

Optimal Distribution of Medical Backpacks and Health Surveillance Assistants in Malawi

Amber G. Kunkel · Elizabeth S. Van Itallie · Duo Wu

Abstract Despite recent progress, Malawi continues to perform poorly on key health indicators such as child mortality and life expectancy. These problems are exacerbated by a severe lack of access to health care. Health Surveillance Assistants (HSAs) help bridge this gap by providing community-level access to basic health care services. However, the success of these HSAs is limited by a lack of supplies and long distances between HSAs and patients. To address this issue, we used large-scale weighted p -median and capacitated facility location problems to create a scalable, three-tiered plan for optimal allocation of HSAs, HSA designated medical backpacks, and backpack resupply centers. Our analysis uses real data on the location and characteristics of hospitals, health centers, and the general population. In addition to offering specific recommendations for HSA, backpack, and resupply center locations, it provides general insights into the scope of the proposed HSA backpack program scale-up. In particular, it demonstrates the importance of local health centers to the resupply network. The proposed assignments are robust to changes in the underlying population structure, and could significantly improve access to medical supplies for both HSAs and patients.

Keywords Malawi · global health · resource allocation · p -median problem · capacitated facility location problem

1 Introduction

Malawi is an African nation whose citizens face severely limited health care access: only 2 physicians/100,000 people, compared to 267 physicians/100,000 people in the United States [3]. With only 7,264 nurses and no midwives in the entire country, Malawi is well below the minimum density of 228 health workers (physicians, nurses, and midwives) per 100,000 people required to provide adequate primary health care to its people [2]. Malawi's performance on key health indicators, including an under-five mortality rate of 13.4% and an infant mortality rate of 7.7% [13], reflects this lack of access to care.

Malawi's well-established network of community health workers, or Health Surveillance Assistants (HSAs), is key to addressing this crisis. Since the 1980s, the Malawi Ministry of Health (MOH) has trained and paid HSAs to provide primary care and health education to remote regions in Malawi [12]. HSAs are responsible for immunization, growth monitoring, health talks, and sanitation in their assigned villages. Some HSAs also provide HIV counseling and testing, family planning services, and treatment of minor diseases. There are currently about 11,000 HSAs in Malawi, each serving an average of 1,200 people [13] and typically working in groups of two or more. Overall, HSAs have played a significant role in bridging the gap between formal preventive health services and Malawian villages [12].

However, the success of the HSA program is threatened by a lack of transport and supplies. HSAs are typically based out of one village, and routinely travel 5 km

A.G. Kunkel
Harvard School of Public Health, Department of Epidemiology,
677 Huntington Ave, Boston, MA 02215
Tel.: 1-804-212-6762
E-mail: agkunkel@gmail.com

E. Van Itallie
Harvard University, Department of Systems Biology, Boston,
MA 02215

D. Wu
Deloitte Consulting LLP, Houston, TX 77002

or more on foot to patients in other villages [12]. Fifty-four percent of mothers sampled in 2001 said lack of access prevented them from consulting HSAs about their children’s health. Furthermore, 88% of HSAs sampled said they had never been issued any drug stocks [12]. To improve patients’ access to HSAs, the Malawi MOH is planning to scale up the current HSA program to achieve a ratio of 1 HSA to 1,000 people (approximately 13,000 HSAs, or an increase of 18%) [13]. Additionally, Rice University’s Beyond Traditional Borders initiative (BTB) has developed a Health Surveillance Assistant backpack over the last four years to supply these HSAs with the tools necessary to perform their jobs. These backpacks contain basic diagnostic, treatment, and prevention tools designed specifically to address the responsibilities of the HSAs. Fourteen of such backpacks have been successfully field tested at St. Gabriel’s Hospital in Namitete, Malawi since 2009, where each backpack is shared by between five and twenty community health workers. BTB hopes to eventually scale up the backpack program to provide every HSA in Malawi with access to a backpack, with no more than six HSAs sharing a single backpack. We assumed that each backpack would be based with one pair of HSAs, who would be responsible for taking the backpack to be resupplied once a month, and that the other HSAs would travel to retrieve and then return the backpack approximately once a week. We expect that the backpacks would travel with HSAs during the day and be stored in a village leader’s house overnight for safekeeping.

To improve the effectiveness of both the HSA and backpack programs, we sought a scale-up plan for the HSA and HSA backpack programs that would optimally

- place HSAs across Malawi, with the total number of HSAs based on a ratio of 1 HSA to 1000 people;
- assign backpacks to HSAs and estimate the total number of backpacks required;
- choose backpack resupply centers from the existing network of hospitals and health centers.

We initially formulated the problem as three facility location problems, specifically a weighted p -median problem for HSA assignment and two capacitated facility location problems (CFLPs) for backpacks and resupply centers. Our primary goals for this stage of the analysis were to determine the scope of the proposed backpack scale-up and the potential impact of the backpacks on the HSAs’ access to medical supplies. We solved the optimization models using extensive real-world data and were able to determine the total cost for providing HSA backpacks across Malawi; quantify the effects the backpacks would have on Malawi’s health

system; and gain insights into the optimal backpack resupply network. Implemented according to our models, the backpack program could decrease the distance between HSAs and their nearest supply sources from 5.2 km to 1.8 km. Our models estimate that 2188 backpacks would be needed to sufficiently cover all of Malawi, at a cost of \$765,800. This high cost persuaded BTB to focus its initial scale-up plans on Lilongwe District, the district in which our model predicted the greatest need for backpacks. We therefore have included the results of our initial models for both Lilongwe District (not including the city of Lilongwe) and the entire country of Malawi in this paper. We provide additional sensitivity analyses and model refinements specifically for Lilongwe District. In particular, we used a weighted p -center problem to assess the fairness of the HSA assignments and the two-echelon CFLP to verify the efficiency of the backpack and resupply center assignments; neither change led to substantial improvement in our initial results. To our knowledge, this is the first analysis of the optimal dissemination of technology among community health workers in Malawi. The quantitative and graphical results included in this paper have already been used by BTB to facilitate talks with Malawian officials, and our analysis will provide a basis for soliciting support from international aid agencies for the scale-up of the HSA backpack program.

2 Optimization Problems in Global Health Resource Allocation

Other researchers have successfully applied operations research techniques to different aspects of the distribution of health resources in developing countries. In particular, several modeling papers have sought to optimally utilize community health workers and locate health care facilities in countries with problems similar to those in Malawi.

Several researchers have examined health system efficiency at the community health worker level. Brunskill and Lesh [5] suggested house visits be treated as a routing and scheduling problem that could be approached as a traveling salesman problem with time windows, with additional complexity caused by the possibility of future follow-up visits. Meanwhile, Parker et al. [16] presented a simple model to prioritize tasks for community health workers in Haiti. Finally, Doerner et al. [7] helped bridge the gap between community health workers and physical health center locations by using a multiobjective combinatorial optimization formulation for mobile health centers and evaluating that model within the Thies region of Senegal. Like these studies, our model is interested in improving the performance

of community health workers; however, it attempts to do so specifically by improving their access to medical supplies. This distinction makes our work more similar to studies focused on locating health facilities than those focused on managing health care workers.

We found a number of such studies on health care facility location within developing countries. For example, Cocking et al. [6] proposed a coverage model for locating new health facilities in Burkina Faso, Africa. Similarly, Reid et al. [19] used a location set covering algorithm to determine optimal locations of medical supply centers in Ecuador. Rahman and Smith [17] provided a review of articles using location-allocation models for locating health centers in developing nations, focusing on their use in locating new sites, measuring effectiveness of past decisions, and improving existing systems. Many but not all of the studies discussed focused on a single tier of facilities and did not consider capacity constraints. While Murawski and Church [15] proposed an integer programming model taking into account possible future improvements in road infrastructure, we retained a static structure within our models, made necessary by the lack of data available about even current road conditions in Malawi. Nevertheless, many of these models followed an integer programming formulation similar to our approach, and they helped provide a conceptual basis for our mathematical analysis.

3 Methods

In this section, we provide a brief summary of the mathematical models used in our analyses along with data and computational details. More detailed modeling descriptions are provided within the Project Assessment and Model Analysis sections.

3.1 Mathematical Models

We selected two mathematical models, the p -median problem and the capacitated facility location problem (CFLP) for our three-tiered distribution problem of HSAs, backpacks, and resupply centers. The p -median problem, which minimizes the average weighted distance between demand points and providers, is appropriate when travel distance is a main concern, such as in locating emergency medical facilities [11], and is one of the most commonly used models for locating health-care facilities in developing countries [17]. We therefore chose this model to assign HSAs to population centers. The CFLP, meanwhile, is used when there are explicit constraints on the number of demand nodes a facility can serve but no designated number of facilities [10].

Within the context of our problem, the CFLP was most appropriate for assigning backpacks and resupply centers, both of which have space and equipment limits which could be exhausted if overextended. We initially maintained separate models for each tier of the analysis to make the problem more computationally tractable and simplify future analyses, should new data become available. Additional details about these models are described in Section 4, Project Assessment. In Section 5, Model Analysis, we compare the results of these models for Lilongwe District to a p -center formulation of the HSA assignment problem and a two-echelon CFLP for backpack and resupply center allocation.

3.2 Data Sources

We obtained geographical information system (GIS) files delineating the location and border of enumeration areas (EAs, the smallest census unit in Malawi) from the International Food Policy Research Institute (IFPRI). This dataset originated from the 1998 Malawi Housing and Population Census conducted by the National Statistical Office in Malawi. To ascertain the accuracy of the data, these files were compared to a set of GIS files donated to us by University of North Carolina at Chapel Hill based on the 2008 Malawi Housing and Population Census [1]. Though the second set of files did not contain EA-level data, which excluded it from our main analysis, the two sets of files were extremely similar at the level of Traditional Authorities (or TAs, an administrative level above EAs). Both contained 367 TAs, and the centroids of the same TA taken from the two datasets were an average of less than 10^{-4} meters of one another. While this result is somewhat reassuring, it could also indicate a common source for the two data sets.

IFPRI also provided us with population data at the EA level. To compensate for the age of the data, the population of each EA was adjusted to approximate 2008 levels according to the “Annual Population Inter-censal Growth Rates and Increase 1998-2008” [1] from the 2008 Malawi Population and Housing Census for the district containing the EA. Additionally, we estimated the number of children under five years of age in each EA using the percentage of the population comprised of under-5 children for the TA containing that EA in the 2008 Malawi Census [1].

The locations of hospitals and health centers were also obtained from IFPRI and were compared with data from the Health Information Systems Programme for accuracy. The IFPRI data came from approximately 1998, based on our correspondence. There was 100%

agreement in the number of central and district hospitals and 92% agreement in the number of total hospitals between the data from IFPRI and that from the Health Information Systems Programme. This discrepancy, while small, reflects the fluidity of the health care system in Malawi and the necessity that BTB follow up our recommendations with their own research on the ground.

With the location of EAs and hospitals, we were then able to calculate the distances between them and use integer programming to optimize the locations of HSAs, backpacks, and resupply centers.

3.3 Implementation

We used ArcGIS, a commercially available GIS data processing program, to process all the geographic and population data used in this project. In particular, ArcGIS was used to find the centroid of each EA; calculate the distance between the centroids of each EA; and calculate the distance from each EA centroid to each of the hospitals and the health centers.

Three C++ routines were then written to interface with Gurobi 4.0.1, a commercially available, large-scale integer program solver, to solve each of the distribution problems [4]. All programs were run on a machine with a memory of 16 GB and a processing speed of 2.83 GHz. After our initial run attempts exceeded the available memory, several strategies were used to limit memory usage. First, we solved the models sequentially rather than simultaneously, the consequences of which are discussed in Section 5. We also set a hard distance threshold of 100 km for each model, allowing us to generate only variables corresponding to connections of less than 100 km. Finally, we set the Gurobi Threads parameter to use only one core, modified the Gurobi NodefileStart parameter to write nodes to disk, and increased the Gurobi MIPGap parameter for our largest models. Gurobi uses a Branch-and-Cut algorithm, whereby it tries to drive down the gap between the best known objective value for a feasible solution and the best bound known. As the optimal solution always lies between these two values, the smaller the gap, the higher confidence we have in the optimality of our results. The gap percentages are listed with the results of each model, and range from no gap to 2%. While heuristic methods exist for both the p -median problem and CFLP [20][10][14], these memory-reducing strategies used allowed us to reach near-optimal solutions to each problem without heuristic methods.

Model outputs were analyzed and visualized using a combination of ArcGIS, Microsoft Excel, and MATLAB.

4 Project Assessment

In this section, we describe our initial assessment of the potential scope and impact of the proposed backpack scale-up based on the p -median problem and CFLP. We provide the results for both the entire country and specifically for Lilongwe District, the location of the pilot backpacks and likely site of the initial backpack scale-up. Note that the distances described throughout represent straight-line distances, as we were unable to obtain sufficiently detailed data about road locations and conditions to include these details within our models.

4.1 HSA Assignments

We used the p -median problem to assign HSA pairs to EAs. HSAs in Malawi are generally assigned to work in pairs to improve their safety and accountability. We represented EA locations by their centroids and used these same spots as possible HSA locations. Our results for Malawi resulted in an average distance of 0.39 km and maximum of 12.7 km between each EA and its assigned HSA. For comparison, a naive strategy that placed HSAs randomly across Malawi and assigned each EA to the nearest HSA resulted in an average distance of 0.62 km and maximum of 20.3 km between EAs and HSAs.

4.1.1 p -Median Problem

In the p -median problem, the goal is to find the p locations that minimize the average distance in a network of n vertices [9]. A p -median problem can be formulated and solved as a binary integer program [18]. Let $W_{ij} = w_i d_{ij}$ be the weighted distance matrix, where w_i represents the weight assigned to each vertex i and d_{ij} is the shortest distance between vertices. Let x_{ij} be an allocation variable that describes whether vertex i is assigned to vertex j , such that $x_{ii} = 1$ if a facility is located at vertex i . We then seek

$$\min \sum_{i=1}^n \sum_{j=1}^n W_{ij} x_{ij}$$

such that

$$\sum_{j=1}^n x_{ij} = 1, \quad \text{for all } i = 1, \dots, n$$

$$\sum_{i=1}^n x_{ii} = p$$

$$x_{ij} \leq x_{jj} \quad \text{for all } i, j = 1, \dots, n$$

$$x_{ij} \in \{0, 1\} \quad \text{for all } i, j = 1, \dots, n$$

For our problem, the nodes represent EA centroids and the associated weights w_i represent the demand assigned to each node. Specifically, the demand at each EA reflects:

– *General and Under-Five Population*

Children under age five were given a weight of 0.55, whereas the rest of the population was given a weight of 0.45. This weighting reflects the design of the HSA backpack, which assumed HSAs would spend 55% of their time on under-five children.

– *Rural or Urban Setting*

We used population density to arrive at a proxy for rural versus urban setting. Rural populations, defined as EAs with population densities below a natural cutoff of 0.0035 people per square meter in the data, were weighted upward by 1.5 to reflect the lack of adequate infrastructure in these areas.

– *Proximity to Health Center*

If an EA was within 1 km of a health center, its demand was reduced to 10% of its original level. That is, we assumed the health centers could provide adequate care to 90% of people within these EAs.

The d_{ij} , distance of shortest path between the nodes, was represented by the shortest distance between centroids of the EAs.

4.1.2 HSA Assignment Results

We first used the p -median problem to assign 615 HSA pairs, based on a 1:1000 ratio of HSAs to population, to the 861 EAs with nonzero population in Lilongwe District. We then calculated the average distance between each EA and the location of the HSA pair serving it. All HSAs were considered to be 0 km from their base EA, and all results represent distance to the centroid of the population areas served rather than actual patient locations. This distinction was necessary given the unknown population distribution within each EA, but as most EAs are only a few square kilometers in area, it does not lessen the value of our results.

Our model resulted in an average distance from EA to HSA pair of 0.46 km and a maximum of 4.4 km, reasonable even for HSAs traveling on foot. A maximum of five EAs were assigned to one HSA pair. These HSA assignments were calculated with no gap in Gurobi and in a runtime of 81 seconds. For comparison, a naive strategy that placed HSAs randomly across Lilongwe District and assigned each EA to the nearest HSA resulted in an average distance of 0.59 km and a maximum distance of 3.9 km between EAs and HSAs. While these

unweighted distances are similar, the average and maximum weighted distances produced by the naive strategy were 582 and 7540, respectively, compared to average and maximum weighted distance of 304 and 1580 from our model.

In the nationwide case, we assigned 6500 HSA pairs to the 9147 EAs in Malawi with nonzero population. Again, the number of HSAs assigned was based on an assumed 1:1000 HSA to population ratio. The distance results for the entire country are similar to those for Lilongwe District and again represent reasonable walking distances. For Malawi, our model resulted in an average distance from each EA centroid to its assigned HSA of 0.39 km and a maximum distance of 12.7 km. For comparison, a naive strategy that placed HSAs randomly across Malawi and assigned each EA to the nearest HSA resulted in an average distance of 0.62 km and maximum of 20.3 km.

The number of EAs served by each HSA pair in Malawi is shown in Table 1. As in Lilongwe District, most HSA pairs are assigned to only one or two EAs. However, our Malawi model assigned some HSAs to serve 10 or more EAs, with the maximum number being 35. We found that the HSAs assigned to serve many EAs were those centered in urban areas and located near health centers. This discrepancy can be explained by the smaller size of EAs located in urban areas along with our assumption that EAs located near a health center would have the majority of their needs met by that health center. These results were calculated with a Gurobi gap percentage of 0.01% in 11.5 hr.

Table 1 Number of Malawi HSAs by the number of EAs they serve.

Number of EAs per HSA	Number of HSA pairs
1	4909
2	1073
3	335
4	95
5–9	71
10–19	14
20+	3
Total	6500

4.2 Backpack and Resupply Center Allocation

The single tier CFLP was used to allocate HSAs to backpacks and select resupply centers from the existing network of health centers and hospitals. For the resupply center problem, we considered two versions of the model: one that allowed both hospitals and health centers to act as resupply centers, and one that allowed only hospital resupply centers. Our results suggest that 2188 backpacks would be necessary to adequately cover the entire country, suggesting that a country-wide scale-up may be cost prohibitive. However, deploying the backpacks thusly could have a substantial impact, reducing the average distance between each HSA and the nearest supply source (healthcare facility or backpack) from 5.2 km to 1.8 km. Finally, our results demonstrate the importance of smaller health centers to the Malawian healthcare network and the need to include them as potential backpack resupply sites.

4.2.1 The Capacitated Facility Location Problem

In the general CFLP, a set of potential facility locations $i = 1, \dots, m$ and a set of customers $j = 1, \dots, n$ are given. The problem is to minimize the total cost of locating facilities and assigning them to customers. The total cost includes the variable travel cost between customers and facilities, c_{ij} , and the fixed cost of opening new facilities, f_i . This problem is “capacitated” because each possible location has an upper limit on its supplying capacity b_i . A CFLP with each customer assigned to only one facility can be formulated as a binary integer program with two decision variables: x_{ij} , whether customer j is served by facility i , and y_i , whether facility i is picked. We seek:

$$\min \sum_{i=1}^m \sum_{j=1}^n c_{ij} x_{ij} + \sum_{i=1}^m f_i y_i$$

such that

$$\sum_{j=1}^n x_{ij} \leq b_i y_i \quad \text{for all } i = 1, \dots, m$$

$$\sum_{i=1}^m x_{ij} = 1 \quad \text{for all } j = 1, \dots, n$$

$$x_{ij} \leq y_i \quad \text{for all } i = 1, \dots, m, \text{ and } j = 1, \dots, n$$

$$x_{ij}, y_i \in \{0, 1\} \quad \text{for all } i = 1, \dots, m, \text{ and } j = 1, \dots, n$$

When assigning backpacks to HSAs, we considered backpacks as our facilities and HSAs as customers. For the backpack to resupply center problems, we considered hospitals and health centers as facilities and backpacks as customers. The parameter values for each of

these models are shown in Table 2. Note that the fixed and variable costs for each backpack, f_i and c_{ij} , are chosen such that the fixed cost of adding a new backpack is equivalent to a 30 km HSA-backpack distance. This measure allowed us to directly compare the financial cost of introducing a new backpack to the accessibility costs of large distances between backpacks and HSAs, and was chosen to reflect the maximum distance between HSAs and health centers in a 2001 survey [12]. The effects of this parameter are explored in more detail in Section 5.2.1.

4.2.2 Optimal Backpack Assignments

We used the capacitated facility location problem to allocate backpacks to HSAs and, in the process, determine the number of packs needed to cover the region. In Lilongwe District, our model assigned 205 backpacks across the 615 HSAs allocated to Lilongwe. These values indicate that 100% of the backpacks were assigned to serve a full capacity of 3 HSA pairs. We explore the importance of this constraint, which BTB felt accurately represented the backpacks’ true capacities, in the Model Analysis section. At an initial cost of \$350 per backpack, 205 backpacks corresponds to a total start-up cost of \$71,750 to cover Lilongwe District.

Under these backpack assignments, the average distance between each HSA pair and the backpack serving it was 1.6 km, with a maximum of 9.1 km. As with the HSA assignments, these distances appear reasonable even for HSAs with no transportation options other than walking. The assignments are displayed visually in Figure 1, which highlights the short average travel distances. The backpack assignments were calculated with a Gurobi gap percentage of 0.01% and a runtime of approximately 2 hours.

The second tier of our model assigned 2188 backpacks to the 6500 HSA pairs across Malawi. This corresponds to a total start-up cost of \$765,800. While this cost is not unreasonable considering the potential of the HSA backpacks to improve health care, it was high enough to encourage BTB to focus on successfully scaling up in Lilongwe District first. Of these 2188 backpacks, 2129 (97.3%) were assigned to serve at the full capacity of 3 HSA pairs, again indicating the importance of this constraint. Distances between backpacks and HSAs were similar to those in the Lilongwe case, with an average distance between each HSA pair and its assigned backpack of 1.8 km and maximum of 21.7 km. These results were calculated with a Gurobi gap of 0.80% in a runtime of 12.5 hr.

Table 2 Parameters for the CFLPs assigning backpack and resupply center locations. These parameters were determined through consultation with BTB, based on their experiences with the pilot backpacks at St. Gabriel’s Hospital in Namitete, Malawi. The number of customers for the resupply center problems were obtained from the output of the backpack problems.

Parameter	Description	Backpacks to HSAs	Resupply Centers to Backpacks
n	Number of Customers (Lilongwe)	615 HSA pairs	205 backpacks
m	Potential Facilities (Lilongwe)	615 HSA pairs	2 hospitals, 37 health centers
n	Number of Customers (Malawi)	6500 HSA pairs	2188 backpacks
m	Potential Facilities (Malawi)	6500 HSA pairs	49 hospitals, 661 health centers
b_i	Facility Capacity	3 HSA pairs	80 (hospitals) or 10 backpacks
D_{ij}	Customer-Facility Distance (km)	-	-
c_{ij}	Variable Cost	D_{ij} km	D_{ij} km
f_i	Fixed Cost	30 km	0 km

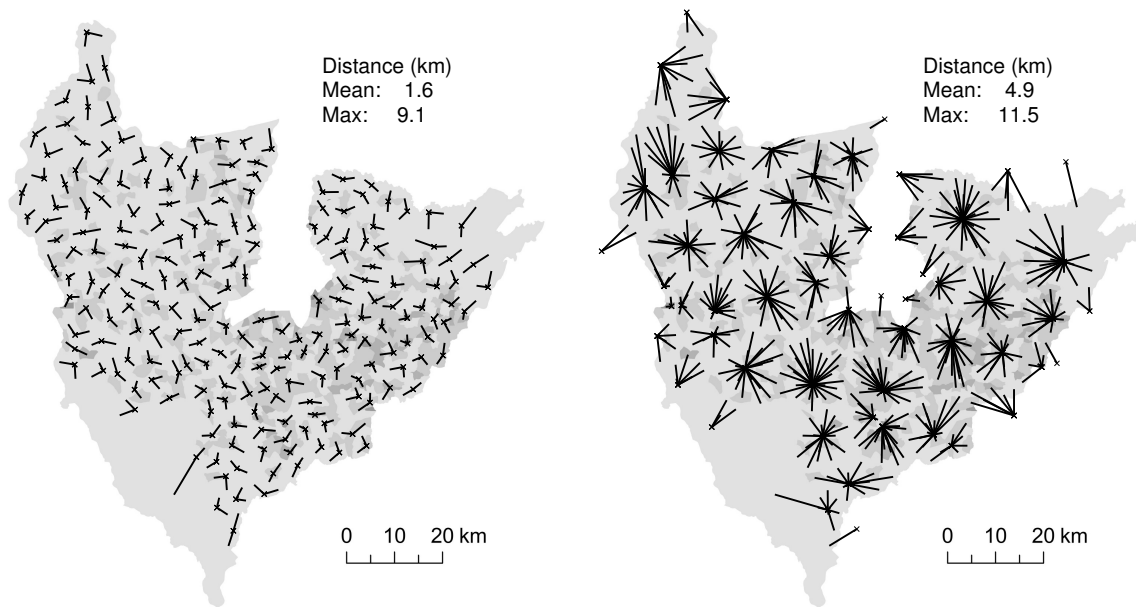


Fig. 1 (Left) Map showing backpack-to-HSA assignments in Lilongwe District. X’s represent the location of each backpack. Lines indicate connections between backpacks and additional HSAs served. (Right) Map showing each HSA in Lilongwe District assigned to the nearest health center or hospital. X’s represent the location of each healthcare facility. Lines indicate connections between healthcare facilities and HSAs. Both plots are overlaid on a map of Lilongwe District, with darker areas representing EAs with higher population density.

4.2.3 Optimal Resupply Center Assignments

Next we assigned the backpacks allocated in the second tier of our model to existing hospitals or health centers which could act as resupply centers for disposables within the packs. We calculated these assignments for two cases: one using hospitals only as resupply centers, and one in which both hospitals and health centers were considered. The advantage of using hospitals only is that they are generally more reliable than smaller health centers, which are prone to frequent stock-outs;

however, our results indicate that this increase in reliability would come at a cost. Only two hospitals exist in rural Lilongwe District, which required us to set the capacity of each above the original 80 backpacks to 150 backpacks. With this arbitrary threshold, our model resulted in one hospital serving 89 backpacks and the other serving 116. While BTB believed these numbers to be feasible, the resulting travel distances between these backpack and hospital pairings are prohibitively high. The hospital-only approach resulted in an average distance between backpacks and resupply centers

of 28 km with a maximum of 62 km. As walking is the only travel option for most HSAs, a one-way distance of 28 km could translate to unreasonably high travel times for even a single backpack refill trip.

When health centers were included in the model, these distances between backpacks and resupply centers fell sharply to an average distance of 5.4 km and a maximum of 25.7 km. These represent much more reasonable travel times for HSAs, who would need to take time away from patients to refill the packs. In the hospital and health center case, 37 of 37 health centers and both hospitals were chosen. The hospitals served 6 and 2 backpacks respectively, while the health centers served a mean of 5.3 backpacks and a maximum of 10.

Results were similar for the nationwide version of our model, shown in Figure 2. We again compared a hospital-only resupply case with a hospital and health center resupply case. In the hospital and health center case, 46 of 49 hospitals and 574 of 661 health centers were chosen as resupply centers. The mean number of backpacks served by each hospital was 4.09, with a maximum of 12; for health centers, the mean was 3.48 backpacks, with a maximum of 10. The large number of facilities chosen led to short distances between each backpack and its resupply site, with a mean of 5.1 km and a maximum of 28.4 km.

In the hospital-only case, all 49 hospitals were chosen as resupply centers. This led to a mean number of backpacks served by each hospital of 40, with a maximum of 80. As Figure 2 shows, the distances HSAs would need to travel in this case were again prohibitively higher than in the health center and hospital case, with an average distance of 19.8 km and a maximum of 70.5 km. The hospital-only case also exposed a few other complications, including the lack of hospitals available to serve HSAs based in the islands of Lake Malawi. Given the poor infrastructure and high fuel cost in Malawi, these HSAs might not be able to refill their backpacks at all if assigned to a hospital resupply center.

Based on these results, BTB determined that they could not plan to use only hospitals as resupply centers when deploying the HSA backpacks. The small number of hospitals and long distances between hospitals and HSAs demonstrate the importance of health centers within the Malawian health care system. However, the potentially irregular supply of health centers remains a concern. Our results suggest that the ideal solution would include hospitals and the most reliable health centers as resupply centers, striking a balance between concerns about distance and supply chain vulnerability. More detailed data about which health centers are most reliable and how often stock-outs occur is needed

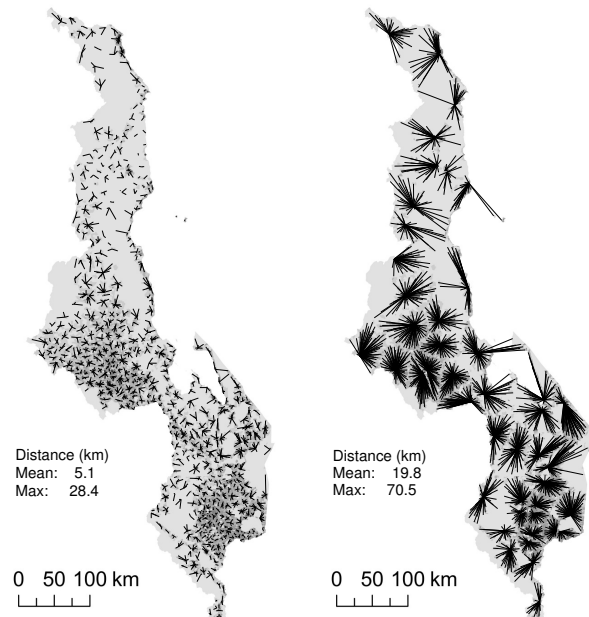


Fig. 2 Resupply Center-to-Backpack assignments across Malawi, when either (left) both hospitals and health centers or (right) hospitals only are allowed as potential resupply sites. These assignments are overlaid on a map of Malawi, with darker areas representing EAs with greater population density.

to determine the combination of resupply centers that best meets these concerns.

4.2.4 Potential Impact of Backpacks

The models above allowed us to determine not only the optimal number and locations of the backpacks, but also their potential impact. Inadequate access to medical supplies is one of the primary factors limiting the current success of HSAs in Malawi. Optimizing the deployment of HSAs using the first tier of our model would not be enough to solve the problem. HSAs deployed in this manner in Lilongwe District would be at an average distance of 4.9 km from the nearest health center or hospital, with a maximum of 11.5 km, as shown in Figure 1. Assuming the HSAs walk at a pace of 3 km/hour [12], this distance corresponds to an average round trip of at least 3 hours to the nearest health center; note that this estimate assumes that the nearest health center is fully stocked, which in many cases may be unrealistic. If we assume that HSAs would travel to the health centers with the same frequency that we assumed they would travel to the backpack locations (twice per week), HSAs would spend at least 6 hours per week retrieving supplies rather than treating patients.

As described earlier, our proposed deployment of the HSA backpacks in Lilongwe would place backpacks at an average distance of 1.6 km from the HSAs they serve, with a maximum of 9.1 km (Figure Figure 1). This would lower the distance that the average HSA needed to travel for supplies by more than half and, under the assumptions described above, free up approximately three hours per week to patient care. While the backpacks would still need to be refilled periodically, these trips could be made once a month, allowing HSAs to spend the majority of their time helping patients rather than searching for supplies. HSAs would see similar gains if the backpacks were extended to cover all of Malawi; instead of walking an average distance of 5.2 km and maximum of 28.4 km to a hospital or health center for supplies, HSAs would need to travel only an average distance of 1.8 km and maximum of 21.7 km to reach their assigned backpacks.

Improved access to supplies would allow HSAs to provide medical care to patients who otherwise may be unable to receive medical treatment. The HSA backpacks contain supplies used to treat acute childhood illnesses such as diarrhea and malaria. For children with these conditions, receiving medications a few hours earlier could mean the difference between life and death. This improved access to lifesaving medical care is especially important in Malawi, which in 2006 had an under-5 mortality rate of 134 deaths per 1,000 live births [13]. Our models therefore provide not only an optimal strategy for deploying the HSA backpacks in Malawi, but also increased evidence of the impact they could have there.

5 Model Analysis

We initially used the p -median problem to determine optimal HSA assignments and the CFLP to assign backpacks to HSAs and resupply centers to backpacks. The results of these analyses provided several insights into the scope and potential impact of the proposed backpack scale-up. However, these analyses relied on several strong modeling assumptions and data sources with potential inaccuracies. In this section, we describe additional analyses used to assess and refine our primary models. Because our initial analyses suggested that the initial backpack scale-up should center on Lilongwe District, we focus our analysis on that region.

5.1 HSA Assignments

In our analysis of the HSA assignment model, we examined two potential concerns: first, that our model

could be sensitive to uncertainties in the EA weights, and second that our model might not fairly distribute HSA workloads. Each of these concerns are addressed below.

5.1.1 Sensitivity Analysis

Within the HSA assignment problem, demand weights were assigned to each EA based on its population, under-five population, urban/rural setting, and proximity to a health center. These variables were chosen in consultation with BTB, but their respective weights were somewhat arbitrary. Furthermore, our estimates of the population and under-five population of each EA were subject to considerable uncertainty. We therefore wished to quantify how much our results depended on each of these factors.

Because our EA-level population estimates were based on the projection of 1998 census data to 2008 levels, the accuracy of our EA-level population estimates was unclear. To determine the potential impact of this uncertainty, we ran a sensitivity analysis on the population distribution within Lilongwe District. The total population of Lilongwe, which was obtained from the 2008 Malawi Census, was assumed fixed. We then varied the population of each EA according to a uniformly distributed random variable with endpoints ranging from $\pm 10\%$ to $\pm 50\%$. The results are shown in Table 3. The columns labeled “Original Objective Value” and “Original Max Weighted Distance” were calculated using the perturbed weights and original HSA assignments; the remaining columns compare the optimal solution given the perturbed weights to these original assignments. The results suggest that the HSA locations are fairly robust to changes in the population distribution, and that our original HSA assignments remain near optimal for moderate changes in the population distribution. The number of backpacks is very robust to these changes, but the optimal backpack locations change dramatically for even small changes in HSA locations. Both of these characteristics can be explained by the importance of the capacity constraint in determining backpack locations. The large changes in backpack locations, however, do not affect the number or location of resupply centers; because there is no fixed cost of opening a new resupply center, the optimal solution consistently uses every potential site to resupply the backpacks.

Next, we analyzed the sensitivity of our results to the estimated under-five population of each EA. Our estimates of the number of under-five children in each EA were based on the percentage of under-five children for the Traditional Authority containing that EA in the 2008 Malawi Census. We assumed the total number of

under-five children in Lilongwe District was fixed and then varied the under-five population of each EA according to a uniformly distributed random variable with endpoints ranging from $\pm 10\%$ to $\pm 50\%$. The results are shown in Table 3. These results suggest that the HSA locations are very robust to changes in the distribution of the under-five population, and that our original HSA assignments remain very near optimal for even large changes in the under-five population distribution. As in the sensitivity analysis for the overall population, the optimal backpack locations are highly sensitive to HSA positions, but the optimal number of backpacks and optimal resupply sites are not.

Our relative weightings of the under-five and over-five population reflects the design of the HSA backpack, which assumed HSAs would spend 55% of their time on under-five children. However, our weights for EAs located near health centers or with low population density were chosen more arbitrarily, in consultation with BTB. We therefore wished to determine the sensitivity of our results to these weights. Changing the weight for rural EAs from 1.5 to 1 or 2 resulted in no change in any aspects of the optimal solution for Lilongwe District; this insensitivity reflected the homogeneity of Lilongwe District, with only 3 urban EAs. The results of changing the weight for EAs near health centers are shown in Table 3. The optimal solution did not change when the health center weight was changed from 1/10 to 1/5 or 1/15. Even when the weight was set to 1, so that EAs near health centers were weighted the same as those that were not, only 2.9% of HSA locations differed between the original and optimal solutions, and the original solution remained within 4.8% of optimal. As in the other sensitivity analyses, the optimal backpack locations are highly sensitive to HSA positions, but the optimal number of backpacks and optimal resupply sites are not.

Overall, our results appear fairly robust to uncertainty in the weights assigned to each EA. Unlike the optimal number of backpacks, however, the optimal backpack locations are highly sensitive to changes in HSA locations. This likely reflects the importance of the capacity constraint to the backpack problem, as is discussed in more detail later. Fortunately, the recommended HSA locations remain near optimal for moderate to large changes in EA demand weights, suggesting that changing HSA (and consequently backpack) locations may be unnecessary.

5.1.2 Fairness

The p -median model, which minimizes the average weighted distance between servers and customers, is among the

models most commonly used for locating healthcare facilities in the developing world [17]. However, the p -median model does not explicitly address the fairness of these locations. We therefore sought to compare our results with those obtained from a model with equity considerations.

One possible approach to improving model fairness is to introduce explicit distance or capacity constraints [17]. The role of distance and capacity thresholds in our backpack assignment model are addressed in the next section. However, we chose not to implement similar thresholds within our HSA assignment model. HSA workload depends on a combination of many factors, including travel distances, the size of the population, and other characteristics of the area served. Therefore, thresholds based on only some of these factors could improve the workload for some HSAs while worsening it for others. Additionally, the results of our original model did not raise any obvious concerns that would need to be addressed by threshold constraints. The maximum distance between HSAs and EAs in our nationwide model was less than the maximum distance between HSAs and their farthest village served in a 2001 survey (17 km), and our maximum distances for Lilongwe District were less than the average distance between HSAs and their farthest village served (6 km) in the same survey [12]. While our model assigned some HSAs to serve many EAs, this aspect of our results reflects the smaller size and increased access to healthcare of many urban EAs and is not necessarily a sign of inequity in overall HSA workload. We therefore doubted whether introducing threshold constraints into our weighted p -median model would improve the equity of our results, and instead chose to reanalyze our data using the weighted p -center problem.

The p -center problem minimizes the maximum weighted distance between servers and customers, as opposed to the average weighted distance. The p -center problem can be solved as a binary integer program, with the traditional formulation highly comparable to that of the p -median problem [8]. As for the p -median problem, let $W_{ij} = w_i d_{ij}$ be the weighted distance matrix between EAs and potential HSA sites. We assume that each of the n EA centroids is also a potential HSA location. Let x_{ij} be an allocation variable indicating whether EA i is allocated to HSA j . Finally, let z represent the maximum weighted distance between EAs and HSAs for a particular feasible solution. We seek

$$\min z$$

such that

$$\sum_{j=1}^n x_{ij} = 1, \quad \text{for all } i = 1, \dots, n$$

Table 3 Robustness of EA-to-HSA assignments. The total and under-5 population of each EA was varied according to a uniformly distributed random variable with endpoints ranging from 10%-50%. Mean values are presented for 10 runs at each level. The weight given to EAs near health centers was varied from the baseline weight of 1/10. All models were solved in Gurobi with a maximum gap of 1%

	Change in Input Values	Original Max Weighted Distance ($*10^3$)	Original Objective Value ($*10^5$)	Optimal Objective Value ($*10^5$)	Change in HSA Locations	Change in Pack Numbers	Change in Pack Locations
Population	0%	1.58	2.62	2.62	-	-	-
	10%	1.69	2.62	2.60	3.7%	0.3	38.8%
	20%	1.80	2.62	2.54	7.7%	0.3	48.5%
	30%	1.98	2.62	2.46	10.0%	0.4	51.5%
	40%	2.10	2.63	2.33	13.0%	0.3	55.9%
	50%	2.27	2.65	2.19	15.4%	0.4	58.9%
Under-5 Population	0%	1.58	2.62	2.62	-	-	-
	10%	1.58	2.62	2.62	0.2%	0.1	16.8%
	20%	1.59	2.62	2.62	0.3%	0	18.8%
	30%	1.58	2.62	2.62	0.5%	0.3	21.8%
	40%	1.60	2.62	2.62	0.7%	0.2	25.3%
	50%	1.58	2.62	2.62	0.9%	0.3	27.1%
Health Center Weight	1/10	1.58	2.62	2.62	-	-	-
	1	3.29	3.02	2.88	2.9%	0	33.8%
	1/5	1.58	2.66	2.66	0%	0	0%
	1/15	1.58	2.61	2.61	0%	0	0%

$$\sum_i^n x_{ii} = p$$

$$x_{ij} \leq x_{jj} \quad \text{for all } i, j = 1, \dots, n$$

$$\sum_{j=1}^n W_{ij} x_{ij} \leq z, \quad \text{for all } i = 1, \dots, n$$

$$x_{ij} \in \{0, 1\} \quad \text{for all } i, j = 1, \dots, n$$

The results of the p -center model were calculated with no Gurobi gap and are shown in Table 4. The solution to the p -center problem differs somewhat from that of the p -median problem, with 10.2% of HSAs located in different EAs, but the outcomes are virtually identical. Despite their differing objective functions, both the p -median and p -center problems result in a maximum of five EAs per HSA, a half-kilometer mean distance between EAs and HSAs, and a maximum EA to HSA distance of 4.4 km. The maximum weighted distance produced by the p -center model was 1560, only slightly smaller than 1580 from the p -median model (compared to the naive value of 7540). The results suggest that our original HSA assignment model was reasonably equitable for Lilongwe District, though equity was not explicitly addressed.

Table 4 Comparison of using the p -center vs p -median problem to assign HSAs to EAs in Lilongwe District

	p -Median	p -Center
Number of EAs (n)	861	861
Number of HSA Pairs (p)	615	615
HSAs Moved	-	10.2%
Max EAs per HSA	5	5
Mean EA-HSA Distance (km)	0.46	0.48
Max EA-HSA Distance (km)	4.4	4.4
Mean Weighted Distance	304	351
Max Weighted Distance	1580	1560

5.2 Backpack and Resupply Center Assignments

We had two primary concerns about the models we had used to allocate backpacks and resupply centers in Lilongwe. First, because all of the backpacks in Lilongwe were assigned to full capacity, we wanted to explore this constraint in more detail. Second, we wanted to see how our recommendations would change if we analyzed the backpack and resupply tiers of our model simultaneously, rather than sequentially. Both of these analyses are described below.

5.2.1 Capacity Constraints

In our original model, 205/205 backpacks assigned to Lilongwe District were at full capacity. We therefore wished to assess the model’s sensitivity to this constraint and how it would change based on the fixed cost of a backpack. First, we compared our original model with a capacity of 3 HSA pairs per backpack to one that allowed 4 HSA pairs per backpack. These results, along with the others described in this section, were calculated in Gurobi with a 0.2% gap stopping criterion and are shown in Table 5. The optimal number of backpacks changes from 205 to 154 when the capacity of each backpack is increased from 3 to 4 pairs of HSAs. This difference shows that the results are highly sensitive to the backpack capacity constraint and, notably, that almost all backpacks remain at full capacity even when the capacity is inflated to 4 HSA pairs.

We hypothesized that these results could be explained by our choice of the fixed cost of a backpack. We initially set the fixed cost of a backpack to be equivalent to 30 km in distance, in essence setting a threshold of 30 km between HSAs and backpacks. Using distance as a cost measure gave us a simple framework for comparing the financial costs of adding new backpacks with the accessibility costs of limiting pack numbers. However, 30 km was a conservative estimate, based on the maximum distance between HSAs and health centers in a 2001 survey [12]. In fact, the distance between HSAs and health centers should be significantly lower than this, as enforced by the capacity constraint in our original model. To reflect this, we lowered the fixed cost of a backpack to 5 km, the average distance between HSAs and health centers in the same survey [12]. The results of the model under this new fixed cost estimate are given in Table 5 for both 3 and 4 pairs of HSAs. Under this new fixed cost and a capacity of 3 HSA pairs per backpack, the number of backpacks needed to serve Lilongwe District adequately increased from 205 to 212 and the maximum HSA to backpack distance decreased from 9.1 km to 4.4 km. These results show that a small increase in the number of backpacks could greatly shorten the distances between HSAs and backpacks. As expected, the new fixed cost also decreases the model’s sensitivity to the capacity constraint, if only moderately. Given the fixed cost of 5 km and capacity of 4 HSA pairs, the model recommended using 176 backpacks to cover Lilongwe District, compared to 154 backpacks under a fixed cost of 30 km.

It is important to note that while the backpack model is quite sensitive to changes in the capacity assigned to each backpack, this parameter is also among the least uncertain ones in our model. The maximum

capacity of three HSA pairs per backpack was determined through interviews with HSAs using the pilot backpacks out of St. Gabriel’s Hospital. Based on this analysis, we believe that the model assuming each backpack has a fixed cost of 5 km and capacity of 3 HSA pairs best captures the tradeoffs between the number of backpacks, travel distances, and model sensitivity.

Table 5 Comparisons of backpack assignments under different capacity constraints and fixed costs

Capacity (HSA pairs)	3	4	3	4
Fixed Cost (km)	30	30	5	5
Number of Backpacks	205	154	212	176
Mean HSA-Pack Distance (km)	1.6	2.0	1.5	1.7
Max HSA-Pack Distance (km)	9.1	9.5	4.4	4.2

5.2.2 Two Echelon CFLP

In our initial model, we solved the backpack and resupply assignment problems separately, using the backpack location outputs from the first CFLP as inputs for the resupply center assignment problem. This approach was necessary to limit the memory usage of our models, and would also make it easier to refine the set of potential resupply centers after backpack locations are fixed. However, solving the two problems independently could result in sub-optimal results, with overall travel times greater than those that would be found by solving them simultaneously. To assess the resulting loss in efficiency, we built a model for backpack and resupply center assignment in Lilongwe District based on the two-echelon, single-source, capacitated facility location problem. This problem describes a system in which second-echelon facilities (such as our backpacks) are used to deliver goods from first echelon facilities (resupply centers) to customers (HSAs) [21].

Let m denote the number of HSAs and potential backpack sites, and n denote the number of potential resupply centers. Let b_i be the capacity of a backpack and a_k be the capacity of a resupply center. Associate a fixed cost g_k with opening each new resupply center; a variable cost f_{ik} with using resupply center k to serve backpack i ; and a variable cost c_{ijk} with using backpack i served by resupply center k to serve HSA j . This problem can be formulated as a binary integer program, with decision variables x_{ijk} indicating

whether backpack i served by resupply center k serves HSA j , y_{ik} indicating whether backpack i exists and is served by resupply center k , and z_k indicating whether resupply center k is in use. The problem is then [21]

$$\min \sum_{i=1}^m \sum_{k=1}^n \sum_{j=1}^m c_{ijk} x_{ijk} + \sum_{i=1}^m \sum_{k=1}^n f_{ik} y_{ik} + \sum_{k=1}^n g_k z_k$$

such that

$$\sum_{j=1}^m x_{ijk} \leq b_i \quad \text{for all } i = 1, \dots, m, k = 1, \dots, n$$

$$\sum_{i=1}^m y_{ik} \leq a_k \quad \text{for all } k = 1, \dots, n$$

$$\sum_{i=1}^m \sum_{k=1}^n x_{ijk} = 1 \quad \text{for all } j = 1, \dots, m$$

$$\sum_{k=1}^n y_{ik} \leq 1 \quad \text{for all } i = 1, \dots, m$$

$$x_{ijk} \leq y_{ik} \quad \text{for all } i, j = 1, \dots, m, k = 1, \dots, n$$

$$y_{ik} \leq z_k \quad \text{for all } i = 1, \dots, m, k = 1, \dots, n$$

$$x_{ijk}, y_{ik}, z_k \in \{0, 1\} \quad \text{for all } i, j = 1, \dots, m, k = 1, \dots, n$$

The parameters used in this problem are described in Table 6. In general, these parameters are similar to those from the one echelon CFLPs. However, we assumed a backpack fixed cost of 5 km, as described under the Capacity Constraints section. We weighted the distances between backpacks and health centers by 1/8, as trips between HSAs and backpacks would likely occur about eight times as often as trips between backpacks and resupply centers (assuming each HSA retrieves and then returns the backpack once per week, and the backpack is resupplied once per month). To limit memory usage, we required that each HSA be within of 10 km of its assigned backpack and each backpack be within 30 km of its assigned resupply center. These thresholds greatly reduced the number of variables in our model (from 1.48×10^7 to 8.06×10^5 , measured after presolve), and as neither threshold was exceeded for the single echelon problems, we did not expect them to change our results. We stopped the simultaneous model after reaching a gap of below 2%, after approximately 21 hours.

As Table 7 shows, the results of the two-echelon models are very similar to those obtained by the bottom-up approach. This may be attributed to the lower weight assigned to distances between health centers and backpacks than to distances between HSAs and backpacks. When the results of the sequential model are compared to the objective function of the simultaneous model, the

difference is <1% for both the health center and hospital scenarios. The two echelon approach only slightly changed the distances between backpacks and resupply centers and the recommended number of backpacks.

Based on these results, we believe our original bottom-up approach is an acceptable alternative to solving the backpack and resupply problems simultaneously. The bottom-up approach is far more computationally tractable than the simultaneous approach, with shorter runtimes and smaller memory demands, and the objective function values are nearly identical for Lilongwe. Using the bottom-up approach allowed us to estimate optimal backpack and resupply center locations for the entire country of Malawi, and will also make it easier to adjust resupply center assignments later in the process of the backpack scale-up. Though we did not have the computational resources to solve the two-echelon model to full optimality, the results of the sequential and simultaneous models are so similar that decreasing the gap tolerance below 2% would be unlikely to change these conclusions.

6 Conclusion

This paper uses facility location models to describe an optimal scale-up plan for Malawi's HSA program and BTB's HSA backpack program. In particular, we used the p -median problem to determine the optimal allocation of HSAs across Malawi and the capacitated facility location problem to assign HSAs to backpacks and backpacks to resupply centers. Using these models, we provided BTB with a backpack scale-up analysis that provides pack numbers, estimates monetary costs, and highlights potential challenges.

Because of the financial and logistical difficulties associated with nationwide scale-up of the HSA backpack program, the backpack deployment will likely take a stepwise form, beginning with Lilongwe District. Under our plan, the HSAs in Lilongwe District would be at an average distance of only 1.6 km from their assigned backpacks, as opposed to 4.9 km if relying directly on a hospital or health center for supplies. This improvement is especially important since most HSAs travel by foot over roads with varying conditions. Giving HSAs easier access to supplies will enable them to deliver care more efficiently and effectively.

Our assessment also allowed BTB to consider the cost of supply chain disruptions in Malawi by comparing a case in which only hospitals were used as backpack resupply centers to one in which both hospitals and health centers were used. The inclusion of health centers led to sharp decreases in the distance between

Table 6 Parameters for two-echelon CFLP for backpacks and resupply centers.

Parameter	Description	Value
m	Number of HSAs/ potential backpack sites	615
n	Number of potential resupply centers	2 (hospitals), 37 (health centers)
b_i	Capacity of backpack (HSA pairs)	3
a_k	Capacity of resupply center (backpacks)	150 (hospitals) or 10
f_{ik}	Cost of assigning resupply center k to backpack i (km)	$5 + D_{ik}/8$
c_{ijk}	Cost of assigning backpack i , supplied by k , to HSA j (km)	D_{ij}
g_k	Fixed cost of resupply center (km)	0
D	Distance (km)	-

Table 7 Comparison of independent and simultaneous backpack assignments. Results were calculated in Gurobi with a stopping criterion of 0.2% gap for the baseline models and 2% gap for the two echelon models

	Baseline (Health Centers)	Baseline (Hospitals)	Two Echelon (Health Centers)	Two Echelon (Hospitals)
Number of Backpacks	212	212	210	206
Number of Resupply Centers	39	2	39	2
Two Echelon Objective Value	2157	2753	2154	2733
Mean HSA to Pack Distance (km)	1.54	1.54	1.56	1.60
Max HSA to Pack Distance (km)	4.38	4.38	4.38	4.90
Mean Pack to Resupply Distance (km)	5.60	28.1	5.38	28.0
Max Pack to Resupply Distance (km)	25.7	64.3	25.7	62.2

backpacks and resupply centers both in Lilongwe District (an average distance of 5.4 km with health centers included compared to an average of 28 km without) and across Malawi. This demonstrates the importance of including health centers as backpack resupply centers, which in turn emphasizes the need for a strong chain of supplies to these small health care facilities.

Our analyses suggest that our results are fairly robust to uncertainty about the population distribution and the overall demand weights assigned to each EA. Although we did not consider fairness of HSA assignments in our initial model, our results are very similar to those produced by the p -center problem, which minimizes the maximum weighted distance between EAs and HSAs rather than the average. Additionally, our initial bottom-up approach to assigning backpacks and resupply centers produced results highly comparable to those obtained by solving the backpack and resupply center problems simultaneously.

Nevertheless, this analysis depends on a number of assumptions. We represented the population of each EA by its centroid and assumed that the population of each EA would be adequately served by at most one HSA pair, even though EAs vary in size and population. Because the HSA program is run by the Malawi gov-

ernment based on its assessments of its people's needs, which may be influenced by a more accurate determination of population characteristics, infrastructure development, and political will power, the actual locations of HSAs will probably differ from those recommended by our model, which in turn could affect both backpack and resupply center assignments. This discrepancy could lead to greater distances between HSAs, backpacks, and resupply centers than those listed here. Furthermore, health centers will have to be evaluated as potential resupply sites on an individual basis, based on information obtained on the ground and contingent on the consistency of their supply chains. In addition, backpack costs will not be limited to an initial investment, as described here, but will continue to grow as items in the backpacks are used or expire. Finally, these results are based on data from 1998, with the population extrapolated to 2008 levels. This means that our model may not fully account for changes to the Malawian demographics between 1998–2008, and surely does not account for those that occurred from 2008–present. BTB will need to consider all of these issues when planning their backpack scale-up in the field.

The techniques used in this paper could be used to plan community health worker supply systems in coun-

tries beyond Malawi. The limitations of the methods used include the large amount of data and computing power necessary to complete a sufficiently thorough analysis. However, this level of detail allowed us to answer general questions, such as how many backpacks were necessary, while also creating a specific plan for the optimal placement and resupply of every individual backpack. Optimization analyses such as this can create plans for supplying community health workers which minimize both travel times and costs, thus allowing them to provide higher quality care to the patients who need them the most.

Conflict of Interest

The authors declare that they have no conflict of interest.

Acknowledgements

A.G.K. was supported in part by National Institute of Health grant T32 AI007535. The majority of this work was conducted while the authors were undergraduates in the Department of Computational and Applied Mathematics at Rice University. We would like to thank Drs. Mark Embree, Illya Hicks, Rebecca Richards-Kortum, and Maria Oden of Rice University for their guidance throughout this project. We would also like to thank Todd Benson of the International Food Policy Research Institute for the GIS data used in this analysis; Dr. Laura McLay of the University of Wisconsin for reviewing a draft of this paper; and our anonymous referees for their helpful comments.

References

- 2008 Malawi population and housing census. National Statistical Office of Malawi
- Working Together for Health: The World Health Report 2006. The World Health Organization. URL <http://www.who.int/whr/2006/en/>
- The world factbook: Malawi. Central Intelligence Agency. URL <https://www.cia.gov/library/publications/the-world-factbook/geos/mi.html>. Last accessed June 2013
- Gurobi optimizer reference manual. Gurobi Optimization, Inc. (2013). URL <http://www.gurobi.com>
- Brunskill, E., Lesh, N.: Routing for rural health: optimizing community health worker visit schedules. Proceedings of AAAI Artificial Intelligence for Development pp. 326–334 (2010)
- Cocking, C., Flessa, S., Reinelt, G.: Locating health facilities in Nouna District, Burkina Faso. In: H. Haasis, H. Kopfer, J. Schönberger (eds.) Operations Research Proceedings 2005, pp. 431–436. Springer Berlin Heidelberg (2006). DOI 10.1007/3-540-32539-5-68
- Doerner, K., Focke, A., Gutjahr, W.: Multicriteria tour planning for mobile healthcare facilities in a developing country. European Journal of Operational Research **179**(3), 1078 – 1096 (2007). DOI 10.1016/j.ejor.2005.10.067
- Elloumi, S., Labbé, M., Pochet, Y.: New formulation and resolution method for the p -center problem (2001). URL http://www.optimization-online.org/DB_HTML/2001/10/394.html
- Hakimi, S.: Optimum distribution of switching centers in a communication network and some related graph theoretic problems. Operations Research **13**(3), 462–475 (1965). DOI 10.1287/opre.13.3.462
- Holmberg, K., Ronnqvist, M., Yu, D.: An exact algorithm for the capacitated facility location problems with single sourcing. European Journal of Operational Research **113**, 544–559 (1999). DOI 10.1016/S0377-2217(98)00008-3
- Jia, H., Ordóñez, F., Dessouky, M.: A modeling framework for facility location of medical services for large-scale emergencies. IIE Transactions **39**, 41–55 (2007). DOI 10.1080/07408170500539113
- Kadzandira, J., Chilowa, W.: The role of health surveillance assistants (HSAs) in the delivery of health services and immunisation in Malawi (2001). URL http://www.unicef.org/evaldatabase/files/MLW_01-04.pdf. UNICEF Evaluation and Research Database
- Katsulukuta, A.: Malawi: Making use of community health workers to improve coverage - opportunities and challenges. Presentation at the World Health Organization Global Immunization Meeting (2010). URL http://www.who.int/immunization_delivery/systems_policy/Making_use_community_health_Malawi.pdf
- Kliniewicz, J., Luss, H.: A lagrangian relaxation heuristic for the capacitated facility location with single-source constraints. Journal of Operational Research Society **37**(5), 495–500 (1986). DOI 10.1057/jors.1986.84
- Murawski, L., Church, R.: Improving accessibility to rural health services: the maximal covering network improvement problem. Socio-Economic Planning Sciences **43**(2), 102–110 (2009). DOI 10.1016/j.seps.2008.02.012
- Parker, B., Stansfield, S., Augustin, A., Boulos, R., Newman, J.: Optimization of task allocation for community health workers in Haiti. Socio-Economic Planning Sciences **22**(1), 3 – 14 (1988). DOI 10.1016/0038-0121(88)90029-8
- Rahman, S., Smith, D.: Use of location-allocation models in health service development planning in developing nations. European Journal of Operational Research **123**(3), 437–452 (2000). DOI 10.1016/S0377-2217(99)00289-1
- Reese, J.: Methods for solving the p -median problem: An annotated bibliography. Tech. rep., Trinity College (2005)
- Reid, R., Ruffing, K., Smith, H.: Managing medical supply logistics among health workers in Ecuador. Social Science and Medicine **22**(1), 9–14 (1986). DOI 10.1016/0277-9536(86)90302-3
- Teitz, M., Bart, P.: Heuristic methods for estimating the generalized vertex median of a weighted graph. Operations Research **16**(5), 955–961 (1968). DOI 10.1287/opre.16.5.955
- Tragantalerngsak, S., Holt, J., Ronnqvist, M.: An exact method for the two-echelon, single-source, capacitated facility location problem. European Journal of Operational Research **123**, 473–489 (2000). DOI 10.1016/S0377-2217(99)00105-8