*—the influx of volunteers to affected areas following disaster events. Volunteer convergence concerns the emergence of *spontaneous volunteers*, who are to be distinguished from emergency management professionals such as police and

emergency medical technicians, and also from volunteers previously affiliated with relief organizations such as the Red Cross. Spontaneous volunteers are individuals who self deploy in an effort to contribute to relief efforts independent of whether or not their services are actually needed (Lowe and Fothergill 2003). They are often untrained in emergency response and tend to contribute on impulse during the early phases of disaster response. A variety of terms have been used by academicians and practitioners in reference to spontaneous volunteers: convergent, disaster, emergent, episodic, informal, unaffiliated, or unofficial volunteers (Cnaan and Handy 2005, Cone et al. 2003, Fulmer et al. 2007, Whittaker et al. 2015).

The convergence of spontaneous volunteers can be both a blessing and a curse. Given the surge in demand for relief supplies and services caused by crisis situations and the importance of satisfying these

demands without delay, perhaps the most obvious advantage of volunteer convergence is that informal volunteers supplement the emergency management infrastructure by providing additional manpower to assist with relief efforts. In fact, emergent volunteers are often on the scene before emergency management professionals arrive, making them the true first responders who actually rescue the majority of disaster survivors (Oberijé 2007). Spontaneous volunteers may perform a variety of relief tasks such as search and rescue, debris clearance, and distribution of relief supplies to disaster survivors (O'Brien and Mileti 1992, Wenger 1991). In addition to providing logistical support, they also contribute socially and psychologically to the overall morale of the relief effort by just "being there" to encourage survivors, responders, and each other (Hamerton et al. 2015, Lowe and Fothergill 2003). Convergence lets survivors know that they are not alone and do not have to feel isolated or abandoned (Barnett and Flint 2005). Furthermore, spontaneous volunteerism can be beneficial to the volunteers themselves. Getting involved with the relief effort can be a very effective coping mechanism to help those affected by disaster begin the process of recovering from the trauma of their experiences (Hamerton et al. 2015, Lowe and Fothergill 2003).

Unfortunately, there are also downsides to volunteer convergence that can be rather significant. This post-disaster phenomenon can be a tremendous burden that impedes relief efforts, so much so that the phrase "disaster within the disaster" has been used by emergency management professionals to describe volunteer convergence (Points of Light Foundation 2002). Perhaps the most problematic and physically observable issue is that affected areas are suddenly overrun by thousands, tens of thousands, or in rare cases, millions of spontaneous volunteers (Allen 1969, Fernandez et al. 2006a). This "mass assault" (a phrase coined by Allen 1969) creates incessant traffic congestion, which in turn denies professional responders unobstructed access to affected areas, and prolongs the delivery of relief supplies and services to the beneficiaries who urgently need them (Fritz and Mathewson 1957, Tierney et al. 2001). Another drawback is that most spontaneous volunteers have no formal training in emergency response when they arrive. Thus instead of attending to critical relief tasks that contribute directly toward the well-being of beneficiaries, professional responders have to deal with multitudes of untrained volunteers who are liabilities to themselves and others (Barsky et al. 2007, Fernandez et al. 2006a, Quarantelli 1998). Furthermore, the professionals may themselves lack training on how to manage spontaneous volunteers and are therefore forced to rely on ad hoc management strategies (Fernandez et al. 2006b). It should be no surprise, then,

that some are of the opinion that spontaneous volunteers are a nuisance and that relief efforts would be better served by excluding them altogether (Barsky et al. 2007).

Accounts of volunteer convergence as well as other forms of convergence¹ date back as far as the 1940s (Fritz and Mathewson 1957). Furthermore, given the examples cited at the very beginning of this study, it is reasonable to assume that volunteer convergence will follow large-scale disasters in the future. Knowing that volunteer convergence is practically inevitable, several guidelines for managing spontaneous volunteers have been proposed over the years. A few of these are (i) encouraging people to become affiliated with a relief organization during times of normalcy in the absence of a crisis situation (Oberijé 2007); (ii) providing volunteer management training for first responders; and (iii) establishing a volunteer reception center to centralize the process of organizing and registering incoming volunteers (Fritz and Mathewson 1957, Points of Light Foundation 2002). It is worth mentioning that imposing hyper-restrictive policies that seek to eliminate volunteer convergence by excluding the majority of emergent volunteers from relief efforts may not achieve desirable outcomes. In addition to losing out on the benefits described above that positively affect both survivors and volunteers, turning away convergent helpers early on can discourage future participation during recovery when the need for volunteer labor is greater. Besides, as described in Lowe and Fothergill (2003), disaster volunteers are often driven by an overwhelming desire, even obsession, to help, and are willing to defy instructions to the contrary in order to find ways to get involved. So, for all practical purposes, the most responsible course of action would be to find ways to integrate spontaneous volunteers into response efforts and work toward addressing the unique challenges they pose (Lowe and Fothergill 2003).

This study concerns the management of disaster volunteers based on a volunteer reception center (VRC) concept. A VRC is a temporary service intended to gain some level of control over the volunteer convergence phenomenon by functioning as a clearinghouse for volunteers to register, earn credentials through training, and receive task assignment instructions. Ideally, VRCs are instituted by local authorities immediately after the occurrence of a ubiquitous disaster event, and are positioned away from affected areas in order to alleviate congestion in those areas and keep volunteers safe. For VRC personnel, the process of assigning volunteers to tasks is very different from labor scheduling decisions in other contexts. The main difference, perhaps, is that volunteer convergence is characterized by uncertainty

in the availability of the entire labor force. Unlike workforce scheduling for paid employees and volunteer labor scheduling under non-disaster-related circumstances, the size and capability of the labor pool change randomly over time. Emergent volunteers show up at uncertain times in uncertain numbers, and they participate in relief efforts for random amounts of time. In addition to frequent and random fluctuations in labor capacity, the task assignment problem faced by VRC personnel is also complicated by the sheer magnitude of convergent volunteers. Generating a feasible set of task assignments for a volunteer workforce would require the availability and perhaps the preferences of each spontaneous volunteer to be taken into account. However, given the inordinate number of volunteers associated with volunteer convergence, the prospect of creating detailed task assignment schedules in real time can be an assiduous task; “big data” analysis of volunteer preferences would be required, and large-scale optimization approaches would likely be intractable. Imagine the Red Cross having to incorporate the individual preferences of the 15,000 volunteers mentioned at the very beginning of this study in an effort to produce a detailed workforce schedule; or imagine having to do so for the 60,000 volunteers following Hurricane Katrina, also referenced earlier in this study. Even if a comprehensive workforce schedule were available, implementation would likely be impractical given the spontaneous behavior of spontaneous volunteers. Instead of attempting to produce a detailed schedule, a more viable option is to simply assign volunteers to tasks at the times they arrive to the VRC.

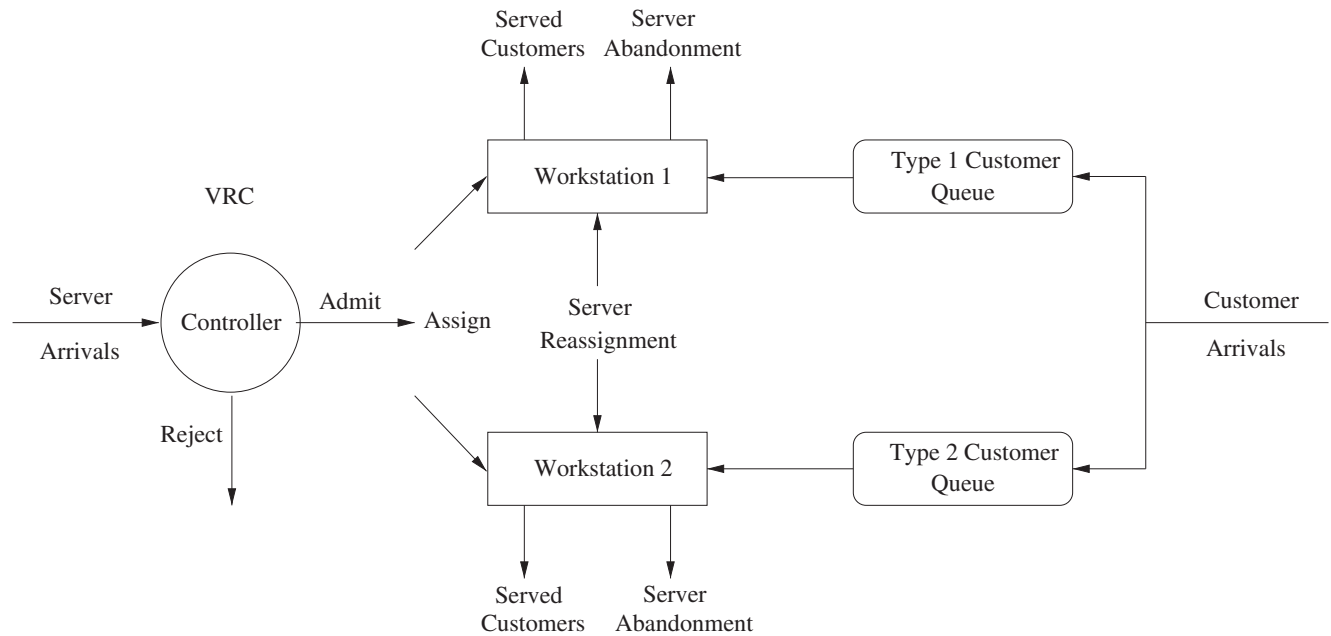
In this study, we represent the VRC and corresponding disaster response as the parallel queueing network shown in Figure 1.² There are Poisson arrival streams to independent workstations in the system, each of which has infinite waiting capacity. “Customers” in this context could be disaster survivors where service represents, for example, receipt of relief supplies or medical attention, which are often done by individual volunteers in scenarios when supplies are small or when medical needs are not emergent (e.g., non-collaborative work); or they could be jobs (i.e., relief tasks) such as unloading in-kind donations from vehicles, where service can be interpreted as the completion of a job (e.g., unloading one of the vehicles). Service times are random and may depend on the workstation and number of servers at each queue. We consider two cases: when servers work collaboratively (i.e., a task/customer is permitted more than one volunteer to work on that task/customer simultaneously) and non-collaboratively (i.e., a task is restricted to one volunteer at a time) to process each customer. We remark that there is no controller who assigns customers to queues; each workstation is

assumed to be a distinct location within the affected region, or customers with different requirements at the same location. In either case, customer arrivals at each queue are exogenous random processes in the sense that they cannot be controlled or influenced in any way. However, we note that there may be scenarios where this is not the case since; for example, beneficiaries may renege or balk when queues for service are long.

To this setup, we add that volunteers, who are the servers in this queueing system, can randomly arrive in batches and randomly depart from the system, which represents a violation of a basic assumption in queueing theory that the set of servers is fixed. Batch arrivals refer to situations in which two or more volunteers arrive at the same time. In the disaster relief context, this is most likely to occur when volunteers are affiliated with the same organization such as Red Cross, a fraternity/sorority, or faith-based organization. However, it is also possible for unaffiliated volunteers to arrive together (e.g., family members or groups of friends). We use the terms “volunteer” and “server” interchangeably throughout the study. In particular, servers arrive to the VRC and the time between consecutive server arrivals is a continuous random variable. Upon arrival, servers are met by a controller (e.g., a VRC manager) who decides whether or not they are admitted into the system³. If the servers are admitted, the controller immediately determines which workstation⁴ the servers are assigned to. We assume that once in the system, servers in the system can be instantaneously re-allocated at any time. In line with what was observed in practice as reported in the Lodree and Davis (2016) case study, we assume that there is no “set-up” cost associated with re-allocation decisions and that once in the system, volunteers cannot be removed. This means that our model is ideally suited for scenarios whereby tasks are relatively close to one another and/or do not require set-ups. Lastly, servers work for a random amount of time before eventually leaving the system. We remark that the time scale is always hours or minutes within a workday, no matter if it is a major response like Hurricane Katrina or a small one. Volunteers work for a random amount of time in a day, and they may or may not return the next day and work for a random time (hours or minutes).

Our goal is to provide server admission and assignment policies that minimize linear holding costs for each customer class, linear holding costs for servers in the system, and server rejection costs. Customer holding costs are used to measure service quality: the longer a customer is in the system, the higher the cost, and hence, the lower the service quality. Server cost rates represent the cost of having the volunteers in this system, who require resources such as

Figure 1 A Queueing System Considered in this Study. In Addition to Random Customer (e.g., disaster survivors, relief tasks) Arrivals and Service Times, there is also Uncertainty in Volunteer (i.e., server) Arrival and Departure Times



equipment, food, and supervision. A useful alternative to linear holding costs is the concept of deprivation cost (Holguín-Veras et al. 2013), which provides a more holistic valuation of post-disaster human suffering based on concepts from welfare economics. However, our study approaches volunteer convergence from a conventional queueing control perspective by analyzing the impact of server admission and allocation policies on congestion (i.e., number of beneficiaries and volunteers in the system) and, hence, on delays to service for beneficiaries. A useful way to do this is to use linear holding costs for beneficiaries and volunteers. In this case, the average holding cost is proportional to beneficiary queue lengths and to the average number of servers in the system. Average beneficiary queue length is proportional to the average delay of beneficiaries due to Little’s Law (Little 1961). We allow for average lump-sum rejection costs since they are proportional to the long-run fraction of servers that are rejected from the system, an important metric for VRC managers. We consider the problem of minimizing expected discounted *and* average costs over an infinite horizon. Allowing for servers to randomly arrive and depart turns out to be important in determining not only what type of policy is optimal, but also what techniques can be used to show structural properties of optimal policies. When server holding costs are linear, we prove that the classic $c\mu$ rule, a server allocation policy that determines where to put servers based on relief task holding costs and processing requirements (c.f., Buyukkoc et al. 1985), is optimal when there are two customer classes. Our

result extends the classic $c\mu$ rule by allowing servers to work collaboratively *or* non-collaboratively on a single job and by allowing batch server arrivals and departures.

One may also conjecture that the admission criteria becomes less (more) strict with higher customer (server) congestion (i.e., the optimality of threshold policies with respect to the admission criteria). However, this turns out not to be so simple. The reason is that a randomly varying service capacity leads to trade-offs between immediately preventing increases in costs (by rejecting arriving servers) and reducing customer holding costs as quickly as possible (by accepting arriving servers). Additionally, we must contend with technical challenges associated with unbounded transition rates that are a consequence of allowing any number of servers in the system, each of which can abandon.

To provide insight into this complex stochastic sequential decision-making scenario, we use a Markov decision process (MDP) framework. We contend that MDPs balance problem complexity (i.e., stochastic arrivals, departures, and service completions) with tractability (i.e., exponentially distributed inter-arrival, inter-departure, and service completion times). Insights (e.g., the optimality of static priority rules for allocating volunteers between tasks) can be used to suggest practical approaches for managing volunteer convergence. To our knowledge, we are the first to use MDPs to determine dynamic server admission and assignment policies for a randomly varying and flexible workforce in a parallel queueing system. We

consider customer linear holding costs, server linear holding and rejection costs, collaborative and non-collaborative service, and batch server arrivals. We also see our work as having the following contributions to queueing control and workforce planning literatures.

Contributions to queueing: First, when there are two customer classes, we show that the classic $c\mu$ rule holds in both the collaborative and non-collaborative model and with batch server arrivals no matter how many workers are available. The optimality of the $c\mu$ rule when servers do not collaborate has not, as far as the authors know, been shown before. Second, we analyze the structure of optimal server admission policy and provide examples showing that seemingly intuitive switching curve policies may not necessarily be optimal for truncated versions of our model, and that, in general, optimal server admission policies may be complex, state-dependent policies that may be difficult to implement in practice. We thus propose a class of admission heuristic policies that depend on the system workload, number of volunteers, and accounts for batch arrivals. We compare their performance with the optimal policy in a numerical experiment when inter-arrival, service, and abandonment times are exponentially distributed and show that, for a given instance of the model (i.e., rates and costs), we can find parameter values for our proposed class of admission heuristics that are close to optimal. We then compare their relative performance in a simulation study with respect to long-run average customer waiting times, number of servers, and number of idling servers in the system when servers arrive and leave in batches and when server departure times are no longer exponential, based on observational data from actual volunteer convergence efforts from Lodree and Davis (2016). Here, we again find parameter values for the admission heuristic that perform well relative to these performance metrics. Lastly, we contend with handling an MDP model with unbounded transition rates due to server abandonments and allowing an unbounded number of servers in the system, which presents a technical challenge. The infinite server population assumption is motivated by the practical considerations described in the Lodree and Davis (2016) case study. In some situations, volunteer managers may need as much help as possible to handle a large number of tasks and/or tasks with long processing times, and therefore not bound the number of volunteers in the system *a priori*. Also for large-scale disasters, volunteers may travel from just about anywhere in the world to participate in relief efforts (e.g., mission groups and student groups from universities may travel to other countries). As such, the population from which servers arrive is effectively infinite. Since we consider average costs, we must

analyze the stability of the system, that is, finite expected queue lengths. When transition rates are unbounded and service is non-collaborative, this requires special care, as we show in Theorem 3.5 (c.f., Blok (2016), Blok and Spieksma (2017)). Additionally, the common technique of uniformization (c.f. Lippman (1975)) is not possible. We thus first truncate the multi-dimensional state space by allowing a fixed and finite number of servers in the system. This allows us to obtain structural properties for the original unbounded as well as approximate the optimal costs of the original unbounded model by first showing the structural properties as well as calculating the optimal costs for each truncated, and hence uniformizable, model, and then taking the limit as the truncation goes to infinity (c.f., Bhulai et al. (2014), Blok and Spieksma (2015)).

Contributions to workforce planning: Sampson (2006) outlines how volunteer scheduling differs from classical workforce scheduling problems that involve paid employees. These unique features have only been considered in a handful of papers that address volunteer scheduling from an operations management perspective. Of these, roughly half do so within the context of Disaster Operations Management (DOM), which according to Falasca and Zobel (2012), has an even more refined set of unique features. The volunteer convergence context addressed in this study differs from all other forms of labor assignment in that worker arrival and departure times are uncertain. Volunteer labor assignment with random volunteer arrivals and abandonments is considered in Mayorga et al. (2017), Abualkhair et al. (2020), and Paret et al. (2020), all of which feature computational approaches to analyzing optimal and/or heuristic assignment policies. We extend these studies by deriving optimal policies analytically for the resulting queue control problem, and also by adding an accept/reject decision for arriving volunteers to the decision process. Finally, we highlight that labor planning studies in DOM are relatively scarce. As such, this study makes an important contribution to the DOM literature as well.

Organization of the study: The remainder of this study is organized as follows. First, selected studies from the literature are discussed in section 2. The contributions of this study relative to three related topical areas within the literature are highlighted. Next, in section 3, the aforementioned continuous-time MDP formulation for the queueing control problem depicted in Figure 1 is presented, along with some preliminary structural results. Further properties of the optimal control policy are derived in section 3, followed by a numerical study in section 4. In the numerical study, we use the insights from the structural results from our MDP model to propose and

analyze a class of practical admission policies under realistic conditions based on the real-world study of Lodree and Davis (2016). Concluding remarks are given in section 5.

2. Literature Review

This study represents a synthesis of three areas of academic literature, namely Disaster Operations Management (DOM), volunteer scheduling, and queueing theory. To our knowledge, Mayorga et al. (2017) is the first study to address the intersection of these three areas. They consider the assignment of spontaneous volunteers to parallel workstations. Each station has a known (i.e., not random) amount of work to be completed, *a priori*, and volunteers are modeled as servers that randomly join and abandon the system. They make two simplifying assumptions: that service is collaborative and that there is fixed and finite upper bound on the total number of volunteers allowed in the system at any time. Their focus is on investigating the effectiveness of various volunteer assignment policies numerically with respect to minimizing the time it takes to complete the tasks at all stations. To do this, they first develop a MDP formulation that maximizes rewards for completed work less a holding cost for unfinished units of work. They then consider the time required to complete all tasks in a simulation study.

The recent study Paret et al. (2020) generalizes Mayorga et al. (2017) by considering stochastic demand streams along with random volunteer arrivals and abandonments. Their approach mirrors that of Mayorga et al. (2017); specifically, they develop a MDP model to optimally assign volunteers to parallel queues, apply a computational approach to generate MDP optimal policies (value iteration), and analyze the effectiveness of several heuristic policies using discrete-event simulation. Another recent study, Abualkhair et al. (2020), is closely related to both Mayorga et al. (2017) and Paret et al. (2020). Like Paret et al. (2020), Abualkhair et al. (2020) consider the assignment of spontaneous volunteers who arrive at random to a parallel queueing system with stochastic demands at each queue, but also consider interdependencies between the two queues. In particular, the stochastic demand at one queue represents donations that are processed by volunteers, while customer arrivals at the other queue are beneficiaries who require the donations that have been processed at the first queue. However unlike Mayorga et al. (2017) and Paret et al. (2020), Abualkhair et al. (2020) does not generate optimal policies; they compare the performance of several heuristic policies through an extensive computational experiment using an agent-based simulation model.

Our work differs from the three above-mentioned studies in several respects. First, in addition to deciding which queue each volunteer is assigned to, our MDP model includes an admission control problem where the controller decides whether to admit or turn away volunteers upon arrival and where each individual rejected volunteer incurs a cost. Second, we consider collaborative and non-collaborative scenarios and do not restrict the total number of volunteers allowed in the system. Third, we consider compound renewal processes in which volunteers arrive and abandon the system in batches, which is an important characteristic of volunteer convergence according to the case study conducted by Lodree and Davis (2016). Fourth, our problem takes place in continuous time over an infinite horizon. Fifth, like Mayorga et al. (2017) and Paret et al. (2020), our cost structure includes customer holding costs for unfinished work (i.e., unserved customers or queue lengths), but in addition, we also include a server cost rate that depends on the number of volunteers in the system to account for the fact that volunteers require resources such as equipment, food, and supervision. Lastly, and perhaps most importantly, we characterize the structure of the optimal policy analytically, which is then used to develop and analyze practical heuristics.

2.1. Disaster Operations Management and Volunteer Scheduling

Classification schemes represent a common, useful, and perhaps necessary feature of DOM survey papers. Gupta et al. (2016) for example, proposes classifying DOM research into two categories: the *solution domain* and the *disaster domain*. The solution domain refers to methodologies used to generate solutions, recommendations, and insights for DOM decision makers; as well as the types of data used to validate results (field/archival, real, or hypothetical). DOM review articles have identified several methodologies such as mathematical programming, decision analysis, game theory, simulation, statistical analysis, queueing theory, heuristics, evolutionary algorithms, and artificial intelligence/expert systems. The disaster domain includes type of disaster (natural or man-made), time phase with respect to disaster occurrence (before, during, or after), and administrative function (decision-making process, prevention and mitigation, evacuation, humanitarian logistics, casualty management, and recovery). Other classification domains have been proposed. One of these is based on which of the four phases of the disaster management cycle they address: mitigation, preparedness, response, or recovery (Altay and Green 2006, Galindo and Batta 2013, Hoyos et al. 2015, Leiras et al. 2014, Ortuño et al. 2013). These describe specific issues addressed within the four phases, including facility location

(Anaya-Arenas et al. 2014, Caunhye et al. 2012, Habib et al. 2016, Hoyos et al. 2015, Leiras et al. 2014, Ortuño et al. 2013), relief distribution and routing (Anaya-Arenas et al. 2014, Caunhye et al. 2012, Habib et al. 2016, Hoyos et al. 2015, Leiras et al. 2014, Ortuño et al. 2013), inventory management (Leiras et al. 2014, Ortuño et al. 2013), casualty transportation (Anaya-Arenas et al. 2014, Caunhye et al. 2012, Hoyos et al. 2015), resource allocation (Caunhye et al. 2012, Hoyos et al. 2015), evacuation (Caunhye et al. 2012, Habib et al. 2016, Ortuño et al. 2013), donations and funding (Burkart et al. 2017, Ortuño et al. 2013), and search and rescue (Hoyos et al. 2015).

Lastly, the DOM literature reveals trends with respect to the solution and problem domains. Mathematical programming is by far the most widely reported methodology among DOM survey papers. For example, Hoyos et al. (2015) identify mathematical programming as the primary methodology in 47% of DOM papers published between 2006 and 2012, the most of any of the methodologies reported in their review. Queueing theory, our main methodological approach, on the other hand, has received much less attention. Hoyos et al. (2015) indicate that queueing models appear in only 4% of the DOM studies during that same period, the least among the OR/MS methods identified in their survey. In terms of problem domain, facility location and relief distribution have received considerably more attention than all the other problem types. By contrast, researchers have identified manpower planning as a crucial element of the disaster management process that has been, for the most part, ignored by the DOM literature (Caunhye et al. 2012, Simpson and Hancock 2009).

There are only a limited number of papers that focus on volunteer scheduling and DOM. Other than the three papers discussed at the beginning of this section (Abualkhair et al. 2020, Mayorga et al. 2017, Paret et al. 2020), the only other studies to examine volunteer scheduling from a DOM perspective are Falasca and Zobel (2012), Lassiter et al. (2015), and Urrea et al. (2019). However, systematic scheduling of volunteer labor has been considered in other settings; an annual folk music festival (Gordon and Erkut 2004), reviewer assignments for an academic conference (Sampson 2006), and a bike sharing program (Kaspari 2010). Several challenges that are unique to volunteer scheduling compared to conventional workforce scheduling (for paid employees) are identified in Sampson (2006). An important one is that traditional labor assignment (TLA) often seeks to minimize labor costs of meeting task demands, whereas for volunteer labor assignment (VLA), labor costs are, for the most part, insignificant. Instead, volunteer preferences such as which tasks they are assigned to and which times slots they are scheduled to work are

of central importance, primarily because volunteer satisfaction is a key factor that determines if a volunteer returns in the future. As such, the objective of VLA is to balance labor shortages among tasks while maximizing volunteer preferences.

While these issues are also relevant to volunteer scheduling in the DOM context, there are other factors that need to be considered. First, although labor costs are not as expensive as they are in TLA, they are not necessarily insignificant as explained in Sampson (2006); humanitarian organizations that facilitate VLA sometimes provide meals and other services for volunteers, and they have a limited budget to do so. Another is that volunteers are often dispersed across different locations geographically, which is typically not the case for non-DOM contexts. Furthermore, many volunteers participate in groups (e.g., families, sports teams, fraternities and sororities, faith-based organizations) in response to disaster events; thus it may be necessary to consider group assignments in volunteer preferences. All of these matters are taken into account in Falasca and Zobel (2012). Another characteristic that is likely to be more pronounced in DOM than in other VLA contexts is the degree of task uncertainty. Lassiter et al. (2015) accommodate task uncertainty by proposing a robust optimization framework that seeks to minimize unmet task demands. Lastly, VLA in DOM is often characterized by a large number of unaffiliated volunteers who lack training and experience, and a smaller number of experienced volunteers who are affiliated with official responding agencies. Using an agent-based simulation model, Urrea et al. (2019) examine polices for pairing inexperienced and experienced volunteers in a relief storehouse while also taking volunteer congestion into account.

All of the above-mentioned volunteer scheduling studies (except for Mayorga et al. 2017, Abualkhair et al. 2020, Paret et al. 2020) assume that volunteer availability is not an issue, and that volunteers always adhere to the shifts and tasks they are assigned to. Consequently, the studies by Falasca and Zobel (2012) and Lassiter et al. (2015), both of which concern VLA from a DOM perspective, are applicable to affiliated volunteers, but are not appropriate for managing the convergence of unaffiliated volunteers. As described in Lodree and Davis (2016), random numbers of unaffiliated volunteers arrive at or near affected areas at random times, and they participate in relief efforts for random amounts of time. Thus, uncertainty in volunteer arrival and abandonment times must be taken into account when making labor assignment decisions within the context of volunteer convergence. We address these uncertainties by modeling volunteer inter-arrival and abandonment times as independent stochastic processes, and to our knowledge, is the

only VLA study besides Paret et al. (2020) to consider uncertainty in both task demands and volunteer availability.

2.2. Queueing Theory

Our proposed framework for accepting and then assigning spontaneous volunteers to relief activities during volunteer convergence is an example of a queueing control problem (c.f., Kitaev and Rykov (1995), Stidham (2002)) that intersects dynamic scheduling of resources (e.g., servers), analysis and control of systems with randomly varying service capacity, and continuous-time MDPs (CTMDPs) with unbounded transitions rates.

Scheduling problems arise when a decision to allocate resources dynamically to competing demands must be made. An early setup of this type of problem consists of customers of different classes that arrive to the system according to mutually independent Poisson processes where they join queues dedicated to their respective classes. Each customer waits for a server on a first-come-first-served (FCFS) basis. Each customer class i takes a random amount of time with mean $1/\mu_i$ where μ_i is the service rate. Customers exit the system after service. Each customer class i incurs a class-dependent linear holding cost c_i for each unit of time that the job is in the system and the objective consists of finding a sequencing rule, or policy, for dynamically allocating the server between queues to minimize costs. For this setup, one of the most celebrated results is that a static priority rule/discipline, termed the $c\mu$ rule, is optimal. Originally proposed by Smith (1956) in a clearing system with deterministic service requirements, and later by Cox and Smith (1961) within the context of a queueing network, this rule says that, except to avoid unforced idling, the server should be assigned to the customer with highest value of $c_i\mu_i$. Extensions have been extensively considered (c.f., Ahn et al. 2002, Baras et al. 1985, Baras et al. 1985, Bell and Williams 2001, Buyukkoc et al. 1985, Buyukkoc et al. 1985, Mandelbaum and Stolyar 2004, Nain 1989, Van Mieghem 1995).

Random server arrivals and departures means that our model is an example of a queueing system with a controlled Markov-modulated service capacity: service capacity can change according to a controlled external environment that is governed by a Markov process. Markov-modulated queues also have been well studied in the literature Neuts (1981), Purdue (1974), Regterschot and De Smit (1986), Mahabhashyam and Gautam (2005), Perel and Yechiali (2008), Thorsdottir and Verloop (2016). We refer the reader to the many references there. These studies (and references therein) focus on analyzing performance measures to evaluate existing or proposed systems when the server allocation and admission policy

are fixed. In the present study, the analysis of a stability condition is required to show the existence of the average cost optimality equations (ACOE). This analysis is akin to the type of analysis in the aforementioned papers.

As far as the authors know, the studies by Kaufman et al. (2005) and, subsequently, by Budhiraja et al. (2014) are the only ones on the control of queueing systems with randomly varying service capacity. Kaufman et al. (2005) consider a system similar ours except that it is a two-station *tandem* queue where customers arrive according to a (single) Poisson process and receive service at both stations before leaving the system, and there is a bound on the total number of workers allowed in the system. They show that all workers should be allocated to one queue or the other and that they should serve exhaustively at one of the queues depending on the direction of an inequality that is akin to the $c\mu$ rule, but for tandem queues.

Our work differs from Kaufman et al. (2005) in several respects. We consider a *parallel* queueing system, and we allow for batch server arrivals. We also consider when servers are allowed to collaborate on a single job and a non-collaborative scenario. Our extension of the $c\mu$ rule to a parallel queueing system with batch server arrivals and non-collaborative settings represent contributions of our work to the queueing literature. We also consider server rejection costs, and lastly, we do not bound the number of the servers allowed. This last point implies that the transition rates for the continuous-time Markov chain under a given policy are unbounded and the problem is not uniformizable, so that traditional solution techniques (e.g., action elimination, successive approximations) can no longer be easily applied.

Applications of MDPs with unbounded rates are becoming widespread (c.f., Bhulai et al. 2014, Down et al. 2011, Guo and Hernández-Lerma 2009, Legros et al. 2014) and are an active research topic because of the technical complications they induce. Although these applications are distinctly different problems from ours, their approaches serve as useful models for how to analyze optimal controls for problems with unbounded rates like ours. For instance, to obtain the structure of the optimal control, we also first truncate the state space before taking limits. This approach can be justified by using recent results from Blok and Spieksma (2017). It allows us to analyze the structure of the optimal policy for the original unbounded model by analyzing the structure of the optimal policy for the uniformizable truncated model.

3. Dynamic Control

In this section, we develop a CTMDP model and derive optimal server assignment policies for the

queuing system shown in Figure 1 with two parallel queues. We also discuss a nonintuitive result regarding the optimal server admission policy, namely that the optimal admission policy is not always monotone with respect to the number of customers nor the number of servers. It is important to note that although the queuing system shown in Figure 1 may not reflect real-world scenarios such as Hurricane Katrina, the two queue model is still relevant in practical situations, even for some large-scale relief efforts. The case study Lodree and Davis (2016) describes a queuing system with two queues for the large-scale response following the 2011 tornado disaster; one queue consisted of beneficiaries who received emergency supply items such as food, clothing, and water, and the other of donors who donated these items.

3.1. Model Description and Preliminaries

Customer arrivals to parallel stations 1 and 2 occur according to independent Poisson processes with rates λ_1 and λ_2 , respectively. We refer to arrivals to station i as class i customers. Independent of the arrival process, customers may receive service before leaving the system. Customers in the same class are served on a FCFS basis and their service requirements are probabilistically the same in the sense that they are exponential with finite rate $\mu_i > 0$. As described in section 1, the system is not equipped with dedicated, permanent servers. Instead, and independently of the customer arrival processes, servers arrive in batches according to a Poisson process of rate $\nu > 0$ where each arriving batch is of size $B = 1, 2, \dots$ with probability p_B , and each admitted server is available for an exponentially distributed amount of time with finite rate $\beta > 0$ before abandoning the system.

The decision-making scenario is the following. At each server arrival time, the manager views the number of customers of each class and the number of available volunteers in the system and based on this and a combination of customer and server costs, decides whether the arriving servers should be accepted or rejected. Individual servers that are rejected incur an immediate one-time cost of $K > 0$ and are lost forever. This one-time cost K can represent opportunity cost associated not having that volunteer at a later, potentially busier time, or, as mentioned in section 1, the emotional cost of turning away a volunteer who is also a victim who could have used the volunteering experience as a coping mechanism. Moreover, after each event (arrival, service completion, or departure), the manager must also decide how each server should be allocated between the different customer classes. We assume throughout our analysis that there is no set-up time or cost associated with each re-allocation decision

and that customers in service can be preempted to reallocate workers. In the following, our primary setting is when servers do not collaborate, but in remarks, we also highlight how all of our results hold in the case when servers do collaborate.

We model the server admission and allocation problem using an MDP formulation. Because inter-arrival and inter-departure times, as well as service requirements are assumed to be exponentially distributed, the state space is given by $X = \{(x_1, x_2, y) \in \prod_{i=0}^2 Z_{\geq 0}\}$, where $Z_{\geq 0}$ is the set of non-negative integers. Coordinates 1 and 2, respectively, denote the number of class 1 and class 2 customers waiting for or in service, and the last coordinate denotes the number of servers in the system. For a given policy π , Let $N(t, \pi)$ be a counting process that counts the number of decision epochs by time t and let σ_n^π represent the corresponding time of the n th epoch. Let h_i be the cost rate at which class i customer holding costs are incurred ($i = 1, 2$), and h_v the cost rate at which server holding costs are incurred. A policy prescribes both whether to admit or reject an arriving server and how many servers should be allocated to each customer class.

For $\alpha > 0$, the finite horizon, α -discounted expected cost for a non-anticipating policy π is given by $v_{t,\alpha}^\pi(x) \equiv \mathbb{E}_x^\pi[\sum_{n=0}^{N(t,\pi)} e^{-\alpha\sigma_n^\pi} k(a_n^\pi)] + \int_0^t e^{-\alpha s} \mathbb{E}_x^\pi[\sum_{i=1}^2 h_i Q_i^\pi(s) + h_v V^\pi(s)] ds$, where $Q_i^\pi(s)$ and $V^\pi(s)$ denote, respectively, the customer class i queue length process and the number of volunteers available at time $s \geq 0$, and a_n^π represents and the type of event seen at the time of the n th decision. The function $k(\cdot)$ denotes the fixed cost; that is, if σ_n corresponds to a volunteer arrival of size B that is rejected, then $k(a_n^\pi) = B \cdot K$ (it is zero otherwise). For fixed $x \in X$, the infinite horizon discounted expected cost under policy π is $v_\alpha^\pi(x) \equiv \lim_{t \rightarrow \infty} v_{t,\alpha}^\pi(x)$. The long-run average cost rate is $\rho^\pi(x) \equiv \limsup_{t \rightarrow \infty} v_{t,0}^\pi(x)/t$. We seek a policy π^* such that $w^\pi(x) = \inf_{\pi \in \Pi} w^\pi(x)$ where Π is the set of all non-anticipating policies and $w = v_\alpha$ or ρ .

We begin our analysis with three intuitive results that are used throughout. The first is a statement on the monotonicity of the value functions, which says that it is better, from a cost standpoint, to have less customers in the system. The second says that there is an optimal policy that does not idle the servers whenever there are customers waiting. The latter is later used to simplify the optimality equations. The third says that if there is no cost associated with having volunteers in the system (i.e., $h_v = 0$), then it is optimal to always accept volunteers. We will see later on that the optimal server admission policy is much more difficult to characterize when this condition does not hold. In the interest of brevity, and because the proofs, based on sample-path arguments, for both the

non-collaborative and collaborative models are straightforward but lengthy, we omit the proofs.

PROPOSITION 3.1. *The following hold:*

1. For all $x \in X$, $v_\alpha(x + e_i) \geq v_\alpha(x)$, $i = 1, 2$, where e_i denotes the i -th standard basis vector in \mathbb{R}^3 ($i = 1, 2, 3$). Similarly, if (g, w) is a solution to the ACOEs (defined below), the previous statement holds with v_α replaced with w .
2. Under the α -discounted cost (finite or infinite horizon) or the average cost criterion, there exists a (Markovian) non-idling policy that is optimal.
3. If $h_v = 0$, then, under the α -discounted cost (finite or infinite horizon) or the average cost criterion, there exists a (Markovian) non-idling policy optimal policy that always accepts arriving volunteers.

To obtain additional properties, we rely on the optimality equations. Proposition 3.1 allows us to restrict attention to non-idling policies. Before we can use the optimality equations, we must first provide conditions that guarantee that the optimality equations under each criterion have a solution. The proofs can be found in the Online Appendix. To simplify notation, let $c(x) = c(x_1, x_2, y) = \sum_{i=1}^2 h_i x_i + h_v y$ denote the cost rate. For any non-negative function v on X , define the mapping T such that

$$\begin{aligned} T v(x_1, x_2, y) = & \lambda_1 v(x_1 + 1, x_2, y) + \lambda_2 v(x_1, x_2 + 1, y) \\ & + y \beta v(x_1, x_2, y - 1) - (\lambda_1 + \lambda_2 + \beta) v(x_1, x_2, y) \\ & + \nu p_B [\min\{v(x_1, x_2, y + B), v(x_1, x_2, y) \\ & + KB\} - v(x_1, x_2, y)] + \min_{0 \leq a \leq y} \{\min\{a, x_1\} \\ & \mu_1 [v(x_1 - 1, x_2, y) - v(x_1, x_2, y)] \\ & + \min\{y - a, x_2\} \mu_2 [v(x_1, x_2 - 1, y) \\ & - v(x_1, x_2, y)]\}. \end{aligned}$$

Mapping T represents the one-step cost associated with current servers, capacity increase/decrease decisions, and a terminal cost v .

THEOREM 3.2. *Suppose $\alpha > 0$. The following hold:*

1. The function v_α satisfies the discounted cost optimality equations (DCOE), that is,

$$\alpha v_\alpha = c + T v_\alpha;$$

2. There exists a stationary, deterministic policy (depending on the discount factor α) that attains the minimum in the right hand side of the optimality equations, and hence, is discounted cost optimal.

PROOF. See Online Appendix A.1. □

THEOREM 3.3. *Suppose $\frac{\lambda_1}{\mu_1} + \frac{\lambda_2}{\mu_2} < \sum_{k=0}^{\infty} k \frac{(v/\beta)^k e^{-v/\beta}}{k!}$. The following hold:*

1. There exists a constant g and function w on the state space such that (g, w) satisfies the ACOEs,

$$g \vec{1} = c + T w,$$

where $\vec{1}$ is the vector of ones, g is the optimal average cost, and hence, is unique, and w , known as the relative value function, is unique up to additive constants.

2. A deterministic stationary policy is average cost optimal if and only if it satisfies the minimum in the ACOEs.

PROOF. See Online Appendixes A.1 and A.2. □

REMARK 1. (THEOREMS 3.2 AND 3.3). The results are stated for the non-collaborative model. The analogous results when servers can collaborate on a single job can be similarly shown. In the latter case, the terms

$$\begin{aligned} & \min_{0 \leq a \leq y} \{\min\{a, x_1\} \mu_1 [v(x_1 - 1, x_2, y) - v(x_1, x_2, y)] \\ & + \min\{y - a, x_2\} \mu_2 [v(x_1, x_2 - 1, y) - v(x_1, x_2, y)]\} \end{aligned}$$

in the mapping T defined above are replaced with

$$\begin{aligned} & \min_{0 \leq a \leq y} \{a \mu_1 [v(x_1 - 1, x_2, y) - v(x_1, x_2, y)] + (y \\ & - a) \mu_2 [v(x_1, x_2 - 1, y) - v(x_1, x_2, y)]\}, \end{aligned}$$

with similar modifications in the expressions for $q(x)$ and $q(x' | x, a)$ in Online Appendix A.1 and A.3 corresponding to the transition rate kernel. Lastly, the same ergodicity condition implies Theorem 3.3 for the collaborative model and is in fact easier to handle (see Remark 2 below).

REMARK 2. (THEOREM 3.3). We require an ergodicity condition $(\frac{\lambda_1}{\mu_1} + \frac{\lambda_2}{\mu_2} < \sum_{k=0}^{\infty} k \frac{(v/\beta)^k e^{-v/\beta}}{k!})$ in order for Theorem 3.3 to hold. Note that the condition is standard in that the “birth” rate is smaller than the “death” rate. The reason for this is that the number of servers in the system is random and, hence, in order for the system to be stable (i.e., have finite expected queue lengths) under at least one policy, the birth rate needs to be smaller than the product of the expected number of servers, given by $\sum_{k=0}^{\infty} k \frac{(v/\beta)^k e^{-v/\beta}}{k!}$, times the “death” rate. In particular, the condition

implies that a non-idling, exhaustive policy yields a stable Markov process, and hence, leads to finite average costs (c.f., Assumptions **D** in Online Appendix A.3). As a result, the vanishing discount approach, a standard technique for dealing with expected average cost by letting the discount rate decrease to zero in the α -discounted cost problem, can be used to obtain Theorem 3.3.

3.2. Volunteer Assignment

In this section, we consider how servers should be allocated to each queue. Our main result is that we provide conditions under which one particular customer class should be prioritized regardless of how many servers are in the system. In other words, we show the existence of optimal policies that are exhaustive in either queue 1 or queue 2. This is stated in Theorem 3.5 below. The following result will be useful.

PROPOSITION 3.4. *The following hold:*

1. For the non-collaborative model, if the number of customers at each queue exceeds the number of servers, there exists a discounted cost optimal control policy that allocates all volunteers to one station or the other (i.e., volunteers are not split between the two queues). Similarly, the result holds in the average cost case.
2. For the collaborative model, there exists a discounted cost optimal control policy that does not split the servers between the two queues. Similarly, the result holds in the average cost case.

PROOF. See Online Appendix A.3. □

THEOREM 3.5. *The following hold:*

1. Under the α -discounted cost criteria, if $\mu_1 h_1 \geq (\leq) \mu_2 h_2$, then it is optimal to serve class 1 (2) customers except to avoid unforced idling.
2. Under the average cost criteria, if $\frac{\lambda_1}{\mu_1} + \frac{\lambda_2}{\mu_2} < \sum_{k=0}^{\infty} k \frac{(v/\beta)^k e^{-v/\beta}}{k!}$ and $\mu_1 h_1 \geq (\leq) \mu_2 h_2$, then it is optimal to serve class 1 (2) customers except to avoid unforced idling.

REMARK 4. (THEOREM 3.5). The inequality $\mu_1 h_1 \geq (\leq) \mu_2 h_2$ implies the $c\mu$ rule for parallel queueing systems when there is a fixed and constant server in the system (Nain 1989). In the constant server case, the choice of which job to serve next is made based on the one that can reduce the cost the fastest—with the highest $c\mu$. The right (left)-hand side of the

inequality is the rate costs can be reduced if the volunteers are allocated to queue 1 (2).

REMARK 5. (THEOREM 3.5). The policy that prioritizes a particular customer class by no means ignores the other customer class. On the contrary, in the case that class 1 customers are prioritized, once all of the class 1 customers are served in the collaborative model, or the number of class 1 customers is less than the number of servers in the systems in the non-collaborative model, then, as a consequence of Propositions 3.1 and 3.4, servers will move to station 2.

REMARK 6. (THEOREM 3.5). We have shown the optimality of the $c\mu$ rule for both the collaborative and non-collaborative service models when there are two queues. Our proof of this result for the non-collaborative case (Online Appendix A.4) also applies if there are $n > 2$ queues; therefore, we conclude that the $c\mu$ rule is an optimal server assignment policy in general for the parallel queue control problem with random server arrivals and abandonments. In contrast, our proof for the collaborative service model does not readily generalize, so we cannot definitively say that the $c\mu$ rule is always an optimal policy for the collaborative service model with $n > 2$ queues (although anecdotally, the $c\mu$ server assignment policy was optimal for a handful of offline numerical examples with $n > 2$, which gives some indication that it might also be an optimal policy in general for the collaborative case).

3.3. Volunteer Admissions

In this section, we consider the problem of when to increase and decrease volunteers. Recall from Proposition 3.1, that if $h_v = 0$, then it is optimal to always accept volunteers. Our next task is to investigate the optimal server admission policy when $h_v > 0$.

One might conjecture that, for both the non-collaborative and collaborative models, if it is optimal to admit a volunteer in state (x_1, x_2, y) , then it is also optimal to admit a volunteer when there are more customers in the system in state $(x_1 + 1, x_2, y)$ or in state $(x_1, x_2 + 1, y)$. Similarly, one may conjecture that if it is optimal to reject a volunteer in state (x_1, x_2, y) , then it is also optimal to reject a volunteer when there are more servers in the system in state $(x_1, x_2, y + 1)$. In short, we would expect the optimal server admission policy to have a threshold or switching curve structure.

A standard approach to showing monotonicity in the queue length (assuming costs and service requirements are identical for both customer classes) is to show submodularity of the value function for the

discounted cost problem and submodularity for the relative value function for the long-run average cost problem. Similarly, to show monotonicity with respect to the number of volunteers (again assuming costs and the service requirements are identical for both customer classes), we can show convexity of the value and/or relative function with respect to the number of volunteers. In other words, submodularity and convexity are sufficient conditions for the results (though not necessary). In Figure 2 and in the Online Appendix A.5, we present examples where these two inequalities do not hold for truncated versions of our model. Figure 2 depicts the optimal policy for a non-collaborative service example; the results for the collaborative service example are similar and shown in Figure A.1 of Online Appendix A.5. As a result, the inequalities do not hold in general for the specific truncated versions of our model (i.e., where the arrival rates are set to 0 once the state space reaches the boundary). Note, however, the fact that this does not necessarily mean that they do not hold for the original, non-truncated model, since the recurrent structures are different as a result of the truncation. However, one approach to obtaining results for the original, non-truncated model is to consider truncated models (c.f., Down et al. (2011), Bhulai et al. (2014), Blok and Spieksma (2015, 2017)) and our examples rule out this particular approach. Additionally, the two examples hint that monotonicity with respect to the queue lengths and number of volunteers may not be true for the truncated model. That is, the optimal server admission policy is a complex state-dependent policy.

Because the optimal server admission policy is complicated, we propose a class of admission heuristics in the numerical study that depend on the number of workers in the system and the remaining system workload (see section 4). We compare our heuristic policies with the optimal one obtained from our MDP formulation and show that they perform well with respect to long-run average costs, waiting times, number of volunteers in the system, and number of volunteers idling in the system over a range of parameter values.

4. Numerical Study

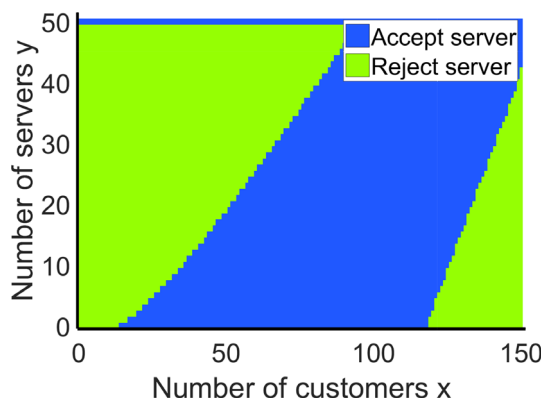
4.1. Rationale

We complement our theoretical work with a numerical study to further identify simple rules for managing volunteers in practical settings. For anyone managing volunteers, the central trade-off is between adding volunteers to reduce waiting times of beneficiaries vs. sending volunteers away to reduce costs. In section 3, we described the optimal *allocation* policy in a CTMDP formulation of a queueing system arising from volunteer convergence efforts, providing conditions under which simple, static priority rules are optimal for the non-collaborative and collaborative models. The CTMDP formulation seeks admission and allocation policies that minimize costs of beneficiary waiting times and number of volunteers. Hence, we used linear holding costs, since average holding costs for volunteers and beneficiaries are proportional to average number of servers and beneficiary queue length, respectively. The latter is proportional to the average delay of beneficiary due to Little’s Law. We allowed rejection costs since long-run average rejection costs for $K = 1$ represents the long-run average server rejection rate, an important performance measure to a VRC manager. For these reasons, allocating volunteers based on simple, static priority rules will perform well with respect to performance metrics (i.e., waiting times, reject rates) that are important when managing volunteers.

We did not, however, fully characterize the optimal server *admission* policy, which can be a complicated state-dependent policy (as examples A.8 and A.9 in Online Appendix A.5 illustrate). Furthermore, some assumptions of the CTMDP formulation are in line with what was observed in real life settings, but others may not be. Our numerical study picks up where our CTMDP model left off: assessing simple rules for admitting volunteers and incorporating more realistic assumptions.

Regarding the latter, recall that our work is motivated by the relief efforts following the 2011 tornado outbreak in Tuscaloosa, Alabama as reported in the case study by Lodree and Davis (2016). Volunteer

Figure 2 Optimal Policy for Truncated System Described in Online Appendix Example 0.8. The Optimal Server Admission Policy is not a Threshold Policy in the Number of Customers or the Number of Servers [Color figure can be viewed at wileyonlinelibrary.com]



Notes: Each box represents a particular state (x,y) of the system (x,y with $x \leq 150$ and $y \leq 50$) and the colors represent the policy at that state and time (blue = accept server; green = reject server). States are organized such that the number of customers (x) varies along the horizontal axis, the number of server (y) varies along the vertical axis.

convergence in Tuscaloosa was significant, and local officials reacted to the mass influx of volunteers by establishing the Tuscaloosa Area Volunteer Reception Center (TAVRC), whose goal was to centralize the process of registering volunteers and assigning them to relief activities taking place at various sites throughout the area. The study analyzed volunteer data collected from the TAVRC for this large-scale disaster relief effort, spanning a four month period following the tornado (May–August 2011). Data included volunteer names, dates of service, locations of relief activities, arrival and departure times each day, and target numbers of volunteers needed. The authors focused on volunteer material handling activities at a single relief warehouse site during the period May 1–May 31, 2011, resulting in approximately 2400 data points.

The following observations from the Lodree and Davis (2016) case study reveal several characteristics of real-world volunteer convergence that we incorporate into our model and explore through numerical experiments and simulation:

1. "...there is some evidence... that the exponential distribution can be used to model volunteer inter-arrival times" (see p. 1129 in Lodree and Davis 2016);
2. "the exponential distribution is not appropriate for modeling volunteer participation [abandonment] time" (see p. 1126 in Lodree and Davis 2016). The authors provide insights for when alternative distributions may be used for model volunteer inter-departure times. For example, they state that "a Weibull distribution with mean 6.134 and standard deviation 2.47 can be used to model volunteer participation times on the Sunday two weeks following the disaster event for individual volunteers who show up in the morning; the Wednesday 2 weeks after the disaster event, also for individual volunteers who arrive in the morning, etc." (see p. 1126 in Lodree and Davis 2016); and
3. "...batch volunteer arrivals should be considered when modeling relief warehouse queuing systems..." "In fact, the average number of group volunteers exceeded the average number of individual volunteers during the peak of the convergence period" (p. 1131 in Lodree and Davis 2016). In addition, we use the means and standard deviations shown in Table 5 on p. 1131 to represent volunteer batch sizes in our numerical experiments.

Lodree and Davis (2016) also found that "the average time between afternoon arrivals was generally greater than the average time between morning arrivals for both individual and group volunteers" (see p. 1128 in Lodree and Davis 2016). In other

words, there may be some time-dependent behavior related to the volunteer arrival process. If the resulting arrival process can be accurately described by non-homogeneous Poisson processes with a piecewise constant rate function, then for each interval where the arrival rate function is constant, the system is akin to a homogeneous Poisson process. In this case, our approach can provide insights specifically tailored to each interval. This methodology has been shown very useful in supporting emergency medical service staffing (c.f., Green et al. 2007).

Lastly, note that the above features are taken into account within the context of queuing systems with $n = 2$ parallel queues. As discussed at the beginning of section 3, the two-queue scenario is plausible from a practical perspective, and also aligns with the presentation of our analytical results.

4.2. Overview

Our numerical study is divided into two parts studies 1 and 2. Study 1 compares heuristic admission policies to the optimal policy from the CTMDP formulation with respect to long-run average costs. This is done to benchmark the proposed policies with the optimal ones. For this part, we truncate the state space so that the resulting state space is finite, calculate optimal policies for this truncated model using modified policy iteration (c.f., Puterman and Shin (1978), Puterman (1994)), and compare the average costs of the optimal policies with those of our proposed heuristics. In particular, we truncate the state space to be a rectangular state space so that class 1 customers, class 2 customers, and servers are blocked from arriving if the number class 1 customers, class 2 customers, and servers, respectively, equals truncation values of 20, 20, and 15.

While minimizing long-run average costs of the system is a primary concern, managers may prefer to increase costs (by allowing more volunteers in the system, say) in favor of having lower waiting times for beneficiaries. For example, minimizing beneficiary waiting times may be desirable to safeguard getting timely access to supplies to those who have serious needs. Lastly, calibration of customer and server holding costs may be difficult. These considerations form the basis for study 2.

In study 2, admission heuristic policies are compared with respect to class 1 and class 2 average waiting times, average volunteer congestion, and average volunteer idling. We extend our CTMDP analysis in three ways based on the observations reported in the Lodree and Davis (2016) case study: (i) server inter-departure times are assumed to be generally distributed; (ii) servers arrive in batches; and (iii) servers leave in batches. Generally distributed inter-departure times implies that the CTMDP framework can no

longer be used since the system is no longer Markovian. Moreover, batch server arrivals and departures yield an intractable state space. As a result, we use a discrete-event simulation. We used a simulation length of six years, using only the last three years for analysis, and performed 100 replications for each set of parameters and policies. We calculated average performance measures for each replication and then averaged these hourly metrics over the 200 simulations and estimated standard errors. All simulations were performed in MATLAB.

For both studies, our volunteer allocation policy is fixed to be the $c\mu$ rule, which we proved was optimal under certain conditions in Theorem 3.5. This allocation policy prioritizes customer classes in decreasing order of $h_i\mu_i$. Due to space considerations and because the non-collaborative model is harder to analyze and simulate, we assume that service is non-collaborative in both studies.

4.3. Heuristic Admission Policies

Having characterized the structure of the optimal allocation policy, we propose simple rules for admitting volunteers. We examine a class of workload-dependent threshold policies for the server admission policy as an alternative to the more complicated state-dependent optimal server admission policy from our CTMDP. A primary advantage of the heuristic policies is that they apply more broadly, to settings where the assumptions of the CTMDP model (such as exponential volunteer participation times) do not hold. The purpose of our study is to analyze how the proposed policies would perform in practice with respect to long-run average performance measures. The central question is whether the proposed policies can achieve relatively low “costs” from providing service, while safeguarding against long waits for beneficiaries. If so, they can then serve to guide how volunteer managers should control volunteer congestion and allocate volunteers between tasks.

We will consider several classes of related server admission policies:

Single threshold policies: Ignore the current workload. Upon a server arrival, accept all servers in the batch if and only if there are fewer than T servers already on hand, otherwise reject the arriving servers.

Two threshold policies: Upon a server arrival, calculate the workload W (i.e., expected time it would take to clear the system if there was only one constant server available) in the system. If $W \leq w$, accept all servers in the batch if and only if there are fewer than T servers on hand. Otherwise if $W > w$, accept all servers in the batch if and only if there are fewer than T_{high} servers on hand. We

refer to these policies as two threshold policies even though w may be viewed as third threshold. *Accounting for batches:* Calculate the workload W . If $W \leq w$, accept servers if there are fewer than T servers on hand so long as the arriving batch size does not increase the total number of servers above $T' \geq T$, otherwise, reject servers. If $W > w$, accept servers only if there are fewer than T_{high} so long as the arriving batch size does not increase the total number of servers over $T'_{\text{high}} \geq T_{\text{high}}$.

For the two threshold policies, the parameters w , T , and T_{high} allow the admission policy to dynamically respond to the workload. When the workload is relatively high, more volunteers will be admitted if $T_{\text{high}} > T$. This structure is similar to the “2-level heuristics” proposed in Kaufman et al. (2005) for tandem queues (without batch arrivals). With batch arrivals, the total number of volunteers cannot necessarily be raised in unit increments. One may have to accept more volunteers at once than what otherwise would be ideal if fractional batches were allowed. The additional parameters T' and T'_{high} are meant to account for this friction. When $T' > T$ and $T'_{\text{high}} > T_{\text{high}}$, we will refer to this as a two threshold policy that “accounts for batches” (even though there are 5 parameters to tune). A special case with $T' > T$ but $w = \infty$ (so that T'_{high} and T_{high} are irrelevant) will be referred to as a single threshold policy that accounts for batches. In any case, it is important to note that both studies assume batch volunteer arrivals, and even single threshold policies (with $T' = T$) actually account for batches when T is tuned.

4.4. Parameters

Volunteer inter-arrival times, abandonment times, and batch sizes are chosen to match the empirical analyses presented in the Lodree and Davis (2016) case study. In particular, we model volunteer inter-arrival and abandonment times using the probability distributions and parameters Lodree and Davis (2016) fit to real data, and we represent mean batch sizes of volunteer arrivals and departures based on their analyses. Model parameters are summarized in Table 1 for both parts of the study with a more detailed discussion to follow.

Beneficiary process and service times—The focus of Lodree and Davis (2016) is on characterizing the behaviour of the volunteer convergence process only. Thus, there are no estimates for beneficiary inter-arrival and service times. Instead, we let class 1 customers be those that arrive more quickly on average (e.g., $\lambda_1 = 12,90$ class 1 customers per hour for the simulation study) and take less time to serve on average (e.g., $\frac{60}{15} = 4$ minutes for the simulation study)

Table 1 Parameters for Comparison with Continuous Time Markov Decision Process Model (study 1) and for Discrete-Event Simulation (study 2)

Parameter	Value(s)	
	Study 1	Study 2
λ_1	$12^a, 90$	$12^a, 90$
λ_2	2	2
μ_1	15	15
μ_2	4	4
ν	3.42	1.71
ρ_B	$\frac{1}{B}, B \in \{1,2,3\}$	$\frac{1}{B}, B \in \{1, \dots, 9\}$
β	$\frac{1}{6.134}$	n/a
h_1	1	n/a
h_2	1,2	n/a
h_ν	0.5,1	n/a
K	0	n/a

^aResults in the appendix.

and class 2 customers those that arrive less quickly on average (e.g., $\lambda_2 = 2$ class 2 customers per hour for the simulation study) and take more time to serve (i.e., $\frac{60}{4} = 15$ minutes for the simulation study). Let $S(\lambda_1, \lambda_2, \mu_1, \mu_2) := \frac{\lambda_1}{\mu_1} + \frac{\lambda_2}{\mu_2}$, the quantity on left-hand side of the inequality in the ergodicity condition in Theorem 3.3. The aforementioned beneficiary arrival process and service time values correspond to values of $S(\lambda_1, \lambda_2, \mu_1, \mu_2) := \frac{\lambda_1}{\mu_1} + \frac{\lambda_2}{\mu_2} = \{1.3, 6.5\}$.

Volunteer arrival process—For some of the cases analyzed in Lodree and Davis (2016), the average time between individual volunteer arrivals (batch size = 1) was approximately 35 minutes, and the average time between group volunteer arrivals (batch size > 1) was about 45 minutes (see Figure 7 in Lodree and Davis (2016)). As such, we assume in Study 1 that volunteer inter-arrival times are exponential with rate $\nu = 2 \cdot \frac{60}{35} \approx 3.42$ volunteers per hour. For Study 2, we assume that volunteer inter-arrival times are exponential with rate $\nu = \frac{60}{35} \approx 1.71$ volunteers per hour. In both studies, we assume that batch sizes are uniformly distributed as no batch distribution was provided in Lodree and Davis (2016). However, the mean batch size was roughly 3 for one of the scenarios they considered. Thus, we assume a support of $\{1,2,3\}$ for the uniform distribution in Study 1 while in Study 2, we extend the support to $\{1, \dots, 9\}$.

Volunteer departure process—For Study 1, we assume that, once admitted, each volunteer remains in the system for an exponentially distributed amount of time with mean 6.134 hours (i.e., there are no batch departures). When we consider the case without batches (i.e., only individual workers arrive to the system), the latter implies that the average number of servers in the system when all individual volunteers are admitted is approximately 10.49 volunteers per hour. For Study 2, we assume that arriving batches also leave in batches. For individuals, a Weibull

distribution with mean 6.134 hours and standard deviation 2.47 is assumed, which was shown to be a good fit for volunteer participation times in one of the scenarios considered in the case study. For batch arrivals greater than 2, a Weibull distribution with mean 5 hours and standard deviation 2.47 is assumed regardless of the batch size, which was also shown to be a good fit in the case study.

Costs and performance metrics—For Study 1, we fix $h_1 = 1$, $h_2 = 1, 2$, and $h_\nu = 0.01, 0.05$ and let $K = 0$. This implies, for instance, that the long-run average holding cost for class 1 customers is the long-run number of class 1 customers per unit time. It also implies that the allocation policy is to prioritize class 1 customers except to avoid unforced idling so that the class 1 average waiting times are always lower than those of class 2. We do not penalize rejections; $K = 0$. For Study 2, our performance metrics include long-run average waiting times for beneficiaries and long-run average number of volunteers idling.

4.5. Study 1: Costs of Admission Heuristic Policies vs. Continuous Time Markov Decision Process Formulation

The purpose of study 1 is to benchmark the heuristic policies against the optimal CTMDP policy. For study 1, values for each of the admission heuristic parameters were chosen as follows. First, we focused on single threshold policies. For each set of parameter values in Table 1, we found the value of T , denoted by T^* , within a finite range of values that minimized long-run average costs. Second, we considered optimal two threshold policies, over a discrete set of w with increments of 0.5. Tables 2 and A.1 present the optimal values in bold. For this study, the optimal policies all happen to have a “gap” $T_{\text{high}} - T$ equal to 1 (though the search space included larger gaps), and either T_{high} or T equal to T^* (though this was not a restriction of the search). The tables also present some nearby policies that depend on workload by varying the workload threshold w with either $T = T^* - 1$ and $T_{\text{high}} = T^*$ or $T = T^*$ and $T_{\text{high}} = T^* + 1$. Third, we consider single threshold policies that account for batches, with $T' = T + 1$. Fourth, we consider two threshold policies that account for batches, with $T' = T + 1$ and $T'_{\text{high}} = T_{\text{high}} + 1$. We remark that searching over a larger set of parameters when accounting for batches could possibly yield even better performance.

We report the percentage deviation of the heuristic admission policies from the optimal policy with respect to long-run average costs. Here, we consider the higher customer arrival rate $\lambda_1 = 90$ in Table 2 for different parameter values of h_2 and h_ν . Additional scenarios are reported in the Online Appendix with a lower customer arrival rate $\lambda_1 = 12$ with essentially the same trends except (Table A.1 in A.6).

Table 2 Percent Away from Optimal for Admission Heuristic Policies in Terms of Average Cost; $\lambda_1 = 90$

w	T	T_{high}	$\lambda_1 = 90$			
			$h_2 = 1$		$h_2 = 2$	
			$h_v = 0.5$	$h_v = 1$	$h_v = 0.5$	$h_v = 1$
n/a	8	n/a	17.936%	7.288%	17.212%	7.072%
n/a	9	n/a	7.987%	2.169%	7.664%	2.104%
n/a	10	n/a	3.636%	0.769%	3.489%	0.746%
n/a	11	n/a	2.249%	1.112%	2.158%	1.079%
n/a	12	n/a	2.072%	1.886%	1.989%	1.830%
n/a	13	n/a	2.211%	2.489%	2.122%	2.415%
n/a	14	n/a	2.339%	2.810%	2.244%	2.727%
n/a	15	n/a	2.402%	2.942%	2.305%	2.854%
n/a	16	n/a	2.402%	2.942%	2.305%	2.854%
n/a	17	n/a	2.402%	2.942%	2.305%	2.854%
2.5	$T^* - 1$	T^*	2.249%	2.168%	2.158%	2.104%
2	$T^* - 1$	T^*	2.245%	2.146%	2.154%	2.083%
1.5	$T^* - 1$	T^*	2.207%	1.903%	2.118%	1.847%
1	$T^* - 1$	T^*	2.099%	1.203%	2.014%	1.167%
0.5	$T^* - 1$	T^*	2.072%	0.769%	1.989%	0.746%
n/a	T^*	T^*	2.072%	0.769%	1.989%	0.746%
2.5	T^*	$T^* + 1$	2.072%	0.769%	1.989%	0.746%
2	T^*	$T^* + 1$	2.071%	0.764%	1.988%	0.741%
1.5	T^*	$T^* + 1$	2.070%	0.727%	1.986%	0.706%
1	T^*	$T^* + 1$	2.121%	0.749%	2.035%	0.727%
0.5	T^*	$T^* + 1$	2.211%	1.112%	2.122%	1.079%

Notes: Top half shows single threshold policies; bottom half shows two threshold policies displayed in terms of the optimal single threshold T^* . Bold indicates deviation in objective function of best performing threshold values from that of the optimal policy.

Note that $T^* = 12(10)$ when $h_v = 0.5(1)$.

Single threshold policies: The top half of Table 2 shows single threshold policies. Note that for each parameter value of h_2 and h_v , there is a parameter value of T for which the corresponding proposed admission policy have average costs that are close to the optimal policy. When $h_2 = 1, h_v = 0.5$, the admission policies with $T = 12$ performs the best and is within 2.072% of the optimal value. The same admission heuristic performs the best and is within 1.989% of the optimal value when h_2 is increased to 2 (i.e., when $h_2 = 2, h_v = 0.5$). When $h_2 = 1, h_v = 1$, the admission policy with $T = 10$ performs the best and is within 0.769% of the optimal value and when h_2 is increased to 2 (i.e., $h_2 = 2, h_v = 1$), the same admission policy performs the best and is within 0.746% of the optimal value. Note that for the single threshold policy with $T = 8$, the system is unstable (i.e., the expected long-run average queue lengths for both customers classes may be arbitrarily large) as a consequence of the relatively extremely high class 1 arrival rate. As a result, we observe extremely high long-run average waiting times for both customer classes, specifically, waiting times of 1095 hours as shown in Table 3. So, single threshold policies may not always perform well. However, when the threshold is relatively near the best single threshold, we find in this study that the best single threshold policies perform very well relative to the optimal policies.

Two threshold policies: The bottom half of Table 2 shows two threshold policies, with the optimal values in bold. When searching over different values w , we see slight improvements over using a single threshold policy. For the various parameter settings, the optimal policies here all have $w = 1.5$ and $T = T^* - 1$, with T^* equal to either 10 or 12. The optimal gap happens to be 1 here, though the search was over larger gaps as well. The optimal policy is within 2.070% of the optimal value when $h_2 = 1, h_v = 0.5$, 1.986% when $h_2 = 2, h_v = 0.5$, 0.727% when $h_2 = 1, h_v = 1$, and 0.706% when $h_2 = 2, h_v = 1$. Other results in the bottom half of Table 2 are similar, indicating that it might not be necessary in practice to find the exact best heuristic parameter settings.

Accounting for batches: We repeated the same procedures that produced the results shown in Tables 2 and A.1 but accounted for batch server arrivals by setting $T' = T + 1$ and $T'_{\text{high}} = T_{\text{high}} + 1$. The results are reported in the Appendix, Tables A.2 and A.3 in A.6. For example, when $\lambda_1 = 90$, the best single threshold policy that accounts for batch server arrivals (top half of Table A.2) outperforms the best single threshold policy that does not account for batch server arrivals (top half of Table 2) in two out the four instances. The other two instances are basically ties. Similarly, the best two threshold policies that account for batches (bottom half of Table A.2) only slightly

Table 3 Long-Run Average (standard error) Waiting Time for Class 1 and Class 2 Customers and Long-Run Average (standard error) Number of Volunteers and Number of Idling Volunteers When $\lambda_1 = 90$

		$\lambda_1 = 90$					
w	T	T_{high}	Waiting times, hrs		Number of volunteers		Objective function Value
			Class 1	Class 2	Total	Idling	
n/a	8	n/a	1095 (0.146)	1094.6 (0.883)	0 (0)	0 (0)	1094.600
n/a	9	n/a	0.991 (0.02)	64.849 (4.793)	6.6 (0.005)	0.098 (0.005)	64.861
n/a	10	n/a	0.397 (0.006)	4.267 (0.083)	7.281 (0.005)	0.784 (0.005)	4.365
n/a	11	n/a	0.244 (0.003)	1.703 (0.024)	7.901 (0.006)	1.4 (0.006)	1.878
n/a	12	n/a	0.184 (0.002)	1.024 (0.012)	8.46 (0.007)	1.963 (0.007)	1.269
n/a	13	n/a	0.152 (0.002)	0.74 (0.007)	8.974 (0.007)	2.475 (0.007)	1.049
n/a	14	n/a	0.133 (0.002)	0.585 (0.007)	9.44 (0.008)	2.937 (0.008)	0.952
n/a	15	n/a	0.124 (0.001)	0.507 (0.005)	9.837 (0.007)	3.335 (0.007)	0.924
n/a	16	n/a	0.116 (0.002)	0.444 (0.005)	10.183 (0.009)	3.681 (0.009)	0.904
n/a	17	n/a	0.11 (0.001)	0.407 (0.004)	10.483 (0.01)	3.984 (0.009)	0.905
4.5	15	19	0.117 (0.001)	0.44 (0.004)	10.065 (0.008)	3.563 (0.008)	0.885

Notes: Top part shows single threshold policies; bottom row shows the optimal heuristic policy. As it turned out, despite T' and T'_{high} being allowed to differ from T and T_{high} , the optimal policy has $T = T' = 15$ and $T_{\text{high}} = T'_{\text{high}} = 19$. Bold indicates objective function values for best performing one and two threshold policies, respectively.

outperform the two threshold policies that do not account for batches (bottom half of Table 2). When $\lambda_1 = 12$, we again find that the threshold policies that account for batches (Table A.3) only slightly outperform the threshold policies that do not account for batches (Table A.1). For this study then, one may conclude that the added complexity of threshold policies that account for batches over those that do not might not be worth it in practice. We explore this point further in study 2 below.

Sensitivity analysis: We conducted a wider sensitivity analysis by varying one parameter at a time from a base set of parameters. The base parameters are $\lambda_1 = 60$, $\lambda_2 = 2$, $\mu_1 = 15$, $\mu_2 = 4$, $v = 3.42$, $\beta = 0.163$, $h_1 = 1$, $h_2 = 2$, and $h_v = 1$. We fixed λ_2 and h_1 and varied the other parameters relative to these two. The results which are presented in Tables A.12–A.15 in Appendix A.6, suggest that both the single threshold and two threshold policies perform well relative to the optimal policy found through value iteration. On average, the best single threshold policies are within an average of 1.12% (max of 2.67%) of the optimal average cost, and the best two threshold policies reduce the percentage slightly by an additional 0.12% on average. Also, single threshold policies may perform nearly as well as two threshold policies at least for some settings. We further explore the performance gap between the single threshold and two threshold policies in study 2 (section 4.6), and find scenarios where the gap is more pronounced.

Alternative stability heuristic: Given that the threshold policies perform well when the thresholds are appropriately tuned, one may wonder whether any type of reasonable server admission policy would perform well. For the sake of comparison, we also

explored a stability heuristic that works as follows: calculate $\lambda_1/\mu_1 + \lambda_2/\mu_2$ and find the minimum value of s , corresponding to the number of servers, such that $\frac{\lambda_1}{\mu_1} + \frac{\lambda_2}{\mu_2} < \sum_{k=0}^s k \frac{(v/\beta)^k e^{-v/\beta}}{k!}$; see Theorem 3.3 and Remark 2. At each server arrival time, if the number of servers is less than s , then accept servers, and otherwise reject. In other words, set the admission threshold to the smallest number of servers that stabilizes the system. The results of the stability heuristic for the sensitivity study parameters are also presented in Tables A.12–A.15 in Appendix A.6. The percent away from optimal ranges from 9.89% to 110.40%, with an average of 49.17%. We conclude that not just any heuristic will perform nearly as well as the threshold policies.

4.6. Study 2: Discrete-Event Simulation

Study 2 incorporates additional characteristics of real-world volunteer convergence by considering observations and empirical results from the Lodree and Davis (2016) case study, some of which violate the assumptions required for the CTMDP framework. Specifically, study 2 considers batch server arrivals, batch server departures, as well as distributions and parameters for volunteer inter-arrival times and non-exponentially distributed participation times. In practice, simulation may be used to tune the parameters of the heuristic policies to optimize desired objectives. Here, we report the long-run average waiting times for class 1 and class 2 customers, the long-run average number of volunteers in the system, and the long-run average number of volunteers idling in the system. We also report an objective value that is the weighted sum of average waiting time of class 2 customers (i.e.,

lower priority customers) and average number of idling volunteers, where weights were chosen so that one idling server on average equaled 7.5 minutes of average class 2 waiting time. We consider the higher customer arrival rate $\lambda_1 = 90$ in Table 3 for different parameter values of h_2 and h_v . Additional scenarios are reported in the Online Appendix with a lower customer arrival rate $\lambda_1 = 12$ (Table A.11 in A.7).

Optimal heuristic policies: We report the performance for single threshold policies and for overall optimal threshold policies where the parameters w , T , T' , T_{high} , and T'_{high} , with $T' - T = T'_{\text{high}} - T_{\text{high}}$, were found through a search on a broad space but with w restricted to increments of 0.5. As it turned out, for both $\lambda_1 = 90$ and $\lambda_1 = 12$ the respective best heuristic policies have $T = T'$ and $T_{\text{high}} = T'_{\text{high}}$, despite being allowed to differ; two threshold policies are best. Relative to the best single threshold policies (top halves of Tables 3 and A.11), the best two threshold policies perform 2.15% and 3.48% better, respectively. As it turned out, the overall optimal threshold policies are two threshold policies that do not account for batches. That is, accounting for batches did not improve performance for this study. From a practical standpoint, the two threshold policies that account for batches have increased complexity in terms of execution and in computation time, since there are five tuning parameters. With simulation being used for performance evaluation, the additional search space may complicate implementation from a computational standpoint enough to avoid them altogether. Luckily, single and two threshold policies, that “do not account for batches,” are very promising. Next we explore the sensitivity of single and two threshold policies in a wider study.

Sensitivity analysis: We again conducted a sensitivity analysis by varying one parameter at a time while keeping the others fixed, finding the best performing single and two thresholds for each instance with respect to our objective value. Results are summarized in Tables A.12–A.15 in Appendix A.7. The base set of parameters are $\lambda_1 = 48$, $\lambda_2 = 2$, $\mu_1 = 15$, $\mu_2 = 4$, $v = 1.71$, and a maximum batch size of 9. We fixed λ_2 and μ_2 , and varied λ_1 , μ_1 , v , and the maximum batch size. The true optimal solution is not known (as value iteration does not apply), so we report the percentage difference between the best single threshold and the best two threshold performances. As in study 1, we find that the two threshold policies are relatively close in performance to the single threshold policies in all cases, but the difference is larger in study 2. The average difference in performance is 3.25% and the maximum difference in performance is 5.88%. Thus the added benefits of the two threshold policies may be worth exploring.

5. Conclusion

Volunteer convergence poses logistical challenges that must be managed effectively to ensure that beneficiaries affected by disaster receive aid as quickly as possible. As such, this study investigates policies for accepting and rejecting spontaneously arriving volunteers into organized relief efforts, and for allocating admitted volunteers among relief tasks. Unlike traditional labor assignment or volunteer scheduling in non-crisis-relief settings, the management of volunteer convergence is characterized by uncertain arrivals and abandonments associated with the labor pool. Volunteer convergence management can also be distinguished from other forms of labor assignment in that worker admission decisions also become relevant.

In this study, we examine volunteer convergence within the context of a queueing system characterized by not only customer arrival and service time uncertainties, but also stochastic server arrivals and abandonments. We first analyze the resulting control problem using an MDP model. We show that the classical $c\mu$ rule is both discounted and average cost optimal under collaborative service but also under non-collaborative service scenarios, and when batch server arrivals are allowed. We also demonstrate that the optimal server admission policy is complex. Our analysis requires us to deal with the technical challenges of unbounded transition rates, which emerge as a result of considering a queueing system with random server abandonments. We do so by truncating the number of servers allowed in the system, analyzing this truncated system, and then letting the truncation go to infinity. Finally, we propose a class of admission heuristic policies that depend on the system workload, number of volunteers, and accounts for batch arrivals. We compare their performance with the optimal policy in a numerical experiment when inter-arrival, service, and abandonment times are exponentially distributed and show that, for a given instance of the model (i.e., rates and costs), we can find parameter values for our proposed class of admission heuristics that are close to optimal. We then compare their relative performance in a simulation study with respect to long-run average customer waiting times, number of servers, and number of idling servers in the system when servers arrive and leave in batches and when server departure times are no longer exponential. Our results show that we can once again find parameter values for the admission heuristic that performs well relative to these performance metrics. These results are encouraging because it shows that a practical and intuitive policy can be used to effectively manage the complexities of volunteer convergence.

The queueing framework presented in this study for addressing volunteer convergence introduces several interesting possibilities for future research, both from theoretical and practical perspectives. On the theoretical side, queueing networks with random server arrivals and abandonments and under complicated server assignments are unique in their own right and are generally not well understood. As such, basic analysis of these queueing systems, namely characterization of the steady state number of customers and servers in the system, would be a logical first step to further advance the queueing literature in this direction. From the control perspective, analyzing and/or approximating optimal control policies when there are set-up times and/or costs for server reallocations is worthy of consideration. In line with what is reported in the Lodree and Davis (2016) case study, we have assumed that there is no “set-up” associated with re-allocation decisions, but there may be other situations where travel or set-up times between tasks could be relevant. Stochastic scheduling for queueing systems with a fixed server capacity and set-ups have been considered (c.f., Hofri and Ross (1987), Van Oyen et al. (1992), Van Oyen and Teneketzis (1994), Duenyas and Van Oyen (1995, 1996), Reiman and Wein (1998)) and are generally hard to analyze. These studies do not directly translate to our setting as a consequence of the volunteer arrival and departure process. We have also assumed throughout that the time volunteers spend in the system once admitted is independent of the task. There may be cases for which volunteers prefer helping on some tasks over others and thus it makes sense to assume that the abandonment rates are task dependent. This added generality comes as the cost of tractability, since the state space would increase from three to four dimensions. Analysis and/or approximation of optimal control strategies when distributions are not longer exponential, and there are more than two customer classes, are all also topics worthy of consideration.

This study can also be extended to include additional features of post-disaster response and relief observed in practice. One such observation is that volunteer convergence does not occur in isolation of other relief activities such as emergency supply positioning or last-mile distribution. In particular, we believe that volunteer convergence is very closely related to material convergence, and therefore studies that consider joint management of volunteer and material convergence would be beneficial. Furthermore, there are very few studies in the academic literature that address material convergence or volunteer convergence from an operations management perspective, so studies that consider both would represent noteworthy research contributions.

Notes

¹Volunteer convergence is a special case of personal convergence, which is defined in Fritz and Mathewson (1957) as the mass movement of people to disaster areas. Besides volunteers (or helpers), the other groups of people that gravitate toward disaster sites are labeled as the returners, the anxious, the curious, and the exploiters.

²Note that although Figure 1 shows two customer queues/classes, there will very likely be more than two queues/classes in many practical situations. The description of the parallel queueing system that follows applies to the general case of two or more parallel queues.

³In recognition that participating in relief efforts can be therapeutic for survivors as mentioned in Lowe and Fothergill (2003) and Hamerton et al. (2015), we assume that volunteers who are not admitted means that they are assigned to ancillary duties that do not contribute holding costs, or, as in large-scale events such as Hurricane Katrina, they go to another part of the city to work.

⁴Depending on the context, a workstation could be an actual queueing station at a relief center or one of many affected areas.

References

- Abualkhair, H., E. J. Lodree, L. B. Davis. 2020. Managing volunteer convergence at disaster relief centers. *Int. J. Prod. Econ.* **220**: 107399.
- Ahn, H.-S., I. Duenyas, M. E. Lewis. 2002. Optimal control of a two-stage tandem queueing system with flexible servers. *Prob. Eng. and Inform. Sci.* **16**(04): 453–469.
- Allen, B. 1969. *Communities in Disaster: A Sociological Analysis of Collective Stress Situations*. Doubleday, Garden City, New York.
- Altay, N., W. G. Green. 2006. OR/MS research in disaster operations management. *Eur. J. Oper. Res.* **175**(1): 475–493.
- Anaya-Arenas, A. M., J. Renaud, A. Ruiz. 2014. Relief distribution networks: A systematic review. *Ann. Oper. Res.* **223**(1): 53–79.
- Baras, J., A. Dorsey, A. Makowski. 1985. Two competing queues with linear costs and geometric service requirements: The μ -rule is often optimal. *Adv. Appl. Prob.* **17**(01): 186–209.
- Barnett, R. V., C. G. Flint. 2005. Community volunteers: The front line of disaster response. *J. Volunt. Adm.* **23**(23): 52–56.
- Barsky, L. E., J. E. Trainor, M. R. Torres, B. E. Aguirre. 2007. Managing volunteers: FEMA’s Urban Search and Rescue programme and interactions with unaffiliated responders in disaster response. *Disasters.* **31**(4): 495.
- Bell, S. L., R. J. Williams. 2001. Dynamic scheduling of a system with two parallel servers in heavy traffic with resource pooling: asymptotic optimality of a threshold policy. *Ann. Appl. Prob.* **11**(3): 608–649.
- Bhulai, S., A. C. Brooms, F. M. Spieksma. 2014. On structural properties of the value function for an unbounded jump Markov process with an application to a processor sharing retrieval queue. *Queueing Syst.* **76**(4): 425–446.
- Blok, H. 2016. Unbounded-rate Markov decision processes: structural properties via a parametrisation approach. PhD thesis, Universiteit Leiden.
- Blok, H., F. Spieksma. 2015. Countable state Markov decision processes with unbounded jump rates and discounted cost: Optimality equation and approximations. *Adv. Appl. Prob.* **47**(4): 1088.
- Blok, H., F. Spieksma. 2017. Structures of optimal policies in MDPS with unbounded jumps: The state of our art. In Richard J.

- Boucherie and Nico M. van Dijk (Eds.), *Markov Decision Processes in Practice*, Cham, Springer, 131–186.
- Budhiraja, A., A. Ghosh, X. Liu. 2014. Scheduling control for Markov-modulated single-server multiclass queueing systems in heavy traffic. *Queueing Syst.* 78(1): 57.
- Burkart, C., M. Besiou, T. Wakolbinger. 2017. The funding–Humanitarian supply chain interface. *Surv. Oper. Res. Manag. Sci.* 21(2): 31–45.
- Buyukkoc, C., P. Variaya, J. Walrand. 1985. $c\mu$ rule revisited. *Adv. Appl. Prob.* 17(1): 237–238.
- Caunhye, A. M., X. Nie, S. Pokharel. 2012. Optimization models in emergency logistics: A literature review. *Socio-Econ. Plan. Sci.* 46(1): 4–13.
- Cnaan, R. A., F. Handy. 2005. Towards understanding episodic volunteering. *Vrijwillige Inzet Onderzocht.* 2(1): 29–35.
- Cone, D. C., S. D. Weir, S. Bogucki. 2003. Convergent volunteerism. *Ann. Emerg. Med.* 41(4): 457–462.
- Cox, D., W. Smith. 1961. *Queues*. Methuen, New York.
- Down, D. G., G. Koole, M. E. Lewis. 2011. Dynamic control of a single-server system with abandonments. *Queueing Syst.* 67(1): 63–90.
- Duenyas, I., M. P. Van Oyen. 1995. Stochastic scheduling of parallel queues with set-up costs. *Queueing Syst.* 19(4): 421–444.
- Duenyas, I., M. P. Van Oyen. 1996. Heuristic scheduling of parallel heterogeneous queues with set-ups. *Management Sci.* 42(6): 814–829.
- Falasca, M., C. Zobel. 2012. An optimization model for volunteer assignments in humanitarian organizations. *Socio-Econ. Plan. Sci.* 46(4): 250–260.
- Fernandez, L. S., J. A. Barbera, J. R. van Dorp. 2006a. Spontaneous volunteer response to disasters: The benefits and consequences of good intentions. *J. Emerg. Manag.* 5(4): 57–68.
- Fernandez, L. S., J. A. Barbera, J. R. van Dorp. 2006b. Strategies for managing volunteers during incident response: A systems approach. *Homel. Secur. Aff.* 2(3): 1–15.
- Fritz, C. E., J. H. Mathewson. 1957. Convergence behavior in disasters: a problem in social control: a special report prepared for the committee on disaster studies. National Academy of Sciences National Research Council.
- Fulmer, T., I. Portelli, G. L. Foltin, R. Zimmerman, E. Chachkes, L. R. Goldfrank. 2007. Organization-based incident management: Developing a disaster volunteer role on a university campus. *Disast. Manag. Resp.* 5(3): 74–81.
- Galindo, G., R. Batta. 2013. Review of recent developments in OR/MS research in disaster operations management. *Eur. J. Oper. Res.* 230(2): 201–211.
- Gordon, L., E. Erkut. 2004. Improving volunteer scheduling for the Edmonton Folk Festival. *Interfaces* 34(5): 367–376.
- Green, L. V., P. J. Kolesar, W. Whitt. 2007. Coping with time-varying demand when setting staffing requirements for a service system. *Prod. Oper. Manag.* 16(1): 13–39.
- Guo, X., O. Hernández-Lerma. 2009. *Continuous-Time Markov Decision Processes*. Springer, Berlin.
- Gupta, S., M. K. Starr, R. Z. Farahani, N. Matinrad. 2016. Disaster management from a POM perspective: Mapping a new domain. *Prod. Oper. Manag.* 25(10): 1611–1637.
- Habib, M. S., Y. H. Lee, M. S. Memon. 2016. Mathematical models in humanitarian supply chain management: A systematic literature review. *Math. Prob. Eng.* 2016: 1–20.
- Hamerton, H., R. J. Sargisson, K. Smith, S. Hunt. 2015. How volunteering reduced the impact of the Rena oil spill: Community responses to an environmental disaster. *Int. J. Mass Emerg. Disasters* 33(2): 1–20.
- Hofri, M., K. W. Ross. 1987. On the optimal control of two queues with server setup times and its analysis. *SIAM J. Comput.* 16(2): 399–420.
- Holguín-Veras, J., N. Pérez, M. Jaller, L. N. Van Wassenhove, F. Aros-Vera. 2013. On the appropriate objective function for post-disaster humanitarian logistics models. *J. Oper. Manag.* 31(5): 262–280.
- Hoyos, M. C., R. S. Morales, R. Akhavan-Tabatabaei. 2015. OR models with stochastic components in disaster operations management: A literature survey. *Comput. Ind. Eng.* 82: 183–197.
- Kaspari, M. 2010. Optimal volunteer assignment with an application to the Denver B-cycle bike sharing program. PhD thesis, University of Colorado Denver.
- Kaufman, D. L., H.-S. Ahn, M. E. Lewis. 2005. On the introduction of an agile, temporary workforce into a tandem queueing system. *Queueing Syst.* 51(1): 135–171.
- Kitae, M. Y., V. V. Rykov. 1995. *Contr. Queueing Syst.* CRC press, Boca Raton.
- Lassiter, K., A. Khademi, K. M. Taaffe. 2015. A robust optimization approach to volunteer management in humanitarian crises. *Int. J. Prod. Econ.* 163: 97–111.
- Legros, B., O. Jouini, G. Koole. 2014. A Uniformization Approach for the Dynamic Control of Multiserver Queueing Systems with Abandonments. Technical report, Working paper.
- Leiras, A., I. de Brito Jr, E. Queiroz Peres, T., Rejane Bertazzo, H. T. Y. Yoshizaki. 2014. Literature review of humanitarian logistics research: Trends and challenges. *J. Humanitarian Logist. Supply Chain Manag.* 4(1): 95–130.
- Lippman, S. A. 1975. Applying a new device in the optimization of exponential queueing systems. *Oper. Res.* 23(4): 687–710.
- Little, J. 1961. A proof of the queueing formula: $L = \lambda w$. *Oper. Res.* 9(3): 383–387.
- Lodree, E. J., L. B. Davis. 2016. Empirical analysis of volunteer convergence following the 2011 tornado disaster in Tuscaloosa, Alabama. *Nat. Hazards.* 84(2): 1109–1135.
- Lowe, S., A. Fothergill. 2003. A need to help: Emergent volunteer behavior after September 11th. Beyond September 11th: An account of post-disaster research, pp. 293–314.
- Mahabhashyam, S. R., N. Gautam. 2005. On queues with Markov modulated service rates. *Queueing Syst.* 51(1): 89–113.
- Mandelbaum, A., A. L. Stolyar. 2004. Scheduling flexible servers with convex delay costs: Heavy-traffic optimality of the generalized $c\mu$ -rule. *Oper. Res.* 52(6): 836–855.
- Mayorga, M. E., E. J. Lodree, J. Wolczynski. 2017. The optimal assignment of spontaneous volunteers. *J. Oper. Res. Soc.* 68(9): 1106–1116.
- Nain, P. 1989. Interchange arguments for classical scheduling problems in queues. *Syst. Control Lett.* 12(2): 177–184.
- Neuts, M. F. 1981. *Matrix-Geometric Solutions in Stochastic Models: An Algorithmic Approach*. Courier Corporation, Chelmsford.
- O'Brien, P. W., D. S. Mileti. 1992. Citizen participation in emergency response following the Loma Prieta earthquake. *Int. J. Mass Emerg. Disast.* 10(1): 71–89.
- Oberijé, N. 2007. Civil response after disasters the use of civil engagement in disaster abatement.
- Ortuño, M., P. Cristóbal J. Ferrer, F. Marín-Campo, S. Muñoz, G. Tirado, B. Vitoriano. 2013. Decision aid models and systems for humanitarian logistics. A survey. *Decision aid Models for Disaster Management and Emergencies*. Springer, Berlin, 17–44.
- Paret, K. E., M. E. Mayorga, E. J. Lodree. 2020. Assigning spontaneous volunteers to relief efforts under uncertainty in task demand and volunteer availability. *Omega*, p. 102228.

- Perel, E., U. Yechiali. 2008. Queues where customers of one queue act as servers of the other queue. *Queueing Syst.* **60**(3): 271–288.
- Perry, R. W., M. K. Lindell, K. J. Tierney. 2001. *Facing the Unexpected: Disaster Preparedness and Response in the United States*. Joseph Henry Press, Washington, D.C.
- Points of Light Foundation. 2002. Preventing a Disaster Within the Disaster: The Effective Use and Management of Unaffiliated Volunteers. Available at <https://www.ncjrs.gov/pdffiles1/Archive/202852NCJRS.pdf> (accessed June 5, 2020).
- Purdue, P. 1974. The m/m/1 queue in a Markovian environment. *Oper. Res.* **22**(3): 562–569.
- Puterman, M. L. 1994. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. Wiley Series in Probability and Mathematical Statistics. John Wiley and Sons, Inc, New York.
- Puterman, M. L., M. C. Shin. 1978. Modified policy iteration algorithms for discounted Markov decision problems. *Management Sci.* **24**(11): 1127–1137.
- Quarantelli, E. L. 1998. Major criteria for judging disaster planning and managing their applicability in developing countries.
- Regterschot, G., De Smit, J. 1986. The queue m–g–1 with Markov modulated arrivals and services. *Math. Oper. Res.* **11**(3): 465–483.
- Reiman, M. I., Wein, L. M. 1998. Dynamic scheduling of a two-class queue with setups. *Oper. Res.* **46**(4): 532–547.
- Sampson, S. E. 2006. Optimization of volunteer labor assignments. *J. Oper. Manag.* **24**(4): 363–377.
- Simpson, N., P. Hancock. 2009. Fifty years of operational research and emergency response. *J. Oper. Res. Soc.* **60**(1): S126–S139.
- Smith, W. E. 1956. Various optimizers for single-stage production. *Nav. Res. Logist. Quart.* **3**(1-2): 59–66.
- Stidham Jr, S. 2002. Analysis, design, and control of queueing systems. *Oper. Res.* **50**(1): 197–216.
- Thorsdottir, H., I. Verloop. 2016. Markov-modulated m/g/1-type queue in heavy traffic and its application to time-sharing disciplines. *Queueing Syst. Theory App.* **83**(1): 29–55.
- Tierney, K. J., M. K. Lindell, R. W. Perry. 2001. *Facing the Unexpected: Disaster Preparedness and Response in the United States*. Joseph Henry Press, Washington, D.C.
- Townsend, F. F. 2006. *The Federal Response to Hurricane Katrina: Lessons Learned*. The White House, Washington, DC.
- Urrea, G., A. J. Pedraza-Martinez, M. Besiou. 2019. Volunteer management in charity storehouses: Experience, congestion and operational performance. *Prod. Oper. Manag.* **28**(10): 2653–2671.
- Van Mieghem, J. A. 1995. Dynamic scheduling with convex delay costs: The generalized c– μ rule. *Ann. Appl. Prob.* **5**(3): 809–833.
- Van Oyen, M., D. Teneketzis. 1994. Optimal stochastic scheduling of forest networks with switching penalties. *Adv. Appl. Prob.* **26**(2): 474–497.
- Van Oyen, M. P., D. G. Pandelis, D. Teneketzis. 1992. Optimality of index policies for stochastic scheduling with switching penalties. *J. Appl. Prob.* **29**(4): 957–966.
- Wenger, D. 1991. Emergent and volunteer behavior during disaster: Research findings and planning implications. Technical report, Texas A&M University, Hazard Reduction and Recovery Center, College Station, TX.
- Whittaker, J., B. McLennan, J. Handmer. 2015. A review of informal volunteerism in emergencies and disasters: Definition, opportunities and challenges. *Int. J. Disast. Risk Red.* **13**: 358–368.

Supporting Information

Additional supporting information may be found online in the Supporting Information section at the end of the article.

Appendix S1. Supplement Material.