VRIJE UNIVERSITEIT AMSTERDAM

Master Thesis

Improving fairness in AED placement models

by

Joeri Kreuk

(2743652)

Supervisor:	Caroline Jagtenberg
Co-supervisor:	Brian Overbeek
Second reader:	Sebastian Kraul

June 29, 2026

Submitted in partial fulfillment of the requirements for
the VU degree of Master of Science in Econometrics and Operations Research

Improving fairness in AED placement models

Joeri Kreuk
Vrije Universiteit Amsterdam
Amsterdam, The Netherlands
j.l.kreuk@student.vu.nl

ABSTRACT

Automated external defibrillators (AEDs) have the potential to save lives, but their placement is strongly correlated with their effectiveness and can often be significantly improved. In New Zealand, Māori and Pacific populations bear a disproportionately high out-of-hospital cardiac arrest (OHCA) burden while receiving comparatively lower AED coverage. This study develops a general fairness-aware AED placement optimization model implemented in Auckland that integrates both ethnic and socioeconomic dimensions of inequality using the Atkinson social welfare framework. Using meshblock-level OHCA demand estimates and a probabilistic coverage function, the model is solved as a mixed-integer program and evaluated through ten-fold cross-validation across four urban zones and multiple budget scenarios. The results show that meaningful reductions in coverage inequality can be achieved at modest efficiency cost in dense urban areas, but that simply expanding the AED network without explicit equity constraints does not reduce, and may worsen, existing disparities. The proposed framework gives policymakers a model for navigating the trade-off between aggregate coverage and equitable access in PAD program design.

1 INTRODUCTION

Out-of-hospital cardiac arrest (OHCA) is a leading cause of mortality worldwide, with survival critically dependent on the speed of intervention. Early cardiopulmonary resuscitation (CPR) and rapid defibrillation with an automated external defibrillator (AED) are the most effective pre-hospital interventions, capable of improving survival rates by 50–75% when applied in time [11, 37, 52]. Public access defibrillation (PAD) programs were introduced to bring AEDs closer to the general public, yet their real-world impact has remained limited. Despite the clinical efficacy of early defibrillation, the proportion of OHCA victims receiving bystander AED defibrillation before emergency medical services (EMS) arrival remains small [16]. A key driver of this gap is the spatial misalignment between where AEDs are placed and where cardiac arrests actually occur [13].

The strategic placement of AEDs has therefore emerged as a critical research and policy challenge. Optimization-based approaches, drawing on classical facility location models such as the maximal covering location problem (MCLP), have demonstrated that data-driven AED deployment can substantially improve OHCA coverage compared to heuristic strategies [12, 14]. However, most existing models optimize for coverage alone, without explicitly accounting for whether coverage improvements are distributed equitably across the population. In practice, this can lead to AED placements that reinforce existing disparities: communities that already have better access to health resources may benefit disproportionately, while those with the highest OHCA incidence rates remain underserved.

These disparities are particularly pronounced in New Zealand, where national registry data reveal significant ethnic and socioeconomic inequities in OHCA incidence, bystander intervention rates, and survival outcomes [22]. Māori and Pacific Peoples experience markedly higher rates of cardiac arrest and lower rates of bystander defibrillation compared to the European/other population, while the existing distribution of publicly accessible AEDs does not reflect this pattern of need. The most deprived communities, which have the highest OHCA incidence rates, currently have fewer AEDs and lower AED-to-OHCA ratios than less deprived areas.

Addressing these inequities requires moving beyond coverage maximization as the sole objective. Fairness-aware optimization frameworks from operations research, such as the Atkinson social welfare function and the related α-fairness framework of Bertsimas et al. [2, 8], provide a principled and mathematically grounded approach to embedding equity directly into optimization problems. By parameterizing the degree of inequality aversion, these frameworks allow for explicit control over the trade-off between total coverage and equitable access across population subgroups. Despite their potential relevance, no prior AED placement study has directly incorporated fairness as an explicit optimization objective.

This study addresses this gap by developing a general equity-aware AED placement optimization model, which is applied to Auckland, New Zealand's largest metropolitan area as a case study. Methodologically, this work extends established facility location models for AED deployment by incorporating multiple, context-specific equity factors within existing mathematical and fairness-aware optimization frameworks. Empirically, it examines how prioritizing fairness influences overall system efficiency by analyzing the trade-offs between equitable access to AEDs and total OHCA coverage under different levels of inequality aversion. This allows us to quantify how strongly fairness considerations affect the performance of AED placement strategies in practice.

From a policy perspective, this study does not aim to prescribe specific AED placement decisions. Rather, it seeks to provide an analytical foundation that demonstrates how equity considerations can be incorporated into optimization-based AED placement strategies. By exploring the relationship between fairness and efficiency, the study aims to contribute to an emerging perspective in which equitable access is treated as an explicit design objective rather than a secondary evaluation criterion. In doing so, it aims to serve as a stepping stone for future research and policy discussions on equitable AED deployment, illustrating that improving access for historically underserved populations does not necessarily require substantial sacrifices in overall coverage.

The remainder of this paper is structured as follows. Section 2 reviews the relevant literature on OHCA disparities, existing AED placement models, and fairness-aware optimization frameworks.

Section 3 describes the data, the OHCA incidence estimation procedure, the probabilistic coverage model, and the equity-aware optimization formulation. Section 4 presents the empirical results across urban areas and budget scenarios. Lastly, Section 5 discusses the findings, limitations, and directions for future research.

2 LITERATURE REVIEW

The time-sensitivity of OHCA survival has been extensively documented. Without intervention, the probability of a favourable neurological outcome declines by approximately 10% per minute [5, 43]. This urgency underpins the clinical rationale for PAD programs and motivates the body of research on optimal AED placement reviewed in this section.

2.1 PAD programs

A public access defibrillation (PAD) program is a public health initiative originally initialized by the American Heart Association designed to improve survival from sudden cardiac arrest by strategically placing AEDs in communities and training lay rescuers to use them before emergency medical services arrive [56]. More PAD initiatives have since then launched globally to improve AED availability and survival rates in OHCAs. Despite the demonstrated effectiveness of early defibrillation, PAD programs have achieved only limited impact in practice. Although PAD initiatives substantially improve survival rates following OHCA when applied [4], they are rarely used, with only a small proportion of victims receiving AED defibrillation before EMS arrival [16]. This gap between clinical efficacy and real-world effectiveness highlights the importance of not only AED availability, but also their spatial accessibility at the time and location of cardiac arrest.

Several factors contribute to the low utilization of public AEDs, including limited public awareness, bystander hesitancy, and challenges in locating nearby devices during emergencies. Importantly, an additional and potentially modifiable factor is the lack of systematic, data-driven guidance in determining where AEDs should be placed [12]. Namely, even otherwise well-organized PAD programs may fail to deliver timely defibrillation if AEDs are not located close to cardiac arrest events. Consequently, the strategic placement of AEDs emerges as a critical determinant of whether their life-saving potential can be realized at the population level.

2.2 Disparities in OHCA and Access to AEDs

Even though AEDs can be life-saving to any individual, disparities in OHCA incidence, access to early defibrillation, and survival outcomes have been observed across population groups and geographic regions in many countries. A recent systematic review by Van Nieuwenhuizen et al. (2019) shows that socioeconomic and demographic factors consistently influence OHCA outcomes and access to public-access defibrillation across many different countries and health systems [54]. Within this broader international context, national registry data from New Zealand reveal substantial differences in both the incidence of OHCA and survival outcomes across ethnic and geographic groups, highlighting important inequities in pre-hospital care systems [22].

Using data from the St. John New Zealand OHCA registry, Dicker et al. (2019) provided a national overview of OHCA incidence and outcomes for approximately 90% of the New Zealand population. The study reported an overall age-adjusted incidence of 105.0 OHCA events per 100,000 person-years during the study period. Among events where resuscitation was attempted by EMS, 30-day survival was 16.7%. However, these aggregate figures mask pronounced disparities. Compared to the European/Others group, Māori and Pacific Peoples experienced significantly higher age-adjusted incidence rates (Māori: 144.4; Pacific Peoples: 113.5; European/Others: 93.8 per 100,000 person-years). Both Māori (adjusted OR 0.61, 95% CI 0.48–0.78, $p < 0.001$) and Pacific Peoples (adjusted OR 0.52, 95% CI 0.37–0.72, $p < 0.001$) had significantly lower adjusted odds of survival to 30 days compared to European/Others, even after adjustment for key clinical and situational factors. Moreover, there was a significantly lower usage rate of defibrillators in Māori communities (8.2%) compared to European/Others communities (11.0%).

Importantly, the majority of OHCAs in New Zealand do not occur in remote or sparsely populated areas but rather in relatively crowded urban and residential environments, where timely defibrillation should, in principle, be feasible if AEDs are appropriately located. Evidence from the same St. John New Zealand OHCA registry shows that around 73% of OHCAs occur in urban settings, with a substantial proportion taking place in metropolitan areas such as Auckland [18]. Similarly, Dicker et al. reported that the majority of resuscitation-attempted OHCAs occurred in residential locations, rather than public or rural settings, underscoring that population density alone does not guarantee access to early defibrillation [22].

However, the predominance of OHCAs in residential settings does introduce another challenge. Many cardiac arrests occur inside private homes, where the presence of bystanders capable of retrieving and using an AED may be limited. Moreover, the response time of EMS often rivals bystander intervention, with defibrillation occurring before EMS arrival in only 8% of OHCAs [18]. This implies that even optimal AED placement cannot guarantee timely defibrillation in all cases.

The observed ethnic disparities in OHCA incidence, survival, and community defibrillation rates are not merely coincidental, but are closely linked to broader social determinants of health, including socioeconomic status, neighborhood characteristics, and access to health-promoting resources. In New Zealand, Māori and Pacific Peoples are disproportionately represented in areas of higher socioeconomic deprivation, a pattern that has been consistently associated with elevated cardiovascular risk and poorer health outcomes [22]. Both international and national evidence indicates that neighborhoods with higher levels of socioeconomic deprivation exhibit systematically lower rates of bystander intervention, including CPR and public-access defibrillation, even after accounting for differences in OHCA incidence [42, 50, 54].

The OHCA incidence rate as reported by the New Zealand national OHCA registry (Figure 7 [21]) has shown that the incidence rate of OHCAs is unevenly distributed across deprivation quintiles in New Zealand; with the most deprived quintile (Q5) having an incidence rate that is more than two times higher than the incidence rate of the least deprived quintile (Q1). Critically, the current distribution of publicly accessible AEDs does not reflect this pattern of need, but rather exhibits an inverse relationship. As Table 7 demonstrates, while OHCA incidence increases with deprivation, the number of existing AEDs in these deprivated areas does not

grow proportionally. The most deprived quintile, with the highest OHCA burden, actually possesses fewer AEDs (613) than Q3 and Q4 (635 and 677 respectively), resulting in a substantially smaller AED-to-OHCA ratio in the meshblocks/communities that need them most. Table 8 reinforces this disparity from an ethnic perspective, revealing that the rate of AEDs per 100,000 persons is lowest for the Māori population (134.9), which controversially bears the highest OHCA burden, compared to the European/other group (162.8). This empirical evidence of misalignment between high OHCA burden among deprived areas and Māori and Pacific populations, and the current inequitable distribution of AEDs, provides the foundational motivation for this study. While a purely utilitarian optimization approach is known to improve overall coverage in high-burden areas such as the higher quintiles [30], it does not inherently guarantee an equitable distribution. Our aim is to move beyond using equity merely as a post-hoc metric of success, and instead integrate both equity and coverage into a formal optimization objective. This allows us to investigate allocation strategies that maximize the number of OHCA events covered while simultaneously ensuring equitable access.

This question of how to balance coverage and equity in AED placement emerges from a broader recognition within the literature that significant disparities in public access defibrillation exist. In New Zealand, most OHCAs occur in residential areas where timely defibrillation is theoretically possible through bystander intervention. However, studies have consistently shown that communities with higher levels of socioeconomic deprivation and larger Māori and Pacific populations experience higher rates of cardiac arrest but lower rates of bystander defibrillation and, often, poorer AED coverage [20, 22]. These disparities in access and outcomes suggest that improving overall survival will require more than just increasing the total number of AEDs. Rather, future AED placement strategies may need to be designed to directly address these spatial and social inequities.

2.3 Current models

In response to the limited real-world effectiveness of these PAD programs, a growing body of research has sought to identify strategies for improving the placement of AEDs. Early efforts to guide public AED deployment relied primarily on heuristic or descriptive approaches rather than formal optimization frameworks.

One common strategy was using historical OHCA data to identify high-risk locations or recurrent cardiac arrest sites, under the implicit assumption that future events would follow similar spatial patterns [51]. By concentrating AEDs in areas with previously observed clusters of cardiac arrests, these approaches aimed to increase the likelihood that an AED would be located near a future OHCA event.

Similar studies focused on deploying AEDs in locations characterized by high foot traffic, such as casinos, airports, and sports venues, emphasizing accessibility and visibility rather than systematic coverage analysis [11, 40, 52]. While such placements were intuitive and logistically convenient, they largely overlooked the substantial proportion of OHCAs occurring in residential settings and the fact that identifying such high-risk locations is not generalizable and very much dependent on the studied cities [14, 57].

Additional work employed fixed distance or time thresholds to guide AED deployment, aiming to ensure availability within a predefined response, such as a three-minute walking distance or a maximum distance of 100 meters [28, 55]. Although these thresholds were supported medically and the observed survival does indeed decline over time, they did not provide a mechanism for prioritizing among multiple candidate locations or for systematically balancing competing placement objectives like budget and coverage.

While these heuristic approaches represented an important step toward data-informed AED placement, they are inherently limited. They often overrely on "hotspots", are unable to combine multiple objectives such as coverage, equity and accessibility, and often interact poorly with operational constraints [10, 32].

Therefore, optimization has been used as a central tool in modeling these location problems, including the application in AED placement. Classical models such as the set covering location problem (SCLP) [39] and the MCLP [15] have been widely used to determine facility locations under resource constraints. In the context of AED placement, these models aim to maximize the coverage of historical OHCA locations within a predefined distance or response time threshold and lay a foundation for all AED placement models to come.

Building on this, literature has proposed extensions to these models that better capture real-world scenarios. These extensions suggest probabilistic service availability, multiple coverage requirements, and distance-dependent effectiveness to reflect declining response performance with increasing travel time [23, 26, 29]. Such extensions are particularly relevant for AED placement, as the probability of successful defibrillation decreases sharply with delays in retrieval and use.

These optimization models have increasingly been applied to public AED deployment. Early applications focused on specific or constrained environments, such as transport systems, indoor facilities, or institutional settings, demonstrating the feasibility of optimization-based AED placement in practice [17, 32]. Subsequent studies showed that mathematically optimized AED placements can substantially improve OHCA coverage compared to population-based strategies and heuristics [12].

More recent work has highlighted the importance of accounting for uncertainty in future cardiac arrest locations and variability in bystander response. Robust and stochastic optimization frameworks have been developed to address spatial uncertainty in demand, while probabilistic coverage models allow multiple AEDs to contribute to the coverage of a single event. These approaches reflect more realistic emergency response behavior and demonstrate that mathematical optimization is a flexible and powerful method for guiding AED deployment by modeling the environment as a facility location problem [14, 46, 49].

Most optimization models in the AED placement literature aim to maximize coverage, typically identifying locations corresponding to previous OHCA events or “OHCA hotspots” [12, 14, 30, 32, 47, 48]. While this strategy improves overall coverage, it may inadvertently undervalue equity across communities: by prioritizing areas with historically high incident rates, these approaches provide comparatively less access to communities with fewer recorded events.

Only a limited number of studies have explicitly addressed equity in AED access. Consequently, purely coverage-focused optimization

may disproportionately favor areas with higher population density or historical OHCA incidence, potentially reinforcing existing disparities in access to life-saving interventions. In the New Zealand context, it is worth noting that a purely utilitarian approach would improve AED access in the most deprived quintiles and thereby reduce the current disparity between high and low deprivation areas [30]. However, while such an approach may improve the balance between OHCA incidences and AED accessibility, it does not explicitly incorporate equity as an objective. Rather, it achieves an AED distribution corresponding to the incidence rate, without guaranteeing that the resulting allocation is equitable across all communities, particularly from an ethnic or cultural perspective.

Equity considerations have, however, been incorporated in broader facility location literature through multi-objective formulations and fairness-centric objective functions. Marsh and Schilling (1994) documented a range of equity measures in facility siting models, emphasizing both the importance and complexity of integrating fairness into optimization [33]. More recent work has proposed frameworks that explicitly prioritize fairness, for example by minimizing disparities between demand points or adopting max–min fairness principles [36, 38]. Application-oriented studies in healthcare facility location further illustrate how efficiency and equity can be balanced, for example using inequality measures such as the Gini index or by explicitly optimizing accessibility for vulnerable populations [9, 24]. Despite these advances, the explicit incorporation of social determinants such as ethnicity and socioeconomic status remains limited in emergency response applications, including AED placement.

Leung et al. (2021) represent a notable exception within AED placement literature, treating equity as an evaluation metric of their model [30]. Their results demonstrate that meaningful improvements in access for socioeconomically deprived areas can be achieved while still showing significant improvements in overall coverage. However, their model did not explicitly optimize for equity, meaning that any improvements in fairness were incidental rather than intentional. This underscores that the trade-off between efficiency and equity remains an important topic for investigation. From a public health and policy perspective, such equity considerations play a critical role because unequal access to AEDs may reinforce existing geographic or demographic disparities in survival outcomes. Communities with fewer historical OHCAs may still face substantial risk, and ensuring that these populations are not systematically underserved is an important societal objective. Consequently, decision makers may wish to consider allocation strategies that balance overall coverage with a more equitable distribution of access.

Fairness-aware optimization frameworks from operations research provide a systematic method to formalize this trade-off. Atkinson [2] introduced an inequality aversion parameter $\varepsilon \geq 0$, where higher values indicate greater concern for reducing inequality. In this framework, resource allocation decisions are made by maximizing a social welfare function that aggregates the utility received by different groups, which in this study translates to OHCA coverage by AEDs. The welfare function is parameterized by a single inequality-aversion parameter ε, which determines how strongly disparities in utility across groups are penalized. Formally, the framework maximizes a constant-elasticity welfare function of the utilities of each group, where the curvature of the function increases with ε. This concavity implies diminishing marginal welfare gains for groups that already receive high utility, meaning that improvements for disadvantaged groups contribute more to the overall objective than equal improvements for already well-served groups.

By adjusting ε, the decision maker can explicitly control the balance between efficiency and equity. When $\varepsilon = 0$, the objective reduces to the utilitarian framework that maximizes total coverage without regard to distribution. When $\varepsilon = 1$, the solution corresponds to proportional fairness (the geometric mean), while as $\varepsilon \to \infty$ the allocation converges to a max-min fairness solution that prioritizes improving outcomes for the worst-off group.

This framework has proven to be very influential and was later adapted by Bertsimas et al. [8] in their work on the α-fairness framework for resource allocation. While Atkinson's original formulation focused on measuring inequality in social welfare contexts, Bertsimas et al. embedded the same mathematical structure directly into optimization problems with explicit resource constraints, coining the term α-fairness where α plays an identical role to Atkinson's ε [7].

Together, these inequality aversion frameworks provide interpretable models to explore the trade-off between efficiency and equity in resource allocation problems. In the context of AED placement optimization, these frameworks are particularly valuable because it enables policymakers to evaluate how prioritizing equitable access across different population subgroups affects overall OHCA coverage, and select an allocation that aligns with societal preferences for health inequality reduction.

2.4 Contributions

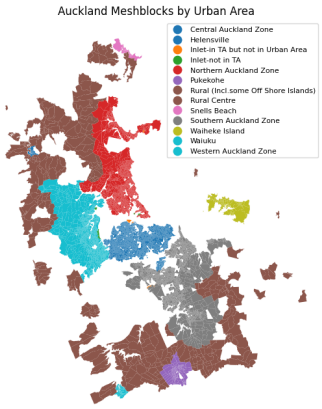
This study contributes to the literature on PAD programs and facility location optimization by developing a general AED placement model that can incorporate equity considerations, which in this application are specified in terms of ethnic and socioeconomic disparities in OHCA risk and access to early defibrillation. While prior optimization-based approaches have demonstrated the effectiveness of data-driven AED deployment, they have largely focused on maximizing coverage or efficiency, with limited attention to equity across population subgroups. Existing fairness- or equity-aware models typically incorporate socioeconomic deprivation but do not consider ethnicity as a distinct and potentially interacting factor [30]. Moreover, these models often treat equity as an evaluation metric rather than embedding it directly into the optimization, leaving the trade-offs between overall coverage and fairness largely unexplored.

Using Auckland as a case study, where robust OHCA registry data show significant ethnic differences in incidence, bystander intervention, and survival outcomes, this research develops a new optimization model that integrates both socioeconomic and ethnic landscapes into AED placement decisions. To operationalize equity within the optimization, we adopt the Atkinson (equivalently, α-fairness) social welfare framework [2, 8] with an inequality aversion parameter $\varepsilon \geq 0$, which allows for explicit control over the trade-off between efficiency (total coverage) and equity (equitable access across population groups). By embedding these factors directly into the objective function, the proposed approach systematically

balances overall OHCA coverage with equitable AED access for high-risk and historically underserved populations. This study extends this framework by incorporating group-based weights to account for heterogeneous group sizes and the statistical precision of coverage estimates, ensuring that the welfare function maintains an individual-level interpretation despite aggregation to ethnicity–deprivation groups.

Concluding, this study makes three specific contributions. First, we extend the Atkinson-α-fairness social welfare framework to the context of AED placement, demonstrating how inequality aversion can be embedded directly into facility location optimization. Second, we develop a population-weighted formulation that accounts for both population size and estimation uncertainty across groups, addressing a methodological gap in group-level equity analysis. Third, we develop and apply an equity-aware optimization framework to Auckland, using local OHCA incidence data, population demographics, and existing AED infrastructure to guide placement decisions, thereby providing policymakers with actionable insights into the balance between efficiency and equity in PAD programs.

Figure 1: Urban areas in Auckland



3 METHODOLOGY

This section presents the methods used in this study. It begins by describing the data on Aucklands demographics, OHCA incidence, and existing AED locations. We then introduce an innovative optimization model that integrates probabilistic coverage and equity considerations building on concepts from previous studies. Finally, we outline the statistical validation of the model and the metrics used to evaluate the resulting efficiency-equity trade-off.

3.1 Study setting

This study was conducted at meshblock level using spatial and demographic data obtained from Statistics New Zealand. Meshblocks are the smallest geographic units at which census data are collected nationwide in New Zealand and therefore provide the highest available spatial resolution for population-level analysis. Meshblocks typically contain between 60 and 120 residents, with an average of 94 residents per meshblock.

The analysis will only focus on Auckland, New Zealand's largest metropolitan area, which contains approximately a third of the country's population. More precisely, the study will focus on Central-, Northern-, Southern- and Western Auckland zone, which are Urban Areas in Auckland shown in Figure 1. These zones will be abbreviated as CAZ, NAZ, SAZ, and WAZ, respectively. This decision is motivated by several considerations. First, the region accounts for a substantial proportion of national AED deployments, resulting in a higher density of existing devices compared to rural areas. Second, these urban areas have a high population density and geographic concentration of cardiac arrest events, which makes them suitable for evaluating spatial accessibility and equity in AED placement. Analyzing New Zealand as a whole would result in an unnecessarily large optimization model with over 40,000 meshblocks, many of which are sparsely populated rural areas with few AEDs and few OHCA incidences. By narrowing the scope to Auckland, the analysis can focus on the region where AED placement decisions have the greatest potential impact on population health outcomes.

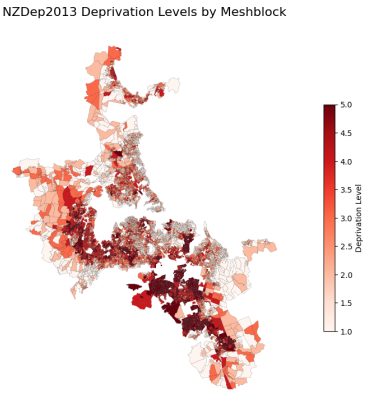
Meshblocks serve as demand nodes in such analyses when we look at it from a facility location problem perspective. Each meshblock represents a geographically defined population cluster to which potential cardiac arrest demand is assigned. Using the smallest available census unit ensures that demand is represented at the highest possible spatial resolution.

If demand were aggregated at a coarser geographic level, such as territorial authorities or larger census areas, local variations in population density and socioeconomic composition would be obscured and could lead to oversimplification and errors that are not negligible [59]. By modeling each meshblock as a separate demand point, the optimization model can more accurately capture how local differences influence optimal AED placement.

This study will use data on the ethnic, demographic and socioeconomic characterstics for each meshblock and were gathered from the 2013 Census and provided by Statistics New Zealand. For this research, data from 2013 will be sufficient, as previous studies have shown that the incidence and geographic distribution of OHCAs remain stable over time [14, 34, 41, 58]. Moreover, the estimates of ethnic proportions in Auckland remains relatively stable over years [45]. This temporal stability implies that patterns observed in the 2013 census data are unlikely to differ substantially from current incidence patterns. Consequently, findings based on this data are likely applicable to present-day decision-making in PAD programs and other organized AED placement efforts.

Socioeconomic status was measured using the NZDep2013 deprivation index, which classifies meshblocks into an ordinal scale ranging from 1 (least deprived) to 10 (most deprived) [3]. The deprivation index combines nine aspects from the 2013 census data including income, employment, housing conditions, and access to transport and communication to reflect eight dimensions of deprivation. To facilitate comparison with the national health statistics published in New Zealand's annual reports, which present incidence rates by deprivation quintile rather than decile, the NZDep2013 deciles were aggregated into quintiles. A visual representation of meshblocks in Auckland and their corresponding deprivation quintile can be found in Figure 2.

Figure 2: NZDep2013 deprivation quintile per meshblock



In addition to socioeconomic status, ethnicity distributions were obtained at meshblock level, including the number of individuals identifying as European, Māori, Pacific Peoples, Asian, and Middle Eastern/Latin American/African (MELAA). As stated in Dicker et al. (2019) [22], Asian patients made up less than 5% of the OHCA registry data, which is the reason for this group to be included in the European/Others group. Moreover, the MELAA community contains only about 1% of the New Zealand population and are therefore also included in the European/Others cohort. The deprivation index and the ethnic composition of each meshblock enable the evaluation of equity in spatial accessibility of AEDs across socioeconomic and ethnic subpopulations.

3.2 OHCA Incidence Estimation on meshblock level

The primary outcome of this study is the coverage of OHCA incidences across New Zealand. Because individual-level OHCA events are not publicly available at meshblock resolution due to privacy constraints, expected OHCA demand was simulated using demographic data and published incidence rates. Population counts at meshblock level were available separately by age group, ethnicity, and sex, but not jointly for the full age × ethnicity × sex cross-classification. To approximate the number of individuals in each demographic category within meshblock i, proportional allocation was applied. Let $P_{i,a}$ denote the number of residents in age group a, $P_{i,e}$ the number in ethnicity e, $P_{i,s}$ the number in sex category s, and P_i the total population of meshblock i. The estimated number of individuals in category (a, e, s) was approximated as:

$$P_{i,a,e,s} = P_{i,a} \times \frac{P_{i,e}}{P_i} \times \frac{P_{i,s}}{P_i}.$$

This approach assumes independence between age, ethnicity, and sex within each meshblock. Although this assumption may not hold exactly, it allows us to construct the most detailed picture possible from the available aggregated data, while remaining consistent with the observed marginal totals.

Age- and ethnicity-specific OHCA incidence rates (per 100,000 person-years) as well as incidence rates stratisfied by sex and ethnicity were obtained from Dicker et al. (2019) [22]. The sum of these age-specific rates equals the crude incidence rate for each ethnic group. However, sex specific incidence rates were not reported jointly by age and ethnicity. Therefore, to derive incidenre rates for each age × ethnicity × sex category, the age-adjusted male and female incidence rates were used to construct proportional scaling factors. For each ethnicity e, let r_e^{male} and r_e^{female} denote the published age-adjusted incidence rates stratisfied by sex and ethnicity. The sex-specific scaling factors were then calculated by:

$$\mathrm{M}_e = \frac{r_e^{\text{male}}}{r_e^{\text{male}} + r_e^{\text{female}}}, \qquad \mathrm{F}_e = 1 - \mathrm{M}_e$$

These factors redistribute cases proportionally according to the observed male-to-female incidence ratio for all ethnicities. Note that it is assumed that the male-to-female ratio is equal for all age-groups.

To estimate OHCA cases at meshblock level, we multiply the estimated population for each age group × ethnicity × sex category ($P_{i,a,e,s}$) by the corresponding incidence rate. However, preliminary calculations revealed that using the estimated rates directly yielded only one-twelfth of the expected total OHCA incidence. This discrepancy arises because the published rates, when summed across the six age groups and the two sex categories, produce the crude incidence rate. Therefore, to obtain a valid estimation of the incidence rate for small-area estimation, each rate must be multiplied by the number of age groups (6) and the number of sex categories (2). This adjustment ensures that the average rate across all age groups matches the overall crude incidence rate per ethnicity, preserving the total expected OHCA count in our meshblock-level estimates. These rates can be interpreted as the expected annual number of OHCAs per 100,000 individuals in demographic group (a, e, s).

This estimation approach relies on two key assumptions. First, we assume that within each meshblock, age, sex, and ethnicity are independently distributed. Second, we assume that the male-to-female incidence ratio is constant across all age groups within each ethnicity. These assumptions introduce some degree of uncertainty into the estimated incidence rates. However, they are necessary to reconcile the published incidence rates with the aggregated population data available at meshblock level.

Alternative approaches such as assuming uniform incidence across all demographic subgroups or applying national incidence rates would fail to capture the known disparities in OHCA risk across ethnic and socioeconomic groups. Such coarse approximations would effectively smooth over the very heterogeneity we aim to examine. By contrast, our approach preserves the marginal distributions from the literature and introduces these demographic variations into the simulated OHCA events. The objective is not to predict individual OHCA events with perfect accuracy, as cardiac arrests are inherently stochastic, but rather to generate plausible small-area estimates that reflect the underlying heterogeneity across subgroups. This allows us to investigate how different AED placement strategies perform when tested against a heteregeneic distribution of OHCA events.

Let $r_{a,e,s}$ denote these final estimated incidence rate per 100,000 population and $P_{i,a,e,s}$ the estimated population in meshblock i belonging to demographic group (a, e, s). The expected number of annual OHCAs in meshblock i is then given by:

$$D_i = \sum_a \sum_e \sum_s \left(\frac{P_{i,a,e,s}}{100{,}000} \times r_{a,e,s} \right) \times \phi_{q(i)},$$

where $\phi_{q(i)}$ is a deprivation-specific correction factor applied based on the NZDep2013 quintile $q(i)$ of meshblock i. These correction factors were derived by comparing simulated quintile-level incidence rates with published deprivation-specific OHCA rates (Figure 7 [21]) and computing correction factors to align the simulated distribution across different deprivation quintiles with the reported deprivation quintiles. The correction factors can be found in Table 1. This adjustment is essential because OHCA incidence is known to vary systematically with socioeconomic deprivation, with higher rates observed in more deprived areas [20]. By incorporating this deprivation gradient, the model accounts for the residual spatial variation in OHCA risk not captured by demographic composition alone. This ensures that the estimated case counts at the meshblock level reflect not only the age, ethnicity, and sex structure of the local population but also the broader socioeconomic context in which individuals live

Table 1: Deprivation quintile correction factors

Deprivation quintile	Correction factor
Q1	0.5577
Q2	0.8297
Q3	0.9477
Q4	1.1489
Q5	1.4376

The quantity D_i represents the expected annual OHCA demand in meshblock i. Aggregating D_i across meshblocks yields total expected national incidence. Based on the age, sex, ethnicity, and deprivation information incorporated into the model, the total estimated OHCA incidence is 1,333 cases. Given a total population of 1,224,900 this corresponds to an overall incidence rate of 108.8 per 100,000 person-years. This estimate closely aligns with the published incidence rates from the St. John Registry [19], providing validation for the model's underlying assumptions. The difference can be attributed to the application of the deprivation-specific correction factors that systematically adjusts incidence upward or downward depending on the socioeconomic profile of each meshblock.

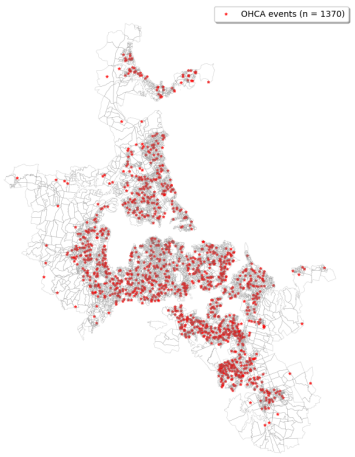
These estimated annual incidence values D_i serve as the input parameters for the spatial OHCA event simulation model. Specifically, each meshblock i is modeled as an independent Poisson process with rate parameter $\lambda_i = D_i$, representing the expected number of OHCA events per year in that meshblock. For a given simulation horizon (e.g., one year), the number of events in meshblock i is sampled from a Poisson distribution with mean λ_i. Once we have simulated how many events occur in a given meshblock, we place them randomly within its boundaries.

The Poisson distribution is a statistically validated model for OHCA events. Research by Skogvoll and Lindqvist (1999) [44] confirms that OHCA events align with the criteria for a Poisson process: the events are independent of one another, the time intervals between them follow an exponential distribution, and their occurrence over a fixed period can be accurately described by the Poisson distribution.

By assuming events are randomly located within each meshblock, we avoid making arbitrary assumptions about exactly where within a neighborhood an OHCA might occur. Since meshblocks are already small geographic areas and we lack more precise location data, this randomly uniform placement within a meshblock is the most neutral approach.

The end result of the OHCA simulation model is a set of events that collectively match the expected totals, reflect the demographic and deprivation patterns in the data, and provide the realistic spatial variation needed to build a model that can optimize AED placement for maximum coverage. Figure 3 shows the OHCA events for a simulation of one year in New Zealand. Even though one simulation year can be sensitive to randomness, the higher density of OHCA events around Auckland and other metropoles in New Zealand is clearly visible in the figure.

Figure 3: Poisson-based event simulation of OHCAs for 1 year in Auckland



Full details of all deprivation quintile characteristics, ethnic group compositions, and the final estimated incidence rates per demographic category are provided in Appendix A.

3.3 Candidate AED locations

National AED registry data were obtained from AED locations NZ ([1]) and cleaned to ensure spatial validity and internal consistency. All locations are validated and projected using the NZTM2000 (EPSG:2193) coordinate system and all records with missing geographic coordinates or coordinates that fall outside of Auckland's boundaries were removed. Figure 4 presents the distribution of AED locations across Auckland.

Each AED was assigned to a meshblock using a two-stage spatial matching procedure. First, AED locations were matched to meshblocks based on strict spatial containment. For AEDs not falling exactly within a meshblock polygon, most often due to minor geographic inaccuracies, a nearest-neighbour assignment was applied,

linking each such AED to its closest meshblock. This ensured that every AED was uniquely associated with exactly one meshblock. AED counts were subsequently aggregated at the meshblock level, and the distribution of meshblocks is shown in Figure 5. In addition to these, 8,949 meshblocks contained no AEDs. The meshblock with the highest concentration of AEDs is located at Auckland Airport, containing 67 AEDs; however, this meshblock was removed as its population was zero. All other meshblocks with a population lower than their AED count were also removed from analysis ($n = 7$). Removing these meshblocks is consistent with previous literature [53].

For each meshblock, a grid-based sampling approach is used. A uniform mesh of 100-meter by 100-meter squares is overlaid across the entire bounding box that encloses the meshblock's irregular shape. The generation of this grid begins at the minimum coordinate boundary of the meshblock and proceeds in increments of 100 meters in both the x and y directions until the maximum boundary is reached. From this grid, candidate points are derived by taking the center, or centroid, of each individual grid cell. All centroids located outside the meshblock's territory are discarded.

The choice of 100-meter resolution is motivated primarily by computational tractability, while remaining consistent with recent theoretical findings on optimal AED spacing. Finer resolutions, such as 50-meter grid cells, would quadratically increase the number of potential AED locations, rendering the optimization problems computationally infeasible at a larger scale. However, this 100-meter grid must also be fine enough to find solutions that approximate continuous-space optima. Lin et al. [31] demonstrate that the optimal distance between AEDs in a linear space is approximately 134 metres, assuming a brisk walking speed of 4 km/h and a one-minute target for AED acquisition. A 100-meter grid resolution ensures that any two AEDs placed in adjacent cells are spaced at a distance reasonably aligned with this theoretical optimum, while a coarser resolution would risk undershooting the coverage potential implied by these findings. Thus, the chosen resolution balances the need for computational feasibility with the ability to capture placement strategies that approximate theoretically optimal configurations.

Moreover, to reduce computation time, all meshblocks with an area bigger than 10km^2 were removed ($n = 21$) because these large meshblocks predominantly correspond to sparsely populated or non-residential areas and would otherwise generate an excessive number of candidate locations.

Figure 4: Auckland AED registry data from AED locations NZ

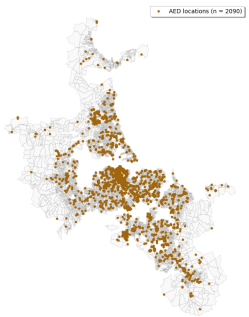
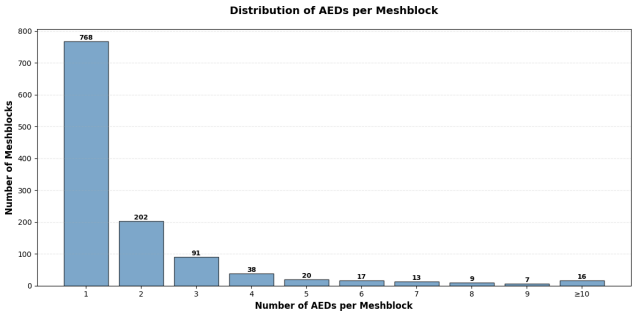
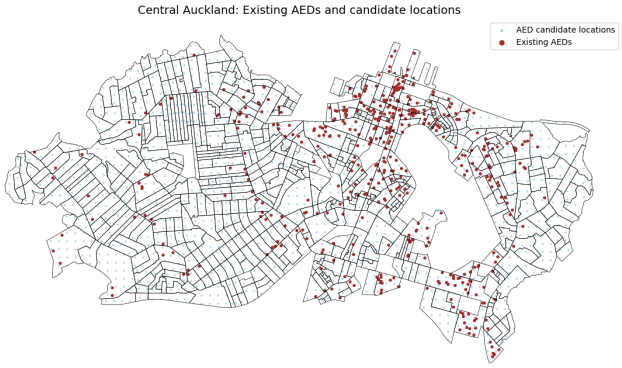


Figure 5: Histogram of the number of meshblocks for each AED count



The existing AED deployment serves as a baseline against which the optimization model is evaluated. After establishing that the optimization model outperforms the current configuration, we are able to extend the analysis to assess the incremental impact of adding N new AEDs at candidate locations, evaluating both coverage and equity improvements. This candidate-based approach divides the study area into a spatially uniform and computationally tractable grid of potential AED locations, which preserves meshblock boundaries and reflects the underlying population distribution. The resulting candidate set defines the feasible decision space for the AED placement optimization model, which maximizes coverage of simulated OHCA events. An example of this candidate grid and the corresponding current AED deployment is shown in Figure 6.

Figure 6: A ward in Central Auckland with the candidate AED locations and the existing registered AEDs.



3.4 Probabilistic coverage

AED placement to cover OHCA events can be formulated as a Maximum Covering Location Problem (MCLP), originally introduced by Church and ReVelle (1974) [15]. In its classical formulation, the MCLP seeks to locate a fixed number K facilities such that the total

covered demand is maximized. Coverage is defined in binary terms: a demand point is considered covered if it lies within a predefined service radius of at least one facility, and uncovered otherwise.

However, Chan et al. (2016) [14] reasoned that covering OHCA events with AEDs can be modelled with a more realistic approach that accounts for the negative relationship between distance and the likelihood of successful deployment of an AED. Their approach uses a nonincreasing coverage function, where the probability of coverage decreases as an AED is located further away from a cardiac arrest. We adopt such a probabilistic coverage function that is consistent with this principle.

Let p_{ij} denote the probability that an OHCA event i is covered by an AED located at j. Coverage is defined using a piecewise exponential decay function:

$$p_{ij} = \begin{cases} 1, & d_{ij} \leq 20, \\ \exp\big(-\alpha(d_{ij} - 20)\big), & 20 < d_{ij} \leq 100, \\ 0, & d_{ij} > 100, \end{cases}$$

where distances are measured in metres and $\alpha > 0$ is a parameter that determines how rapidly the coverage declines. Consistent with the approach in Chan et al. (2016) [14], we set $\alpha = 0.05$ to reflect a reasonably fast decay. Full coverage is assumed within 20 metres, reflecting the average line of sight. Between 20 and 100 metres, the probability of effective AED use decreases exponentially with distance, capturing the challenges of finding and retrieving the device. Beyond 100 metres, coverage is assumed negligible and set to zero, aligning with the diminishing likelihood of a bystander retrieving an AED from such a distance in a timely manner [4].

To integrate the coverage function into the final optimization model, we adopt the single-responder best-case formulation proposed by Chan et al. (2016) [14]. This specification assumes that exactly one bystander retrieves an AED and selects the closest available device. Accordingly, the probability that cardiac arrest j is covered is defined as

$$f_j(z) = \max_{i \in I} z_{ij},$$

where z_{ij} denotes the coverage probability of cardiac arrest j by an AED located at site i, and I is the set of candidate locations. In the optimization model, this probability depends on the AED placement decision. Let y_i be a binary decision variable indicating whether an AED is installed at location i. The effective coverage contribution of location i to event j is therefore $z_{ij} = y_i p_{ij}$, where p_{ij} is the distance-based coverage probability defined above. This will be used when formulating the optimization model.

This modeling choice is supported by recent evidence on bystander behavior. Kono et al. (2024) [27] show that the presence of multiple bystanders does not necessarily increase the likelihood of AED retrieval, as social dynamics may inhibit individual action. Their findings suggest that assuming multiple independent responders simultaneously searching for an AED may overstate effective coverage. In contrast, modeling a single active responder is more consistent with observed behavioral patterns in emergency situations.

Moreover, technological advancements further strengthen the justification for the best-case model. Nielsen et al. (2022) [35] show that when volunteer responders are activated through smartphone-based dispatch systems, they receive real-time information about the cardiac arrest location and nearby AEDs. These systems are advancing rapidly and make it increasingly likely that volunteers and other bystanders can access the right information to find the nearest AED instead of the furthest.

In the standard single-responder best-case formulation, only the closest device contributes to the coverage probability for a given cardiac arrest event. That means that the presence of additional devices at the same location does not further increase coverage. This modelling choice is substantiated by Lin et al. (2023) [31], who demonstrate that to maximize overall rescue benefit, AEDs should avoid overlapping coverage, as overlapping hot zones represent inefficient deployment. Their optimal spacing principle implies that any two devices placed closer than their suggested 134 meters yield diminishing marginal returns. A 100 m × 100 m grid resolution ensures that AEDs placed in the same cell would be somewhat inefficiently close according to this theoretical optimum, making binary treatment of cells a reasonable simplification aligned with maximizing spatial coverage. Moreover, the primary goal of this study is to investigate the fairness in the efficiency-equity trade-off, where the spatial distribution of AEDs is a more important driver of the outcome than the density of AEDs.

3.5 Incorporating Equity Considerations

To incorporate equity into our model, we use intersections of the ethnic groups and deprivation quintiles. By creating intersections, we can have one model that will optimize for equity across socioecomic status and ethnicity. To do this, let $E \in \{\text{European/other}, \text{Maori}, \text{Pacific}\}$ denote the set of ethnic groups and $D \in \{1, 2, 3, 4, 5\}$ the set of deprivation quintiles. The set of intersectional groups is then defined as

$$G = E \times D.$$

A brief overview of the size of all groups is shown in Figure 7. It can be seen that the groups are highly heterogeneous in size. The (European/other, Q1) group for example contains approximately 36 times more people than the (Maori, Q1) group. This heterogeneity will be important to account for when evaluating the efficiency-equity trade-off, as larger groups might need a bigger emphasis in the model than relatively smaller groups to keep individual-level equity.

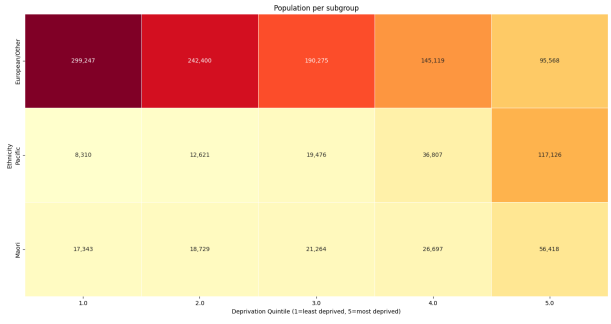


Figure 7: The population count per subgroup in Auckland

To use the intersectional groups, we need to assign each simulated OHCA event to one or more of these groups to calculate the average coverage per group and the corresponding fairness amongst these groups. Let J denote the set of simulated OHCA events. Each event $j \in J$ occurs in a meshblock $m(j) = i$, which determines the deprivation quintile of the event.

While the location and corresponding meshblock of each simulated OHCA event is known, the ethnicity of the individual experiencing the OHCA is not directly observed. Instead, we estimate the probability that an event belongs to a particular ethnicity using the demographic composition of the meshblock combined with the incidence rates that were calculated before.

The expected annual number of OHCAs in meshblock i for demographic group (a, e, s) can be estimated by:

$$\lambda_{i,a,e,s} = \frac{P_{i,a,e,s}}{100{,}000}\, r_{a,e,s}.$$

Notice that when we sum these values over all ethnicities, age groups and sex categories and multiply it with the deprivation correction factor, we again obtain the expected number of total OHCA events for meshblock i.

By aggregating across age groups and sex categories, we obtain the expected number of OHCAs in meshblock i attributable to ethnicity e:

$$\lambda_{i,e} = \sum_{a} \sum_{s} \lambda_{i,a,e,s}.$$

The total expected number of OHCAs in meshblock i without accounting for the socioeconomic factor is then:

$$\lambda_i = \sum_{e \in E} \lambda_{i,e}.$$

Using these quantities, we estimate the probability that an OHCA event occurring in meshblock i (with deprivation quintile d) belongs to ethnicity e as

$$p_{i,e} = \frac{\lambda_{i,e}}{\lambda_i}.$$

Thus, each simulated event j occurring in meshblock i contributes probabilistically to the ethnic groups according to $p_{i,e}$. To illustrate the implications of this probabilistic assignment, we compute the implied distribution of OHCA events across the intersectional subgroups.

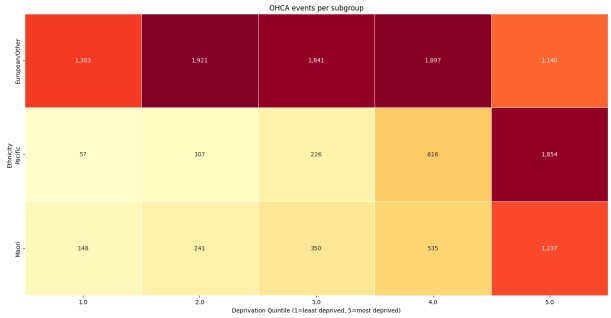


Figure 8: The number of expected OHCA events per subgroup in Auckland

Figure 8 shows substantial variation in the expected number of OHCA events across intersectional subgroups. Notably, some smaller population groups, such as (Pacific, Q5), experience a higher number of expected OHCA events than larger groups such as (European/other, Q1), despite having a substantially smaller population size. This shows our emphasis on the underlying differences in OHCA incidence rates between different ethnicities and socioeconomic statuses.

Next, let $f_j(y)$ denote the AED coverage probability for event j under AED placement decision y. The average AED coverage for each intersectional group $(e, d) \in G$ is then computed as

$$F_{e,d}(y) = \frac{\sum_{j \in J : d(j)=d} p_{m(j),e}\, f_j(y)}{\sum_{j \in J : d(j)=d} p_{m(j),e}}, \quad \forall \{e, d\} \in G.$$

In this formulation, each simulated OHCA event contributes proportionally to ethnic groups according to the expected incidence implied by the meshblock demographics, while remaining fully associated with the deprivation quintile of the meshblock in which it occurs. The resulting $F_g(y)$ therefore represents the average AED accessibility experienced by group g across all simulated OHCA events.

3.6 An efficiency-equity trade-off function

As discussed in the literature review, Bertsimas et al. [8] introduced the α-fairness framework for resource allocation problems, which has become foundational in operations research for balancing efficiency and equity. In their formulation, consider a set of n individuals, each receiving a utility allocation $u_j \geq 0$ from a feasible utility set $U \subset \mathbb{R}^n_+$. The α-fair social welfare function is defined as:

$$W_\alpha(u) = \begin{cases} \displaystyle\sum_{j=1}^{n} \frac{u_j^{1-\alpha}}{1-\alpha}, & \alpha \geq 0, \alpha \neq 1, \\ \displaystyle\sum_{j=1}^{n} \log(u_j), & \alpha = 1, \end{cases}$$

where $\alpha \geq 0$ governs the trade-off between efficiency and equity. When $\alpha = 0$, the function reduces to the utilitarian objective $W_0(u) = \sum_j u_j$. As α increases, more weight is placed on individuals with lower utility. At $\alpha = 1$, the objective corresponds to proportional fairness. In the limit $\alpha \to \infty$, the solution converges

to max-min fairness, which prioritizes improving the utility of the worst-off individual above all other individuals.

While this framework is very useful, applying this framework directly to our AED placement problem presents two challenges. First, α-fairness operates at the individual level, requiring a utility value u_j for every person in the population. In our setting, the vast majority of individuals would have $u_j = 0$ because there are simply not enough AEDs to cover everyone. Consequently, for $\alpha > 1$, the terms $0^{1-\alpha}$ become infinite and $\log(0)$ is undefined when $\alpha = 1$. Thus, the α-fairness objective is not well-defined for our problem. Second, even when defined, the objective is unbounded for $\alpha < 1$, making it difficult to interpret the magnitude of welfare improvements or compare outcomes across different population sizes.

To address these challenges, we turn to the Atkinson social welfare function [2], originally developed for income inequality measurement but well-suited to settings where outcomes are naturally aggregated across groups. Notably, in banking and finance, regulators evaluate risk or service quality across heterogeneous customer segments rather than at the individual level. Since we do not have hetergenous group sizes however (see Figure 7), we evaluate the efficiency-equity trade-off using a population-weighted Atkinson social welfare function.

For this social welfare function, we use the coverage values $F_{e,d}(y)$ as defined before to represent the expected probabilistic coverage for individuals in that subgroup under AED placement strategy y. To aggregate these group-level outcomes into a single welfare measure while preserving an individual-level interpretation of inequality aversion, we use the weighted power mean with population weights $w_{e,d}$. For $\varepsilon \geq 0$, the Atkinson social welfare function is defined as:

$$W_\varepsilon(\mathbf{F}(y)) = \begin{cases} \left(\sum_{g \in G} w_g \cdot F_g(y)^{1-\varepsilon} \right)^{\frac{1}{1-\varepsilon}}, & \varepsilon \geq 0, \varepsilon \neq 1, \\ \prod_{g \in G} F_g(y)^{w_g}, & \varepsilon = 1, \end{cases}$$

where $\mathbf{F}(y) = (F_{1,1}(y), \ldots, F_{3,5}(y))$ and the population weights $w_{e,d}$ are defined as:

$$w_{e,d} = \frac{N_{e,d}}{\sum_{\{e',d'\} \in G} N_{e',d'}} \quad \text{with} \quad N_{e,d} = \sum_{j \in J : d(j) = d} p_{m(j),e}.$$

Here, $N_{e,d}$ represents the expected number of OHCA events for group (e, d) over the simulation period. The population-weights that are added to Atkinsons original formulation serve two essential purposes regarding equity and social welfare. First, they ensure equal consideration of individuals. Without weights, a group of 50 individuals would have the same influence as a group of 5,000. Second, they account for statistical precision. Groups with higher OHCA incidence produce more reliable coverage estimates, and weights prevent volatile estimates from small groups from driving optimization outcomes arbitrarily.

The welfare value $W_\varepsilon(\mathbf{F}(y))$ is bounded between 0 and 1 and admits a natural interpretation as the equally distributed equivalent coverage: the hypothetical coverage level that, if received uniformly by every individual, would yield the same social welfare as the actual unequal distribution [2]. For $\varepsilon > 0$, this value is always less than or equal to the average coverage $\bar{F} = \sum_i w_i F_i$, with equality only when all groups have identical coverage. The difference $\bar{F} - W_\varepsilon(\mathbf{F}(y))$ represents the welfare loss due to inequality, and the Atkinson index $A_\varepsilon = 1 - W_\varepsilon / \bar{F}$ quantifies this loss as a proportion of average coverage.

The parameter $\varepsilon \geq 0$ governs the strength of inequality aversion, which acts similar to α in Bertsimas' framework. When $\varepsilon = 0$, the function reduces to the utilitarian objective of maximizing expected average coverage and as ε increases, greater weight is placed on improving coverage for groups with lower expected coverage. In the limit $\varepsilon \to \infty$, the objective approaches max-min fairness. For $\varepsilon > 0$, the welfare function exhibits constant relative inequality aversion. This means that the marginal social welfare contribution of an individual in group (e, d) is proportional to $F_{e,d}^{-\varepsilon}$, indicating that improvements to groups with lower coverage values ($F_{e,d}$) are increasingly valued as ε grows.

The marginal rate of substitution (MRS) between an individual in group (e, d) and an individual in group (e', d') measures the rate at which society is willing to trade coverage from one individual to another while keeping overall welfare constant. For the population-weighted Atkinson welfare function, the MRS between two individuals is given by:

$$\text{MRS}_{(e,d),(e',d')} = \frac{\partial W / \partial F_{e,d}}{\partial W / \partial F_{e',d'}} = \frac{w_{e,d} \cdot F_{e,d}^{-\varepsilon}}{w_{e',d'} \cdot F_{e',d'}^{-\varepsilon}}.$$

The elasticity of substitution σ measures the percentage change in the ratio of coverages divided by the percentage change in the MRS. Formally:

$$\sigma = \left| \frac{d \ln(F_{e,d} / F_{e',d'})}{d \ln(\text{MRS})} \right|.$$

To derive σ, we first rewrite the denominator as follows

$$\ln(\text{MRS}) = \ln(w_{e,d}) - \ln(w_{e',d'}) - \varepsilon \ln(F_{e,d}) + \varepsilon \ln(F_{e',d'}).$$

Differentiating with respect to $\ln(F_{e,d}/F_{e',d'})$:

$$d \ln(\text{MRS}) = -\varepsilon \, d \ln(F_{e,d}) + \varepsilon \, d \ln(F_{e',d'}) = -\varepsilon \, d \ln \left(\frac{F_{e,d}}{F_{e',d'}} \right).$$

Note that the terms $\ln(w_{e,d})$ and $\ln(w_{e',d'})$ are constants, so their derivatives are zero. Therefore, the population weights do not affect the elasticity. We obtain the elasticity of substition now by taking absolute values and rearranging:

$$\sigma = \left| \frac{d \ln(F_{e,d} / F_{e',d'})}{d \ln(\text{MRS})} \right| = \left| \frac{d \ln(F_{e,d} / F_{e',d'})}{-\varepsilon \, d \ln \left(F_{e,d} / F_{e',d'} \right)} \right| = \frac{1}{\varepsilon}.$$

Thus, the elasticity of substitution is constant and given by $\sigma = 1/\varepsilon$. When ε is small (close to 0), the elasticity is large, meaning that coverage can be easily traded between individuals with little welfare loss. As ε increases, the elasticity decreases, making it harder to compensate for low coverage in one group with high coverage in another. In the extreme case $\varepsilon \to \infty$, the elasticity approaches 0, implying that no substitution is possible. In this case, the welfare function focuses solely on the worst-off individual.

Importantly, despite the different domain of application, the Atkinson framework is mathematically equivalent to α-fairness

when applied to group-level data with population weights. To see this, observe that for $\varepsilon \neq 1$, the welfare function $W_\varepsilon(\mathbf{F}(y))$ is monotonically increasing in the quantity $\sum_{g\in G} w_g \cdot F_g(y)^{1-\varepsilon}$. Thus, maximizing $W_\varepsilon(\mathbf{F}(y))$ is ordinally equivalent to maximizing:

$$\max_{y\in\mathcal{Y}} \sum_{g\in G} w_g \cdot F_g(y)^{1-\varepsilon}.$$

This objective takes the identical functional form as α-fairness with $\alpha = \varepsilon$, where individual utility u_j is replaced by group coverages $F_{e,d}(y)$ and uniform weights $1/n$ are replaced by population weights w_g. The only differences are that we aggregate at the group level rather than the individual level, and we include population weights to ensure each individual's welfare is counted equally. This connection implies that the theoretical properties established by Bertsimas et al. [8], including convergence to max-min fairness as $\varepsilon \to \infty$ and the existence of efficient solution methods for convex problems, carry over to our setting.

3.7 Optimization Problem Formulation

The objective of our model is to determine where to place a limited number of AEDs in order to maximize the social welfare function defined above.

Let I denote the set of candidate AED locations and let y_i be a binary decision variable indicating whether an AED is placed at location $i \in I$. The total number of deployable AEDs is limited to K.

Let p_{ij} denote the probability that OHCA event j is covered by an AED located at site i, as defined by the probabilistic coverage function described previously. As stated before, we adopt the single-responder best-case assumption from Chan et al. (2016) [14], which assumes that the effective coverage probability for event j is determined by the available AED that provides the highest coverage probability. Consequently, the coverage of event j is defined as

$$f_j(y) = \max_{i\in I: p_{ij}>0} \{y_i p_{ij}\}.$$

Substituting this event-level coverage into the group-level coverage measure $F_{e,d}(y)$ yields the following optimization problem:

$$\begin{aligned}
\text{Maximize} \quad & W_\varepsilon(\mathbf{F}(y)) && (1)\\
\text{subject to} \quad & f_j(y) = \max_{i\in I: p_{ij}>0} \{y_i p_{ij}\}, \quad \forall j \in J, && (2)\\
& \sum_{i\in I} y_i \leq K, && (3)\\
& y_i \in \{0,1\}, \quad \forall i \in I. && (4)
\end{aligned}$$

The max operator in Constraint (2) makes this formulation nonlinear and therefore difficult to solve. Similar to the models of Berman et al. (2003) [6] and Karasakal and Karasakal (2004) [26], we obtain a tractable formulation by introducing binary variables w_{ij} indicating whether cardiac arrest event j is assigned to an AED placed at location i.

Let

$$I_j = \{i \in I \mid p_{ij} > 0\}$$

denote the set of candidate locations capable of covering event j.

Using these assignment variables, the max operator can now be rewritten to make the model easier to solve. Namely, the optimization problem can be reformulated as the following mixed-integer optimization model:

$$\begin{aligned}
\text{Maximize} \quad & W_\varepsilon(\mathbf{F}(y)) && (5)\\
\text{subject to} \quad & f_j = \sum_{i\in I_j} p_{ij} w_{ij}, \quad \forall j \in J, && (6)\\
& \sum_{i\in I} y_i \leq K, && (7)\\
& w_{ij} \leq y_i, \quad \forall i \in I_j, \forall j \in J, && (8)\\
& \sum_{i\in I_j} w_{ij} \leq 1, \quad \forall j \in J, && (9)\\
& w_{ij} \in \{0,1\}, \quad \forall i \in I_j, \forall j \in J, && (10)\\
& y_i \in \{0,1\}, \quad \forall i \in I. && (11)
\end{aligned}$$

In this formulation, the variables w_{ij} assign each cardiac arrest event to at most one AED location, consistent with the single-responder assumption. Constraint (8) ensures that an event can only be assigned to a location if an AED is installed there. Constraint (9) guarantees that each event is covered by at most one AED, thereby capturing the best-case coverage assumption.

The resulting event-level coverage values f_j are subsequently aggregated into the group-level coverage measures $F_{e,d}(y)$ defined earlier, which enter the ε-fair welfare objective. In this study, the model is implemented in Python and solved using Gurobi Optimizer version 13.0.1. Although the welfare function is nonlinear in the group-level coverage variables, due to the smoothness and monotonicity of the function, modern solvers like Gurobi are able to handle such nonlinear structures. In particular, recent versions of Gurobi support general nonlinear expressions and mixed-integer nonlinear programming (MINLP) via spatial branch-and-bound and successive linear outer-approximation techniques [25].

To improve computational tractability, the optimization is performed only on the inner aggregation term of the Atkinson welfare function,

$$\sum_{g\in G} w_g F_g(y)^{1-\varepsilon},$$

while the outer transformation is evaluated post-optimization. Because the Atkinson transformation is monotonic in this inner term for $\varepsilon \neq 1$, this reformulation preserves the welfare ordering of solutions.

For $\varepsilon < 1$, the inner summation is maximized, whereas for $\varepsilon > 1$, it is minimized, reflecting the sign change induced by the exponent $1-\varepsilon$. The case $\varepsilon = 1$, which corresponds to the logarithmic (geometric mean) form of the Atkinson index, is approximated numerically by setting $\varepsilon = 1.01$ in the implementation. This approximation ensures consistency with the power-form nonlinear structure used in the optimization model while remaining arbitrarily close to the logarithmic limit.

3.8 Statistical model validation

To validate the proposed optimization model, we compare the AED deployment suggested by the optimization model with the existing AEDs, both constrained to have the same number of AEDs to ensure

a fair comparison. We evaluate these models based on overall probabilistic coverage (the sum of f_j across all OHCA events), which is independent of the inequality aversion parameter ε.

To do this, we use a pure utilitarian approach for our model, which maximizes total expected OHCA coverage without regard to equity across population subgroups. This is achieved by setting the inequality aversion parameter to $\varepsilon = 0$ in our population-weighted welfare function $W_\varepsilon(\mathbf{F}(y))$. To see that this recovers the utilitarian objective, recall that $W_\varepsilon(\mathbf{F}(y))$ with $\varepsilon = 0$ is:

$$W_0(\mathbf{F}(y)) = \sum_{e,d} w_{e,d} \cdot F_{e,d}(y).$$

Substituting $w_{e,d} = N_{e,d}/N_{\text{total}}$ and $F_{e,d}(y) = \frac{1}{N_{e,d}} \sum_{j \in (e,d)} f_j(y)$ yields:

$$W_0(\mathbf{F}(y)) = \sum_{e,d} \frac{N_{e,d}}{N_{\text{total}}} \cdot \frac{1}{N_{e,d}} \sum_{j \in (e,d)} f_j(y) = \frac{1}{N_{\text{total}}} \sum_{j \in J} f_j(y).$$

Since N_{total} is constant for a given urban area, maximizing $W_0(\mathbf{F}(y))$ is equivalent to maximizing $\sum_{j \in J} f_j(y)$.

In this scenario, the model maximizes total probabilistic coverage by relocating all AEDs to new locations on the candidate solutions grid, providing an upper bound on the achievable coverage for the current number of AEDs. This utilitarian approach serves as a benchmark against which both the existing AED network and equity-aware deployments (with $\varepsilon > 0$) can be evaluated. The difference between the coverage achieved by the current deployment and the utilitarian optimum reflects the potential improvement that could be realized through strategic relocation, while the trade-off between coverage and equity is examined by comparing deployments across a range of ε values.

To obtain reliable performance estimates and reduce potential overfitting to historical cardiac arrest locations, we employ a 10-fold cross-validation procedure similar to Chan et al. [12]. To ensure a sufficient number of OHCA events relative to the number of AEDs being placed, we simulate ten years of OHCA events. This provides an adequately sized dataset for training and testing. The set of historical OHCA events is then randomly partitioned into ten equally sized subsets. In each iteration, nine subsets (90% of the events) are used as a training set to determine the optimal AED locations, while the remaining subset (10%) is used as a testing set to evaluate the resulting coverage. This process is repeated ten times so that each subset serves once as the testing set, after which the coverage results are aggregated across all testing sets.

Given the computational complexity of solving the optimization model at the full Auckland scale, the analysis is conducted separately for the four urban areas: Central, Northern, Southern, and Western Auckland Zones. This approach avoids an unnecessarily large optimization problem that would be computationally intractable if solved for the entire region at once. For each urban area, the model is run independently using the meshblocks, candidate AED locations, and OHCA events specific to that area. The resulting coverage probilities are then aggregated across the four areas. To obtain an overall estimate of model performance for Auckland as a whole, we compute a weighted average of the results from each urban area, using the total number of OHCA events in each area as the weight. This ensures that the areas with higher cardiac arrest incidence contribute proportionally more to the aggregate performance.

Let U denote the set of urban areas, and let n_u represent the total number of OHCA events in urban area $u \in U$. The weighted average probabilistic coverage is given by:

$$\overline{W}_\varepsilon(\mathbf{F}(y)) = \frac{\sum_{u \in U} n_u \cdot W_\varepsilon^{(u)}(\mathbf{F}(y))}{\sum_{u \in U} n_u},$$

where $W_\varepsilon^{(u)}(\mathbf{F}(y))$ is the population-weighted sociale welfare for urban area u.

Since there are only four urban areas to analyze, statistical tests have limited significant power. Therefore, to assess whether differences in probabilistic coverage between the current AED deployment and the optimized deployments are statistically significant, we use bootstrappinging. Specifically, we resample the OHCA events with replacement 10,000 times, computing the mean coverage difference (optimized minus existing) for each bootstrap sample to calculate a 95% confidence interval. If this interval does not contain zero, the improvement is considered statistically significant at the $\alpha = 0.05$ level.

In addition to evaluating average probabilistic coverage, we also assess the equity of AED deployments using the Gini coefficient across all demographic groups $g \in G$. This can be useful to obtain a better understanding of the current AED deployment, but also the results of a simple utilitarian approach.

The Gini coefficient is computed over the vector of group-level coverages and takes values between 0 and 1, where 0 indicates perfect equality and 1 indicates maximal inequality. Formally, it is defined as:

$$\text{Gini}(y) = \frac{\sum_g \sum_{g'} \left| F_g(y) - F_{g'}(y) \right|}{2|G| \sum_g F_g(y)}.$$

This metric is particularly useful in the validation setting, as it allows us to assess equity for a single deployment (e.g., the existing network or the utilitarian solution with $\varepsilon = 0$), where trade-off-based measures are not applicable.

3.9 Budget scenarios: Expanding the Number of AEDs

Although the optimized utilitarian approach described above should yield the largest increase in overall coverage, such a full relocation of existing AEDs is not desirable for several reasons. First, improvements in total coverage do not necessarily translate into more equitable access across groups. Moreover, the relocation scenario considered in the statistical model validation is not a practical implementation strategy, as relocating existing AEDs is both costly and logistically challenging, and may also reduce accessibility in locations where devices could already be well-placed and frequently used. For these reasons, the next step is to consider a setting in which AEDs are added rather than relocated. That means that existing AEDs will be fixed in the model and cannot be relocated and the optimization model will focus on expanding the network through additional AEDs. This should preserve the benefits of the current infrastructure while targeting remaining coverage and equity gaps.

We expand the network using a percentage-based addition for the four urban areas rather than a fixed global budget K.

Let K_u^{existing} denote the number of AEDs currently located in urban area $u \in U$. For each scenario, we define an expansion percentage $\rho \in \{10\%, 20\%, 50\%, 100\%\}$, which determines the percentage of AEDs that can be added to the existing network. The total number of AEDs to be placed in area u is then given by:

$$K_u(\rho) = K_u^{\text{existing}} + \left\lceil \rho \cdot K_u^{\text{existing}} \right\rceil.$$

This approach ensures that areas with larger existing AED networks receive proportionally more AEDs. This approach aims to reflect realistic public health investment strategies where funding is often allocated as a percentage increase to the existing structure. The optimization model is solved independently for each urban area using its respective budget $K_u(\rho)$, and the resulting coverage metrics are aggregated using the weighted average defined before. The addition of new AEDs and fixing the current AEDs does require a slight adaption of the optimization model defined in 3.7. Let $I^{\text{existing}} \subset I$ denote the set of indices corresponding to existing AED locations, and let $I^{\text{candidate}} = I \setminus I^{\text{existing}}$ be the set of candidate locations available for new AED placements. The binary decision variables y_i are defined such that $y_i = 1$ for all $i \in I^{\text{existing}}$, i.e., existing AEDs are fixed in the model and cannot be removed. Only y_i for $i \in I^{\text{candidate}}$ are treated as decision variables.

The optimization problem for a given urban area u and expansion percentage ρ is formulated as follows:

$$\text{Maximize} \quad W_\alpha(\mathbf{F}(y)) \tag{12}$$

$$\text{subject to} \quad f_j = \sum_{i \in I_j} p_{ij} w_{ij}, \quad \forall j \in J, \tag{13}$$

$$\sum_{i \in I^{\text{candidate}}} y_i \le K_u^{\text{new}}(\rho), \tag{14}$$

$$w_{ij} \le y_i, \quad \forall i \in I_j, \forall j \in J, \tag{15}$$

$$\sum_{i \in I_j} w_{ij} \le 1, \quad \forall j \in J, \tag{16}$$

$$y_i = 1, \quad \forall i \in I^{\text{existing}}, \tag{17}$$

$$w_{ij} \in \{0, 1\}, \quad \forall i \in I_j, \forall j \in J, \tag{18}$$

$$y_i \in \{0, 1\}, \quad \forall i \in I^{\text{candidate}}. \tag{19}$$

The objective function (12) remains the same as in the original formulation, maximizing the social welfare function $W_\alpha(\mathbf{F}(y))$. The event coverage constraint (13) is also the same as in the original formulation.

Constraint (14) limits the number of new AEDs placed in the candidate set to $K_u^{\text{new}}(\rho) = \lceil \rho \cdot K_u^{\text{existing}} \rceil$, where K_u^{existing} is the number of existing AEDs in urban area u and ρ is the expansion percentage. Notice that we use the ceiling function in order to ensure that $K_u^{\text{new}}(\rho)$ is always a positive integer value.

Constraints (15)-(19) are similar to the original formulation with the addition of the fixed AED locations, where $y_1 = 1 \ \ \forall i \in I^{\text{existing}}$.

For each expansion level ρ, the model is solved under the utilitarian objective ($\alpha = 0$) to establish an upper bound on achievable coverage, as well as under a range of $\alpha > 0$ values to examine how the efficiency-equity trade-off evolves. By adding AEDs proportionally, we investigate whether marginal improvement in coverage diminish as more AEDs are added. Moreover, we explore whether or not increasing number of AEDs reduce or increase inequities across demographic groups. Lastly, we look into how the price of fairness changes when resource constraints are relaxed.

3.10 Visualizing the Efficiency-Equity Trade-off

Before evaluating the trade-off, we first illustrate how the choice of $\varepsilon \in \{0, 0.5, 1, 2, 4, \infty\}$ shapes the objective function itself to justify that this range of ε best captures the solution space. Figure 9 shows the behavior of the social welfare function W_ε under controlled settings. Figure 9a considers a simple two-group scenario where F_2 is fixed at 0.5 and F_1 varies, with equal population weights. Figure 9b extends this to 15 groups, with 14 groups fixed at 0.5 and one group varying.

When $\varepsilon = 0$, the welfare function is linear, meaning coverage gains across groups are perfectly substitutable. This corresponds to a purely utilitarian objective that maximizes average coverage without regard for its distribution. As ε increases to 0.5, 1, 2, and 4, the function becomes progressively more concave, reflecting increasing aversion to inter-group inequality. Improvements in lower-covered groups receive greater weight, while marginal gains in already well-served groups contribute less to overall welfare. At $\varepsilon = 1$, the welfare function reduces to the weighted geometric mean, which provides a natural intermediate benchmark. For $\varepsilon > 1$, the curvature becomes more pronounced, and the function increasingly prioritizes the worst-off groups. This effect is stronger with more groups, as seen in Figure 9b compared to the simpler two-group setting. In the limit as $\varepsilon \to \infty$, the welfare function converges to the minimum group coverage.

Having established how ε shapes the objective function, we now translate this into observable allocation outcomes. To visualize the resulting efficiency-equity frontier, we use two metrics. First, average probabilistic coverage $\bar{F} = \sum_g w_g F_g$ measures efficiency. Second, the Gini coefficient of the group-level coverages $\{F_g\}$ measures equity. Together, these define the efficiency-equity frontier: each ε gives a point in the (coverage, Gini) plane, and the frontier shows how much coverage must be sacrificed to achieve a given reduction in inequality. We will plot this frontier for different values of ρ to see whether the trade-off changes with budget.

To ensure that observed differences reflect genuine allocation effects rather than noise, we use Common Random Numbers (CRN) across five independent simulation seeds. Within each seed, every ε is evaluated on the exact same set of simulated OHCA events and cross-validation folds. This means any difference in coverage or Gini between two ε values comes from ε itself, not from random variation. After running all five seeds, we aggregate by taking the mean across seeds to obtain the final reported values.

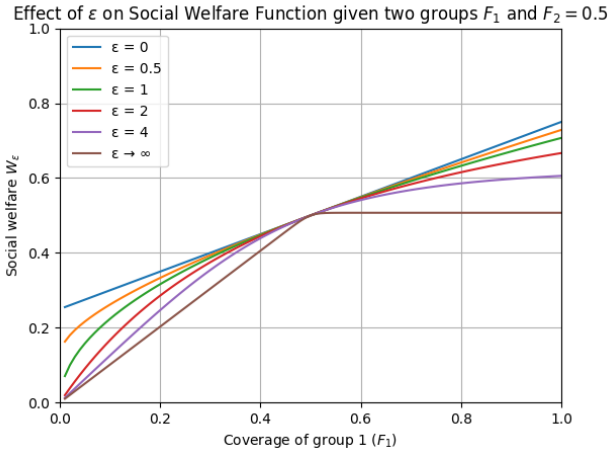
Because we use the same random draws, we report paired-difference standard errors rather than independent standard errors. For each fold and seed we compute:

$$\Delta\bar{F} = \bar{F}(\varepsilon) - \bar{F}(0), \qquad \Delta\,\text{Gini} = \text{Gini}(\varepsilon) - \text{Gini}(0),$$

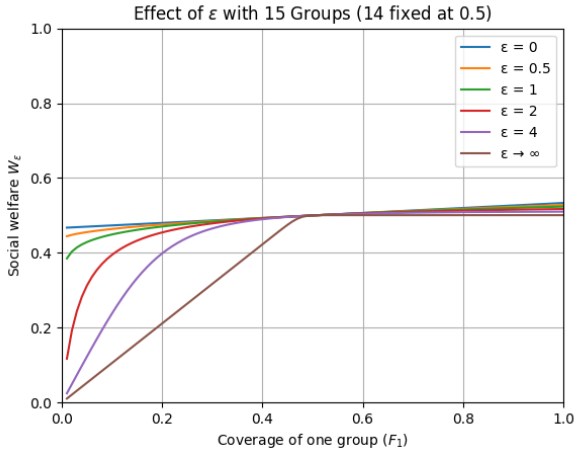
and estimate the standard error of these differences directly. Since $\bar{F}(\varepsilon)$ and $\bar{F}(0)$ come from the same events, they are highly correlated, and the variance of their difference is:

$$\text{Var}\big(\bar{F}(\varepsilon) - \bar{F}(0)\big) = \text{Var}\big(\bar{F}(\varepsilon)\big) + \text{Var}\big(\bar{F}(0)\big) - 2\,\text{Cov}\big(\bar{F}(\varepsilon),\, \bar{F}(0)\big)\,.$$

Figure 9: Effect of ε on the social welfare function.



(a) Two groups with $F_2 = 0.5$.



(b) 15 groups with 14 fixed at $F = 0.5$.

The covariance term shrinks the variance considerably, giving us tighter confidence intervals on the frontier so we can tell whether different ε values lead to meaningfully different outcomes.

3.11 Evaluating the Efficiency-Equity Trade-off

The frontier described above provides a visual summary of the trade-off, but we also require quantitative metrics to further substantiate policy decisions based on the efficiency-equity trade-off. Therefore, we use three different metrics to give a more complete assessment of each allocation's performance.

Following Bertsimas et al. [7], the price of fairness (POF) measures the relative loss in total coverage when moving from the utilitarian solution ($\varepsilon = 0$) to a more equity-oriented allocation:

$$\mathrm{POF}(\varepsilon) = \frac{\bar{F}(y^*(0)) - \bar{F}(y^*(\varepsilon))}{\bar{F}(y^*(0))}.$$

A small POF indicates that equity gains can be achieved at low efficiency cost. Symmetrically, the price of efficiency (POE) measures the relative loss in coverage of the worst-off subgroup compared to the most equitable allocation ($\varepsilon \to \infty$):

$$\mathrm{POE}(\varepsilon) = \frac{\min_g F_g(y^*(\infty)) - \min_g F_g(y^*(\varepsilon))}{\min_g F_g(y^*(\infty))}.$$

A small POE indicates that the utilitarian solution already provides near-maximal coverage to the worst-off subgroup. Together, POF and POE anchor the trade-off to two interpretable reference points: the purely utilitarian optimum and the most equitable allocation. These metrics are reported for all values of ε and ρ.

While the POF and POE describe the extremes, they do not identify a preferred ε for a given scenario as the relative importance of efficiency and equity is up to a policy maker. However, because AEDs are by nature meant to save lives, coverage is inherently the more critical objective. To support the selection of ε while respecting this priority, we introduce a weighted benefit-cost ratio.

Let $\lambda \geq 1$ be a weight reflecting the priority of coverage relative to equity. The efficiency-equity ratio is then defined as:

$$R(\varepsilon, \lambda) = \frac{\Delta\,\mathrm{Gini}(\varepsilon)/\,\mathrm{Gini}(0)}{\lambda \cdot \Delta\bar{F}(\varepsilon)/\,\bar{F}(0)},$$

where $\Delta\,\mathrm{Gini}(\varepsilon) = \mathrm{Gini}(0) - \mathrm{Gini}(\varepsilon) \geq 0$ is the absolute reduction in the Gini coefficient and $\Delta\bar{F}(\varepsilon) = \bar{F}(0) - \bar{F}(\varepsilon) \geq 0$ is the coverage sacrifice, both expressed as fractions of their utilitarian baseline. The weight λ penalizes coverage losses by a factor of λ relative to equity gains: at $\lambda = 1$, coverage and equity are treated as equally costly on their respective relative scales, whereas larger values of λ make the optimum reluctant to accept coverage sacrifices.

The optimal ε for a given policy weight is then:

$$\varepsilon^*(\lambda) = \arg\max_{\varepsilon > 0} \; R(\varepsilon, \lambda).$$

This criterion has two useful properties. First, $\varepsilon = 0$ can emerge as optimal: if the frontier is steep from the outset, meaning even a small increase in ε sacrifices a disproportionate amount of coverage relative to the equity gained, then $R(\varepsilon, \lambda)$ will be low for all $\varepsilon > 0$, and the utilitarian solution remains preferred. Second, $\varepsilon = \infty$ can be selected if the frontier is nearly flat throughout, meaning equity gains remain cheap all the way to the max-min solution. By reporting $\varepsilon^*(\lambda)$ for a range of λ values, we try and provide a transparent analysis where policymakers can choose their own settings based on their priorities.

As with all other results, ten-fold cross-validation is used to ensure out-of-sample performance rather than in-sample fit. For each fold, the model is trained on 90% of simulated OHCA events and evaluated on the remaining 10%. All reported metrics (POF, POE, and $R(\varepsilon, \lambda)$) are computed on the held-out test folds to avoid overfitting.

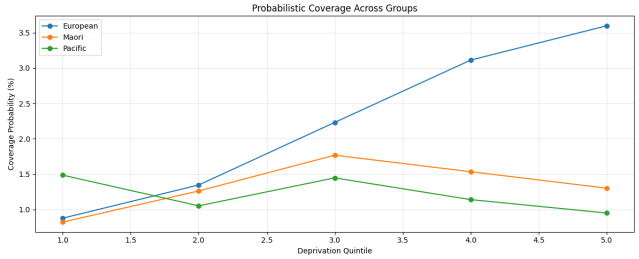
4 RESULTS

This section will discuss the results from the experiment that has been described before. First, we will go over the details of the existing AED network to establish the baseline. Next, we will validate our model by comparing the existing network with our model under a utilitarian objective. Lastly, we will broadly investigate the different budget scenarios and the corresponding efficiency-equity trade-offs that occur.

4.1 Current situation

Before investigating the improvements achievable through optimal AED placement, we first examine the performance of the existing AED network. Figure 10 illustrates the average probabilistic coverage ($F_{e,d}(y)$) for all subgroups G.

Figure 10: Performance of the existing AED deployment over 10 simulated years of OHCA events



It can be seen that in the current situation, the coverage by AEDs in the least deprived quintile is the lowest, which is in line with the number of AEDs in this quintile. Moreover, coverage varies substantially across ethnic groups, with the European/other population consistently achieving higher coverage than Māori and Pacific populations across most deprivation quintiles. The overall weighted average probabilistic coverage across all subgroups is 1.98%.

4.2 Importance of subgroup disaggregation

To assess whether our disaggregation into subgroups is necessary, we examine the results when considering ethnicity and socioeconomic status separately.

Table 2 reports weighted average probabilistic coverage by ethnicity when aggregating over deprivation quintiles. Table 3 reports weighted average coverage by deprivation quintile when aggregating over ethnicities. For both tables, we used the weighted average based on the total number of expected OHCA events per subgroup shown in Figure 8.

Table 2: Weighted average probabilistic coverage by ethnicity (aggregated over deprivation quintiles).

Ethnicity	Coverage (%)
European/other	2.25
Māori	1.38
Pacific	1.08

Table 3: Weighted average probabilistic coverage by deprivation quintile (aggregated over ethnicities).

Deprivation quintile	Coverage (%)
Quintile 1	0.91
Quintile 2	1.29
Quintile 3	2.04
Quintile 4	2.16
Quintile 5	1.74

At first sight, these aggregated results suggest relatively smooth differences between ethnic or socioeconomic groups. In particular, deprivation-based averages indicate a general increase in coverage from lower to higher deprivation quintiles, while ethnicity-based averages suggest higher coverage for the European/other population compared to Māori and Pacific populations.

However, these aggregated statistics conceal substantial heterogeneity that becomes visible only when the population is divided into the 15 ethnicity X deprivation subgroup combinations shown in Figure 10. For example, the European/other population exhibits a strong increase in coverage with deprivation, rising from 0.9% in quintile 1 to 3.6% in quintile 5, whereas Pacific populations display the opposite trend, decreasing from 1.5% to 1.0%. Māori populations show relatively stable coverage across deprivation levels.

These contrasting subgroup-specific patterns are obscured when averaging across either ethnicity or deprivation alone. Consequently, using only three ethnic groups or five deprivation groups would mask important intersectional disparities in AED accessibility. This substantiates the decision to model the population using 15 intersectional subgroups, allowing the optimization framework to explicitly account for inequities that would otherwise remain hidden in aggregated analyses.

4.3 Model validation and baseline comparison

To validate the proposed optimization model, we compare the performance of the optimized AED deployment with the current AED network. The existing AED locations, as obtained from the national registry, are first mapped onto the candidate location grid to ensure consistency with the optimization framework.

Table 4 presents the average probabilistic coverage across different urban areas in Auckland for both the existing and optimized AED deployments in the relocation scenario. The weighted average coverage increases from 0.0198 in the existing network to 0.0456 in the optimized deployment, representing a mean improvement of 0.0258 (95% CI: 0.0192, 0.0325). Note that the optimized average probabilistic coverage corresponds directly to $W_0(\mathbf{F}(y))$ and are therefore equivalent.

Statistical comparison has little power, but the fact that our 95% confidence interval, obtained from bootstrapping, does not contain zero suggests that the relocation of all AEDs does improve the coverage significantly. Altogether, these results show that the probabilistic coverage can be improved with the optimization model under the utilitarian approach for all of the four urban areas, providing validation for the utilitarian formulation of the model ($\varepsilon = 0$).

The optimized AED deployment produces a more evenly distributed network compared to the existing one, extending coverage

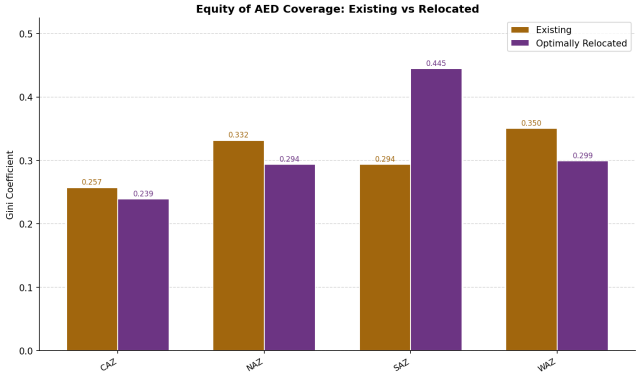
Table 4: Coverage comparison and OHCA events by urban area

Urban Area	Existing	Optimized	# Events
CAZ	0.0434	0.0771	4258
NAZ	0.0128	0.0280	2469
SAZ	0.00899	0.0388	4768
WAZ	0.00630	0.0208	2328
Weighted Average	0.0198	0.0456	13823

toward areas with relatively high OHCA incidence that are currently under-served. The existing network is largely concentrated around city centres and highly frequented public locations, consistent with prior research emphasizing placement in areas with high pedestrian traffic and visibility [11, 40, 52], whereas the optimized placements spread across residential areas where OHCA incidence is highest. Both networks, however, tend to avoid the more rural areas where very few OHCA events occur. The spatial distribution of both networks across all urban areas is shown in Appendix B.

In addition to overall coverage, we examine how the utilitarian approach affects the distribution of coverage across demographic groups. Figure 11 presents the Gini coefficient for both the existing and optimized deployments.

Figure 11: The Gini coefficient across all subgroups G in four urban areas of Auckland for the existing and optimally relocated network.



The results indicate that improvements in overall coverage do not necessarily translate into improvements in equity. While the Gini coefficient remains relatively stable or improves slightly in some urban areas, there is a substantial increase in the Southern Auckland Zone where the equity has worsened considerably. This mixed pattern is consistent with a purely utilitarian approach, which maximizes overall coverage without considering equity considerations.

Although the optimized utilitarian approach yields the largest increase in overall coverage, such a full relocation of AEDs is not desirable. In practice, existing AEDs are valuable regardless of their marginal contribution in the optimization model, and removing them from established public locations is generally not desirable. The following section focuses on a more realistic and policy-relevant setting in which AEDs are added rather than relocated.

4.4 Incremental AED deployment: adding new AEDs

This section will discuss the optimization model for various scenarios where we have a budget of $K_u^{\text{new}}(\rho)$ new AEDs to add to the existing AED network. This means that the existing AED network is fixed and cannot be relocated. To start the analysis, we will again look into the utilitarian approach ($\varepsilon = 0$) to formalize an upper bound on probabilistic coverage and a lower bound on equity with respect to other values of ε. Then, we will continue by looking into the efficiency-equity frontier and the POF and POE metrics for various scenarios of ρ.

4.4.1 A utilitarian baseline. The optimization model is run for different values of $\rho \in \{10\%, 20\%, 50\%, 100\%\}$. Newly added AEDs are primarily placed in areas not yet well covered by the existing network. At lower values of ρ, the model is forced to select the most critical areas first, revealing the locations with the highest marginal impact on coverage. As the number of additional AEDs increases, the spatial distribution becomes more dispersed, gradually extending coverage across the entire area. The spatial progression across all budget scenarios for the Central Auckland Zone, which has the highest number of AEDs among the urban areas and therefore serves as the most informative example, is provided in Appendix B.

To quantify the impact of adding AEDs, Figure 12 presents the mean probabilistic coverage across all urban areas as a function of $K_u^{\text{new}}(\rho)$. As expected, coverage increases monotonically with the number of additional AEDs. However, surprisingly, the rate of improvement remains relatively stable as $K_u^{\text{new}}(\rho)$ increases, suggesting that the marginal gains from additional AEDs do not diminish substantially over the considered range. This indicates that even a doubling of the current AED network could still yield meaningful improvements in overall coverage.

Despite these coverage gains, adding AEDs under a utilitarian objective does not translate into a more equitable distribution. Figure 13 shows that the Gini coefficient does not improve substantially as ρ increases, and even worsens in the Northern Auckland Zone. This is explained by the spatial structure of these zones: the worst-off subgroups are located in areas that are harder to reach and have lower population density, meaning a coverage-maximizing model still prioritizes denser areas with higher absolute OHCA counts even when the per-capita incidence rate is lower. Figure 14 makes this concrete: the coverage of the already best-served group (European, Quintile 5) increases the most across all values of ρ, while the least-served groups see comparatively little improvement. These findings provide a solid foundation for questioning whether a purely utilitarian approach is the optimal policy for improving AED coverage across the entire population.

4.4.2 The efficiency-equity trade-off. A pure utilitarian approach can lead to more inequity among subgroups. Therefore, we want to increase the social welfare parameter ε to analyze different policies in which there is a trade-off between equity and overall coverage or efficiency.

Figure 12: The mean probabilistic coverage over 10 folds for different areas of Auckland and various scenarios of ρ, expressed in the number of newly added AEDs ($K_u^{\text{new}}(\rho)$)

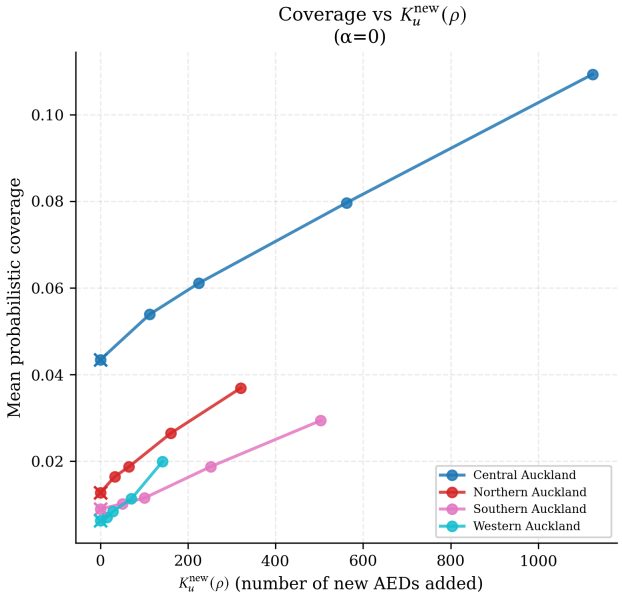
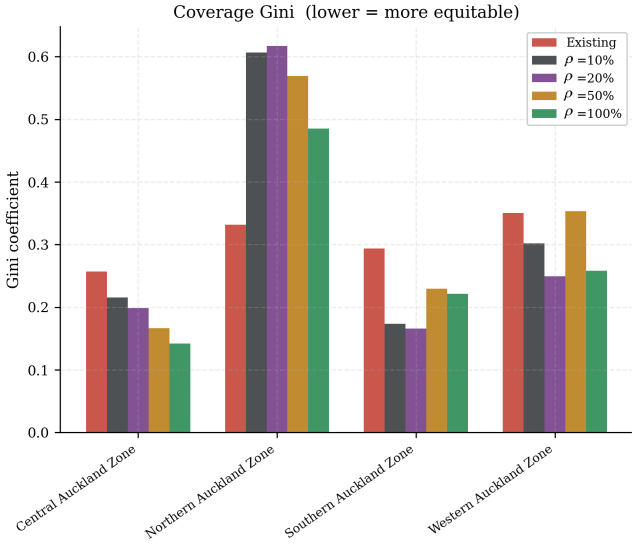
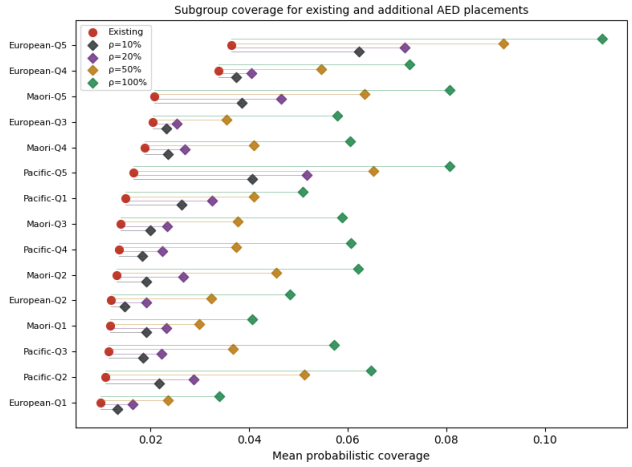


Figure 13: The Gini coefficient for four urban areas in Auckland for different values of ρ



To start, we examine the two extreme cases $\varepsilon = 0$ and $\varepsilon = \infty$ and look for spatial differences. Figure 15 shows the placement for a budget of $\rho = 10\%$ in the Central Auckland Zone as its high population and AED density make spatial differences most visible. The utilitarian policy ($\varepsilon = 0$) concentrates new AEDs in the already-dense northern cluster, reinforcing coverage where OHCA incidence is highest. In contrast, the max-min policy ($\varepsilon = \infty$) disperses new

Figure 14: Per-subgroup weighted average coverage across all four urban areas in Auckland



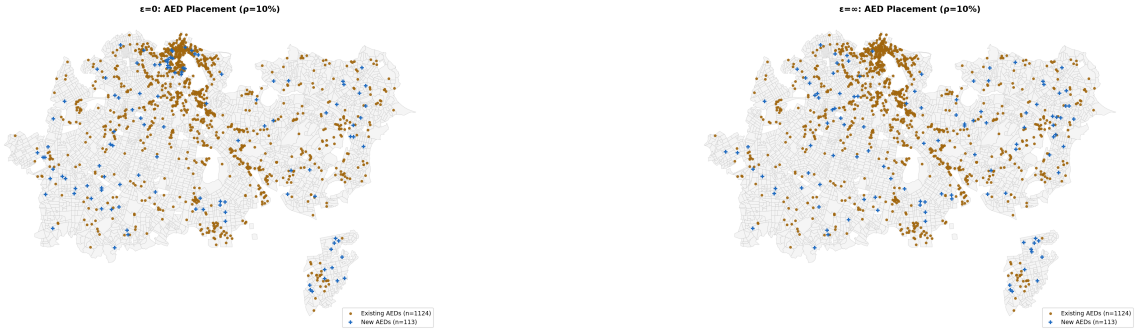
AEDs across the rest of Central Auckland, targeting areas that have little coverage under the existing AED deployment. This spatial difference confirms that the two policies pursue fundamentally different objectives.

The efficiency–equity trade-off is quantified in Figure 16 and Table 5. As ε increases from 0 to ∞, mean coverage $\bar{F}$ falls from 0.0563 to 0.0497 (11.7%), while the Gini coefficient improves from 0.404 to 0.331 (-18.1%). Most equity gains occur at low ε: moving from $\varepsilon = 0$ to $\varepsilon = 1$ reduces the Gini by 3.3 percentage points with a POF of only 3.8%, whereas the final step ($\varepsilon = 4$ to ∞) yields just 1.3 additional Gini points at a POF of 3.2%. The POE column confirms this from the opposite direction: at $\varepsilon = 0$, the worst-off group lags the max-min benchmark by 24.1% (POE = +24.1%); by $\varepsilon = 2$, the gap closes to -1.4%, indicating that $\varepsilon = 2$ approximates the max-min distributional outcome while preserving 6.3% more aggregate coverage (POF = 6.3%).

To support the selection of a preferred ε, Table 5 reports the weighted efficiency-equity ratio $R(\varepsilon, \lambda)$ for $\lambda \in 1, 2, 5$. Whereas POE focuses specifically on the worst-off group, $R(\varepsilon, \lambda)$ uses the Gini coefficient to capture distributional equity across the full population; this difference in equity metrics explains why they tell slightly different stories about the optimal ε. Recall that $R(\varepsilon, \lambda) > 1$ indicates that the relative equity gain exceeds the λ-weighted relative coverage sacrifice, so moving to ε is worthwhile under that weight, whereas $R(\varepsilon, \lambda) < 1$ indicates that a policymaker is best off going with the utilitarian approach.

At $\lambda = 1$, treating coverage and equity losses as equally costly on their relative scales, the ratio is maximized at $\varepsilon = 1$ with $R = 2.14$, confirming that this step on the frontier delivers the most equity per unit of coverage sacrificed. All values of $\varepsilon > 0$ yield $R > 1$ at $\lambda = 1$, meaning that under equal weighting no value of ε is worse than the utilitarian baseline in benefit-cost terms. At $\lambda = 2$, reflecting a policy context in which a percentage point of coverage loss is considered twice as costly as a percentage point of Gini reduction, $\varepsilon = 1$ remains the preferred choice with $R = 1.07$, though the advantage over $\varepsilon = 0$ is now slim. Notably, $\varepsilon = 2$, 4, and ∞ all fall

Figure 15: Optimal placement of AEDs in the Central Auckland Zone for a budget of $\rho = 10\%$

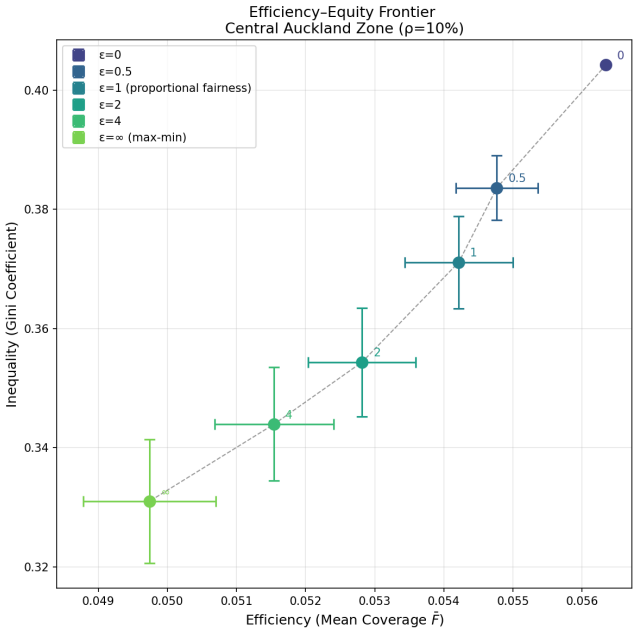


(a) $\varepsilon = 0$ **(b)** $\varepsilon = \infty$

Table 5: Central Auckland Zone ($\rho = 10\%$)

ε	$\bar{F}$ (±SE)	Gini (±SE)	POF	POE	$R(\varepsilon, 1)$	$R(\varepsilon, 2)$	$R(\varepsilon, 5)$
0	0.0563 (±0.0013)	0.4042 (±0.0081)	ref	+24.1%	—	—	—
0.5	0.0548 (±0.0013)	0.3836 (±0.0085)	+2.8%	+14.3%	1.85	0.92	0.37
1	0.0542 (±0.0014)	0.3711 (±0.0082)	+3.8%	+7.0%	2.14	1.07	0.43
2	0.0528 (±0.0013)	0.3544 (±0.0085)	+6.3%	−1.4%	1.71	0.86	0.34
4	0.0515 (±0.0013)	0.3440 (±0.0086)	+8.5%	−3.1%	1.44	0.72	0.29
∞	0.0497 (±0.0012)	0.3310 (±0.0087)	+11.7%	ref	1.21	0.61	0.24

Figure 16: The efficiency-equity frontier in the Central Auckland Zone with a budget of $\rho = 10\%$

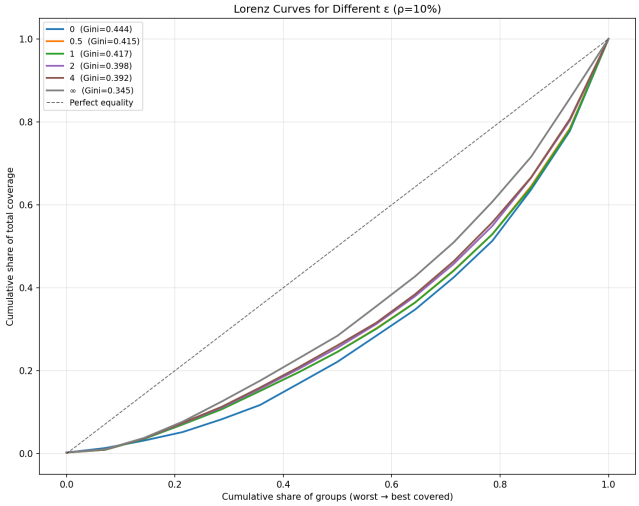


below $R = 1$ at $\lambda = 2$, meaning they are no longer justifiable under this weight. At $\lambda = 5$, representing a strongly coverage-prioritizing policymaker, all values of ε yield $R < 1$, which means no deviation from the utilitarian solution can be justified because the coverage cost of any equity improvement is simply too high relative to its benefit.

The Lorenz curves in Figure 17 provide a distributional view that complements these metrics. As ε increases the Lorenz curve shifts upwards towards the line of perfect equality. Most notably, the improvement in the Gini coefficient is not made in the lower tail of the distribution, but more in the middle of the distribution. This suggests that higher values of ε primarily redistribute coverage away from moderately well-served groups toward those in the middle, rather than exclusively targeting the absolute worst-off. The groups in the lower tail, those with the least access under any policy, appear to be structurally constrained by the existing AED deployment and the spatial distribution of OHCA incidence, such that even the max-min objective cannot substantially lift their coverage within the available budget of $\rho = 10\%$. In other words, lifting the floor of the Lorenz curve requires a larger budget for the Central Auckland Zone as an equity-aware optimization strategy is insufficient.

4.4.3 Optimal ε across zones and budget scenarios. The choice of ε determines the balance between aggregate coverage and distributional equity. However, it remains an open question whether the same value of ε performs well across different urban areas and budget levels, or whether the optimal trade-off is context-dependent. To investigate this, we solve the model for all combinations of urban zones and budget levels $\rho \in \{10\%, 20\%, 50\%, 100\%\}$, and evaluate the resulting efficiency-equity trade-offs using the associated welfare and inequality metrics. Detailed frontiers, Lorenz curves, and supplementary result tables are reported in Appendix B.

Figure 17: The Lorenz curve for different values of ε in the Central Auckland Zone with a budget of $\rho = 10\%$



To summarize the results, Table 6 reports the optimal values of ε for different values of λ that maximize $R(\varepsilon, \lambda)$. A first observation is that some of the zone-budget combinations do not yield a feasible solution. This can be understood from the structure of the social welfare function: when the coverage of the least-covered group becomes extremely small, the term $F_g(y)^{1-\varepsilon}$ grows rapidly for large values of ε, causing numerical instability and intractability in the optimization problem.

The results further reveal a broader structural pattern across Auckland. In CAZ, minimum group coverage grows consistently and substantially with budget, reaching relatively high levels at $\rho = 100\%$. This reflects the relatively dense existing AED infrastructure, which provides a meaningful coverage floor even for the worst-off groups. In contrast, NAZ, SAZ, and WAZ exhibit persistently low minimum group coverage across all feasible budgets, indicating that the most disadvantaged groups in these zones remain structurally difficult to reach regardless of how many AEDs are added. This pattern is also reflected in the underlying distributional outcomes, where improvements for the worst-off groups remain limited even under strongly equity-oriented objectives.

A second and central finding from Table 6 is that the optimal ε is highly context-dependent. In CAZ, equity-aware optimization with $\varepsilon > 0$ is consistently justifiable at $\lambda = 1$ across all budgets, with the optimal ε ranging from 0.5 to 2. This reflects that the dense existing AED infrastructure in CAZ creates meaningful room for redistribution without large coverage sacrifices. At $\lambda = 2$, equity-aware optimization remains worthwhile in CAZ at $\rho = 20\%$ and $\rho = 50\%$ ($\varepsilon^* = 0.5$), but the advantage narrows considerably at both higher and lower budget levels. At $\lambda = 5$, the utilitarian solution is always preferred in CAZ, as coverage losses are penalized heavily.

Outside CAZ, the utilitarian approach dominates in almost all zones, budgets, and values of λ. The utilitarian solutions are often substantially separated from the alternative equity-aware solutions in terms of both aggregate coverage and distributional outcomes. This may be attributed to the smaller number of AEDs in these regions, where even modest relocation decisions can lead to large shifts in coverage distributions.

A notable exception is NAZ at $\rho = 100\%$ and $\varepsilon = 2$, which achieves $R(\varepsilon, 1) = 1.78$. This suggests that at sufficiently large budgets, even NAZ can exhibit a favorable efficiency-equity trade-off at moderate levels of inequality aversion, as enough AEDs become available to meaningfully differentiate between utilitarian and equity-aware placements. This conclusion should, however, be interpreted cautiously given the relatively high standard errors in NAZ, which stem from the smaller number of OHCA events observed in the zone.

Taken together, these results indicate that the value of equity-aware AED optimization is not uniform across Auckland. It is most clearly justified in CAZ, where the existing infrastructure is dense enough to allow meaningful redistribution, and where small to moderate values of ε, particularly $\varepsilon \in \{0.5, 1, 2\}$, deliver the best overall balance between coverage and fairness across budget scenarios. In the remaining zones, the utilitarian objective is generally preferable unless substantially larger budgets become available to overcome the structural constraints that currently limit access for the worst-off groups.

5 CONCLUDING REMARKS

This study developed and applied an equity-aware AED placement optimization model for Auckland, integrating both ethnic and socioeconomic dimensions of inequality into the objective function using the Atkinson social welfare framework. The results provide several substantive insights into the nature of the efficiency-equity trade-off in AED deployment, while also raising broader questions about how fairness should be conceptualized and measured in emergency response optimization.

The key empirical finding from the Central Auckland Zone (CAZ) is that meaningful reductions in coverage inequality, as measured by the Gini coefficient, can be achieved with modest sacrifices in total probabilistic coverage. Moving from the utilitarian solution ($\varepsilon = 0$) to proportional fairness ($\varepsilon = 1$) in CAZ reduced the Gini coefficient by 3.3 percentage points at a price of fairness (POF) of only 3.8%. This result demonstrates that equity and efficiency are not necessarily in strong conflict in areas with sufficiently dense AED infrastructure. In CAZ, the relatively high density of existing AEDs creates meaningful room for redistribution, allowing the model to reallocate coverage toward underserved subgroups without forgoing substantial aggregate benefit.

Outside CAZ, the picture is considerably more mixed. In the Northern, Southern, and Western Auckland Zones, the utilitarian approach is preferable in nearly all budget and weighting scenarios. Two structural factors explain this pattern. First, the number of existing AEDs in these zones is substantially smaller, meaning that any alternative placement strategy produces larger swings in both coverage and distributional outcomes, making it difficult to achieve stable equity improvements. Second, minimum group coverage in these zones remains extremely low even under the max-min policy ($\varepsilon \to \infty$) and at the largest considered budget ($\rho = 100\%$), indicating that the worst-off subgroups in these regions are structurally constrained in a way that additional AEDs alone

Table 6: Optimal ε maximizing $R(\varepsilon, \lambda)$ for each zone and budget.

Zone	ρ	ε^*		
		$\lambda = 1$	$\lambda = 2$	$\lambda = 5$
CAZ	10%	1	1	0
	20%	0.5	0.5	0
	50%	0.5	0.5	0
	100%	2	0	0
NAZ	20%	0	0	0
	50%	2	0	0
	100%	2	0	0
SAZ	20%	0	0	0
	50%	0.5	0	0
	100%	0	0	0
WAZ	50%	0	0	0
	100%	0	0	0

cannot easily resolve. The implication is that equity-aware optimization is most valuable where the existing infrastructure already provides a meaningful coverage floor for the worst-off groups. This also the most likely explanation as to why an increase in AEDs does not always lead to a more equitable distribution. The expectation was that the underserved groups, which were also groups with higher incidence rates, would be served by the additional AEDs to maximize expected coverage. However, the same results as before applied, where these underserved groups are harder to reach and therefore a utilitarian approach still preferred the denser areas even though the incidence rate per capita was lower. These findings together underscore that simply scaling up the AED network without explicit equity constraints could amplify rather than narrow existing disparities. Policy makers should therefore be thoughtful about how to reach these underserved groups. An option could be to invest in portable AEds that could serve a larger area than the 20 to 100 meters.

I would like to end with several limitations to consider. First, the analysis relies on simulated OHCA demand derived from published incidence rates and census demographics rather than individual-level event data. While the simulation is validated against national registry estimates and incorporates demographic heterogeneity at the meshblock level, the underlying assumptions, particularly the independence of age, sex, and ethnicity within meshblocks, and the constant male-to-female incidence ratio across age groups, introduce some degree of uncertainty into the estimated demand distribution. Moreover, the use of census data implicitly assumes a static residential population, whereas individuals are not always present at their registered home locations due to commuting, work, and mobility patterns. Future work could benefit from access to geo-referenced individual-level OHCA data to improve reliability and improve applicability to real-life scenarios.

Second, this study uses the Gini coefficient and the Atkinson social welfare function as the primary measures of equity. As Marsh and Schilling [33] document extensively, a wide variety of equity measures exist in the facility location literature, each encoding different assumptions about what constitutes a fair distribution. Bertsimas et al. [7] further show that the choice of fairness criterion materially affects where the price of fairness is paid across the solution space, while Ogryczak [36] demonstrates that different inequality measures can yield conflicting rankings of the same set of allocations. It is therefore plausible that the conclusions drawn here, in particular the optimal value of ε and the shape of the efficiency-equity frontier, would differ under alternative fairness criteria. Investigating the sensitivity of results to the choice of equity metric is an important direction for future research, as it would clarify how robust the identified trade-offs are and whether different measures align in their policy recommendations.

Lastly, the model does not account for the spatial relationship between candidate AED locations and existing EMS infrastructure. In practice, placing an AED in the immediate vicinity of an ambulance station offers limited marginal benefit, as EMS responders are likely to arrive with a defibrillator before a bystander could retrieve and deploy the AED. Incorporating EMS station locations and estimated response times as constraints or penalties in the coverage function would allow the model to better differentiate between AED placements that genuinely add value to the pre-hospital response chain and those that merely overlap with existing EMS coverage. This is a meaningful extension that could improve both the realism and the policy relevance of the optimization framework.

REFERENCES

[1] AED Locations NZ. 2025. AED locations: find your nearest AED. https://aedlocations.co.nz/ Accessed: [Insert access date here].

[2] Anthony B Atkinson et al. 1970. On the measurement of inequality. *Journal of economic theory* 2, 3 (1970), 244–263.

[3] June Atkinson, Clare Salmond, and Peter Crampton. 2014. *NZDep2013 Index of Deprivation User's Manual.* Technical Report 069929. Department of Public Health, University of Otago, Wellington, Wellington, New Zealand. https://www.otago.ac.nz/wellington

[4] Tom Aufderheide, Mary Fran Hazinski, Graham Nichol, Scott S Steffens, Amy Buroker, Ryan McCune, Eric Stapleton, et al. 2006. Community lay rescuer automated external defibrillation programs. *Circulation* 113, 9 (2006), 1260–1270.

[5] Jocelyn Berdowski, Robert A. Berg, Jan G. P. Tijssen, and Rudolph W. Koster. 2010. Global Incidences of Out-of-Hospital Cardiac Arrest and Survival Rates: Systematic Review of 67 Prospective Studies. *Resuscitation* 81, 11 (2010), 1479–1487. https://doi.org/10.1016/j.resuscitation.2010.08.006

[6] Oded Berman, Dmitry Krass, and Zvi Drezner. 2003. The gradual covering decay location problem on a network. *European Journal of Operational Research* 151, 3 (2003), 474–480.

[7] Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. 2011. The price of fairness. *Operations research* 59, 1 (2011), 17–31. https://doi.org/10.1287/opre.1100.0865

[8] Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. 2012. On the efficiency-fairness trade-off. *Management Science* 58, 12 (2012), 2234–2250.

[9] Víctor Blanco and Ricardo Gázquez. 2022. Fairness in Maximal Covering Location Problems. *arXiv preprint arXiv:2204.06446* (2022).

[10] Robin Buter, Hans van Schuppen, Remy Stieglis, Hendrik Koffijberg, and Derya Demirtas. 2024. Increasing cost-effectiveness of AEDs using algorithms to optimise location. *Resuscitation* 201 (2024), 110300.

[11] Sherry L. Caffrey, Paula J. Willoughby, Paul E. Pepe, and Lance B. Becker. 2002. Public Use of Automated External Defibrillators. *New England Journal of Medicine* 347, 16 (2002), 1242–1247. https://doi.org/10.1056/NEJMoa020932 arXiv:https://www.nejm.org/doi/pdf/10.1056/NEJMoa020932

[12] Timothy C.Y. Chan, Heyse Li, Gerald Lebovic, Sabrina K. Tang, Joyce Y.T. Chan, Horace C.K. Cheng, Laurie J. Morrison, and Steven C. Brooks. 2013. Identifying Locations for Public Access Defibrillators Using Mathematical Optimization. *Circulation* 127, 17 (2013), 1801–1809. https://doi.org/10.1161/CIRCULATIONAHA.113.001953 arXiv:https://www.ahajournals.org/doi/pdf/10.1161/CIRCULATIONAHA.113.001953

[13] Timothy CY Chan, Zuo-Jun Max Shen, and Auyon Siddiq. 2018. Robust defibrillator deployment under cardiac arrest location uncertainty via row-and-column generation. *Operations Research* 66, 2 (2018), 358–379.

[14] Timothy C. Y. Chan, Derya Demirtas, and Roy H. Kwon. 2016. Optimizing the Deployment of Public Access Defibrillators. *Management Science* 63, 3 (March 2016), 732–757. https://doi.org/10.1287/mnsc.2015.2362

[15] Richard Church and Charles R Velle. 1974. The maximal covering location problem. *Papers in regional science* 32, 1 (1974), 101–118.

[16] Linda L Culley, Thomas D Rea, John A Murray, Brian Welles, Carol E Fahrenbruch, Michele Olsufka, Mickey S Eisenberg, and Michael K Copass. 2004. Public access defibrillation in out-of-hospital cardiac arrest. *Circulation* 109, 15 (2004), 1859–1863.

[17] Thi Hong Diep Dao, Yuhong Zhou, Jean-Claude Thill, and Eric Delmelle. 2012. Spatio-temporal location modeling in a 3D indoor environment: the case of AEDs as emergency medical devices. *International Journal of Geographical Information Science* 26, 3 (2012), 469–494.

[18] Bridget Dicker, Paul Davey, Tony Smith, and Ben Beck. 2018. Incidence and outcomes of out-of-hospital cardiac arrest: A New Zealand perspective. *Resuscitation* 129 (2018), 43–49. https://doi.org/10.1016/j.resuscitation.2018.05.028 This is the BibTeX entry generated based on the provided PDF content. The specific journal details (volume, pages, DOI) are inferred from common citation patterns for this paper. If you have the exact publication details, they should be verified.

[19] Bridget Dicker, Bernard Doo, Jeannette Perez, Colin Tan, Graham Howie, Tony Ward, and Paul Davey. 2014. *Out-of-Hospital Cardiac Arrest Registry Annual Report 2013/2014.* Technical Report. St John New Zealand, Auckland, New Zealand. https://www.stjohn.org.nz ISSN 2382-1528.

[20] Bridget Dicker, Nick Garrett, Samuel Wong, Helen McKenzie, John McCarthy, Gareth Jenkin, Tony Smith, Jonathan R Skinner, Tammy Pegg, Gerry Devlin, et al. 2019. Relationship between socioeconomic factors, distribution of public access defibrillators and incidence of out-of-hospital cardiac arrest. *Resuscitation* 138 (2019), 53–58.

[21] Bridget Dicker, Verity Todd, Graham Howie, Pablo Callejas, Glen Stewart, Andy Swain, and Abhishek Ranjan. 2022. *Out-of-Hospital Cardiac Arrest Registry: Aotearoa New Zealand, National Report 2020/21.* Technical Report. St John New Zealand. https://www.stjohn.org.nz/globalassets/documents/health--wellbeing/cardiac-register/st-john-ohca-national-report-2020-21_final_web.pdf Clinical Audit and Research Enquiries: CART@stjohn.org.nz.

[22] Bridget Dicker, Verity F Todd, Bronwyn Tunnage, Andy Swain, Kate Conaglen, Tony Smith, Michelle Brett, Chris Laufale, and Graham Howie. 2019. Ethnic disparities in the incidence and outcome from out-of-hospital cardiac arrest: A New Zealand observational study. *Resuscitation* 145 (2019), 58–82. https://doi.org/10.1016/j.resuscitation.2019.09.026

[23] Michel Gendreau, Gilbert Laporte, and Frédéric Semet. 1997. Solving an ambulance location model by tabu search. *Location science* 5, 2 (1997), 75–88.

[24] Mengxia Gu, Shangbo Li, and Guoquan Zheng. 2025. Optimizing spatial accessibility and equity of hierarchical older adult care facilities using a multimodal two-step floating catchment area method: a case study of Lin'an District, Hangzhou. *Frontiers in Public Health* 13 (2025), 1559463.

[25] Gurobi Optimization, LLC. 2026. Nonlinear Constraints in Gurobi Optimizer Reference Manual. https://docs.gurobi.com/projects/optimizer/en/current/features/nonlinear.html. Accessed: 2026-05-28.

[26] Orhan Karasakal and Esra K Karasakal. 2004. A maximal covering location model in the presence of partial coverage. *Computers & Operations Research* 31, 9 (2004), 1515–1526.

[27] Hideko Kono, Koichi Takaishi, Masaya Onuma, Michi Fukushima, and Ryosuke Takeuchi. 2024. Bystanders' willingness to assist using automated external defibrillators during cardiac arrest. *Heliyon* 10, 17 (2024).

[28] Mary P. Larsen, Mickey S. Eisenberg, Richard O. Cummins, and Alfred P. Hallstrom. 1993. Predicting Survival from Out-of-Hospital Cardiac Arrest: A Graphic Model. *Annals of Emergency Medicine* 22, 11 (Nov. 1993), 1652–1658. https://doi.org/10.1016/S0196-0644(05)81302-4

[29] Youngho Lee, Hanif D Sherali, Ikhyun Kwon, and Seongin Kim. 2006. A new reformulation approach for the generalized partial covering problem. *Naval Research Logistics (NRL)* 53, 2 (2006), 170–179.

[30] Ka Chun Leung, Roy H Kwon, and Timothy C Y Chan. 2021. Socioeconomically equitable public defibrillator placement using mathematical optimization. *Management Science* 67, 10 (2021), 6116–6136.

[31] Chih-Hao Lin, Kuan-Chao Chu, Jung-Ting Lee, and Chung-Yao Kao. 2023. Optimal deployment of automated external defibrillators in a long and narrow environment. *Plos one* 18, 2 (2023), e0264098.

[32] Marvin B Mandell and Les R Becker. 1996. A model for locating automatic external defibrillators. *Socio-Economic Planning Sciences* 30, 1 (1996), 51–66.

[33] Michael T Marsh and David A Schilling. 1994. Equity measurement in facility location analysis: A review and framework. *European journal of operational research* 74, 1 (1994), 1–17.

[34] Owen McBride, Amy Poel, Catherine R Counts, Megin Parayil, Camilla Osborne, Chris Drucker, Mickey Eisenberg, David Murphy, Peter Kudenchuk, Michael Sayre, et al. 2025. Temporal patterns in out-of-hospital cardiac arrest incidence and outcome. *JAMA cardiology* 10, 9 (2025), 922–931.

[35] Christian Gantzel Nielsen, Fredrik Folke, Linn Andelius, Carolina Malta Hansen, Ulla Væggemose, Erika Frischknecht Christensen, Christian Torp-Pedersen, Annette Kjær Ersbøll, and Mads Christian Tofte Gregers. 2022. Increased bystander intervention when volunteer responders attend out-of-hospital cardiac arrest. *Frontiers in Cardiovascular Medicine* 9 (2022), 1030843.

[36] Włodzimierz Ogryczak. 2000. Inequality measures and equitable approaches to location problems. *European Journal of Operational Research* 122, 2 (2000), 374–391.

[37] Richard L. Page, José A. Joglar, Robert C. Kowal, Jason D. Zagrodzky, Lynn L. Nelson, Karthik Ramaswamy, Samuel J. Barbera, Mohamed H. Hamdan, and David K. McKenas. 2000. Use of Automated External Defibrillators by a U.S. Airline. *New England Journal of Medicine* 343, 17 (2000), 1210–1216. https://doi.org/10.1056/NEJM200010263431702

[38] Bozidar Radunovic and Jean-Yves Le Boudec. 2007. A unified framework for max-min and min-max fairness with applications. *IEEE/ACM Transactions on networking* 15, 5 (2007), 1073–1083.

[39] Charles ReVelle, Constantine Toregas, and Louis Falkson. 1976. Applications of the Location Set-covering Problem. *Geographical Analysis* 8, 1 (1976), 65–76. https://doi.org/10.1111/j.1538-4632.1976.tb00529.x arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/j.1538-4632.1976.tb00529.x

[40] Femke van de Sandt and Victor Umans. 2009. Acute cardiac events and deployment of emergency medical teams and automated external defibrillators in large football stadiums in the Netherlands. *European journal of cardiovascular prevention and rehabilitation* 16, 5 (10 2009), 571–575. https://doi.org/10.1097/HJR.0b013e32832d1927 arXiv:https://academic.oup.com/eurjpc/article-pdf/16/5/571/34358333/eurjpc0571.pdf

[41] Comilla Sasson, Carla C Keirns, Dylan Smith, Michael Sayre, Michelle Macy, William Meurer, Bryan F McNally, Arthur L Kellermann, Theodore J Iwashyna, and CARES (Cardiac Arrest Registry to Enhance Survival) Study Group. 2010. Small area variations in out-of-hospital cardiac arrest: does the neighborhood matter? *Annals of internal medicine* 153, 1 (2010), 19–22.

[42] Comilla Sasson, David J. Magid, Paul Chan, Elisabeth D. Root, Bryan F. McNally, Arthur L. Kellermann, and Jason S. Haukoos. 2012. Association of Neighborhood Characteristics with Bystander-Initiated CPR. *New England Journal of Medicine* 367, 17 (2012), 1607–1615. https://doi.org/10.1056/NEJMoa1110700

arXiv:https://www.nejm.org/doi/pdf/10.1056/NEJMoa1110700
[43] Comilla Sasson, Mary A. M. Rogers, Jason Dahl, and Arthur L. Kellermann. 2010. Predictors of Survival from Out-of-Hospital Cardiac Arrest: A Systematic Review and Meta-Analysis. *Circulation: Cardiovascular Quality and Outcomes* 3, 1 (2010), 63–81. https://doi.org/10.1161/CIRCOUTCOMES.109.889576
[44] Eirik Skogvoll and Bo Henry Lindqvist. 1999. Modeling the occurrence of cardiac arrest as a Poisson process. *Annals of emergency medicine* 33, 4 (1999), 409–417.
[45] Paul A. Smith, Peter G. van der Heijden, Maarten Cruyff, Francesco Pantalone, Hannes Diener, and Kim Dunstan. 2023. *Estimating population size of four ethnic groupings in New Zealand.* Technical Report. Stats NZ Tatauranga Aotearoa, Wellington, New Zealand. https://www.stats.govt.nz ISBN 978-1-99-104943-8 (online).
[46] Lawrence V Snyder. 2006. Facility location under uncertainty: a review. *IIE transactions* 38, 7 (2006), 547–564.
[47] Christopher LF Sun, Derya Demirtas, Steven C Brooks, Laurie J Morrison, and Timothy CY Chan. 2016. Overcoming spatial and temporal barriers to public access defibrillators via optimization. *Journal of the American College of Cardiology* 68, 8 (2016), 836–845.
[48] Christopher LF Sun, Lena Karlsson, Christian Torp-Pedersen, Laurie J Morrison, Fredrik Folke, and Timothy CY Chan. 2018. Spatiotemporal AED optimization is generalizable. *Resuscitation* 131 (2018), 101–107.
[49] Renata Turkeš, Kenneth Sörensen, and Daniel Palhazi Cuervo. 2021. A matheuristic for the stochastic facility location problem. *Journal of Heuristics* 27, 4 (2021), 649–694.
[50] C. Vaillancourt, A. Lui, V.J. De Maio, G.A. Wells, and I.G. Stiell. 2008. Socioeconomic status influences bystander CPR and survival rates for out-of-hospital cardiac arrest victims. *Resuscitation* 79, 3 (2008), 417–423. https://doi.org/10.1016/j.resuscitation.2008.07.012
[51] Terence D Valenzuela, Denise J Roe, Shan Cretin, Daniel W Spaite, and Mary P Larsen. 1997. Estimating effectiveness of cardiac arrest interventions: a logistic regression survival model. *Circulation* 96, 10 (1997), 3308–3313.
[52] Terence D. Valenzuela, David J. Roe, Graham Nichol, Lani Clark, Daniel W. Spaite, and Ronald G. Hardman. 2000. Outcomes of Rapid Defibrillation by Security Officers After Cardiac Arrest in Casinos. *New England Journal of Medicine* 343, 17 (2000), 1206–1209. https://doi.org/10.1056/NEJM200010263431701
[53] Pieter L van den Berg, Shane G Henderson, Caroline J Jagtenberg, and Hemeng Li. 2025. Modeling the impact of community first responders. *Management Science* 71, 2 (2025), 992–1008.
[54] Benjamin P. van Nieuwenhuizen, Jacoline van de Ven, Ingrid Oving, Irene van Valkengoed, Mardien de Waard, Esther de Vries, Désirée de la Fosse, Monique van Dijk, M. Christien van der Linden, Barbara van der Nagel, and Marieke T. Blom. 2019. Socio-Economic Differences in Incidence, Bystander Cardiopulmonary Resuscitation and Survival from Out-of-Hospital Cardiac Arrest: A Systematic Review. *Resuscitation* 141 (Aug. 2019), 44–62. https://doi.org/10.1016/j.resuscitation.2019.05.018
[55] Reinier A. Waalewijn, Rien de Vos, Jan G.P. Tijssen, and Rudolph W. Koster. 2001. Survival models for out-of-hospital cardiopulmonary resuscitation from the perspectives of the bystander, the first responder, and the paramedic. *Resuscitation* 51, 2 (Nov. 2001), 113–122. https://doi.org/10.1016/S0300-9572(01)00407-5
[56] Myron L Weisfeldt, Richard E Kerber, R Pat McGoldrick, Arthur J Moss, Graham Nichol, Joseph P Ornato, David G Palmer, Barbara Riegel, and Sidney C Smith Jr. 1995. Public access defibrillation: a statement for healthcare professionals from the American Heart Association Task Force on Automatic External Defibrillation. *Circulation* 92, 9 (1995), 2763–2763.
[57] Myron L. Weisfeldt, Colleen M. Sitlani, Joseph P. Ornato, Thomas Rea, Tom P. Aufderheide, Daniel Davis, Jonathan Dreyer, Erik P. Hess, Jonathan Jui, Justin Maloney, George Sopko, Judy Powell, Graham Nichol, Laurie J. Morrison, and null null. 2010. Survival After Application of Automatic External Defibrillators Before Arrival of the Emergency Medical System. *JACC* 55, 16 (2010), 1713–1720. https://doi.org/10.1016/j.jacc.2009.11.077 arXiv:https://www.jacc.org/doi/pdf/10.1016/j.jacc.2009.11.077
[58] Michael KY Wong, Laurie J Morrison, Feng Qiu, Peter C Austin, Sheldon Cheskes, Paul Dorian, Damon C Scales, Jack V Tu, P Richard Verbeek, Harindra C Wijeysundera, et al. 2014. Trends in short-and long-term survival among out-of-hospital cardiac arrest patients alive at hospital arrival. *Circulation* 130, 21 (2014), 1883–1890.
[59] Jing Yao and Alan T Murray. 2014. Serving regional demand in facility location. *Papers in Regional Science* 93, 3 (2014), 643–663.

6 APPENDIX A: BACKGROUND DATA ON DEPRIVATION QUINTILES AND INCIDENCE RATES

Table 7: Characteristics of NZDep2013 deprivation quintiles

	NZDep2013 quintile (1 = least deprived, 5 = most deprived)					
	1	**2**	**3**	**4**	**5**	**Total**
Number of meshblocks	2,023	2,026	1,961	1,991	1,983	9,984
Population	244,500	245,200	233,800	244,900	256,500	1,224,900
Area (km²)	305.3	258.5	184.2	127.3	101.2	976.5
Population density (residents/km²)	800.9	948.5	1269.3	1923.8	2534.6	1254.4
Number of suspected OHCAs	151	224	236	310	412	1,333
Number of existing AEDs	265	343	452	549	458	2,067

Table 8: Population characteristics and OHCA incidences by ethnicity in Auckland

Ethnicity	Population (2013)	% of stated population	AEDs (total)	AEDs per 100,000 residents	Suspected OHCAs (total)	Suspected OHCAs per 100,000 person-years
European/other	904,167	73.8%	1,666	184.3	821	90.8
Māori	129,666	10.6%	205	157.9	243	187.4
Pacific	191,055	15.6%	219	114.4	284	148.6

Table 9: Original age-adjusted OHCA incidence rates per 100,000 person-years from Dicker et al. (2019)

Age Group	European/Other	Māori	Pacific Peoples
0-4	0.8	6.0	5.3
5-14	0.5	2.4	0.2
15-24	3.0	11.0	4.0
25-44	8.0	24.0	18.0
45-64	24.5	56.0	41.0
65+	57.0	45.0	45.0
Total	93.8	144.4	113.5

Table 10: Estimated female OHCA incidence rates per 100,000 Person-Years

Age Group	European/Other	Māori	Pacific Peoples
0-4	2.94	25.14	23.28
5-14	1.80	10.08	0.90
15-24	10.92	46.14	17.58
25-44	29.16	100.68	79.14
45-64	89.40	234.84	180.30
65+	207.96	188.76	197.94
Average	57.03	100.94	83.19

Table 11: Estimated male OHCA incidence rates per 100,000 Person-Years

Age Group	European/Other	Māori	Pacific Peoples
0-4	6.66	46.86	40.32
5-14	4.20	18.72	1.50
15-24	25.08	85.86	30.42
25-44	66.84	187.32	136.86
45-64	204.60	437.16	311.70
65+	476.04	351.30	342.12
Average	130.57	187.87	143.82

7 APPENDIX B: DETAILED RESULTS PER BUDGET PER ZONE

Table 12: Central Auckland Zone (ρ = 20%)

ε	$\bar{F}$ (±SE)	Gini (±SE)	POF	POE	$R(\varepsilon, 1)$	$R(\varepsilon, 2)$	$R(\varepsilon, 5)$
0	0.0633 (±0.0015)	0.3796 (±0.0084)	ref	+27.1%	1.00	1.00	**1.00**
0.5	0.0626 (±0.0015)	0.3535 (±0.0079)	+1.0%	+9.9%	**2.57**	**1.28**	0.51
1	0.0614 (±0.0014)	0.3382 (±0.0081)	+3.0%	+2.1%	1.69	0.84	0.34
2	0.0608 (±0.0015)	0.3348 (±0.0080)	+3.8%	+1.6%	1.34	0.67	0.27
4	0.0599 (±0.0014)	0.3225 (±0.0078)	+5.4%	−1.9%	1.14	0.57	0.23
∞	0.0570 (±0.0013)	0.3007 (±0.0080)	+9.9%	ref	1.00	0.50	0.20

Table 13: Central Auckland Zone (ρ = 50%)

ε	$\bar{F}$ (±SE)	Gini (±SE)	POF	POE	$R(\varepsilon, 1)$	$R(\varepsilon, 2)$	$R(\varepsilon, 5)$
0	0.0832 (±0.0017)	0.2990 (±0.0086)	ref	+15.7%	1.00	1.00	**1.00**
0.5	0.0826 (±0.0017)	0.2822 (±0.0085)	+0.6%	+7.2%	**2.69**	**1.34**	0.54
1	0.0810 (±0.0016)	0.2732 (±0.0088)	+2.5%	+4.2%	1.19	0.60	0.24
2	0.0815 (±0.0016)	0.2710 (±0.0086)	+2.0%	+2.0%	1.43	0.71	0.29
4	0.0808 (±0.0016)	0.2657 (±0.0089)	+2.8%	+2.7%	1.23	0.61	0.25
∞	0.0794 (±0.0015)	0.2558 (±0.0089)	+4.6%	ref	0.99	0.50	0.20

Table 14: Central Auckland Zone (ρ = 100%)

ε	$\bar{F}$ (±SE)	Gini (±SE)	POF	POE	$R(\varepsilon, 1)$	$R(\varepsilon, 2)$	$R(\varepsilon, 5)$
0	0.1153 (±0.0019)	0.2619 (±0.0076)	ref	+12.2%	1.00	**1.00**	**1.00**
0.5	0.1147 (±0.0019)	0.2563 (±0.0075)	+0.5%	+9.5%	1.04	0.52	0.21
1	0.1136 (±0.0018)	0.2525 (±0.0072)	+1.5%	+9.4%	0.68	0.34	0.14
2	0.1137 (±0.0019)	0.2447 (±0.0072)	+1.4%	+6.2%	**1.15**	0.57	0.23
4	0.1124 (±0.0018)	0.2388 (±0.0073)	+2.5%	+5.4%	0.92	0.46	0.18
∞	0.1094 (±0.0018)	0.2190 (±0.0067)	+5.1%	ref	0.85	0.42	0.17

Table 15: Northern Auckland Zone (ρ = 20%)

ε	$\bar{F}$ (±SE)	Gini (±SE)	POF	POE	$R(\varepsilon, 1)$	$R(\varepsilon, 2)$	$R(\varepsilon, 5)$
0	0.0177 (±0.0008)	0.5543 (±0.0159)	ref	−22.7%	**1.00**	**1.00**	**1.00**
0.5	0.0173 (±0.0007)	0.5488 (±0.0159)	+2.2%	−2.8%	0.44	0.22	0.09
1	0.0172 (±0.0007)	0.5499 (±0.0161)	+2.4%	+4.8%	0.33	0.17	0.07
2	0.0171 (±0.0007)	0.5470 (±0.0167)	+3.0%	+12.0%	0.40	0.20	0.08
4	0.0169 (±0.0007)	0.5439 (±0.0174)	+4.1%	+7.6%	0.38	0.19	0.08
∞	0.0171 (±0.0007)	0.5516 (±0.0167)	+3.0%	ref	0.21	0.10	0.04

Table 16: Northern Auckland Zone (ρ = 50%)

ε	$\bar{F}$ (±SE)	Gini (±SE)	POF	POE	$R(\varepsilon, 1)$	$R(\varepsilon, 2)$	$R(\varepsilon, 5)$
0	0.0269 (±0.0008)	0.5038 (±0.0165)	ref	+9.2%	1.00	**1.00**	**1.00**
0.5	0.0263 (±0.0008)	0.4887 (±0.0146)	+1.9%	−1.9%	0.73	0.37	0.15
1	0.0260 (±0.0009)	0.4870 (±0.0149)	+3.1%	+16.5%	0.48	0.24	0.10
2	0.0265 (±0.0009)	0.4917 (±0.0165)	+1.3%	−5.6%	**1.27**	0.64	0.25
4	0.0266 (±0.0009)	0.4975 (±0.0159)	+1.0%	+2.6%	0.73	0.36	0.15
∞	0.0264 (±0.0010)	0.4844 (±0.0155)	+1.8%	ref	0.97	0.48	0.19

Table 17: Northern Auckland Zone (ρ = 100%)

ε	$\bar{F}$ (±SE)	Gini (±SE)	POF	POE	$R(\varepsilon, 1)$	$R(\varepsilon, 2)$	$R(\varepsilon, 5)$
0	0.0382 (±0.0012)	0.4214 (±0.0138)	ref	+2.9%	1.00	**1.00**	**1.00**
0.5	0.0378 (±0.0012)	0.4159 (±0.0132)	+1.1%	−1.6%	0.97	0.48	0.19
1	0.0373 (±0.0012)	0.4130 (±0.0132)	+2.4%	−5.4%	0.62	0.31	0.12
2	0.0379 (±0.0012)	0.4087 (±0.0129)	+0.7%	−1.5%	**1.78**	0.89	0.36
4	0.0379 (±0.0012)	0.4109 (±0.0130)	+0.8%	+0.2%	1.22	0.61	0.24
∞	0.0375 (±0.0012)	0.4147 (±0.0137)	+1.9%	ref	0.49	0.24	0.10

Table 18: Southern Auckland Zone (ρ = 20%)

ε	$\bar{F}$ (±SE)	Gini (±SE)	POF	POE	$R(\varepsilon, 1)$	$R(\varepsilon, 2)$	$R(\varepsilon, 5)$
0	0.0161 (±0.0006)	0.5129 (±0.0144)	ref	+78.4%	**1.00**	**1.00**	**1.00**
0.5	0.0151 (±0.0005)	0.4756 (±0.0149)	+6.4%	+63.4%	0.68	0.34	0.14
1	0.0151 (±0.0005)	0.4638 (±0.0156)	+6.8%	+50.6%	0.71	0.35	0.14
2	0.0148 (±0.0005)	0.4617 (±0.0152)	+8.1%	+39.6%	0.64	0.32	0.13
4	0.0150 (±0.0005)	0.4539 (±0.0152)	+7.4%	+34.7%	0.78	0.39	0.16
∞	0.0149 (±0.0006)	0.4584 (±0.0200)	+7.6%	ref	0.71	0.36	0.14

Table 19: Southern Auckland Zone (ρ = 50%)

ε	$\bar{F}$ (±SE)	Gini (±SE)	POF	POE	$R(\varepsilon, 1)$	$R(\varepsilon, 2)$	$R(\varepsilon, 5)$
0	0.0269 (±0.0008)	0.4699 (±0.0116)	ref	+73.8%	1.00	**1.00**	**1.00**
0.5	0.0252 (±0.0007)	0.4053 (±0.0107)	+6.1%	+22.9%	**1.05**	0.52	0.21
1	0.0242 (±0.0007)	0.3955 (±0.0112)	+10.1%	+3.8%	0.78	0.39	0.16
2	0.0244 (±0.0007)	0.3940 (±0.0102)	+9.3%	+6.7%	0.83	0.42	0.17
4	0.0240 (±0.0007)	0.3883 (±0.0100)	+10.6%	−0.7%	0.81	0.41	0.16
∞	0.0237 (±0.0007)	0.3868 (±0.0100)	+11.7%	ref	0.75	0.38	0.15

Table 20: Southern Auckland Zone (ρ = 100%)

ε	$\bar{F}$ (±SE)	Gini (±SE)	POF	POE	$R(\varepsilon, 1)$	$R(\varepsilon, 2)$	$R(\varepsilon, 5)$
0	0.0433 (±0.0009)	0.4160 (±0.0113)	ref	+55.8%	**1.00**	**1.00**	**1.00**
0.5	0.0397 (±0.0010)	0.3606 (±0.0106)	+8.3%	+24.3%	0.90	0.45	0.18
1	0.0391 (±0.0009)	0.3593 (±0.0104)	+9.6%	+26.3%	0.76	0.38	0.15
2	0.0391 (±0.0009)	0.3522 (±0.0104)	+9.5%	+20.1%	0.84	0.42	0.17
4	0.0390 (±0.0009)	0.3502 (±0.0100)	+9.8%	+19.9%	0.85	0.42	0.17
∞	0.0386 (±0.0009)	0.3352 (±0.0101)	+10.7%	ref	0.93	0.47	0.19

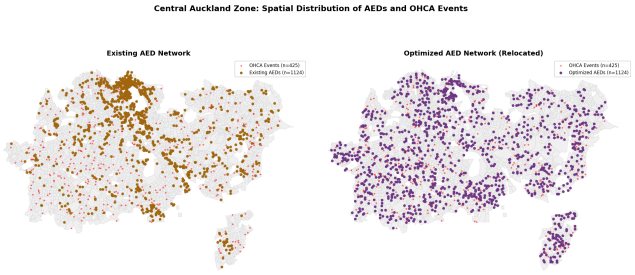
Table 21: Western Auckland Zone (ρ = 50%)

ε	$\bar{F}$ (±SE)	Gini (±SE)	POF	POE	$R(\varepsilon, 1)$	$R(\varepsilon, 2)$	$R(\varepsilon, 5)$
0	0.0155 (±0.0009)	0.5294 (±0.0146)	ref	+40.0%	**1.00**	**1.00**	**1.00**
0.5	0.0153 (±0.0009)	0.5218 (±0.0139)	+1.1%	+35.6%	0.46	0.23	0.09
1	0.0149 (±0.0009)	0.5191 (±0.0143)	+4.0%	+20.7%	0.31	0.16	0.06
2	0.0147 (±0.0009)	0.5219 (±0.0132)	+5.1%	+20.1%	0.18	0.09	0.04
4	0.0147 (±0.0009)	0.5173 (±0.0135)	+5.2%	+4.4%	0.27	0.14	0.05
∞	0.0139 (±0.0012)	0.5221 (±0.0198)	+10.4%	ref	0.09	0.04	0.02

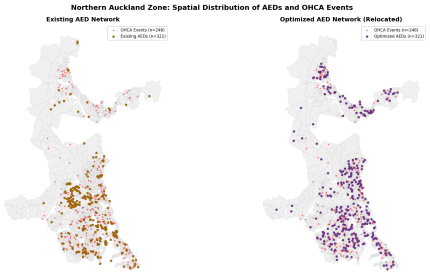
Table 22: Western Auckland Zone ($\rho = 100\%$)

ε	$\bar{F}$ (±SE)	Gini (±SE)	POF	POE	$R(\varepsilon, 1)$	$R(\varepsilon, 2)$	$R(\varepsilon, 5)$
0	0.0237 (±0.0011)	0.4965 (±0.0136)	ref	+23.7%	**1.00**	**1.00**	**1.00**
0.5	0.0229 (±0.0011)	0.4778 (±0.0133)	+3.6%	−7.9%	0.48	0.24	0.10
1	0.0225 (±0.0010)	0.4736 (±0.0131)	+5.1%	−12.4%	0.46	0.23	0.09
2	0.0223 (±0.0011)	0.4731 (±0.0123)	+6.2%	−2.4%	0.40	0.20	0.08
4	0.0218 (±0.0010)	0.4657 (±0.0117)	+8.2%	−7.7%	0.40	0.20	0.08
∞	0.0216 (±0.0009)	0.4663 (±0.0109)	+8.9%	ref	0.37	0.19	0.07

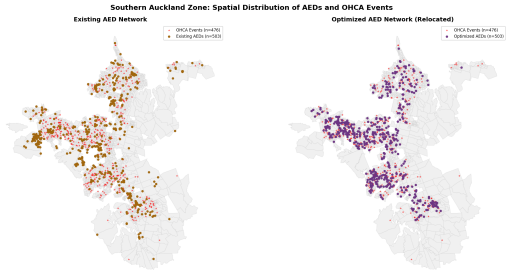
Figure 18: Optimized AED placements across Auckland's urban areas



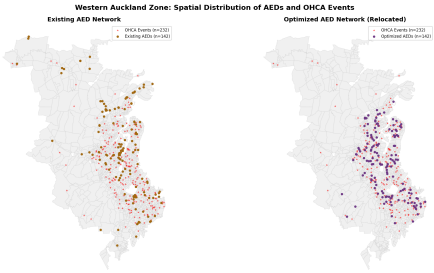
(a) Central Auckland Zone (CAZ)



(b) Northern Auckland Zone (NAZ)



(c) Southern Auckland Zone (SAZ)



(d) Western Auckland Zone (WAZ)

Figure 19: The min-group coverage under the max-min policy ($\min_g F_g(y^*(\varepsilon = \infty))$) for different budget scenarios and urban areas

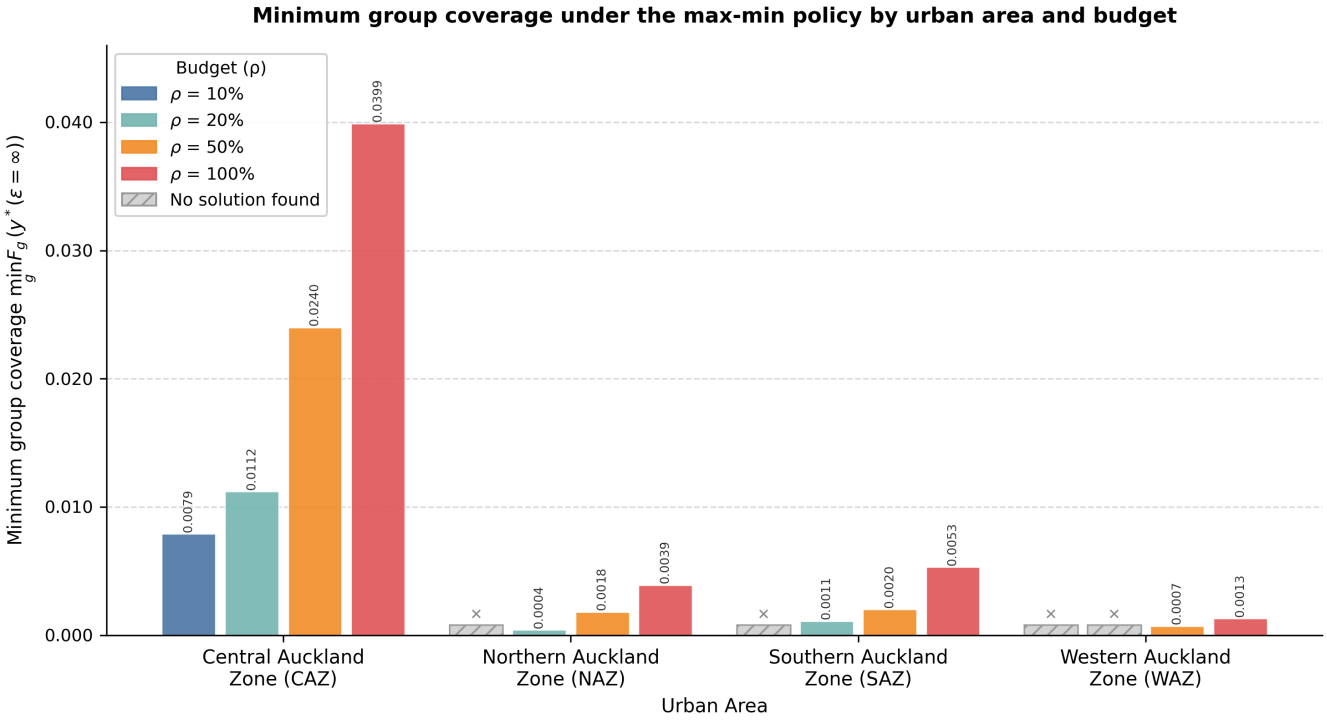
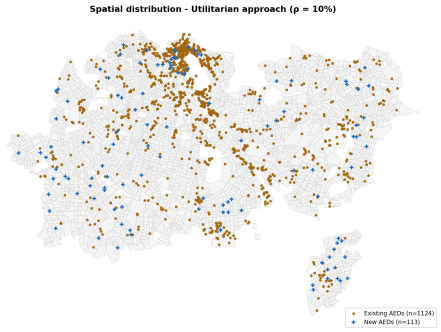
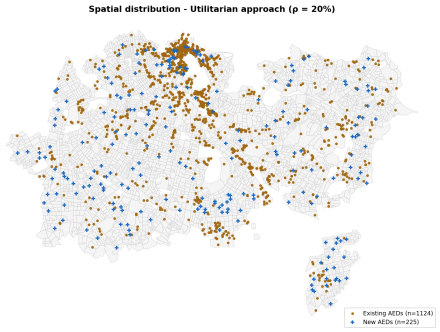


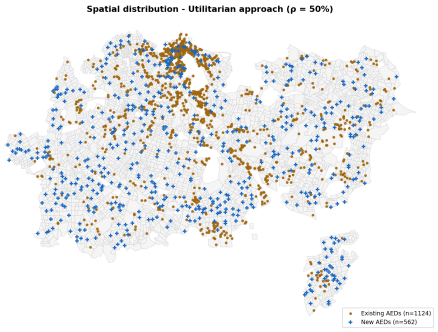
Figure 20: Optimized AED placements under the utilitarian objective for the Central Auckland Zone



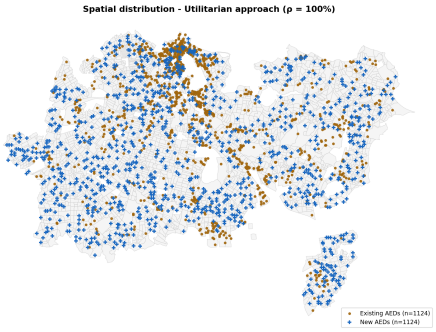
(a) $\rho = 10\%$



(b) $\rho = 20\%$



(c) $\rho = 50\%$



(d) $\rho = 100\%$

Figure 21: The efficiency-equity frontier for all combinations of urban area and budget. Optimal ε values are indicated using colored circles: green ($\lambda = 1$), orange ($\lambda = 2$) and red ($\lambda = 5$)

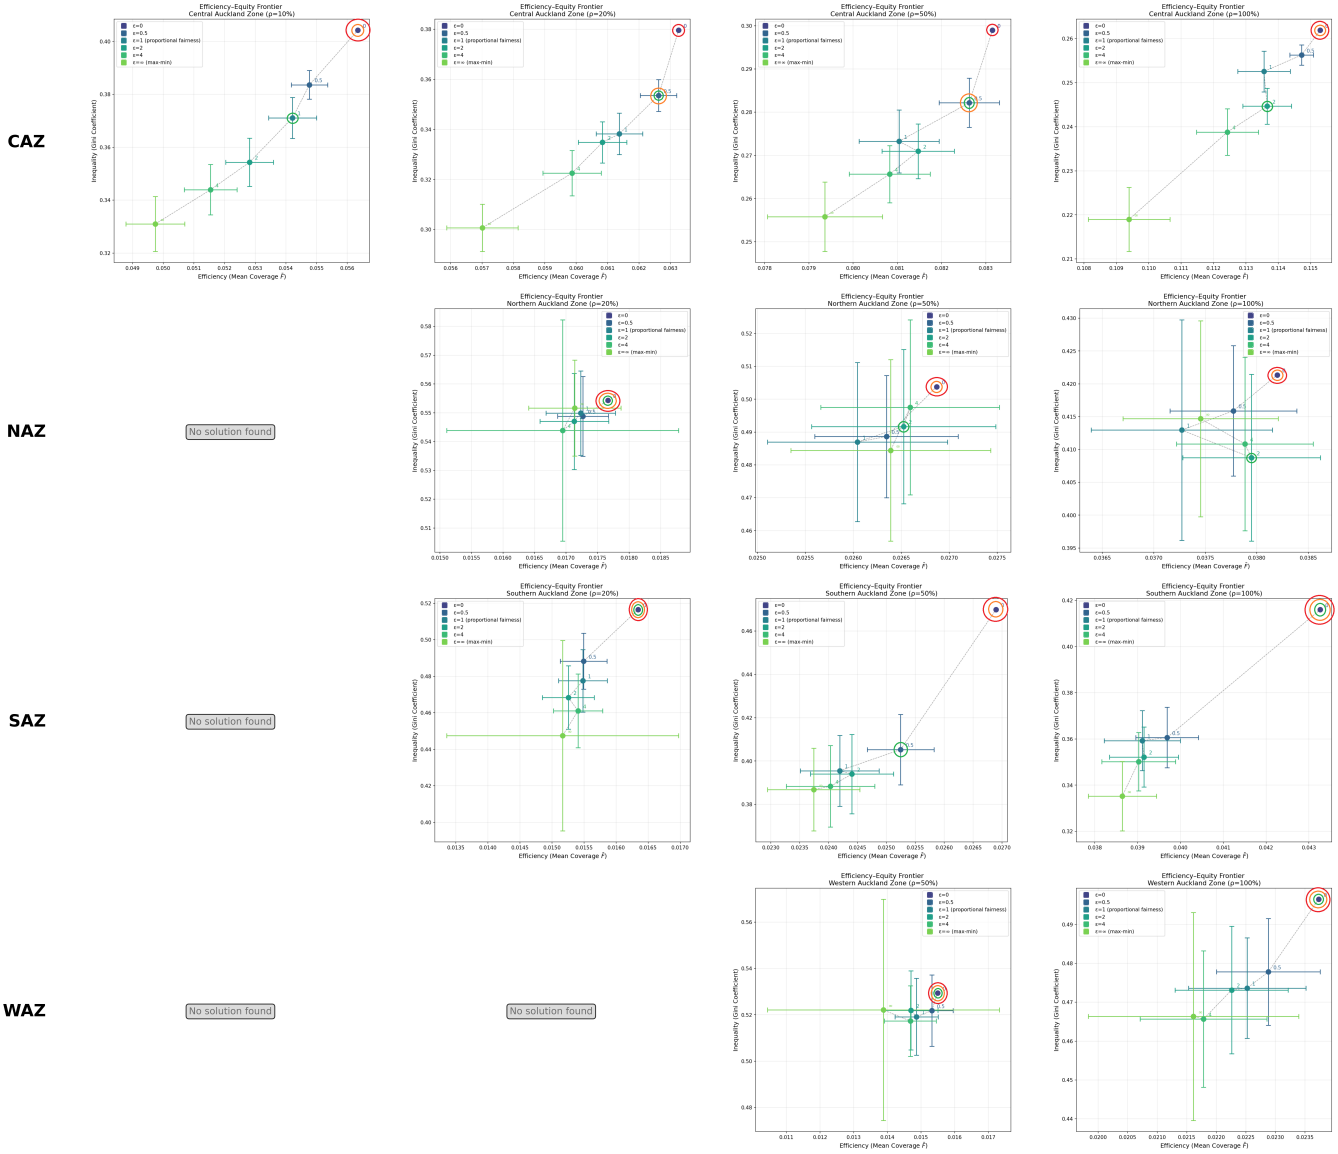


Figure 22: The efficiency-equity frontier for all combinations of urban area and budget. Optimal ε values are indicated using colored circles: green ($\lambda = 1$), orange ($\lambda = 2$) and red ($\lambda = 5$)

