DEPARTMENT OF ENGINEERING MANAGEMENT

**Pre-positioning of emergency supplies:
Does putting a price on human life help to save lives?**

Renata Turkeš, Daniel Palhazi Cuervo & Kenneth Sörensen

UNIVERSITY OF ANTWERP
Faculty of Applied Economics

City Campus
Prinsstraat 13, B.226
B-2000 Antwerp
Tel. +32 (0)3 265 40 32
Fax +32 (0)3 265 47 99
www.uantwerpen.be

FACULTY OF APPLIED ECONOMICS

DEPARTMENT OF ENGINEERING MANAGEMENT

Pre-positioning of emergency supplies: Does putting a price on human life help to save lives?

Renata Turkeš, Daniel Palhazi Cuervo & Kenneth Sörensen

RESEARCH PAPER 2016-012
NOVEMBER 2016

University of Antwerp, City Campus, Prinsstraat 13, B-2000 Antwerp, Belgium
Research Administration – room B.226
phone: (32) 3 265 40 32
fax: (32) 3 265 47 99
e-mail: joeri.nys@uantwerpen.be

**The research papers from the Faculty of Applied Economics
are also available at www.repec.org
(Research Papers in Economics - RePEc)**

D/2016/1169/012

Pre-positioning of emergency supplies: Does putting a price on human life help to save lives?

Renata Turkeš, Daniel Palhazi Cuervo, Kenneth Sörensen

University of Antwerp Operations Research Group ANT/OR

Prinsstraat 13, 2000 Antwerp, Belgium

Abstract

The number of people affected by natural disasters or displaced by conflict, persecution, violence or human rights violations has been steadily increasing, doubling in a decade and reaching 130.5 million in 2016. Fortunately, such trends have been accompanied by a growing research interest in the field of humanitarian logistics that investigates mechanisms which can improve assistance to disaster-affected communities and thus minimize human suffering. In spite of acknowledging a major difference between such an objective and the priorities of business logistics, many authors still adopt disaster relief problem formulations that aim to minimize costs. In this paper, we list a number of issues with the cost-minimizing approach, placing emphasis on the significant challenge of determining the controversial economic value of human suffering that is usually a part of such formulations. In order to circumvent these issues, we formulate an alternative mathematical model that maximizes response directly. The discussion about the cost-minimizing and our alternative approach is illustrated with the problem of increasing emergency preparedness by pre-positioning relief items at strategic locations. We evaluate the two formulations of the pre-positioning problem using a number of randomly generated instances and a case study focused on hurricane threat in the Gulf Coast area of the United States. The optimal solution of our model always meets at least the same percentage of demand as the cost-minimizing model, and is obtained in comparable computation time. Our study therefore suggests that putting a price on human life can and ergo should be avoided.

Keywords: Humanitarian logistics, Community operational research, Stochastic programming

1. Motivation and literature review

Every year, natural and man-made disasters have devastating effects on millions of people around the world. An estimated 400 000 people lost their lives and more than 4 000 000 were affected in only two recent major natural disasters, the Indian Ocean earthquake and tsunami, and the Haiti earthquake [1]. What is worse, both natural and man-made disasters are expected to increase another five-fold over the next forty years due to environmental degradation, rapid urbanization and the spread of HIV/AIDS in the developing world [2]. This highlights the need for effective and efficient disaster response operations. Humanitarian logistics specializes in rapid provision of relief (such as emergency shelter, food, water or medicine) to affected areas, so as to minimize human suffering and death [3]. For an introduction to humanitarian logistics or surveys of relevant literature, see [4], [5], [6].

Most of the research in humanitarian logistics focuses on the post-event distribution of aid to people in need of assistance, e.g., [7], [8], [9]. Aid distribution can be formulated as an optimization problem that decides which vehicles to employ to transport the relief items to which cities or villages that are affected by the disaster, with the goal of minimizing, for example, transportation cost or response time. However, disaster response can be further enhanced in

Email address: renata.turkes@uantwerpen.be (Renata Turkeš)

URL: <http://antor.uantwerpen.be/members/renata-turkes/> (Renata Turkeš)

the preparedness phase, before the occurrence of an emergency. Improving preparedness for natural and man-made disasters is one of the challenges of the EU Research and Innovation Horizon 2020 programme [10]. Some research suggests that every dollar invested into disaster preparedness saves seven dollars in disaster aftermath [11].

One mechanism to increase preparedness is advance procurement and pre-positioning of relief inventory at strategic locations. This allows to additionally speed up emergency assistance and save more lives by reaching areas that could be otherwise inaccessible. In turn, it makes the optimization problem much more complicated as it increases the number of decisions that need to be made, before even knowing the scale and particularities of the disaster. Indeed, next to deciding about the transportation of supplies, one must also determine the location of storage facilities and the amount of relief items to be pre-positioned, all under uncertainty about the impacts of the disaster. The relatively little research that has been conducted on this topic fails to account for a number of complex aspects inherent to the problem, such as multiple commodity types, limited vehicle capacity or uncertainties about some of the emergency characteristics. An example of a formulation that is better able to capture most of these complexities can be found in [12]. The authors formulate a mathematical model that decides on the number, location and size of storage facilities, the quantities of various types of emergency supplies stocked in each facility and the distribution of the supplies to demand locations after an event, under uncertainty about demands, survival of pre-positioned supplies and transportation network availability. The objective is to minimize a sum of different types of costs.

Minimizing costs is a typical objective in business logistics, but it does not necessarily reflect the aforementioned priority of emergency response that is to minimize immediate risks to human health and survival [13]. Even though many authors have explicitly addressed the fundamental distinctions between commercial supply chains and humanitarian relief chains ([13], [8], [2], [14], [15], [9]), cost-minimization is still a common trend in (both pre- and post-disaster) humanitarian logistics literature, e.g., [7], [16], [17], [18]. Due to a shortage of critical supplies in humanitarian relief chains which does not allow the demand to be completely satisfied, an economic value of human suffering is usually added to the cost-minimizing objective function of the mathematical models that describe the problem. Not only is this cost of human suffering controversial due to its ethical implications, it is also challenging to determine, while the quality of emergency response obtained by the model is extremely sensitive to this value.

The question naturally arises as to whether such proxies for human suffering are a necessary component of the objective function, i.e., whether they enable computing better emergency plans compared to formulations that do not incorporate such values. A study about the appropriate objective function in humanitarian logistics models [19] argues that this is indeed the case: the authors conclude that the sum of logistics and deprivation costs is the most preferred objective, where the deprivation costs represent an economic value of human suffering that is a function of time an individual has been deprived of service. Therefore, the authors decide which locations receive which amounts of aid, at a given moment in time. This means that a multi-period model is used to describe the problem, where locations with equal demands might have different deprivation times. Although we acknowledge the strengths of considering deprivation time in analytical formulations of humanitarian logistics problems, as it is better able to capture their many complexities, it is important to note that considering such assumptions has a strong effect on the outcomes of the study. This point is amply discussed in Section 5.

In this paper, we therefore aim to challenge the recommendations given in [19] by also investigating simpler single-period models that are more common in the literature. To do so, we compare the cost-minimizing model that is commonly used to describe the pre-positioning problem with an alternative formulation that, respecting the resource limitations, directly maximizes met demand. We show that our alternative model is able to achieve the same, if not better quality of emergency response, without assigning a price to human life. This helps us to expand the discussion on the appropriate choice of the objective function in a number of different settings. We hereby hope that this paper motivates further research on the development of humanitarian logistics models that are more realistic and easy to use.

We start the remainder of this paper with a formal description of the pre-positioning problem in Section 2, where we also introduce the case study and the random instances used throughout the paper. Section 3 presents the two different approaches to modelling the problem. Firstly, the common model that minimizes costs is presented, and is followed by a discussion about the issues with this formulation. These issues motivate us to present an alternative model in the

second part of the section, that also illustrates how such a formulation offers more insights and flexibility to decision makers. Numerical experiments presented in Section 4 allow us to compare the two different approaches. In Section 5, we discuss a few other alternative analytical formulations used in the literature in order to identify the features that help to make our models user-friendly for practitioners in the field. The paper ends with a summary of key findings and possible directions for future work in Section 6.

2. Problem description

This section describes our representation of an instance and a solution of the problem of pre-positioning emergency supplies, and it also includes a number of problem instances that will be used throughout this paper. An example of a definition that captures many complexities of the pre-positioning problem is found in [12]. We join the many authors that adopt this problem definition and therefore look at pre-positioning strategies that determine the number, location and size of storage facilities, the quantities of various types of emergency supplies stocked in each facility and the distribution of the supplies to demand locations after a disaster, under uncertainty about demands, survival of pre-positioned supplies and transportation network availability. These uncertainties are modelled through the use of a set of scenarios S .

2.1. Instance

Let $\{G_s = (V, E_s) \mid s \in S\}$ be a set of directed graphs, where the set of vertices V represents the cities, villages or communities that might be potential facility and/or demand locations, and the set of edges E_s represents the roads that connect them in scenario $s \in S$, with the weight of an edge (i, j) being the distance L_{ij}^s from vertex $i \in V$ to vertex $j \in V$ in scenario $s \in S$. At any vertex $i \in V$ that is a potential facility location (indicated by a binary coefficient F_i), a storage facility of a number of different categories $q \in Q$ might be open. Facility categories differ in capacity M_q and opening cost A_q . We consider a set of different commodities $k \in K$ (e.g., food, water, medicine, blankets, clothing) with V^k, B^k, C^k denoting their unit volume, unit acquisition cost and unit transportation cost respectively. These commodities may be pre-positioned at open storage facilities, and distributed further to the demand locations after the disaster strikes. In the preparedness phase, these demands are still uncertain. Next to the uncertainty about the demands, we also consider uncertainty about the survival of pre-positioned supplies and transportation network availability.

The uncertainties that are considered can be represented in a number of different ways. The two most common approaches in the literature are robust optimization that only assumes that uncertain values lie in given bounded uncertainty sets, and stochastic programming that assumes that uncertain values follow given probability distributions. The argument that is often used to speak in favour of robust optimization is that practitioners might not feel comfortable with probability distributions, whereas they would have no difficulty in, e.g., identifying the intervals that encapsulate the uncertain values. However, we argue that practitioners undoubtedly have a better knowledge of how likely certain values are to occur and in which combinations. There is no need to lose such valuable data that can help to make more informed and better pre-positioning strategies. This explicit exploitation of all available information about uncertain values is the advantage of stochastic programming that prompted us to adopt the stochastic description of uncertainty. An interesting case study that validates such a representation in disaster response optimization can be found in [7]. Therefore, in this paper, we model the uncertainties as a random vector with a finite number of possible realizations, called scenarios $s \in S$, with respective probability masses P^s . Every scenario is therefore defined by

- demand by commodity and location D_i^{ks} ,
- proportion of pre-positioned aid that remains usable, by commodity and location, R_i^{ks} , and
- transportation link usability L_{ij}^s .

This information can be obtained from historical data or some reasonable predictions. For example, FEMA's Hazus software [20] is an assessment tool that can estimate losses from potential hurricanes, floods and earthquakes.

The notation for an instance of the pre-positioning problem is summarized in Table 1.

Sets	
Q	set of facility categories
K	set of commodities
S	set of scenarios
V	set of vertices
E_s	set of edges in scenario $s \in S$
Coefficients	
F_i	$\begin{cases} 1, & \text{if a facility (of any category) can be open at vertex } i \in V \\ 0, & \text{otherwise} \end{cases}$
M_q	volume capacity of a facility of category $q \in Q$
A_q	opening cost of a facility of category $q \in Q$
V^k	unit volume of commodity $k \in K$
B^k	unit acquisition cost of commodity $k \in K$
C^k	unit transportation cost of commodity $k \in K$
P^s	probability of scenario $s \in S$
D_i^{ks}	demand of commodity $k \in K$ at vertex $i \in V$ in scenario $s \in S$
R_i^{ks}	proportion of pre-positioned commodity $k \in K$ that remains usable at vertex $i \in V$ in scenario $s \in S$
L_{ij}^s	$\begin{cases} -1, & \text{if } (i, j) \notin E_s \\ \text{distance from vertex } i \in V \text{ to vertex } j \in V \text{ in scenario } s \in S, & \text{otherwise} \end{cases}$

Table 1: Instance of the pre-positioning problem.

In addition, using the instance information, we can pre-calculate a few other coefficients that will make the models in the next section more readable. For every scenario $s \in S$, it is easy to compute the shortest path distance matrix $\Lambda^s = [\Lambda_{ij}^s]$ from the distance matrix $L^s = [L_{ij}^s]$, e.g., by employing the well-known Dijkstra algorithm [21]. Furthermore, we are interested in calculating the transportation cost between every pair of vertices in every scenario. Since the unit transportation cost varies across commodities, the transportation cost from one vertex to another will also depend on the amounts that are transported on this path. This makes it impossible to calculate the cost of traversing a path in the general case. Due to our choice of problem formulation (see Section 2.2), however, we will only be interested in the cost of traversing the shortest path from a vertex $i \in V$ to a vertex $j \in V$, transporting the amounts of aid that correspond to the demands of the destination vertex in the given scenario $s \in S$:

$$C_{ij}^s = \sum_{k \in K} C^k D_j^{ks} \Lambda_{ij}^s.$$

If there is no path from vertex $i \in V$ to vertex $j \in V$ in scenario $s \in S$, we define this transportation cost to be -1 . The notation for the auxiliary coefficients is listed in Table 2.

Auxiliary coefficients

$$\Lambda_{ij}^s \begin{cases} -1, & \text{if there is no path from vertex } i \in V \text{ to vertex } j \in V \text{ in graph } G_s = (V, E_s) \\ \text{shortest path distance from vertex } i \in V \text{ to vertex } j \in V \text{ in scenario } s \in S, & \text{otherwise} \end{cases}$$

$$C_{ij}^s \begin{cases} -1, & \text{if there is no path from vertex } i \in V \text{ to vertex } j \in V \text{ in graph } G_s = (V, E_s) \\ \text{cost of transporting the demands of vertex } j \in V \text{ via the shortest path from vertex } i \in V \text{ to vertex } j \in V \\ \text{in scenario } s \in S, & \text{otherwise} \end{cases}$$

Table 2: Auxiliary coefficients that can be derived from an instance of the pre-positioning problem.

We conclude the definition of a problem instance with a description of a number of instances that will be used throughout this study. All the instances are available from the authors on request.

2.1.1. Toy instance

A toy example of an instance with 3 cities, 2 facility categories, 2 commodities and 2 scenarios is given in Tables 3 and 4, and Figure 1. We assume that every vertex is a potential facility location, i.e., $F_1 = F_2 = F_3 = 1$.

Facility category q	Volume capacity M_q	Opening cost A_q
1 (small facility)	5000	15000
2 (big facility)	200000	125000

Table 3: Toy instance facility categories that might be open at potential facility locations.

Commodity k	Unit volume V^k	Unit acquisition cost B^k	Unit transportation cost C^k
1 (water)	100	600	3
2 (food)	60	5000	5

Table 4: Toy instance different types of aid that are in demand and might be pre-positioned at open storage facilities.

Figure 1: Toy instance graphs G_1 and G_2 represent three cities and the road network that connects them in two possible disaster scenarios, with probabilities of occurring P^1 and P^2 respectively. Every scenario $s \in S$ is defined by the demand D_i^{ks} and proportion of aid that remains usable R_i^{ks} for every commodity $k \in K$ and every vertex $i \in V$, together with the availability of every edge that is indicated in the graph.

The distance matrices in two possible disaster scenarios are the following:

$$L^1 = \begin{pmatrix} 0 & 300 & 400 \\ 300 & 0 & 400 \\ 400 & 400 & 0 \end{pmatrix}, \quad L^2 = \begin{pmatrix} 0 & 300 & 400 \\ 300 & 0 & -1 \\ 400 & -1 & 0 \end{pmatrix}.$$

2.1.2. Case study

In this study, we consider the case study addressing hurricane risks in the south-eastern area of the United States which is introduced in [12]. This instance consists of a graph of 30 vertices and 58 edges (Figure 2), 3 facility categories, 3 commodities and 51 scenario that are constructed based on historical records from a sample of 15 hurricanes in the region.

Figure 2: A case study focused on hurricane threat in the south-eastern area of the United States is used to compare the two different mathematical models that describe the pre-positioning problem.

2.1.3. Random instances

In order not to base our conclusions on a single case study, we also introduce an instance generator that returns a random instance $Random.instance_|V|_|Q|_|K|_|S|$ for a given number of vertices, facility categories, commodities and scenarios. The vertices are randomly generated and lie in the square $[0,100] \times [0,100]$, all being potential facility locations. There exists an edge between every vertex and the three vertices that are closest to it. The properties of different facility categories and commodities are initialized to some reasonable values. Scenario 0 corresponds to no disaster (zero demands and no damages), while the other scenarios have a landfall that is a random vertex, and a magnitude of the disaster that is a random integer in $[1, 5]$. Most of the vertices in each scenario have zero demands. Non-zero demand of each commodity at every vertex in each scenario is directly proportional to the magnitude of the disaster, and inversely proportional to the distance from the landfall and the unit space requirement of the commodity. In each scenario, only the aid stored at the landfall vertex would be (completely) destroyed, and for every edge there is a very small probability that it is destroyed. Probabilities of scenarios are randomly generated.

2.2. Solution

To solve the pre-positioning problem is to develop a strategy that determines:

- the number, location and category of storage facilities to open, represented with binary variables x_{iq} that indicate whether a facility of category $q \in Q$ is open at vertex $i \in V$,
- the amount y_i^k of commodity $k \in K$ to pre-position at every open facility at $i \in V$, and
- the aid distribution strategy, represented by binary variables z_{ij}^s that indicate whether a facility open at vertex $i \in V$ serves the demands of vertex $j \in V$ in scenario $s \in S$.

In fact, aid distribution decisions are not truly a part of a solution of the pre-positioning problem. Indeed, pre-positioning strategies are developed before a disaster and therefore do not need to incorporate decisions about the transportation of the pre-positioned aid from the storage facilities to the demand locations in every possible disaster scenario. In the preparedness phase, the practitioners only wish to know where to open the facilities and what to store there. However, we also determine how the pre-positioned aid would be distributed as this allows us to assess the number of people that could be reached, and thus evaluate the quality of the location and inventory decisions we consider implementing.

In our formulation, we consider the aid distribution as an assignment rather than a network flow or a routing problem: the decision variable z_{ij}^s indicates if the vertex $j \in V$ is assigned to (be fully served by) a facility open at vertex $i \in V$ in scenario $s \in S$. Often in the literature, the aid distribution is modelled as a network flow problem (e.g., [7], [12], [9]), since there exist extremely efficient algorithms to solve it. However, these formulations tend to over-simplify the distribution problem as they do not allow to take into account the capacity nor the number of vehicles needed to transport different commodities. On the other hand, if aid distribution were to be considered as a routing problem, we would have to decide on a routing scheme for each scenario $s \in S$, which is computationally demanding and is not necessarily more realistic: once a disaster happens, it is highly unlikely that it will completely match one of the considered scenarios, implying that the routing decisions made would be entirely useless.

The decision to model the aid distribution problem as an assignment problem is further motivated by the following considerations. First, serving a demand location from multiple facilities is operationally overly complex for a chaotic setting after a disaster, e.g., carrying out a plan where 20% of demand of one commodity and 73% of demand of another commodity of a vertex are served by one facility, and the remainder by another (or more) facilities. Second, making the decisions binary (a vertex is assigned to a facility or not) is also more suitable for heuristic procedures that are necessary for complex real-world problems like the problem of pre-positioning emergency supplies. Ultimately, for the purpose of this study, the choice of aid distribution formulation is irrelevant, as it will be adopted by both models of the pre-positioning problem that we wish to analyse and compare.

We assume that the aid will be transported using the shortest path from the facility to the demand vertex. When evaluating the costs, we do not worry about the vehicles returning back to their respective facilities. Indeed, this sometimes might not be possible, as a path from the vertex to the facility does not necessarily exist; but we assume that the priority is to reach a demand vertex. The number of vehicles placed at a facility will be the minimum number of vehicles needed to cover the demand of all vertices assigned to the facility across all scenarios. In a different version of the problem, we could introduce the number of vehicles to be placed at facilities as additional decisions to be made, or assume a constant number of vehicles at each facility category and impose constraints on the number of utilized vehicles.

A solution of the pre-positioning problem is summarized in Table 5.

Decision variables

x_{iq}	$\begin{cases} 1, & \text{if a facility of category } q \in Q \text{ is open at vertex } i \in V \\ 0, & \text{otherwise} \end{cases}$
y_i^k	amount of commodity $k \in K$ pre-positioned at vertex $i \in V$
z_{ij}^s	$\begin{cases} 1, & \text{if the facility open at vertex } i \in V \text{ fully meets the demands of vertex } j \in V \text{ in scenario } s \in S \\ 0, & \text{otherwise} \end{cases}$

Table 5: Solution of the pre-positioning problem.

3. Choice of objective(s)

In this section, we present two different approaches to modelling the problem of pre-positioning emergency supplies. In both cases, it is formulated as a two-stage stochastic mixed-integer programming problem, but with an important difference in the objective function. We start the section with a description of the more common approach that minimizes costs, and follow it by a discussion about what we consider as the major issues of such an approach. These issues lead us to an alternative model that, respecting the budget constraints, maximizes response directly. Following a description of our alternative model, we give a few examples that show how it also offers more flexibility and insights to decision makers.

3.1. Minimize costs

Although most researchers acknowledge many differences between humanitarian and business logistics, an interesting trend can be seen in humanitarian logistics literature: the common objective is still to minimize costs. Since it often cannot be imposed that all people in need must receive assistance, simply minimizing logistics costs would in this case yield no service provided whatsoever. In order to circumvent this, certain penalty costs for unmet demand are generally introduced and added to the objective. The motivation of an emergency plan then becomes to minimize the sum of logistics costs and these costs for unmet demand (amongst others).

The model (1)-(11) below is an adapted version of one such model that is introduced in [12] and widely adopted in the literature. In the original model, the aid distribution is formulated as a network flow problem. We adjusted this part to match the assignment problem formulation adopted earlier for the aid distribution problem. As it was suggested, next to the instance and solution information that is given in the previous section, additional coefficients U^k , O^k , and decision variables u_i^{ks} , o_i^{ks} , are introduced to help calculate the penalty costs of under- and also oversupply.

Additional coefficients

U^k unit undersupply (unmet demand) penalty cost of commodity $k \in K$
 O^k unit oversupply (holding) penalty cost of commodity $k \in K$

Additional decision variables

u_i^{ks} undersupply of commodity $k \in K$ at vertex $i \in V$ in scenario $s \in S$
 o_i^{ks} oversupply of commodity $k \in K$ pre-positioned at vertex $i \in V$ in scenario $s \in S$

Table 6: Additional coefficients and decision variables that can be introduced as a part of an instance and solution of the pre-positioning problem respectively, to facilitate minimizing costs in the formulation (1)-(11).

The model is now as follows.

$$\min \sum_{i \in V} \sum_{q \in Q} A_q x_{iq} + \sum_{i \in V} \sum_{k \in K} B^k y_i^k + \sum_{s \in S} P^s \left(\sum_{i \in V} \sum_{j \in V} C_{ij}^s z_{ij}^s + \sum_{i \in V} \sum_{k \in K} U^k u_i^{ks} + \sum_{i \in V} \sum_{k \in K} O^k o_i^{ks} \right) \quad (1)$$

$$\sum_{q \in Q} x_{iq} \leq F_i \quad (i \in V) \quad (2)$$

$$\sum_{k \in K} V^k y_i^k \leq \sum_{q \in Q} M_q x_{iq} \quad (i \in V) \quad (3)$$

$$D_i^{ks} \sum_{j \in V} z_{ji}^s + u_i^{ks} = D_i^{ks} \quad (i \in V)(k \in K)(s \in S) \quad (4)$$

$$\sum_{j \in V} D_j^{ks} z_{ij}^s + o_i^{ks} = R_i^{ks} y_i^k \quad (i \in V)(k \in K)(s \in S) \quad (5)$$

$$x_{iq} \in \{0, 1\} \quad (i \in V)(q \in Q) \quad (6)$$

$$y_i^k \geq 0 \quad (i \in V)(k \in K) \quad (7)$$

$$z_{ij}^s \in \{0, 1\} \quad (i \in V)(j \in V)(s \in S) \quad (8)$$

$$z_{ij}^s \leq 1 + \Lambda_{ij}^s \quad (i \in V)(j \in V)(s \in S) \quad (9)$$

$$u_{ik}^s \geq 0 \quad (i \in V)(k \in K)(s \in S) \quad (10)$$

$$o_{ik}^s \geq 0 \quad (i \in V)(k \in K)(s \in S) \quad (11)$$

The objective of the model is to minimize the expected costs over all scenarios, resulting from opening the facilities, aid procurement, transportation of the supplies to the demand locations, unmet demand penalty costs and holding cost for unused aid (1):

$$\min \underbrace{\sum_{i \in V} \sum_{q \in Q} A_q x_{iq}}_{\text{cost of opening facilities}} + \underbrace{\sum_{i \in V} \sum_{k \in K} B^k y_i^k}_{\text{acquisition cost}} + \sum_{s \in S} P^s \left(\underbrace{\sum_{i \in V} \sum_{j \in V} C_{ij}^s z_{ij}^s}_{\text{transportation cost}} + \underbrace{\sum_{i \in V} \sum_{k \in K} U^k u_i^{ks}}_{\text{undersupply penalty costs}} + \underbrace{\sum_{i \in V} \sum_{k \in K} O^k o_i^{ks}}_{\text{oversupply penalty costs}} \right)$$

$\min$ cost of opening facilities acquisition cost transportation cost undersupply penalty costs oversupply penalty costs

The first set of constraints (2) limits the number of open facilities at a vertex to one or zero, depending if the vertex is a potential facility location or not. The second set of constraints (3) ensures that aid is only stored at vertices where a facility is open, and limits the amount of pre-positioned aid to the capacity of that facility. The next set of constraints (4), together with constraints (8) and (10), makes certain that each vertex is assigned to at most one facility and at the same time defines the variables that represent undersupply. The set of constraints (5), together with (11), limits the amount of aid transported from a vertex to the amount of aid that remained usable, while the remaining commodities define the oversupply variables. In other words, there must be enough usable aid at each facility to fully serve the demands of all vertices that are assigned to it. Constraints (9) make sure that a demand vertex is only assigned to a facility if it can be reached from it, while the remaining (6)-(11) are positivity and integrality constraints.

An objective function that resembles the objective of the above presented model has some major drawbacks. To start with, cost minimization does not properly reflect priorities of disaster relief. In particular, pre-positioning aid at strategic locations in the preparedness phase might even end up being more costly than simply procuring the required relief items after the disaster, due to maintenance costs of storage facilities and inventory holding and spoilage costs. However, pre-positioning strategies allow speeding up emergency assistance and saving more lives by reaching areas that could be otherwise inaccessible. This corresponds to the ultimate goal of an emergency response, that is to provide assistance to as many people as possible, as soon as possible.

Incorporation of penalty costs for unmet demand

$$\sum_{s \in S} P^s \sum_{i \in V} \sum_{k \in K} U^k u_i^{ks}$$

might be perceived as somehow equivalent to minimizing unmet demand, i.e., saving more lives. However, determining the penalty costs U^k that monetize unmet demand is a significant challenge. In an extreme case where U^k is zero for every commodity $k \in K$, model (1)-(11) becomes underspecified and produces a trivial solution with no open facilities, no pre-positioned aid and no people reached. Reversely, very high unmet demand penalty costs will ensure that more people receive the needed assistance, but the logistics costs will also increase drastically, often beyond available or even reasonable budget limits. We therefore end up with an unusable emergency plan, as we lack resources to actually carry it out. In general, for different penalty costs U^k and O^k , the model (1)-(11) produces solutions with logistics costs that might be much lower or greater than a reasonable budget and thus require experimenting with a pool of different possibilities to obtain a feasible emergency strategy that uses all the available resources. In most papers, these or similar penalty cost values are merely introduced for a single case-study they are tailored to, without further elaboration or guidance on how to define them for other problem instances. For example, for the case study they present in [12], the authors take $U^k = \alpha B^k$ and $O^k = \beta B^k$, where $(\alpha, \beta) = (10, 0.25)$, but for other instances these values would likely lead to a very poor solution. Figure 3 shows how the quality of emergency response varies for different values of penalty costs U^k, O^k , for a few problem instances. For example, $(\alpha, \beta) = (10, 10)$ proposes an emergency plan for the case study that has no open facilities, no pre-positioned aid and that therefore provides no assistance to the people in need. The pair $(\alpha, \beta) = (10, 0)$ yields a solution with the unmet demand that is less than 20%, that is even further decreased with $\alpha = 100$, although the logistics costs also unpredictably increase. For another problem instance, however, $\alpha = 100$ still produces the trivial solution. Next to the difficulty in determining the penalty costs, assigning an economic value to human suffering is controversial due to its ethical implications.

It is important to note that minimizing transportation cost in general also differs greatly from minimizing response time. In this paper, we consider the response time as the sum of arrival times, but it can also be represented by the average arrival time or maximum, i.e., the latest arrival time. These objectives favour employing all available vehicles and shorter routes in order to provide assistance as soon as possible. Strong evidence that significantly better service times to beneficiaries are possible than those created by traditional cost-minimizing routing problems and algorithms is provided in [8]. It is crucial to keep in mind that the pressure of time in the relief chain is not a question of money but a difference between life and death [14].

Figure 3: The quality of emergency response changes with respect to changes in under- and oversupply penalty costs $U^k = \alpha B^k$ and $O^k = \beta B^k$ respectively (recall that B^k is the unit acquisition cost of commodity $k \in K$), for a single instance and across different instances.

3.2. Minimize unmet demand

In order to address the issues listed in the previous section, we formulate an objective function that drops the penalty costs as a proxy to maximize the response, but that directly aims to provide assistance to the greatest number of people possible, as soon as possible. This yields an alternative mathematical model that, respecting the budget constraints, minimizes unmet demand and response time in lexicographic order.

Lexicographic order of unmet demand and response time means that any emergency plan that reaches more people is always considered as superior, and only when we consider a number of plans with an equal percentage of unmet

demand, the one with lowest response time is preferred. For example, we prefer a plan where 1000 people are reached in 100 hours to a plan that reaches 900 people in 1 hour. We might also opt for a scheme that provides assistance to a slightly lower number of people, but significantly faster (where we define what slightly or significantly means for us). For instance, we might permit a percentage of unmet demand that is within 1% of the minimum unmet demand, if the response time is considerably reduced, and thus favour a solution that reaches 999 people in 1 hour rather than a solution that reaches 1000 people in 50 hours. We choose the lexicographic order of the objectives in order to avoid the complex multi-objective optimization and the issue of determining the weights if we would somehow congregate the objectives into a single objective function.

For simplicity's sake, we present here a model that only minimizes unmet demand (or, equivalently, maximizes met demand), without considering response time. We expect this to improve the readability of the section by drawing attention to the contrast between the two presented models. Note that minimizing response time further in strict lexicographic order will not make the solution worse with respect to unmet demand.

In the alternative model, the budget for opening the facilities, aid acquisition and transportation become constraints, rather than the objective. This approach seems to be much more suitable for non-profit organizations [13]. Next to the instance information that is given in the previous section, the user must therefore also provide the available budgets A , B and C for opening the facilities, aid acquisition and transportation respectively. We also introduce importance factor I^k of commodity $k \in K$ to be able to prioritize over commodities. In the cost-minimizing model, these priorities over commodities can be implicitly manipulated through different under- and oversupply penalty costs for different commodities. Finally, to improve the readability of the model, we can easily pre-calculate the fraction of the total demand of the given commodity in the given scenario that the given vertex accounts for:

$$\Delta_i^{ks} = \frac{D_i^{ks}}{\sum_{j \in V} D_j^{ks}}.$$

We set $\Delta_i^{ks} = 1$ if the total demand of commodity $k \in K$ in scenario $s \in S$ is zero.

Auxiliary coefficients

$$\Delta_i^{ks} = \begin{cases} 1, & \text{if } \sum_{j \in V} D_j^{ks} = 0 \\ \text{proportion of total demand of commodity } k \in K \text{ in scenario } s \in S \text{ that is accounted at vertex } i \in V, & \text{otherwise} \end{cases}$$

Additional coefficients

I^k	importance factor of commodity $k \in K$, $I^k \in [0, 1]$
A	total budget for opening the facilities
B	total budget for aid acquisition
C	total budget for transportation

Table 7: Auxiliary coefficients can be derived from an instance and additional coefficients can be introduced as a part of an instance of the pre-positioning problem, to facilitate minimizing unmet demand in the formulation (12)-(23).

The model is now as follows.

$$\max \sum_{s \in S} P^s \frac{1}{|K|} \sum_{k \in K} I^k \sum_{i \in V} \Delta_i^{ks} \sum_{j \in V} z_{ji}^s \quad (12)$$

$$\sum_{q \in Q} x_{iq} \leq F_i \quad (i \in V) \quad (13)$$

$$\sum_{k \in K} V^k y_i^k \leq \sum_{q \in Q} M_q x_{iq} \quad (i \in V) \quad (14)$$

$$\sum_{j \in V} z_{ji}^s \leq 1 \quad (i \in V)(s \in S) \quad (15)$$

$$\sum_{j \in V} D_j^{ks} z_{ij}^s \leq R_i^{ks} y_i^k \quad (i \in V)(k \in K)(s \in S) \quad (16)$$

$$\sum_{i \in V} \sum_{q \in Q} A_q x_{iq} \leq A \quad (17)$$

$$\sum_{i \in V} \sum_{k \in K} B^k y_i^k \leq B \quad (18)$$

$$\sum_{i \in V} \sum_{j \in V} C_{ij}^s z_{ij}^s \leq C \quad (s \in S) \quad (19)$$

$$x_{iq} \in \{0, 1\} \quad (i \in V)(q \in Q) \quad (20)$$

$$y_i^k \geq 0 \quad (i \in V)(k \in K) \quad (21)$$

$$z_{ij}^s \in \{0, 1\} \quad (i \in V)(j \in V)(s \in S) \quad (22)$$

$$z_{ij}^s \leq 1 + \Lambda_{ij}^s \quad (i \in V)(j \in V)(s \in S) \quad (23)$$

The objective of the model is to maximize the expected weighted average proportion of met demand (12). The sets of constraints (13), (14), (20)-(23) were already included and explained in the previous model (1)-(11). The set of constraints (15), together with (22), ensures that every vertex is served by at most one facility, and constraints (16) limit the aid that can be transported from a facility to the amount of pre-positioned aid that remained usable. Finally, constraints (17)-(19) represent the facility, acquisition and transportation budget limitations, respectively.

To further optimize response time in lexicographic order, we can solve an additional model that minimizes total expected response time subject to the same constraints (13)-(23), with an additional constraint that guarantees that the unmet demand is not greater than the optimal value obtained from (12)-(23). To do this, additional time matrices for every scenario $s \in S$ must be given as a part of the instance of the pre-positioning problem.

The definition of the lexicographic order can, as already noted, be relaxed by the practitioners to better describe what they believe constitutes as superior emergency plan. This can be done by replacing the additional constraint that guarantees that the unmet demand is not greater than the optimal value obtained from (12)-(23) with the constraint that ensures that the unmet demand is within a given percentage from the minimum unmet demand. By allowing the emergency strategy to meet a slightly lower percentage of demand, the response time might be significantly reduced. The time reduction might be of crucial importance, as it is able to ensure that the relief items are provided within a given time limit that avoids starvation, dehydration or an outbreak of epidemics. Figure 3 shows this correlation for a random instance with 10 vertices, 2 facility categories, 2 commodities and 5 scenarios. For a given problem at hand, this simple analysis of different strictness levels of lexicographic order can therefore also offer flexibility to practitioners as it produces a number of diverse solutions to choose from. In a way, without employing the complex tools of multi-objective optimization, the alternative model can be used to produce a set of solutions that mirror the

Pareto frontier, with the unmet demand and response time as the problem objectives.

Figure 4: The alternative model allows to easily investigate the effect of the lexicographic order definition on the response time of an emergency plan, for a given problem instance. For a random instance with 10 vertices, 2 facility types, 2 commodities and 5 scenarios, allowing the percentage of unmet demand to be 0, 1, 2, 3, 4 and 5% greater than the optimal unmet demand may reduce the total response time from 41.34 to 22.84 hours.

Another type of examination that is made straightforward by the alternative model is the study of the effect of different budget limitations on the performance of emergency response. Besides the assurance that a formulation that considers these budget limitations as constraints rather than the objective is more suitable for non-profit organizations, it offers the decision makers further insights about the problem at hand by simply running the alternative model for different interesting values of budgets A , B and C . The results of such an investigation can help to guide fund-raising efforts. For instance, if the budget for opening the facilities proves to be the most restrictive, it would be worthwhile to lobby for low or no cost storage sites offered by, e.g., government officials. The results can also be used to provide information to donors on how much spending a little more can influence the number of people reached. In addition, these simple experiments also give more flexibility as they provide a number of different solutions to choose from. Since adjusting the budgets is much more straightforward than manipulating the completely intangible and controversial penalty cost for unmet demand, we consider this model to be more user-friendly for practitioners in the field. The results of such an experiment for a random instance with 10 vertices, 2 facility types, 2 commodities and 5 scenarios can be found in Table 8 and Figure 5. We start by increasing the pre-disaster budgets A and B by a given percentage (since increasing them individually rarely improves the quality of emergency plan), proceeding to increase the post-disaster budget C , and finally increasing all three budgets simultaneously. For the given instance, we can see that increasing the budgets $(A, B) = (300000, 60000)$ even to $(A, B) = (550000, 11000)$, while keeping the budget C fixed, cannot improve the percentage of met demand (although the response time is decreased). Similarly, if the pre-disaster budgets A and B are fixed, increasing the budget C beyond 3000 does not improve the quality of the emergency plan. Also, the budgets $(A, B, C) = (150000, 30000, 3000)$ are sufficient to satisfy the demands completely, although increasing them further reduces the time that is required to achieve this. A detailed investigation of these effects can be found in [3].

Budget for facilities A (€)	Acquisition budget B (€)	Transportation budget C (€)	Unmet demand (%)	Response time (h)
50000	10000	1000	71.12	66.50
60000	12000	1000	69.13	70.68
75000	15000	1000	69.01	31.61
100000	20000	1000	59.32	35.20
150000	30000	1000	59.32	12.80
300000	60000	1000	50.88	2.65
550000	110000	1000	50.88	1.10
50000	10000	1200	70.66	66.61
50000	10000	1500	70.20	68.06
50000	10000	2000	62.57	44.66
50000	10000	3000	61.45	71.32
50000	10000	6000	61.45	71.32
50000	10000	11000	61.45	71.2
60000	12000	1200	68.18	42.16
75000	15000	1500	19.62	49.47
100000	20000	2000	8.43	40.47
150000	30000	3000	0	29.14
300000	60000	6000	0	4.72
550000	110000	11000	0	3.16

Table 8: The alternative model allows to easily investigate the effect of different pre- and post-disaster budgets A , B and C on the emergency plan's performance, for a given problem instance. The table shows the effects of a 20, 50, 100, 200, 500 and 1000% increase in pre- and/or post-disaster budgets for a random instance with 10 vertices, 2 facility types, 2 commodities and 5 scenarios.

Figure 5: The alternative model allows to easily investigate the effect of different pre- and post-disaster budgets A , B and C on the emergency plan's performance, for a given problem instance. The graphs show the effects of the 20, 50, 100 and 200% increase in pre- and/or post-disaster budgets for a random instance with 10 vertices, 2 facility types, 2 commodities and 5 scenarios.

Besides circumventing the issues of the cost-minimizing model, our model therefore also gives practitioners the possibility to easily gain some further insights and explore a number of diverse emergency plans. In this section we gave a few examples of an inspection of a single problem instance, but the same experiments on an extensive set of instances could provide general managerial implications to guide comprehensive emergency response planning.

We note, however, that the above alternative model (12)-(23) is symmetric, as the objective function (12) only registers if a demand vertex is assigned to be served or not, without taking into account which facility meets its demand (unless restricted by a tight transportation budget constraint). It was especially for this reason that it became important to also compare the computation time needed to solve the two models in the numerical experiments that were carried out, presented in the next section.

4. Numerical comparison

We compare the quality of solutions obtained using both presented models, on the toy instance, a case study focused on hurricane threat in the Gulf Coast area of the United States, and two randomly generated instances that were introduced in 2.1.1, 2.1.2 and 2.1.3.

We evaluated the two different mathematical models given in the previous section by comparing the optimal solutions obtained by CPLEX of a few problem instances. We compared the percentage of unmet demand of the emergency response strategies obtained using the two different approaches, together with the computation time needed to find those solutions. To obtain a meaningful comparison, we proceeded as follows. Firstly, to be able to solve the first model, recall that we must explicitly define some of the weights in the objective function, i.e., under- and oversupply penalty costs U^k and O^k . We defined them to be $U^k = \alpha B^k$, $O^k = \beta B^k$ (B^k is the unit acquisition cost of commodity $k \in K$), where $\alpha \in \{0, 10, 100\}$, $\beta \in \{10, 1, 0.25, 0\}$. For each pair (α, β) , we found the solution of the model (1)-(11) that minimizes the sum of logistics costs and costs for under- and oversupply. The minimum logistics costs for opening the facilities, aid acquisition and transportation were then set to be the budget limitations A , B and C respectively, that are necessary to solve the second model. Finally, after making a choice about how to (not) prioritize over commodities (we set the importance factor $I^k = 1$ for every commodity $k \in K$), we can proceed to solve the second model (12)-(23) that minimizes unmet demand.

The overview of the obtained results for the aforementioned instances is listed in Table 9. The optimal solution of our alternative model never has a greater (but often considerably smaller) percentage of unmet demand than the cost-minimizing model, and is obtained in comparable computation time.

This is to be expected as the objective function of our model directly minimizes unmet demand, but in theory, this does not have to be the case. Indeed, the transportation cost obtained from the solution of the model (1)-(11) actually corresponds to the expected transportation cost across all scenarios. The actual transportation cost in some scenarios might be significantly greater than this expected value. Constraints (19) of the model (12)-(23) prohibit these situations, as they limit the transportation cost in every scenario by the obtained expected value. In our model, we opted for these constraints as it makes more sense in practice to require that the transportation budget constraint is not violated in any disaster scenario. However, to ensure a sensible comparison of the two models with respect to the unmet demand, we would have to substitute the transportation budget constraints for every scenario (19) by a single aggregated constraint

$$\sum_{s \in S} P^s \sum_{i \in V} \sum_{j \in V} C_{ij}^s z_{ij}^s \leq C. \quad (24)$$

This allows the optimal solution of the alternative model to also have a transportation cost in some scenarios that is greater than C , as long as the expected value across all scenarios is within the specified limit. In this case, the solution of the alternative model can never meet a lower percentage of unmet demand than the cost-minimizing model, since it must respect the same constraints, while its objective is precisely the percentage of unmet demand.

Instance	Minimize costs				Minimize unmet demand	
	α	β	Unmet demand (%)	CPU time (s)	Unmet demand (%)	CPU time (s)
Toy instance	0	10	100	0.006482	100	0.003021
	0	1	100	0.002583	100	0.002487
	0	0.25	100	0.001141	100	0.000626
	0	0	100	0.000614	100	0.000628
	10	10	9.95026	0.004714	9.95026	0.001289
	10	1	9.95026	0.003749	9.95026	0.001366
	10	0.25	9.95026	0.003467	9.95026	0.001329
	10	0	9.95026	0.00347	9.95026	0.001383
	100	10	9.95026	0.003068	9.95026	0.001108
	100	1	9.95026	0.003288	9.95026	0.001079
	100	0.25	9.95026	0.003205	9.95026	0.001115
	100	0	9.95026	0.003262	9.95026	0.001098
Case study	0	10	98.26	26.8593	98.26	10.9486
	0	1	98.26	26.8639	98.26	10.9479
	0	0.25	98.26	26.8741	98.26	10.9664
	0	0	98.26	27.3035	98.26	10.8745
	10	10	98.26	27.2705	98.26	10.6422
	10	1	35.8457	3938.08	33.3831	206.01
	10	0.25	21.4873	37986.8	16.6059	8961.43
	10	0	15.2479	4323.6	13.5397	11732.8
	100	10	16.3197	11803.9	13.7236	11640.8
	100	1	5.85022	10711.8	5.1139	347.934
	100	0.25	5.46569	136618	5.19812	359.686
	100	0	5.38199	8622.23	5.20076	371.8
Random_instance_5_3_3_15	0	10	95.3662	0.015701	95.3662	0.005453
	0	1	95.3662	0.014167	95.3662	0.005234
	0	0.25	95.3662	0.0142	95.3662	0.00512
	0	0	95.3662	0.013985	95.3662	0.005186
	10	10	95.3662	0.022274	95.3662	0.005578
	10	1	95.3662	0.022703	95.3662	0.005254
	10	0.25	95.3662	0.022476	95.3662	0.005462
	10	0	95.3662	0.021025	95.3662	0.005181
	100	10	95.3662	0.022171	95.3662	0.005247
	100	1	95.3662	0.022467	95.3662	0.0055
	100	0.25	95.3662	0.022332	95.3662	0.00525
	100	0	95.3662	0.020564	95.3662	0.005467
Random_instance_25_3_3_15	0	10	97.2491	1.32098	97.2491	0.680048
	0	1	97.2491	1.3204	97.2491	0.67947
	0	0.25	97.2491	1.32072	97.2491	0.680363
	0	0	97.2491	1.32406	97.2491	0.677249
	10	10	97.2491	1.3739	97.2491	0.690039
	10	1	97.2491	1.38665	97.2491	0.702131
	10	0.25	97.2491	1.41157	97.2491	0.675797
	10	0	97.2491	1.37154	97.2491	0.67713
	100	10	55.9609	366.997	50.0074	1166.98
	100	1	53.0074	120.767	48.6323	881.433
	100	0.25	52.1057	101.069	47.9934	1050.06
	100	0	52.0968	71.154	47.9934	1505.76

Table 9: For any under- and oversupply penalty costs $U^k = \alpha B^k$ and $O^k = \beta B^k$, the optimal solution of the model that minimizes unmet demand always has lower or equal percentage of unmet demand than the model that minimizes costs, and is obtained in comparable times using CPLEX.

Our model is therefore able to achieve the same (or better) results as the cost-minimizing model, without considering any proxies that monetize human life.

5. Alternative modelling approaches

Although we proposed an alternative to the common cost-minimizing models that describes the pre-positioning problem, we acknowledge that minimizing unmet demand might not be the most suitable objective for every problem definition or setting. For example, if the available budget is not known, minimizing costs seems to be a reasonable objective to assess operational needs (although this could also be done by minimizing unmet demand subject to different budget limitations, see Table 8). We do encourage, however, to make our mathematical models more user-friendly for practitioners by separating values that do not share a common standard of measurement. Only commensurable quantities with the same fundamental dimension (e.g., units of time, money) may be compared, equated, added or subtracted, as the most basic rule of dimensional analysis instructs. Mathematical models that describe the pre-positioning problem and respect this rule of dimensional homogeneity can be formulated in a number of different ways.

One way of doing this is to formulate the problem as a multi-objective one, although this increases its complexity. Alternatively, the objectives can be considered in lexicographic order or simply separated into the objective function and constraints of the model, although there are also multiple ways to do this. In the example of assessing operational needs, for instance, the objective can be to minimize logistics costs, but respecting a constraint that ensures that at least a certain percentage of demand is met. This again allows the decision maker to change this percentage and investigate how it influences the costs that would enable such a quality of service. Adjusting the aimed percentage of demand to meet is much more straightforward than changing completely intangible and controversial unmet demand penalty costs.

It is interesting to note that, although comparing, equating, adding or subtracting incommensurable quantities violates the rule of dimensional homogeneity, the dimensions form a group under multiplication. This implies that a formulation of the pre-positioning problem that, e.g., maximizes met demand over response time respects that rule. However, note that this objective does not differentiate an emergency strategy that reaches 10 people in 10 hours to a strategy that reaches 1 person in 1 hour, and that it favours a strategy that reaches 10 people in 9 hours to a strategy that reaches 100 people in 100 hours. A constant could be added to one of the objectives to determine its relative importance, but that brings us back to the issue of determining such a constant.

Goal programming is another mechanism that allows to congregate multiple incommensurable objectives, referred to as attributes, into a single objective function without violating the rule of dimensional homogeneity. A good example of a goal programming model that describes the problem of humanitarian aid distribution can be found in [9]. The objective of the goal programming model is to minimize the weighted sum of normalized percentage deviations of the attributes from their target values. However, this model involves confronting a formidable challenge of determining the target values for each attribute. A common recommendation for the target of an attribute is the optimal value obtained by considering the attributes independently, but it requires to individually optimize each attribute. Merely minimizing costs to obtain the cost target would produce a trivial emergency plan that reaches no people, implying that an additional constraint that, e.g., ensures a certain quality of service would be imperative. Such a model is equivalent to the simple aforementioned alternative, and could be used as a standalone optimizer that describes the problem altogether.

Another study that stems from the same research question about the appropriate choice of the objective function in humanitarian logistics can be found in [19]. It is interesting to see that the conclusions of this paper seemingly contradict our findings: the authors suggest that minimizing unmet demand is the least meaningful objective found in the literature, while they promote minimizing the sum of logistic costs and (variable) costs for unmet demand. These costs for unmet demand are referred as deprivation costs and represent an economic value of the human suffering caused by lack of access to a good or service. Deprivation costs are therefore a function of the time an individual has been deprived of service, and their socio-economic characteristics (e.g., age, gender). However, it is crucial to note that

the authors aim to decide in what order, how much and when to deliver to demand vertices during a given planning horizon (a number of deliveries has to be made at each demand location during the horizon), thus using a multi-period model to describe their problem. They also consider different urgency with which supplies may be needed at different locations. This means that vertices with equal demands can have different deprivation times, which obviously makes minimizing unmet demand a poor choice of the objective. Calculating deprivation costs bears similar problems as calculating the penalty costs for unmet demand: they are controversial since they put a price on human suffering and they are very difficult to calculate, what the authors themselves acknowledge. This becomes even more difficult if we consider multiple types of commodities and/or uncertainties about the demands (the paper only focuses on post-event distribution of one type of commodity). In the numerical experiments presented in the paper, these costs are calculated using willingness to pay for a commodity and statistical value of life. It is probably too optimistic to assume that practitioners would have this kind of information for any type of commodity and any country in the world. In addition, such a formulation does not allow for much flexibility: if the total logistic costs turn out to be too large, the decision maker is not able to make straightforward modifications to the model in order to explore a number of possibilities. We believe that the introduction of deprivation times is definitely a valuable contribution to the field, as it is better able to capture the complexities of many humanitarian logistics problems. This is especially true if the deprivation time is properly adjusted to take into account the non-linearity of human suffering (e.g., a two-day shortage of relief items is for an individual not necessarily "twice as bad" as a one-day shortage of the same amounts of supplies). However, we do not find it necessary to further transform such a deprivation time into a cost. Considering logistic costs and deprivation time separately, we can opt for one of the previously proposed alternatives. For example, the objective could be to minimize the total deprivation time, while respecting given budget constraints, what is similar to the alternative model we propose in this paper, as it aims to optimize the response directly. Reversely, logistics costs could be minimized, while ensuring that the deprivation time is under a certain limit.

Separating incommensurable values in the analytical formulations of humanitarian logistics problems eliminates the challenging task of determining intangible values that are necessary to congregate the different objectives of emergency response. The resulting mathematical models are therefore much more easy-to-use for practitioners in the field.

6. Conclusions and future research

In this paper, we aimed to challenge the most common choice of the objective function in humanitarian logistics literature. For the problem of pre-positioning emergency supplies that we used to illustrate our discussion, the common objective is to minimize costs that result from logistics costs (opening the facilities, aid acquisition and distribution) and unmet demand (and holding) penalty costs. Besides the fact that defining these unmet demand penalty costs forces the decision makers to assign an economic value to human suffering, we demonstrate that the quality of the emergency response that is obtained by such a model is extremely sensitive to these intangible values.

The issues of the cost-minimizing models led us to propose an alternative formulation that directly minimizes unmet demand. In our experiments, the optimal solution of the latter model never met a lower percentage of demand, and the solution was obtained in comparable computation time. This shows the existence of an alternative that is able to remedy the aforementioned problematic issues of the cost-minimizing model, without having additional drawbacks. What is more, we give a few examples that show how the alternative model offers more flexibility and insights about the problem at hand. The results therefore suggest that putting a price on human life can and thus should be avoided.

Nevertheless, we acknowledge that there do not exist only two polar opposites when it comes to tackling humanitarian logistics problems, nor that it is easy to make recommendations about the best choice of the objectives regardless of the context or the problem definition. For this reason, we extended the discussion further to other choices of objectives and analytical formulations. We argue that is possible and recommendable to make our mathematical models more user-friendly for practitioners by eliminating any abstract and imperceptible weights or penalties.

Hopefully, this discussion leastwise invites researchers to elaborate the choice of specific objectives and other modelling decisions in their writing and opens the floor for further dialogue on the topic. For instance, a promising

objective seems to be deprivation time, as it is better able to capture the complexities of humanitarian relief chains by considering the demands together with the urgency with which the supplies may be needed. So far, the standard approach that considers the deprivation time aims to minimize the sum of logistics costs and the deprivation cost that is function of the deprivation time. It would therefore be interesting to compare such a formulation to a mathematical model that aims to minimize the deprivation time directly, without taking the next step that translates the deprivation time into a cost to add to logistics costs.

Acknowledgements

This research was supported by the Interuniversity Attraction Poles (IAP) Programme on Combinatorial Optimization: Metaheuristics and Exact Methods (COMEX) initiated and funded by the Belgian Science Policy Office (BELSPO).

References

- [1] <http://www.emdat.be/>.
- [2] A. S. Thomas, L. R. Kopczak, From logistics to supply chain management: the path forward in the humanitarian sector, Fritz Institute 15 (2005) 1–15.
- [3] B. Balcik, B. M. Beamon, Facility location in humanitarian relief, International Journal of Logistics 11 (2) (2008) 101–121.
- [4] N. Altay, W. G. Green, OR/MS research in disaster operations management, European Journal of Operational research 175 (1) (2006) 475–493.
- [5] A. M. Caunhye, X. Nie, S. Pokharel, Optimization models in emergency logistics: A literature review, Socio-Economic Planning Sciences 46 (1) (2012) 4–13.
- [6] M. Çelik, O. Ergun, B. Johnson, P. Keskinocak, A. Lorca, P. Pekgün, J. Swann, Humanitarian logistics, appears in Tutorials in Operations Research, Mirchandani P.(ed.), INFORMS, Maryland (2012) 18–49.
- [7] G. Barbarosolu, Y. Arda, A two-stage stochastic programming framework for transportation planning in disaster response, Journal of the Operational Research Society 55 (1) (2004) 43–53.
- [8] A. M. Campbell, D. Vandenbussche, W. Hermann, Routing for relief efforts, Transportation Science 42 (2) (2008) 127–145.
- [9] B. Vitoriano, M. T. Ortuño, G. Tirado, J. Montero, A multi-criteria optimization model for humanitarian aid distribution, Journal of Global Optimization 51 (2) (2011) 189–208.
- [10] <https://ec.europa.eu/programmes/horizon2020/>.
- [11] http://www.undp.org/content/undp/en/home/ourwork/get_involved/ActNow.html.
- [12] C. G. Rawls, M. A. Turnquist, Pre-positioning of emergency supplies for disaster response, Transportation Research Part B: Methodological 44 (4) (2010) 521–534.
- [13] B. M. Beamon, B. Balcik, Performance measurement in humanitarian relief chains, International Journal of Public Sector Management 21 (1) (2008) 4–25.
- [14] L. N. Van Wassenhove, Humanitarian aid logistics: supply chain management in high gear, Journal of the Operational Research Society 57 (5) (2006) 475–489.
- [15] L. N. Van Wassenhove, A. J. Pedraza Martinez, Using or to adapt supply chain management best practices to humanitarian logistics, International Transactions in Operational Research 19 (1-2) (2012) 307–322.
- [16] J. Chapman, L. B. Davis, F. Samanlioglu, X. Qu, Evaluating the effectiveness of prepositioning policies in response to natural disasters, International Journal of Operations Research and Information Systems 2014 (5 (4)).
- [17] L. B. Davis, F. Samanlioglu, X. Qu, S. Root, Inventory planning and coordination in disaster relief efforts, International Journal of Production Economics 141 (2) (2013) 561–573.
- [18] I. de Brito Junior, A. Leiras, H. T. Y. Yoshizaki, Stochastic optimization applied to the prepositioning of disaster relief supply decisions in Brazil11, Technical Report POMS 2013.
- [19] J. Holguín-Veras, N. Pérez, M. Jaller, L. N. Van Wassenhove, F. Aros-Vera, On the appropriate objective function for post-disaster humanitarian logistics models, Journal of Operations Management 31 (5) (2013) 262–280.
- [20] <http://www.fema.gov/hazus>.
- [21] E. W. Dijkstra, A note on two problems in connexion with graphs, Numerische mathematik 1 (1) (1959) 269–271.