informs
https://pubsonline.informs.org/journal/mnsc

MANAGEMENT SCIENCE
*Articles in Advance*, pp. 1–17
ISSN 0025-1909 (print), ISSN 1526-5501 (online)

# Modeling the Impact of Community First Responders

**Pieter L. van den Berg,[a] Shane G. Henderson,[b] Caroline J. Jagtenberg,[c,*] Hemeng Li[b]**

[a] Rotterdam School of Management, Erasmus University, 3062 PA Rotterdam, Netherlands; [b] Cornell University, Ithaca, New York 14850; [c] Vrije Universiteit Amsterdam, 1081 HV Amsterdam, Netherlands

*Corresponding author

**Contact:** vandenberg@rsm.nl, https://orcid.org/0000-0002-8550-6769 (PLvdB); sgh9@cornell.edu, https://orcid.org/0000-0003-1004-4034 (SGH); c.j.jagtenberg@vu.nl, https://orcid.org/0000-0002-0742-7707 (CJJ); hl2359@cornell.edu, https://orcid.org/0000-0002-9004-516X (HL)

**Received:** December 16, 2022
**Revised:** September 6, 2023; January 24, 2024
**Accepted:** January 25, 2024
**Published Online in Articles in Advance:** April 26, 2024

**https://doi.org/10.1287/mnsc.2022.04024**



**Abstract.** In community first responder (CFR) systems, traditional emergency service response is augmented by a network of trained volunteers who are dispatched via an app. A central application of such systems is out-of-hospital cardiac arrest (OHCA), where a very fast response is crucial. For a target performance level, how many volunteers are needed, and from which locations should they be recruited? We model the presence of volunteers throughout a region as a Poisson point process, which permits the computation of the response-time distribution of the first-arriving volunteer. Combining this with known survival-rate functions, we deduce survival probabilities in the cardiac arrest setting. We then use convex optimization to compute a location distribution of volunteers across the region that optimizes either the fraction of incidents with a fast response (a common measure in the industry) or patient survival in the case of OHCA. The optimal location distribution provides a bound on the best possible performance with a given number of volunteers. This can be used to determine whether introducing a CFR system in a new region is worthwhile or can serve as a guide for additional recruitment in existing systems. Effective target areas for recruitment are not always obvious because volunteers recruited from one area may be found in various areas across the city depending on the time of day; we explicitly capture this issue. We demonstrate these methods through an extended case study of Auckland, New Zealand.

**History:** Accepted by Carri Chan, healthcare management.
**Funding:** This research was financed in part by the Netherlands Organization for Scientific Research (NWO) in the form of a Rubicon grant (019.172EN.016) to C. J. Jagtenberg and a Veni grant (VI.Veni.191E.005) to P. L. van den Berg. S. G. Henderson and H. Li were supported in part by National Science Foundation [Grant CMMI-2035086]. Vrije Universiteit Amsterdam and Erasmus University received funding from TKI Dinalog [Grant 2023-1-307TKI]. Some of this work has been executed under the TKI Dinalog [Grant 2023-1-307TKI]. Part of this work was completed during a research visit made possible by a [Distinguished Visitor Award] from the University of Auckland.
**Supplemental Material:** The online appendix and data files are available at https://doi.org/10.1287/mnsc.2022.04024.

**Keywords:** Poisson point process • community first responders • volunteer alert • out-of-hospital cardiac arrest

## 1. Introduction

Certain medical emergencies require such a fast response that it can be helpful to supplement traditional ambulance services with community first responders (CFRs). CFRs are volunteers available to be dispatched by an ambulance control center and can be members of the public who have received training or off-duty medical professionals.

The predominant application is out-of-hospital cardiac arrest (OHCA) for which, in Europe alone, 19 of 29 countries have a CFR system in place (Oving et al. 2019). OHCA is a major cause of mortality around the world, and the probability of survival significantly improves if patients have early access to treatment such as cardiopulmonary resuscitation (CPR) (Nichol et al. 1999, Berdowski et al. 2010). Alternative applications where CFRs are dispatched to general medical emergencies are surveyed in Phung et al. (2017) and include Lives, NHS North West Ambulance Service, and Northern Ireland Ambulance Service Health and Social Care Trust.

CFR systems are activated by an ambulance control center in parallel with traditional emergency medical services (EMS). In the past, CFRs were dispatched based on their self-reported home and work addresses (HartslagNu 2020), whereas in recent years, an increase in smartphone usage has enabled applications (apps) that alert CFRs based on their real-time GPS location. Such systems include PulsePoint in the United States (PulsePoint 2020) and GoodSAM in several countries, including the United Kingdom, Australia, and New Zealand (Smith et al. 2017, GoodSAM Platform 2020). These systems monitor the real-time location of CFRs who have the app running on their smartphones, though these data are not logged for privacy reasons. CFRs near an incident are alerted, and some random subset of them accept the alert and proceed to the scene. CFRs typically do not have trouble accessing the patient, as the person who made the emergency phone call is informed about their dispatch. The CFR provides first aid to the level that their training allows until the ambulance crew, which is dispatched independently, takes over.

The central goal of this paper is to determine how the availability of CFRs and their distribution throughout the city affect response times. Faster responses are often equated with better care, and consequently, response times are commonly used as a proxy for a system's effectiveness. The response time largely depends on the CFR density, which, in turn, depends on the number of registered CFRs and their distribution over the city. The number of registered CFRs ranges from 0.1% of the population in a newly introduced system in New Zealand (Global Resuscitation Alliance 2019) to more than 1% in a mature system in the Netherlands (Wikipedia 2020). Another major factor is the acceptance probability of CFRs; Brooks et al. (2016) found that PulsePoint CFRs accepted, on average, 23% of the alerts they receive, though this value varies considerably from one city to another. One might be concerned that CFRs could receive simultaneous alerts, leading to queueing effects, but the rate at which CFRs receive alerts is typically very low; for example, Pijls et al. (2019) report an average of 1.3 alerts per CFR per year.

One might attempt to understand CFR response by looking at how recorded CFR response times differ as a function of location. Doing so is unlikely to give good results because of data sparsity. Indeed, the number of incidents per location is very small, leading to few observations even after collecting years of data. Even if such data were sufficient to accurately estimate how the number of incidents differs from location to location, much more data would be needed for predicting the distribution of response times.

In contrast with these retrospective approaches, we offer a prospective method by estimating the performance of CFR systems using a stochastic model. Besides being less sensitive to statistical error, an additional benefit of this method is that it can be applied to either existing or prospective CFR systems. The model we propose uses a Poisson point process to estimate response-time distributions based on CFR densities and is combined with traditional EMS response and survival functions from literature. The Poisson point process model enables the use of convex optimization to obtain bounds on the potential performance of various CFR deployments. We can then answer the following questions:

1. In what areas is introducing a CFR app an efficient way to reach certain response-time targets? (Section 6.1)
2. What is the benefit of introducing a CFR app for a base of already trained volunteers, with known home addresses? (Section 6.2)
3. What is the benefit of a CFR app with a given number of registered volunteers when you do not know where they live or spend their time? (Section 6.3)
4. Where should one recruit additional CFRs for an existing app with a known current CFR base? (Section 6.4)

All of these questions arose in real-life use cases. Question 1 was of interest to the Fire Department of Amsterdam, which adopted a moonshot goal of having a CFR app and a large CFR base by 2030 (Brandweer Amsterdam-Amstelland 2021, p. 6). Question 2 was posed by a national Red Cross organization that considered extending the use of its existing volunteer base via an app. Both of these parties executed a case study based on the models in a preliminary version of this paper. Questions 3 and 4 are relevant to the St. John Ambulance Service that operates GoodSAM in Auckland, and we performed corresponding calculations that are partially included in this paper.

We view the primary contributions of this work to be (1) the introduction of a new class of CFR problems to the operations research community, and (2) the selection of, and interplay between, Poisson point processes to model volunteers and convex optimization to perform spatial optimization. The combination of these techniques enables powerful and transparent analyses that can inform, and have informed, decisions that save lives.

The remainder of this paper is structured as follows. Section 2 reviews the existing literature, and Section 3 introduces our model that quantifies the impact of a *given* CFR base. Building upon this model, Section 4 introduces optimization models to determine the ideal geographical distribution of CFRs. We demonstrate our methods within the city of Auckland, New Zealand, in Sections 5 and 6. We conclude by answering the four questions given and discussing the managerial implications of this work in Section 7. An Online Appendix provides supporting details.

## 2. Literature

We discuss performance metrics for CFR systems in Section 2.1 and review studies that evaluate the effectiveness of such systems in Section 2.2. Section 2.3 summarizes CFR literature that is specific for OHCA patients, and Section 2.4 reviews studies on first-responder dispatching. Section 2.5 places our work in a broader context.

### 2.1. Performance Metrics

As an essential goal of CFR systems is to reach patients quickly, their performance can be measured in terms of response times: the duration between the moment a call arrives in the emergency call center and the moment the first responder arrives on scene. The same is true for EMS providers, who often have targets that are a function of response times, phrased as service-level agreements (responding to $x$ percent of all calls in $y$ minutes).

One way of translating response times into a medical outcome–based metric is to convert them to a survival probability through a so-called survival function. This is highly dependent on the medical condition of the patient. Survival functions have been used in work on ambulance operations in the absence of volunteer schemes; see Erkut et al. (2008), McLay (2009), Bandara et al. (2012), and Zaffar et al. (2016).

## 2.2. Retrospective Studies on Effectiveness

In contrast to our work, existing literature focuses on retrospective analysis of data to determine the impact of CFR systems.

Early work from Sweden (Ringh et al. 2011) analyzed the impact of a CFR app using both a physical experiment and a small retrospective study in which volunteers arrived before ambulances approximately half the time. For a Dutch text-based volunteer CFR system, Zijlstra et al. (2014) compared response times in a system with volunteers to a system with ambulances only, reporting a reduction of 159 seconds in the time to defibrillation. Pijls et al. (2016) focused on survival rates: the survival rate for OHCA increased from 16% to 27% when at least one volunteer responded.

Several studies specifically investigated the impact of volunteer density. Both Jansma (2014) and Stieglis et al. (2020) investigated how density relates to response times via an empirical study and computer simulation, respectively. Pijls et al. (2019) show a positive correlation between the fraction of inhabitants registered as volunteers and patient survival.

## 2.3. CFR for OHCA

Some CFR systems are designed explicitly for OHCA patients, perhaps because OHCA survival is known to significantly improve with early CPR administered by trained individuals (Nichol et al. 1999, Sasson et al. 2010, Yan et al. 2020). This relationship between time to CPR and survival has been made explicit in multiple studies that have followed a group of patients for which the time to CPR has been recorded, for example, Valenzuela et al. (1997), Waalewijn et al. (2001), and De Maio et al. (2003).

Cardiac arrest patients are known to benefit from receiving defibrillation by automated external defibrillators (AEDs). CFR systems are therefore often designed to include information on AED locations. The optimization of AED placement is explored in Folke et al. (2009) and Chan et al. (2016, 2018); however, they treat this question without explicitly incorporating the dynamics of a CFR system. Like those studies, we view as out of scope the detailed modeling of AED-CFR dispatch coordination, such as when one CFR is sent directly to the patient while another is sent to retrieve the nearest AED. Whereas such considerations are important for an individual OHCA, such detailed modeling seems unnecessary for the high-level questions of recruitment that we consider here.

## 2.4. Dispatching

At the time of an incident, real-time decisions must be made on which volunteers are dispatched by sending a push notification to their phones. This decision is typically based on the observed locations of nearby volunteers and should balance overburdening volunteers with response-time benefits (Henderson et al. 2022).

Advanced dispatch methods may vary the number and time of alerts and may consider a range of tasks, for example, picking up medical equipment (Nazarian 2018, Matinrad et al. 2019). In this context, Matinrad et al. (2021) discuss uncertain task compliance.

A recent trend is to have medical supplies delivered by drones, for example, AEDs (Boutilier et al. 2017, Chu et al. 2021, Boutilier and Chan 2022) or blood (Nisingizwe et al. 2022). Such a design disconnects equipment delivery decisions from CFR coordination, hence further confirming our choice to leave such questions out of scope of this paper.

Recent unpublished work (Liu et al. 2022, Shin et al. 2022) considers AED drone management under ambulance as well as bystander response. Both papers model bystander availability through a Bernoulli distribution, and the bystanders' response time does not depend on the number of available responders. This work, besides evincing very recent interest in the topic, highlights the unavailability of a detailed CFR response model, a gap that we aim to fill here.

## 2.5. Time-Sensitive Volunteers

More broadly, our work relates to the literature on time-sensitive volunteers and crowdsourcing, for example, Ata et al. (2019), McElfresh et al. (2020), and Manshadi and Rodilitz (2022). The key differences with our work are that, in that setting, (1) requests that are relevant for any fixed volunteer are likely to be more frequent, (2) volunteer disengagement is a central consideration and is modeled in a variety of ways, and (3) the proximity of volunteers is relevant but not critical, as in our setting. There is a slight resemblance between our work and the literature on matching and crowdsourcing, for example, Özkan and Ward (2020), Tafreshian et al. (2020), and Johari et al. (2021); queueing, pricing, and learning are central in this literature but not in our work. The use of spatial Poisson processes to model the location of servers (CFRs in our setting) could potentially be relevant in many applications, including the modeling of available drivers in ride hailing; see, for example, Qin et al. (2020).

## 3. Modeling CFR Response

In this section, we argue that CFR locations are well modeled by a Poisson point process and thereby derive the distribution of their response times. This distribution, particularly when complemented with the distribution of EMS response times, can be translated to an expected health gain in the population because of CFRs.

Throughout the paper, we use the term "city" to refer to the overall area in which the CFR response is modeled. However, our modeling approach applies whether the area under consideration is a city, a county, a state, or even a nation, at least in principle.

### 3.1. Using a Poisson Point Process

We say that CFRs are *available* if they are present in the city, have the app running on their phone, and will accept a notification should it be sent to them. We assume that available CFRs are distributed throughout the city according to a spatial Poisson point process. We assume that the set of available CFRs does not depend on call volume, which is reasonable because, typically, CFRs are called out on the order of once per year (Pijls et al. 2019), at least in the context of OHCA.

Modeling available CFRs as a spatial Poisson point process is reasonable because of results that justify approximating certain spatial point processes by Poisson point processes. For a general introduction to Poisson point processes, including theoretical results that justify Poisson point process modeling in applications in great generality; see Kingman (1993) and Barbour et al. (1992). We justify the Poisson point process assumption with Proposition 1 in Online Appendix 1, which shows that a Poisson point process arises in a regime where the number of CFRs $n$ is large. That result is not the most general result possible, nor will it be surprising to those versed in point process theory, but we provide it to support our contention that the Poisson point process model is a good one in our setting. This assertion is further supported by the prevalence of Poisson point process modeling in related settings; see Larson and Odoni (1981), for example.

Proposition 1 shows that even when CFRs have a unique availability probability and location distribution, the location of available CFRs is well modeled by a Poisson point process with a certain mean measure $\mu$. Thus, it is not necessary to assume CFRs are a homogeneous group: responders may have their own probability of accepting alerts and own distribution of time spent in each part of the city. The exact way in which individual CFRs contribute to the aggregate mean measure $\mu$ is not important for our analysis; the overall measure $\mu$ is what is important.

It is therefore sufficient to henceforth assume that $\mu = n\alpha\nu$, where $n$ indicates the overall number of CFRs who have the *same* availability probability $\alpha \in (0,1)$ and the same location probability distribution $\nu$, conditional on being available.

### 3.2. Response-Time Distribution

The response time to an incident is the minimum of the time until a CFR or ambulance arrives at the scene, with these intervals measured from the moment the call to EMS is initiated.

As is common in CFR apps (Smith et al. 2017), we assume that the closest available CFR is dispatched to an incident as long as the CFR is within a given maximal distance from the patient. We are interested in the distribution of the response time of this closest CFR, which depends on two components: (1) the response delay, and (2) walking time. The response delay is the time that passes between the call initiation and the responder starting to travel. This consists of the interval between the time the call is made and the time a CFR is dispatched (triage and dispatch delay) and the interval between the CFR dispatch and the time a CFR starts walking (CFR acceptance delay). For simplicity, the total response delay for a CFR is assumed to be constant and is denoted by $\kappa$. For any location $x$, define the ball $B(x, t)$ as the area surrounding $x$ within which a CFR can reach the location within $t$ minutes, including response delay.

We assume that ambulance bases are given and ambulances respond from their bases and not from the road. This is an approximation that is reasonable when the ambulances are not too heavily loaded. Thus, ambulances incur the first two components of the response delay along with a potentially different (from CFR) delay from dispatch to the time they begin traveling. As in many standard models of ambulance operations, for example, Daskin (1983), we assume that each ambulance is busy with probability $q \in (0,1)$ and that ambulances are busy or not, independent of one another. We also assume that ambulance locations and availabilities are independent of those for CFRs. For any given call location $x$, let $y(x, t)$ denote the number of ambulances that can reach the location $x$ within $t$ minutes of the incident when stationed at their base.

The response time, including response delay, of the closest responder, whether CFR or ambulance, to a patient at location $x$ is a random variable, which we denote $T(x)$. The number of CFRs that are available within $t$ minutes from a patient at location $x$ is, because of the Poisson point process assumption, Poisson distributed with mean $\mu(B(x,t)) = n\alpha\nu(B(x,t))$. As is standard for Poisson point processes, the probability $\mathbb{P}(T(x) > t)$ that the response time of the closest responder is greater than $t$ is then

$$\mathbb{P}(T(x) > t) = q^{y(x,t)}\exp(-\mu(B(x,t))). \tag{1}$$

Here, we used the fact that the response time is greater than $t$ if and only if all ambulances stationed at nearby bases are busy and there are no CFRs in the set $B(x,t)$. The number of CFRs in the set $B(x, t)$ is a Poisson-distributed random variable, $Z$ say, with mean $\gamma = \mu(B(x,t))$, and for such random variables, $\mathbb{P}(Z=0) = e^{-\gamma}$.

To illustrate how one might perform these calculations, suppose a CFR walks with a constant pace of $w$ km/min, and let $d_t$ be the distance in kilometers for which the response delay plus walking time equals $t$ for $t \geq \kappa$. We can convert response times $t \geq \kappa$ to distances using $d_t = w(t-\kappa)$. We further assume that $t$ is so small that the CFR density may be assumed to be constant throughout the ball $B(x,t)$. Letting $\omega$ denote the value of the assumed-constant available CFR density in the ball $B(x,t)$, we get

$$\begin{aligned}\mathbb{P}(T(x) \leq t) &= 1 - q^{y(x,t)}\exp(-\mu(B(x,t))) \\ &= 1 - q^{y(x,t)}\exp(-\omega\pi d_t^2). \qquad (2)\end{aligned}$$

The number of sufficiently close ambulances $y(x, t)$ is readily computed for any fixed $x$ and $t$ given deterministic travel time models. Online Appendix 2 shows the CFR response time distributions (not counting ambulances) that follow from our model.

Our discussion thus far assumes a fixed call location $x$. The unconditional response time, $T$ say, is a mixture of the call location–specific response-time random variables. Assume that calls arise over the city according to a probability distribution $\Lambda$, where $\Lambda(B)$ is the probability that a given call arises within the ball $B$. Then, $\mathbb{P}(T \leq t) = \int \mathbb{P}(T(x) \leq t)\Lambda(dx)$.

For computational convenience, we subdivide the city into a finite number of regions indexed by $l \in \mathcal{L} = \{1, 2, \ldots, \ell\}$. The location of a generic call lies in region $l$ with probability $\lambda_l$, $l \in \mathcal{L}$, so that $\sum_{l\in\mathcal{L}}\lambda_l = 1$. For all call locations $x$ within region $l$, we assume that response times $T(x)$ are identically distributed as in (2), with distribution function $P(T(x) \leq t) = 1 - q^{z(l,t)} \exp(-\mu_l\pi d_t^2)$. Here, $z(l, t)$ is the location-discrete equivalent of $y(x, t)$ and yields a region-specific number of nearby ambulances, and $\mu_l$ is a region-specific CFR density. Thus, $\mathbb{P}(T \leq t) = 1 - \sum_{l\in\mathcal{L}}\lambda_l q^{z(l,t)}\exp(-\mu_l\pi d_t^2)$.

This model for the CFR response assumes that the density of CFRs and the number of nearby ambulances are constant within each region $l \in \mathcal{L}$ and that the response-time distribution does not vary, as the call location varies within the region. This is an imperfection in the model because calls that arise close to the boundary of a region may receive a CFR response from a neighboring region where the CFR density is different, but this effect is not captured. This model is therefore plausible when regions are large enough that CFRs outside the region are unlikely to materially impact survival rates within the region or in the situation when neighboring regions have similar CFR densities. Online Appendix 3 numerically checks the effect of the constant density assumption on our case study and concludes that, overall, it is minor. A similar discretization effect occurs for the modeling of the ambulance response, in line with existing literature modeling ambulance response.

## 4. The Ideal CFR Distribution

In view of (1), it is clear that for any fixed ball $B$, we want $\mu(B)$ to be as large as possible. Recall that $\mu(B) = n\alpha\nu(B)$, where $n$ is the number of CFRs, $\alpha$ is the probability that a CFR is available, and $\nu(B)$ is the conditional probability that a given CFR can be found in the ball $B$ given that the CFR is available. Accordingly, to make $\mu(B)$ large, we can increase the number of CFRs $n$, we can increase the probability $\alpha$ that a given CFR is available, or we can influence the probability distribution $\nu$ that describes where a given available CFR is found across the city. We assume that $\alpha$ is exogenous to our model, though reality is potentially more complicated.

The probability distribution $\nu$ is not under our direct control. Nevertheless, it is worth identifying the *choice* of probability distribution $\nu$ that is most beneficial for two reasons:

1. An optimal CFR distribution can guide efforts in recruiting new volunteers by indicating populations of potential CFRs that are likely to have a large impact.
2. The performance of an optimal CFR distribution can be used to determine a lower bound on the number $n$ of volunteers needed to reach response-time goals.

With these reasons in mind, we now identify the optimal probability distribution $\nu$ for a given demand distribution $\lambda$.

First, we investigate the ideal CFR distribution across the city, where regions in the city are linked only through a bound on the total CFR mass to be allocated. Then, we incorporate the reality that an individual recruited from region $l$ has a region-specific *profile*, that is, a contribution of its CFR mass to all regions across the city.

As in Section 3, we assume that the locations in a city are partitioned into a finite list of neighborhoods or regions indexed by $l \in \mathcal{L}$ as earlier, with $\lambda_l$ denoting the conditional probability that a call originates within region $l$, given that a call arises. Rather than optimizing over all possible probability measures $\nu$, which is an infinite-dimensional optimization problem, we restrict attention to probability distributions $\nu$ that are uniform within a region so that they can be parameterized by the probability mass $\nu_l$ associated with region $l$, with $\sum_{l\in\mathcal{L}}\nu_l = 1$. This greatly reduces the complexity of determining an optimal probability distribution because the calculation becomes finite dimensional. This is also reasonable in practice because our goals relate to the overall distribution of CFRs across a city and not to the small-scale detail of exactly how they are distributed within small areas. To quantify the impact of dividing CFRs into a finite set of homogeneous regions, we compare our results with those of a more common approach: a very fine-grained but discrete network. This analysis in Online Appendix 3 confirms that our approach is reasonable.

We consider the contribution of CFRs to both (1) coverage with respect to a fixed response-time threshold, and (2) a function of the full response-time distribution. The first is a standard objective in EMS; the second is more general and can capture, for example, the average response time and the probability of patient survival (under the assumption that it is known how survival declines with response time). The details for the first objective are more straightforward, so we begin there. Moreover, insights obtained from the structure of the solution based on the first objective can help explain the solutions with respect to more complex objectives. Henceforth, when we discuss or illustrate the second type of objective, we will phrase it as "survival," although those analyses have a broader applicability and can be performed for essentially any monotone function of the response time, including the mean.

### 4.1. Optimizing Coverage

Coverage is defined as the fraction of demand that can be served within a predefined threshold time $\tau$. For our purpose, it is slightly more convenient to work with the probability that the response time *exceeds* $\tau$, conditioning on the region $l$ in which the call occurs and using (2). Let $\nu_l$ denote the CFR probability mass within region $l$, and let $a_l$ be the area of region $l$ in square kilometers. Thus, within region $l$, the density of available CFRs is $n\alpha\nu_l/a_l$. Recalling that $d_\tau$ is the distance within which a CFR can reach a call within the time threshold $\tau$ and $z(l,\tau)$ is the number of ambulances that can reach a call in region $l$ within time $\tau$, we get

$$\begin{aligned}\mathbb{P}(T>\tau) &= \sum_{l} \lambda_l \, \mathbb{P}(T(l)>\tau)\\ &= \sum_{l\in\mathcal{L}} \lambda_l q^{z(l,\tau)} \exp(-\pi d_\tau^2 n\alpha\nu_l/a_l). \qquad (3)\end{aligned}$$

Expression (3) is (jointly) convex in the probabilities $(\nu_l : l \in \mathcal{L})$, yielding

**Proposition 1.** *The probability that the response time is greater than any fixed quantity $\tau$ is a convex function of the probabilities $(\nu_l : l\in\mathcal{L})$.*

A consequence of Proposition 1 is that we can use convex optimization methods to minimize the probability that the response time is greater than the time threshold $\tau$, $\mathbb{P}(T>\tau)$, subject to the simplicial constraints that $\sum_{l\in\mathcal{L}}\nu_l = 1$ and $\nu_l \geq 0$ for all $l\in\mathcal{L}$. For notational simplicity, let the *residual call probability* $\tilde{\lambda}_l = \lambda_l q^{z(l,\tau)}$, which is the probability a call arises in region $l$ and all nearby ambulances are busy, and let $\theta_l = \pi d_\tau^2 n\alpha/a_l$ for $l\in\mathcal{L}$. The optimization problem we want to solve is to

$$\begin{aligned}\min_{\nu} \quad & \sum_{l\in\mathcal{L}} \tilde{\lambda}_l e^{-\theta_l\nu_l},\\ \text{s/t} \quad & \sum_{l\in\mathcal{L}} \nu_l = 1,\\ & \nu_l \geq 0, l\in\mathcal{L}.\end{aligned}$$

The objective function is separable, and each term is convex and decreasing in $\nu_l$. Accordingly, a standard application of the Karush-Kuhn-Tucker conditions ensures that a greedy allocation is optimal. To describe this allocation, let us call $\tilde{\lambda}_l \exp(-\theta_l\nu_l)$ the *late rate from region $l$*. We define the *marginal benefit* of adding CFR mass to region $l$ to be the absolute value of the derivative of the late rate with respect to $\nu_l$, namely, $\tilde{\lambda}_l\theta_l e^{-\theta_l\nu_l}$. Starting from $\nu_i = 0$ for all $i$, we begin by increasing CFR mass across all regions $l$ in the set, $I$ say, of regions having maximal marginal benefit. In doing so, the marginal increase in CFR mass in region $l$ should be inversely proportional to $\theta_l$, and, in particular, proportional to the area $a_l$ so that the marginal benefit remains equal across all regions $l\in I$. In other words, the CFR *density* in all regions $l\in I$ increases at the same rate. This continues until either the total CFR mass 1 is used up or the maximal marginal benefit is reduced to the point where a new region $j$ joins the set $I$ of regions attaining the maximal marginal benefit. A more precise description of the algorithm can be found in Online Appendix 4.

The value of $\nu_l$ for each $l\in I$ at which a new region joins the index set $I$ can be computed in closed form. Let $j\in \arg\max\{\tilde{\lambda}_i\theta_i : i\notin I\}$ be the maximal marginal benefit of any region that is yet to receive positive CFR mass, and let $l\in I$ be arbitrary. We solve $\tilde{\lambda}_l\theta_l e^{-\theta_l\nu_l} = \tilde{\lambda}_j\theta_j$. Using the definition of $\theta_l$ and $\theta_j$, canceling common terms, and solving for $\nu_l$ yields

$$\nu_l = \frac{1}{-\theta_l}\ln\left(\frac{\tilde{\lambda}_j/a_j}{\tilde{\lambda}_l/a_l}\right) = \frac{a_l}{\pi d_\tau^2 n\alpha}\ln\left(\frac{\tilde{\lambda}_l/a_l}{\tilde{\lambda}_j/a_j}\right). \qquad (4)$$

A proviso is that we do not reach this new value if we "hit" the constraint $\sum_{l\in I}\nu_l = 1$ first. In this case, we do not reach the point where $\nu_j$ becomes positive.

Expression (4) yields several insights that are reinforced in a two-region example in Online Appendix 5. First, regions are not guaranteed to receive positive probability mass; excluded regions are those with the lowest residual call densities $\tilde{\lambda}_l/a_l$. Second, regions receiving positive probability mass do so roughly in proportion to the product of their area and the log of their residual call densities. Because the density of CFRs is proportional to $\nu_l/a_l$ and thus divided by area, it follows that regions receiving positive CFR density do so roughly in proportion to the log of their residual call densities. We use the term "roughly" because (4) is not the precise form of the optimal allocations, as mentioned earlier. Third, the expected number of CFRs allocated to a region $l$ that receives positive mass is roughly $n\nu_l$, which, from (4), does not depend on $n$ because $\nu_l$ is proportional to $n^{-1}$, where "roughly" has the same sense as before. Finally, the optimal allocations are roughly proportional to the inverse of the area, $\pi d_\tau^2$, of the circle within which CFRs can respond within time $\tau$. For smaller $\tau$ (more stringent targets), the optimal allocations are thus more concentrated in high-call-density regions.

### 4.2. Optimizing Survival

The same approach works for more complex performance metrics, which we illustrate by showing that maximizing the probability of survival for OHCA is a convex optimization problem in the CFR location probabilities $(\nu_l, l \in \mathcal{L})$.

**Proposition 2.** *If the probability of death is increasing in the response time, then the probability of survival is a concave function of the probabilities* $(\nu_l : l \in \mathcal{L})$.

**Proof.** We show that the probability of death is convex in the CFR location probabilities $(\nu_l, l \in \mathcal{L})$. We assumed that the death probability is an increasing function, $\bar{f}(\cdot)$ say, of the response time $T$, which includes the response delay $\kappa$. Increasing functions have at most a countable number of discontinuities, so we can adjust $\bar{f}$ to an increasing, right-continuous function $f$ by redefining $\bar{f}$ in at most a countable number of points. Thus, we now work with the increasing, right-continuous function $f$. Let $\underline{d} = f(\kappa) \in [0,1]$ be the minimal death probability assuming immediate ambulance response. Let $\bar{d} \in [0,1]$ be the maximal death probability assuming response time is infinite. Using the fact that the mean of a nonnegative random variable $Y$ can be expressed as $\int_0^\infty \mathbb{P}(Y > u)\,du$, the death probability as a function of $\nu$, $d(\nu)$ say, is

$$d(\nu) = \mathbb{E}f(T) = \int_0^1 \mathbb{P}(f(T) > u)\,du$$
$$= \underline{d} + \int_{\underline{d}}^{\bar{d}} \mathbb{P}(T > f^{-1}(u))\,du,$$

where $f^{-1}(u) = \inf\{r : f(r) \geq u\}$ for $u \in [0,1]$, and the final step follows from Serfling (1980, lemma 1.1.4(iii)). Hence, conditioning on region $l$ of the call,

$$d(\nu) = \underline{d} + \sum_{l\in\mathcal{L}} \lambda_l \int_{\underline{d}}^{\bar{d}} \mathbb{P}(T(l) > f^{-1}(u))\,du$$
$$= \underline{d} + \sum_{l\in\mathcal{L}} \lambda_l \int_{\underline{d}}^{\bar{d}} q^{z(l,f^{-1}(u))} \exp\{-|B(l,f^{-1}(u))|n\alpha\nu_l/a_l\}\,du. \tag{5}$$

The integrand in (5) is convex in $\nu_l$. □

Thus, we can efficiently maximize the survival probability over probability distributions $\nu$ on $\mathcal{L}$. In fact, a greedy approach that adds CFRs to regions where they yield the highest marginal reduction in death rate gives an optimal solution. In contrast to the greedy approach for late arrivals, the marginal allocation over the different regions does not remain constant. Therefore, an expression similar to (4) to compute the *exact* CFR distribution does not hold. Instead, we use a discretized step size $\epsilon$, which is optimal as $\epsilon \to 0$. We implement this method and give numerical results in Section 6.

### 4.3. More Complex Survival Functions for OHCA

So far, we have focused on survival functions that depend only on the response time, whether by CFR or ambulance. More complex survival functions for OHCA patients, such as model 1 in Waalewijn et al. (2001), depend on both the time to CPR ($t_{\text{CPR}}$) and the time to EMS arrival ($t_{\text{EMS}}$), where $t_{\text{CPR}} = t_{\text{EMS}}$ when an ambulance arrives before volunteers, and $t_{\text{CPR}}$ equals the response time of the volunteer otherwise. Let $f(t_{\text{CPR}}, t_{\text{EMS}})$ be the probability of death as a function of the time to CPR and the time to EMS arrival. Model 1 of Waalewijn et al. (2001) uses

$$f(t_{\text{CPR}}, t_{\text{EMS}}) = 1 - (1 + e^{0.04+0.3t_{\text{CPR}}+0.14(t_{\text{EMS}}-t_{\text{CPR}})})^{-1}. \tag{6}$$

These times are measured from patient collapse instead of call initiation. Waalewijn et al. (2001) estimate the difference between these two moments to be one minute (median). Recall that $z(l,\tau)$ is the number of ambulances that can reach region $l$ from their home base within $\tau$ minutes of cardiac arrest. As before, let $d(\nu)$ be the probability of death for a single call as a function of the CFR location distribution $\nu$. Let $a_{jl}$ denote the response time of the $j$th closest ambulance to region $l$ when all ambulances are available and at their respective bases. Let $T_v(l)$ be the (random) response time of the closest volunteer, the distribution of which can be found by dropping the $q^{y(x,t)}$ term in (2). Conditioning on the location of the call and the closest available ambulance to region $l$, we get

$$d(\nu) = \sum_{l\in\mathcal{L}} \lambda_l\, \mathbb{E}f(T_{\text{CPR}}(l), T_{\text{EMS}}(l))$$
$$= \sum_{l\in\mathcal{L}} \lambda_l \Bigg( q^k \mathbb{E}f(T_v(l), \infty) + \sum_{j=1}^{k} q^{j-1}(1-q)\mathbb{E}f(\min\{T_v(l), a_{jl}\}, a_{jl}) \Bigg). \tag{7}$$

The term $f(T_v(l),\infty)$, corresponding to the situation where all ambulances are busy, is zero according to model 1 of Waalewijn et al. (2001) but, more generally, is convex in $\nu$ as long as $f$ is increasing in its first argument, as we have already established. As to $\mathbb{E}f(\min\{T_v(l), a_{jl}\}, a_{jl})$, as long as $f(t_{\text{CPR}}, t_{\text{EMS}})$ is increasing in $t_{\text{CPR}}$ for any fixed value of $t_{\text{EMS}}$, which we henceforth assume, then our previous arguments extend to ensure convexity in $\nu$. Thus, we have established

**Proposition 3.** *Suppose that $f(t_{CPR}, t_{EMS})$ is increasing in $t_{CPR}$ for each fixed $t_{EMS} \in [0, \infty]$. Then, the probability of survival is a concave function of the probabilities $(\nu_l : l \in \mathcal{L})$, even when accounting for the impact of ambulances providing advanced care in addition to CPR.*

### 4.4. Profile Optimization

Until now, we have assumed that CFRs spend all their time in a single region, yet in reality, a person likely spends one's time in a few different regions during different times of day. Therefore, instead of having one specific location probability distribution $\nu$, we have time-dependent CFR location distributions $\nu^{(t)}$, where $t$ is assumed to fall into one of a small number of discrete time blocks $\mathcal{T}$, for example, $\mathcal{T}$ = (weekday, weeknight, weekend day, weekend night). A CFR provides a contribution in each time $t \in \mathcal{T}$, and, for any fixed $t$, can contribute to several elements of $\nu^{(t)}$. We model these contributions through *time-dependent profiles* in which people with the same profile are assumed to spend their time in the same way, at least probabilistically speaking. That is, we focus on the aggregate contribution of all CFRs of a certain profile.

Assume CFRs take one of $m$ different profiles. For each time block $t \in \mathcal{T}$, let $\mathcal{P}^{(t)} \in [0, 1]^{|\mathcal{L}| \times m}$ denote the given CFR profile matrix in which $\mathcal{P}^{(t)}_{li}$ is the probability, in time block $t$, that a CFR with profile $i$ appears in region $l$. We have $\sum_{l \in \mathcal{L}} \mathcal{P}^{(t)}_{li} \le 1$ for all $t \in \mathcal{T}$, with strict inequality when CFRs temporarily leave the city.

Let $x_i$ be a decision variable denoting the proportion of CFRs recruited with profile $i$ for $i = 1, 2, \ldots, m$ such that $\sum_{i=1}^{m} x_i = 1, x_i \ge 0$. The resulting CFR location distribution per time block can then be calculated as

$$\nu^{(t)} = \mathcal{P}^{(t)} x, \quad t \in \mathcal{T}. \tag{8}$$

Our goal is to optimize either coverage or patient survival obtained from this $\nu^{(t)}$ with decision variables $(x_i : i = 1, 2, \ldots, m)$ and Constraints (8) ensuring the link between recruitment and CFR location.

**Proposition 4.** *For both the coverage and survival objectives, the optimization problem with profiles remains a convex optimization problem under the same conditions stated in Proposition 2.*

To prove Proposition 4, observe that the objective function becomes a weighted sum of (3) or (5) over the set of time blocks $\mathcal{T}$, which remain convex. Moreover, the constraints we add are linear. Hence, the problem remains convex under both objectives.

With profiles, the optimization problems are no longer separable, so a greedy approach is no longer optimal. To solve the optimization problem with profiles, we implemented a version of the away-step Frank-Wolfe algorithm in Bomze et al. (2020) with an adaptive line search and a stopping criterion based on the (optimality) gap between the objective value of the current iterate and a lower bound obtained through a cutting plane method that relies on convexity of the objective function. See Online Appendix 6 for more details and the pseudocode of the implemented Frank-Wolfe algorithm.

## 5. Case Description

Auckland is the largest city in New Zealand, with the greater Auckland region having a population of 1.4 million and an area of 5,000 square kilometers. We focus on the most urban area by excluding seven out of 21 local boards with a population density of less than 1,000 people per square kilometer. The resulting area has a population of 1.1 million, covers 500 square kilometers, and is divided into 287 so-called area units, yielding the set of regions $\mathcal{L}$. These area units are the second smallest unit at which Statistics New Zealand collects data (Stats New Zealand Geographic Data Service 2016).

### 5.1. OHCA Incidence Rates

The St. John Ambulance Service (SJAS) has provided us with the time and location of all OHCAs in Auckland since 2013. One could use the empirical OHCA rate to estimate the demand vector $\lambda$, but low counts in some area units have very high variability relative to their mean. Instead, we estimate the OHCA incident rate by combining demographics with socioeconomic factors as in Dicker et al. (2019).

Figure 1 shows both the empirical incident rate of different area units and a scatter plot of the number of OHCA incidents since 2013 versus the number of estimated OHCA incidents using Dicker et al. (2019), where each point represents one area unit. We see a general linear trend as expected, but the OHCA incident counts are dispersed relative to a linear trend. The three area units highlighted in Figure 1(b) are two area units covering the central business district (CBD) and the area unit containing the airport. We replaced the estimated rates in these three area units with the empirical counts in SJAS data because these area units were extreme outliers in the original fit. The input demand rate $\lambda$ to our model is then the normalized estimated OHCA counts after the CBD and airport adjustments. More details and discussion on this estimation process can be found in Online Appendix 7.

### 5.2. Response Process

We assume that an ambulance always responds and that volunteers are alerted if they are close to the incident. Of interest is the time it takes for the first responder to arrive on site. All reported response times are measured from call initiation, meaning that they do not include the time between cardiac arrest and call initiation. This is in line with classical performance measures for EMS systems and is the only time known to dispatchers. A response time, whether by ambulance or by volunteer, consists of

**Figure 1.** (Color online) OHCA Incidence for Auckland, New Zealand

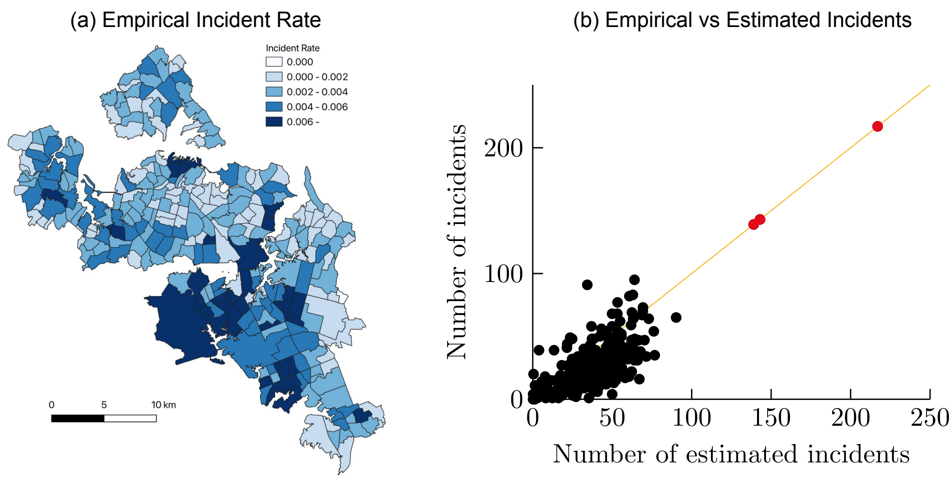


*Notes.* (a) The empirical incident rate per area unit in Auckland, New Zealand. (b) Scatter plot of the actual and estimated number of cardiac arrests between 2013 and 2020 over all area units.

a response delay and a travel time. We base our estimates on GoodSAM data whenever available and reliable; in other cases, we resort to estimates in the literature.

**5.2.1. CFR Response.** The local CFR app GoodSAM is run by the SJAS. We have obtained data from the SJAS on CFR operations since the inception of the program in December 2017. Between December 2017 and November 2020, 6,749 CFR alerts were sent out for 2,827 OHCA incidents. We provide an exploratory data analysis in Online Appendix 8 that suggests that our model is reasonably consistent with the very noisy data at hand.

For the probability that a volunteer accepts an alert, we use the historical value as observed in the data: $\alpha = 0.14$. We take this value to be equal for all CFRs rather than assuming an individual-specific acceptance probability because our data are insufficient to accurately estimate such values. We assume that accepting CFRs respond on foot, and we use the most conservative estimate of Slaa (2020) in the Netherlands to model their speed at six kilometers per hour. We assume a constant response delay, which consists of a triage and system activation time as well as a delay for CFR acceptance. For triage and system activation, we assume 1.5 minutes; this is in line with Slaa (2020), who estimated this time at 105 seconds before a system update that automates and thereby speeds up system activation. For the time it takes for a volunteer to accept or reject (CFR acceptance delay), we take the mode of the GoodSAM data: 30 seconds. The sum gives a response delay of two minutes. We further assume that only CFRs within one kilometer of the patient are dispatched, in line with current practice. This implies that the response time of CFRs, conditioned on finding at least one available CFR within one kilometer, is between two and 12 minutes.

The assumed 30-second CFR acceptance delay is not perfect because volunteers may accept an alert but not initiate their departure immediately, or they may accept the alert after they start traveling. To see the impact of a shorter or longer acceptance delay, one can refer to adjacent columns in Table 1.

**5.2.2. Ambulance Response.** The SJAS operates 15 bases from which they could realistically respond to the part of the city that we analyze. We model the ambulance pretrip delay as two minutes; Ridler et al. (2022) estimate this value to be 2.1 minutes in Auckland. We used driving times for EMS vehicles between each base and the centroid of each area unit as obtained from an open-source ambulance simulation package (Ridler 2020); see also Ridler et al. (2017, 2022).

Because OHCA patients only constitute a small portion of all the requests served by SJAS, it is not reasonable to assume the ambulances are placed optimally with respect to OHCA response targets. Instead, we assume that ambulances are placed to meet the SJAS target of responding within 12 minutes for at least 95% of their patients (Ministry of Health 2020). We obtained the corresponding ambulance locations as follows. We assume that each ambulance has a utilization of 0.44, as used in Ridler et al. (2022), and distribute ambulances across the existing bases according to the optimal solution of MEXCLP (Daskin 1983). We use 25 ambulances, which is just sufficient to achieve the mentioned target (95% within 12 minutes, including a pretrip delay of two minutes).

## 5.3. CFR Profiles

For the optimization that includes CFR movement as introduced in Section 4.4, we need data on profiles. These profiles within a given time block can be interpreted

**Table 1.** Required Density of Available CFRs (per Square Kilometer) to Meet a Range of Response-Time Targets, Assuming a Two-Minute Pretrip Delay and a Walking Speed of Six Kilometers per Hour

| | Response-time target (minutes) | | | | | | | | | |
|---|---|---|---|---|---|---|---|---|---|---|
| Coverage | 3 | 4 | 5 | 6 | 7 | 8 | 9 | 10 | 11 | 12 |
| 0.5 | 22.06 | 5.52 | 2.45 | 1.38 | 0.88 | 0.61 | 0.45 | 0.34 | 0.27 | 0.22 |
| 0.7 | 38.32 | 9.58 | 4.26 | 2.40 | 1.53 | 1.06 | 0.78 | 0.60 | 0.47 | 0.38 |
| 0.9 | 73.29 | 18.32 | 8.14 | 4.58 | 2.93 | 2.04 | 1.50 | 1.15 | 0.90 | 0.73 |

as giving the probability that a certain CFR appears in a given area unit in that time block. We introduce a profile for each area unit and assign potential CFRs a profile based on their home address. We distinguish two time blocks: working hours (weekdays from 8 a.m. until 10 p.m.) and all hours outside of that, which we will refer to as "nights and weekends." Thus, the interpretation of $\mathcal{P}_{li}^{(t)}$ is the probability that a CFR living in area unit $i$ appears in area unit $l$ in time block $t$. Although we defined profiles per area unit, it is also possible to have more—or fewer—profiles than area units, for example, based on other location aggregation levels, for example, people who work at the airport or other properties on which one can distinguish recruitment, for example, students.

To obtain profile data for working hours, we combine the population of age 15 years or older from the 2013 census (Stats New Zealand Geographic Data Service 2015) with the commuting routes for the resident population (Stats New Zealand Geographic Data Service 2020). These data give the work location of residents from each area unit in New Zealand. We then assume that working potential CFRs spend their time in the area unit of their work location and nonworking potential CFRs spend their time in the area unit in which they live. The resulting profiles have, on average, 62% of people on the diagonal, so a CFR recruited from a certain area unit will, on average, contribute to another part of the city for 38% of the "working hours" time block. For the profiles during nights and weekends, we assume that everyone stays in the area unit where they live.

This fitting process is not perfect, but we consider the approach sufficiently realistic to illustrate our models. Indeed, a closely related approach was used in Cont et al. (2021) to model population mobility in a study of intervention policies during the COVID-19 pandemic. Alternative sources were infeasible for this case study because they either were not sufficiently fine-grained (Facebook Data For Good 2021) or were not available for Auckland (Safegraph 2021). We discuss profile estimation impact on performance and provide a sensitivity analysis in Online Appendix 9.

### 5.4. Performance Metrics

To measure late arrivals, we consider the two response-time targets under which the SJAS operates: 95% of their incidents should be reached within 12 minutes, and 50% should be reached within six minutes (Ministry of Health 2020). The other metric we consider is patient survival, which we quantify using (6). This function was fitted on data where both ambulances and volunteers responded and volunteers did not carry a defibrillator. This function takes as input $t_{CPR}$ and $t_{EMS}$, which are measured from patient collapse, unlike response time, which is measured from call initiation. In line with Waalewijn et al. (2001)'s estimate, we add one minute (median time to call initiation) to account for this difference.

## 6. Numerical Results

In this section, we answer the four questions posed in the introduction: (1) in what areas is introducing a CFR app an efficient way to reach certain response-time targets? (Section 6.1); (2) what is the benefit of introducing a CFR app for a base of already trained volunteers, with known home addresses? (Section 6.2); (3) what is the benefit of a CFR app with a given number of registered volunteers when you do not know where they live or spend their time? (Section 6.3); and (4) where to recruit additional CFRs for an existing app with a known current CFR base? (Section 6.4). The first two questions are answered using the models from Section 3; the last two questions relate to the optimization models presented in Section 4. For clarity, ambulances are omitted in Section 6.1. The results in all other sections incorporate ambulance response as outlined in Section 5.2.2. We illustrate our calculations with data from the case study described in Section 5.

### 6.1. CFR Density Requirements

We next analyze the response-time distributions derived in Equation (2) and, for clarity, omit ambulances. The CFR response-time distributions are plotted in Figure 1 in Online Appendix 2. These allow us to calculate the required CFR density to reach a given response-time target in closed form. This can be used to investigate in what areas introducing a CFR app would be an efficient way to meet a certain target. Table 1 gives the required density of CFRs to meet a target that is defined by a combination of a time limit and a compliance level.

The results presented in Table 1 are generic in the sense that they apply to any location. They assume a fixed walking speed and pretrip delay. The table gives the required density of *available* CFRs; the required number of *registered* CFRs can be obtained by dividing the given density by the acceptance probability $\alpha$.

We illustrate the application of these results for the city of Auckland, where the observed value of $\alpha$ is 0.14. For each area unit, we can compute the fraction of the area-unit population that needs to be registered in order to satisfy the response-time requirements. We take both requirements (90% in 12 minutes and 50% in six minutes) and observe that the latter requires more CFRs: reaching this target requires a density of $1.38/0.14 = 9.19$ registered CFRs per square kilometer.

Looking at the population density of Auckland, we see that this density translates to 0.42% of the population registered as CFRs. However, the population density fluctuates over the city, so we need different registration rates in different parts of the city. A uniform registration rate of 0.42% of the population would meet the target only in 217 out of 287 area units. The required registration rate per area unit fluctuates between 0.08% and 130% of the population. Under the assumption that registration rates will not exceed 1% (a number currently achieved in countries with high participation), at least 23 area units (covering 3.7% of the population) will not be able to meet the specified target. Thus, even in an urban area like Auckland, the population density is not sufficiently consistent to obtain good response times throughout. Given that organizations are evaluated based on aggregate performance, one might observe different CFR densities in different area units that, together, do achieve an overall response-time requirement, which we explore next.

### 6.2. Quantifying the Benefit of a Given CFR Base

Next, we quantify the benefit of introducing an alert app in a scenario with a known base of already trained CFRs and a given ambulance distribution. As already noted, we have been contacted by a party facing exactly this question. They have a base of trained CFRs that is normally used at events such as festivals or parades, and they are considering introducing an alert system to expand CFR impact.

When we say a given CFR base is known, we mean that $n$ and $\nu$ are known, for example, because CFRs registered their home and/or work address or because they allow GPS tracking. We quantify how such a known CFR base would contribute to the fraction of late arrivals and the survival probability, and we explore how this depends on the acceptance rate $\alpha \in [0, 0.5]$. We vary $\alpha$ because it is, to some extent, controllable, and there is uncertainty about its true value. This range covers values that seem plausible given Auckland data and conversations with app developers.

As the party that raised this question wants to remain anonymous, we illustrate our calculations for Auckland. We consider the case where 1% of the population of Auckland are trained CFRs and a probability measure $\nu$ that is proportional to the estimated demand. In contrast to the results from Table 1 for which it was sufficient to consider every area unit separately, this question requires us to look at the entire city with varying CFR and incident densities, aggregating results over the city of Auckland.

Figure 2(a) shows that bringing down response times that exceed the 12-minute target from 4.8% (without using CFRs) to 1% requires an $\alpha$ of 9% (1,000 available CFRs). Although not depicted in the figure, to satisfy the requirement of reaching 50% of the incidents within six minutes, $\alpha \approx 2.7\%$ suffices. In Figure 2(b), we see that even for low values of $\alpha$, CFRs contribute substantially to survival. Convergence to the upper limit of survival is very slow and cannot be achieved by only increasing $\alpha$. Additional recruitment would be needed, which we discuss later.

**Figure 2.** (Color online) Impact of Increasing the CFR Availability $\alpha$ on Fraction of Late Arrivals and Survival, Assuming 1% of the Population Is Signed Up as CFRs, Distributed over the Area Units Proportionally to Demand

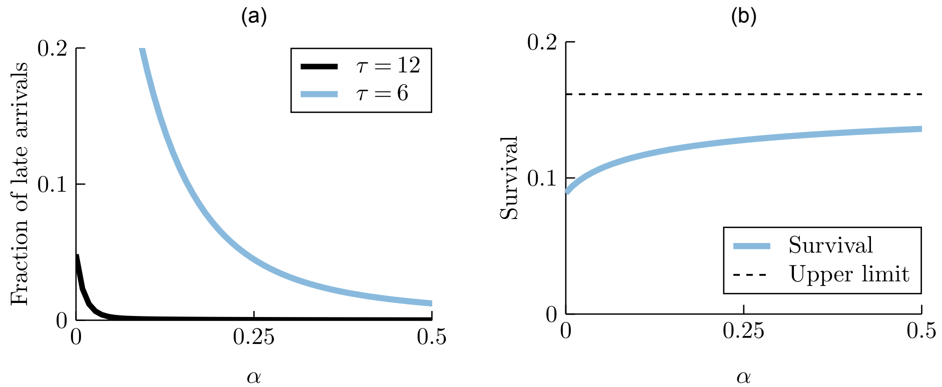


*Notes*. (a) Late arrivals. (b) Survival.

### 6.3. Quantifying the Impact of CFRs' Geographical Spread

In the previous section, we assumed that the CFR distribution $\nu$ was known; in contrast, we now investigate the scenario in which it is unknown where CFRs spend their time. Practitioners in Auckland are facing this scenario and would like to quantify the benefit of their app alongside the existing ambulance response. Whereas a retrospective study is possible, this is difficult for recently introduced CFR systems, as they have limited data available, and accurately estimating the response-time distribution requires many data points per area unit.

First, we consider two scenarios: one where CFRs are uniformly distributed over the city and another where CFRs are distributed proportionally to demand. The uniform distribution serves as a useful reference, whereas the proportional distribution is plausible. Second, we bound the performance of a CFR base of given size, using the optimization models from Section 4. These result in a lower bound for the fraction of late arrivals and an upper bound for the survival probability with a given total number of CFRs.

For a city that operates a CFR app, the values of $n$ and $\alpha$ are likely known. However, to provide more insight, we perform this case study for different values of $n\alpha$, the *expected number of available CFRs*. As we shall see, different values result in different optimal CFR allocations to area units, requiring us to solve the optimization problem separately for each value of $n\alpha$. We run our calculations up to 5,000 available CFRs, which, with the observed $\alpha$ of 0.14, corresponds to approximately 3.2% of the overall population of 1.1 million.

Recall that for both objectives—late arrivals and patient survival—a greedy approach is optimal. When optimizing late arrivals, an exact implementation (up to numerical precision) of the greedy method exists, using Equation (4). In contrast, when optimizing patient survival, the greedy method requires a step size $\epsilon$, which we set to 1/10,000. This means that we allocate the total number of CFRs in portions of size $n/10{,}000$ to the area unit where they would most improve the objective.

Figure 3 depicts the performance of the three CFR distributions for two different objectives. Figure 3 shows that the impact of the system depends on the geographical distribution of the CFRs. For example, in Figure 3(a), we can see that to bring the late arrivals down to 10%, one needs either 1,400 optimally distributed CFRs, 1,900 uniformly distributed CFRs, or 1,800 proportionally distributed CFRs. In particular, in this part of the graph, the proportional distribution outperforms the uniform distribution. In contrast, for a large number of CFRs, the uniform distribution outperforms the proportional one and approaches the optimal performance. This observation is even more apparent in the case with a homogeneous ambulance response, for which the results are given in Online Appendix 10. Not only is the performance similar but the optimal volunteer distribution also closely resembles the uniform distribution, as illustrated in Figure 4. This figure compares the geographical distribution of an optimal $\nu$ for both a limited number of CFRs ($n\alpha = 500$) and a larger number of CFRs ($n\alpha = 5{,}000$). Distinguishing two objectives (late arrivals and patient survival) yields four scenarios. In particular, Figure 4(c) shows a distribution of CFRs that resembles a uniform distribution, only distorted by the ambulances. This is in line with our earlier observations regarding the greedy algorithm. Figure 4(a) shows that for a limited number of CFRs, area units with a low call density do not receive any CFR mass.

For survival (Figure 4, (b) and (d)), the uniform distribution does not emerge as $n\alpha$ grows; it remains optimal to focus on high-demand areas even with 5,000 available CFRs, in accord with the close-to-optimal performance seen in Figure 3 for the proportional distribution. For a high number of CFRs, low-demand areas receive slightly more CFRs than a proportional allocation would assign to them, as is further detailed in Online Appendix 11.

**Figure 3.** (Color online) Impact of Increasing Number of Available CFRs ($n\alpha$) for Different CFR Distributions ($\nu$) on the Fraction of Late Arrivals ($\tau = 6$ min) and the Survival Probability

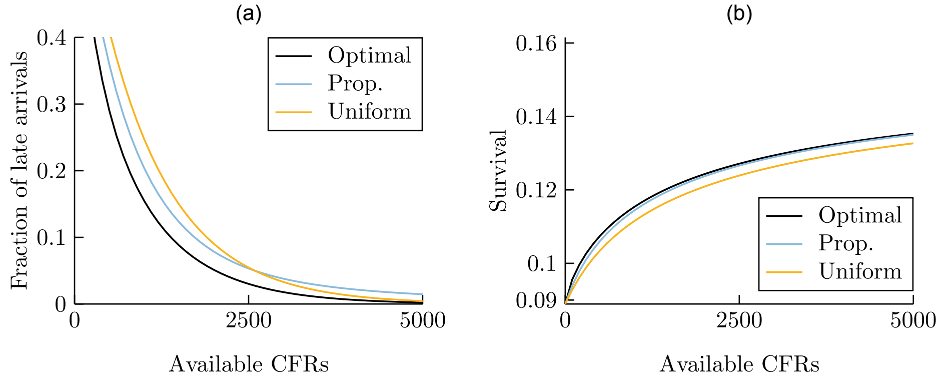


*Notes.* The optimal lines are each formed by the solution of 50 optimization problems. (a) Late arrivals $\tau = 6$ min. (b) Survival.

**Figure 4.** (Color online) Optimal CFR Distribution (Fraction of CFRs in That Region Divided by Surface Area)

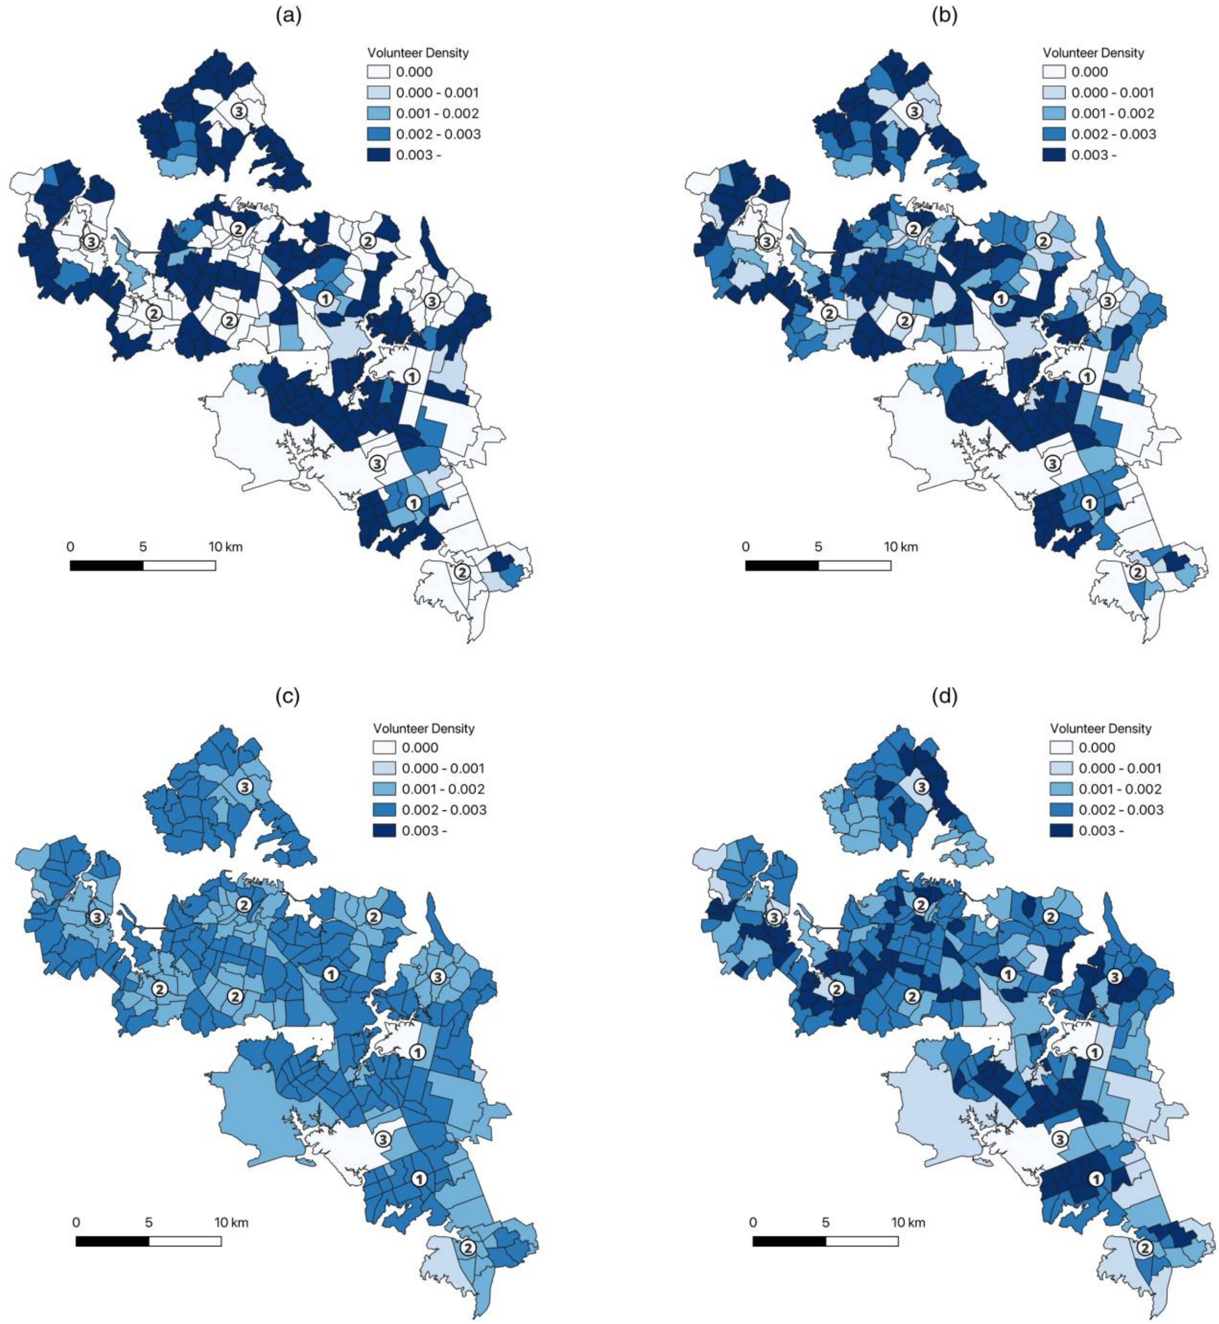


*Notes*. (a) and (c) Minimizing late arrivals for $\tau = 6$ min. (b) and (d) Maximizing survival, with $n\alpha = 500$ or $n\alpha = 5{,}000$. Darker colors represent higher CFR density. The white circles represent the location of the ambulance bases, and the numbers within them indicate the number of ambulances stationed at each base. (a) Late arrival, $n\alpha = 500$. (b) Survival, $n\alpha = 500$. (c) Late arrival, $n\alpha = 5{,}000$. (d) Survival, $n\alpha = 5{,}000$.

## 6.4. Optimal Recruitment Efforts

Next, we investigate efficient recruitment strategies while taking into account the current ambulance bases. We explore *profile recruitment*, where we take into account CFR presence across the city depending on where they live. In the previous upper bound calculations, we optimized for the available CFR location distribution $\nu$, which, in a scenario without profiles, is equal to $x$. When CFR locations are driven by profiles, the relationship between $x$ and $\nu$ is described by Equation (8). Not every desired CFR probability mass $\nu$ is attainable by adjusting recruitment mass $x$ per area unit. To find the optimal recruitment strategy, we use the formulation given in Section 4.4 and the method for obtaining profiles described in Section 5.3.

The impact of profiles on the optimal recruitment strategy can best be illustrated when we consider a homogeneous ambulance response throughout the city.

We contrast the difference in recruitment per area unit with and without profiles in Online Appendix 12. There, we see that some area units require no recruitment under profiles because of the influx of CFRs from other area units. Some of these area units, most notably the CBD, have such a high influx of CFRs that, even without recruitment from that area unit, the density of CFRs exceeds the optimal density without profiles. Most other area units have slightly higher recruitment under profiles to compensate for the outflux of CFRs.

**6.4.1. Time-Dependent Profiles.** Thus far, we have ignored a time component so we could isolate the importance of capturing CFR movement. Next, we consider the more realistic case where profiles have both a time and location component, as modeled in Section 5.3. We apply the models from Section 4.4 to investigate optimal recruitment under these time-dependent profiles. We illustrate our results for a survival objective and the equivalent of 500 available CFRs. Because the value of $\alpha$ is now time dependent, we cannot fix the value of $n\alpha$ throughout the day but, instead, fix $n$ to $n = 500/0.14 = 3,571$.

Figure 5 shows the amount of recruitment as well as the resulting CFR density per area unit. The area around the CBD is easily recognized as an area with high influx (dark blue in Figure 5(b)), so it requires little to no recruitment. Only in the CBD itself, a minor recruitment effort is still advised because of the disproportionately high demand there.

The solution in Figure 5 results in a 24-hour-average survival rate of 10.58% city-wide, which breaks down to 10.74% for day and 10.21% at night. Optimizing for daytime only, the survival rate during the day would be 10.75%, but this improvement would not outweigh the lives sacrificed during the night.

Our case study demonstrates the possibility of capturing profiles in recruitment. Ignoring profiles altogether (both their time and location component) would lead to a minor reduction in survival: 10.56% instead of 10.58%. This difference, though small, is unnecessary. To provide some perspective, a survival difference of this magnitude is the equivalent of recruiting and training approximately 80 extra CFRs (2% of the CFR base).

## 7. Conclusion and Discussion

We introduced models to quantify the impact of CFR systems alongside the existing ambulance response. These models were subsequently used to answer four questions that are central in practice in Sections 6.1–6.4, using data from the city of Auckland, New Zealand. The results yield important insights, such as how the performance depends on the number and location of CFRs. Whereas the results and insights are then specific to Auckland, the same models and approaches may be used to explore the potential impact of such an app in new cities or to guide the recruitment of additional CFRs in existing systems.

**Figure 5.** (Color online) Time-Dependent Profile Recruitment That Maximizes Survival for $n = 500/0.14 = 3,571$

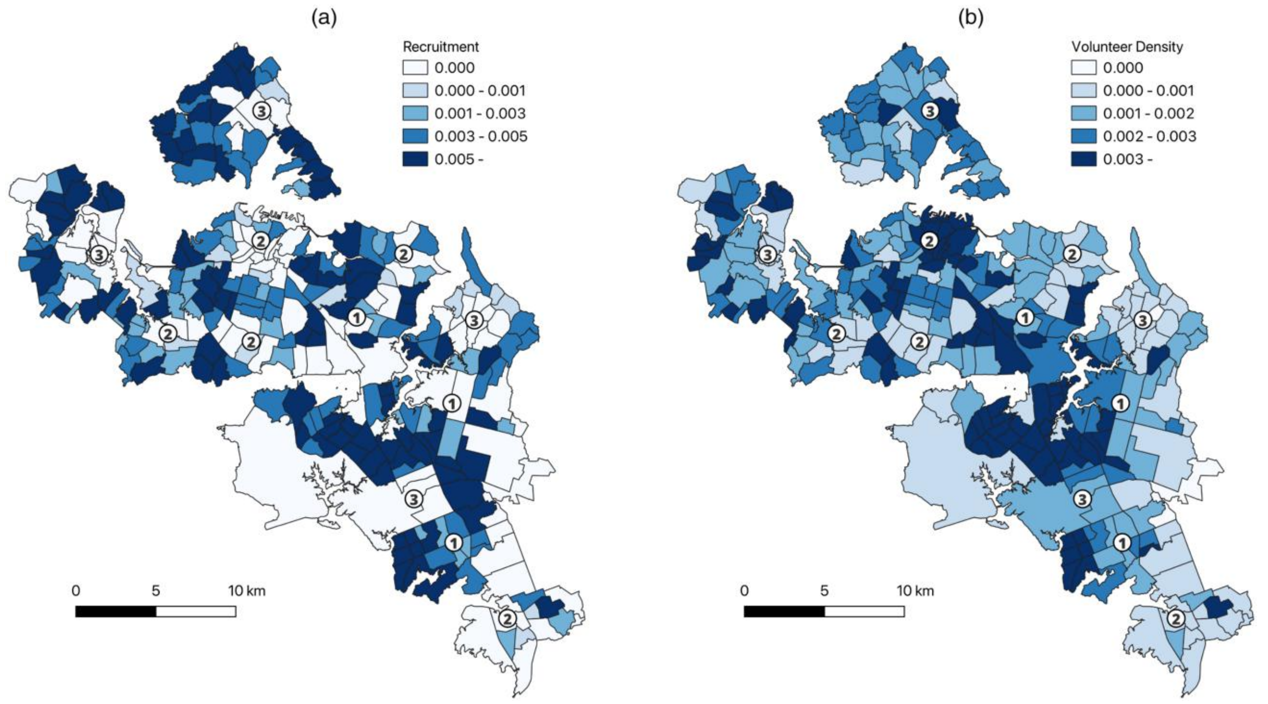


*Notes*. The white circles represent the location of the ambulance bases, and the numbers within them indicate the number of ambulances stationed at each base. (a) Optimal recruitment ($x_i$ values). (b) Resulting daytime density ($\nu_{l2}$).

Figures 2 and 3 show the steady improvement in two performance metrics (late arrival and patient survival) over a wide range of numbers of available CFRs that complement the existing ambulance contribution. The marginal benefit of additional CFRs remains practically important, even beyond optimistic registration rates (~1%). Besides increasing the number of CFRs $n$, one may take measures to increase the availability probability $\alpha$. The importance of increasing $\alpha$ is underlined by the insight that doubling $\alpha$ has the same impact as doubling the number of CFRs. Thus, systems must be designed to maximize $\alpha$.

Every additional CFR improves the system-wide response time, but where those CFRs spend their time matters. Thus, it makes sense to steer recruitment efforts to the right areas. These areas can be determined with the models from Section 4. An important empirical observation from Section 6 is that optimal CFR densities heavily depend on the objective. To maximize cardiac arrest survival, a volunteer distribution proportional to demand is nearly optimal. Conversely, for late arrivals, we see clear differences between uniform volunteer densities (if the time threshold is loose) and clustered volunteers (if the time threshold is strict). Thus, the objective heavily influences the equity of the final solution. A uniform distribution is, at least in the absence of ambulances, perhaps the most equitable. However, having equal volunteer densities everywhere would require an impractically high sign-up rate in sparsely populated areas. If equity needs to be explicitly captured, our model can be extended by specifying a lower bound on the performance in each area without overly complicating the solution method.

For OHCA, we consider maximizing survival an appropriate objective, and under that assumption, a first conclusion appears that incident rates can guide recruitment efforts. However, CFRs do not spend all of their time in the area unit where they live, and these effects typically do not even out. Consequently, even without *any* recruitment from the city center, our daytime scenario showed an overresourcing of CFRs in the city center. We expect similar effects in other cities. The results from Section 6.4 show that this indeed makes a significant difference in terms of where to best recruit. If one knows how volunteers tend to move through the city, these dynamics can be incorporated in the form of profiles, and our method remains able to find the optimal recruitment. An additional use of the profile concept is to quantify the marginal benefit of training a new CFR based on that person's profile, which could inform, for example, an organization's willingness to fund the CFR's training. In this manner, our model allows personalized advice and ensures that a limited training budget can be spent in an effective way.

As mentioned in the introduction, the practical value of our work has already been demonstrated through studies with the Fire Department Amsterdam-Amstelland in the Netherlands and a Red Cross organization (anonymous, but at a national level). Similar ideas are applicable in settings with many uncontrollable servers that can thus be modeled as randomly situated throughout an area. One of the authors has used these ideas in that manner, but confidentiality agreements prevent the disclosure of the setting.

There are many important directions for future research. First, a detailed analysis of data from CFR schemes is needed to understand CFR responses. For example, how do acceptance probabilities depend on the distance to the patient or personal attributes of the CFR such as age or mobile phone specifications? Second, how should dispatch strategies be designed, for example, given the distances of nearby CFRs, which should be notified in a first batch of alerts, and what is the benefit to retaining notifications for very close CFRs after a more distant CFR accepts? Should one avoid overutilization of CFRs and navigate a trade-off between notifying CFRs for the current task and saving them for future ones? This question is posed in more detail in Henderson et al. (2022) and explored in a more general setting in Manshadi and Rodilitz (2022).

It may appear that we have ignored the impact of bystanders on survival rates. Even though bystanders are unlikely to be trained in CPR, De Maio et al. (2003) found that most patients who survived received some sort of bystander intervention. This bystander help was already included in the patient data that was used to train survival functions. In the absence of more comprehensive data, we have tried to incorporate the bystander effect to the best of our ability by using survival functions that account for bystander involvement.

In the context of OHCA, whether an AED is available has a significant impact on survival. For example, what is the additional benefit of taxi drivers carrying AEDs in their cars and being dispatched to nearby incidents, even if they do not have medical training? What survival benefits can be expected if an AED gets delivered to the scene by a drone? How valuable would it be to provide a small fraction of volunteers with their own AED? A mix of volunteers carrying AEDs supplemented with drone-delivered AEDs could be optimized by extending our model. Some of these questions may require a model with multiple classes of CFRs, which is an interesting direction for future work.

## Acknowledgments

The authors thank Bridget Dicker and Dave Richards from St John Ambulance Services New Zealand for helpful discussions and data. The authors also thank the organizers of the First International Workshop on Planning of Emergency Services, which included a discussion of Jansma (2014) that partly inspired this work. The authors are grateful to Verity Todd and Nick Garrett for providing them with a detailed relationship between demographics and OHCA rates in New Zealand. The authors thank Sam Ridler for providing them travel times for EMS vehicles in Auckland, the Research Software Engineering & Consulting (RSEC) team at RSM for data storage support, and M. Eng. students Jui-an Wang and Lars Kouwenhoven for coding assistance. Some of this work has been executed under the TKI Dinalog grant 2023-1-307TKI.

## References

Ata B, Lee D, Sönmez E (2019) Dynamic volunteer staffing in multi-crop gleaning operations. *Oper. Res.* 67(2):295–314.

Bandara D, Mayorga ME, McLay LA (2012) Optimal dispatching strategies for emergency vehicles to increase patient survivability. *Inernat. J. Oper. Res.* 15(2):195–214.

Barbour AD, Holst L, Janson S (1992) *Poisson Approximation*. Oxford Studies in Probability, vol. 2 (Clarendon Press, Oxford, UK).

Berdowski J, Berg RA, Tijssen JG, Koster RW (2010) Global incidences of out-of-hospital cardiac arrest and survival rates: Systematic review of 67 prospective studies. *Resuscitation* 81(11):1479–1487.

Bomze IM, Rinaldi F, Zeffiro D (2020) Active set complexity of the away-step Frank–Wolfe algorithm. *SIAM J. Optim.* 30(3):2470–2500.

Boutilier JJ, Brooks SC, Janmohamed A, Byers A, Buick JE, Zhan C, Schoellig AP, Cheskes S, Morrison LJ, Chan TCY (2017) Optimizing a drone network to deliver automated external defibrillators. *Circulation* 135(25):2454–2465.

Boutilier JJ, Chan TCY (2022) Drone network design for cardiac arrest response. *Manufacturing Service Oper. Management* 24(5):2407–2424.

Brandweer Amsterdam-Amstelland (2021) Beleidsplan 2021–2024. https://openresearch.amsterdam/image/2021/7/7/eindconcept_vraa_beleidsplan_brandweer.pdf.

Brooks SC, Simmons G, Worthington H, Bobrow BJ, Morrison LJ (2016) The PulsePoint Respond mobile device application to crowdsource basic life support for patients with out-of-hospital cardiac arrest: Challenges for optimal implementation. *Resuscitation* 98:20–26.

Chan TC, Demirtas D, Kwon RH (2016) Optimizing the deployment of public access defibrillators. *Management Sci.* 62(12):3617–3635.

Chan TC, Shen ZJM, Siddiq A (2018) Robust defibrillator deployment under cardiac arrest location uncertainty via row-and-column generation. *Oper. Res.* 66(2):358–379.

Chu J, Leung KB, Snobelen P, Nevils G, Drennan IR, Cheskes S, Chan TC (2021) Machine learning-based dispatch of drone-delivered defibrillators for out-of-hospital cardiac arrest. *Resuscitation* 162:120–127.

Cont R, Kotlicki A, Xu R (2021) Modelling COVID-19 contagion: Risk assessment and targeted mitigation policies. *R. Soc. Open Sci.* 8(3):201535.

Daskin M (1983) A maximum expected location model: Formulation, properties and heuristic solution. *Transportation Sci.* 7:48–70.

De Maio V, Stiell I, Wells G, Spaite D (2003) Optimal defibrillation for maximum out-of-hospital cardiac arrest survival rates. *Ann. Emergency Medicine* 42:242–250.

Dicker B, Garrett N, Wong S, McKenzie H, McCarthy J, Jenkin G, Smith T, et al. (2019) Relationship between socioeconomic factors, distribution of public access defibrillators and incidence of out-of-hospital cardiac arrest. *Resuscitation* 138:53–58.

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

Facebook Data For Good (2021) Colocation maps. Accessed August 24, 2021, https://dataforgood.facebook.com/dfg/tools/colocation-maps.

Folke F, Lippert FK, Nielsen SL, Gislason GH, Hansen ML, Schramm TK, Sørensen R, et al. (2009) Clinical perspective. *Circulation* 120(6):510–517.

Global Resuscitation Alliance (2019) GoodSAM—A smart phone application to crowdsource CPR in New Zealand. Accessed June 6, 2023, https://www.globalresuscitationalliance.org/wp-content/uploads/2019/12/St_John_NZ_Smart_Tech.pdf.

GoodSAM Platform (2020) GoodSAM. Accessed February 29, 2020, https://www.goodsamapp.org/.

HartslagNu (2020) HartslagNu. Accessed December 14, 2020, https://hartslagnu.nl/.

Henderson SG, Berg PL, Jagtenberg CJ, Li H (2022) How should volunteers be dispatched to out-of-hospital cardiac arrest cases? *Queueing Systems* 100:437–439.

Jansma T (2014) Extending the emergency medical services network for out-of-hospital cardiac arrest victims: An explorative study for the province of Drenthe. Master's thesis, University of Groningen, Groningen, Netherlands.

Johari R, Kamble V, Kanoria Y (2021) Matching while learning. *Oper. Res.* 69(2):655–681.

Kingman JFC (1993) *Poisson Processes* (Oxford University Press, Oxford, UK).

Larson RC, Odoni AR (1981) *Urban Operations Research* (Prentice Hall, Englewood Cliffs, NJ).

Liu W, Sun Q, Ching Tang L, Ye Z (2022) Robust data-driven design of a smart cardiac arrest response system. Preprint, submitted October 2, http://dx.doi.org/10.2139/ssrn.4590433.

Lives. Emergency responders. Accessed October 24, 2022, https://www.lives.org.uk/what-we-do/emergency-responders/.

Manshadi V, Rodilitz S (2022) Online policies for efficient volunteer crowdsourcing. *Management Sci.* 68(9):6572–6590.

Matinrad N, Granberg TA, Angelakis V (2021) Modeling uncertain task compliance in dispatch of volunteers to out-of-hospital cardiac arrest patients. *Comput. Indust. Engrg.* 159:107515.

Matinrad N, Granberg T, Ennab Vogel N, Angelakis V (2019) Optimal dispatch of volunteers to out-of-hospital cardiac arrest patients. *Proc. 52nd Hawaii Internat. Conf. System Sci.* (Association for Information Systems, Atlanta), 4088–4097.

McElfresh DC, Kroer C, Pupyrev S, Sodomka E, Sankararaman KA, Chauvin Z, Dexter N, Dickerson JP (2020) Matching algorithms for blood donation. *Proc. 21st ACM Conf. Econom. Comput.* (Association for Computing Machinery, New York), 463–464.

McLay LA (2009) Emergency medical service systems that improve patient survivability. Cochran JJ, Cox LA, Keskinocak P, Kharoufeh JP, Smith JC, eds. *Wiley Encyclopedia of Operations Research and Management Science* (John Wiley & Sons, Hoboken, NJ).

Ministry of Health (2020) Emergency ambulance service national performance report. Accessed May 12, 2021, https://www.health.govt.nz/system/files/documents/pages/naso-march-report-2019-2020.pdf.

Nazarian A (2018) Optimizing the deployment of automated external defibrillators by a data-driven algorithmic approach. Master's thesis, University of Twente, Enschede, Netherlands.

NHS North West Ambulance Service. Community first responders (CFR). Accessed October 24, 2022, https://www.nwas.nhs.uk/get-involved/volunteering/community-first-responder/.

Nichol G, Stiell IG, Laupacis A, Pham B, De Maio VJ, Wells GA (1999) A cumulative meta-analysis of the effectiveness of defibrillator-capable emergency medical services for victims of out-of-hospital cardiac arrest. *Ann. Emergency Medicine* 34(4):517–525.

Nisingizwe MP, Ndishimye P, Swaibu K, Nshimiyimana L, Karame P, Dushimiyimana V, Musabyimana JP, Musanabaganwa C, Nsanzimana S, Law MR (2022) Effect of unmanned aerial vehicle (drone) delivery on blood product delivery time and wastage in Rwanda: A retrospective, cross-sectional study and time series analysis. *Lancet Global Health* 10(4):564–569.

Northern Ireland Ambulance Service Health and Social Care Trust. Community first responder scheme. Accessed April 9, 2024, https://nias.hscni.net/calling-999/who-will-treat-you/first-responders/.

Oving I, Masterson S, Tjelmeland IB, Jonsson M, Semeraro F, Ringh M, Truhlar A, et al. (2019) First-response treatment after out-of-hospital cardiac arrest: A survey of current practices across 29 countries in Europe. *Scandinavian J. Trauma Resuscitation Emergency Medicine* 27(1):112.

Özkan E, Ward AR (2020) Dynamic matching for real-time ride sharing. *Stochastic Systems* 10(1):29–70.

Phung VH, Trueman I, Togher F, Orner R, Siriwardena AN (2017) Community first responders and responder schemes in the United Kingdom: Systematic scoping review. *Scandinavian J. Trauma Resuscitation Emergency Medicine* 25(1):58.

Pijls RW, Nelemans PJ, Rahel BM, Gorgels AP (2016) A text message alert system for trained volunteers improves out-of-hospital cardiac arrest survival. *Resuscitation* 105:182–187.

Pijls RW, Nelemans PJ, Rahel BM, Gorgels AP (2019) Characteristics of a novel citizen rescue system for out-of-hospital cardiac arrest in the Dutch province of Limburg: Relation to incidence and survival. *Netherlands Heart J.* 27(2):100–107.

PulsePoint (2020) PulsePoint. Accessed February 29, 2020, https://www.pulsepoint.org/.

Qin Z, Tang X, Jiao Y, Zhang F, Xu Z, Zhu H, Ye J (2020) Ride-hailing order dispatching at DiDi via reinforcement learning. *INFORMS J. Appl. Anal.* 50(5):272–286.

Ridler S (2020) Emergency medical services research at the University of Auckland, Department of Engineering Science. Accessed April 9, 2024, https://github.com/uoa-ems-research.

Ridler S, Mason AJ, Raith A (2017) A simulation package for emergency medical services. *Proc. 51st Annual Conf. ORSNZ* (Operations Research Society of New Zealand, Hamilton, New Zealand).

Ridler S, Mason AJ, Raith A (2022) A simulation and optimisation package for emergency medical services. *Eur. J. Oper. Res.* 298(3):1101–1113.

Ringh M, Fredman D, Nordberg P, Stark T, Hollenberg J (2011) Mobile phone technology identifies and recruits trained citizens to perform CPR on out-of-hospital cardiac arrest victims prior to ambulance arrival. *Resuscitation* 82(12):1514–1518.

Safegraph (2021) Safegraph. Accessed March 14, 2023, https://www.safegraph.com/.

Sasson C, Rogers MA, Dahl J, Kellermann AL (2010) Predictors of survival from out-of-hospital cardiac arrest: A systematic review and meta-analysis. *Circulation Cardiovascular Quality Outcomes* 3(1):63–81.

Serfling RJ (1980) *Approximation Theorems of Mathematical Statistics.* Wiley Series in Probability and Mathematical Statistics (Wiley, New York).

Shin J, Marla L, Boutilier JJ (2022) Heterogeneous facility location under coordination and collaboration: The case of ambulance-bystander-drone coordination. Preprint, submitted June 6, https://ssrn.com/abstract=3519243.

Slaa G (2020) Increasing cardiac arrest survival by improving the volunteer alerting algorithm. Accessed November 30, 2023, http://essay.utwente.nl/80432/.

Smith CM, Wilson MH, Ghorbangholi A, Hartley-Sharpe C, Gwinnutt C, Dicker B, Perkins GD (2017) The use of trained volunteers in the response to out-of-hospital cardiac arrest–The GoodSAM experience. *Resuscitation* 121:123–126.

Stats New Zealand Geographic Data Service (2015) Age by meshblock (2013 census). Accessed October 25, 2021, https://datafinder.stats.govt.nz/layer/8447-age-by-meshblock-2013-census/.

Stats New Zealand Geographic Data Service (2016) Area units 2013. Accessed October 28, 2020, https://datafinder.stats.govt.nz/layer/25743-area-unit-2013/.

Stats New Zealand Geographic Data Service (2020) 2018 census main means of travel to work by statistical area 2. Accessed October 25, 2021, https://datafinder.stats.govt.nz/table/104720-2018-census-main-means-of-travel-to-work-by-statistical-area-2/.

Stieglis R, Zijlstra JA, Riedijk F, Smeekes M, van der Worp WE, Koster RW (2020) AED and text message responders density in residential areas for rapid response in out-of-hospital cardiac arrest. *Resuscitation* 150:170–177.

Tafreshian A, Masoud N, Yin Y (2020) Frontiers in service science: Ride matching for peer-to-peer ride sharing: A review and future directions. *Service Sci.* 12(2–3):44–60.

Valenzuela TD, Roe DJ, Cretin S, Spaite DW, Larsen MP (1997) Estimating effectiveness of cardiac arrest interventions: A logistic regression survival model. *Circulation* 96(10):3308–3313.

Waalewijn RA, de Vos R, Tijssen JG, Koster RW (2001) Survival models for out-of-hospital cardiopulmonary resuscitation from the perspectives of the bystander, the first responder, and the paramedic. *Resuscitation* 51(2):113–122.

Wikipedia (2020) Burgerhulpverlening. Accessed December 14, 2020, https://nl.wikipedia.org/wiki/Burgerhulpverlening.

Yan S, Gan Y, Jiang N, Wang R, Chen Y, Luo Z, Zong Q, Chen S, Lv C (2020) The global survival rate among adult out-of-hospital cardiac arrest patients who received cardiopulmonary resuscitation: A systematic review and meta-analysis. *Critical Care* 24(1):61.

Zaffar MA, Rajagopalan HK, Saydam C, Mayorga M, Sharer E (2016) Coverage, survivability or response time: A comparative study of performance statistics used in ambulance location models via simulation optimization. *Oper. Res. Health Care* 11:1–12.

Zijlstra JA, Stieglis R, Riedijk F, Smeekes M, Van der Worp WE, Koster RW (2014) Local lay rescuers with AEDs, alerted by text messages, contribute to early defibrillation in a Dutch out-of-hospital cardiac arrest dispatch system. *Resuscitation* 85(11):1444–1449.