

Robust Defibrillator Deployment Under Cardiac Arrest Location Uncertainty via Row-and-Column Generation

Timothy C. Y. Chan,^a Zuo-Jun Max Shen,^{b,c} Auyon Siddiq^b

^aDepartment of Mechanical and Industrial Engineering, University of Toronto, Toronto, Ontario M5S 1A1, Canada; ^bDepartment of Industrial Engineering and Operations Research, University of California, Berkeley, Berkeley, California 94720; ^cDepartment of Civil and Environmental Engineering, University of California, Berkeley, Berkeley, California 94720

Contact: tcychan@mie.utoronto.ca (TCYC); maxshen@berkeley.edu,  <http://orcid.org/0000-0003-4538-8312> (Z-JMS); auyon.siddiq@berkeley.edu (AS)

Received: October 1, 2015

Revised: December 15, 2016; April 9, 2017

Accepted: May 3, 2017

Published Online in Articles in Advance:
October 25, 2017

Subject Classifications: Facilities: location;
programming: integer

Area of Review: Operations and Supply Chains

<https://doi.org/10.1287/opre.2017.1660>

Copyright: © 2017 INFORMS

Abstract. Sudden cardiac arrest is a significant public health concern. Successful treatment of cardiac arrest is extremely time sensitive, and use of an automated external defibrillator (AED) where possible significantly increases the probability of survival. Placement of AEDs in public locations can improve survival by enabling bystanders to treat victims of cardiac arrest prior to the arrival of emergency medical responders, thus shortening the time between collapse and treatment. However, since the exact locations of future cardiac arrests cannot be known a priori, AEDs must be placed strategically in public locations to ensure their accessibility in the event of an out-of-hospital cardiac arrest emergency. In this paper, we propose a data-driven optimization model for deploying AEDs in public spaces while accounting for uncertainty in future cardiac arrest locations. Our approach involves discretizing a continuous service area into a large set of scenarios, where the probability of cardiac arrest at each location is itself uncertain. We model uncertainty in the spatial risk of cardiac arrest using a polyhedral uncertainty set that we calibrate using historical cardiac arrest data. We propose a solution technique based on row-and-column generation that exploits the structure of the uncertainty set, allowing the algorithm to scale gracefully with the total number of scenarios. Using real cardiac arrest data from the City of Toronto, we conduct an extensive numerical study on AED deployment public locations. We find that hedging against cardiac arrest location uncertainty can produce AED deployments that outperform an intuitive sample average approximation by 9%–15% and cuts the performance gap with respect to an ex post model by half. Our findings suggest that accounting for cardiac arrest location uncertainty can lead to improved accessibility of AEDs during cardiac arrest emergencies and the potential for improved survival outcomes.

Funding: This research was partially supported by the National Science Foundation [CMMI 1265671], the National Natural Science Foundation of China [Grants 71210002 and 71332005], the Natural Sciences and Engineering Research Council of Canada, and the Ontario Ministry of Training, Colleges and Universities.

Supplemental Material: The electronic companion is available at <https://doi.org/10.1287/opre.2017.1660>.

Keywords: robust optimization • row-and-column generation • healthcare operations • facility location

1. Introduction

Sudden cardiac arrest is a leading cause of death in North America and is responsible for more than 400,000 deaths each year (Heart and Stroke Foundation 2013, Mozaffarian et al. 2016). Chances of survival have been estimated to decrease by as much as 10% with each minute of delay in treatment (Larsen et al. 1993), making treatment of cardiac arrest extremely time sensitive. Currently, the probability of survival from an out-of-hospital cardiac arrest is low, with only 5%–10% of out-of-hospital cardiac arrest victims surviving to hospital discharge (Weisfeldt et al. 2010).

Treatment of cardiac arrest involves cardiopulmonary resuscitation (CPR) and electrical shocks—known as defibrillation—from an automated external

defibrillator (AED). AEDs are portable devices that can automatically assess heart rhythms and perform defibrillation if necessary, making their use by nonprofessionals (lay responders) in a cardiac arrest emergency viable (Gundry et al. 1999). Treatment with a defibrillator following cardiac arrest is vital: the probability of survival from cardiac arrest has been estimated to be as high as 40%–70% if defibrillation is administered within three minutes of the victim's collapse (Valenzuela et al. 2000, Page et al. 2000, Caffrey et al. 2002). As a consequence, there is growing interest in the development of public access defibrillation (PAD) programs, which are organized efforts to place defibrillators in public areas, such as malls, coffee shops, restaurants, and subway stations, with the hopes of

improving easy access to defibrillators in the event of a cardiac arrest emergency. With appropriate positioning, public AEDs can enable lay responders to treat cardiac arrest victims while awaiting the arrival of professional emergency medical responders. Indeed, PAD programs have been shown to decrease the time to treatment and improve survival outcomes (Public Access Defibrillation Trial Investigators 2004), and it is estimated that their widespread implementation may save between 2,000 and 4,000 lives annually in the United States (Hazinski et al. 2005). Thus, the strategic positioning of public AEDs has a crucial role to play in strengthening the overall response to cardiac arrest.

Despite evidence that supports improved chances of survival with PAD programs, AEDs are used infrequently. For example, only 8% of public location cardiac arrests in Toronto, Canada involve the use of an AED by a bystander (Sun et al. 2016). Although the reasons for this low usage rate are multifaceted, it is clear that AEDs that are poorly positioned may be located too far from the locations of cardiac arrest emergencies for them to be accessed and used by lay responders. Recently, several mobile phone applications have been developed that aim to improve AED usage during public location cardiac arrests. These applications notify volunteer lay responders of the location of nearby AEDs so that they can be retrieved and used quickly (e.g., PulsePoint 2015). To be effective, these emerging technologies—and public access defibrillation in general—rely critically on AEDs being placed strategically in public locations. The question of where to place public access AEDs remains relatively open in the medical literature (Portner et al. 2004, Folke et al. 2009).

In this paper, we present a data-driven optimization model for deploying AEDs in public locations. We consider a setting where cardiac arrests occur continuously over a bounded service area. Our model addresses a key challenge in the AED deployment problem, which is the uncertainty in the locations of future cardiac arrests during the deployment of AEDs. Since the chances of survival from cardiac arrest diminish rapidly with each passing minute, it is critical that witnesses of cardiac arrest are able to quickly retrieve and use nearby AEDs when needed. Because of this extreme time sensitivity, protecting against cardiac arrest location uncertainty can help improve survival outcomes by shortening the distance (and thus the response time) to the nearest AED during a cardiac arrest emergency.

Although cardiac arrests are not uncommon within the general population, they are extremely rare events for any *particular* building or geographic location. Thus, any estimate of the risk of cardiac arrest at a specific location is likely to be subject to considerable uncertainty. To account for this uncertainty while still leveraging historical data, we take a distributionally robust optimization approach in which service

is optimized for the worst-case spatial distribution of cardiac arrests. Our approach involves discretizing the continuous service area into a large set of scenarios, where the probability vector that describes the probability of the next cardiac arrest arriving at each scenario is only known to reside within a polyhedral uncertainty set. We construct the uncertainty set by representing the service area as the union of a set of *uncertainty regions*, where only the aggregate probability of the next cardiac arrest arriving in each uncertainty region (i.e., group of scenarios) is known. An important consequence of modeling uncertainty in this manner is that it induces sparsity in the worst-case distribution (with respect to the number of scenarios), which we show can be exploited through an efficient row-and-column algorithm to obtain optimal solutions to the robust problem.

For a given AED deployment and a set of cardiac arrest locations, one can construct a *distance distribution*, which describes the histogram of distances between the cardiac arrests and the nearest AED in the given deployment. Inspired by the use of quantile-based targets in emergency medical services (Pell et al. 2001, Pons and Markovchick 2002), we use a conditional value-at-risk (CVaR) objective function to directly optimize the tail of the worst-case distance distribution, which allows us to mitigate the risk of unacceptably long distances between cardiac arrest locations and their nearest AEDs. An additional benefit of a CVaR-based objective function is that it permits more flexibility with respect to selecting a performance metric to optimize; for example, we can minimize the mean or tail of the worst-case distribution by using an appropriate parameterization of the model.

We summarize our contributions as follows:

1. We present a robust optimization model for AED deployment that accounts for uncertainty in the spatial distribution of cardiac arrests. Our model permits cardiac arrests to arrive anywhere within a continuous service area, which we approximate through a fine discretization. We use historical cardiac arrest data to construct a distributional uncertainty set that consists of a set of generally shaped uncertainty regions, where only the probability of the next cardiac arrest event occurring within each uncertainty region is known. We also present an auxiliary result that shows that our choice of uncertainty set and parameters unifies the canonical p -median and p -center problems, as well as their robust analogues.

2. We present a bound on the approximation error induced by the discretization of the service area, which depends on the granularity of the discretization and the configuration of the uncertainty regions. The bound can be used to quantify the suboptimality of the solutions produced by our discretized model with respect to the underlying continuous problem.

3. We propose a row-and-column generation algorithm for solving the robust AED deployment problem. Our proposed algorithm operates by exploiting the structure of the uncertainty set, which allows it to substantially outperform a standard Benders-based algorithm and a single-stage mixed integer programming (MIP) reformulation. Notably, our algorithm decouples the size of the master problem from the number of scenarios used to discretize the service area, which allows us to efficiently solve problems where the discretization is extremely fine. Combined with the bounds on the discretization error, the row-and-column generation algorithm allows us to efficiently obtain provably near-optimal solutions to the underlying continuous problem.

4. We present results from an extensive numerical study on the placement of public AEDs using real cardiac arrest data from Toronto, Canada. Our results demonstrate that accounting for uncertainty in cardiac arrest locations can decrease the distance between cardiac arrest victims and the nearest AED by 9%–15%, under both typical and worst-case realizations of the uncertainty.

2. Related Literature

Our paper primarily contributes to the medical literature on cardiac arrest and public access defibrillation, and the operations research literature on facility location under uncertainty. We also briefly review the connection between our model and the general literature on optimization under uncertainty.

2.1. Public Access Defibrillation

Most of the literature in the medical community on AED location has focused on identifying building types that have a high incidence of cardiac arrest, such as shopping malls (Becker et al. 1998, Engdahl and Herlitz 2005, Brooks et al. 2013). However, relying on building type alone to guide AED locations has several drawbacks. First, it does not account for cardiac arrests that occur outdoors, since these events cannot be assigned to a building type. Second, a high-risk building type in one city may not be high risk in another. Finally, identifying high risk buildings does not address the question of how to optimally deploy a limited number of defibrillators. By contrast, our proposed approach is guided by historical cardiac arrest data and is agnostic to building type, which makes it generalizable to any city.

There is a recent but limited body of work on optimization-based approaches to placing AEDs in public locations (Siddiq et al. 2013; Chan et al. 2013, 2016; Sun et al. 2016). Our paper differs from and extends this literature in a few substantive ways. First, previous approaches have solely optimized over exact historical cardiac arrest locations. While these approaches serve

as a good first step toward improving public access defibrillation, they do not explicitly account for uncertainty in future cardiac arrest locations. Our paper is the first to explicitly incorporate the uncertainty in cardiac arrest locations into the optimization model. We find that accounting for cardiac arrest location uncertainty leads to improved outcomes based on several performance metrics, compared with a baseline method of optimizing over historical locations only. Second, existing AED models focus on maximizing cardiac arrest *coverage*, which is the number of cardiac arrests that are within some fixed distance (e.g., 100 m) of an AED. While this approach can be effective for the *median* cardiac arrest patient, it can have the adverse effect of neglecting cardiac arrest victims at the tail of the distance distribution. By contrast, our modeling approach involves a risk-based objective function (inspired by ambulance response time guidelines), which allows us to mitigate the risk of unacceptably large distances between cardiac arrest victims and the nearest AED. Finally, existing work on AED location has focused on identifying general hotspots for cardiac arrest within a large service area (e.g., Chan et al. 2016 identify historical clusters of cardiac arrests within a city). We complement this work by taking a more tactical approach to the AED deployment problem, by focusing on the precise placement of AEDs within a small service area consisting of several city blocks.

2.2. Facility Location Under Uncertainty

Facility location problems under uncertainty have received considerable attention, particularly with respect to uncertainty in demand node weights and edge lengths (Owen and Daskin 1998, Snyder 2006, Baron et al. 2011), and more recently in the risk of service disruptions at the facilities (Lim et al. 2010, Cui et al. 2010, Shen et al. 2011). Our model differs from the existing literature in that we consider the arrival of demand points over a continuous, planar service area, rather than being restricted to a small set of locations. As a result, we do not aggregate demand into a small set of known locations, which is a common approach in facility location models related to emergency response (Cho et al. 2014, Erkut et al. 2008, Brotcorne et al. 2003) and in general (Francis et al. 2000). Previous work on demand location uncertainty in facility location problems is limited and has generally focused on the placement of a single facility (Cooper 1978, Drezner 1989, Averbakh and Bereg 2005). To the best of our knowledge, this paper is the first to extend the literature on demand location uncertainty to the general case of siting multiple facilities, which we show can be done in a computationally tractable manner.

Our paper is most closely related to work by Baron et al. (2011), who also present a robust optimization model for facility location. In their paper, demand is again restricted to a set of known locations, but the

vector of demands over all locations is only known to live within a given uncertainty set. By contrast, we model total demand as being known (i.e., normalized to 1) and consider uncertainty in how this demand is distributed in space.

Our work is also related to facility interdiction problems, which focus on protecting against potential disruptions at facilities. Facility interdiction models, which, like ours, take the form of min-max optimization problems, were first considered by Church and Scaparra (2007) and later by Scaparra and Church (2008), Liberatore et al. (2011), and Losada et al. (2012). With respect to CVaR, the only other work to use it in a facility location context is the *mean-excess* model presented in Chen et al. (2006), which focuses on demand-level uncertainty.

2.3. Optimization Under Uncertainty

Our main formulation takes the form of a robust optimization problem in which a large set of scenario probabilities are uncertain and only known to reside within a given polyhedron. In the special case where there is no uncertainty (i.e., the uncertainty set is a singleton representing the true distribution), our formulation simplifies to an extensive-form stochastic program with a large set of scenarios (Birge and Louveaux 2011). Other papers that model uncertainty in probabilities using polyhedral uncertainty sets include Chan and Mišić (2013), who consider radiation therapy with uncertainty in patient breathing patterns, and Farias et al. (2013), who consider uncertainty in consumer choice probabilities. Ben-Tal et al. (2013) present a general framework that uses ϕ -divergences to model distributional uncertainty sets, and they obtain polyhedral sets as a special case. For a review of the broader robust optimization literature, we refer the reader to Ben-Tal et al. (2009) and Bertsimas et al. (2011).

Since the source of uncertainty in our model stems from an unknown discrete probability distribution, our model can also be interpreted as an instance of distributionally robust optimization, which refers to optimization problems in which only a partial description of a probability distribution is available (Scarf 1958, Dupačová 1980, Birge and Wets 1987, Shapiro and Kleywegt 2002, Delage and Ye 2010, Goh and Sim 2010, Xu et al. 2012). Recently, Wiesemann et al. (2014) proposed a canonical framework for distributionally robust convex optimization, and they show that it subsumes well-known robust optimization models. A common approach to modeling distributional uncertainty sets is to impose constraints on the moments of the distribution (e.g., Ghaoui et al. 2003, Delage and Ye 2010). Since we instead model distributional uncertainty using polyhedral sets, from a technical perspective, our model is more aligned with the literature on robust linear optimization (e.g., Bertsimas and Sim 2004).

Because of our use of a CVaR objective function, our main formulation takes the form of a two-stage (min-max-min) robust optimization model, where the inner minimization problem is used to compute CVaR. Gabrel et al. (2014) and Zeng and Zhao (2013) also present two-stage models for facility location problems, where the focus is on uncertainty in the demand level rather than location. Other settings in which two-stage robust optimization models have been considered include network flows (Atamturk and Zhang 2007, Ordóñez and Zhao 2007), power systems (Bertsimas et al. 2013, Zhao and Zeng 2012), and military applications, including the defense of critical infrastructure (Brown et al. 2006, Alderson et al. 2011) and network interdiction (Brown et al. 2009).

Finally, our paper has some similarities with work by Carlsson and Delage (2013) in that they also consider a distributionally robust optimization problem in a spatial setting where demand points arrive in a continuous service area. Assuming the mean and covariance of the demand distribution is known, they focus on a vehicle routing problem where the goal is to partition the service area into a set of subregions such that the worst-case load for any vehicle across all subregions is minimized. By comparison, we assume that a fixed set of subregions is given, and the only distributional information available is the probability of an arrival in each subregion.

3. Model

Let $\mathcal{A} \subset \mathbb{R}^2$ represent the continuous service area over which cardiac arrests arrive. Let \mathcal{J} be an index set for m candidate sites, and let $\mathbf{y} \in \{0, 1\}^m$ be a binary vector, where $y_i = 1$ indicates the presence of an AED in the i th candidate site. Let the function $d: \mathcal{A} \times \mathcal{A} \rightarrow \mathbb{R}_+$ measure the distance between two locations in \mathcal{A} . With a slight abuse of notation, let $\mathbf{y}(a) \in \mathcal{A}$ represent the nearest facility to a point $a \in \mathcal{A}$, so that $d(a, \mathbf{y}(a))$ is the distance between a point $a \in \mathcal{A}$ and its nearest facility.

Let ξ be a random vector representing the location of the next cardiac arrest event within \mathcal{A} , with distribution μ . The distribution μ can be interpreted as describing the risk of a cardiac arrest event occurring over \mathcal{A} . Without loss of generality, we assume $P_\mu(\xi \in \mathcal{A}) = 1$. For a set of AED locations \mathbf{y} , $d(\xi, \mathbf{y}(\xi))$ represents the distance between the next demand arrival and its nearest facility. Since $d(\xi, \mathbf{y}(\xi))$ is itself random because of ξ , we shall refer to the distribution of $d(\xi, \mathbf{y}(\xi))$ as the *distance distribution* induced by \mathbf{y} . We can think of the distribution of $d(\xi, \mathbf{y}(\xi))$ as fully describing the performance of the AED deployment.

As discussed in Section 1, we wish to place AEDs in a manner that mitigates the risk of large distances between cardiac arrest victims and the nearest AED. Restated, our goal is to identify an AED deployment \mathbf{y} that controls the right tail of the distance distribution,

$d(\xi, \mathbf{y}(\xi))$. Since placing an AED in every possible candidate site is prohibitively expensive, we assume throughout that up to P AEDs are available for deployment. If the distribution μ is known, then the problem of minimizing the CVaR of the distance distribution is given by

$$\underset{\mathbf{y} \in \mathbf{Y}}{\text{minimize}} \text{CVaR}_{\mu}[d(\xi, \mathbf{y}(\xi))], \quad (1)$$

where

$$\mathbf{Y} = \left\{ \mathbf{y} \in \{0, 1\}^m \mid \sum_{i \in \mathcal{J}} y_i = P \right\}.$$

CVaR is closely related to the value-at-risk (VaR) measure, both of which have origins in the finance literature (see Rockafellar and Uryasev 2000, 2002). For a given loss distribution, β -VaR represents the smallest value α such that the loss does not exceed α with probability $1 - \beta$. By contrast, β -CVaR represents the expected loss conditional on the loss exceeding β -VaR.

In practice, the distribution of cardiac arrest locations is unlikely to be known precisely. Since—for a given location—cardiac arrests are low probability events, optimizing AED locations solely with respect to limited historical data (effectively a sample average approximation) may result in a poor deployment of AEDs. Moreover, any estimate of μ from historical data is also likely to be subject to uncertainty, especially if a limited amount of data are available.

To account for this uncertainty, we assume instead that the distribution μ is only known to belong to an uncertainty set, \mathcal{U} . A key question at this juncture is how to structure the set \mathcal{U} . To obtain a formulation that is both tractable and effective at capturing uncertainty in μ , we take the following approach. Suppose the service area \mathcal{A} can be divided into (possibly overlapping) *uncertainty regions* $\mathcal{A}_1, \mathcal{A}_2, \dots, \mathcal{A}_{|\mathcal{J}|}$, where $\mathcal{A} = \mathcal{A}_1 \cup \mathcal{A}_2 \cup \dots \cup \mathcal{A}_{|\mathcal{J}|}$. Suppose now that the only available information about the distribution μ is the probability of the next cardiac arrest occurring in each of the regions $\mathcal{A}_1, \dots, \mathcal{A}_{|\mathcal{J}|}$, which we denote by $\lambda_1, \dots, \lambda_{|\mathcal{J}|}$. The worst-case CVaR can now be minimized by solving the problem

$$\min_{\mathbf{y} \in \mathbf{Y}} \max_{\mu \in \mathcal{U}} \text{CVaR}_{\mu}[d(\xi, \mathbf{y}(\xi))], \quad (2)$$

where

$$\mathcal{U} = \{ \mu \mid \mathbb{P}_{\mu}(\xi \in \mathcal{A}_1) = \lambda_1, \dots, \mathbb{P}_{\mu}(\xi \in \mathcal{A}_{|\mathcal{J}|}) = \lambda_{|\mathcal{J}|} \}$$

is the distributional uncertainty set. Note that since the uncertainty regions are permitted to overlap, the probabilities $\lambda_1, \dots, \lambda_{|\mathcal{J}|}$ may sum to a value greater than 1. We assume throughout that the uncertainty regions are given. In practice, the design of the uncertainty regions \mathcal{A}_j should be guided by several considerations such as tractability, availability of historical data, and

model interpretability (we discuss these considerations in greater detail in Section 5).

There are several reasons behind our choice of the uncertainty set \mathcal{U} . First, we posit that obtaining estimates of the cardiac arrest risk at an aggregate (i.e., uncertainty region) level is less onerous than estimating the distribution μ directly, especially if data on historical cardiac arrest locations are limited. Estimating the parameters $\lambda_1, \dots, \lambda_{|\mathcal{J}|}$ in lieu of μ itself allows us to partially capture the underlying distribution while accounting for uncertainty. This approach is well aligned with the notion that aggregate forecasts tend to be more reliable than point forecasts when dealing with uncertainty (Simchi-Levi et al. 2004, Sheffi 2005, Bertsimas and Thiele 2006, Nahmias and Cheng 2009). In addition, our choice for the structure of \mathcal{U} lends itself to a tractable mathematical programming formulation of (2). The key to this tractability is the observation that for any \mathbf{y} , the worst-case distribution in \mathcal{U} is supported on a relatively small number of locations in \mathcal{A} . This sparsity can, in turn, be exploited through a decomposition algorithm, which allows us to solve the robust optimization problem efficiently (see Section 4). Our choice of \mathcal{U} also provides a natural generalization of classical facility location models, which have a rich and extensive history. In particular, classical models typically assume demand is concentrated at n weighted demand points (e.g., a weight of λ_j for demand point $j = 1, \dots, n$), which is recovered from our model by simply shrinking each \mathcal{A}_j to a singleton. Thus, a natural interpretation of our distributionally robust model is as an extension of classical models where each demand point lives in a region \mathcal{A}_j , wherein the true demand location will be realized after facility siting. We discuss the connection between our model and the classical location literature in more detail in Section EC.2 of the electronic companion. In particular, we show that our model unifies the classical p -median and p -center problems, as well as their robust analogues.

We note here that our choice of uncertainty set bears similarities to the uncertainty set discussed in the paper by Wiesemann et al. (2014), who present a canonical modeling approach for distributionally robust optimization problems. The authors formulate the uncertainty set using constraints on the probability that a random vector is realized within various convex sets. They show that within their framework, simply checking whether the distributional uncertainty set is empty can be a strongly NP-hard problem, unless the constituent convex sets satisfy a particular nesting condition. This hardness follows from the generality of the uncertainty set, which can be shown to subsume integer programming, in the absence of the nesting condition. Our setting is considerably more structured. By discretizing the service area into a large number

of scenarios, we obtain a representation of \mathcal{U} that is polyhedral. As a consequence, the worst-case distribution can be identified by solving a linear program.

In the remainder of this section, we show how our discretization approach allows us to formulate the optimization problems (1) and (2) as tractable mathematical programs. We then propose a bound on the error introduced by the discretization.

3.1. Known Cardiac Arrest Distribution

For ease of exposition, we first consider the case where μ is known. A standard approach for formulating stochastic programs is to discretize the outcome space into a finite set of *scenarios* and to assign a probability to each scenario. Since in our setting ξ is the random location of the next cardiac arrest event, each scenario (i.e., realization of ξ) has a one-to-one correspondence with a particular geographic location in \mathcal{A} . Accordingly, let $\Xi \subset \mathcal{A}$ be a discrete set of locations that approximate the continuous service area \mathcal{A} , so that Ξ is a discrete approximation of the outcome space of ξ , and each element of Ξ corresponds to a distinct location in \mathcal{A} . Letting \mathcal{K} be an index set for the scenarios in Ξ , we can write $\Xi = \{\xi_1, \xi_2, \dots, \xi_{|\mathcal{K}|}\}$ to represent the set of all possible realizations of ξ . Let u_k be the probability that the next cardiac arrest will occur at location ξ_k . Since in this discrete setting the possible cardiac arrest locations are restricted to Ξ , we have $\sum_{k \in \mathcal{K}} u_k = 1$.

Let the parameter d_{ik} represent the distance between candidate site i and scenario ξ_k based on the distance function $d(\cdot, \cdot)$. Let z_{ik} be an assignment variable that is equal to 1 if a cardiac arrest realized at location ξ_k is assigned to an AED at location i . The feasible set of assignments is given by

$$\mathbf{Z}(\mathbf{y}) = \left\{ z_{ik} \geq 0, i \in \mathcal{I}, k \in \mathcal{K} \mid \sum_{i \in \mathcal{I}} z_{ik} = 1, k \in \mathcal{K}; \right. \\ \left. z_{ik} \leq y_i, i \in \mathcal{I}, k \in \mathcal{K} \right\}.$$

Note that on the basis of the discretization, we have $d(\xi_k, \mathbf{y}(\xi_k)) = \min_{z \in \mathbf{Z}(\mathbf{y})} \sum_{i \in \mathcal{I}} d_{ik} z_{ik}$. Hence, for fixed \mathbf{y} , $\text{CVaR}_\mu[d(\xi, \mathbf{y}(\xi))]$ can be expressed as (Rockafellar and Uryasev 2000, 2002)

$$\text{CVaR}_\mu[d(\xi, \mathbf{y}(\xi))] \approx \min_{z, \alpha} \left\{ \alpha + \frac{1}{1 - \beta} \right. \\ \left. \cdot \sum_{k \in \mathcal{K}} u_k \max \left\{ \sum_{i \in \mathcal{I}} d_{ik} z_{ik} - \alpha, 0 \right\} \right\}. \quad (3)$$

Strictly speaking, the quantity in (3) is known as β - CVaR^- , which is an approximation of β - CVaR . This is a commonly used approximation for the case of discrete distributions; see Rockafellar and Uryasev (2000) for further details. The nominal optimization problem (1) for a known distribution can now be written as

$$\underset{\mathbf{y}, z, \alpha}{\text{minimize}} \left\{ \alpha + \frac{1}{1 - \beta} \right. \\ \left. \cdot \sum_{k \in \mathcal{K}} u_k \max \left\{ \sum_{i \in \mathcal{I}} d_{ik} z_{ik} - \alpha, 0 \right\} \right\} \quad (\text{N-AED})$$

$$\text{subject to } \mathbf{z} \in \mathbf{Z}(\mathbf{y}), \quad \mathbf{y} \in \mathbf{Y}, \\ \alpha \geq 0,$$

which is a mixed integer linear program after linearizing the maximization term in the objective.

3.2. Distributionally Robust Model for Unknown Cardiac Arrest Distribution

We now extend formulation N-AED to accommodate uncertainty in the distribution μ . Given the uncertainty regions $\mathcal{A}_1, \dots, \mathcal{A}_{|\mathcal{J}|}$, we construct scenario sets $\Xi_1, \Xi_2, \dots, \Xi_{|\mathcal{J}|}$, where

$$\Xi_j = \{\xi_k, k \in \mathcal{K} \mid \xi_k \in \mathcal{A}_j\}$$

and $\bigcup_{j=1}^n \Xi_j = \Xi$. Let \mathcal{K}_j index the scenarios in Ξ_j , and let \mathcal{J} index the set of uncertainty regions. Since in the discretized setting we assume a cardiac arrest arriving in \mathcal{A}_j can only be realized at one of the locations in Ξ_j , we have $P_\mu(\xi \in \Xi_j) = P_\mu(\xi \in \mathcal{A}_j) = \lambda_j$. The worst-case CVaR of the distance distribution given a feasible \mathbf{y}, \mathbf{z} is then given by optimal value of the following max-min problem:

$$\max_{\mathbf{u}} \min_{\alpha} \left\{ \alpha + \frac{1}{1 - \beta} \sum_{k \in \mathcal{K}} u_k \max \left\{ \sum_{i \in \mathcal{I}} d_{ik} z_{ik} - \alpha, 0 \right\} \right\} \\ \text{subject to } \sum_{k \in \mathcal{K}} u_k = 1, \\ \sum_{k \in \mathcal{K}_j} u_k = \lambda_j, \quad j \in \mathcal{J}, \\ u_k \geq 0, \quad k \in \mathcal{K}, \\ \alpha \geq 0. \quad (4)$$

For conciseness, define

$$\mathbf{U} = \left\{ \mathbf{u} \in \mathbb{R}_+^{|\mathcal{K}|} \mid \sum_{k \in \mathcal{K}} u_k = 1, \sum_{k \in \mathcal{K}_j} u_k = \lambda_j, j \in \mathcal{J} \right\}.$$

Note that the set \mathbf{U} may be empty if the parameters $\lambda_1, \dots, \lambda_{|\mathcal{J}|}$ are selected arbitrarily (e.g., consider two regions $\mathcal{A}_1, \mathcal{A}_2$ where $\mathcal{A}_1 \cap \mathcal{A}_2 = \emptyset$ but $\lambda_1 = \lambda_2 = 1$). However, if λ is estimated appropriately from historical cardiac arrest data, then nonemptiness of \mathbf{U} is guaranteed (see Proposition 2). We discuss the design of uncertainty regions and the estimation of λ further in Section 5.2. Note also that our discretization scheme allows us to model generally shaped uncertainty regions. This flexibility allows physical obstacles to be incorporated into the uncertainty regions, such as certain private buildings where public cardiac arrests cannot occur.

Minimizing the worst-case CVaR over \mathbf{y}, \mathbf{z} , we now arrive at the following two-stage robust optimization problem:

$$\min_{\mathbf{y}, z} \max_{\mathbf{u}} \min_{\alpha} \left\{ \alpha + \frac{1}{1 - \beta} \sum_{k \in \mathcal{K}} u_k \max \left\{ \sum_{i \in \mathcal{I}} d_{ik} z_{ik} - \alpha, 0 \right\} \right\} \\ \text{subject to } \mathbf{u} \in \mathbf{U}, \mathbf{z} \in \mathbf{Z}(\mathbf{y}), \mathbf{y} \in \mathbf{Y}, \\ \alpha \geq 0. \quad (\text{R-AED})$$

With respect to the size of R-AED, in practice we expect \mathcal{K} to be large (on the order of 10^3 or larger) so that our discretization procedure reasonably approximates the underlying continuous problem. For example, in our case study (see Sections 5 and 6), we take $|\mathcal{K}| = 2,600$. Similarly, the number of candidate sites $|\mathcal{I}|$ in our case study is 120, and the number of uncertainty regions $|\mathcal{J}|$ is 15.

It is worth discussing how formulation R-AED compares to other robust location models with demand uncertainty. The model closest to ours is the one proposed by Baron et al. (2011), who consider demand-level uncertainty at each of a fixed set of demand locations. Our model can also be viewed as modeling demand-level uncertainty, although there are some key differences with Baron et al. (2011). In formulation R-AED, the discrete uncertainty set \mathbf{U} forces the aggregate demand from the locations \mathcal{K}_j to be λ_j and normalizes the total demand over all locations in \mathcal{K} to be 1. This is in contrast to the approach of Baron et al. (2011), who model the uncertainty by allowing the vector of demands over all locations to reside within a box or ellipsoidal uncertainty set. This difference in how demand uncertainty is modeled turns out to be consequential. Based on our choice of the (polyhedral) uncertainty set, the worst-case distribution in \mathbf{U} will generally be supported on a relatively small number of locations, compared with the total number of scenarios. This means that, in the worst-case distribution, the demand u_k in the vast majority of locations in \mathcal{K} will be 0. Moreover, increasing the size of the scenario \mathcal{K} has no effect on the number of locations that support the worst-case distribution in \mathbf{U} . This sparsity and independence from the total number of scenarios in \mathcal{K} is exploited by our row-and-column generation algorithm to obtain solutions to R-AED (see Section 4). By contrast, in Baron et al. (2011), the uncertain demand at each location typically has a nonzero lower bound, which leads to all locations having nonzero demand, and thus precludes the application of our decomposition technique.

In the most general case, formulation R-AED can be interpreted as a facility location model with a single demand point (i.e., the location of the “next” cardiac arrest) and stochastic edge lengths, where d_{ik} gives the length of the edge between the demand point and candidate site i under scenario k . Under this interpretation, λ_j represents the probability that one of the scenarios in the set \mathcal{K}_j occurs. Our model differs from the existing literature on stochastic edge lengths, since we assume that the vector of probabilities \mathbf{u} is itself uncertain and only known to reside within a polyhedral uncertainty set.

Formulation R-AED can also be shown to unify the well-known p -median and p -center location problems,

as well as their robust analogues under demand location uncertainty. We formalize this connection in Section EC.2 of the electronic companion.

3.3. Discretization Error Bounds

The model R-AED is an approximation in the sense that it restricts the continuous demand distribution to be supported on a discrete set. As a result, the “true” worst-case CVaR in the continuous setting may be underestimated by the optimal value of formulation R-AED. In this section, we present bounds on this underestimation, which we refer to as the *discretization error*. The discretization error can be interpreted as the optimality gap associated with an optimal solution of formulation R-AED with respect to the underlying continuous problem (2).

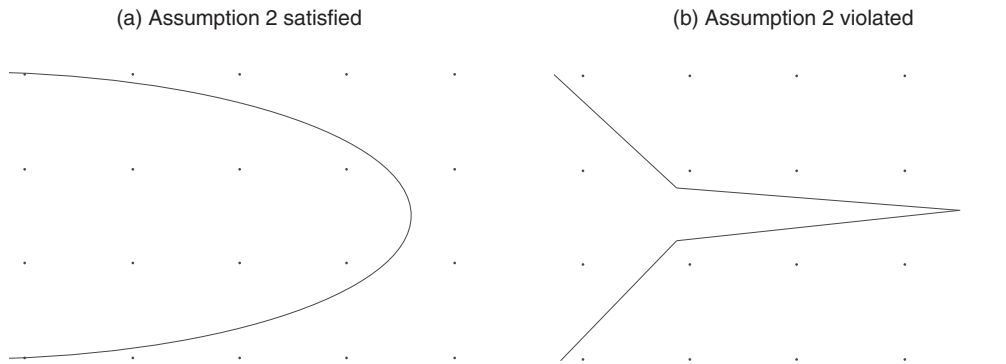
To define and prove the discretization error bounds, we first construct a partition of the service area \mathcal{A} . The partitioning is required for technical purposes only and can be applied to any general setting with overlapping uncertainty regions. Note that any configuration of the uncertainty regions \mathcal{A}_j implies a partitioning of the full service area \mathcal{A} into potentially more than n “subregions.” For example, suppose $\mathcal{A} = \mathcal{A}_1 \cup \mathcal{A}_2$, where $\mathcal{A}_1 \cap \mathcal{A}_2 \neq \emptyset$. Then we consider three disjoint subregions: $\mathcal{A}'_1 = \mathcal{A}_1 \setminus \mathcal{A}_2$, $\mathcal{A}'_2 = \mathcal{A}_2 \setminus \mathcal{A}_1$, and $\mathcal{A}'_3 = \mathcal{A}_1 \cap \mathcal{A}_2$. In general, let $\mathcal{A}'_1, \dots, \mathcal{A}'_{|\mathcal{J}|}$ represent these subregions, and let $\Xi'_1, \dots, \Xi'_{|\mathcal{J}|}$ be the discrete counterparts to the \mathcal{A}'_r . Let \mathcal{R} be the index set for the subregions, and let \mathcal{K}_r index the scenarios in subregion r . Since by definition the \mathcal{A}'_r form a partition of \mathcal{A} , it follows that $\Xi'_r \cap \Xi'_{r'} = \emptyset$ for all $r \neq r'$. Note that in the most general case, we may have $|\mathcal{R}| = 2^{|\mathcal{J}|}$, if every uncertainty region partially overlaps with every other uncertainty region. However, this extreme case is unlikely to arise naturally in the context of AED deployment and other facility location problems. Note also that if there is no overlap among the \mathcal{A}_j , then $|\mathcal{R}| = |\mathcal{J}|$. To construct the bound, we first require the following assumptions.

Assumption 1. Assume that $\Xi'_r = \mathcal{A}'_r \cap \mathcal{L}$ for all $r \in \mathcal{R}$, where \mathcal{L} is a square lattice in the plane with a grid spacing of length σ .

Assumption 2. The shape of the subregion \mathcal{A}'_r is such that for any point $a \in \mathcal{A}'_r \setminus \Xi'_r$, at least one of the four points in \mathcal{L} that define the smallest square containing a is inside \mathcal{A}'_r .

Assumption 1 states that the uncertainty regions are discretized using a square lattice, which simplifies the analysis considerably without loss of generality. Assumption 2 is meant to exclude pathologically shaped uncertainty regions relative to the size of the discretization (see Figure 1), and it can be satisfied with a sufficiently granular discretization of the uncertainty region. We now define a parameter $l(\sigma)$ with the

Figure 1. Examples for Assumption 2



following property:

$$l(\sigma) \geq \max_{a \in \mathcal{A}'_r} \min_{\xi \in \Xi'_r} d(a, \xi), \quad \text{for all } r \in \mathcal{R}. \quad (5)$$

Here, $l(\sigma)$ is an upper bound on the maximum distance between any point in \mathcal{A}'_r and its closest scenario location in Ξ'_r . To make the expression in (5) concrete, we provide some examples of the parameter $l(\sigma)$ in Remark 1, depending on the choice of the distance metric.

Remark 1. Let Assumptions 1 and 2 hold. Let $\|\cdot\|_1$, $\|\cdot\|_2$, and $\|\cdot\|_\infty$ be the Euclidean, rectilinear, and maximum metrics, respectively, on \mathbb{R}^2 :

- (1) If $d(a, b) = \|a - b\|_1$, then $l(\sigma) = 2\sigma$.
- (2) If $d(a, b) = \|a - b\|_2$, then $l(\sigma) = \sqrt{2}\sigma$.
- (3) If $d(a, b) = \|a - b\|_\infty$, then $l(\sigma) = \sigma$.

Note that setting $l(\sigma) = 2\sigma$ is also valid for any distance function based on a p -norm with $p \geq 1$, since $\|\cdot\|_p$ is nonincreasing in p . Let Z_D be the optimal value of formulation R-AED, let Z_C be the optimal value of the associated continuous problem (2), and let

$$\Delta := \frac{Z_C - Z_D}{Z_C}$$

be the discretization error. We now present the main bound.

Theorem 1. For any $\varepsilon > 0$, if $l(\sigma) \leq \varepsilon(1 - \beta)Z_D$, then $\Delta \leq \varepsilon$.

Theorem 1 characterizes how finely the uncertainty regions must be discretized to achieve a certain discretization error. Since $l(\sigma)$ vanishes with σ , the discretization error can be made arbitrarily small by selecting an appropriate value for σ . Thus, we can obtain near-optimal solutions to the underlying continuous uncertainty problem by solving formulation R-AED with a sufficiently fine discretization (and thus a sufficiently large number of scenarios). We can also rearrange the expression in Theorem 1 to calculate an upper bound on Z_C as a function of σ and Z_D . The

upper bound \bar{Z}_C is constructed as follows:

$$Z_C \leq \bar{Z}_C = Z_D(1 + \varepsilon), \quad \text{where } \varepsilon = \frac{l(\sigma)}{Z_D(1 - \beta)}. \quad (6)$$

In the next section, we outline a decomposition technique that enables us to compute a tight bound \bar{Z}_C by solving instances of R-AED with an extremely large number of scenarios.

4. Row-and-Column Generation Algorithm

It is well known that two-stage robust optimization problems with the structure of (8) can be solved using Benders-based decomposition algorithms (see, e.g., Brown et al. 2006, 2009; Alderson et al. 2011). These approaches decompose the formulation into a master problem, which is a relaxation of the original problem, and a subproblem, which generates constraints to be added to the master problem. The master and subproblems are solved iteratively until convergence to a provably optimal solution. However, when applied to (8), the number of constraints and decision variables of the (integer) master problem using these approaches depends on the number of scenarios $|\mathcal{K}|$. As a consequence, these decomposition approaches can fail to be tractable if $|\mathcal{K}|$ is even moderately sized. We overcome this issue by proposing an algorithm that decouples the size of the master problem from the number of scenarios. A key consequence of this decoupling is that the tractability of the resulting formulation no longer depends on the granularity of the discretization, which, when combined with the bounds in Section 3.3, allows us to approximate the continuous problem using an extremely large number of scenarios. A similar row-and-column generation algorithm for solving min-max optimization problems was independently proposed by Zeng and Zhao (2013). A key difference is that unlike the model in Zeng and Zhao (2013), which has an integer subproblem, our subproblem is a linear program.

We compare our proposed method with two other solution approaches that are natural candidates for

solving formulation **R-AED**—a Benders-based row generation algorithm and an MIP reformulation of **R-AED**.

4.1. Algorithm Overview

We now provide an overview of the row-and-column generation algorithm. First, observe that the inner minimization problem of (4) is a linear program. Letting \mathbf{p} denote the dual vector for this linear program, the dual polyhedron—which depends on the distribution \mathbf{u} —can be written as

$$\mathbf{P}(\mathbf{u}) = \left\{ p_k \geq 0, k \in \mathcal{K} \mid \sum_{k \in \mathcal{K}} p_k \leq 1, p_k \leq \frac{u_k}{1-\beta}, k \in \mathcal{K} \right\}. \quad (7)$$

The optimization problem (4) can now be reformulated as

$$\begin{aligned} \min_{\mathbf{y}, \mathbf{z}} \max_{\mathbf{u}, \mathbf{p}} \quad & \sum_{i \in \mathcal{J}} \sum_{k \in \mathcal{K}} d_{ik} z_{ik} p_k \\ \text{subject to} \quad & \mathbf{p} \in \mathbf{P}(\mathbf{u}), \quad \mathbf{u} \in \mathbf{U}, \\ & \mathbf{z} \in \mathbf{Z}(\mathbf{y}), \quad \mathbf{y} \in \mathbf{Y}. \end{aligned} \quad (8)$$

By introducing the auxiliary variable t , we obtain the following epigraph representation of (8):

$$\text{minimize}_{\mathbf{y}, \mathbf{z}, t} \quad t \quad (9a)$$

$$\text{subject to} \quad t \geq \sum_{i \in \mathcal{J}} \sum_{k \in \mathcal{K}} d_{ik} z_{ik} p_k, \quad \mathbf{p} \in \mathbf{P}(\mathbf{u}), \quad \mathbf{u} \in \mathbf{U}, \quad (9b)$$

$$\mathbf{z} \in \mathbf{Z}(\mathbf{y}), \quad \mathbf{y} \in \mathbf{Y}. \quad (9c)$$

Note that formulation (9a) is a semi-infinite optimization problem, since the constraint (9b) is enforced over all possible \mathbf{p} . To obtain a feasible solution to (9a), we may instead solve a relaxation where the constraint (9b) is enforced over a finite subset of \mathbf{P} . Let this subset be $\mathbf{P}_{|\mathcal{S}|} = \{\mathbf{p}^1, \dots, \mathbf{p}^{|\mathcal{S}|}\}$, and let \mathcal{S} be an index set for the vectors in $\mathbf{P}_{|\mathcal{S}|}$. To obtain a relaxation of (9b) with a finite number of constraints, we replace the constraint set represented by (9b) with the following set of $|\mathcal{S}|$ constraints:

$$t \geq \sum_{i \in \mathcal{J}} \sum_{k \in \mathcal{K}} d_{ik} z_{ik} p_k^s, \quad s \in \mathcal{S}. \quad (10)$$

Although replacing the constraint (9b) with (10) in formulation (9a) produces a mixed integer linear program, the resulting formulation may still be intractably large, as a result of the dependence of the number of constraints and variables on the size of the scenario set, \mathcal{K} . Consider now the set

$$\mathcal{K}_+ = \left\{ k \in \mathcal{K} \mid \sum_{s \in \mathcal{S}} p_k^s > 0 \right\}, \quad (11)$$

and note that constraint (10) can be written as

$$t \geq \sum_{i \in \mathcal{J}} \sum_{k \in \mathcal{K}_+} d_{ik} z_{ik} p_k^s + \sum_{i \in \mathcal{J}} \sum_{k \in \mathcal{K} \setminus \mathcal{K}_+} d_{ik} z_{ik} p_k^s, \quad s \in \mathcal{S}. \quad (12)$$

Since by definition we have $\sum_{k \in \mathcal{K} \setminus \mathcal{K}_+} p_k^s = 0$ for all $s \in \mathcal{S}$, we may drop the second term on the right-hand side of (10), which further simplifies the constraint to

$$t \geq \sum_{i \in \mathcal{J}} \sum_{k \in \mathcal{K}_+} d_{ik} z_{ik} p_k^s, \quad s \in \mathcal{S}. \quad (13)$$

Observe that if $k \notin \mathcal{K}_+$, then the value of the corresponding assignment variable z_{ik} has no impact on the objective function, since its coefficient is 0 for all $s \in \mathcal{S}$. We can therefore remove all z_{ik} variables and the constraints $z_{ik} \leq y_i$ and $z_{ik} \geq 0$ for $k \in \mathcal{K} \setminus \mathcal{K}_+$. This leads to the following relaxation of (9a):

$$\text{minimize}_{\mathbf{y}, \mathbf{z}, t} \quad t$$

$$\text{subject to} \quad t \geq \sum_{i \in \mathcal{J}} \sum_{k \in \mathcal{K}_+} d_{ik} z_{ik} p_k^s, \quad s \in \mathcal{S},$$

$$\sum_{i \in \mathcal{J}} y_i = P,$$

$$\sum_{i \in \mathcal{J}} z_{ik} = 1, \quad k \in \mathcal{K}_+, \quad (\text{R-AED-MP})$$

$$z_{ik} \leq y_i, \quad i \in \mathcal{J}, k \in \mathcal{K}_+,$$

$$z_{ik} \geq 0, \quad i \in \mathcal{J}, k \in \mathcal{K}_+,$$

$$y_i \in \{0, 1\}, \quad i \in \mathcal{J}.$$

The relaxed problem **R-AED-MP** yields a lower bound on the optimal value of the original problem (8). To tighten the relaxation, we can add a new vector \mathbf{p} to $\mathbf{P}_{|\mathcal{S}|}$ and re-solve **R-AED-MP**, with the aim of obtaining an improved lower bound. Let $\bar{\mathbf{y}}$ be a deployment obtained at a solution of **R-AED-MP**. We then construct a feasible assignment vector $\bar{\mathbf{z}}$ that assigns each scenario to its nearest AED as follows. Let $i^*(k) = \arg \min_{\{i \in \mathcal{J} \mid \bar{y}_i = 1\}} \{d_{ik}\}$, which represents the closest AED to location k , given the deployment $\bar{\mathbf{y}}$. Then for each $k \in \mathcal{K}$, set

$$\bar{z}_{i^*(k), k} = 1 \quad \text{and} \quad \bar{z}_{ik} = 0 \quad \text{for all } i \neq i^*(k). \quad (14)$$

It is straightforward to show that the $\bar{\mathbf{z}}$ given by (14) is also optimal for **R-AED-MP**. To identify a new \mathbf{p} , we then solve the following subproblem, where $\bar{\mathbf{y}}, \bar{\mathbf{t}}$ is obtained as the incumbent solution to **R-AED-MP** and $\bar{\mathbf{z}}$ is given by (14):

$$\text{maximize}_{\mathbf{u}, \mathbf{p}} \quad \sum_{i \in \mathcal{J}} \sum_{k \in \mathcal{K}} d_{ik} \bar{z}_{ik} p_k \quad (\text{R-AED-SP})$$

$$\text{subject to} \quad \mathbf{p} \in \mathbf{P}(\mathbf{u}), \quad \mathbf{u} \in \mathbf{U}.$$

To understand why the construction of $\bar{\mathbf{z}}$ using (14) is necessary, note that since **R-AED-MP** minimizes CVaR, the optimal assignment variables \mathbf{z} obtained at a solution to **R-AED-MP** may not assign every location k to its nearest AED, since only the locations in the tail of the distance distribution affect the optimal objective function value. As a consequence, solving **R-AED-SP** using the $\bar{\mathbf{z}}$ obtained directly from **R-AED-MP** may produce

a solution to R-AED-SP that does not correspond to a worst-case distribution (with respect to the incumbent deployment \mathbf{y}). While in theory the algorithm converges in finite time without reconstructing $\bar{\mathbf{z}}$ according to (14), in practice we found that including this additional assignment step results in fewer iterations and a significant speed improvement.

Observe that formulation R-AED-SP is a linear program. Now let $\mathbf{u}^*, \mathbf{p}^*$ be an optimal solution to R-AED-SP, and let Z_{SP} be the optimal value. If $\bar{t} \geq Z_{SP}$, then it follows that $\bar{t} \geq \sum_{i \in \mathcal{J}} \sum_{k \in \mathcal{K}} d_{ik} \bar{z}_{ik} p_k^*$ for all $\mathbf{p} \in \mathbf{P}(\mathbf{u})$, $\mathbf{u} \in \mathbf{U}$, and thus the incumbent solution is certifiably optimal to the original problem. If $\bar{t} < Z_{SP}$, then we introduce the variables z_{ik} , $k \in \mathcal{K}^*$ to the master problem, as well as the constraints $z_{ik} \geq 0$, $z_{ik} \leq y_i$, $k \in \mathcal{K}^*$, and $t \geq \sum_{i \in \mathcal{J}} \sum_{k \in \mathcal{K}} d_{ik} z_{ik} p_k^*$, where $\mathcal{K}^* = \{k \in \mathcal{K} \mid p_k^* > 0\}$. We then solve R-AED-MP again to obtain a new solution $\bar{\mathbf{y}}, \bar{\mathbf{z}}, t$, which begins a new iteration of the algorithm. Since there are a finite (but possibly large) number of rows and columns that can be generated, convergence after a finite number of iterations is guaranteed. The algorithm is summarized in Algorithm 1.

Algorithm 1 (Row-and-Column Generation)

Initialize set $\mathcal{S} = \{0\}$, $s = 1$, and $\epsilon > 0$. Let $\mathcal{K}_+ = \emptyset$.

1. Solve R-AED-MP to obtain solution $(\bar{\mathbf{y}}, \bar{t})$ and objective value Z_{MP} .
2. Construct $\bar{\mathbf{z}}$ according to (14).
3. Solve R-AED-SP with fixed $\bar{\mathbf{z}}$ to obtain $\bar{\mathbf{p}}^s$ and objective value Z_{SP} .
4. **If** $(Z_{SP} - Z_{MP})/Z_{SP} \leq \epsilon$, terminate and return optimal deployment $\bar{\mathbf{y}}$ and worst-case CVaR \bar{t} .
else
 (i) Construct index set

$$\mathcal{K}^* = \{k \in \mathcal{K} \setminus \mathcal{K}_+ \mid \bar{p}_k^s > 0\}$$

and set $\mathcal{K}_+ \leftarrow \mathcal{K}_+ \cup \mathcal{K}^*$.

- (ii) Add variables z_{ik} , $i \in \mathcal{J}$, $k \in \mathcal{K}^*$ to master problem R-AED-MP.
- (iii) Add constraints

$$t \geq \sum_{i \in \mathcal{J}} \sum_{k \in \mathcal{K}_+} d_{ik} z_{ik} \bar{p}_k^s,$$

$$z_{ij}^k \leq y_i, \quad i \in \mathcal{J}, k \in \mathcal{K}^*,$$

$$\sum_{i \in \mathcal{J}} z_{ij}^k = 1, \quad k \in \mathcal{K}^*,$$

to master problem R-AED-MP.

- (iv) Set $\mathcal{S} \leftarrow \mathcal{S} \cup \{s\}$. Increment s and return to step 1.

The intuition behind the algorithm is as follows. At a given iteration, the master problem only includes the subset of scenarios that are assigned a nonzero probability mass in the worst-case distribution for at least one of the previous iterations. Based on the incumbent solution, the subproblem then identifies the worst-case

distribution using the entire scenario set. The scenarios supporting the worst-case distribution (obtained by solving the subproblem) are then added to the master problem on an as-needed basis, along with the relevant assignment variables and related constraints. Because of this decomposition, at any iteration of the algorithm the size of the master problem R-AED-MP depends only on the number of candidate sites, $|\mathcal{J}|$, and the size of the restricted scenario set, \mathcal{K}_+ , which may be substantially smaller than \mathcal{K} . In fact, because of the structure of the uncertainty set \mathbf{U} , the size of the restricted scenario set \mathcal{K}_+ can be shown to depend only on the configuration of the uncertainty regions and the number of iterations completed by the algorithm, and it is otherwise independent of the total number of scenarios in \mathcal{K} . This result is formalized in the following proposition.

Proposition 1. Assume that for any k and k' such that $k \neq k'$, we have $d_{ik} \neq d_{ik'}$ for all $i \in \mathcal{J}$. Then the number of constraints and variables in the master problem R-AED-MP in the s th iteration of Algorithm 1 is $O(|\mathcal{J}|(|\mathcal{J}| + 1)s)$.

The assumption that $k \neq k'$ implies $d_{ik} \neq d_{ik'}$ for all $i \in \mathcal{J}$ is mild, and it ensures that for any candidate site i , there are no exact ties for the closest location in \mathcal{K} . (Indeed, this assumption can be relaxed if we assume instead that the solution to R-AED-SP is always a basic feasible solution, which is the case if the simplex method is used to solve R-AED-SP.) It is worth noting that the size of the master problem R-AED-MP may still depend indirectly on \mathcal{K} if the total number of scenarios is small (i.e., if $|\mathcal{R}|s > |\mathcal{K}|$). Additionally, Proposition 1 only guarantees that the size of the master problem in any given iteration is independent of \mathcal{K} but does not guarantee that the total number of iterations required will be independent of \mathcal{K} . However, our numerical results indicate that increasing the number of scenarios generally has a minimal impact on the performance of the row-and-column generation algorithm, because the worst-case distribution is generally sparse with respect to \mathcal{K} . Note that unlike the master problem, the subproblem R-AED-SP still depends on the size of \mathcal{K} . However, the subproblem is a linear program with $|\mathcal{K}|$ decision variables, and we found that it can be solved efficiently for practically sized problems, even if the set \mathcal{K} is quite large.

For illustrative purposes, in Section EC.4 of the electronic companion we present a plot that depicts the total number of scenarios and the associated optimality gap for an example instance of the row-and-column generation algorithm.

4.2. Benchmark Solution Approaches

We compare the performance of the row-and-column generation algorithm to two benchmark approaches: a mixed integer linear programming reformulation of (2) and a standard row generation algorithm for min-max problems.

4.2.1. Mixed Integer Linear Programming Reformulation. By formulating the dual problem of the inner maximization problem in (8), we obtain the following mixed integer linear program:

$$\begin{aligned}
 & \text{minimize}_{\mathbf{y}, \mathbf{z}, \mathbf{w}, \eta, \alpha, \gamma} \left\{ \eta + \alpha + \sum_{j \in \mathcal{J}} \lambda_j w_j \right\} \\
 & \text{subject to } \alpha + \gamma_k \geq \sum_{i \in \mathcal{J}} \sum_{k \in \mathcal{K}} d_{ik} z_{ik}, \quad k \in \mathcal{K}, \\
 & \quad \eta - \frac{\gamma_k}{1 - \beta} + \sum_{\{j | k \in \mathcal{K}_j\}} w_j \geq 0, \quad k \in \mathcal{K}, \\
 & \quad \mathbf{z} \in \mathbf{Z}(\mathbf{y}), \quad \mathbf{y} \in \mathbf{Y}, \\
 & \quad \gamma_k \geq 0, \quad k \in \mathcal{K}, \\
 & \quad \alpha \geq 0.
 \end{aligned} \tag{15}$$

If \mathcal{K} is relatively small, then the MIP formulation (15) can be solved using standard integer programming solvers. However, this formulation can become computationally intractable if a large number of scenarios are used to discretize the uncertainty regions.

4.2.2. Row Generation. A well-known approach to solving min-max problems such as (2) is to use a Benders-based row generation algorithm (Brown et al. 2006, 2009; Alderson et al. 2011). The row generation algorithm is conceptually similar to our row-and-column generation method outlined in Section 4.1 in that it also decomposes the problem into a master and subproblem, which provide lower and upper bounds, respectively. However, a key difference with our approach is that the row generation algorithm carries the full set of columns (i.e., all scenarios in \mathcal{K}) in the master problem in each iteration. The number of variables and constraints in the master problem of the

row generation algorithm in the s th iteration is thus $O(|\mathcal{J}||\mathcal{K}|)$ instead of $O(|\mathcal{J}||\mathcal{K}|s)$. As a result, the row generation algorithm may have poor performance if $|\mathcal{K}|$ is large.

4.3. Comparison of Solution Approaches

To demonstrate the performance of the row-and-column generation algorithm, we generated and solved several random problem instances using rectangular uncertainty regions and a square lattice for the discretization. We use Euclidean distances for all parameters. We solve each of these instances using the three solution methods discussed above, and we focus on the impact that the number of scenarios has on the total solution time. All problems were implemented in MATLAB R2011a using YALMIP as the modeling language and CPLEX 12.1 with default parameter settings as the solver, on a 2.4 GHz quad-core CPU.

In Table 1, we report numerical results on how the solution times of the MIP formulation (“MIP”), row generation algorithm (“Row”) and row-and-column generation algorithm (“R + C”) are impacted by the number of scenarios. We report results only for the cases where $\beta = 0$ and $\beta = 0.9$, and we note that similar trends were observed for other values of β . We also compute the error bound ϵ and the associated upper bound on the continuous problem, \bar{Z}_C , using (6). Note that the error bounds for the $\beta = 0.9$ cases are larger than the $\beta = 0$ case, which is expected based on the bound expression given in (6).

Unlike the MIP and row generation algorithm, the performance of the row-and-column generation algorithm scales extremely well with the number of scenarios, and in many cases it outperforms the benchmark

Table 1. Computational Results Showing Effect of Number of Scenarios ($|\mathcal{K}|$) on Error Bound and Solution Times (in CPU Seconds) of MIP, Row Generation, and Row-and-Column Generation Algorithm

\mathcal{J}	\mathcal{I}	P	\mathcal{K}	σ	$\beta = 0$						$\beta = 0.9$					
					Z_D	$\epsilon(\%)$	\bar{Z}_C	MIP	Row	R + C	Z_D	$\epsilon(\%)$	\bar{Z}_C	MIP	Row	R + C
25	5	2	456	5.0	124.6	5.7	131.7	1	2	2	180.0	39.3	250.7	12	35	44
25	5	2	1,864	2.5	126.0	2.8	129.5	2	3	2	182.4	19.4	217.8	56	45	42
25	5	2	7,446	1.3	126.4	1.4	128.2	9	9	3	182.8	9.7	200.4	107	135	49
25	5	2	29,621	0.6	126.9	0.7	127.8	93	311	8	183.4	4.8	192.2	10,000 ^a	5,359	64
50	10	5	886	5.0	105.2	6.7	112.3	6	12	4	178.0	39.7	248.7	37	342	230
50	10	5	3,563	2.5	107.4	3.3	110.9	94	104	13	181.8	19.4	217.2	390	1,257	290
50	10	5	14,242	1.3	108.3	1.6	110.0	10,000 ^a	10,000 ^a	33	183.1	9.7	200.8	10,000 ^a	10,000 ^a	414
50	10	5	57,016	0.6	108.5	0.8	109.4	10,000 ^a	10,000 ^a	49	181.3	4.9	190.2	10,000 ^a	10,000 ^a	603
75	15	10	1,263	5.0	83.6	8.5	90.7	23	1,875	174	139.3	50.8	210.0	142	293	340
75	15	10	5,112	2.5	85.1	4.2	88.6	2,415	10,000 ^a	523	138.9	25.5	174.2	2,088	9,799	614
75	15	10	20,470	1.3	85.4	2.1	87.1	10,000 ^a	10,000 ^a	614	139.1	12.7	156.8	10,000 ^a	10,000 ^a	1,534
75	15	10	81,940	0.6	85.8	1.0	86.7	10,000 ^a	10,000 ^a	778	140.2	6.3	149.0	10,000 ^a	10,000 ^a	1,610
100	20	20	1,511	5.0	60.0	11.8	67.0	30	228	53	84.6	83.6	155.3	9	174	87
100	20	20	6,158	2.5	61.4	5.8	64.9	2,617	2,579	485	84.7	41.7	120.1	129	454	172
100	20	20	24,743	1.3	62.1	2.8	63.9	10,000 ^a	10,000 ^a	585	85.7	20.6	103.4	10,000 ^a	4,255	2,175
100	20	20	99,032	0.6	62.7	1.4	63.5	10,000 ^a	10,000 ^a	1,739	73.0	12.1	81.8	10,000 ^a	10,000 ^a	3,252

^aInstance did not solve in 10,000 CPU seconds.

approaches by one to two orders of magnitude with respect to the solution time. Note that the performance improvement is most pronounced at larger values of $|\mathcal{S}|$, as a result of the poor scalability of the MIP and row generation approaches. These results suggest that, when using the row-and-column generation algorithm, if a problem can be solved to optimality using a small number of scenarios, then the problem tends to remain tractable when using a fine discretization and an extremely large number of scenarios (e.g., just under 100,000 scenarios are used in the final row of Table 1).

5. AED Deployment Study: Setup

In this section, we discuss the setup of our numerical study on the placement of AEDs in public locations, including an overview of the data used, calibration of the uncertainty set, benchmark models, and a simulation-based validation method to evaluate the performance of the AED deployments produced by each model.

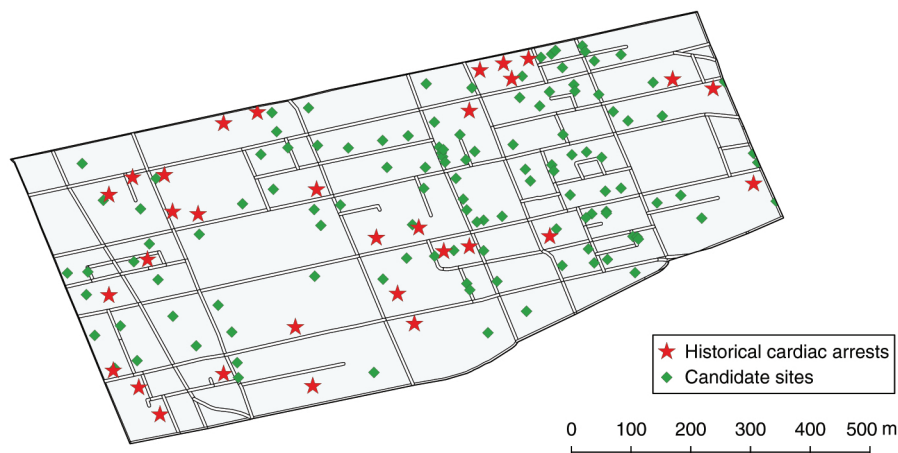
5.1. Data

5.1.1. Historical Cardiac Arrests. We consider the AED location problem in a densely populated region in the downtown core of Toronto, Canada, which roughly corresponds to the city’s financial district. We obtained a cardiac arrest data set through the Resuscitation Outcomes Consortium (ROC), which is a collaborative network of research institutions across the United States and Canada that collects cardiac arrest data for the purpose of improving clinical practice (Morrison et al. 2008). The ROC data set contains the geographic coordinates of all cardiac arrests in the City of Toronto from January 2006 to April 2013. During this period, a total of 43 cardiac arrests occurred in public locations within the study region. These cardiac arrest locations are used to estimate the λ_j parameters in formulation R-AED.

5.1.2. Candidate Sites. We constructed a set of candidate sites for AED placement by randomly selecting 120 public spaces (restaurants, cafés, shops, etc.) within the downtown region we consider. The addresses of these public spaces were obtained through the City of Toronto Employment Survey (City of Toronto 2010), which collects data on all Toronto businesses on an annual basis. We use publicly accessible buildings such as restaurants and cafés since they represent the type of location where one might reasonably expect an AED to be accessible to the public through a public access defibrillation program. Figure 2 provides a visualization of the study region, historical cardiac arrests, and the candidate sites for AED placement.

5.1.3. Distance Parameters. To calculate the distance parameters, we converted the addresses of the candidate sites into Universal Transverse Mercator (UTM) coordinates. Coordinates in the UTM system are given in units of meters (unlike latitude and longitude coordinates), which allows for straightforward calculation of distances within the study region. We constructed the scenario set \mathcal{S} by intersecting the study region with a square grid with a spacing of 20 meters, which resulted in approximately 2,600 scenarios. We generated the d_{ik} parameters by calculating the Euclidean distance between each candidate site and scenario pair. Although the Euclidean distance is only an approximation of the one-way distance a lay responder would have to travel to retrieve an AED, previous studies have shown Euclidean distance to be highly correlated with actual road distance, especially in urban areas where the road network is relatively dense (Bach 1981, Jones et al. 2010, Boscoe et al. 2012). Moreover, the use of Euclidean distances is particularly justified in the context of AED placement, since the typical distance between a cardiac arrest and the nearest AED is low enough that lay responders can often take shortcuts that do not necessarily appear on the road network

Figure 2. (Color online) Locations of Historical Cardiac Arrests and Candidate Sites for AED Placement



(e.g., through buildings, unmarked pedestrian paths, or parking lots).

As a robustness check, we repeat a subset of experiments using rectilinear distances and actual walking distances. The walking distances are obtained using the Google Maps Distance Matrix Application Programming Interface (API), implemented in Python. The Distance Matrix API is a query tool that accepts pairs of locations as input (as latitude-longitude coordinates or street addresses) and returns the distance between the two locations for a given mode of transportation (e.g., walking, driving) (Google 2016).

5.2. Uncertainty Regions and Estimation of λ_j

One of the key modeling decisions in our approach is to construct a set of uncertainty regions from the full service area. One simple approach is to partition the service area into a set of identically shaped cells and allow each cell to represent an uncertainty region. We propose that in the context of public access defibrillation, it may be most intuitive from a policy-making perspective to use preexisting geographic divisions within the city, such as city blocks, postal/zip codes, or census tracts. Furthermore, in settings where historical cardiac arrest data are unavailable, one may have to rely on alternate demographic information that are only recorded in aggregate as a proxy for cardiac arrest incidence. For example, the City of Toronto collects socioeconomic data at the neighbourhood level for planning purposes (City of Toronto 2016), which naturally implies a set of uncertainty regions.

Once a set of uncertainty regions is selected, the corresponding probabilities $\lambda_1, \dots, \lambda_{|\mathcal{J}|}$ must be estimated from historical cardiac arrest data. We propose the following approach, which produces consistent (i.e., asymptotically optimal) estimates of the uncertainty region probabilities.

Proposition 2. Let $a_1, \dots, a_n \in \mathcal{A}$ be the locations of n cardiac arrest events, and assume that a_1, \dots, a_n are drawn independent and identically distributed (i.i.d.) from the distribution μ . For each $j \in \mathcal{J}$, let

$$\hat{\lambda}_j^n = \frac{1}{n} \sum_{c=1}^n \mathbb{I}\{a_c \in \mathcal{A}_j\}. \quad (16)$$

Then $|\hat{\lambda}_j^n - \lambda_j| \rightarrow 0$ as $n \rightarrow \infty$ almost surely. Furthermore, if the sets $\Xi'_1, \dots, \Xi'_{|\mathcal{J}|}$ are all nonempty, then the uncertainty set \mathbf{U} constructed from $\hat{\lambda}_j^n$, $j \in \mathcal{J}$ is guaranteed to be nonempty.

The estimator in Proposition 2 is simply the fraction of cardiac arrest events that have occurred in each uncertainty region. Since the uncertainty regions are fixed, the convergence result in Proposition 2 follows directly from the strong law of large numbers. (To be precise, note that the arrival of cardiac arrests in each region can, with some modification, be modeled as a multinomial trials process, which has a

well-known maximum likelihood estimator; see Bickel and Doksum 2015.) The assumption that the cardiac arrest locations are drawn i.i.d. from a common distribution is a standard statistical assumption and can be interpreted as requiring the probability of a cardiac arrest event in different parts of the city to be stable over time. Fortunately, cardiac arrests in Toronto have been observed to exhibit this temporal stability (Chan et al. 2016).

Note also that these properties hold for *any* chosen configuration of the uncertainty regions. In other words, once the uncertainty regions are selected, the estimate $\hat{\lambda}^n$ is guaranteed to converge to the true arrival probabilities. This convergence is generally known as statistical *consistency*, which is considered a fundamental requirement for an estimator (Bickel and Doksum 2015).

In addition to estimating the arrival probabilities, we can also use the Wilson score method as a heuristic to estimate an associated 95% confidence interval for each arrival probability (Wilson 1927). These confidence intervals can also be used to guide the design of the uncertainty regions. For example, we might require that the uncertainty regions be selected such that estimation error of each of the arrival probabilities $\lambda_1, \dots, \lambda_{|\mathcal{J}|}$ is no greater than some threshold δ with 95% confidence. We emphasize, however, that the Wilson score method is only a heuristic for estimating the confidence intervals, meaning the estimation error is not always guaranteed to be less than δ with 95% confidence. As an alternative, we also note that the problem of designing uncertainty regions based on data is closely related to the well-studied problem in statistics of optimally selecting histogram bin widths (e.g., Wand 1997), although we do not formally explore that connection in this paper.

The service area in our numerical study spans only two census tracts and three postal code zones. Since using too few uncertainty regions may fail to capture important aspects of the cardiac arrest distribution, we instead identified uncertainty regions by using the road network to divide the study area into 15 approximately equal-area uncertainty regions. The estimates of the arrival probabilities for the 15 uncertainty regions obtained using this approach are reported in Section EC.5 of the electronic companion, along with their approximate 95% confidence intervals, calculated using the Wilson score method (Wilson 1927).

5.3. Benchmarks: Sample Average Approximation and Ex Post Models

We compare the performance of the AED deployment produced by the robust formulation R-AED with two benchmark models. The first benchmark is a *nominal* model, which optimizes only over historical cardiac arrest locations without taking uncertainty into account. The nominal model represents the sample

average approximation approach to the AED problem, since it optimizes directly with respect to a historical sample. The nominal model is given by formulation **N-AED**, where the scenarios $\{\xi_1, \dots, \xi_{|\mathcal{K}|}\}$ are given by the historical cardiac arrest locations, and $u_k = 1/|\mathcal{K}|$ for all $k \in \mathcal{K}$. The second benchmark is an ex post model, which optimizes directly over a set of simulated cardiac arrests and thus has perfect foresight. The ex post model is also given by formulation **N-AED**, except instead of optimizing over historical cardiac arrest locations, the ex post model optimizes the AED deployment with respect to a simulated set of cardiac arrest locations, against which we also evaluate the performance of the nominal and robust models. Thus, the ex post model can be interpreted as the “best possible” model we could solve if the locations of all future cardiac arrests were known a priori. In all instances, we set $\beta = 0.9$ and the number of AEDs (P) to 30. All instances for all models were implemented using MATLAB R2011a via YALMIP and solved using CPLEX 12.1 with default parameter settings on a single node of a computing cluster with a 2.9 GHz quad-core CPU.

5.4. Model Validation


To assess the performance of the AED deployments produced by **R-AED** and **N-AED**, we used simulation to generate hypothetical cardiac arrest events, and we computed various performance metrics associated with the distance distribution induced by each AED deployment and the set of simulated cardiac arrests. We then solve the ex post model based on the simulated cardiac arrest locations to gain insight into the best possible deployment for the given model parameterization and objective function.

To construct the validation sets, we first used kernel density estimation to estimate the underlying demand density, which is a well-known semiparametric technique for estimating a probability density function from a finite sample (Sheather and Jones 1991, Terrell and Scott 1992). Intuitively, kernel density involves centering a continuous density function (the kernel) at

each data point and then aggregating and normalizing all kernel functions to obtain a single probability density function. Kernel density estimation requires the specification of a kernel function as well as a parameter h , known as the *bandwidth*, which is typically proportional to the standard deviation of the kernel distribution. In short, the bandwidth represents the degree to which the data set is smoothed to obtain the estimated density, with a higher bandwidth resulting in more aggressive smoothing. We note that using kernel density estimation to estimate the distribution of future cardiac arrests is supported by the fact that cardiac arrest locations in Toronto have been shown to exhibit temporal stability (Chan et al. 2016).

To assess model performance under different values of the bandwidth parameter h , we constructed four different validation sets. Specifically, we used Gaussian kernels with $h = 10, 50, 100,$ and 150 meters to estimate four demand densities, each of which is used to simulate a separate validation set. The case where $h = 10$ can be interpreted as a “low uncertainty” environment, since the low bandwidth causes a majority of simulated events to fall near a historical cardiac arrest location. Conversely, the case where $h = 150$ can be interpreted as a “high uncertainty” environment, since the larger bandwidth leads to some cardiac arrests being simulated far from any historical location. By varying the bandwidth parameter from 10 to 150, we are able to evaluate the performance of the AED deployment under a variety of possible demand distributions. For each of the four estimated densities, we simulated a total of 50 validation sets containing 100 cardiac arrests each. We do not impose any constraints on where cardiac arrests may be simulated (i.e., a small fraction may occur beyond the service area). Figure 3 illustrates the effect of the bandwidth on the estimated demand density. Note that the distribution of the simulated cardiac arrests is notably more uniform for a bandwidth of 100 meters compared with the distribution for a bandwidth of 50 meters, as a result of increased smoothing.

Figure 3. Locations of 1,000 Simulated Cardiac Arrests Based on Kernel Density Estimation Using Bandwidths of 50 m and 100 m



6. AED Deployment Study: Results

The numerical results are organized as follows. In Section 6.1, we compute the average value (over 50 validation sets) of four performance metrics: the mean, VaR, CVaR, and maximum of the distance distribution induced by each AED deployment. We also report worst-case performance of each deployment, which is given by the maximum value of each performance metric over 50 validation sets. In Section 6.2, we evaluate the models using a different but related measure of AED deployment performance, known as cardiac arrest coverage. Finally, in Section 6.3, we do a robustness check by presenting results under alternate distance measures—namely, actual walking distances obtained via Google Maps and the rectilinear distance metric.

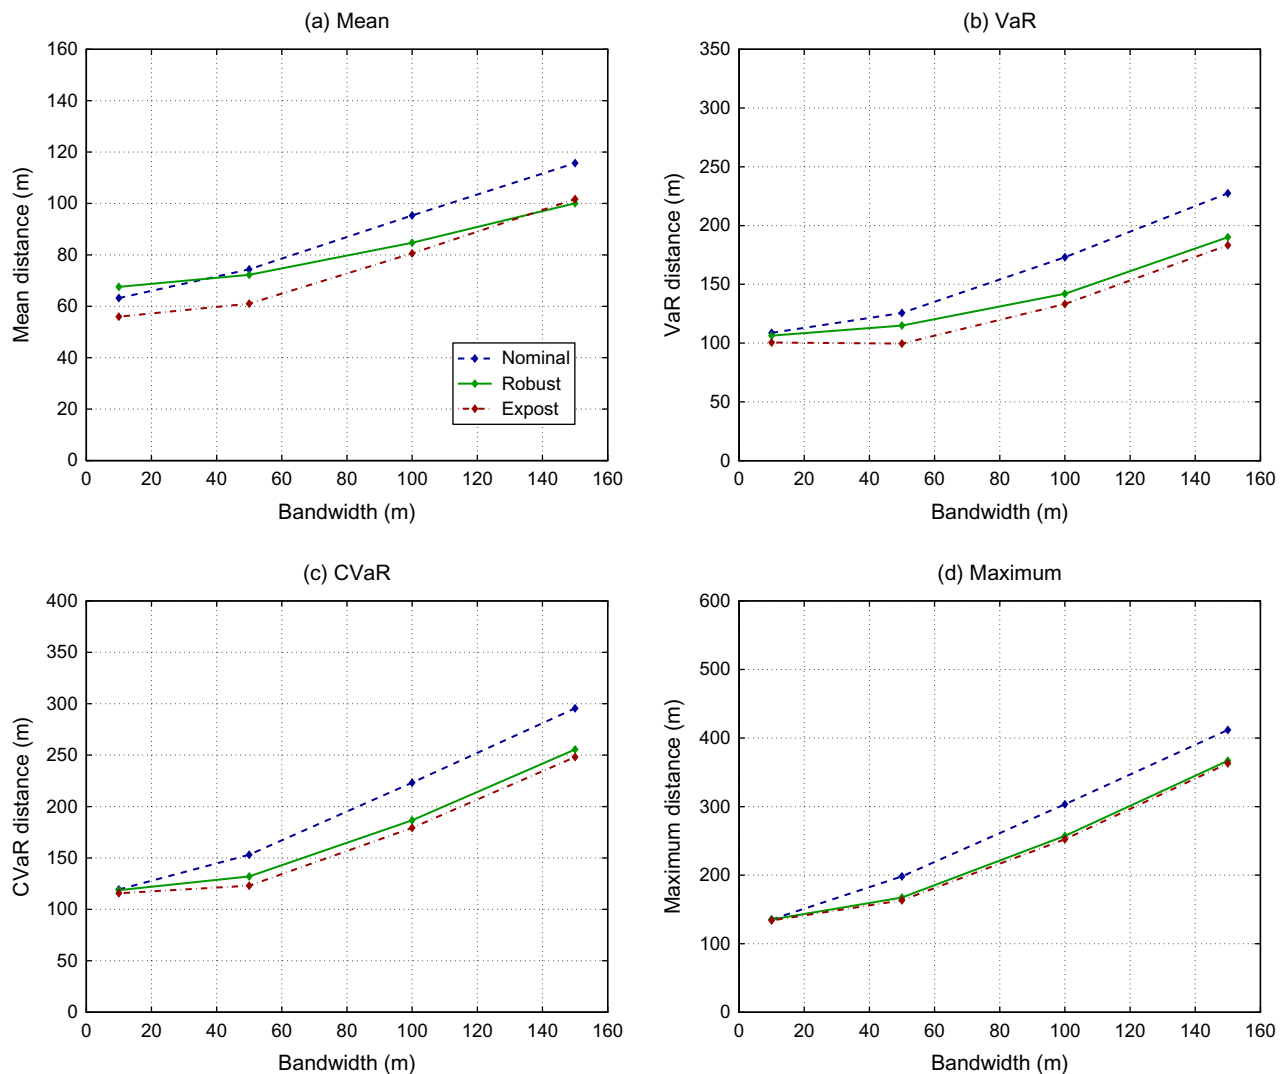
6.1. Performance Metrics

Figure 4 shows the mean, VaR, CVaR, and maximum value of the distance distribution for the nominal, robust, and ex post models.

robust, and ex post models (note that the y axis differs between the plots). The performance metrics are computed for each of the four demand distributions (i.e., four bandwidth parameters) and are averaged over 50 validation sets. Since Figure 4 shows the average value of each performance metric, we can interpret it as depicting the performance of each model under “typical” cardiac arrest locations.

In the cases where $h = 10$, the AED deployments produced by all three models perform similarly. This is not a surprising observation, since in the $h = 10$ case, which represents a low uncertainty environment, the vast majority of simulated cardiac arrests occur within the vicinity of one of the 43 historical cardiac arrest locations. However, it is worth noting that the robust solution performs at least as well as the nominal solution in the $h = 10$ case. In other words, the performance of the robust model does not appear to suffer when the level of uncertainty is low, despite optimizing AED

Figure 4. (Color online) Average Performance over 50 Validation Sets of Nominal, Robust, and Ex Post Solutions Under Increasing Kernel Bandwidth ($\beta = 0.9$)



locations to protect against uncertainty. Moreover, in the cases where $h \geq 50$, the robust solution substantially outperforms the nominal solution on all four performance metrics. For example, in the case where $h = 100$, the robust solution outperforms the nominal solution on the mean distance metric by 9% (86 m versus 95 m) and on the CVaR metric by 15% (223 m versus 189 m). The performance difference between the robust and nominal models was found to be statistically significant ($p < 0.05$) in all comparisons where $h \geq 50$. Examples of AED deployments under each model for $\beta = 0$ and $\beta = 0.9$ are given in Section EC.3 of the electronic companion.

Note that the ex post model achieves the best performance in all instances, which is expected since it optimizes the AED deployment directly over the simulated cardiac arrest locations (note that the ex post model is solved independently for each of the 50 trials, since it takes the simulated cardiac arrest locations as input). The gap in performance between the nominal and ex post models may therefore be interpreted as the degradation in performance as a result of the uncertainty in cardiac arrest locations. Figure 4 shows that the robust deployment substantially closes this performance gap in many cases. For example, for the CVaR metric with $h = 100$, using the robust deployment decreases the performance gap with the ex post model from 44 meters (223 m versus 179 m) down to 10 meters (189 m versus 179 m), representing a gap improvement of 77%. In this setting, it may be convenient to think of the performance gap as being analogous to the usual notion of optimality gap. However, we emphasize that the ex post solution represents a theoretical lower bound on performance only for the CVaR metric, since all models optimize for CVaR only. For the remaining three performance metrics, it is possible for the nominal or robust model to outperform the ex post model by chance.

Table 2 shows the performance gap achieved by the nominal and robust deployment with respect to the ex post deployment. For example, the first row and column of Table 2 indicates that the performance of the nominal solution on the mean distance metric in the $h = 10$ case is within 12% of the performance achieved by the ex post solution, whereas the gap from the robust solution is within 7% (averaged over 50 trials).

Table 2. Average Performance Gap of Nominal and Robust Solutions with Respect to Ex Post Solution

h	Mean (%)		VaR (%)		CVaR (%)		Max (%)	
	Nom.	Rob.	Nom.	Rob.	Nom.	Rob.	Nom.	Rob.
10	12	7	8	5	3	3	1	1
50	18	11	21	13	20	8	18	2
100	15	6	23	9	20	5	17	3
150	12	1	19	4	16	4	12	3

The difference in performance between the nominal and robust solutions is most pronounced at the tail of the distribution, which is unsurprising since the models optimize for CVaR. Note that for the CVaR and maximum distance metrics, the performance gap of the robust solutions is at most 8%, whereas the gap of the nominal solution can be as high as 20%.

We also consider the worst-case performance of each model by computing the maximum value of each performance metric over the 50 trials. As shown in Figure 5, the robust model generally outperforms the nominal model in the majority of high uncertainty instances ($h \geq 50$) and performs no worse than the nominal model in the low uncertainty instances ($h = 10$). Table 3 presents the worst-case performance gaps, which are obtained by computing the maximum relative gap in performance over 50 trials. As in Table 2, the robust solution substantially closes the worst-case performance gap in the vast majority of instances. For example, on the CVaR metric with $h = 100$, the worst-case performance gap is improved by 41% (27% down to 16%).

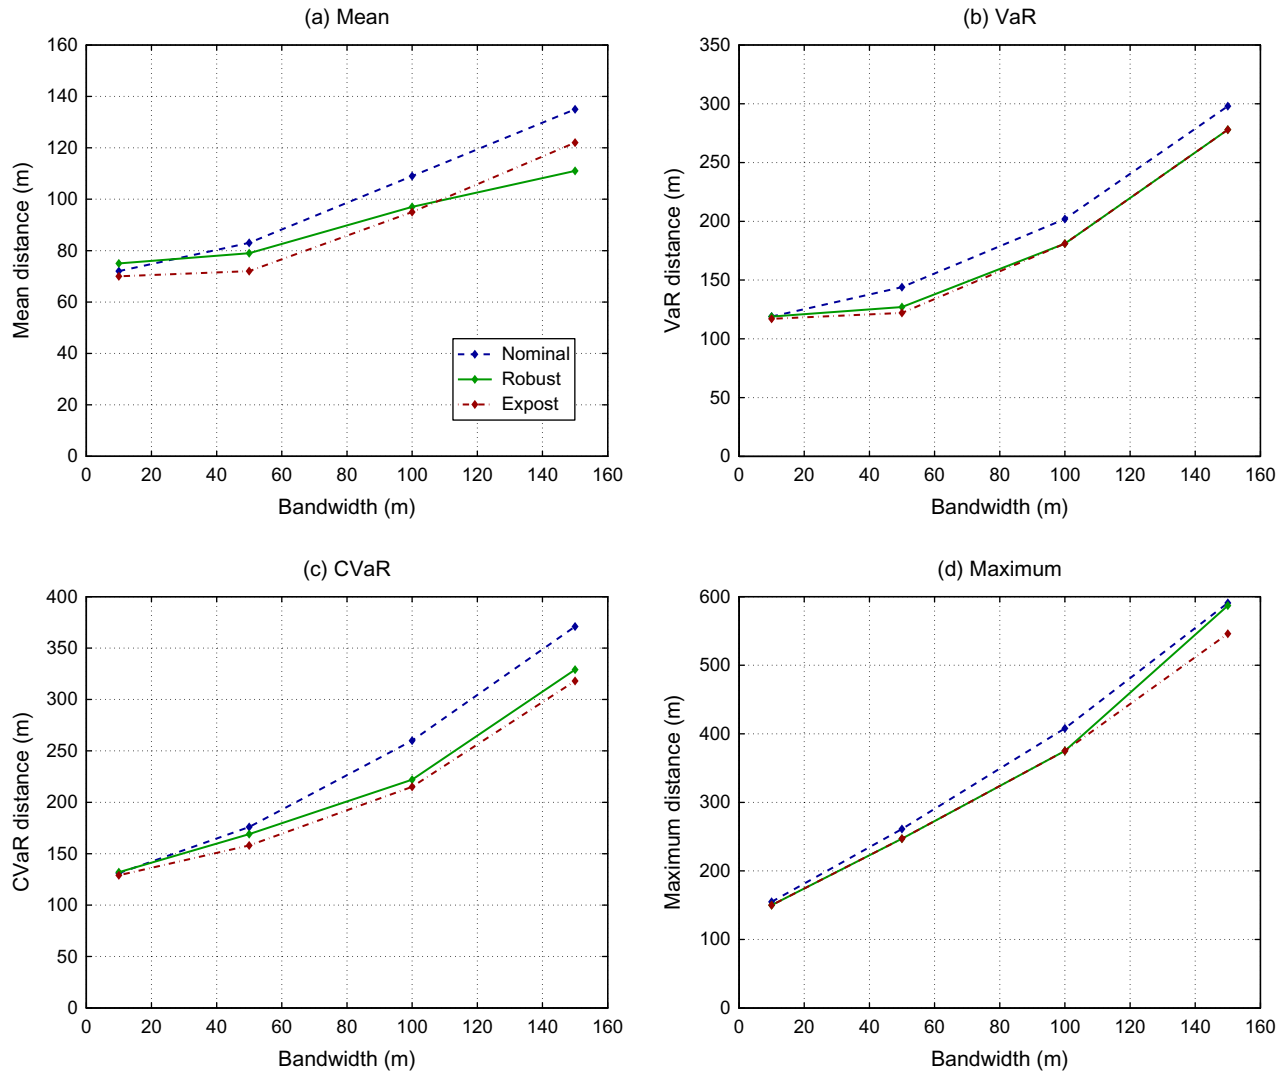
6.2. Cardiac Arrest Coverage

An alternate but related measure of performance in public access defibrillation is cardiac arrest *coverage*, which is the fraction of cardiac arrests that occur within a given distance of at least one AED (Chan et al. 2013). Figure 6 shows the cardiac arrest coverage under the nominal, robust, and ex post solutions at all distances between 0 and 250 meters, for each of the four bandwidths. The curves shown in Figure 6 are obtained by computing the average coverage at each distance over all 50 trials. Note that in the case where $h = 10$, all three models attain similar coverage at all distances, which is consistent with the results shown in Figure 4. Note that the robust solution dominates the nominal solution at all bandwidths, in the sense that it achieves a higher coverage level at all distances. Moreover, the coverage of the robust model is very close to the coverage obtained by the ex post model, especially for high bandwidths. The improvement of the robust deployment over the nominal deployment appears most pronounced at higher distances (≥ 100 m), which is expected behavior since the AED locations are optimized for the tail of the distance distribution.

6.3. Alternate Distance Measures

We repeated a set of experiments using two alternate distance measures: the rectilinear distance, and actual walking distances based on city roads and walking paths. Table 4 shows the performance of each model when Euclidean, rectilinear, and walking distances are used to compute the distances (for both the d_{ik} parameters and the performance metrics). When measured using actual walking distances, the robust and nominal

Figure 5. (Color online) Worst-Case Performance of Nominal, Robust, and Ex Post Solutions on Four Distance Metrics Under Increasing Kernel Bandwidth



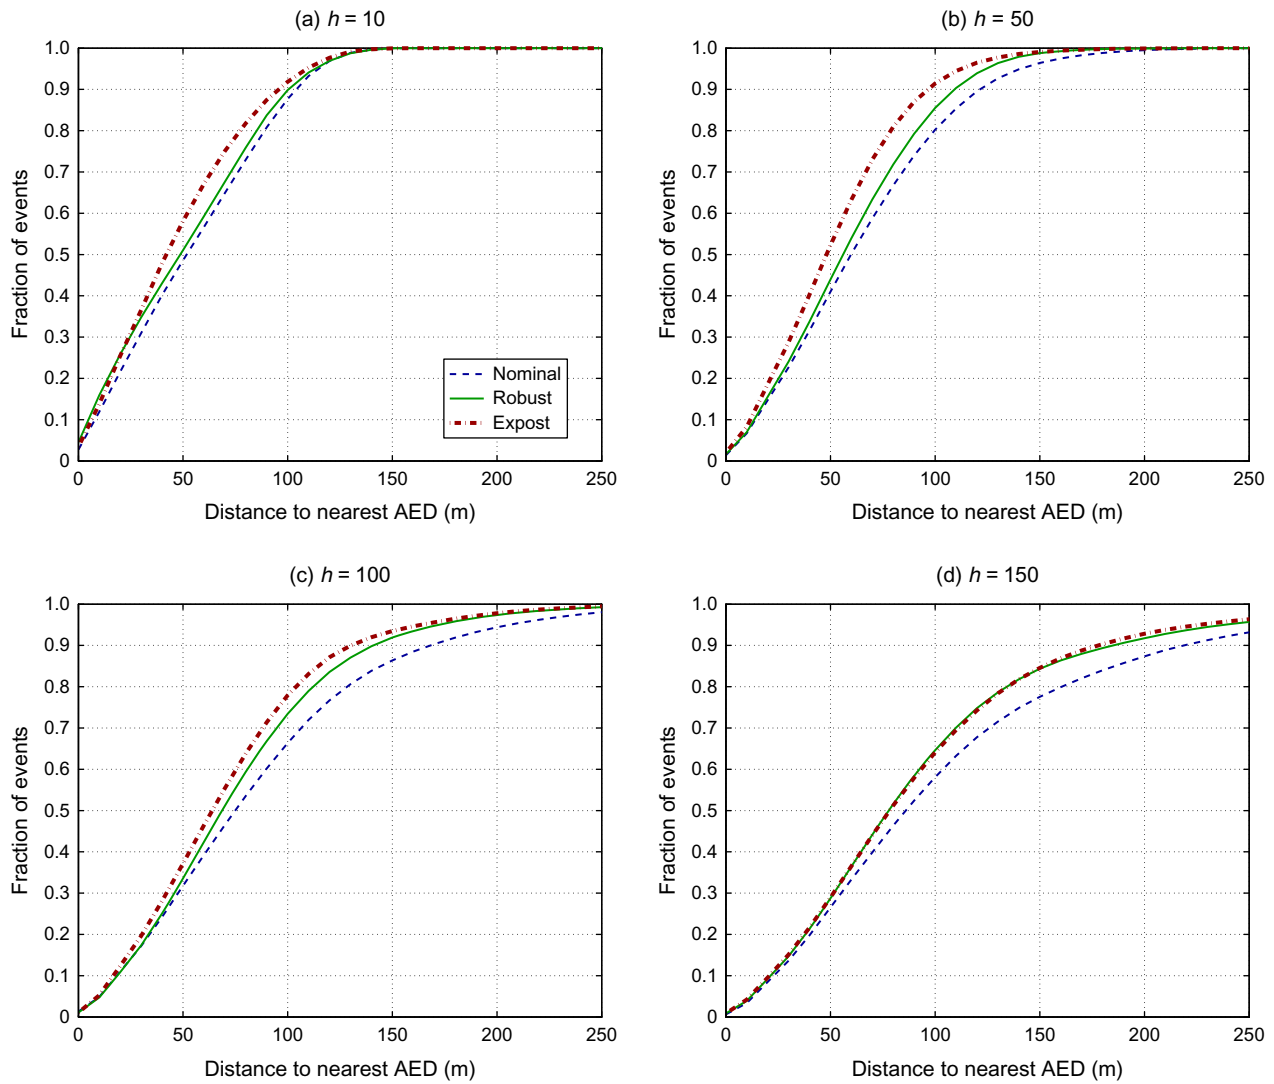
deployments perform similarly on the mean distance metric. However, the robust model attains a lower VaR and CVaR, indicating improved performance by the robust deployment at the tail of the distance distribution. Table 5 provides the accompanying average and worst-case performance gaps of the nominal and robust solutions with respect to the ex post solution.

Table 3. Worst Case Performance Gap of Nominal and Robust Solutions with Respect to Ex Post Solution

h	Mean (%)		VaR (%)		CVaR (%)		Max (%)	
	Nom.	Rob.	Nom.	Rob.	Nom.	Rob.	Nom.	Rob.
10	30	24	28	19	10	7	6	6
50	29	20	34	29	31	19	41	20
100	29	17	40	26	27	16	31	24
150	20	10	31	20	24	12	26	26

We found the results to be consistent with our main finding in Section 6.1: the AED deployments produced by the robust model outperform the nominal model on all four performance metrics, both on average and in the worst-case realization of cardiac arrest locations. On the CVaR performance metric (which we emphasize is the most relevant point of comparison since all three models optimize for CVaR), the robust model cuts the performance gap in half when compared with the nominal model. It should be noted that the walking distances from the Google Maps API are likely conservative estimates of the actual distance a lay responder would have to travel, since it constrains the lay responder to walk on recognized pedestrian paths. In practice, a lay responder may be able to take shortcuts that are unmarked in Google Maps (e.g., parking lots, buildings, unmarked crosswalks). We therefore expect the actual lay responder distance to be

Figure 6. (Color online) Cardiac Arrest Coverage for Nominal, Robust, and Ex Post Solutions Under Four Kernel Bandwidths



somewhere between the Euclidean and Google Maps distances.

As an additional check to compare Euclidean and actual walking distances, we performed the following experiment. We randomly generated 1,000 pairs of

locations within the downtown region of the City of Toronto, which includes the service area in our study. For each pair of locations, we computed the Euclidean distance (based on the UTM coordinates of the locations) and the walking distance (via the Google Maps

Table 4. Average and Worst-Case Performance of Nominal, Robust, and Ex Post Solutions Over 50 Simulated Validation Sets Using Different Distance Measures ($\beta = 0.9, h = 100$)

	Mean			VaR			CVaR			Max		
	Nom.	Rob.	Ex.	Nom.	Rob.	Ex.	Nom.	Rob.	Ex.	Nom.	Rob.	Ex.
Average												
Euclidean	95	86	81	173	147	133	223	189	179	202	260	252
Rectilinear	121	107	101	216	190	164	280	248	223	386	342	316
Walking	134	133	119	233	221	202	281	269	256	367	365	355
Worst-case												
Euclidean	109	96	95	202	181	181	260	222	215	408	375	375
Rectilinear	136	120	119	263	241	215	345	311	283	553	553	498
Walking	150	148	137	256	243	236	326	315	307	558	558	572

Table 5. Average and Worst-Case Performance Gap of Nominal and Robust Solutions with Respect to Ex Post Solution Under Alternate Distance Measures

	Mean (%)		VaR (%)		CVaR (%)		Max (%)	
	Nom.	Rob.	Nom.	Rob.	Nom.	Rob.	Nom.	Rob.
Average								
Euclidean	15	6	23	9	20	5	17	3
Rectilinear	17	6	24	13	21	10	18	8
Walking	12	11	13	8	9	5	7	3
Worst-case								
Euclidean	29	17	40	26	27	16	31	24
Rectilinear	28	18	41	27	32	20	41	23
Walking	21	20	31	24	20	12	24	22

Distance Matrix API). We then performed an ordinary least squares regression on the distance data, which yielded the following model:

$$Dist_{Walk} = 1.341 \times Dist_{Euclid} + 21.1.$$

The R^2 coefficient of this model was 0.904, suggesting a strong linear relationship between Euclidean and actual walking distances. Both coefficients were statistically significant ($p < 0.05$). This finding is consistent with previous studies that have also found very high correlation between Euclidean and road network distances (Bach 1981, Jones et al. 2010, Boscoe et al. 2012). These results suggest that Euclidean distances serve as a reasonable (constant factor) approximation for actual walking distances.

7. Discussion and Policy Implications

Interestingly, our results indicate that the robust deployment performed better than the nominal deployment under *typical* demand realizations (see Figure 4) in addition to the worst-case realizations (see Figure 5). This might seem like a counterintuitive result, given that robust optimization models can perform poorly in typical realizations of the uncertainty because of their emphasis on optimizing a worst-case objective function. However, in the context of AED deployment (and facility location more generally), the observed performance gap between the robust and nominal models can likely be attributed to the nominal model overfitting to the set of historical cardiac arrest locations. In other words, the nominal model represents a naive sample average approximation in which AED placement is optimized only with respect to a small set of historical cardiac arrest locations. Our results suggest that if there is a nontrivial amount of uncertainty in the locations of future cardiac arrest, a sample average approximation based on limited historical data can lead to AED deployments that perform poorly when measured against out-of-sample cardiac arrests. On the other hand, our robust formulation produces an improved deployment by accounting for the

possibility of cardiac arrests occurring in new locations not reflected in the historical data.

7.1. Survival Outcomes

Our results suggest that a robust approach to AED deployment can improve cardiac arrest response by improving bystander accessibility to AEDs, especially for those patients who collapse far away from the nearest defibrillator. While a rigorous analysis of survival outcomes is beyond the scope of this paper, the time-sensitive nature of cardiac arrest treatment suggests that even modest improvements in AED accessibility (through shorter distances) can have a significant impact on survival outcomes. We also emphasize that all numerical results presented in Section 6 are with respect to the one-way distance to the nearest AED. However, since retrieval of an AED by a lay responder likely requires a round-trip, decreasing the distance between the cardiac arrest victim and the nearest AED by 30 meters implies that the trip made by the lay responder is shortened by 60 meters, which can have a material impact on the total time between collapse and treatment.

7.2. Comparison to Best Practices

Historically, best practices for AED deployment (e.g., as prescribed by the American Heart Association (AHA)) have been purely retrospective and do not account for the uncertainty in future cardiac arrest locations, which our results suggest can help improve the accessibility of AEDs. Existing guidelines also emphasize the importance of identifying building *types* that are at risk of cardiac arrest, as well as placing AEDs in those locations. Our proposed approach is markedly different from the current AHA recommendations. We instead focus on characterizing the distributional uncertainty in cardiac arrest locations and optimally deploying AEDs in a manner that does not depend on building type, as long as the AED can be accessed by a lay responder. We also note that placing AEDs in a small set of high-risk buildings ignores the shape of the distance distribution and may lead to AED deployments where a large number of cardiac arrests occur far away from any AED. By contrast, our approach considers the entirety of the (worst-case) distance distribution and allows the modeler to directly optimize various aspects of the distribution via the CVaR objective function.

7.3. Crowdsourcing Lay Responders

The existing focus on identifying high-risk buildings for AED placement is partially due to the current low rate of bystander intervention during cardiac arrests and a prevailing belief that an AED must be located at the precise location of a cardiac arrest in order to be used. However, recent advances in mobile phone applications show promise in recruiting lay responders (who may be up to 500 m away) to perform CPR

on victims of cardiac arrest (Ringh et al. 2015; Brooks et al. 2016, 2014). This “crowdsourcing” approach to cardiac arrest interventions can be extended to incorporate AEDs as well, by notifying lay responders of the locations of both the cardiac arrest victim and a nearby AED. We posit that in a setting where lay responders are recruited to intervene during a cardiac arrest, it is even more critical that AEDs are tactically located throughout an urban area, so that they can be quickly located and transported to a victim of cardiac arrest by a lay responder. Thus, our approach to strategically deploying AEDs can work synergistically with other innovations in the chain of survival to improve the overall response to cardiac arrest.

7.4. Risk Measures in Public Access Defibrillation

Our model also highlights the potential role that risk measures such as CVaR can play in public access defibrillation. While VaR-based response time targets are common in ambulance dispatching (Pons and Markovchick 2002), no similar guidelines exist in public access defibrillation. Given the potential impact on survival of protecting against large distances between cardiac arrest victims and nearby AEDs, it may be worthwhile to explicitly incorporate risk-based targets in the decision-making process when determining AED locations. This paper formally introduces the notion of risk to the AED deployment problem by allowing decision makers to emphasize the tail of the distance (or response time) distribution when planning AED locations.

8. Conclusion

We develop a data-driven optimization approach for deploying AEDs in public spaces while accounting for uncertainty in the locations of future cardiac arrests. Our approach involves constructing a distributional uncertainty set that consists of several uncertainty regions, where the parameters of the uncertainty set are calibrated based on historical cardiac arrest data. To approximate the arrival of cardiac arrests in continuous space, we discretize the service area into an extremely large set of scenarios, and we develop a row-and-column generation algorithm that exploits the structure of the uncertainty set and scales gracefully in the number of scenarios. As an auxiliary result, we show that our formulation subsumes a large class of facility location problems and unifies the classical p -median and p -center problems, as well as their robust analogues.

Our numerical results suggest that hedging against cardiac arrest location uncertainty can lead to improved AED accessibility under both typical and worst-case realizations of the uncertainty. In particular, we found that our robust AED deployment outperformed a nominal (sample average approximation) deployment by between 9% and 20% with respect to AED

retrieval distances, depending on the performance metric and underlying demand distribution. Furthermore, we found that in many cases, the robust deployment performed nearly as well as an ex post model with perfect foresight, and in many instances it improved the performance gap with respect to the ex post model by 40%–70%. This finding suggests that our robust approach manages to avoid the performance loss due to uncertainty that is suffered by the sample average approach.

We highlight a few potential directions for future work. First, we have assumed throughout this paper that the set of uncertainty regions is given. It may be fruitful to investigate data-driven approaches to identifying an effective configuration of the uncertainty regions. Similarly, we note that the uncertainty set proposed in this paper can be generalized to permit the arrival probabilities to be uncertain as well. For example, we might allow each λ_j to reside within a 95% confidence interval that is also estimated from the data, which would retain the polyhedral structure of the uncertainty set. Our model can also be extended to incorporate additional considerations in public access defibrillation, such as lay responder behavior. However, we expect that our main finding in this paper—that accounting for cardiac arrest location uncertainty improves AED accessibility—would persist under such modifications to the model. Finally, the main features of our modeling approach—a distributional uncertainty set that induces sparse worst-case distributions coupled with an efficient row-and-column generation algorithm—may find relevance in other applications with a similar problem structure as well.

Acknowledgments

The authors thank area editor Chung Piaw Teo, the associate editor, and three anonymous referees for providing valuable comments that significantly improved the paper. The authors also thank Laurie Morrison and the Resuscitation Outcomes Consortium for the cardiac arrest data.

References

- Alderson DL, Brown GG, Carlyle WM, Wood RK (2011) Solving defender-attacker-defender models for infrastructure defense. *12th INFORMS Comput. Soc. Conf.* (INFORMS, Hanover, MD), 28–49.
- Atamturk A, Zhang M (2007) Two-stage robust network flow and design under demand uncertainty. *Oper. Res.* 55(4):662–673.
- Averbakh I, Berez S (2005) Facility location problems with uncertainty on the plane. *Discrete Optim.* 2(1):3–34.
- Bach L (1981) The problem of aggregation and distance for analyses of accessibility and access opportunity in location-allocation models. *Environ. Planning A* 13(8):955–978.
- Baron O, Milner J, Naseraldin H (2011) Facility location: A robust optimization approach. *Production Oper. Management* 20(5): 772–785.
- Becker L, Eisenberg M, Fahrenbruch C, Cobb L (1998) Public locations of cardiac arrest implications for public access defibrillation. *Circulation* 97(21):2106–2109.
- Ben-Tal A, Ghaoui LE, Nemirovski A (2009) *Robust Optimization* (Princeton University Press, Princeton, NJ).

- Ben-Tal A, Den Hertog D, De Waegenare A, Melenberg B, Rennen G (2013) Robust solutions of optimization problems affected by uncertain probabilities. *Management Sci.* 59(2):341–357.
- Bertsimas D, Sim M (2004) The price of robustness. *Oper. Res.* 52(1):35–53.
- Bertsimas D, Thiele A (2006) Robust and data-driven optimization: Modern decision making under uncertainty. Johnson MP, Norman B, Secomandi N, eds. *Models, Methods, and Applications for Innovative Decision Making*, INFORMS Tutorials in Operations Research (INFORMS, Hanover, MD), 95–122.
- Bertsimas D, Brown D, Caramanis C (2011) Theory and applications of robust optimization. *SIAM Rev.* 53(3):464–501.
- Bertsimas D, Litvinov E, Sun XA, Zhao J, Zheng T (2013) Adaptive robust optimization for the security constrained unit commitment problem. *IEEE Trans. Power Systems* 28(1):52–63.
- Bickel PJ, Doksum KA (2015) *Mathematical Statistics: Basic Ideas and Selected Topics*, Vol. 2 (CRC Press, Boca Raton, FL).
- Birge JR, Louveaux F (2011) *Introduction to Stochastic Programming* (Springer Science and Business Media, New York).
- Birge JR, Wets RJB (1987) Computing bounds for stochastic programming problems by means of a generalized moment problem. *Math. Oper. Res.* 12(1):149–162.
- Boscoe FP, Henry KA, Zdeb MS (2012) A nationwide comparison of driving distance versus straight-line distance to hospitals. *Professional Geographer* 64(2):188–196.
- Brooks SC, Hsu JH, Tang SK, Jeyakumar R, Chan TCY (2013) Determining risk for out-of-hospital cardiac arrest by location type in a Canadian urban setting to guide future public access defibrillator placement. *Ann. Emergency Medicine* 61(5):530–538.
- Brooks SC, Simmons G, Worthington H, Bobrow BJ, Morrison LJ (2016) The PulsePoint respond mobile device application to crowdsourcing basic life support for patients with out-of-hospital cardiac arrest: Challenges for optimal implementation. *Resuscitation* 98:20–26.
- Brooks SC, Worthington H, Gonedalles T, Bobrow B, Morrison LJ (2014) Implementation of the PulsePoint smartphone application for crowd-sourcing bystander resuscitation. *Critical Care* 18(Suppl. 1):P484.
- Brotcorne L, Laporte G, Semet F (2003) Ambulance location and relocation models. *Eur. J. Oper. Res.* 147(3):451–463.
- Brown G, Carlyle M, Salmeron J, Wood K (2006) Defending critical infrastructure. *Interfaces* 36(6):530–544.
- Brown GG, Carlyle WM, Harney RC, Skroch EM, Wood RK (2009) Interdicting a nuclear-weapons project. *Oper. Res.* 57(4):866–877.
- Caffrey SL, Willoughby PJ, Pepe PE, Becker LB (2002) Public use of automated external defibrillators. *New England J. Medicine* 347(16):1242–1247.
- Carlsson JG, Delage E (2013) Robust partitioning for stochastic multivehicle routing. *Oper. Res.* 61(3):727–744.
- Chan TCY, Mišić VV (2013) Adaptive and robust radiation therapy optimization for lung cancer. *Eur. J. Oper. Res.* 231(3):745–756.
- Chan TCY, Demirtaş D, Kwon R (2016) Optimizing the deployment of public access defibrillators. *Management Sci.* 62(12):3617–3635.
- Chan TCY, Li H, Lebovic G, Tang SK, Chan JYT, Cheng HCK, Morrison LJ, Brooks SC (2013) Identifying locations for public access defibrillators using mathematical optimization. *Circulation* 127(17):1801–1809.
- Chen G, Daskin MS, Shen ZJM, Uryasev S (2006) The α -reliable mean-excess regret model for stochastic facility location modeling. *Naval Res. Logist.* 53(7):617–626.
- Cho SH, Jang H, Lee T, Turner J (2014) Simultaneous location of trauma centers and helicopters for emergency medical service planning. *Oper. Res.* 62(4):751–771.
- Church RL, Scaparra MP (2007) Protecting critical assets: The R-interdiction median problem with fortification. *Geographical Anal.* 39(2):129–146.
- City of Toronto (2010) Toronto employment survey (for 2009). Accessed December 15, 2016, https://www1.toronto.ca/wps/portal/contentonly?vgnextoid=1e68f40f9aae0410VgnVCM10000071d60f89RCRD&WT.rd_id=demographics.
- City of Toronto (2016) Demographics—Neighbourhood profiles. Accessed December 15, 2016, <http://www1.toronto.ca/wps/portal/>.
- Cooper L (1978) Bounds on the Weber problem solution under conditions of uncertainty. *J. Regional Sci.* 18(1):87–92.
- Cui T, Ouyang Y, Shen ZJM (2010) Reliable facility location design under the risk of disruptions. *Oper. Res.* 58(4):998–1011.
- Delage E, Ye Y (2010) Distributionally robust optimization under moment uncertainty with application to data-driven problems. *Oper. Res.* 58(3):595–612.
- Drezner Z (1989) Stochastic analysis of the Weber problem on the sphere. *J. Oper. Res. Soc.* 40(12):1137–1144.
- Dupačová J (1980) On minimax decision rule in stochastic linear programming. Prékopa A, ed. *Stud. Math. Programming* (Akadémiai Kiadó, Budapest), 47–60.
- Engdahl J, Herlitz J (2005) Localization of out-of-hospital cardiac arrest in Göteborg 1994–2002 and implications for public access defibrillation. *Resuscitation* 64(2):171–175.
- Erkut E, Ingolfsson A, Erdoğan G (2008) Ambulance location for maximum survival. *Naval Res. Logist.* 55(1):42–58.
- Farias VF, Jagabathula S, Shah D (2013) A nonparametric approach to modeling choice with limited data. *Management Sci.* 59(2):305–322.
- Folke F, Lippert F, Nielsen S, Gislason G, Hansen M, Schramm T, Sorensen R (2009) Location of cardiac arrest in a city center strategic placement of automated external defibrillators in public locations. *Circulation* 120(6):510–517.
- Francis RL, Lowe TJ, Tamir A (2000) Aggregation error bounds for a class of location models. *Oper. Res.* 48(2):294–307.
- Gabrel V, Lacroix M, Murat C, Remli N (2014) Robust location transportation problems under uncertain demands. *Discrete Appl. Math.* 164(1):100–111.
- Ghaoui LEI, Oks M, Oustry F (2003) Worst-case value-at-risk and robust portfolio optimization: A conic programming approach. *Oper. Res.* 51(4):543–556.
- Goh J, Sim M (2010) Distributionally robust optimization and its tractable approximations. *Oper. Res.* 58(4, Part 1):902–917.
- Google (2016) Google Maps distance matrix API. Accessed December 15, 2016, <https://developers.google.com/maps/documentation/distance-matrix/>.
- Gundry JW, Comess KA, DeRook FA, Jorgenson D, Bardy GH (1999) Comparison of naive sixth-grade children with trained professionals in the use of an automated external defibrillator. *Circulation* 100(16):1703–1707.
- Hazinski MF, Idris AH, Kerber RE, Epstein A, Atkins D, Tang W, Lurie K (2005) Lay rescuer automated external defibrillator (“public access defibrillation”) programs: Lessons learned from an international multicenter trial: Advisory statement from the American Heart Association Emergency Cardiovascular Committee; the Council on Cardiopulmonary, Perioperative, and Critical Care; and the Council on Clinical Cardiology. *Circulation* 111(24):3336–3340.
- Heart and Stroke Foundation (2013) Statistics. Accessed December 15, 2016, https://www.heart.org/HEARTORG/General/Heart-and-Stroke-Association-Statistics_UCM_319064_SubHomePage.jsp.
- Jones SG, Ashby AJ, Momin SR, Naidoo A (2010) Spatial implications associated with using Euclidean distance measurements and geographic centroid imputation in health care research. *Health Services Res.* 45(1):316–327.
- Larsen M, Eisenberg M, Cummins RO, Hallstrom AP (1993) Predicting survival from out-of-hospital cardiac arrest: A graphic model. *Ann. Emergency Medicine* 22(11):1652–1658.
- Liberatore F, Scaparra MP, Daskin MS (2011) Analysis of facility protection strategies against an uncertain number of attacks: The stochastic R-interdiction median problem with fortification. *Comput. Oper. Res.* 38(1):357–366.
- Lim M, Daskin M, Bassamboo A, Chopra S (2010) A facility reliability problem: Formulation, properties, and algorithm. *Naval Res. Logist.* 57(1):58–70.

- Losada C, Scaparra MP, Church RL, Daskin MS (2012) The stochastic interdiction median problem with disruption intensity levels. *Ann. Oper. Res.* 201(1):345–365.
- Morrison LJ, Nichol G, Rea TD, Christenson J, Callaway CW, Stephens S, Pirralo RG, Atkins DL, Davis DP, Idris AH, Newgard C, et al. (2008) Rationale, development and implementation of the Resuscitation Outcomes Consortium Epistry—Cardiac Arrest. *Resuscitation* 78(2):161–169.
- Mozaffarian D, Benjamin EJ, Go AS, Arnett DK, Blaha MJ, Cushman M, Das SR, de Ferranti S, Després JP, Fullerton HJ, Howard VJ, et al. (2016) Executive summary: Heart disease and stroke statistics—2016 update: A report from the American Heart Association. *Circulation* 133(4):447–454.
- Nahmias S, Cheng Y (2009) *Production and Operations Analysis*, Vol. 5 (McGraw-Hill, New York).
- Ordóñez F, Zhao J (2007) Robust capacity expansion of network flows. *Networks* 50(2):136–145.
- Owen SH, Daskin MS (1998) Strategic facility location: A review. *Eur. J. Oper. Res.* 111(3):423–447.
- Page RL, Joglar JA, Kowal RC, Zagrodzky JD, Nelson LL, Ramaswamy K, Barbera SJ, Hamdan MH, McKenas DK (2000) Use of automated external defibrillators by a US airline. *New England J. Medicine* 343(17):1210–1216.
- Pell JP, Sirel JM, Marsden AK, Ford I, Cobbe SM (2001) Effect of reducing ambulance response times on deaths from out of hospital cardiac arrest: Cohort study. *British Medical J.* 322(7299):1385–1388.
- Pons PT, Markovchick VJ (2002) Eight minutes or less: Does the ambulance response time guideline impact trauma patient outcome? *J. Emergency Medicine* 23(1):43–48.
- Portner ME, Pollack ML, Schirk SK, Schlenker MK (2004) Out-of-hospital cardiac arrest locations in a rural community: Where should we place AEDs? *Prehospital Disaster Medicine* 19(4):352–355.
- Public Access Defibrillation Trial Investigators (2004) Public-access defibrillation and survival after out-of-hospital cardiac arrest. *New England J. Medicine* 351(7):637–646.
- PulsePoint (2015) PulsePoint AED. Accessed December 15, 2016, <http://www.pulsepoint.org/pulsepoint-aed/>.
- Ringh M, Rosenqvist M, Hollenberg J, Jonsson M, Fredman D, Nordberg P, Järnbert-Pettersson H, Hasselqvist-Ax I, Riva G, Svensson L (2015) Mobile-phone dispatch of laypersons for CPR in out-of-hospital cardiac arrest. *New England J. Medicine* 372(24):2316–2325.
- Rockafellar RT, Uryasev S (2000) Optimization of conditional value-at-risk. *J. Risk* 2(3):21–42.
- Rockafellar RT, Uryasev S (2002) Conditional value-at-risk for general loss distributions. *J. Banking Finance* 26(7):1443–1471.
- Scaparra MP, Church RL (2008) A bilevel mixed-integer program for critical infrastructure protection planning. *Comput. Oper. Res.* 35(6):1905–1923.
- Scarf H (1958) A min-max solution of an inventory problem. Scarf H, Arrow KJ, Karlin S, eds. *Studies in the Mathematical Theory of Inventory and Production* (Stanford University Press, Stanford, CA), 201–209.
- Shapiro A, Kleywegt A (2002) Minimax analysis of stochastic problems. *Optim. Methods Software* 17(3):523–542.
- Sheather S, Jones M (1991) A reliable data-based bandwidth selection method for kernel density estimation. *J. Roy. Statist. Soc. Ser. B (Methodological)* 53(3):683–690.
- Sheffi Y (2005) *The Resilient Enterprise: Overcoming Vulnerability for Competitive Advantage* (MIT Press, Cambridge, MA).
- Shen ZJM, Zhan RL, Zhang J (2011) The reliable facility location problem: Formulations, heuristics, and approximation algorithms. *INFORMS J. Comput.* 23(3):470–482.
- Siddiq AA, Brooks SC, Chan TC (2013) Modeling the impact of public access defibrillator range on public location cardiac arrest coverage. *Resuscitation* 84(7):904–909.
- Simchi-Levi D, Kaminsky P, Simchi-Levi E (2004) *Managing the Supply Chain: The Definitive Guide for the Business Professional* (McGraw-Hill, New York).
- Snyder LV (2006) Facility location under uncertainty: A review. *IIE Trans.* 38(7):547–564.
- Sun CL, Demirtas D, Brooks SC, Morrison LJ, Chan TCY (2016) Overcoming spatial and temporal barriers to public access defibrillators via optimization. *J. Amer. College Cardiol.* 68(8):836–845.
- Terrell GR, Scott DW (1992) Variable kernel density estimation. *Ann. Statist.* 20(3):1236–1265.
- Valenzuela TD, Roe DJ, Nichol G, Clark LL, Spaite DW, Hardman RG (2000) Outcomes of rapid defibrillation by security officers after cardiac arrest in casinos. *New England J. Medicine* 343(17):1206–1209.
- Wand MP (1997) Data-based choice of histogram bin width. *Amer. Statistician* 51(1):59–64.
- Weisfeldt ML, Sitlani CM, Ornato JP, Rea T, Aufderheide TP, Davis D, Dreyer J, et al. (2010) Survival after application of automatic external defibrillators before arrival of the emergency medical system evaluation in the resuscitation outcomes consortium population of 21 million. *J. Amer. College Cardiol.* 55(16):1713–1720.
- Wiesemann W, Kuhn D, Sim M (2014) Distributionally robust convex optimization. *Oper. Res.* 62(6):1358–1376.
- Wilson EB (1927) Probable inference, the law of succession, and statistical inference. *J. Amer. Statist. Assoc.* 22(158):209–212.
- Xu H, Caramanis C, Mannor S (2012) A distributional interpretation of robust optimization. *Math. Oper. Res.* 37(1):95–110.
- Zeng B, Zhao L (2013) Solving two-stage robust optimization problems using a column-and-constraint generation method. *Oper. Res. Lett.* 41(5):457–461.
- Zhao L, Zeng B (2012) Robust unit commitment problem with demand response and wind energy. 2012 *IEEE Power and Energy Society General Meeting* (IEEE, Piscataway, NJ), 1–8.

Timothy C. Y. Chan is the Canada Research Chair in Novel Optimization and Analytics in Health, an associate professor in the Department of Mechanical and Industrial Engineering, and the director of the Centre for Healthcare Engineering at the University of Toronto. His primary research interests are in robust and inverse optimization and the application of optimization methods to problems in healthcare and sports.

Zuo-Jun Max Shen is a Chancellor's Professor in the Department of Industrial Engineering and Operations Research and the Department of Civil and Environmental Engineering at the University of California, Berkeley. He has been active in the following research areas: integrated supply chain design and management, design and analysis of optimization algorithms, energy systems, and transportation system planning and optimization.

Auyon Siddiq is a PhD candidate in the Department of Industrial Engineering and Operations Research at the University of California, Berkeley. His current research interests are in data-driven optimization, healthcare operations, and incentive design.

Electronic companion for “Robust Defibrillator Deployment Under Cardiac Arrest Location Uncertainty via Row-and-Column Generation”.

EC.1. Proofs.

We first present two lemmas that are helpful in the proof of Theorem 1.

LEMMA EC.1. *Let \mathbf{y} be fixed. Define $\mathcal{U}^* = \operatorname{argmax}_{\mu \in \mathcal{U}} \operatorname{CVaR}_\mu[d(\xi, \mathbf{y}(\xi))]$. Then for any set of facility locations \mathbf{y} , there exists a worst-case distribution $\tilde{\mu} \in \mathcal{U}^*$ supported on $a_1, \dots, a_{|\mathcal{R}|}$, where $a_r = \operatorname{argmax}_{a \in \mathcal{A}'_r} d(a, \mathbf{y}(a))$, for all $r = 1, \dots, |\mathcal{R}|$.*

Proof. For any $\bar{\mathcal{A}} \subset \mathcal{A}$, define

$$\lambda(\mu, \bar{\mathcal{A}}) = P_\mu(\xi \in \bar{\mathcal{A}}) = \int_{\bar{\mathcal{A}}} f_\mu(\xi) d\xi, \quad (\text{EC.1})$$

where f_μ is the density function associated with distribution μ . Pick a distribution $\mu^* \in \mathcal{U}^*$. Now construct $\tilde{\mu}$ to be the distribution with mass $\lambda(\mu^*, \mathcal{A}'_1), \dots, \lambda(\mu^*, \mathcal{A}'_{|\mathcal{R}|})$ placed at the locations $a_1, \dots, a_{|\mathcal{R}|}$. We now show that $\tilde{\mu}$ is also in \mathcal{U}^* . We do so by first showing that $\tilde{\mu}$ is in \mathcal{U} , and then by showing that the CVaR of the distance distribution under the constructed $\tilde{\mu}$ is at least as large as the CVaR under μ^* . For the first step, let \mathcal{R}_j index the subregions that comprise \mathcal{A}_j . Now note that for each $j \in \mathcal{J}$, we have

$$P_{\tilde{\mu}}(\xi \in \mathcal{A}_j) = \sum_{r \in \mathcal{R}_j} P_{\tilde{\mu}}(\xi \in \mathcal{A}'_r) = \sum_{r \in \mathcal{R}_j} \lambda(\tilde{\mu}, \mathcal{A}'_r) = \sum_{r \in \mathcal{R}_j} P_{\mu^*}(\xi \in \mathcal{A}'_r) = P_{\mu^*}(\xi \in \mathcal{A}_j) = \lambda_j.$$

The first equality is due to the fact that $\bigcup_{r \in \mathcal{R}_j} \mathcal{A}'_r = \mathcal{A}_j$. The second equality follows from the definition of $\lambda(\mu, \bar{\mathcal{A}})$ in (EC.1). The third equality follows by the construction of $\tilde{\mu}$. The fourth equality is again due to $\bigcup_{r \in \mathcal{R}_j} \mathcal{A}'_r = \mathcal{A}_j$. The final equality holds since $\mu^* \in \mathcal{U}$. Since from the above equations, $P_{\tilde{\mu}}(\xi \in \mathcal{A}_j) = \lambda_j$, we have $\tilde{\mu} \in \mathcal{U}$. It remains to show that $\operatorname{CVaR}_{\mu^*} \leq \operatorname{CVaR}_{\tilde{\mu}}$. Writing CVaR explicitly, we have

$$\begin{aligned} \operatorname{CVaR}_{\mu^*}[d(\xi, \mathbf{y}(\xi))] &= \min_{\alpha} \alpha + \frac{1}{1-\beta} \int_{\mathcal{A}} [d(\xi, \mathbf{y}(\xi)) - \alpha]^+ f_{\mu^*}(\xi) d\xi \\ &= \min_{\alpha} \alpha + \frac{1}{1-\beta} \sum_{r \in \mathcal{R}} \int_{\mathcal{A}'_r} [d(\xi, \mathbf{y}(\xi)) - \alpha]^+ f_{\mu^*}(\xi) d\xi \\ &\leq \min_{\alpha} \alpha + \frac{1}{1-\beta} \sum_{r \in \mathcal{R}} \int_{\mathcal{A}'_r} [d(a_r, \mathbf{y}(a_r)) - \alpha]^+ f_{\mu^*}(\xi) d\xi \\ &= \min_{\alpha} \alpha + \frac{1}{1-\beta} \sum_{r \in \mathcal{R}} [d(a_r, \mathbf{y}(a_r)) - \alpha]^+ \int_{\mathcal{A}'_r} f_{\mu^*}(\xi) d\xi \\ &= \min_{\alpha} \alpha + \frac{1}{1-\beta} \sum_{r \in \mathcal{R}} [d(a_r, \mathbf{y}(a_r)) - \alpha]^+ \lambda(\tilde{\mu}, \mathcal{A}'_r) \\ &= \operatorname{CVaR}_{\tilde{\mu}}[d(\xi, \mathbf{y}(\xi))]. \end{aligned}$$

The first line follows from the definition of CVaR. The second line is due to the fact that the subregions \mathcal{A}'_r , $r \in \mathcal{R}$ form a partition of \mathcal{A} . The third line follows from the definition of a_r as a point in \mathcal{A}'_r that maximizes distance to the nearest AED. The fourth line follows since $[d(a_r, \mathbf{y}(a_r)) - \alpha]^+$ is constant with respect to the variable ξ . The fifth line follows from the definition of $\lambda(\tilde{\mu}, \mathcal{A}'_r)$ as the total probability mass in \mathcal{A}'_r . The final line follows by construction of $\tilde{\mu}$ and the definition of CVaR. \square

For the second lemma, define $z_C(\mathbf{y}) = \max_{\mu \in \mathcal{U}} \text{CVaR}_\mu[d(\xi, \mathbf{y}(\xi))]$ to be the worst-case CVaR under the solution \mathbf{y} in the continuous case. Similarly, let $z_D(\mathbf{y})$ be the worst-case CVaR in the discrete case (i.e., the optimal value of (4) given \mathbf{y}). The next lemma bounds the difference between the the worst-case continuous and discrete CVaR, for a given deployment \mathbf{y} . Note that Lemma 1 is used in the proof.

LEMMA EC.2. *For any \mathbf{y} , $z_C(\mathbf{y}) - z_D(\mathbf{y}) \leq \ell(\sigma)/(1 - \beta)$.*

Proof. Let μ^* be a worst-case continuous distribution with mass $\lambda'_1, \dots, \lambda'_{|\mathcal{R}|}$ on the locations $p_1, \dots, p_{|\mathcal{R}|}$ (cf Lemma EC.1). Construct a discrete distribution by placing mass $\lambda'_1, \dots, \lambda'_{|\mathcal{R}|}$ on the locations $\xi_1, \dots, \xi_{|\mathcal{R}|}$, where $\xi_r = \underset{\xi \in \Xi'_r}{\operatorname{argmin}} d(\xi, p_r)$ for all r . We can now write

$$\begin{aligned} z_C(\mathbf{y}) - z_D(\mathbf{y}) &= \left(\min_{\alpha} \alpha + \frac{1}{1 - \beta} \sum_{r \in \mathcal{R}} \lambda'_r [d(a_r, \mathbf{y}(a_r)) - \alpha]^+ \right) - z_D(\mathbf{y}) \\ &\leq \left(\min_{\alpha} \alpha + \frac{1}{1 - \beta} \sum_{r \in \mathcal{R}} \lambda'_r [d(\xi_r, \mathbf{y}(\xi_r)) + \ell(\sigma) - \alpha]^+ \right) - z_D(\mathbf{y}) \\ &\leq \left(\min_{\alpha} \alpha + \frac{1}{1 - \beta} \sum_{r \in \mathcal{R}} \lambda'_r [d(\xi_r, \mathbf{y}(\xi_r)) - \alpha]^+ \right) + \frac{1}{1 - \beta} \sum_{r \in \mathcal{R}} \lambda'_r \ell(\sigma) - z_D(\mathbf{y}) \\ &= z_D(\mathbf{y}) + \frac{1}{1 - \beta} \sum_{r \in \mathcal{R}} \lambda'_r \ell(\sigma) - z_D(\mathbf{y}) \\ &= \frac{\ell(\sigma)}{1 - \beta} \end{aligned}$$

The first line follows from the definition of $z_C(\mathbf{y})$. The second line follows from the triangle inequality and since $d(\xi_r, a_r) \leq \ell(\sigma)$ by definition of $\ell(\sigma)$. The third line follows from pulling $\ell(\sigma)$ out of the CVaR expression. The fourth line follows by definition of $z_D(\mathbf{y})$. The fifth line follows from cancelling out $z_D(\mathbf{y})$ terms and noting that $\sum_{r \in \mathcal{R}} \lambda'_r = 1$. \square

Proof of Theorem 1.

Note that $z_D(\mathbf{y}) \leq z_C(\mathbf{y})$ for any \mathbf{y} , since $\Xi'_r \subset \mathcal{A}'_r$ for all $r \in \mathcal{R}$. Now define $\mathbf{y}_C \in \underset{\mathbf{y}}{\operatorname{argmin}} z_C(\mathbf{y})$ and $\mathbf{y}_D \in \underset{\mathbf{y}}{\operatorname{argmin}} z_D(\mathbf{y})$. Using Lemma EC.2, we can now write

$$z_D(\mathbf{y}_D) \leq z_D(\mathbf{y}_C) \leq z_C(\mathbf{y}_C) = Z_C \leq z_C(\mathbf{y}_D) \leq z_D(\mathbf{y}_D) + \ell(\sigma)/(1 - \beta) = Z_D + \ell(\sigma)/(1 - \beta),$$

which implies $Z_C - Z_D \leq \ell(\sigma)/(1 - \beta)$. Since $Z_D \leq Z_C$, we have

$$\Delta = \frac{Z_C - Z_D}{Z_C} \leq \frac{Z_C - Z_D}{Z_D} \leq \frac{\ell(\sigma)}{Z_D(1 - \beta)}.$$

It follows that if $\ell(\sigma) \leq \varepsilon(1 - \beta)Z_D$, then $\Delta \leq \varepsilon$. \square

Proof of Proposition 1.

Note that the number of variables in the master problem formulation R-AED-MP is $|\mathcal{I}|(|\mathcal{K}_+| + 1) + 1$, and the number of constraints is $(2|\mathcal{I}| + 1)|\mathcal{K}_+| + |\mathcal{I}| + |\mathcal{S}| + 2$. Hence the number of variables and constraints in iteration s is $O(|\mathcal{I}||\mathcal{K}_+|)$. It remains to show that $|\mathcal{K}_+| \leq (|\mathcal{J}| + 1)s$ in the s^{th} iteration of the row-and-column generation algorithm. Since the new variables and constraints that are added to R-AED-MP in iteration s correspond directly to non-zero elements of the vector \mathbf{p}_s obtained at a solution to R-AED-SP (cf Algorithm 1), it suffices to show that any \mathbf{p} obtained at a solution to R-AED-SP contains at most $|\mathcal{J}| + 1$ non-zero elements. We first prove that any \mathbf{u} obtained at an optimal solution to R-AED-SP is also an optimal solution to the following problem:

$$\begin{aligned} & \underset{\mathbf{u}}{\text{maximize}} && \sum_{k \in \mathcal{K}} d_{ik} \bar{z}_{ik} u_k \\ & \text{subject to} && \sum_{k \in \mathcal{K}} u_k = 1, \\ & && \sum_{k \in \mathcal{K}_j} u_k = \lambda_j, \quad j \in \mathcal{J}, \\ & && u_k \geq 0. \end{aligned} \tag{EC.2}$$

By way of contradiction, suppose a solution \mathbf{p}, \mathbf{u} is optimal to R-AED-SP but \mathbf{u} is suboptimal to (EC.2). Then there exists k_1 and k_2 such that $\sum_{i \in \mathcal{I}} d_{ik_2} \bar{z}_{ik_2} > \sum_{i \in \mathcal{I}} d_{ik_1} \bar{z}_{ik_1}$. Further, there must exist $\epsilon > 0$ and a vector $\tilde{\mathbf{u}} \in \mathcal{U}$ such that $\tilde{u}_{k_1} = u_{k_1} - \epsilon$, $\tilde{u}_{k_2} = u_{k_2} + \epsilon$, and $\tilde{u}_k = u_k$ for all $k \in \mathcal{K} \setminus \{k_1, k_2\}$. Let $\tilde{\mathbf{p}}$ be the solution obtained by solving R-AED-SP with $\tilde{\mathbf{u}}$ as a parameter. Now note in R-AED-SP, since $\sum_{i \in \mathcal{I}} d_{ik_2} \bar{z}_{ik_2} > \sum_{i \in \mathcal{I}} d_{ik_1} \bar{z}_{ik_1}$, the constraints $p_{k_2} \leq u_{k_2}/(1 - \beta)$ and $\tilde{p}_{k_2} \leq \tilde{u}_{k_2}/(1 - \beta)$ must be binding (otherwise we could slightly increase p_{k_2} while slightly decreasing p_{k_1} by the same amount to improve the objective function, with an identical argument applying to \tilde{p}_{k_2}). Since $p_{k_2} = u_{k_2}/(1 - \beta)$ and $\tilde{p}_{k_2} = \tilde{u}_{k_2}/(1 - \beta)$, by the construction of $\tilde{\mathbf{u}}$ it follows that $\tilde{p}_2 - p_2 \geq \epsilon/(1 - \beta)$ and $p_1 - \tilde{p}_1 \leq \epsilon/(1 - \beta)$. Since $\sum_{i \in \mathcal{I}} d_{ik_2} \bar{z}_{ik_2} > \sum_{i \in \mathcal{I}} d_{ik_1} \bar{z}_{ik_1}$, the solution $\tilde{\mathbf{p}}, \tilde{\mathbf{u}}$ improves the objective by $(\sum_{i \in \mathcal{I}} d_{ik_2} \bar{z}_{ik_2} - \sum_{i \in \mathcal{I}} d_{ik_1} \bar{z}_{ik_1}) \frac{\epsilon}{1 - \beta} > 0$ over the solution \mathbf{p}, \mathbf{u} , which contradicts the optimality of \mathbf{p}, \mathbf{u} with respect to R-AED-SP. Hence \mathbf{u} must be optimal to (EC.2) as well.

Since (EC.2) is a standard form linear program with $|\mathcal{J}| + 1$ equality constraints and $|\mathcal{K}|$ variables, at most $|\mathcal{J}| + 1$ elements of a basic feasible solution to (EC.2) can be non-zero (Bertsimas and Tsitsiklis 1997). Further, since by assumption there are no ties in the d_{ik} parameters, every

optimal solution to (EC.2) must be a basic feasible solution. It follows that every \mathbf{u} obtained at a solution of R-AED-SP has no more than $|\mathcal{J}| + 1$ non-zero elements. Since $p_k > 0$ implies $u_k > 0$ for any feasible \mathbf{p}, \mathbf{u} , it follows that any \mathbf{p} obtained at a solution to R-AED-SP cannot have more than $|\mathcal{J}| + 1$ non-zero elements. \square

Proof of Proposition 2.

Let the subregions $\mathcal{A}'_1, \dots, \mathcal{A}'_{|\mathcal{R}|}$ be the partitioning of \mathcal{A} implied by the uncertainty regions $\mathcal{A}_1, \dots, \mathcal{A}_{|\mathcal{J}|}$. Let $\lambda'_r = P(\xi \in \mathcal{A}'_r)$ for all $r \in \mathcal{R}$ under the true distribution. Since the subregions are disjoint, the number of historical cardiac arrests that fall in each of the regions $\mathcal{A}'_1, \dots, \mathcal{A}'_{|\mathcal{R}|}$ is simply the outcome of n independent trials from a multinomial distribution with parameters $\lambda'_1, \dots, \lambda'_{|\mathcal{R}|}$. Now define

$$\hat{\lambda}'_r = \frac{1}{n} \sum_{c=1}^n \mathbb{1}\{a_c \in \mathcal{A}'_r\}, \quad r \in \mathcal{R},$$

is a maximum likelihood estimator of $\lambda'_1, \dots, \lambda'_r$, and that $|\hat{\lambda}'_r - \lambda'_r| \rightarrow 0$ with probability 1 as $n \rightarrow \infty$ (e.g., Example 1.6.7. in Bickel and Doksum (2015)). Now let \mathcal{R}_j be the set of subregions that comprise region \mathcal{A}_j . We now have $\hat{\lambda}_j = \sum_{r \in \mathcal{R}_j} \hat{\lambda}'_r$, and thus $|\hat{\lambda}_j - \lambda_j| = |\sum_{r \in \mathcal{R}_j} (\hat{\lambda}'_r - \lambda'_r)| \leq \sum_{r \in \mathcal{R}_j} |\hat{\lambda}'_r - \lambda'_r|$ for each $j \in \mathcal{J}$, which implies $|\hat{\lambda}_j - \lambda_j| \rightarrow 0$ with probability 1 for each $j \in \mathcal{J}$ as $n \rightarrow \infty$. To see that $\hat{\lambda}$ ensures that \mathbf{U} is non-empty, construct a distribution $\mathbf{u} \in \mathbf{U}$ as follows. First, assign each of the cardiac arrest locations a_1, \dots, a_n to the nearest scenario location within the same subregion. Then for each $k \in \mathcal{K}$, set u_k equal to the proportion of cardiac arrests that are assigned to location k . Thus \mathbf{u} is non-negative and $\sum_{k \in \mathcal{K}} u_k = 1$ by construction. Lastly, for each $r \in \mathcal{R}$, $\sum_{k \in \mathcal{K}'_r} u_k$ is equal to the proportion of cardiac arrests within subregion \mathcal{A}'_r , which is also equal to $\hat{\lambda}'_r$ by definition. Thus $\sum_{k \in \mathcal{K}_j} u_k = \sum_{r \in \mathcal{R}_j} \sum_{k \in \mathcal{K}'_r} u_k = \sum_{r \in \mathcal{R}_j} \hat{\lambda}'_r = \hat{\lambda}_j$. \square

EC.2. Unifying the p -median and p -center problems

The p -median and p -center problems are two of the most well-studied location models, having served as the foundation for a significant portion of the existing facility location literature (Owen and Daskin 1998, Snyder 2006, Melo et al. 2009). In this section, we discuss how a non-stochastic interpretation of R-AED specializes to robust variants of the p -median and p -center problems, depending on how β is selected. A straightforward corollary of this result is that if the \mathcal{A}_j are all singletons, then formulation R-AED unifies the classical p -median and p -center problems (i.e. without any uncertainty).

We begin by defining robust variants of the p -median and p -center problems, and then prove an equivalence between these problems and formulation R-AED. Let \mathcal{J} index a finite and fixed set of

demand points. In the classical p -median problem, one wishes to site P facilities so to minimize the (weighted) total distance between all demand points and their nearest facilities. Suppose now that the j^{th} demand point is only known to reside within an uncertainty region \mathcal{A}_j . We formulate a generalization of the p -median problem that aims to minimize the *worst-case* total distance. Let each uncertainty region be discretized into a set of locations $\Xi_j \subset \mathcal{A}_j$, so that the demand point j is only known to be at one of the locations in Ξ_j . Let \mathcal{K}_j index the locations in Ξ_j . Let $\mathbf{Z}(\mathbf{y})$ and \mathbf{Y} retain their definitions from Section 3. Let λ_j be the usual notion of “weight” placed on demand point j . Now define $\mathbf{X} = \mathbf{X}_1 \times \dots \times \mathbf{X}_{|\mathcal{J}|}$, where

$$\mathbf{X}_j = \left\{ x_{jk} \geq 0, k \in \mathcal{K} \mid \sum_{k \in \mathcal{K}_j} x_{jk} = 1 \right\}, \quad j \in \mathcal{J}. \quad (\text{EC.3})$$

We can now write the *robust p -median problem* to minimize the worst-case total distance as

$$\begin{aligned} \min_{\mathbf{y}, \mathbf{z}} \max_{\mathbf{x}} \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}_j} \sum_{j \in \mathcal{J}} \lambda_j d_{ik} z_{ik} x_{jk} \\ \text{subject to } \mathbf{x} \in \mathbf{X}, \mathbf{z} \in \mathbf{Z}(\mathbf{y}), \mathbf{y} \in \mathbf{Y}. \end{aligned} \quad (\text{EC.4})$$

In formulation (EC.4), the set \mathbf{X} can be interpreted as the decision space of an adversary which, after observing the facility locations, realizes each demand point within its uncertainty region at a location that maximizes the distance to the nearest facility. Using the same definitions as above, we can also define the *robust p -center problem*, which minimizes the worst-case *maximum* distance between a demand point and its nearest facility:

$$\begin{aligned} \min_{\mathbf{y}, \mathbf{z}} \max_{\mathbf{x}} \min_t t \\ \text{subject to } t \geq \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}_j} d_{ik} z_{ik} x_{jk}, \quad j \in \mathcal{J}, \\ \mathbf{x} \in \mathbf{X}, \mathbf{z} \in \mathbf{Z}(\mathbf{y}), \mathbf{y} \in \mathbf{Y}. \end{aligned} \quad (\text{EC.5})$$

We now show that the robust p -median and p -center problems are unified by formulation R-AED.

PROPOSITION EC.1. *Assume that each scenario k belongs to exactly one of the sets $\mathcal{K}_1, \dots, \mathcal{K}_{|\mathcal{J}|}$, and $\lambda_1 = \dots = \lambda_{|\mathcal{J}|} = 1/n$. If $\beta = 0$, then formulation R-AED is equivalent to the robust p -median problem (EC.4) with equally weighted demand points. If $\beta \geq (n-1)/n$, then formulation R-AED is equivalent to the non-weighted robust p -center problem (EC.5).*

Proof of Proposition EC.1.

Since each scenario k belongs to exactly one set \mathcal{K}_j , in the context of formulation R-AED this implies that the uncertainty regions are disjoint. The set \mathcal{U} is then separable in j , meaning the uncertainty

in the discrete distribution \mathbf{u} can be interpreted equivalently as uncertainty in $|\mathcal{J}|$ *conditional* distributions of \mathbf{u} – one for each uncertainty region. We may now interpret x_j^k as the unknown probability that the next cardiac arrest arrives at location k , conditioned on the event $\{\xi \in \Xi_j\}$. Note that due to our choice of the uncertainty set \mathcal{U} , each vector \mathbf{x}_j is only known to reside in the $|\mathcal{K}_j - 1|$ -dimensional probability simplex, which is given precisely by the sets $\mathbf{X}_1, \dots, \mathbf{X}_{|\mathcal{J}|}$. Also, since each scenario k only belongs to one of the sets $\mathcal{K}_1, \dots, \mathcal{K}_{|\mathcal{J}|}$, it follows that $u_k = \sum_{j \in \mathcal{J}} \lambda_j x_{jk}$, for each $k \in \mathcal{K}$. By substituting $\sum_{j \in \mathcal{J}} \lambda_j x_{jk}$ for u_k , we can write formulation R-AED equivalently as

$$\begin{aligned} \min_{\mathbf{y}, \mathbf{z}} \max_{\mathbf{x}} \min_{\alpha} \quad & \alpha + \frac{1}{(1-\beta)} \sum_{j \in \mathcal{J}} \sum_{k \in \mathcal{K}_j} \lambda_j x_{jk} \max \left\{ \sum_{i \in \mathcal{I}} d_{ik} z_{ik} - \alpha, 0 \right\} \\ \text{subject to } \quad & \mathbf{x} \in \mathbf{X}, \mathbf{z} \in \mathbf{Z}(\mathbf{y}), \mathbf{y} \in \mathbf{Y}, \\ & \alpha \geq 0. \end{aligned} \tag{EC.6}$$

Due to the separability of the scenario sets in j , we may take the x_{jk} terms inside the max operator in the objective to obtain another reformulation of R-AED:

$$\begin{aligned} \min_{\mathbf{y}, \mathbf{z}} \max_{\mathbf{x}} \min_{\alpha} \quad & \alpha + \frac{1}{(1-\beta)} \sum_{j \in \mathcal{J}} \lambda_j \max \left\{ \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}_j} d_{ik} z_{ik} x_{jk} - \alpha, 0 \right\} \\ \text{subject to } \quad & \mathbf{x} \in \mathbf{X}, \mathbf{z} \in \mathbf{Z}(\mathbf{y}), \mathbf{y} \in \mathbf{Y}, \\ & \alpha \geq 0. \end{aligned} \tag{EC.7}$$

We now show equivalence of (EC.7) and the robust p -median problem. Let $(\mathbf{y}^*, \mathbf{z}^*, \mathbf{x}^*)$ be an optimal solution for formulation (EC.7). Let $Z(\alpha)$ be the objective value of (EC.7) for the solution $(\mathbf{y}^*, \mathbf{z}^*, \mathbf{x}^*, \alpha)$. Since the feasible sets of (EC.7) and (EC.4) are identical, all that remains is to show that their optimal values are equal. For conciseness define $g(\mathbf{x}_j, \mathbf{z}_j) := \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}_j} d_{ik} z_{ik} x_{jk}$. We first observe that for $\beta = 0$ and any $\alpha \leq \min_{j \in \mathcal{J}} g(\mathbf{x}_j^*, \mathbf{z}_j^*)$, the objective value of (EC.7),

$$\begin{aligned} Z(\alpha) &= \alpha + \sum_{j \in \mathcal{J}} \lambda_j (g(\mathbf{x}_j^*, \mathbf{z}_j^*) - \alpha) \\ &= \alpha + \sum_{j \in \mathcal{J}} \lambda_j g(\mathbf{x}_j^*, \mathbf{z}_j^*) - \frac{1}{n} (n\alpha) \\ &= \sum_{j \in \mathcal{J}} \lambda_j g(\mathbf{x}_j^*, \mathbf{z}_j^*) \\ &= \sum_{i \in \mathcal{I}} \lambda_j \sum_{j \in \mathcal{J}} \sum_{k \in \mathcal{K}_j} z_{ik}^* d_{ik} x_{jk}^*, \end{aligned}$$

is equivalent to the objective value of (EC.4). Note that in the second equality above we use the fact that $\lambda_j = 1/n$ for all $j \in \mathcal{J}$. Thus to show equivalence in the optimal values of formulations

R-AED and (EC.4) when $\beta = 0$, it suffices to show that an optimal α^* satisfies $\alpha^* \leq \min_{j \in \mathcal{J}} g(\mathbf{x}_j^*, \mathbf{z}_j^*)$. Suppose that $\alpha^* > \min_{j \in \mathcal{J}} g(\mathbf{x}_j^*, \mathbf{z}_j^*)$. Then,

$$\begin{aligned} Z(\alpha^*) &= \alpha^* + \sum_{j \in \mathcal{J}} \lambda_j \max \{g(\mathbf{x}_j^*, \mathbf{z}_j^*) - \alpha^*, 0\} \\ &> \alpha^* + \sum_{j \in \mathcal{J}} \lambda_j (g(\mathbf{x}_j^*, \mathbf{z}_j^*) - \alpha^*) \\ &= \sum_{j \in \mathcal{J}} \lambda_j g(\mathbf{x}_j^*, \mathbf{z}_j^*) \\ &= \sum_{i \in \mathcal{I}} \sum_{j \in \mathcal{J}} \sum_{k \in \mathcal{K}_j} \lambda_j z_{ik}^* d_{ik} x_{jk}^*, \end{aligned}$$

which contradicts the optimality of α^* . For the robust p -center problem, let $(\mathbf{y}^*, \mathbf{z}^*, \mathbf{x}^*)$ again be an optimal solution to formulation (EC.7). Let $Z(\alpha)$ be the objective value of (EC.7) for the solution $(\mathbf{y}^*, \mathbf{z}^*, \mathbf{x}^*, \alpha)$. If $\alpha = \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\}$, then the objective of formulation (EC.7),

$$\begin{aligned} Z(\alpha) &= \alpha + \frac{1}{(1-\beta)} \sum_{j \in \mathcal{J}} \lambda_j \max \{g(\mathbf{x}_j^*, \mathbf{z}_j^*) - \alpha, 0\} \\ &= \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\} + \frac{1}{(1-\beta)} \sum_{j \in \mathcal{J}} \lambda_j \max \left\{ g(\mathbf{x}_j^*, \mathbf{z}_j^*) - \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\}, 0 \right\} \\ &= \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\} \\ &= \max_{j \in \mathcal{J}} \left\{ \sum_{i \in \mathcal{I}} \sum_{k \in \mathcal{K}_j} z_{ik}^* d_{ik} x_{jk}^* \right\}, \end{aligned}$$

is equivalent to the objective of (EC.5). We now wish to show that an optimal α^* satisfies $\alpha^* = \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\}$ for $\beta > (n-1)/n$. We do so by showing that neither $\alpha^* > \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\}$ nor $\alpha^* < \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\}$ can hold. First, suppose $\alpha^* > \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\}$. Then,

$$\begin{aligned} Z(\alpha^*) &= \alpha^* + \frac{1}{(1-\beta)} \sum_{j \in \mathcal{J}} \lambda_j \max \{g(\mathbf{x}_j^*, \mathbf{z}_j^*) - \alpha^*, 0\} \\ &= \alpha^* \\ &> \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\}, \end{aligned}$$

which contradicts the optimality of α^* . Now suppose $\alpha^* < \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\}$, and let $\delta = \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\} - \alpha^*$. It follows that $\sum_{j \in \mathcal{J}} \max \{g(\mathbf{x}_j^*, \mathbf{z}_j^*) - \alpha^*, 0\} \geq \delta$. Since $\lambda_j = 1/n$ for all $j \in \mathcal{J}$, we can write

$$\begin{aligned} Z(\alpha^*) &= \alpha^* + \frac{1}{(1-\beta)n} \sum_{j \in \mathcal{J}} \max \{g(\mathbf{x}_j^*, \mathbf{z}_j^*) - \alpha^*, 0\} \\ &\geq \alpha^* + \frac{1}{(1-\beta)n} \delta \end{aligned}$$

$$\begin{aligned}
&= \alpha^* + \delta + \frac{1}{(1-\beta)n} \delta - \delta \\
&= \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\} + \left(\frac{1}{(1-\beta)n} - 1 \right) \delta \\
&> \max_{j \in \mathcal{J}} \{g(\mathbf{x}_j^*, \mathbf{z}_j^*)\},
\end{aligned}$$

where the final strict inequality follows from $\beta > (n-1)/n$. Again, this contradicts the optimality of α^* . \square


Based on Proposition EC.1, we can interpret formulation R-AED as subsuming a spectrum of location models with the robust p -median problem at one end ($\beta = 0$) and the robust p -center problem at the other (β close to 1). While the specialization of CVaR to the mean and maximum of a general discrete distribution is well known, this explicit unification of the p -median and p -center problems via a CVaR objective does not appear to have been discussed previously in the facility location literature. We note also that the assumption in Proposition EC.1 that each k belong to exactly one \mathcal{K}_j is needed to ensure the equivalence between the main formulation R-AED and the intermediate model (EC.6) used in the proof. If this assumption is relaxed, then formulation (EC.6) still serves as a unifying model for the p -median and p -center problems, although the equivalence with R-AED no longer holds. The *ordered median problem* (Kalcsics et al. 2002, Nickel and Puerto 1999) also unifies median and center problems, but requires an ordering of the demands according to their distance to the sited facilities so that a set of appropriate parameters can be identified. By contrast, our approach involves the adjustment of a single parameter and requires no such ordering.

To the best of our knowledge, (EC.4) and (EC.5) are the first robust formulations of the general p -median and p -center problem with non-stochastic uncertainty in the demand location (i.e., in the spirit of robust optimization rather than scenario-based stochastic optimization). Previous work in robust facility location has considered only one facility (Averbakh and Berman 2000, Averbakh 2003) or uncertainty in demand level rather than location (Baron et al. 2011). As discussed in Section 3.2, the formulation in (Baron et al. 2011) is well equipped to handle uncertainty in the level of demand at each of a set of known demand points. The robust p -median and p -center formulations presented above may be more appropriate in applications where there is continuous uncertainty in future demand locations (e.g., emergency response or peer-to-peer ridesharing applications).


EC.3. Visualization of AED deployments

Figure EC.1 depicts an instance of the AED deployments produced by the nominal (SAA), robust and ex-post models, along with a set of simulated cardiac arrests (corresponding to the ex-post


deployment) and the 43 historical cardiac arrest locations. Note that in this instance, the nominal model produces small “clusters” of AEDs, whereas the robust and ex-post models appear to position AEDs more efficiently with respect to the simulated cardiac arrests. This behavior is expected, since the nominal model only optimizes the tail of the empirical distance distribution, whereas the robust model accounts for potential deviation of the simulated cardiac arrests from the historical locations. Note also that the ex-post model achieves the most effective deployment of AEDs with respect to the simulated cardiac arrests, since it optimizes directly with respect to the simulated data.



(a) Nominal.



(b) Robust.



(c) Ex-post.

Figure EC.1 Example of nominal, robust and ex-post AED deployments ($P = 30$, $h = 100$, $\beta = 0.9$).

EC.4. Row-and-column generation example: Total scenarios and optimality gap

The plot below depicts an example of the total number of scenarios generated in the master problem R-AED-MP and the associated optimality gap at each iteration of the row-and-column generation algorithm. The plot corresponds to the instance given in the second-last row of Table 1 ($\mathcal{I} = 100$, $\mathcal{J} = 20$, $P = 20$, $\mathcal{K} = 24,743$).




Figure EC.2 Example of total number of scenarios generated (blue squares) and optimality gap (green circles) for each iteration of the row-and-column generation algorithm ($|\mathcal{I}| = 100$, $|\mathcal{J}| = 20$, $P = 20$, $|\mathcal{K}| = 24,743$).

