NORTHWESTERN UNIVERSITY

Operations Management of Food Recovery Programs

A DISSERTATION

SUBMITTED TO THE GRADUATE SCHOOL

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS

for the degree

DOCTOR OF PHILOSOPHY

Field of Industrial Engineering and Management Sciences

By

Joseph Warfel

EVANSTON, ILLINOIS

December 2017

ABSTRACT

Operations Management of Food Recovery Programs

Joseph Warfel

Food recovery programs (FRPs) divert potential waste at grocery stores so that it can be distributed to people who do not have enough food. FRPs are administered by food banks, nonprofit organizations dedicated to the alleviation of hunger. The primary purpose of FRP is to collect donations. Eventually, the food is distributed to other nonprofit organizations (referred to as “agencies”) which in turn provide it to families and individuals. A few food banks include agencies on FRP routes, a practice that is becoming more common. This innovation presents opportunities and challenges: the presence of agencies allows the food bank to reduce the required vehicle capacity and more quickly distribute perishable food, but donations are random, so it is difficult to provide consistent service to the agencies.

In this dissertation, we study three closely related models of FRP operations.

The *one-commodity pickup and delivery allocation problem* (1-PDA) models allocation decisions for a given FRP route. The objective of the 1-PDA is to minimize the required

vehicle capacity. We develop a simple three-step algorithm, the *MILB algorithm*, that obtains an optimal solution to the 1-PDA.

We augment the 1-PDA with agency selection and node sequencing decisions to formulate the *selective 1-PDTSP with stochastic supply* as a mixed-integer linear program (MILP). It is possible to solve the problem with a MILP solver, but the solution time is prohibitive for many realistic instances. Therefore, we propose a heuristic procedure, the *capacity reuse insertion heuristic* (CRIH), based on inserting agencies into existing FRP routes. In a case study based on data provided by Northern Illinois Food Bank, we obtain insights regarding agency selection and node sequencing for FRP. We also demonstrate that CRIH provides near-optimal solutions.

To model FRP operations at food banks where routing is inflexible and the food obtained from FRP is crucial to agency operations, we generalize the 1-PDA to model the one-commodity pickup and delivery allocation problem for agency-supporting FRP (the *1-PDA-as*). The 1-PDA-as differs from the 1-PDA by including parameters that specify additional service requirements at donors and agencies. The objective of the 1-PDA-as is to maximize total donations collected for a given route. By applying several reformulations, we develop an optimal solution procedure for the 1-PDA-as that relies on solving a series of linear programs; however, this solution procedure cannot be applied to many realistic instances due to issues of numerical precision. Therefore, we propose a heuristic solution procedure based on the MILB algorithm. In a case study, we obtain insights about node parameters and node sequencing. We also demonstrate that the heuristic generates near-optimal solutions.

Dedication

Kini nga buhat gipahinungod ngadto ni Birhen Maria ug ni Señor Santo Niño.

Table of Contents

List of Tables

List of Figures

CHAPTER 1

Introduction

> Starvation is the characteristic of some people not *having* enough food to eat. It is not the characteristic of there not *being* enough food to eat.
>
> Amartya Sen
>
> *Poverty and Famines*

Abundant food is available in the United States: "3,900 calories per day for every man, woman, and child in the country, whereas the average adult needs only a bit more than half that amount, and children much less" [46]. Despite this abundance, millions of Americans experience hunger. The United States Department of Agriculture (USDA) reported that 12.7% of households (15.8 million) were food insecure in 2015, meaning that they "had difficulty at some time during the year providing enough food for all members due to a lack of resources." Furthermore, 5.0% of households (6.3 million) had very low food insecurity, a more severe state in which "the food intake of some household members was reduced and normal eating patterns were disrupted at times during the year due to limited resources" [14].

Simultaneously, enormous quantities of food are wasted: Between a quarter and a half of the food in the United States (160 to 320 billion pounds annually) is wasted on farms, in transit, at grocery stores, and in homes and restaurants [8]. At grocery stores, edible and nutritious food is wasted because of store policies that require culling items before their sell-by date; ordering far more of an item than demanded due to manager inexperience or data entry errors; and, most often (because grocery store managers feel obligated to have everything that the customer wants), overstocking or overproducing due to the fear of missing a sale [8, 55].

In this work, we study *food recovery programs* (FRPs), programs which divert potential waste at grocery stores so that it can be distributed to people who do not have enough food. FRPs are administered by food banks, nonprofit organizations dedicated to the alleviation of hunger. In the most limited sense, a food bank is a warehouse where donated and recovered food is stored until it is distributed to other nonprofit organizations (referred to as "agencies") which in turn provide it to families and individuals. In practice, food banks have also assumed much of the responsibility for transporting donated food to their warehouses and for distributing food from warehouses to agencies. There are approximately two hundred food banks in the United States. In 2014, Feeding America (the national organization of food banks) reported that 46.5 million unique clients receive food assistance through food banks each year [19].

The primary purpose of FRP is to collect donations of food from grocery stores. However, a few food banks include agencies on their FRP routes, a practice that is becoming more common. This innovation presents challenges and opportunities, which we study in this work. For motivation and context, we focus on the operations of two food banks: Northern Illinois Food Bank (NIFB) and the Greater Chicago Food Depository (GCFD).

In Chapter 2, we review the literature relevant to our work.

In Chapter 3, we describe the framework we use to model FRP. We state our modeling assumptions regarding routes, donors, and agencies, and define notation used for models in subsequent chapters.

The purpose of including agencies on FRP routes differs among food banks. At NIFB, agencies are included to extend the capacity of the vehicle. Although agencies value the food they obtain through FRP, those allocations are supplemental to food provided through other NIFB programs. At GCFD, FRP allocations play a larger role in the operations of agencies. For some GCFD agencies, FRP may be the only consistent source of some types of perishable food. We develop and analyze models for FRP at both types of food banks: those where agencies are included primarily to extend vehicle capacity, and those where FRP plays an essential role in supporting agency operations.

In Chapter 4, we introduce the selective one-commodity pickup and delivery traveling salesman problem with stochastic supply (the *selective 1-PDTSP with stochastic supply*). This problem models FRP at food banks (such as NIFB) that serve a large region, in which agencies are included on FRP routes primarily to allow smaller vehicles to be used. As an initial step in our formulation of the selective 1-PDTSP with stochastic supply, we analyze the allocation decision for a fixed route, which we deem the one-commodity pickup and delivery allocation problem (the 1-PDA). We present an algorithm that obtains an optimal solution for the 1-PDA. We then use the structure of that algorithm as the basis for a linear program formulation of the selective 1-PDTSP with stochastic supply.

To solve the selective 1-PDTSP with stochastic supply, we must choose which agencies to include on the route; find a tour that satisfies a travel time constraint; and determine how much food to place in the vehicle before leaving the food bank (the initial load) and how much food to distribute at each agency (the allocation policy). In making these decisions, we seek to minimize the vehicle capacity while accounting for the randomness in donations. Although the linear program formulation can be solved to optimality, food bank staff do not have access to or experience with sophisticated optimization tools; furthermore, many realistic instances cannot be solved in a reasonable amount of time. Therefore, we also propose a heuristic that obtains solutions by inserting agencies into existing routes of donors. We conclude with a case study based on NIFB data.

In Chapter 5, we generalize the 1-PDA to model the one-commodity pickup and delivery allocation problem for agency-supporting FRP (the *1-PDA-as*) as a stochastic program. The 1-PDA-as models the allocation decision at food banks (such as GCFD) where routing is inflexible and the food obtained from FRP is crucial to agency operations. By reformulating the problem as a constrained Markov decision process (CMDP), we obtain an optimal solution procedure, but due to issues of numerical precision, it cannot be applied to general 1-PDA-as instances. We also propose a heuristic solution procedure for the 1-PDA-as. We conclude with a case study to obtain insights about FRP routes and to evaluate the quality of heuristic solutions.

In Chapter 6, we summarize the contributions of this research and outline possible directions of future work regarding the selective 1-PDSTP with stochastic supply and the 1-PDA-as, as well as a more general model of food bank logistics. We conclude by highlighting the potential benefits of further collaboration between food banks and Operations Research.

CHAPTER 2

Literature review

In §2.1, we discuss articles in the operations research literature with a focus similar to our work. In the remaining sections, we review the literature relevant to the two models that are our primary objects of study. In §2.2, we review the literature relevant to the selective 1-PDTSP with stochastic supply, in particular the 1-PDTSP with stochastic demands defined and studied by Louveaux and Salazar-González [38]. In §2.3, we describe problems similar to the 1-PDA-as.

2.1. Charitable food distribution in the OR literature

Our work joins a growing literature in the application of operations research to the challenges of charitable food distribution.

Several authors have formulated and studied models of FRP or similar systems. Lien et al. model allocation policies for the FRP at GCFD [37]. They model the problem as a *sequential resource allocation* problem (SRA), but with an objective that considers equity (SRA-e). Simultaneously, they seek to reduce the amount of donated food left unallocated at the end of the route. Balçık et al. extend the single-route analysis of Lien et al. to consider multiple routes, which they solve with a decomposition-based heuristic [5]. Güneş et al. model Three Rivers Table, a program similar to FRP at Greater Pittsburgh Community Food Bank, as a one-commodity pickup-and-delivery vehicle routing problem

(1-PDVRP) [28]. They apply three solution techniques to the problem: mixed integer programming, constraint programming, and constraint-based local search.

Nair et al. formulate a periodic unpaired pickup-and-delivery vehicle routing problem (PU-PDVRP) to model FRP at OzHarvest, a food rescue organization in Sydney, Australia, which they solve heuristicially [42]. A later article expands the model by considering perishable and nonperishable food, and proposes a more efficient meta-heuristic based on Tabu search [43]. In another paper about FRP at OzHarvest, Nair et al. propose the Food Rescue Allocation and Routing Problem (FRARP), a bi-objective formulation that simultaneously optimizes routing cost and the equity of the allocation decision [45]. They solve the FRARP with a two-stage goal programming approach.

Davis et al. model a version of FRP in which the first node on a route may be a "Food Delivery Point" (FDP), a site to which several agencies travel to receive food from the food bank warehouse *and* where a FRP donation is collected [16]. (The FDP is a grocery store serving as a satellite delivery location.) After visiting the FDP, the vehicle can visit additional donors, but no further deliveries are made. They model the problem using a two-stage approach: a capacitated set covering problem to assign agencies to FDPs, then a periodic vehicle pick-up and delivery problem with backhauls to determine the collection and delivery schedule over a finite planning horizon. They present a case study for operations at the Second Harvest Food Bank of Northwest North Carolina.

A separate strand of research focuses on forecasting FRP donations. Brock and Davis evaluate four approximation methods (including an artificial neural network and multiple linear regression) to estimate donations from grocery stores, using data provided the Food Bank of Central and Eastern North Carolina (FBCENC) [10]. They also estimate the transportation costs associated with using each approximation method, which they compare to FBCENC's actual transportation costs. Nair et al. evaluate several techniques (including neural networks and multiple linear regression) to forecast OzHarvest's donations from grocery stores, restaurants, and bakeries [44]. Phillips et al. develop a statistical model for the quantity of food available from donors through the FRP at Community Food Share (CFS), a food bank in north central Colorado [50]. They use their model in a Monte Carlo simulation to study the conditions under which FRP donations suffice to meet the demand of CFS's agencies, and at what transportation cost.

Models of FRP vary in whether they focus on donors or agencies. In the existing research motivated by FRP at GCFD, the allocation decision at agencies is modeled intricately, but donations are considered deterministic within the context of a single problem instance [5, 37]. Conversely, the research on forecasting FRP donations focuses exclusively on donors [10, 44, 50]. The present models of FRP that include a routing decision define donors as nodes with a deterministic supply and agencies as nodes with a deterministic demand [16, 28, 42, 43, 45]. Our work is unique because it incorporates the randomness of donations while providing a rich model of the allocation decision at agencies.

Researchers have also studied the distribution of food via scheduled deliveries from the food bank warehouse. The vehicle routing with demand allocation problem (VDRAP), first formulated by Ghoniem et al. as a location routing problem (LRP) variant, is motivated by a distribution system used at some food banks in which, instead of delivering directly to agencies, food bank vehicles visit intermediate delivery points where agencies pick up food [25]. The objective of the VDRAP is to minimize the weighted sum of the total distance traveled by food bank and agency vehicles. Several heuristics have been proposed for the VDRAP: Ghoniem et al. apply a relax-and-fix technique and a specialized column generation approach [25]; Solak et al. develop a two-stage sequential solution approach and two procedures based on Benders decomposition [59]; and Reihaneh and Ghoniem propose a multi-start heuristic with local improvement and perturbation schemes that yields near-optimal solutions within a few seconds [53].

Örgüt et al. study the problem of equitable food allocation at FBCENC [48]. FBCENC must distribute food to each of the 34 counties it serves in proportion to the number of people living in poverty. However, counties (to be precise, the agencies in each county) are limited in their capacity to accept and store food. Örgüt et al. formulate and solve the Food Distribution Model, which minimizes the total amount of food not distributed while meeting FBCENC's requirements. They also formulate the Capacity Allocation Model to determine how additional storage capacity should be allocated to counties to improve the objective of the Food Distribution Model. Fianu and Davis further develop the Food Distribution Model to include stochasticity in supply, proposing a discrete-time, discrete-state (DTDS) Markov decision process (MDP) in which county-level allocation decisions are made monthly from random donations and transfers of food [22]. They compare several allocation policies for the DTDS MDP.

Gleaning, the collection of unharvested food from farms by volunteers, is an important source of food for some food banks. Sönmez et al. model the problem of making a gleaning schedule at the Food Bank of the Southern Tier in New York as a stochastic optimization problem, which they solve with a simulation-optimization approach [60]. Ata et al. approach the problem of scheduling gleaning events by developing a queuing model, from which they obtain a dynamic staffing policy [3].

The OR literature includes papers about various other aspects of food bank and agency operations. Chou et al. study optimal delivery routes for a charity in Singapore that organizes volunteers to gather donations of unsold bread from bakeries and transport it to homes for the elderly [13]. Cuevas-Ortuño and Gómez-Padilla propose a mixed integer program to model the problem of preparing boxes using a mixture of donated and purchased foods at the Banco de Alimentos de México (the food bank of Mexico City) [15]. The boxes are personalized by family and must meet a set of nutritional requirements. Mohan et al. develop a simulation model to improve warehouse operations and recommend changes to food distribution practices at the Society of Saint Vincent de Paul in Phoenix, Arizona [41]. Yıldız et al. propose a LRP variant to model the Meals on Wheels program in Allegheny County, Pennsylvania [69]. They develop a memetic algorithm to solve the problem. Hemmelmayr et al. study a variant of the periodic location routing problem that models the disposal of cardboard by anti-hunger agencies in Chicago, Illinois [29].

2.2. Literature review for the selective 1-PDTSP with stochastic supply

In routing problems, vehicle capacity is often modeled as a resource (i.e., a constraint) that limits the set of feasible solutions. However, the required capacity can also be interpreted as a value impacted by a set of logistic decisions, as is the case in pickup and delivery problems (PDPs). In a PDP, each node either supplies or demands a set of commodities. For many-to-many PDPs [6], including the one-commodity case, intelligently planned routes permit vehicle capacity to be "re-used:" the demand at one node can be satisfied with the supply picked up at a node visited earlier, freeing vehicle capacity to accept additional supply.

In the one-commodity pickup-and-delivery travelling salesman problem (1-PDTSP), every node other than the depot has either a positive or negative demand of a homogeneous commodity [30, 31, 32]. Nodes with positive demand are *pickups*, while those with negative demand are *deliveries*. Both types of demand are deterministic. A feasible solution to the 1-PDTSP is a tour that departs the depot (possibly with some initial load), visits each customer exactly once to satisfy its demand, then returns to the depot, maintaining a feasible vehicle load throughout.

In the literature, the model most similar to the selective 1-PDTSP with stochastic supply is a generalization of the 1-PDTSP by Louveaux and Salazar-González [38]. They formulate the *1-PDTSP with stochastic demands* by defining the demand at every node as a random variable with a discrete probability distribution. For the 1-PDTSP with stochastic demands, the vehicle capacity is not a parameter; rather, the vehicle capacity is minimized through the choice of the initial load and the tour.

Louveaux and Salazar-González define three versions of feasibility for the 1-PDTSP with stochastic demands: *sufficiency*, *adaptability*, and *survivability*. A *scenario* is the particular sample path that results when each of the random demands is observed. The three versions differ in whether the initial load and tour must be chosen before or after the scenario is known. If neither the tour nor the initial load must be decided before demands are observed, then the vehicle capacity is *sufficient* if a feasible tour and initial load can be found for every possible scenario. If the tour must be decided before demands are observed, but the initial load may be decided afterwards, then a tour is *adaptable* if a feasible initial load can be found for every possible scenario. If both the tour and the initial load must be decided before demands are observed, then a tour is *survivable* if there exists an initial load that is feasible for every possible scenario.

The selective 1-PDTSP with stochastic supply is closely related to the 1-PDTSP with stochastic demands in that we seek to minimize capacity through the same decisions (the initial load and the route). For our motivating context, survivability is the relevant version of feasibility, since both the initial load and the node sequence of an FRP route must be chosen before the donations are observed. We extend the 1-PDTSP with stochastic demands by incorporating two additional decisions: allocation and node selection.

In the 1-PDTSP with stochastic demands, the demand at any node is a single value which must be satisfied exactly; in the selective 1-PDTSP with stochastic supply, an allocation decision must be made at each agency. Demand at an agency is as a range based on the population the agency serves and its ability to store perishable foods. Agencies' flexibility in accepting a range of allocations can be used by the food bank to mitigate the impact of the randomness in donations.

The selective 1-PDTSP with stochastic supply further differs from the 1-PDTSP with stochastic demands because, although all donors must be visited, the food bank chooses which agencies to include in the FRP route. In standard PDPs, all nodes are included on the route [6]. An exception is the single-vehicle routing problem with unrestricted backhauls [62], in which all deliveries are compulsory, but a set of optional pickups is available. Each pickup provides some profit, which can be used to offset the total travel time of the route. In a sense, this problem is complementary to the setting of FRP, in which pickups are compulsory and deliveries are optional. However, unlike the FRP, the single-vehicle routing problem with unrestricted backhauls is a one-to-many-to-one PDP (using the classification of Berbeglia et al. [6]), in which the commodities picked up cannot be used to satisfy demand at delivery nodes.

Although the node selection decision is rarely considered in the PDP literature, there are contexts in which such a decision is relevant. Several well-known TSP variants, collectively called *TSPs with profits*, include a node selection decision [21]. Some examples are the orienteering problem (also known as the selective TSP) [36, 65] and the prize-collecting TSP [4, 7]. In TSPs with profits, a profit value is associated with each node. The objective is to find a circuit that maximizes total collected profit subject to a constraint on maximum total travel time; or to find a circuit that minimizes total travel time subject to a constraint on minimum total collected profit; or to minimize total travel time less total collected profit. There are a few variants of TSPs with profits in which including some subset of the nodes on the route is compulsory (akin to FRP donors) while the remaining nodes are optional (akin to FRP agencies) [24, 61].

The multiple-vehicle equivalent of 1-PDTSP, the one-commodity pickup-and-delivery vehicle routing problem (1-PDVRP), is used by Güneş et al. to model a program similar to FRP at the Greater Pittsburgh Community Food Bank [28]. They optimize vehicle routes for the Three Rivers Table program based on deterministic estimates of demand and supply obtained from historical data. van Anholt et al. formulate a problem similar to the 1-PDTSP, the *inventory-routing problem with pickups and deliveries* (IRPPD), to find routes for a fleet of armored trucks that collect and deliver cash at recirculation automated teller machines [63]. They obtain optimal solutions for realistically sized problems using an exact branch-and-cut algorithm. Although both of these papers focus on applications similar to FRP, neither model addresses randomness in supply, an essential property of FRP donations.

The selective 1-PDTSP with stochastic supply contributes to the literature on PDPs by introducing more flexible operating constraints regarding demand. In particular, we allow delivery nodes to accept a range of allocations, and we incorporate node selection into the routing aspect of the problem. Our work extends the survivability version of the 1-PDTSP with stochastic demands, but the additional flexibility in our problem requires additional decisions: not only must we obtain a route and initial load that minimize the required vehicle capacity, but we must also choose which agencies to include on the route and determine an allocation policy at each agency.

2.3. Literature review for the 1-PDA-as

The 1-PDA-as resembles stochastic variants of the *inventory routing problem* (IRP) and the *vehicle routing problem* (VRP), with several key differences. The IRP models the distribution of single commodity from a single facility to a set of demand nodes over a (possibly infinite) planning horizon [12]. Some IRP variants include consideration of the allocation decision (which is absent in the 1-PDTSP); however, even in stochastic variants, all allocation decisions and routing decisions are made before the vehicle leaves the depot [33, 34]. Some variants of the stochastic vehicle routing problem (SVRP) consider uncertainty in demand [26, 47, 58]; however, in these, it is possible to make routing changes in response to stochastic information on route. Despite the resemblance between the context of our research and that of other routing problems, the mixture of supply and demand nodes requires fundamentally different modeling choices, so results regarding the IRP and SVRP are not directly applicable to the 1-PDA-as.

In the literature, the problem that most closely resembles the 1-PDA-as is the SRA-e proposed by Lien et al. to model allocation policies for FRP at GCFD [2014] and extended to the multiple-vehicle case by Balçık et al. [5]. The 1-PDA-as differs from their model primarily due to two modeling choices.

First, we consider a wider variety of route structures for FRP. Lien et al. considered only routes in which all donors were visited before all agencies. (At the time of their observations, this was the only type of route used at GCFD, and it is still the most common FRP route structure.) However, this type of route structure does not allow agencies to be used to extend the capacity of the vehicle. Therefore, the 1-PDA-as models FRP routes in which the donors and agencies are visited in any order.

As a consequence of the greater variety of route structures, the 1-PDA-as models donors differently. Although Lien et al. recognize that donations are random, there is no need to model donors for the SRA-e: decisions are only made at agencies, which are visited after all donors have been visited. In contrast, the 1-PDA-as models the randomness of donations explicitly.

Second, we model the allocation decision at agencies differently. In the SRA-e, agencies express their demand according to a known probability distribution. For the 1-PDA-as, the need for food at agencies is not modeled as random. Rather, it is expressed through a set of parameters that represent acceptable service (described in §3.2.3). In essence, this modeling choice represents a different conception about the role of food provided through FRP in agency operations and the amount of information available to the food bank about its agencies.

CHAPTER 3

Modeling framework

In §4 and §5, we propose and analyze several models related to FRP operations. Although each model has a distinct scope and purpose, they share a common framework and notation, which we describe in the present chapter.

Our notation for the components of a FRP route is based on the decomposition of the route into "segments," which we define in §3.1. We then present our modeling approach for each type of node: the food bank warehouse (§3.2.1), donors (§3.2.2), and agencies (§3.2.3). In §3.3, we define a set of summary values for FRP routes.

3.1. Segments

The fundamental structure of a FRP route is the dyad of a donor followed by an agency. The donor provides food that is allocated at the agency, and that allocation frees capacity in the vehicle to accept further donations. For agencies to serve the role of freeing capacity, a donor must be immediately followed by an agency somewhere in the FRP route. We refer to this fundamental structure as a "segment." That is, a *segment* consists of a donor immediately followed by an agency.

We decompose any FRP route into n segments, indexed $i \in \{1, \ldots, n\}$. Donors and agencies are designated by their segment; that is, Donor i and Agency i comprise Segment i. When necessary, "dummy" nodes (a donor with no supply or an agency with no demand) are introduced to induce the segment structure. For example, consider the following FRP route, in which donors are represented as D and agencies as A:

$$D\ A\ D\ D\ A\ A\ D$$

We augment the route with dummy nodes, denoted D^0 and A^0, to create the pattern of donors followed by agencies:

$$D\ A\ D\ A^0\ D\ A\ D^0\ A\ D\ A^0$$

The route then consists of five segments:

$$D\ A \mid D\ A^0 \mid D\ A \mid D^0\ A \mid D\ A^0$$

The number of segments in a route depends not only on the number of donors and agencies, but also on their order. For a route of $|\mathcal{D}|$ donors and $|\mathcal{A}|$ agencies, the number of segments n is bounded by:

$$\max\{|\mathcal{D}|, |\mathcal{A}|\} \leq n \leq |\mathcal{D}| + |\mathcal{A}| \quad (3.1)$$

The concept of a segment is not necessary to model FRP. We choose to model FRP with segments because it provides a standardized modeling framework. Unlike similar problems (such as the the 1-PDTSP with stochastic demands [38]), donor and agency nodes are different not only in that one has positive "demand" and the other has negative "demand"; rather, they are materially different types of nodes, at which different decisions must be made. By modeling the FRP route as a series of segments, a donor is always followed by an agency. This simplifies our problem formulations and other algebraic statements, because we need not make conditional statements that depend on whether the prior or subsequent node is an agency or a donor.

Although we choose to model FRP routes with segments primarily for expository purposes, it is a modeling choice that can impact problem size. However, problem size is only a concern for formulations that are solved using optimization software. In our work, there are only two such formulations: Formulation 4.9, the selective 1-PDTSP with stochastic supply, and Formulation 5.12, the LP solved repeatedly to obtain a solution to the 1-PDA-as.

Formulation 4.9, the selective 1-PDTSP with stochastic supply, models the route as a set of nodes without reference reference to segments. (A formulation with segments is not convenient for that model, since one component of its solution is the node sequence.)

Formulation 5.12 is modeled with segments. In principle, this leads to the creation of unnecessary decision variables when the route contains consecutive donors, because each donor is modeled individually (with a dummy agency for all but the last donor in the sequence) although a decision is only made between the final donor in the sequence and the following agency. However, if preprocessing were applied to consolidate consecutive donors, the segment-based model would not include these unnecessary decision variables. We have not implemented donor consolidation because the average solution time for the 1-PDA-as is modest. Furthermore, maintaining the separation between consecutive donors allows more flexibility to potentially adapt the model to a context in which a decision regarding collection must be made at each donor.

3.2. Modeling nodes

A FRP route comprises several types of nodes, each with different characteristics: the food bank warehouse (the depot), donors, and agencies.

The models of §4 and §5 model donors and agencies differently. In particular, additional parameters present in §5 allow for a richer model of collection and allocation decisions. In the present section, we describe the notation relevant to all models, making note of those used only in §5.

3.2.1. Food bank warehouse

The food bank warehouse has ample food available to be placed in the vehicle before starting the FRP route. This *initial load* may be important if the donors on the route are unlikely to provide enough food to satisfy the demand of the agencies; however, its presence reduces the capacity available to collect donations. We denote the initial load $\mathbf{S}_0$.

We denote the vehicle capacity $\mathbf{Q}$ when it is a decision variable (in §4) and Q when it is a parameter (in §5). Food banks operate a variety of vehicles that differ in capacity, ranging from tractor trailers to small trucks called "cube vans." It is preferable to use the smallest vehicle possible for FRP because smaller vehicles are less expensive to operate; they consume less fuel – a significant expense because the vehicles are refrigerated – and can be operated by a driver with a more common type of license. Furthermore, supporting FRP with smaller vehicles allows larger vehicles to be used for specialized tasks, such as picking up large donations from wholesalers or making scheduled deliveries to agencies located far from the food bank warehouse.

We assume that the initial load and the vehicle capacity are integers.

3.2.2. Donors

Donors supply a random amount of food. Food banks maintain extensive records of food donations from all sources. As such, the food bank has data that can be used to estimate the probability distribution of the quantity of food donated at each visit. Therefore, we model the quantity of food available at Donor i as a discrete independent bounded integer-valued non-negative random variable denoted D_i. The minimum and maximum donations are respectively denoted $d_i^{\min} = \min \operatorname{supp}(D_i)$ and $d_i^{\max} = \max \operatorname{supp}(D_i)$. (By $\operatorname{supp}(X)$ we refer to the minimal support of the random variable X.) We define two extreme donation scenarios: the *minimum scenario*, in which all donors supply their minimum donation (that is, $D_i = d_i^{\min}\ \forall i$), and the *maximum scenario*, in which all donors supply their maximum donation (that is, $D_i = d_i^{\max}\ \forall i$).

We assume that D_i is discrete and integer-valued because we have observed donation amounts measured as an integer number of boxes. Our assumption that the D_i are independent is based on conversations with FRP drivers and staff at several food banks.

The quantity of donated food collected at Donor i, also a random variable, is denoted $\mathbb{C}_i$. Clearly, $\mathbb{C}_i \preccurlyeq D_i$.

We assume that the driver collects as much of the donated food as the available vehicle capacity allows. We term this the *Maximum Collection Policy* because the driver collects the maximum possible quantity of food. This policy reflects our observations of FRP operations. The primary purpose of FRP is to collect donations. Collecting as much of the donation as possible also fulfills the expectations of donors. When any part of the donation is not collected, the donor must seek another entity that can take the remainder of the donation, a process that compromises the primary advantage of working with the food bank: having a single destination for all donations [70]. Alternatively, the donor could discard food that is not picked up. However, this causes the donor to lose the tax benefit of the refused donation, which may be an important incentive [1].

The load upon arrival to Donor i is a discrete random variable denoted $\mathbb{S}_i^D$. The relationship between the donation collected at Donor i and the vehicle capacity is expressed by:

$$\mathbb{S}_i^D + \mathbb{C}_i \leq Q \ \forall i \in \{1, \ldots, n\} \tag{3.2}$$

In §5, we model each donor with a *guaranteed collection* value denoted c_i. That is, upon arrival to Donor i, the vehicle must have sufficient capacity to accept a donation of c_i. Therefore, the Maximum Collection Policy is still applied, but with respect to a value that may be less than the vehicle capacity.

3.2.3. Agencies

The load upon arrival to Agency i is a discrete random variable denoted $\mathbb{S}_i^A$. In terms of the notation introduced to model donors, $\mathbb{S}_i^A = \mathbb{S}_i^D + D_i$.

Upon arrival to Agency i, food is removed from the vehicle and delivered to the agency. The amount of food is the *allocation*, which we denote a. We define three values at each agency that constrain the allocation decision: the minimum, maximum, and sustaining allocations.

- The *minimum allocation* is the least amount of food that may be allocated to the agency. The minimum allocation at Agency i is denoted $a_i^{\min}$. Agencies must recruit volunteers or requisition staff to unload, inspect, and store the food from FRP. Therefore, there must be a guarantee of receiving, at minimum, some substantial amount. Although the minimum allocation is sufficient to justify the agency's participation in FRP, it is not enough to provide a variety of perishable foods to the entire population served by the agency, or it is not enough to offer such a variety throughout the period between FRP visits.

- The *maximum allocation* is the greatest amount of food that may be allocated to the agency. The maximum allocation at Agency i is denoted $a_i^{\max}$. This value is determined by the storage capacity of the agency, or by the greatest amount of perishable food it can distribute between FRP visits.
- The *sustaining allocation* (used only in §5) is the least amount of food that the agency needs in order to offer a variety of perishable foods to the entire population it serves throughout the period between FRP visits. The sustaining allocation at Agency i is denoted a_i^{sust}. Ideally, every agency would receive at least its sustaining allocation every time it is visited, so that the population served by the agency would consistently have access to nutritious perishable foods. Due to the randomness of FRP donations, this may not be possible. Therefore, we introduce the parameter α_i, the minimum probability that Agency i receives at least its sustaining allocation.

We assume that $a_i^{\min}$, $a_i^{\max}$, and a_i^{sust} are integers. We describe the allocation decision in terms of two distinct but closely related concepts: the *allocation policy* and the *long-run allocation*. The allocation policy is a function that states how much food is allocated for a particular realization of $\mathbb{S}_i^A$, the load upon arrival to Agency i. The long-run allocation is a random variable that represents the quantity of food allocated over all possible realizations of $\mathbb{S}_i^A$.

Definition 3.2.1. *The* ***allocation policy*** *at Agency i, denoted $\mathbf{A}_i(s)\ \forall\ s \in \operatorname{supp}(\mathbb{S}_i^A)$, is a discrete integer-valued random variable which is a function on $\operatorname{supp}(\mathbb{S}_i^A)$. $\mathbf{A}_i(s)$ represents the allocation if the load is s upon arrival to Agency i.*

The range of the allocation policy at Agency i is a space of discrete random variables. For each $s \in \operatorname{supp}(\mathbb{S}_i^A)$, the allocation policy $\mathbf{A}_i(s)$ is a discrete random variable defined by a set of probabilities $\{p_0^s, p_1^s, ..., p^s_{\min\{a_i^{\max},s\}}\}$ such that

$$\Pr\{\mathbf{A}_i(s) = a|\mathbb{S}_i^A = s\} = p_a^s\ \forall\ a \in \{a_i^{\min}, ..., \min\{a_i^{\max}, s\}\}$$

and

$$\sum_{a=a_i^{\min}}^{\min\{a_i^{\max},s\}} p_a^s = 1.$$

If any of the $p_a^s = 1$, then $\mathbf{A}_i(s)$ is deterministic for the realization $\mathbb{S}_i^A = s$.

The long-run allocation at Agency i is defined in terms of the allocation policy:

Definition 3.2.2. *The* ***long-run allocation*** *at Agency i, denoted $\mathbb{A}_i$, is the amount allocated at Agency i in the long run. It is a discrete non-negative integer-valued bounded random variable. It is the outcome of applying $\mathbb{A}_i(\cdot)$ to each $s \in \operatorname{supp}(\mathbb{S}_i^A)$, with the probability $\Pr\{\mathbb{S}_i^A = s\}$. That is,*

$$\mathbb{A}_i = \mathbf{A}_i(s)\ w.p.\ \Pr\{\mathbb{S}_i^A = s\}, \tag{3.3}$$

or, equivalently,

$$\mathbb{A}_i = \mathbf{A}_i(\mathbb{S}_i^A). \tag{3.4}$$

The relationship between the long-run allocation at Agency i and the parameters of Agency i is described by a set of inequalities, one for each allocation value:

$$\mathbb{A}_i \geq a_i^{\min} \tag{3.5a}$$

$$\mathbb{A}_i \leq a_i^{\max} \tag{3.5b}$$

$$\Pr\{\mathbb{A}_i \geq a_i^{sust}\} \geq \alpha_i \tag{3.5c}$$

3.3. Summary values

We define three summary values with respect to a FRP solution:

C Expected total collection

A Expected total allocation

T Expected terminal load (food in the vehicle after allocation at Agency n)

In terms of the notation introduced in §3.2, the expected total collection and expected total allocation are:

$$C = \sum_{i=1}^{n} \mathbf{E}[\mathbb{C}_i] \tag{3.6a}$$

$$A = \sum_{i=1}^{n} \mathbf{E}[\mathbb{A}_i] \tag{3.6b}$$

The terminal load is the amount in the vehicle after all collections and allocations; therefore, it is computed in terms of the other summary values and $\mathbf{S}_0$ as:

$$T = \mathbf{S}_0 + C - A \tag{3.7}$$

Equivalently, Equation (3.7) can be expressed:

$$\mathbf{S}_0 + C = A + T \tag{3.8}$$

In Equation (3.8), the left side comprises all food that enters the vehicle (the initial load and collected donations) and the right side comprises all food that leaves the vehicle (allocations at agencies on the route and food stored in the warehouse at the end of the route).

CHAPTER 4

The selective 1-PDTSP with stochastic supply

Our formulation of the selective 1-PDTSP with stochastic supply is motivated by operations at Northern Illinois Food Bank (NIFB). NIFB serves thirteen counties in northeastern Illinois, comprising urban, suburban, and rural areas. The primary purpose of FRP at NIFB is to bring donations to a food bank warehouse. Agencies are included on FRP routes at NIFB to extend the capacity of the vehicle, allowing relatively smaller vehicles to be used. Participating in FRP is beneficial for agencies, especially those located far from a warehouse, because it provides frequent deliveries of perishable food. However, the food from FRP plays a primarily *supplemental* role for NIFB agencies; the bulk of their food is provided by scheduled deliveries from a food bank warehouse.

In §4.1, we define the *one-commodity pickup and delivery allocation problem* (1-PDA), which seeks the minimum capacity such that a given FRP route is survivable, as a stochastic program. It is not possible to solve the stochastic program directly, so we develop a simple three-step algorithm that obtains an optimal solution. In §4.2, we use a linear program reformulation of the 1-PDA as the basis for a mixed-integer linear program formulation of the *selective 1-PDTSP with stochastic supply.* We also provide analytical results regarding solutions to the problem. In §4.3, we propose the capacity reuse insertion heuristic (CRIH), a simple heuristic that can be applied by food banks to insert agencies into existing FRP routes. In §4.4, we present a case study based on NIFB data to provide managerial insights about FRP operations and to evaluate CRIH.

4.1. Allocation decisions for a fixed route

In §4.1.1, we formulate the 1-PDA as a stochastic program. The solution to the 1-PDA has three components: the initial load, the allocation policy at each agency, and the vehicle capacity. In §4.1.2, we develop an exact solution algorithm for the 1-PDA.

4.1.1. Stochastic program formulation of the 1-PDA

The parameters and decision variables used to formulate the 1-PDA are defined in Table 4.1.

Table 4.1. Parameters and decision variables of the 1-PDA

Parameters	
n	number of segments, indexed by $i \in \{0, 1, \ldots, n\}$; "0" represents the depot
D_i	donation from Donor i, a discrete random variable
$a_i^{\min}$	minimum allocation at Agency i
$a_i^{\max}$	demand at Agency i; the maximum allocation at Agency i
Decision variables	
$\mathbf{Q}$	capacity of the vehicle, a non-negative integer
$\mathbf{S}_0$	initial load, a non-negative integer
$\mathbf{A}_i(\cdot)$	allocation policy at Agency i; defined for all $s \in \mathrm{supp}(\mathbb{S}_i^A)$
$\mathbb{S}_i^D$	load upon arrival to Donor i, a discrete integer-valued random variable
$\mathbb{S}_i^A$	load upon arrival to Agency i, a discrete integer-valued random variable

Figure 4.1 depicts the relationship among the variables associated with Segment i.

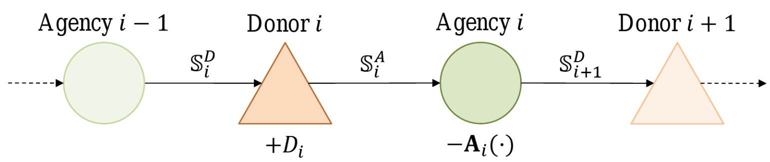

Figure 4.1 Segment i with corresponding variables

$$\text{Minimize} \quad \mathbf{Q} \tag{4.1a}$$

$$\text{subject to} \quad \mathbb{S}_i^A = \mathbb{S}_i^D + D_i \qquad i \in \{1, \ldots, n\} \tag{4.1b}$$

$$\mathbb{S}_{i+1}^D = \mathbb{S}_i^A - \mathbf{A}_i(\mathbb{S}_i^A) \qquad i \in \{1, \ldots, n\} \tag{4.1c}$$

$$\mathbf{A}_i(s) \leq s \qquad \forall s \in \text{supp}(\mathbb{S}_i^A), i \in \{1, \ldots, n\} \tag{4.1d}$$

$$a_i^{\min} \leq \mathbf{A}_i(s) \leq a_i^{\max} \qquad \forall s \in \text{supp}(\mathbb{S}_i^A), i \in \{1, \ldots, n\} \tag{4.1e}$$

$$\mathbb{S}_i^D + D_i \leq \mathbf{Q} \qquad i \in \{1, \ldots, n\} \tag{4.1f}$$

$$\mathbb{S}_1^D = \mathbf{S}_0 \tag{4.1g}$$

$$\mathbf{S}_0 \in \mathbb{Z}_0^+ \tag{4.1h}$$

$$\mathbf{Q} \in \mathbb{Z}^+ \tag{4.1i}$$

The objective function (4.1a) minimizes the required capacity.

Constraints (4.1b) and (4.1c) describe the movement of food in and out of the vehicle. Constraints (4.1b) calculate the load after visiting a donor in terms of the load upon arrival and the donation collected. Constraints (4.1c) express the load after visiting an agency in terms of the load upon arrival and the allocation policy at the agency.

Constraints (4.1d) and (4.1e) restrict the allocation policy $\mathbf{A}_i(\cdot)$. Constraints (4.1d) limit allocations to be no more than the load. Constraints (4.1e) ensure that allocations be no less than the minimum allocation and no more than the maximum allocation.

Constraints (4.1f) ensure that the capacity suffices to collect the entire donation at each donor. Constraint (4.1g) defines $\mathbb{S}_1^D$ as a random variable equal to the initial load $\mathbf{S}_0$ with probability 1. Constraint (4.1h) declares the initial load $\mathbf{S}_0$ to be a non-negative integer. Constraint (4.1i) declares the capacity $\mathbf{Q}$ to be a positive integer.

Solving the 1-PDA to optimality is not straightforward. Approached directly as a stochastic program, the problem is intractable for any but the most trivial instances. In the following section, we develop a simple and efficient three-step algorithm to solve the problem to optimality. We utilize key concepts that underlie the algorithm to formulate the selective 1-PDTSP with stochastic supply in §4.2. (It is also possible to solve the 1-PDA with dynamic programming (DP), as described in Appendix B.)

4.1.2. Solution approach for the 1-PDA

The allocation policy for FRP must address two conflicting constraints. In order to collect donations, capacity must be freed by allocating to agencies. However, if too much is allocated at an agency and donations later in the route are low, the load may not satisfy demand at subsequent agencies. We develop an exact solution method that balances these two constraints by determining the load required at every segment for the minimum scenario, and which therefore ensures survivability over all scenarios.

Minimum intermediate load. At each segment, there must exist a minimum load that ensures all subsequent agencies receive at least the minimum allocation, even if only the minimum donation is collected at every subsequent donor. We formalize this idea as the *minimum intermediate load*:

Definition 4.1.1. *The* ***minimum intermediate load*** l_i^* *at Segment* i *is the minimum load required upon arrival to Segment* i *to provide at least the minimum allocation to Agency* i *and all subsequent agencies.*

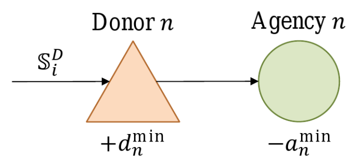

Figure 4.2 Final segment under the minimum scenario

Consider the final segment of the route under the minimum scenario, depicted in Figure 4.2. If the minimum donation is less than the minimum allocation (that is, if $d_n^{\min} < a_n^{\min}$), then to satisfy the minimum allocation requirement at the agency, the load must be at least $a_n^{\min} - d_n^{\min}$ upon arrival to Donor n. We isolate and generalize this value as the *supply gap*:

Definition 4.1.2. *The* ***supply gap*** g_i *is the amount by which the minimum supply available in Segment* i *fails to meet the minimum demand in Segment* i*. It is the difference between the minimum collection at Donor* i *and the minimum allocation at Agency* i*:*

$$g_i = a_i^{\min} - d_i^{\min} \tag{4.2}$$

Figure 4.2 only depicts the final segment of the route, Segment n. At Segment $n-1$, the situation is similar, but the requirements of Segment n must also be considered. That is, upon arrival to Donor $n-1$, the load must be enough to overcome a potential shortfall in donations at Donor $n-1$ and, after allocation at Agency $n-1$, the load must be at least g_n in order to serve Agency n. In general, the minimum intermediate load of Segment i can be computed in terms of the supply gap at Segment i (g_i) and the minimum intermediate load at Segment $i+1$ (l^*_{i+1}):

$$l^*_i = (l^*_{i+1} + g_i)^+ \tag{4.3}$$

Since allocation policies are indexed by segments, we have defined the minimum intermediate load in terms of segments; however, it could be defined in terms of individual nodes. We use the node-based approach to formulate the selective 1-PDTSP with stochastic supply in §4.2.

MILB algorithm. Since our solution algorithm for the 1-PDA is based on the concept of the minimum intermediate load, we deem it the *minimum intermediate load-based* (MILB) algorithm.

A solution to the 1-PDA comprises the initial load $\mathbf{S}_0$, the allocation policy $\mathbf{A}_i(\cdot)\ \forall i \in \{1,\ldots,n\}$, and the vehicle capacity $\mathbf{Q}$. Each of the three steps of the MILB algorithm obtains one of the three parts of a 1-PDA solution.

Step 1: Calculate minimum intermediate loads and obtain initial load. After visiting Agency n, no further allocations are made, so the minimum intermediate load after the final segment is $l^*_{n+1} = 0$. Therefore, starting from Segment n, the minimum intermediate load for every prior segment is calculated recursively by applying Equation (4.3).

Any feasible initial load $\mathbf{S}_0$ must be at least l_1^*. Although an initial load greater than l_1^* would be feasible, capacity is minimized by choosing the minimum feasible initial load. We denote this value $\mathbf{S}_0^{l^*}$ to emphasize its relationship to the MILB algorithm: $\mathbf{S}_0^{l^*} = l_1^*$.

Step 2: Define allocation policy. We define the minimum intermediate load-based (MILB) allocation policy $\mathbf{A}_i^{l^*}(\cdot)$ in terms of l_i^*:

$$\mathbf{A}_i^{l^*}(s) = \min\{s - l_{i+1}^*, a_i^{\max}\} \; \forall \; s \in \operatorname{supp}(\mathbb{S}_i^A), \; \forall \; i \in \{1, \ldots, n\} \tag{4.4}$$

Lemma 4.1.3. *Given a scenario* $(d_1, d_2, \ldots, d_n)$ *and initial load* $\mathbf{S}_0$*, no allocation policy allocates more in total than the MILB allocation policy* $\mathbf{A}_i^l(\cdot) \; \forall \; i \in \{1, \ldots, n\}$.

A formal proof of Lemma 4.1.3 (and all analytical results) is provided in Appendix A, but the result can be demonstrated intuitively: At Agency i, the MILB allocation policy $\mathbf{A}_i^{l^*}(\cdot)$ allocates $a_i^{\max}$, unless doing so would reduce the load to less than the minimum intermediate load of the next segment, in which case $\mathbf{A}_i^{l^*}(\cdot)$ allocates $s - l_{i+1}^*$. For a given load s, it is not possible to define a feasible policy that allocates more food to the agency: allocating more than $a_i^{\max}$ would violate Constraint (4.1e) of the 1-PDA, while allocating more than $s - l_{i+1}^*$ would cause the load to be less than the minimum intermediate load of the next segment.

Step 3: Calculate maximum intermediate loads and capacity. The *MILB maximum intermediate load* $u_i^{l^*}$ is the maximum possible load upon arrival to Segment i if the initial load is $\mathbf{S}_0^{l^*} = l_1^*$ and the allocation policy is the MILB allocation policy $\mathbf{A}_i^{l^*}(\cdot)$. Note that $u_i^{l^*}$ is *not* an upper bound on the load for every feasible 1-PDA solution; it is only relevant to the MILB solution.

For Segment 1, $u_1^{l^*} = \mathbf{S}_0^{l^*} = l_1^*$. The MILB maximum intermediate load for every subsequent segment is then calculated recursively:

$$u_{i+1}^{l^*} = \max\left\{u_i^{l^*} + d_i^{\max} - a_i^{\max}, l_{i+1}^*\right\} \tag{4.5}$$

The structure of Equation (4.5) results from applying the MILB allocation policy $\mathbf{A}_i^{l^*}(\cdot)$, Equation (4.4). The load upon arrival to Segment $i+1$ is maximized when the maximum load is present upon arrival to the previous segment ($u_i^{l^*}$) and the maximum donation ($d_i^{\max}$) is collected at Donor i. The MILB allocation policy allocates $a_i^{\max}$, unless doing so would cause the load to decrease below the minimum intermediate load of the next segment. Therefore, the load after allocating at Agency i is the maximum of $u_i^{l^*} + d_i^{\max} - a_i^{\max}$ and l_{i+1}^*.

We calculate the MILB maximum intermediate loads as a step to obtaining the capacity of the MILB solution, $\mathbf{Q}^{l^*}$. The maximum load over the course of the route results when the sum of $u_i^{l^*}$ and the donation collected at Donor i is greatest; that is:

$$\mathbf{Q}^{l^*} = \max_i \left(u_i^{l^*} + d_i^{\max}\right) \tag{4.6}$$

Theorem 4.1.4. *The MILB solution (consisting of initial load $\mathbf{S}_0^{l^*}$, allocation policy $\mathbf{A}_i^{l^*}(\cdot)\ \forall\ i \in \{1, \ldots, n\}$, and capacity $\mathbf{Q}^{l^*}$) is an optimal solution to the 1-PDA.*

Intuitively, the MILB solution is an optimal solution to the 1-PDA because no allocation policy could free more capacity to accept donations: the vehicle begins the route with the minimum feasible load, and at each agency, no greater amount can be allocated while maintaining feasibility.

4.2. Formulation of the selective 1-PDTSP with stochastic supply

In §4.2.1, we introduce the implicit 1-PDA, an alternate formulation of the 1-PDA that does not explicitly model the allocation decision. In §4.2.2, we augment the implicit 1-PDA with agency selection and node sequencing decisions to formulate the selective 1-PDTSP with stochastic supply.

4.2.1. Implicit 1-PDA

To solve the 1-PDA, we need not model the allocation decision explicitly. Rather, we demonstrate that it suffices to model the range of possible vehicle load values through the *implicit 1-PDA*. The minimum load at each segment is the only input (other than parameters) required by the MILB algorithm, which provides an optimal solution, so an optimal solution in terms of the range of load values implies an optimal solution to the 1-PDA. The implicit 1-PDA is not intrinsically useful, but since it can be stated as a linear program, it is a crucial intermediate step in the formulation of the selective 1-PDTSP with stochastic supply.

We have demonstrated the primacy of the extreme scenarios in solving the 1-PDA. The minimum scenario determines the minimum intermediate load at each segment and, through this, the minimum feasible initial load; the maximum scenario and the minimum intermediate loads determine the capacity. Therefore, we define two sets of variables which represent the load under the extreme scenarios: the l_i and u_i respectively represent the minimum and maximum load upon arrival to Segment i. By describing the relationship among these sets of variables, we formulate the implicit 1-PDA as a linear program with $2n + 4$ variables and $5n + 3$ constraints.

$$
\begin{aligned}
&\text{Minimize} \quad && \mathbf{Q} && && \text{(4.7a)}\\
&\text{subject to} \quad && l_i \geq 0 && i \in \{1, \ldots, n+1\} && \text{(4.7b)}\\
& && l_{i-1} \geq l_i - d_{i-1}^{\min} + a_{i-1}^{\min} && i \in \{2, \ldots, n+1\} && \text{(4.7c)}\\
& && u_i \geq u_{i-1} + d_{i-1}^{\max} - a_{i-1}^{\max} && i \in \{2, \ldots, n+1\} && \text{(4.7d)}\\
& && u_i \geq l_i && i \in \{1, \ldots, n+1\} && \text{(4.7e)}\\
& && \mathbf{Q} \geq u_i + d_i^{\max} && i \in \{1, \ldots, n\} && \text{(4.7f)}\\
& && \mathbf{S}_0 = l_1 && && \text{(4.7g)}
\end{aligned}
$$

The objective (4.7a) is the minimization of vehicle capacity. Constraints (4.7b) constrain the minimum load to be nonnegative. Constraints (4.7c) ensure that the minimum load suffices for the minimum allocation. Constraints (4.7d) constrain the maximum load in relation to the maximum donations and allocations. Constraints (4.7e) ensure that the maximum load be greater than the minimum load. Constraints (4.7f) relate the capacity to the maximum load under the maximum scenario. Constraint (4.7g) defines the initial load.

Since Formulation 4.7 minimizes capacity, to demonstrate that it provides an optimal solution to the 1-PDA, we need only demonstrate that its solution implies a feasible allocation policy. This is addressed by Theorem 4.2.1.

Theorem 4.2.1. *Given a solution to the implicit 1-PDA, the allocation policy:*

$$\mathbf{A}_i^l(s) = \min\{s - l_{i+1}, a_i^{\max}\} \ \forall\ s \in \operatorname{supp}(\mathbb{S}_i^A),\ \ \forall\ i \in \{1, \ldots, n\} \tag{4.8}$$

is feasible.

4.2.2. MILP formulation of the selective 1-PDTSP with stochastic supply

Louveaux and Salazar-González [38] formulate the 1-PDTSP with stochastic demands as a MILP in terms of the minimum and maximum scenarios. We use the l_i and u_i of the implicit 1-PDA to formulate our problem in a similar way. The parameters and decision variables of the selective 1-PDTSP with stochastic supply are summarized in Tables 4.2 and 4.3:

Table 4.2. Decision variables of the selective 1-PDTSP with stochastic supply

Decision variables	
$\mathbf{Q}$	capacity of the vehicle
Y_{jk}	binary variable indicating whether node j is assigned to position k
L_k	amount of food sufficient to provide the minimum allocation to agencies in positions beyond position k
U_k	maximum possible load upon departure from the node in position k
$X_{j_1j_2}$	binary variable indicating whether the edge from j_1 to j_2 is used in the solution

Table 4.3. Parameters of the selective 1-PDTSP with stochastic supply

Parameters	
$\mathcal{D}$	set of donors, all of which must be visited
$\mathcal{A}$	set of agencies, from which a (possibly empty) subset is chosen for inclusion
$\mathcal{N}$	set of all donor and agency nodes, indexed by $j \in \mathcal{N}$; $\mathcal{N} = \mathcal{D} \cup \mathcal{A}$
$\mathcal{N}^0$	set of all donor and agency nodes *and* the depot (denoted by "0"); $\mathcal{N}^0 = \mathcal{N} \cup \{0\}$
p	number of positions to which the nodes in $\mathcal{N}^0$ are assigned, indexed by $k \in \{0, \dots, p\}$; $p = \lvert\mathcal{N}\rvert$
$\tilde{l}_j$	least amount by which the load can change due to visiting node j: if node j is a donor, $\tilde{l}_j$ is its minimum donation, $d^{\min}$; if node j is an agency, $\tilde{l}_j$ is the negative of its minimum allocation, $-a^{\min}$; at the depot, $\tilde{l}_0 = 0$
$\tilde{u}_j$	greatest amount by which the load can change due to visiting node j: if node j is a donor, $\tilde{u}_j$ is its maximum donation, $d^{\max}$; if node j is an agency, $\tilde{u}_j$ is the negative of its maximum allocation, $-a^{\max}$; at the depot, $\tilde{u}_0 = 0$
$t_{j_1 j_2}$	travel time of the edge from node j_1 to node j_2; should be interpreted as the sum of travel time between the nodes and the time required for collection or allocation at j_2
T	upper bound on total travel time

Since the route is not fixed a priori, we cannot use segment-based indexing. Instead, we consider a set of donors $\mathcal{D}$, the elements of which are denoted $D^1, D^2, \dots, D^{|\mathcal{D}|}$ and indexed by $D \in \mathcal{D}$, as well as a set of agencies $\mathcal{A}$, the elements of which are denoted $A^1, A^2, \dots, A^{|\mathcal{A}|}$ and indexed by $A \in \mathcal{A}$. The set of all nodes (other than the depot) is

$\mathcal{N} = \mathcal{D} \cup \mathcal{A}$; the set of all nodes including the depot is $\mathcal{N}^0 = \mathcal{N} \cup \{0\}$. The symbols $d^{\min}$, $d^{\max}$, $a^{\min}$, and $a^{\max}$ are associated with nodes instead of segments.

Let the binary decision variable Y_{jk} equal 1 if node j is assigned to position $k \in \{1, \ldots, p\}$, where $p = |\mathcal{N}|$. Since including agencies is optional, the route may not use all p positions. In this case, unused positions at the end of the route are assigned to the depot. Two parameters, denoted $\tilde{l}_j$ and $\tilde{u}_j$, are associated with each node, respectively representing the least and greatest amounts by which the load can change at node j. If node j is a donor, $\tilde{l}_j = d_j^{\min}$ and $\tilde{u}_j = d_j^{\max}$. If node j is an agency, $\tilde{l}_j = -a_j^{\min}$, and $\tilde{u}_j = -a_j^{\max}$. For the depot, $\tilde{l}_0 = \tilde{u}_0 = 0$.

In the implicit 1-PDA, l_i represents the minimum load upon arrival to Segment i. In the selective 1-PDTSP with stochastic supply, the variable L_k represents a load that ensures survivability for the nodes in positions beyond k. The variable U_k represents the maximum load upon departure from the node in position k.

An optimal solution to the 1-PDTSP with stochastic supply minimizes capacity subject to the constraint that the total travel time not exceed the parameter T. Let the binary decision variable $X_{j_1 j_2}$ equal 1 if edge (j_1, j_2) is used in the route. Let $t_{j_1 j_2}$ denote the travel time of edge (j_1, j_2), which includes the time required for collection or allocation at j_2.

$$
\begin{aligned}
\text{Minimize} \quad & \mathbf{Q} && && \text{(4.9a)}\\
\text{subject to} \quad & \sum_{k=1}^{p} Y_{jk} = 1 && j \in \mathcal{D} && \text{(4.9b)}\\
& \sum_{k=1}^{p} Y_{jk} \leq 1 && j \in \mathcal{A} && \text{(4.9c)}\\
& \sum_{j \in \mathcal{N}^0} Y_{jk} = 1 && k \in \{1, \ldots, p\} && \text{(4.9d)}\\
& Y_{0k} \leq Y_{0(k+1)} && k \in \{1, \ldots, p-1\} && \text{(4.9e)}\\
& L_{k-1} \geq L_k - \sum_{j \in \mathcal{N}} Y_{jk} \tilde{l}_j && k \in \{1, \ldots, p\} && \text{(4.9f)}\\
& U_k \geq U_{k-1} + \sum_{j \in \mathcal{N}} Y_{jk} \tilde{u}_j && k \in \{1, \ldots, p\} && \text{(4.9g)}\\
& 0 \leq L_k \leq U_k \leq \mathbf{Q} && k \in \{0, \ldots, p\} && \text{(4.9h)}\\
& X_{0j} \geq Y_{j1} && j \in \mathcal{N} && \text{(4.9i)}\\
& X_{j0} \geq Y_{jp} && j \in \mathcal{N} && \text{(4.9j)}\\
& X_{j_1 j_2} \geq Y_{j_1 k} + Y_{j_2 (k+1)} - 1 && k \in \{1, \ldots, p-1\}; j_1 \in \mathcal{N}, j_2 \in \mathcal{N}^0 : j_1 \neq j_2 && \text{(4.9k)}\\
& \sum_{j_1 \in \mathcal{N}^0} \sum_{j_2 \in \mathcal{N}^0 : j_1 \neq j_2} t_{j_1 j_2} X_{j_1 j_2} \leq T && && \text{(4.9l)}\\
& Y_{jk} \in \{0, 1\} && j \in \mathcal{N}^0, k \in \{1, \ldots, p\} && \text{(4.9m)}\\
& X_{j_1 j_2} \in \{0, 1\} && j_1 \in \mathcal{N}^0, j_2 \in \mathcal{N}^0 : j_1 \neq j_2 && \text{(4.9n)}
\end{aligned}
$$

Objective (4.9a) minimizes the vehicle capacity.

Constraints (4.9b) through (4.9e) govern the assignment of nodes to positions in the route. Constraints (4.9b) require each donor to be assigned to exactly one position. Constraints (4.9c) require each agency be assigned to at most one position. Constraints (4.9d) require each position in the route to be assigned to exactly one node, which may be a donor, an agency, or the depot. Constraints (4.9e) ensure that, once the vehicle has returned to the depot, it persists at the depot. Persistence at the depot represents the end of the route.

Constraints (4.9f) calculate an upper bound on the minimum intermediate load at each position. Constraints (4.9f) adapt Constraints (4.7c) of the implicit 1-PDA to the node-based structure of the selective 1-PDTSP with stochastic supply. If the node in position k is a donor, the constraint is:

$$L_{k-1} \geq L_k - d^{\min} \tag{4.10}$$

That is, the minimum intermediate load at position $k-1$ can be up to $d^{\min}$ less than the minimum intermediate load at position k because at least $d^{\min}$ is collected at position k. If the node in position k is an agency, the constraint is:

$$L_{k-1} \geq L_k + a^{\min} \tag{4.11}$$

That is, the minimum intermediate load at position $k-1$ must be at least $a^{\min}$ greater than the minimum intermediate load at position k because at least $a^{\min}$ must be allocated to the agency at position k.

Constraints (4.9g) calculate an upper bound on the maximum load at each position. Constraints (4.9g) adapt Constraints (4.7d) of the implicit 1-PDA to the node-based structure of the selective 1-PDTSP with stochastic supply. If the node in position k is a donor, the constraint is:

$$U_k \geq U_{k-1} + d^{\max} \tag{4.12}$$

That is, the maximum possible load at position k must be at least $d^{\max}$ greater than at position $k-1$. If the node in position k is an agency, the constraint is:

$$U_k \geq U_{k-1} - a^{\max} \tag{4.13}$$

That is, the maximum possible load can decrease no more than $a^{\max}$ at position k, because the agency cannot accept more.

Constraints (4.9h) ensure that, when allocating at an agency, U_k does not decrease below the minimum intermediate load L_k. Furthermore, they ensure that the minimum intermediate load is non-negative. Constraints (4.9h) also establish the relationship between the U_k and the capacity $\mathbf{Q}$.

Constraints (4.9i), (4.9j), and (4.9k) relate position assignments to edges. Constraints (4.9i) determine which edge is used to leave the depot by identifying the node assigned to in the first position. Constraints (4.9j) relate the edge used to return to the depot with the node in the final position of the node sequence. Constraints (4.9k) determine all other edges used in the node sequence by examining which nodes are visited in consecutive positions. Note that, if not all agencies are included, Constraints (4.9k) determine which edge is used to return to the depot instead of Constraints (4.9j).

Constraint (4.9l) restricts the total travel time of the node sequence to be no more than T.

Constraints (4.9m) and (4.9n) declare the Y_{jk} and the $X_{j_1j_2}$, respectively, as binary decision variables.

4.2.3. Analytical results

In this subsection, we present analytical results regarding optimal solutions to the selective 1-PDTSP with stochastic supply. We use these properties in the case study in §4.4. Theorem 4.2.2 provides a guideline regarding the insertion of agencies.

Theorem 4.2.2. *Inserting an agency adjacent to the depot cannot decrease the minimum capacity.*

Inserting an agency is beneficial only if it allows capacity to be reused under the maximum scenario. That is, under the maximum scenario, if visiting an agency allows donations that have already been collected to be removed from the vehicle *and* if the freed capacity is occupied by donations collected later in the route, then the presence of the agency on the route reduces the required capacity by as much as the amount of capacity reused.

Agencies adjacent to the depot cannot allow capacity to be reused because they are visited either before or after all donations are collected. Therefore, we can eliminate from consideration solutions in which an agency is the first or last node visited. Due to Theorem 4.2.2, we do not consider agency insertions at the beginning or end of FRP routes in the insertion heuristics presented in §4.3 and Appendix D.

With Theorem 4.2.3, we obtain a lower bound on the optimal vehicle capacity for a given set of donors and agencies. Since agencies reduce the required capacity through reuse, adding an agency cannot reduce the required capacity by more than its maximum allocation.

Theorem 4.2.3. *For any route of the set of donors $\mathcal{D}$ and the set of agencies $\mathcal{A}$, a bound on the minimum required capacity is:*

$$\mathbf{Q}^* \geq \sum_{D \in \mathcal{D}} d_D^{\max} - \sum_{A \in \mathcal{A}} a_A^{\max} \tag{4.14}$$

For a particular instance of the selective 1-PDTSP with stochastic supply, the restriction on maximum total travel time, Constraint (4.9l), may limit the number of agencies that can be inserted in a route. In that case, Corollary 4.2.4 provides a tighter bound.

Corollary 4.2.4. *If at most ρ agencies of the set $\mathcal{A}$ may be inserted into an FRP route of the set of donors $\mathcal{D}$, let $\mathcal{A}^\rho \subset \mathcal{A}$ represent the ρ agencies with the highest values of $a_j^{\max}$. Then,*

$$\mathbf{Q}^* \geq \mathbf{Q}^{lb}(\rho) = \sum_{D \in \mathcal{D}} d_D^{\max} - \sum_{j \in \mathcal{A}^\rho} a_j^{\max} \tag{4.15}$$

We refer to $\mathbf{Q}^{lb}(\rho)$ as the *ρ-agency lower bound on capacity.* In the case study, we demonstrate that, for realistic instances, the ρ-agency lower bound is tight. Thus, we can use it to estimate the benefit of inserting agencies on a FRP route and to identify instances in which the heuristic solution may be poor.

4.3. Heuristic method for inserting agencies into existing routes

We now consider the restricted problem of inserting agencies into existing FRP routes. This problem interests us for two reasons:

First, food banks and donors are accustomed to the routes that already exist. The food bank would prefer to continue to visit donors at approximately the same time as in the current schedule, so it would be preferable to add agencies to the route without changing the order in which donors are visited.

Second, although it is possible to solve instances of the selective 1-PDTSP with stochastic supply with a MILP solver, food bank staff do not have optimization tools available, nor do they have experience using them. Even if a solver were available, many realistic instances cannot be solved in a reasonable amount of time (see §4.4.4).

We develop a simple heuristic method for inserting agencies into existing FRP routes that requires no use of technology and is not sensitive to problem size. Our heuristic approach is based on the intuition that including agencies permits the capacity reduction by allowing capacity reuse, so we call it the *capacity reuse insertion heuristic* (CRIH). In the computational study in §4.4.4, we evaluate the "cost" of not solving the full problem in terms of the difference in required capacity between CRIH and optimal solutions.

In §4.3.1, we define the *estimated reuse contribution of an agency*, the value that is used in the CRIH to rank potential insertions. We formally present the steps of the CRIH in §4.3.2. CRIH is a myopic heuristic that inserts one agency at a time into the existing route. It does not reoptimize the route, nor does it consider potential future insertions when it chooses an insertion.

4.3.1. Estimated reuse contribution of an agency

For the inclusion of an agency to reduce capacity, donations must be collected before *and* after visiting the agency. In the present subsection, we first consider how to quantify agencies' contribution to capacity reduction when the entire route is given. Then, we develop an expression to estimate an agency's contribution before it is inserted into a route.

The amount of capacity reuse due to an included agency is limited not only by the maximum allocation at the agency, but also by the minimum intermediate load of the next segment. Since calculating the minimum intermediate loads and the initial load is cumbersome analytically (although easy numerically), we introduce the *Zero Minimum Intermediate Load (ZMIL) condition*.

Definition 4.3.1. *The* ***Zero Minimum Intermediate Load (ZMIL) condition*** *is satisfied for a route if the minimum intermediate load is zero (that is,* $l_i^* = 0$*) at all segments.*

The ZMIL condition is obviously satisfied if $a^{\min} = 0$ for the agency being inserted. Less restrictively, it is also fulfilled if the combined minimum donation of the donors immediately preceding the agency is at least $a^{\min}$. Even if the ZMIL condition is not satisfied, for realistic instances of the selective 1-PDTSP with stochastic supply, the minimum intermediate load is often near zero because parameter values for donors are generally much greater than the corresponding values for agencies (as demonstrated in §4.4.3).

For a given route, the *maximum reuse contribution of a set of consecutive agencies*, denoted $R_{\mathcal{A}_{i,i+k}}$ for the set of consecutive agencies $\mathcal{A}_{i,i+k} = \{\text{Agency } i, \text{Agency } i +$

$1, \ldots,$ Agency $i+k\}$ represents the maximum possible contribution of a set of consecutive agencies to the reduction of $\mathbf{Q}$ by capacity reuse. (Since the agencies are consecutive, it must be that Donors $i+1, \ldots, i+k$ are dummy donors.) The maximum reuse contribution of a set of consecutive agencies can be no more than the sum of their maximum allocations $\sum_{\iota=i}^{i+k} a_\iota^{\max}$. However, if (under the maximum scenario) fewer than $\sum_{\iota=i}^{i+k} a_\iota^{\max}$ units of food are collected before Agency i or after Agency $i+k$, the maximum reuse contribution of the set of consecutive agencies is less.

Definition 4.3.2. *If the ZMIL condition is satisfied, then the* ***maximum reuse contribution*** *of the set of consecutive agencies $\mathcal{A}_{i,i+k}$ is defined as the maximum amount by which the set can decrease the required capacity by allowing capacity reuse):*

$$R_{\mathcal{A}_{i,i+k}} = \min\left\{\sum_{\iota=1}^{i+k} a_\iota^{\max}, \sum_{\iota=1}^{i} d_\iota^{\max} - \sum_{\iota=1}^{i-1} \mathbf{A}_\iota^{\max}, \sum_{\iota=i+k+1}^{n} d_\iota^{\max} - \sum_{\iota=i+k+1}^{n} \mathbf{A}_\iota^{\max}\right\} \tag{4.16}$$

where $\mathbf{A}_\iota^{\max}$ represents the allocation at Agency ι under the maximum scenario.

Since the maximum reuse contribution of an agency can be no more than the maximum allocation, if $R_{\mathcal{A}_{i,i+k}} = \sum_{\iota=i}^{i+k} a_\iota^{\max}$ for all sets of consecutive agencies on the route, then no other node sequence allows for a further reduction in capacity.

Theorem 4.3.3. *For a given route of donors and agencies, if the ZMIL condition is satisfied and $R_{\mathcal{A}_{i,i+k}} = \sum_{\iota=i}^{i+k} a_\iota^{\max}$ for all sets of consecutive agencies, then the node sequence in the route obtains the minimum vehicle capacity $\mathbf{Q}^*$ possible for that set of nodes.*

Theorem 4.3.3 motivates us to seek an insertion method that maximizes the reuse contribution of agencies. However, the definition of the maximum reuse contribution assumes that the entire route is known, so using $R_{\mathcal{A}_{i,i+k}}$ directly as the basis of a heuristic would require that all agencies be inserted simultaneously. To myopically insert agencies one at a time, the CRIH makes decisions using a estimated version of $R_{\mathcal{A}_{i,i+k}}$ for individual agencies. The metric used for the CRIH is an estimate because it assumes that the ZMIL condition is satisfied and that $\mathbf{A}_i^{\max} = a_i^{\max}$ (that is, that the allocation at Agency i under the maximum scenario is its maximum allocation).

Definition 4.3.4. *For a given route of donors and agencies, the **estimated reuse contribution** of agency A if inserted on edge e, denoted $\tilde{R}_{Ae}$, is the approximate amount by which the presence of A would decrease the required capacity if it were inserted on edge e:*

$$\tilde{R}_{Ae} = \min\left\{ a_A^{\max}, \sum_{\iota=1}^{i^-(\mathcal{A}_e)} d_\iota^{\max} - \sum_{\iota=1}^{i^+(\mathcal{A}_e)-1} a_\iota^{\max}, \sum_{\iota=i^+(\mathcal{A}_e)}^{n} d_\iota^{\max} - \sum_{\iota=i^-(\mathcal{A}_e)}^{n} a_\iota^{\max} \right\} \tag{4.17}$$

Where $\mathcal{A}_e$ represents the set of consecutive agencies adjacent to edge e, $i^-(\mathcal{A}_e)$ is the index of the last donor in the route before $\mathcal{A}_e$, and $i^+(\mathcal{A}_e)$ is the is the index of the first donor in the route after $\mathcal{A}_e$. If neither endpoint of e is an agency, then $\mathcal{A}_e = \emptyset$, $i^-(\mathcal{A}_e)$ is the index of the tail of e, and $i^+(\mathcal{A}_e)$ is the index of the head of e.

4.3.2. CRIH algorithm

A feasible route including all donors is provided to the CRIH as an input. The heuristic repeatedly identifies the agency insertions with the greatest value of $\tilde{R}_{Ae}$, then myopically applies the one with the lowest insertion cost. The values of $\tilde{R}_{Ae}$ for agencies that have not yet been inserted are updated after each insertion, since the ability of an agency to allow capacity reuse depends on the agencies before and after it. The heuristic continues until no feasible insertions with positive $\tilde{R}_{Ae}$ values remain. We summarize the steps of the CRIH below.

Step 1: *Update:* Calculate C, the total travel time of the current route.

Step 2: *Identify best insertion point for each agency:* For every agency k not on the current route, find edge e that maximizes $\tilde{R}_{ke}$ such that the insertion cost $c_{ke} \leq T - C$; denote it e_k^*. If there are multiple such edges, let e_k^* be one with minimum c_{ke} and go to Step 3. If there are no such edges for any agency, return the current route as the heuristic solution.

Step 3: *Identify best agency to insert:* Let $k^* = \text{argmax}_k \tilde{R}_{ke_k^*}$. Insert agency k^* on edge $e_{k^*}^*$. If there are multiple such agencies, insert one with minimum insertion cost $c_{k^*e_{k^*}^*}$. Go to Step 1.

4.4. NIFB case study

To quantify the benefits of including agencies on FRP routes and to provide implementation guidance, we present a case study based on data provided by NIFB. We use optimal solutions for more than 3,000 instances of the selective 1-PDTSP with stochastic supply to obtain insights regarding the insertion of agencies into FRP routes and evaluate the performance of CRIH.

4.4.1. NIFB operations

Our experiment is based on data provided by NIFB for donations received and food delivered to agencies from April 2010 to April 2011. From the donations data, we obtain minimum and maximum donations for each donor. For each agency, we first determine if the agency can accept weekly donations from FRP (due to its service schedule and ability to store perishable food). Of NIFB's more than 600 partner agencies, only 139 can potentially be included on FRP routes. Although NIFB presently includes agencies on most FRP routes, they were not doing so during the time period for which data was provided, and they do not formally define a minimum and maximum allocation for each agency; thus, we use the delivery data to infer minimum and maximum allocations.

As depicted in Figure 4.3, although donor parameters cover a wide range of values ($d^{\text{min}} = 0, \ldots, 31$, $d^{\text{max}} = 6, \ldots, 138$) agency parameters are much more limited ($a^{\text{min}} = 1, \ldots, 10$, $a^{\text{max}} = 3, \ldots, 50$). Furthermore, the minimum and maximum allocations at an agency tend to be strongly correlated, since agencies that can accept a large maximum allocation have to dedicate resources to participating in FRP, so they require a larger minimum allocation to justify their participation.

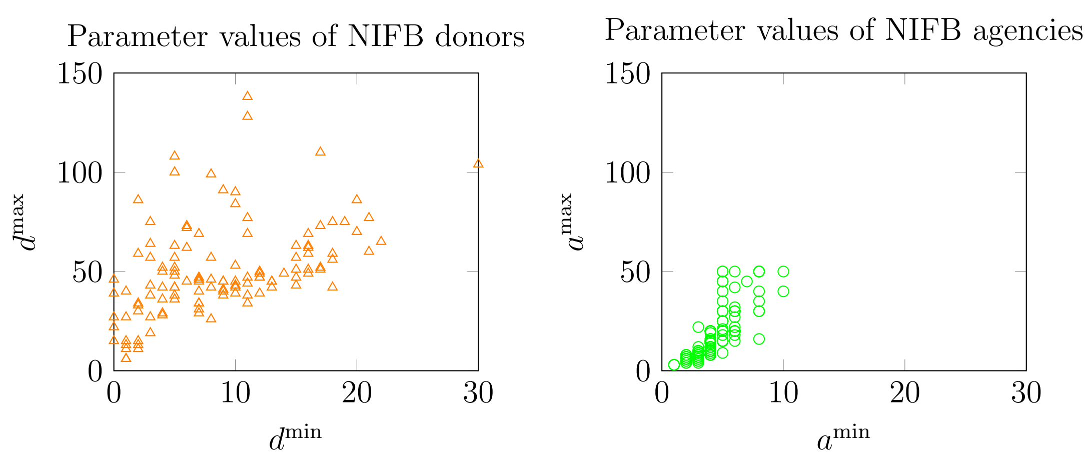

Figure 4.3 Parameter values of NIFB nodes

The NIFB service area is divided into four regions, referred to as North Suburban (NS), South Suburban (SS), West Suburban (WS), and Northwest (NW). We summarize

Table 4.4. Summary statistics for NIFB regions

	NS	SS	WS	NW
Number of donors ($\|\mathcal{D}\|$)	34	23	49	15
Number of agencies ($\|\mathcal{A}\|$)	38	27	48	26
Average travel time from the depot to a donor (minutes)	31.63	54.99	24.12	17.05
Average travel time between donors (minutes)	33.93	32.84	28.60	24.29
Average number of agencies no more than five minutes from a donor	0.76	1.22	1.35	1.53
Total maximum donations per total maximum allocations ($\sum d^{\max} / \sum a^{\max}$)	2.28	1.24	2.43	0.49

key statistics of the regions in Table 4.4. Maps of the regions are provided in Appendix C. Three factors distinguish the regions from one another: the availability of donations, the geographic dispersion of donors, and the potential for inserting agencies into FRP routes.

Availability of donations. For NS, SS, and WS, donors and agencies are roughly equal in number (that is, $|\mathcal{D}| \approx |\mathcal{A}|$), and the donations available are sufficient to satisfy the total demand from agencies, as measured by the ratio of total maximum donations to total maximum allocations in Table 4.4. This should not be interpreted as a surplus of donated food in the region: the maximum donation is not the average donation, only perishable food is considered, and only agencies that can potentially participate in FRP are included.

In NW, the situation is reversed: there are many more agencies than donors, and their combined demand is greater than the total supply. Therefore, adding even a single agency can dramatically reduce the capacity by a FRP route.

Geographic dispersion of donors. As reflected by the maps in Appendix C, every NIFB region contains both urban and rural areas. The range of travel time from the depot to donors is similar for WS and NW. However, donors in NW tend to be closer to the depot since that region includes only one large urban area, containing nearly all of the region's donors, while WS has donors throughout many urban areas.

Donors in NS and SS are more dispersed, as demonstrated by the average travel time between donors in Table 4.4. However, on average, donors in SS are about 22 minutes farther from the depot than those in NS, because SS is served by a depot that lies outside the region.

Potential for inserting agencies into FRP routes. Every NIFB region contains agencies that could participate in FRP, but to improve a route by inserting agencies, agencies must lie near the route. In most regions, agencies are dispersed in approximately the same areas as donors, providing ample opportunities for insertion. This is reflected in the average number of agencies within 5 minutes of a donor in Table 4.4. The exception to this is NS, in which most agencies lie in one cluster where there are no donors. Therefore, although $|\mathcal{D}| \approx |\mathcal{A}|$ for the region as a whole, over most of the region (all of it outside the cluster of agencies) there are many more donors than agencies.

4.4.2. NIFB challenges

NIFB includes agencies on its FRP routes to extend the capacity of its vehicles, and thus to reduce the total number of vehicles that must be used for FRP. Currently, NIFB includes at least one agency on each FRP route. The agencies are chosen and inserted in an ad hoc manner. With this research, we seek to formalize the method by which agencies are chosen and inserted by NIFB, and to inform other food banks that may wish to use agency insertion to address the challenge of limited vehicle capacity for FRP.

4.4.3. Experiment design

We use the NIFB data to construct 3,148 realistic instances of the selective 1-PDTSP with stochastic supply. An instance comprises a set of donors $\mathcal{D}$, a set of agencies $\mathcal{A}$, and a maximum total travel time T. To construct instances, we first construct sets of donors. We then associate up to five sets of agencies (one for each possible delivery day) with each donor set. Finally, for each combination of donor set and agency set, we generate separate instances for each of four maximum total travel time values.

Donor sets. We construct sets of donors by applying techniques observed at food banks to group donors into FRP routes. We refer to the donor set construction techniques as *Zone*, *Chain*, and *Destination*. The donors in a Zone set are located near one another. The set may contain all the donors in one large city or in several adjacent small cities or, if a part of the region is sparsely populated, all of the donors in a county. The donors included in a Chain set are all part of the same supermarket chain. Destination sets contain all the donors that lie along a possible route from the depot to a distant donor.

In Table 4.5, we list the number of distinct donor sets chosen for each region as well as the construction technique used to generate them. Some construction techniques are more readily applied in certain regions; for example, although there are few donors in NW, many of them are located far from the depot, so there are relatively more ways to apply the Destination technique. Since the vehicles used for FRP have a capacity of 200 boxes of donations, we only consider donor sets such that $\mathbf{Q}^D > 200$, since those are the only routes for which NIFB has an incentive to reduce the required capacity by inserting agencies.

Table 4.5. Donor sets by region and construction technique

	Total	Zone	Chain	Destination
NS	39	18	15	6
SS	31	11	10	10
WS	71	32	28	11
NW	13	4	3	6

Agency sets. We associate agency sets with the donor sets by identifying all agencies near the donors. We then use data about the agencies' availability to receive deliveries to generate up to five distinct agency sets for each donor set, one for each possible delivery day. The agency sets vary in size from 1 to 16 agencies, and the resulting sets of donors and agencies vary in size from 5 to 24 nodes, as summarized in Table 4.6.

Table 4.6. Composition of node sets in NIFB experiment, by number of donors and agencies

		$\lvert\mathcal{A}\rvert$															
		1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	Total
$\lvert\mathcal{D}\rvert$	4	4	4	6	10	1	6	2	0	3	3	1	0	4	2	0	46
	5	15	24	29	20	11	9	2	5	3	2	2	0	1	0	0	123
	6	12	15	24	14	14	10	7	7	6	7	4	2	2	0	0	124
	7	14	20	19	26	18	13	9	5	5	8	4	2	7	3	4	157
	8	12	16	29	20	16	17	14	4	3	4	2	3	4	6	0	150
	9	4	7	16	10	21	14	6	5	6	3	5	1	3	0	1	102
	10	2	4	9	9	15	17	11	9	3	3	1	0	2	0	0	85
	Total	63	90	132	109	96	86	51	35	29	30	19	8	23	11	5	787

Maximum total travel time. For each set of donors and agencies, we calculate the travel time of the TSP tour of the donors (denoted T^D) and generate instances with four values of the maximum total travel time: $T = T^D + 30, T^D + 60, T^D + 90, T^D + 120$. We assume that the stopping time at any node is 20 minutes, so increasing the maximum travel time by 30 minutes is roughly equivalent to allowing the insertion of one nearby agency.

4.4.4. Computational results

We solve the set of instances described in §4.4.3 using CPLEX 12.5.1 with one hour of CPU time, the same time limit used by Louveaux and Salazar-González [2009]. Only 61% of the instances can be solved to optimality within this time limit. Instance size (measured by the number of nodes) and total available travel time are the factors that we find to most influence solution time.

Table 4.7. Fraction of instances solved to optimality within time limit

	Maximum total travel time T			
$\lvert\mathcal{N}\rvert$	T^D+30	T^D+60	T^D+90	T^D+120
5	100%	100%	100%	100%
6	100%	100%	100%	100%
7	100%	100%	100%	100%
8	100%	100%	100%	100%
9	100%	100%	100%	100%
10	100%	100%	100%	100%
11	100%	99%	100%	100%
12	67%	51%	60%	91%
13	48%	22%	30%	54%
14	35%	17%	12%	35%
15	19%	12%	4%	5%
16	27%	20%	5%	0%
17	34%	20%	11%	0%
18	25%	14%	7%	0%
19	17%	8%	0%	0%
20	11%	0%	0%	0%
21	0%	0%	0%	0%
22	0%	0%	0%	0%
23	0%	0%	0%	0%
24	0%	0%	0%	0%

As summarized in Table 4.7, it is generally possible to obtain an optimal solution for instances with 6 to 10 nodes within the time limit, but as the number of nodes increases, the ability of CPLEX to obtain an optimal solution rapidly diminishes. CPLEX fares better for instances with $T = T^D + 30$; such instances have a smaller feasible region due to the stricter limit on total travel time.

As a consequence of the properties of the instances for which an optimal solution is available, we restrict our attention to sets of donors and agencies such that $|\mathcal{N}| \leq 11$ and such that an optimal solution was obtained for all four values of T. There are 367 such sets in the experiment, comprising 1,468 total instances.

Benefits of including agencies. We first consider a set of questions regarding the benefits of including agencies on FRP routes.

- Under which circumstances does including agencies on a FRP route yield a significant benefit?
- Which agencies should be recruited to participate in FRP?

As depicted in Figure 4.3, maximum allocation values are generally much smaller than maximum donation values, so one might expect only a modest decrease in required capacity by including agencies on a route. Furthermore, despite the randomness of donations, agencies need consistency in order to participate in FRP. Ensuring the minimum allocation for an agency may require an initial load, potentially increasing the required capacity. However, we find that the benefits of including agencies can be substantial, even with a modest increase in total travel time, if many agencies are available for inclusion.

Observation 4.4.1. *Including agencies is an effective intervention to reduce required capacity. Agency inclusion is more beneficial as the total available travel time increases and as the number of available agencies increases.*

We summarize the average benefit for varying maximum total travel time T and number of available agencies $|\mathcal{A}|$ in Table 4.8. As expected, the average reduction in capacity improves as T increases. This only fails to occur for very small values of $|\mathcal{A}|$, since there are no additional agencies to include with the available travel time. The reduction in required capacity also improves as the number of available agencies increases; clearly, if there are more agencies from which to choose, a better solution can be obtained, even if the travel time constraint only permits the inclusion of one or two.

Table 4.8. Average reduction in capacity as a percentage of $\mathbf{Q}^D$ for varying maximum total travel time T and number of available agencies $|\mathcal{A}|$

	$|\mathcal{A}|$						
T	1	2	3	4	5	6	7
$T^D + 30$	9%	10%	11%	10%	11%	11%	17%
$T^D + 60$	9%	14%	17%	17%	20%	20%	29%
$T^D + 90$	9%	14%	20%	22%	26%	30%	38%
$T^D + 120$	9%	14%	21%	23%	29%	36%	43%

To estimate the capacity of the optimal solution, $\mathbf{Q}^*$, we can calculate the ρ-agency lower bound on capacity, $\mathbf{Q}^{lb}(\rho)$, introduced in Corollary 4.2.4. This bound is an approximation that uses no information other than maximum distribution and allocation values, ignoring geographic information such as likely routes among the donors and the availability of agencies to include on donor routes.

One would expect the efficacy of the the ρ-agency lower bound to be related to the geography of the donor set or the region. However, we find that $\mathbf{Q}^{lb}(\rho)$ is an effective predictor of $\mathbf{Q}^*$ in all geographic settings.

Observation 4.4.2. *NIFB can obtain a high-quality estimate of the decrease in required capacity by calculating* $\mathbf{Q}^{lb}(\rho)$*. This measure uses no information about the available agencies beyond the* ρ *largest maximum allocation values. We observe this across all types of donor sets and in all regions.*

For each instance, we compare the solution obtained with maximum total travel time $T \in \{T^D + 30, T^D + 60, T^D + 90, T^D + 120\}$ to the ρ-agency lower bound on capacity obtained by $\rho \in \{1, 2, 3, 4\}$, respectively. Although $\mathbf{Q}^{lb}(\rho)$ is a lower bound on $\mathbf{Q}^*$ for $\rho = 1$ and $\rho = 2$ (since the stopping time is 20 minutes at all nodes and all edge travel times are positive), this is not the case for $\rho = 3$ and $\rho = 4$. In Table 4.9, we summarize the average absolute value of the gap between $\mathbf{Q}^*$ and $\mathbf{Q}^{lb}(\rho)$ for each of the regions and donor set construction techniques. We do not report a gap for regions and construction techniques for which there were no sets of donors and agencies such that an optimal solution was available for all four values of T. This occurs for all sets in NW because many agencies are clustered near the warehouse, so they are included in nearly all of the instances for the region, causing the number of nodes to be too large to obtain an optimal solution within the time limit.

As reflected in Table 4.9, the gap between $\mathbf{Q}^*$ and $\mathbf{Q}^{lb}(\rho)$ is small and shows little variation among regions or donor set construction techniques. Although we investigated several geographic characteristics of instances as potential predictors of the benefits of

Table 4.9. Average absolute gap between $\mathbf{Q}^*$ and $\mathbf{Q}^{lb}(\rho)$ as a percentage of $\mathbf{Q}^D$ by region and donor set construction technique

	Donor set construction technique		
Region	Zone	Chain	Destination
NS	0.5%	0.5%	0.8%
SS	0.3%	1.1%	2.1%
WS	0.1%	0.3%	0.0%
NW	—	—	—

including agencies, such as node density and the average distance of nodes from the depot, we found none of these to have a significant impact, either positive or negative, on the benefit of reducing capacity by including agencies or on the predictive ability of $\mathbf{Q}^{lb}(\rho)$.

It is possible to construct instances for which geography inhibits agency inclusion as an effective intervention. For example, if every agency in an instance were at least $T - T^D$ away from the nearest donor, it would be impossible to include even one agency on the FRP route. However, our numerical results demonstrate that, in realistic settings, geography has no discernible impact on average.

Table 4.8 clearly demonstrates that it is beneficial to have many agencies available for inclusion. This implies that NIFB should consider all eligible agencies for inclusion. However, the selective 1-PDTSP with stochastic supply becomes more difficult to solve as the number of nodes increases, which is precisely the context in which solving the problem to optimality is most valuable. Fortunately, we find that NIFB can realize the benefits of agency inclusion while considering relatively few agencies, chosen a priori by their maximum allocation.

Observation 4.4.3. *NIFB should emphasize the inclusion of agencies that can accept a large maximum allocation.*

We support Observation 4.4.3 by comparing $\mathbf{Q}^*$ and $\mathbf{Q}^{lb}(\rho)$ for instances with many agencies. Since $\mathbf{Q}^{lb}(\rho)$ uses no information about agencies other than their maximum allocations, the small gap we observe between $\mathbf{Q}^*$ and $\mathbf{Q}^{lb}(\rho)$ reflects the primacy of the value of $a^{\max}$ in predicting the benefit of agency inclusion.

In 85% of instances, $\mathbf{Q}^* = \mathbf{Q}^{lb}(\rho)$; on average, the absolute value of the difference between $\mathbf{Q}^*$ and $\mathbf{Q}^{lb}(\rho)$ as a percentage of $\mathbf{Q}^D$ is less than 1%. We summarize the average absolute gap between $\mathbf{Q}^*$ and $\mathbf{Q}^{lb}(\rho)$ in Table 4.10. For instances with many available agencies and a high maximum total travel time, the average absolute gap between $\mathbf{Q}^*$ and $\mathbf{Q}^{lb}(\rho)$ is larger because, for some instances in which agencies are located very close to donors or along likely routes between donors, it is possible to insert four agencies if $T = T^D + 90$ or five agencies if $T = T^D + 120$. In practice, instead of considering many agencies, food banks should consider a few large agencies, especially those that lie close to donors or along likely routes between donors.

Table 4.10. Average absolute gap between $\mathbf{Q}^*$ and $\mathbf{Q}^{lb}(\rho)$ as a percentage of $\mathbf{Q}^D$ for varying maximum total travel time T and number of available agencies $|\mathcal{A}|$

	$\lvert\mathcal{A}\rvert$						
T	1	2	3	4	5	6	7
$T^D + 30$	0.5%	0.1%	0.6%	0.6%	0.6%	0.0%	0.0%
$T^D + 60$	0.0%	0.2%	0.5%	0.2%	0.5%	0.0%	0.0%
$T^D + 90$	0.0%	0.0%	0.4%	0.8%	1.8%	3.0%	0.6%
$T^D + 120$	0.0%	0.0%	0.0%	0.0%	2.5%	3.3%	2.4%

Performance of CRIH. We propose CRIH to give decision-makers specific advice regarding the insertion of agencies into existing FRP routes. We use the computational study to investigate the following questions about CRIH:

- How does CRIH perform compared to the optimal solution?
- How does CRIH perform compared to other heuristic insertion techniques applied to the same initial route of donors?
- Under what conditions could CRIH provide a poor solution?

CRIH is a myopic heuristic for inserting agencies into existing FRP routes. We compare the optimal solution of the selective 1-PDSTP with stochastic supply with CRIH applied to an inital donor route designed to minimize distance. Since CRIH is deliberately designed to be applied by humans to routes that they have developed without optimization tools, we obtain initial donor routes by implementing a heuristic developed in cognitive science to model how humans solve the TSP [40].

To test the "cost" of not including the donor routing decision and of myopically inserting agencies, we consider the performance of CRIH on the instances for which an optimal solution was obtained.

Observation 4.4.4. *For the NIFB case study, CRIH provides near-optimal solutions with an average optimality gap of less than 0.5%.*

We summarize the performance of CRIH in Tables 4.11 and 4.12. The average optimality gap of CRIH applied to the heuristic donor routes is 0.46%. CRIH performs worse on instances with many agencies and high maximum total travel time. This occurs because the set of feasible solutions to the selective 1-PDTSP with stochastic supply increases with T, but CRIH explores only a small part of that space because (i) it starts from a single initial route and (ii) given a set of insertions that have the same estimated reuse, it chooses the one with the least insertion cost even if the available travel time is plentiful. The optimality gap is less than 1% in 89% of the instances in the case study, and it is less than 5% in 98% of instances. In Table 4.11, we see that significant gaps occur, on average, when there are many agencies to choose from; however, considering so many agencies is inefficient (see Observation 4.4.3).

In Appendix D, we describe several other insertion algorithms that we find to have worse performance than CRIH.

Table 4.11. Average optimality gap of CRIH applied to heuristically-generated donor routes for varying maximum total travel time T and number of available agencies $|\mathcal{A}|$

T	$\lvert\mathcal{A}\rvert$ 1	2	3	4	5	6	7
$T^D + 30$	0.5%	0.5%	0.5%	0.3%	0.5%	0.4%	0.0%
$T^D + 60$	0.1%	0.5%	0.6%	0.5%	2.0%	0.4%	0.7%
$T^D + 90$	0.1%	0.2%	0.3%	0.7%	0.9%	0.9%	2.2%
$T^D + 120$	0.1%	0.2%	0.2%	0.4%	1.1%	2.0%	5.0%

Table 4.12. Maximum optimality gap of CRIH applied to heuristically-generated donor routes for varying maximum total travel time T and number of available agencies $|\mathcal{A}|$

T	$\lvert\mathcal{A}\rvert$ 1	2	3	4	5	6	7
$T^D + 30$	15%	16%	16%	5%	11%	3%	0%
$T^D + 60$	1%	12%	20%	12%	21%	2%	1%
$T^D + 90$	1%	5%	4%	6%	7%	4%	4%
$T^D + 120$	1%	5%	4%	6%	6%	9%	9%

Some instances with the lowest value of maximum total travel time ($T = T^D + 30$) have large gaps because the travel time of the initial route is much greater than T^D, the total travel time of the optimal TSP solution. Since the stopping time at all agencies is 20 minutes, $T = T^D + 30$ implies that 10 minutes are available to divert from the optimal TSP donor route to travel to an agency. If the initial donor route is more than 10 minutes longer than T^D, it is impossible to insert an agency for $T = T^D + 30$. If the difference is less than 10 minutes, agency insertion is feasible, but fewer agencies are reachable than would be if 10 minutes were available. For instances with more than 30 minutes of additional travel time, the initial donor route had little impact.

Table 4.13. Average absolute gap between $\mathbf{Q}^*$ and $\mathbf{Q}^{lb}(\rho)$ as a percentage of $\mathbf{Q}^D$ for instances in which the optimality gap of CRIH is at least 5%

	$\lvert\mathcal{A}\rvert$						
T	1	2	3	4	5	6	7
$T^D + 30$	0.0%	0.3%	0.4%	1.5%	0.0%	—	—
$T^D + 60$	—	0.0%	0.0%	0.0%	0.0%	—	—
$T^D + 90$	—	0.0%	—	1.9%	0.0%	—	—
$T^D + 120$	—	0.0%	—	0.0%	3.1%	3.3%	3.0%

Due to its ease of implementation, CRIH is well-suited to the motivating context of our problem. Although it is myopic, it works well for the NIFB case study because the minimum donation values are generally greater than the minimum allocation values, so the ZMIL condition is approximately satisfied. However, there exist instances in which a more sophisticated approach is needed to obtain a high-quality solution. To identify such instances, we recommend that practitioners compare the CRIH solution to $\mathbf{Q}^{lb}(\rho)$: if $\mathbf{Q}^{lb}(\rho)$ is much less than the capacity of the CRIH solution, then the CRIH solution is potentially far from optimal.

As stated in Observation 4.4.3, $\mathbf{Q}^{lb}(\rho)$ is an excellent estimator of the optimal capacity. In Table 4.13, we report the average absolute optimality gap of $\mathbf{Q}^{lb}(\rho)$ as a percentage of $\mathbf{Q}^{D}$ for instances where the CRIH solution exhibits a large optimality gap (at least 5%). For the entire range of such instances, $\mathbf{Q}^{lb}(\rho)$ is close to $\mathbf{Q}^{*}$, with an average gap of less than 1%.

The computational results presented in this section demonstrate that agency inclusion is an effective and practical intervention for FRP. For a wide variety of realistic instances, the strategic inclusion of agencies allows for a significant reduction in vehicle capacity. Furthermore, we provide food banks with the tools necessary to implement this intervention: Our analytical results help identify which agencies should be considered for insertion, while CRIH makes near-optimal agency insertion decisions for existing FRP routes through a simple procedure.

CHAPTER 5

The 1-PDA-as

The Greater Chicago Food Depository (GCFD), the food bank that serves Chicago and suburban Cook County, Illinois, has developed a novel approach to provide its agencies with frequent deliveries of perishable foods. Few GCFD agencies own a large vehicle, so most contract a truck to make deliveries from the food bank warehouse approximately once per month. This limits the agencies' ability to offer perishable food, since any perishable food they receive must be distributed to clients quickly or, if possible, refrigerated or frozen and then rationed until the next delivery. To provide frequent deliveries of perishable food to such agencies, GCFD includes agencies on FRP routes.

Note the contrast between the implementation of FRP at NIFB and GCFD. As at NIFB, the primary purpose of FRP at GCFD is to collect donations. However, GCFD places greater emphasis on allocating donations to agencies. At NIFB, FRP allocations are supplemental to scheduled deliveries from the food bank, while at GCFD, FRP allocations are essential to agency operations.

We contend that the system developed by GCFD, which we deem "agency-supporting FRP," could be beneficial for other food banks, in particular those that serve densely-populated urban areas. We formulate a problem, the *one-commodity pickup and delivery allocation problem for agency-supporting FRP* (1-PDA-as), which finds an allocation policy for a given FRP route that maximizes the expected donations collected, while fulfilling specified service requirements at donors and agencies. Unlike the prevailing system at

GCFD, in which all donors are usually visited before any agencies, we consider general FRP routes in which donors and agencies can be visited in any order.

In §5.1, we formally define the 1-PDA-as as a stochastic program. In §5.2, we develop a procedure to solve the 1-PDA-as to optimality by solving a series of linear programs. In §5.3, we present analytical results regarding the optimal solution of a class of 1-PDA-as instances, then use these results as the basis for a heuristic procedure. In §5.4, we analyze computational results to obtain insights about donor and agency parameters, route structure, and the quality of solutions generated by our heuristic.

5.1. Formulation of the 1-PDA-as

In the present section, we formulate the 1-PDA-as as a stochastic program. The objective of the problem is to maximize the expected total donations collected for a given route. A solution to the 1-PDA-as comprises the initial load and the allocation policy at each agency. The 1-PDA-as is a generalization of the 1-PDA (Formulation (4.1)) with three additional parameters (described in §3.2): the guaranteed collection c_i, the sustaining allocation a_i^{sust}, and α_i, the minimum probability of receiving the sustaining allocation. Furthermore, note that since vehicle capacity is a parameter, we denote it Q (instead of $\mathbf{Q}$). In Tables 5.1 and 5.2, we summarize the parameters and decision variables of the 1-PDA-as.

Table 5.1. Parameters of the 1-PDA-as

Parameters	
Q	capacity of the vehicle, a positive integer
n	number of segments, indexed by $i \in \{0, 1, \ldots, n\}$; "0" represents the depot
D_i	donation from Donor i, a discrete random variable
c_i	guaranteed collection at Donor i
$a_i^{\min}$	minimum allocation at Agency i
$a_i^{\max}$	maximum allocation at Agency i
a_i^{sust}	sustaining allocation at Agency i
α_i	minimum probability that a_i^{sust} be allocated at Agency i

Table 5.2. Decision variables of the 1-PDA-as

Decision variables	
$\mathbf{S}_0$	initial load, a non-negative integer
$\mathbb{C}_i$	quantity of food collected from Donor i, a discrete integer-valued random variable; $\mathbb{C}_i \preccurlyeq D_i$
$\mathbf{A}_i(\cdot)$	allocation policy at Agency i; defined for all $s \in \text{supp}(\mathbb{S}_i^A)$
$\mathbb{A}_i$	long-run allocation at Agency i; $\mathbb{A}_i = \mathbf{A}_i(\mathbb{S}_i^A)$
$\mathbb{S}_i^D$	load upon arrival to Donor i, a discrete integer-valued random variable
$\mathbb{S}_i^A$	load upon arrival to Agency i, a discrete integer-valued random variable

The philosopher and game designer Ian Bogost remarks, "While we often think that rules always limit behavior, the imposition of constraints also creates expression" [9]. Such is the case when comparing the 1-PDA-as with the 1-PDA: the additional parameters present in the 1-PDA-as allow GCFD to *express* operational goals that are crucial to the success of its version of FRP. Unlike the 1-PDA, vehicle capacity is limited in the 1-PDA-as; specifying a guaranteed collection ensures that donations are collected consistently. The additional agency parameters are important because GCFD agencies depend on FRP allocations. To be clear, we do not suggest that a food bank employing the 1-PDA-as enter into a contract-like agreement with donors or agencies based on these parameters. Rather, their presence in the model provides the food bank an opportunity to improve the experience of participating in FRP for its donor and agency partners.

$$
\begin{aligned}
\text{maximize} \quad & \left[\sum_{i=1}^{n} \mathbf{E}\, \mathbb{C}_i\right] & & \quad (5.1a)\\
\text{subject to} \quad & \mathbb{C}_i \geq \min\{D_i, c_i\} & i \in \{1,\dots,n\} & \quad (5.1b)\\
& \mathbb{S}_i^A \leq Q & i \in \{1,\dots,n\} & \quad (5.1c)\\
& \mathbb{S}_i^A = \mathbb{S}_i^D + \mathbb{C}_i & i \in \{1,\dots,n\} & \quad (5.1d)\\
& \mathbb{S}_{i+1}^D = \mathbb{S}_i^A - \mathbb{A}_i & i \in \{1,\dots,n\} & \quad (5.1e)\\
& \mathbf{A}_i(s) \leq s \quad \forall s \in \operatorname{supp}(\mathbb{S}_i^A) & i \in \{1,\dots,n\} & \quad (5.1f)\\
& a_i^{\min} \leq \mathbb{A}_i \leq a_i^{\max} & i \in \{1,\dots,n\} & \quad (5.1g)\\
& \Pr\{\mathbb{A}_i \geq a_i^{sust}\} \geq \alpha_i & i \in \{1,\dots,n\} & \quad (5.1h)\\
& \mathbb{S}_1^D = \mathbf{S}_0 & & \quad (5.1i)\\
& \mathbf{S}_0 \in \mathbb{Z}_0^+ & & \quad (5.1j)
\end{aligned}
$$

The objective function (5.1a) maximizes the expected total donations collected on the FRP route. (This is the summary value C defined in §3.3.)

Constraints (5.1b) ensure that the guaranteed collection is accepted if the donor offers at least that amount of food. Otherwise, the collection is determined by the Maximum Collection Policy, considering the guaranteed collection at later donors (as explained in full §5.2.2). Constraints (5.1c) bound the load after collecting the donation by the vehicle capacity.

Constraints (5.1d) and (5.1e) describe the movement of food in and out of the vehicle. Constraints (5.1d) calculate the load after visiting the donor in terms of the load upon arrival to the donor and the collection. Constraints (5.1e) express the load after visiting the agency in terms of the load upon arrival to the agency and the long-run allocation.

Constraints (5.1f) state that allocation at an agency may not exceed the load upon arrival to the agency. Constraints (5.1g) and (5.1h) govern allocations by constaining the long-run allocation. Constraints (5.1g) force allocations to be no less than the minimum allocation and no more than the maximum allocation. Constraints (5.1h) ensure that each agency receives its sustaining allocation with at least the required minimum probability.

Constraint (5.1i) defines $\mathbb{S}_1^D$ as a random variable equal to the initial load $\mathbf{S}_0$ with probability 1. This definition is necessary for Constraints (5.1d) to be well-defined. Constraint (5.1j) declares the initial load $\mathbf{S}_0$ to be a non-negative integer.

5.2. Solving the 1-PDA-as to optimality

Dynamic programming is often applied to stochastic programs that resemble Formulation (5.1). In §5.2.1, we explain why that method and similar ones cannot be applied to the 1-PDA-as.

Since we cannot solve the 1-PDA-as directly, we have developed another solution technique, which we develop in the remainder of the present section. In §5.2.2, we eliminate the need to explicitly model the guaranteed collection through the concept of "effective capacity." In §5.2.3, we demonstrate that if $\mathbf{S}_0$ is fixed, the 1-PDA-as can be reformulated as a constrained Markov decision process (CMDP). In §5.2.4, we define the minimum intermediate load (from §4.1.2) in the context of the 1-PDA-as. In §5.2.5, we reformulate the CMDP from §5.2.3 as a linear program (LP). Thus, it is possible to solve the 1-PDA-as to optimality by solving a series of LPs (one for each feasible value of $\mathbf{S}_0$).

Our solution approach is novel. Each of the steps we apply to solve the 1-PDA-as is well-known; for example, the technique of reformulating a MDP as a LP was first described by d'Epenoux in 1960 [17]. However, the combination of these techniques is, as far as we are aware, unique in the operations research literature.

5.2.1. Inapplicability of dynamic programming and similar methods

Unlike the 1-PDA-as, the 1-PDA can be solved to optimality analytically because the lack of the sustaining allocation constraint allows the model to consider only the extreme values of supp(D_i) (see §4.1). A similar problem in the literature, the SRA-e (see §2.3), can be solved to optimality using dynamic programming (DP) [37]. However, DP cannot be applied to the 1-PDA-as because it is not possible to define states appropriately.

One option is to define the state space at a segment as the set of random variables that can represent the quantity of food in the vehicle. Therefore, the set of terminal states would be all possible values of $\mathbb{S}_n^D$, which is the set of all discrete random variables with support lying in $\{0, 1, \ldots, Q\}$. This set is uncountably infinite. Furthermore, the action space of this conceptualization of the problem would consist of all feasible randomized allocation policies, another uncountably infinite set. Such a problem is intractable for DP.

Another option is to define the set of states at a segment as all possible vehicle loads. However, this conceptualization is also not amenable to DP because the constraints regarding the sustaining allocation are global constraints whose satisfaction cannot be ensured through recursion. For this conceptualization of the problem, the actions are the possible random allocations from a given supply quantity; that is, if the vehicle contains s units of food upon arrival to Agency i, the feasible actions are all discrete random variables with support lying in $\left\{a_i^{\min}, \ldots, \min\{s, a_i^{\max}\}\right\}$. The actions taken at each $s \in \text{supp}(\mathbb{S}_i^A)$ comprise the allocation policy $\mathbf{A}_i(\cdot)$ for Segment i. However, whether the allocation policy satisfies the requirement to allocate at least a_i^{sust} to Agency i with probability α_i depends on the probability of each possible supply value s. That is, determining whether the allocation policy $\mathbf{A}_i(\cdot)$ is feasible requires calculating the long-run allocation $\mathbb{A}_i$, for which we must know $\Pr\{\mathbb{S}_i^A = s\}\ \forall s \in \text{supp}(\mathbb{S}_i^A)$. We cannot determine those probabilities without knowledge of the initial load and the actions at all prior segments, rendering the application of DP solution techniques impossible.

Several methods have been developed to solve stochastic programs and chance-constrained programs (CCP) with discrete random variables, but none we found can be applied to the 1-PDA-as. For example, Sample Average Approximation (SAA) consists of choosing realizations of the random variables in a problem, then finding values of the decision variables for which the chance constraints hold for sufficiently many of the realizations [54]. However, the applications of SAA to stochastic programming are generally limited to one-stage [49] or two-stage stochastic programs [2, 66]; since uncertainty is revealed at each segment in the 1-PDA-as, it is a multistage stochastic program. There are CCP solution approaches that rely on mixed integer linear program formulations, in which each possible outcome of the random vector is represented by a binary variable [39, 67], but they cannot be applied to the 1-PDA-as for the same reason. Although Shapiro is able obtain solutions to a multistage stochastic program through SAA [57], that approach (which is similar to DP) cannot be applied to the 1-PDA-as because it does not permit chance constraints nor the use of random variables as decision variables.

5.2.2. Effective capacity

The Maximum Acceptance Policy determines the amount of food collected at donors in the 1-PDA-as; however, due to the guaranteed collection c_i, its application is not as straightforward as for the 1-PDA. In some cases, less than the entire donation is collected so that capacity is available for the guaranteed collection of later donors.

In such cases, the Maximum Acceptance Policy still determines the amount of food collected, but with respect to a value less than the vehicle capacity. We term that value the "effective capacity."

Definition 5.2.1. *The* ***effective capacity*** *at Segment i, denoted q_i, is the maximum load upon departure from Donor i. It is defined recursively:*

$$q_n = Q \tag{5.2a}$$

$$q_i = \min\left\{q_{i+1} - c_{i+1} + a_i^{\max}, Q\right\} \tag{5.2b}$$

At Segment n, the effective capacity is the vehicle capacity; there are no later donors, so there are no later guaranteed collection values to consider. At prior segments, the effective capacity is determined by the effective capacity and guaranteed collection of the next segment and the maximum allocation of the current segment. The guaranteed collection in the next segment reduces the effective capacity in the current segment because it represents capacity that must be reserved to collect the donation at the next donor. The maximum allocation of the current segment increases the effective capacity of the current segment because it represents capacity that can be freed through allocation.

5.2.3. Constrained Markov decision process

If the initial load $\mathbf{S}_0$ is fixed, the 1-PDA-as can be reformulated as a constrained Markov decision process (CMDP).

The state space of the CMDP in Segment i is the set of possible load values upon reaching Agency i:

$$S_i = \{0, \ldots, q_i\} \tag{5.3}$$

Using the notation of Formulation (5.1), $S_i = \operatorname{supp}(\mathbb{S}_i^A)$.

The action set $A_i(s)$ from state s in Segment i consists of the possible allocations:

$$A_i(s) = \{a_i^{\min}, ..., \min\{s, a_i^{\max}\}\}$$

Using the notation of Formulation (5.1), $A_i(s) = \text{supp}(\mathbf{A}_i(s))$.

We define the initial state of the CMDP through the use of a Segment "0" that represents the load before visiting any donors or agencies (that is, the initial load). The state space of Segment 0 consists of only the initial load:

$$\Pr\{S_0 = \mathbf{S}_0\} = 1 \tag{5.4}$$

We must also define an action space for Segment 0. The sole action available is to allocate 0 units of food, since allocation only occurs at agencies, and Segment 0 does not contain an agency:

$$A_0(\mathbf{S}_0) = \{0\} \tag{5.5}$$

The transition probabilities from state s to state s' under action (allocation) a are denoted $p_i(s'|s,a)$. That is, $p_i(s'|s,a) = \Pr\{S_{i+1} = s'|S_i = s, A_i = a\}$. They are derived from the donor distributions:

$$p_i(s'|s,a) = \Pr\{D_{i+1} = s' - s + a\} \; \forall s' \in \{0, \ldots, q_{i+1} - 1\} \tag{5.6a}$$

$$p_i(q_{i+1}|s,a) = \Pr\{D_{i+1} \geq q_{i+1} - s + a\} \tag{5.6b}$$

The constraints of the CMDP are those imposed by the requirement to allocate at least a_i^{sust} to Agency i with probability α_i. This is expressed through Constraints (5.7). The left side represents the probability that allocations made from the states $s \in S_i$

provide at least the sustaining allocation. For a particular state s, this probability is $\Pr\{A_i(s) \geq a_i^{sust}\}$; the linear combination on the left side weights that probability by the probability that $S_i = s$ over all possible states.

$$\sum_{s \in S_i} \Pr\{A_i(s) \geq a_i^{sust}\} \cdot \Pr\{S_i = s\} \geq \alpha_i \ \forall i \in \{1, \ldots, n\} \tag{5.7}$$

The reward function of the CMDP, denoted by C, is the collection from the donor. It is the same for all segments:

$$C(s, a, s') = s' - (s - a) \tag{5.8}$$

5.2.4. Minimum intermediate load

In the solution method for the 1-PDA described in §4.1.2, the *minimum intermediate load*, denoted l_i^* plays a crucial role.

Our solution procedure for the 1-PDA-as does not require the calculation of the minimum intermediate load, but the value can be used to reduce the number of variables and constraints in the LP formulation. Since the LP must be solved many times to obtain the optimal solution to a single 1-PDA-as instance, any step we take to reduce its size is worthwhile.

As in our exposition in §4.1.2, we first calculate the supply gap g_i:

$$g_i = a_i^{\min} - d_i^{\min} \tag{5.9}$$

The minimum intermediate load at each segment is calculated recursively:

$$l_{n+1} = 0 \tag{5.10a}$$

$$l_i = (l_{i+1} + g_i)^+ \tag{5.10b}$$

Note that we denote the minimum intermediate load without an asterisk; that is, here we use l_i, whereas in §4.1.2, we use l_i^*. In the context of the 1-PDA, the minimum intermediate load is a component of the optimal solution; for the 1-PDA-as, the minimum intermediate load is simply a cut.

5.2.5. Linear program reformulation

The reformulation of a Markov decision process (MDP) as a linear program was first described by d'Epenoux [17]. The approach was first applied to a CMDP by Derman [18], who introduced the *state-action frequency approach* [52]. The LP solution method for CMDPs has been applied to hospital admission scheduling [35], highway maintenance [27], building maintenance [64], and the management of spectrum in wireless networks [71]. Each of these papers starts from a CMDP, which is then expressed as a LP and solved to obtain an optimal solution to the original problem. Our transformation of a stochastic program into a set of CMDPs, which are then expressed as LPs and solved in series to obtain an optimal solution, is unique in the literature.

State-action frequency variables represent the probability of occupying a state (s, i) and choosing action a. In the context of the 1-PDA-as, each state is an amount of food in the vehicle upon arrival to a specific agency: state (s, i) represents arriving to Agency i with load s. The actions are the possible allocations. Choosing action a from state (s, i)

represents the driver deciding to allocate a units of food from the s units in the vehicle to Agency i. Therefore, the probability represented by the state-action frequency variable is that of a joint event, which we denote $X^i_{sa} = \Pr\{\mathbb{S}^A_i = s \text{ and } \mathbf{A}_i(s) = a\}$.

Since solving an instance of the 1-PDA-as to optimality ultimately requires solving a series of LPs, any reduction in solution time for the LP is potentially valuable. We can reduce the size of the LP by eliminating variables which are not used in any feasible solution. In particular, the state-action frequency variables, X^i_{sa}, are only needed for those s in the minimal support of $\mathbb{S}^A_i$.

In general, there is no a priori way to determine $\text{supp}(\mathbb{S}^A_i)$, since $\mathbb{S}^A_i$ an auxiliary decision variable of the 1-PDA-as. Clearly, $\text{supp}(\mathbb{S}^A_i) \subseteq \{0, \ldots, Q\}$, but by applying the concepts developed in §5.2.2 and §5.2.4, we can obtain tighter bounds for the minimum and maximum values of $\text{supp}(\mathbb{S}^A_i)$:

Lemma 5.2.2. *Let $\mathcal{S}^D_i$ and $\mathcal{S}^A_i$ be supersets of* $\text{supp}(\mathbb{S}^D_i)$ *and* $\text{supp}(\mathbb{S}^A_i)$*, respectively. For Donor i,* $\text{supp}(\mathbb{S}^D_1) \subseteq \mathcal{S}^D_1 = \{\mathbf{S}_0\}$*. All other $\mathcal{S}^D_i$ and $\mathcal{S}^A_i$, for $i = 1, 2, \ldots, n$, are computed recursively:*

$$\text{supp}(\mathbb{S}^A_i) \subseteq \mathcal{S}^A_i = \left\{\min\{s^D + d, q_i\} \;\middle|\; s^D \in \mathcal{S}^D_i, d \in \text{supp}(D_i)\right\} \tag{5.11a}$$

$$\text{supp}(\mathbb{S}^D_{i+1}) \subseteq \mathcal{S}^D_{i+1} = \left\{s^A - a \;\middle|\; s^A \in \mathcal{S}^A_i, a \in \{a^{\min}_i, \ldots, a^{\max}_i\}, l_{i+1} \le s^A - a \le q_{i+1} - c_{i+1}\right\} \tag{5.11b}$$

Equation (5.11a) defines $\mathcal{S}^A_i$, the superset of $\text{supp}(\mathbb{S}^A_i)$, in terms of $\mathcal{S}^D_i$, the superset of $\text{supp}(\mathbb{S}^D_i)$. The elements of $\mathcal{S}^A_i$ are all possible sums of the load upon arrival to Donor i

($s^D \in \mathcal{S}_i^D$) and the donation from Donor i ($d \in \text{supp}(D_i)$), such that the sum not exceed the effective capacity.

Equation (5.11b) defines $\mathcal{S}_{i+1}^D$, the superset of $\text{supp}(\mathbb{S}_{i+1}^D)$, in terms of $\mathcal{S}_i^A$, the superset of $\text{supp}(\mathbb{S}_i^A)$. The elements of $\mathcal{S}_{i+1}^D$ are all possible differences of the load upon arrival to Agency i ($s^A \in \mathcal{S}_i^A$) and the allocation to Agency i ($a \in \{a_i^{\min}, \ldots, a_i^{\max}\}$), such that this difference is no less than the minimum intermediate load of Segment $i+1$ *and* no more than the effective capacity at Donor $i+1$ less the guaranteed collection at Donor $i+1$ (that is, $q_{i+1} - c_{i+1}$), which ensures that the guaranteed collection can be accepted without exceeding the effective capacity.

In Table 5.3, we summarize the additional notation presented in this section that appears in the formulation. We include the initial load $\mathbf{S}_0$ to emphasize that, for each LP, its value is fixed.

Table 5.3. Additional parameters and constraints for the LP reformulation of the 1-PDA-as

Additional Parameters for the LP	
$\mathcal{S}_i^A$	a superset of $\text{supp}(\mathbb{S}_i^A)$, a subset of $\{0, \ldots, Q\}$
$\mathbf{S}_0$	the initial load, a *fixed* non-negative integer
Decision Variables of the LP	
X_{sa}^i	the probability that $\mathbb{S}_i^A = s$ and $\mathbf{A}_i(s) = a$

Objective (5.12a) of the LP maximizes the expected total collection from all donors.The form of Objective (5.12a) is a consequence of solving Equation (3.8) for C (the expected total collection) to obtain $C = A + T - \mathbf{S}_0$. The first term computes the total expected allocation A by calculating the expected allocation at each segment, then summing over all segments. The second term computes the expected terminal load T by calculating the expected value of the difference between the load upon arrival to the final agency and the allocation to the final agency. The last term, subtracted from the sum of the first two, is simply the initial load $\mathbf{S}_0$.

Constraints (5.12b) and (5.12c) define a network among the X^i_{sa} and ensure that probability flow is conserved among them. Constraints (5.12b) govern events in which, even if the entire donation were collected, the load would be less than the effective capacity at that donor; Constraints (5.12c) govern events in which the load upon departing the donor is exactly the effective capacity. The left side of each of these constraints calculates the probability that $\mathbb{S}^A_i = s$ in terms of the X^{i-1}_{sa} and the possible combinations of donations d and allocations a'. The right side is the sum of the X^i_{sa} over all possible actions from state (s, i), which is simply $\Pr\{\mathbb{S}^A_i = s\}$.

Constraints (5.12d) ensure that each agency receives its sustaining allocation with at least the required probability. Constraint (5.12e) defines the initial load. Constraints (5.12f) restrict the X^i_{sa} to $[0, 1]$, since the state-action frequency variables represent probabilities.

$$\text{Maximize} \quad \sum_{i=1}^{n} \sum_{s \in \mathcal{S}_i^A} \sum_{a=a_i^{\min}}^{\min\{a_i^{\max},s\}} a \cdot X_{sa}^i + \sum_{s \in \mathcal{S}_n^A} \sum_{a=a_n^{\min}}^{\min\{a_n^{\max},s\}} (s-a) \cdot X_{sa}^n - \mathbf{S}_0 \tag{5.12a}$$

$$\text{subject to} \quad \sum_{d=d_i^{\min}}^{d_i^{\max}} \left(\Pr\{D_i = d\} \cdot \sum_{a'=a_{i-1}^{\min}}^{a_{i-1}^{\max}} X_{s-d+a',a'}^{i-1} \right) = \sum_{a=a_i^{\min}}^{a_i^{\max}} X_{sa}^i \tag{5.12b}$$

$$\forall s \in \mathcal{S}_i^A \setminus q_i, \forall i = \{1, \ldots, n\}$$

$$\sum_{d=d_i^{\min}}^{d_i^{\max}} \left(\Pr\{D_i \geq d\} \cdot \sum_{a'=a_{i-1}^{\min}}^{a_{i-1}^{\max}} X_{q_i-d+a',a'}^{i-1} \right) = \sum_{a=a_i^{\min}}^{a_i^{\max}} X_{q_i a}^i \tag{5.12c}$$

$$\forall i = \{1, \ldots, n\}$$

$$\sum_{s \in \mathcal{S}_i^A} \sum_{a \geq a_i^{sust}} X_{sa}^i \geq \alpha_i \tag{5.12d}$$

$$\forall i = \{1, \ldots, n\}$$

$$X_{\mathbf{S}_0 0}^0 = 1 \tag{5.12e}$$

$$X_{sa}^i \in [0,1] \tag{5.12f}$$

$$\forall a \in \{a_i^{\min}, \ldots, \min\{a_i^{\max}, s\}\}, \forall s \in \mathcal{S}_i^A, \forall i = \{1, \ldots, n\}$$

The only variables of the LP are the X_{sa}^i, of which there is one for each element in $\mathcal{S}_i^A$ for each segment, as well as the additional variable $X_{\mathbf{S}_0 0}^0$ to represent the initial load. Therefore, the LP has $\sum_{i=1}^{n} |\mathcal{S}_i^A| + 1$ variables. Since $|\mathcal{S}_i^A| \leq Q + 1$, the LP has at most $(Q+1)n + 1$ variables.

The LP has $\sum_{i=1}^{n} |\mathcal{S}_i^A| + 2n + 1$ constraints. Each segment has $|\mathcal{S}_i^A|$ of Constraints (5.12b), for a total of $\sum_{i=1}^{n} |\mathcal{S}_i^A|$ constraints conserving probability flow when the donation does not exceed the effective capacity. Each segment has one of the Constraints (5.12c), for a total of n constraints conserving probability flow when the donation *does* exceed the effective capacity. Each segment has one of the Constraints (5.12d), for a total of n constraints regarding the sustaining allocation. The lone Constraint (5.12e) defines the initial load. The remaining constraints merely describe the range of the X^i_{sa}. Since $|\mathcal{S}^A_i| \le Q + 1$, the LP has at most $(Q + 3)n + 1$ constraints.

5.3. One-supersegment routes

In the present section, we narrow our focus to a class of routes with a particular structure: those that consist of a set of donors followed by a set of agencies. We term this structure a "supersegment;" therefore, this type of FRP route can equivalently be referred to as a *one-supersegment* route. This route structure is a useful object of study because it is large enough to present complexity in the allocation decision, but small enough to obtain strong analytical results.

Furthermore, a general FRP route can be conceptualized as a sequence of supersegments. For example, consider the following route, in which donors are represented as D and agencies as A:

$$D \ A \ D \ D \ A \ A \ D$$

The route comprises three supersegments:

$$D \ A \mid D \ D \ A \ A \mid D$$

In §5.3.1, we define a myopic allocation policy, $\mathbf{A}^m$, for one-supersegment routes. In §5.3.2, we prove analytical properties of the myopic allocation policy, including conditions that guarantee its optimality. In §5.3.3, we employ a procedure based on the myopic allocation policy $\mathbf{A}^m$ to adapt the MILB solution procedure from §4.1 as a heuristic solution method for the 1-PDA-as.

5.3.1. Myopic allocation policy

At Agency i, any amount between $a_i^{\min}$ and $a_i^{\max}$ may be allocated. However, Constraints (5.1g) and (5.1h) establish minimum requirements only on allocations of at least $a_i^{\min}$ and a_i^{sust}, respectively. Therefore, a feasible allocation policy need only dictate the conditions under which to allocate $a_i^{\min}$ and a_i^{sust}. Likewise, a feasible allocation policy need only ensure that a_i^{sust} be allocated with a probability of exactly α_i. We term such an allocation policy, in which constraints regarding allocation are satisfied exactly at their minimum values, a *minimal allocation policy*.

There is an obvious incentive to allocate more than the minimum required by constraints – namely, to free capacity in order to collect additional donations. We address this in §5.3.3 through the "discretionary allocation," but at present, note that applying a minimal allocation policy does not impact the total collection for a one-supersegment route, since all donations are collected before any allocations are made.

The load upon arrival to Agency i is an observation s of the random variable $\mathbb{S}_i^A$. Intuitively, it is reasonable to allocate more food when the vehicle contains more food. That is, when s is greater than some value, a_i^{sust} should be allocated; otherwise, $a_i^{\min}$ should be allocated. If the "cutoff" on s is chosen so that a_i^{sust} is allocated when the upper α_i of the distribution of $\mathbb{S}_i^A$ is observed, then the constraints on allocation at Agency i are met exactly. In Definition 5.3.1, we formalize this intuition.

Definition 5.3.1. ***Myopic allocation policy*** *Let $s_i^{1-\alpha}$ denote the $(1-\alpha_i)$ quantile of $\mathbb{S}_i^A$; that is, $s_i^{1-\alpha}$ is the greatest $s \in \operatorname{supp}(\mathbb{S}_i^A)$ such that $\Pr\{\mathbb{S}_i^A \geq s\} \geq \alpha_i$. The myopic allocation policy at Agency i, denoted $\mathbf{A}_i^m(s)$, is the allocation at Agency i when s units of food are available:*

$$\mathbf{A}_i^m(s) = \begin{cases} a_i^{sust} & s > s_i^{1-\alpha} \quad (5.13a) \\ a_i^{min} & s < s_i^{1-\alpha} \quad (5.13b) \\ \begin{cases} a_i^{sust} & w.p. \ \frac{\Pr\{\mathbb{S}_i^A > s_i^{1-\alpha}\} - (1-\alpha_i)}{\Pr\{\mathbb{S}_i^A = s_i^{1-\alpha}\}} \\ a_i^{min} & w.p. \ \frac{(1-\alpha_i) - \Pr\{\mathbb{S}_i^A < s_i^{1-\alpha}\}}{\Pr\{\mathbb{S}_i^A = s_i^{1-\alpha}\}} \end{cases} & s = s_i^{1-\alpha} \quad (5.13c) \end{cases}$$

Cases (5.13a) and (5.13b), respectively, state that a_i^{sust} be allocated when the load is greater than the $(1-\alpha)$ quantile of $\mathbb{S}_i^A$ and that $a_i^{\min}$ be allocated when the load is less. Case (5.13c) addresses the allocation when the load is exactly equal to the $1-\alpha_i$ quantile of $\mathbb{S}_i^A$. If $a_i^{\min}$ were allocated in this case, the probability of receiving the sustaining allocation would be less than α_i, violating Constraints (5.1h); therefore, the definition establishes a random allocation under which a_i^{sust} is allocated with the minimum probability necessary so that the probability of receiving the sustaining allocation is exactly α_i.

5.3.2. Analytical results

By considering the one-supersegment case, we obtain several analytical results about the 1-PDA-as and the myopic allocation policy. Theorem 5.3.2 establishes a relationship between the initial load and total collection. This result is useful for proving later results, because, in the context of a proof, it is generally easier to compute the initial load than to compute the total collection. The proof of Theorem 5.3.2 and the other theorems in this subsection are provided in Appendix A.

Theorem 5.3.2. *For a one-supersegment route, any solution that minimizes the initial load* $\mathbf{S}_0$ *is optimal.*

Note that minimizing the initial load is more than a mathematical convenience. Food that remains in the food bank warehouse is available to agencies that own or lease vehicles. Therefore, reducing the initial load required for FRP has intrinsic value in food bank operations.

The intuitive appeal of the myopic allocation policy is borne out by our analytical results regarding simple route structures.

Theorem 5.3.3. *For a one-supersegment route with a single agency, the myopic allocation policy* $\mathbf{A}^m$ *is optimal.*

The proof of Theorem 5.3.3 demonstrates optimality by showing that the initial load obtained by applying the myopic allocation policy, $\mathbf{S}_0^m$, is the minimum feasible initial load. There may exist other allocation policies with initial load $\mathbf{S}_0^m$, which, by Theorem 5.3.2, are also optimal solutions. Corollary 5.3.4 states conditions under which the myopic allocation policy is the unique optimal minimal allocation policy.

Corollary 5.3.4. *For a one-supersegment route with a single agency, if* $\Pr\{\mathbb{S}_n^A \geq s_n^{1-\alpha}\} = \alpha_n$ *and* $s_n^{1-\alpha} = a_n^{sust}$*, then the myopic allocation policy* $\mathbf{A}^m$ *is the unique optimal minimal allocation policy.*

With Theorem 5.3.5, we generalize Theorem 5.3.3 to a one-supersegment route with any number of agencies, if a condition on the order of the agencies and the differences between their sustaining and minimum allocations is satisfied. In essence, this condition states that the allocation decision at each agency is of greater magnitude than all subsequent allocation decisions combined.

Theorem 5.3.5. *If* $(a_k^{sust} - a_k^{\min}) \geq \sum_{i=k+1}^{n}(a_i^{sust} - a_i^{\min})$ *for all* $k \in \{|\mathcal{D}|, \ldots, n\}$, *then the myopic allocation policy* $\mathbf{A}^m$ *is optimal.*

When determining the allocation policy within a supersegment, the order of the agencies is mutable. In reality, the FRP vehicle must visit the agencies in the order stipulated, but no additional information is obtained as the agencies are visited; therefore, the allocation decision for each agency can be made before visiting any. That is, after collecting the final donation, the agencies can be "sorted" by decreasing $(a_i^{sust} - a_i^{\min})$ to determine the allocation at each agency, even though the agencies are not visited in that order. If, after "sorting," the decreasing differences condition of Theorem 5.3.5 is satisfied, then the myopic allocation policy is optimal. A one-supersegment route with only two agencies must satisfy this condition when sorted, so the myopic allocation policy is always optimal in that case:

Corollary 5.3.6. *For a one-supersegment route with two agencies, if "sorting" is applied, then the myopic allocation policy* $\mathbf{A}^m$ *is optimal.*

5.3.3. Heuristic adaptation of MILB algorithm

In §4.1.2, we presented the Minimum Intermediate Load-Based (MILB) algorithm to solve the 1-PDA to optimality. Here, we adapt that algorithm to develop a heuristic procedure for the 1-PDA-as. Instead of the MILB allocation policy $\mathbf{A}_i^{l^\star}$, we apply the myopic allocation policy $\mathbf{A}^m$; therefore, we refer to the heuristic as the Myopic MILB or "MMILB" procedure.

The MMILB procedure has essentially the same steps as the MILB algorithm, but applied to supersegments (indexed by $j \in \{1, \ldots, \tilde{n}\}$) instead of segments. Each supersegment is treated as an independent one-supersegment route. They are first solved from the end of the route to the beginning (to obtain minimum intermediate loads and the initial load), then from the beginning to the end (to obtain the allocation policy, and hence the 1-PDA-as solution).

Step 1: Calculate minimum intermediate loads and obtain initial load. As for the MILB algorithm, we obtain the MMILB minimum intermediate loads $\tilde{l}_j$ recursively, starting with $\tilde{l}_{\tilde{n}+1} = 0$. The minimum intermediate load at the final supersegment, $\tilde{l}_{\tilde{n}}$, is the minimum feasible initial load for the last supersegment considered as an independent route with the myopic allocation policy $\mathbf{A}^m$ applied. For supersegments prior to the final one, the minimum intermediate load is computed in the same way (as the minimum feasible initial load of the supersegment as an independent route), except that $\mathbf{A}^m$ is applied to the total donations less the minimum intermediate load of the next supersegment.

Step 2: Define allocation policy. The myopic allocation policy is a minimal allocation policy: it allocates no more than the amount needed to exactly satisfy the constraints regarding allocation. However, the objective of the 1-PDA-as is to maximize collection, which is aided by freeing vehicle capacity. From the perspective of the model, there is no incentive to keep excess food in the vehicle. Furthermore, in consideration of the motivating context of the problem, one of GCFD's primary motivations for including agencies on FRP routes is the opportunity to deliver perishable food to them quickly.

Therefore, the MMILB allocation policy has two components: the myopic allocation policy $\mathbf{A}^m$ and the discretionary allocation. The discretionary allocation is the difference between the supply (less the MMILB minimum intermediate load of the next supersegment) and $\mathbf{A}^m$. As noted in §5.3.2, the allocation decision is made for all agencies in the supersegment simultaneously, so this difference is known before visiting the first agency.

Deliberately, we provide no guidance regarding how the discretionary allocation should be distributed among the agencies (except, of course, that the allocation at an agency not exceed its maximum allocation). The decision is left to the FRP driver's discretion. The heuristic includes this design choice as a recognition of the shortcomings of the 1-PDA-as and, indeed, of any model. For example, we have assumed that food is a homogeneous commodity, but clearly it is not; certain agencies may be better able to make use of certain types of food, due to characteristics of the community they serve or the types of programs they administer. This information is available to the FRP driver, as well as other information about the agencies' operations. The digital rhetorician James J. Brown, Jr. observes that computational procedures often lack "the flexibility required for the ethical predicaments of hospitality" [11], and that this undermines their effectiveness and makes human users less likely to apply them; allowing the FRP driver to distribute the discretionary allocation introduces needed flexibility to our 1-PDA-as solution.

Step 3: Apply $\mathbf{A}^m$ to determine feasibility and evaluate objective. For the 1-PDA, vehicle capacity is a decision variable, so the existence of a feasible solution is guaranteed. This is not the case for the 1-PDA-as; hence, this step of the MMILB procedure is slightly different from that of the MILB algorithm.

The initial load $\mathbf{S}_0$ is simply the minimum intermediate load of the first supersegment, $\tilde{l}_1$. By applying the MMILB allocation policy (which consists of the myopic allocation policy $\mathbf{A}^m$ and the discretionary allocation) iteratively to each supersegment, we determine whether it provides a feasible solution and, if so, we calculate the expected total allocation.

5.4. Computational results

In §5.4.1, we describe the experimental design used to study the properties of the 1-PDA-as. In §5.4.2, we summarize insights obtained from analysis of optimal solutions to the problem. In §5.4.3, we evaluate solutions obtained with the MMILB heuristic procedure.

5.4.1. Experiment design

Although operations at GCFD motivate our formulation of the 1-PDA-as, we do not have data from GCFD for our computational study. Therefore, we use the NIFB data described in §4.4, with some adaptations.

Vehicle. For all instances, we use a vehicle capacity of $\mathbf{Q} = 200$, the approximate capacity in boxes of the refrigerated straight trucks used for FRP by GCFD.

Donors. For the 1-PDA, only the minimum and maximum donation amounts $d^{\min}$ and $d^{\max}$ are relevant. The 1-PDA-as requires the full distribution D. However, the state-action frequency variables of Formulation (5.12) are products of probabilities, which poses a challenge due to the arithmetic precision of CPLEX (or any other solver). The minimum precision available in CPLEX is 1e−9; therefore, to include routes that contain up to 9 consecutive agencies, the least probable donation amount must have a probability of at least 0.1. In consideration of this issue, we obtain distributions from the NIFB data with a specific structure. The distribution is defined in terms of the minimum, first quartile, median, third quartile, and maximum of the observed data:

$$D = \begin{cases} d^{\min} & \text{w.p. } 0.1 \\ Q1 & \text{w.p. } 0.25 \\ Q2 & \text{w.p. } 0.35 \\ Q3 & \text{w.p. } 0.25 \\ d^{\max} & \text{w.p. } 0.1 \end{cases} \tag{5.14}$$

Several values are used for the guaranteed collection c through the experiment design, as described below.

Agencies. For the 1-PDA, only the minimum and maximum allocation amounts $a^{\min}$ and $a^{\max}$ are relevant. The 1-PDA-as requires a third value, the sustaining allocation, which we compute as the ceiling (to obtain an integer) of the mean of the minimum and maximum allocations:

$$a^{sust} = \left\lceil \frac{a^{\min} + a^{\max}}{2} \right\rceil \tag{5.15}$$

Several values are used for α, the minimum probability of receiving the sustaining allocation, through the experiment design, as described below.

Routes. The experiment consists of a large set of randomly generated FRP routes. First, we generate one hundred random sets of nodes for each problem size from 6 to 15 nodes, for a total of one thousand node sets. Then, from each node set, we generate twenty random routes (with the restriction that the first node be a donor). Then, from each route, we obtain instances based on different levels of the donor and agency parameters, denoted "Low," "Moderate," and "High."

The donor levels are based on the guaranteed collection value, with a higher level corresponding to a donor whose entire donation is collected more frequently:

Donor Low: $c = Q1$

Donor Moderate: $c = Q2$

Donor High: $c = Q3$

The agency levels are based on the probability of receiving the sustaining allocation, with a higher level corresponding to an agency that receives its sustaining allocation more frequently:

Agency Low: $\alpha = 0.6$

Agency Moderate: $\alpha = 0.9$

Agency High: $\alpha = 1$

For the Agency High level, note that $\alpha = 1$ implies that $a^{\min} = a^{sust}$.

There are nine combinations of the donor and agency levels for each route; therefore, the experiment consists of $180,000$ instances. Routes of 6 to 10 nodes were solved to optimality. All instances were solved heuristically with the MMILB procedure. The experiment was run on a Windows 7 Pro virtual machine with 32 GB of RAM using CPLEX 12.5.1.

5.4.2. Analysis of optimal solutions

Of the $90,000$ instances with 6 to 10 nodes, 94% are feasible. As the number of nodes increases, the fraction of feasible instances decreases, which is unsurprising since the vehicle capacity does not change. Likewise, the average solution time increases as the number of nodes (an approximation of the problem size) increases. These relationships are apparent in Table 5.4. Although the optimal solution procedure requires solving a series of LPs, each LP takes a relatively short time to solve (likely in part due to the simple form of the donor distributions chosen due to issues of numerical precision), so the average solution time is low.

Number of nodes	Feasible instances	Average solution time
6	98.7%	0.54
7	97.8%	1.21
8	91.8%	3.01
9	94.4%	5.12
10	86.5%	10.62

Table 5.4. Fraction of instances feasible and average CPLEX solution time (in CPU seconds) for instances solved to optimality

Donor and agency levels. Each route in the experiment is run nine times, once for each combination of the donor and agency levels described in §5.4.1. Of these combinations, Donor Low with Agency Low is the least restrictive (i.e., the instance with the largest feasible region), while Donor High with Agency High is the most restrictive. Thus, as the level increases from Low to Moderate to High for either donors or agencies, the fraction of instances that are feasible cannot increase.

The impact of donor and agency levels on feasibility is summarized in Table 5.5. Increasing the agency level has little impact on feasibility; in fact, only the change from Agency Moderate to Agency High, which eliminates the distinction between the minimum and sustaining allocations, registers a decrease of more than 1% in the fraction of instances feasible, and even then only for some donor levels. In contrast, increasing the donor level has a pronounced impact on feasibility. Although a high guaranteed collection provides better service to a donor, reserving excessive capacity for collections makes it impossible to adequately serve agencies when donations are low, substantially decreasing the fraction of instances feasible.

	Agency Low	Agency Moderate	Agency High
Donor Low	100%	100%	100%
Donor Moderate	97%	97%	96%
Donor High	86%	86%	84%

Table 5.5. Fraction of instances feasible, by combinations of donor and agency levels

Since it restricts the feasible region, increasing the donor or agency level can also impact the objective. There are $8,340$ routes that are feasible for all combinations of the donor and agency levels. (That is, 84% of the $10,000$ routes, corresponding to those that are feasible for the combination of Donor High and Agency High.) In Tables 5.6 and 5.7, we compare the objective value for each combination of levels with the objective value for the least restrictive combination, Donor Low and Agency Low.

Table 5.6 indicates the fraction of routes for which the objective value is less than that for Donor Low and Agency Low. As observed above for fraction of instances feasible, the difference between Agency Low and Agency Moderate is negligible. However, the difference of Agency High is more marked, likely because the lack of flexibility at agencies requires an increase in the initial load. The impact of donor levels is similar to that observed for fraction of instances feasible.

	Agency Low	Agency Moderate	Agency High
Donor Low	0%	1%	11%
Donor Moderate	7%	7%	15%
Donor High	11%	12%	18%

Table 5.6. Fraction of instances for which the objective value is less than that for Donor Low and Agency Low

Although a substantial fraction of routes have a lower objective value for the more restrictive donor and agency levels, the difference is slight. In Table 5.7, we report the *maximum* gap from the objective for Donor Low and Agency Low. The maxima are very low; in fact, the average gap for every combination of levels is 0%. More stringent requirements at donors and agencies reduce the capacity available to collect donations (either because that capacity is reserved for later donors or is occupied by an increased initial load), but the impact on the objective value is limited because the objective function is an expectation, and it is improbable that the additional capacity available in the least restrictive case is utilized.

Route structure. As described in §5.4.1, twenty routes are generated for each node set. We compare the solutions for these twenty routes to obtain insights about the impact of node sequencing on the 1-PDA-as. (Clearly, more than twenty sequences are possible for

	Agency Low	Agency Moderate	Agency High
Donor Low	0%	0%	1%
Donor Moderate	0%	0%	1%
Donor High	2%	2%	2%

Table 5.7. Maximum gap in objective value compared to that for Donor Low and Agency Low

a set of 6 to 10 nodes; the twenty-route sample is a limitation of our experiment.) To isolate our attention to the impact of node sequencing, we only consider instances with the levels Donor Low and Agency Low.

Number of nodes	Gap exists	Average gap
6	17%	0.2%
7	19%	0.2%
8	43%	0.9%
9	43%	0.9%
10	57%	2.5%

Table 5.8. Gaps in objective value between the best and worst routes for each node set, for instances with the levels Donor Low and Agency Low

In Table 5.8, for each node set, we compare the route with the highest objective value with the rest of the routes. As the number of nodes increases, the impact of node sequencing increases. For only 6 or 7 nodes, all twenty routes have the same objective value in more than 80% of node sets, but for 10 nodes, fewer than half of node sets have the same property. The average gap between the best and worst routes also increases as the number of nodes increases.

To interpret these relationships, it is important to recall that the vehicle capacity is the same for all instances. Therefore, as the number of nodes increases, it is more likely that donations exceed capacity, and therefore that donors and agencies should be ordered judiciously to free capacity.

Number of nodes	Gap exists	Average gap
6	8%	0.0%
7	13%	0.0%
8	35%	0.3%
9	24%	0.4%
10	48%	0.8%

Table 5.9. Gaps in objective value between the best and worst routes for each node set, for instances with the levels Donor Low and Agency Low; only includes routes that end in a donor

Upon examination of particular instances, it is apparent that the best routes had agencies dispersed throughout the route. In most cases, the worst sequence for a node set ended with one or several agencies. A primary motivation for including agencies on FRP routes is to free capacity to accept more donations, but an agency at the end of the route does not fulfill this purpose. (The same is true of an agency at the beginning of the route, but by design, no such routes were generated.) The importance of agency position is reflected in Table 5.9, which compares the best and worst routes for each node set, but only among routes that end with a donor. Compared to Table 5.8, gaps between the best and worst routes occur less often, and the average gap in objective value is substantially less.

5.4.3. Performance of MMILB procedure

We test two implementations of the MMILB procedure. As noted in §5.3.2, the agencies in a supersegment can be "sorted" by decreasing value of $a_i^{sust} - a_i^{\min}$ before applying the myopic procedure. (This "sorting" does not correspond to a change in the FRP route, but merely a change in the order in which allocation decisions are made, since all the information pertinent to allocation is available before the first agency in the supersegment is visited.) Therefore, for each instance, we apply the myopic procedure to the given node sequence (denoted "myopic") *and* to the node sequence with the agencies sorted by decreasing $a_i^{sust} - a_i^{\min}$ in each supersegment (denoted "sorted myopic").

In Table 5.10, we summarize the performance of each implementation of the MMILB procedure by the number of nodes in the route. As expected for a heuristic, the solution time is faster than the optimal procedure (see Table 5.4). However, our primary motivation for developing a heuristic is not the solution time required for an optimal solution (which is modest), but the limitation on the instances that can be solved with the optimal procedure due to issues of numerical precision. Unlike the optimal procedure, the MMILB procedure can be applied to routes of any length, without restrictions on the form of the D_i.

	Solution time	
Number of nodes	Myopic	Sorted myopic
6	0.04	0.03
7	0.04	0.04
8	0.05	0.05
9	0.05	0.05
10	0.06	0.06
11	0.06	0.07
12	0.07	0.08
13	0.07	0.08
14	0.09	0.11
15	0.09	0.09

Table 5.10. Average solution time for implementations of the MMILB procedure, in seconds

	Myopic		Sorted myopic	
Number of nodes	Average gap	Maximum gap	Average gap	Maximum gap
6	0.0%	1.3%	0.0%	1.3%
7	0.0%	2.9%	0.0%	2.9%
8	0.0%	6.1%	0.0%	6.1%
9	0.0%	8.1%	0.0%	8.1%
10	0.0%	11.4%	0.0%	11.4%

Table 5.11. Average and maximum optimality gaps for implementations of the MMILB procedure

The two implementations of the MMILB procedure have nearly identical performance in terms of optimality gaps, as summarized in Table 5.11. The sorted myopic implementation has a slight advantage in the rate at which a gap exists. In Tables 5.12 and 5.13, respectively, we summarize the fraction of instances for which each implementation of the MMILB procedure has a gap, by donor and agency level. (These statistics are restricted to routes for which every combination of the levels is feasible.) For several combinations of donor and agency level, the sorted myopic implementation exhibits an optimality gap at a slightly lower rate than the myopic implementation.

Tables 5.12 and 5.13 reveal a pattern in the instances for which the MMILB solution (of either implementation) has an optimality gap: gaps are more common for instances with the Agency Moderate level. In terms of the feasibility of the instance (see Table 5.5) and the impact on the objective (see Tables 5.6 and 5.7), Agency Moderate is indistinguishable from Agency Low; here, in contrast, the impact of the difference is stark. This occurs because the other agency levels are relatively simple from the perspective of the heuristic: for Agency Low, since the sustaining allocation is very low, an initial load is probably not needed; for Agency High, the minimum and sustaining allocation are the same, so no allocation decision need be made. For Agency Moderate, the allocation decision is the most challenging, so MMILB obtains an optimal solution less frequently (although the solutions it finds are still quite good; the average optimality gap for such instances is 0%).

	Agency Low	Agency Moderate	Agency High
Donor Low	0.09%	4.45%	0.00%
Donor Moderate	0.09%	4.50%	0.00%
Donor High	0.11%	4.83%	0.00%

Table 5.12. Fraction of instances for which the myopic implementation of the MMILB procedure exhibits an optimality gap

	Agency Low	Agency Moderate	Agency High
Donor Low	0.07%	4.44%	0.00%
Donor Moderate	0.07%	4.48%	0.00%
Donor High	0.08%	4.79%	0.00%

Table 5.13. Fraction of instances for which the sorted myopic implementation of the MMILB procedure exhibits an optimality gap

In Table 5.14, we report the average gap between the minimum initial load associated with the optimal solution and the initial load mandated by each implementation of the MMILB procedure. The sorted myopic implementation has a clear advantage over the myopic implementation in terms of initial load. This is a consequence of Theorem 5.3.5; the proof of that result demonstrates conditions under which ordering agencies by the metric used for sorting minimizes the initial load for a one-supersegment route. Although minimizing the initial load is not the objective of the 1-PDA-as, it is advantageous for the food bank to keep food in the warehouse for other means of distribution, so the sorted myopic solution is preferable.

Number of nodes	Initial load gap Myopic	Sorted myopic
6	3.2	3.0
7	4.8	4.5
8	3.9	3.7
9	4.9	4.5
10	4.3	4.0

Table 5.14. Average gap in initial load for implementations of the MMILB procedure

CHAPTER 6

Conclusion

Too many food bankers get hung up on trucks, docks, and warehouses, and they forget that their real goal is to get families on their feet.

Willy Elliot-McCrea

Executive Director, Second Harvest Food Bank of Santa Cruz County

In this work, we studied two models of FRP operations: the selective 1-PDTSP with stochastic supply and the 1-PDA-as. We demonstrated that the inclusion of agencies on FRP routes can be beneficial to the food bank, donors, and agencies. This is achieved by structuring routes and allocation decisions so that agencies can be used to mitigate the challenges posed by randomness in donations, permitting vehicle capacity to be freed and reused. Several opportunities exist to enhance the models we have formulated.

The selective 1-PDTSP with stochastic supply would be more useful if it could be applied to multiple vehicles, i.e., if it were a generalization of the 1-PDVRP. In practice, several vehicles are used each day to collect donations in each NIFB region, and each agency can be assigned to only FRP route, so a model that considers only one vehicle at a time is of limited value. Additional research should also be done regarding heuristic approaches for solving the selective 1-PDTSP with stochastic supply. Although CRIH applied to heuristic routes of donors works well in the context of NIFB, its success relies heavily on the ZMIL condition being approximately fulfilled. In another context (for example, a food bank with a high minimum allocation at every agency), most FRP routes could have nonzero initial and intermediate loads, causing the simple $\tilde{R}_i$ used by CRIH to be a poor estimate of capacity reuse.

An unfortunate aspect of the 1-PDA-as is that, due to issues of numerical precision, it can only be solved to optimality for a restricted class of instances. The MMILB heuristic performs well, but more research is warranted to find a different solution technique that guarantees optimality without modeling joint probabilities as decision variables, the underlying cause of the issues of numerical precision. Furthermore, the 1-PDA-as is a generalization of a part of the selective 1-PDTSP with stochastic supply, so it would be useful to add the other components of that problem (node sequencing and routing), and ideally to consider multiple vehicles as well. At present, since it can only be applied to a single route, the 1-PDA-as is not useful when total donations on a route are less than the vehicle capacity; in that case, there is no reason to consider including agencies. However, a multiple-vehicle analogue of the 1-PDA-as could combine existing FRP routes, utilizing agency inclusion to ensure that the routes produced satisfy the capacity constraint.

The eventual goal of this line of research would be an integrated model considering all of the food collection and distribution tasks of the food bank. The vehicles used for FRP routes are also those that make scheduled deliveries at agencies; that collect large donations from food manufacturers and wholesalers; and that support "mobile pantries," events in which food is distributed from a truck directly to people (generally in an area that is not served by a food pantry). Including all of these decisions in the same model would allow for extremely efficient solutions that combine multiple tasks, such as an FRP route that begins after a scheduled delivery to a distant agency. Such a model could also include the option for some agencies to directly collect FRP donations from donors, a recent innovation at NIFB.

Although the selective 1-PDTSP with stochastic supply and the 1-PDA-as apparently consider only the logistics of collecting and distributing donated food, the potential contribution of Operations Research to the campaign against hunger is more profound.

Recently, a growing chorus of anti-hunger advocates have criticized the emphasis that Feeding America and its member food banks place on growth, particularly as measured by pounds of food distributed and number of people served. They argue that food banks perpetuate their own expansion by creating a co-dependent relationship among the food bank and food recipients, in which increasing the supply of food only increases demand. They also criticize the poor nutritional quality of much food distributed by food banks, the lack of diversity in food bank leadership, and that food banks in general manage hunger as a symptom rather than attempting to reduce poverty [23, 68].

Feeding America and its member food banks have responded to these critiques. For example, food banks are working to improve food quality by increasing the quantity of produce they distribute [20], and Feeding America has recently created a senior level position tasked with "ending hunger" [23]. A few food banks have innovative programs that directly address the root causes of poverty, such as the community food security programs at Oregon Food Bank, the community-supported agriculture (CSA) programs supported by the Capital Area Food Bank and the Western Massachusetts Food Bank, and the use of warehouse space to catalyze a new catering business at Foodlink (in Rochester, New York) [23], but the ability of food banks in general to emulate these programs or try other ideas is hampered by the effort required to manage the logistics of their operations. The core work of food banks remains the collection and distribution of food, and as sociologist Janet Poppendieck points out, "because food programs are logistically demanding, their maintenance absorbs the attention and energy of many of the people most concerned about the poor, distracting them from the larger issues of distributional politics" [51].

Operations Research is the tool that food bank leadership lacks. Through partnership with OR experts, food banks can reduce the effort required to address their current logistics challenges and utilize their existing resources more efficiently to support innovative programs that address the root causes of poverty. The extensive recent OR research in charitable food collection and distribution outlined in §2.1, including this work, are early steps in what will hopefully become an extensive and mutually beneficial collaboration between food banks and the OR community.

References

[1] C. Adams and M. Tabacchi. Perishable-food-rescue programs: Status and prospects. *Cornell Hotel and Restaurant Administration Quarterly*, 38(2):62–67, 1997.

[2] S. Ahmed and A. Shapiro. The sample average approximation method for stochastic programs with integer recourse. Technical report, Georgia Tech ISyE, 2002. URL `http://www2.isye.gatech.edu/~sahmed/saasip.pdf`.

[3] B. Ata, D. Lee, and E. Sönmez. Dynamic staffing of volunteer gleaning operations. SSRN Working Paper, 2016. URL `https://ssrn.com/abstract=2873250`.

[4] E. Balas. The prize collecting traveling salesman problem. *Networks*, 19(6):621–636, 1989.

[5] B. Balçık, S. Iravani, and K. Smilowitz. Multi-vehicle sequential resource allocation for a nonprofit distribution system. *IIE Transactions*, 46(12):1279–1297, 2014.

[6] G. Berbeglia, J.-F. Cordeau, I. Gribkovskaia, and G. Laporte. Static pickup and delivery problems: a classification scheme and survey. *TOP*, 15(1):1–31, 2007.

[7] D. Bienstock, M. X. Goemans, D. Simchi-Levi, and D. Williamson. A note on the prize collecting traveling salesman problem. *Mathematical Programming*, 59(1-3): 413–420, 1993.

[8] J. Bloom. *American wasteland: how America throws away nearly half of its food (and what we can do about it)*. Da Capo Press, 2010.

[9] I. Bogost. *Persuasive games: the expressive power of videogames.* The MIT Press, 2007.

[10] L. Brock and L. Davis. Estimating available supermarket commodities for food bank collection in the absence of information. *Expert Systems with Applications*, 42(7): 3450–3461, 2015.

[11] J. Brown. *Ethical programs : hospitality and the rhetorics of software.* University of Michigan Press, 2015.

[12] A. Campbell, L. Clarke, A. Kleywegt, and M. Savelsbergh. *Fleet Management and Logistics*, chapter The inventory routing problem, pages 95–113. Citeseer, 1998.

[13] M.C. Chou, C.P. Teo, and H. Zheng. Process flexibility: design, evaluation, and applications. *Flexible Services and Manufacturing Journal*, 20(1):59–94, 2008.

[14] A. Coleman-Jensen, C. Gregory M. Rabbit, and A. Singh. Household Food Security in the United States in 2015. USDA ERS Report No. 215, September 2016.

[15] J. Cuevas-Ortuño and A. Gómez-Padilla. Un modelo de asignación-empaque de despensas personalizadas para bancos de alimentos: un sistema sujeto a condiciones nutricionales y logísticas. *DYNA-Ingeniería e Industria*, 88(5), 2013.

[16] L. Davis, I. Sengul, J. Ivy, L. Brock, and L. Miles. Scheduling food bank collections and deliveries to ensure food safety and improve access. *Socio-Economic Planning Sciences*, 48(3):175–188, 2014.

[17] F. d'Epenoux. Sur un problème de production et de stockage dans l'aléatoire. *Revue Française de Recherche Opérationnelle*, 14:3–16, 1960.

[18] C. Derman. *Finite State Markovian Decision Processes.* Academic Press, Inc., 1970.

[19] Feeding America. Hunger in America 2014, August 2014.

[20] Feeding America. 2016 Annual Report: Nourishing Healthy Futures, 2016.

[21] D. Feillet, P. Dejax, and M. Gendreau. Traveling salesman problems with profits. *Transportation Science*, 39(2):188–205, 2005.

[22] S. Fianu and L. Davis. A Markov decision process model for equitable distribution of supplies under uncertainty. *European Journal of Operational Research*, 2017.

[23] A. Fisher. *Big hunger: the unholy alliance between corporate America and anti-hunger groups*. The MIT Press, 2017.

[24] M. Gendreau, G. Laporte, and F. Semet. A branch-and-cut algorithm for the undirected selective traveling salesman problem. *Networks*, 32(4):263–273, 1998.

[25] A. Ghoniem, C. Scherrer, and S. Solak. A specialized column generation approach for a vehicle routing problem with demand allocation. *Journal of the Operational Research Society*, 64(1):114–124, 2013.

[26] R. Giesen, H.S. Mahmassani, and P. Jaillet. Logistics in real time inventory routing operations under stochastic demand. *Innovations in Distribution Logistics*, pages 109–148, 2009.

[27] K. Golabi, R.B. Kulkarni, and G.B. Way. A statewide pavement management system. *Interfaces*, pages 5–21, 1982.

[28] C. Güneş, W.J. van Hoeve, and S. Tayur. *Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems*, chapter Vehicle Routing for Food Rescue Programs: A Comparison of Different Approaches, pages 176–180. Springer, 2010.

[29] V. Hemmelmayr, K. Smilowitz, and L. de la Torre. A periodic location routing problem for collaborative recycling. *IISE Transactions*, 49(4):414–428, 2017.

[30] H. Hernández-Pérez and J.J. Salazar-González. The one-commodity pickup-and-delivery travelling salesman problem. In Michael Juenger, Gerhard Reinelt, and Giovanni Rinaldi, editors, *Combinatorial Optimization – Eureka, You Shrink!*, volume 2570 of *Lecture Notes in Computer Science*, pages 89–104. Springer Berlin / Heidelberg, 2003.

[31] H. Hernández-Pérez and J.J. Salazar-González. A branch-and-cut algorithm for a traveling salesman problem with pickup and delivery. *Discrete Applied Mathematics*, 145(1):126–139, 2004.

[32] H. Hernández-Pérez and J.J. Salazar-González. Heuristics for the one-commodity pickup-and-delivery traveling salesman problem. *Transportation Science*, 38(2):245–255, 2004.

[33] L.M. Hvattum and A. Løkketangen. Using scenario trees and progressive hedging for stochastic inventory routing problems. *Journal of Heuristics*, 15(6):527–557, 2009.

[34] A.J. Kleywegt, V.S. Nori, and M.W.P. Savelsbergh. The stochastic inventory routing problem with direct deliveries. *Transportation Science*, 36(1):94, 2002.

[35] P. Kolesar. A Markovian model for hospital admission scheduling. *Management Science*, pages 384–396, 1970.

[36] G. Laporte and S. Martello. The selective travelling salesman problem. *Discrete Applied Mathematics*, 26(2):193–207, 1990.

[37] R. Lien, S. Iravani, and K. Smilowitz. Sequential resource allocation for nonprofit operations. *Operations Research*, 62(2), 2014.

[38] F. Louveaux and J.J. Salazar-González. On the one-commodity pickup-and-delivery traveling salesman problem with stochastic demands. *Mathematical Programming*,

119(1):169–194, 2009.

[39] J. Luedtke, S. Ahmed, and G.L. Nemhauser. An integer programming approach for linear programs with probabilistic constraints. *Mathematical Programming*, 122(2): 247–272, 2010.

[40] J. N. MacGregor, T. C. Ormerod, and E. P. Chronicle. A model of human performance on the traveling salesperson problem. *Memory & Cognition*, 28(7):1183–1190, 2000.

[41] S. Mohan, M. Gopalakrishnan, and P.J. Mizzi. Improving the efficiency of a non-profit supply chain for the food insecure. *International Journal of Production Economics*, 143(2):248–255, 2013.

[42] D. J. Nair, H. Grzybowska, D. Rey, and V. V. Dixit. Food rescue and delivery: Heuristic algorithm for periodic unpaired pickup and delivery vehicle routing problem. *Transportation Research Record: Journal of the Transportation Research Board*, (2548):81–89, 2016.

[43] D. J. Nair, H. Grzybowska, Y. Fu, and V. V. Dixit. Scheduling and routing models for food rescue and delivery operations. *Socio-Economic Planning Sciences*, 2017.

[44] D. J. Nair, T. H. Rashidi, and V. V. Dixit. Estimating surplus food supply for food rescue and delivery operations. *Socio-Economic Planning Sciences*, 57:73–83, 2017.

[45] D. J. Nair, D. Rey, and V. V. Dixit. Fair allocation and cost-effective routing models for food rescue and redistribution. *IISE Transactions*, pages 1–17, 2017.

[46] M. Nestle. *What to eat.* North Point Press, 2006.

[47] C. Novoa and R. Storer. An approximate dynamic programming approach for the vehicle routing problem with stochastic demands. *European Journal of Operational*

Research, 196(2):509–515, 2009.

[48] İ. Ş. Örgüt, J. Ivy, R. Uzsoy, and J. Wilson. Modeling for the equitable and effective distribution of donated food under capacity constraints. *IIE Transactions*, 48(3): 252–266, 2016.

[49] B. Pagnoncelli, S. Ahmed, and A. Shapiro. Sample average approximation method for chance constrained programming: Theory and applications. *Journal of Optimization Theory and Applications*, 142:399–416, 2009.

[50] C. Phillips, R. Hoenigman, B. Higbee, and T. Reed. Understanding the sustainability of retail food recovery. *PloS one*, 8(10):e75530, 2013.

[51] J. Poppendieck. *Sweet charity? : emergency food and the end of entitlement*. Viking Penguin, 1998.

[52] M.L. Puterman. *Markov Decision Processes: Discrete Stochastic Dynamic Programming*. John Wiley & Sons, Inc., 1994.

[53] M. Reihaneh and A. Ghoniem. A multi-start optimization-based heuristic for a food bank distribution problem. *Journal of the Operational Research Society*, pages 1–17, April 2017.

[54] R. Rossi. An introductory note to sample average approximation. *ArXiv e-prints*, October 2011. URL `http://arxiv.org/abs/1110.1892`.

[55] M. Ruhlman. *Grocery: the buying and selling of food in America*. Abrams Press, 2017.

[56] A. Sen. *Poverty and famines*. Oxford University Press, 1981.

[57] A. Shapiro. On complexity of multistage stochastic programs. *Operations Research Letters*, 34(1):1–8, 2006.

[58] Z. Shen, F. Ordóñez, and M. Dessouky. The minimum unmet demand stochastic vehicle routing problem. Working paper, University of Southern California, 2006.

[59] S. Solak, C. Scherrer, and A. Ghoniem. The stop-and-drop problem in nonprofit food distribution networks. *Annals of Operations Research*, 221(1):407–426, 2014.

[60] E. Sönmez, D. Lee, M. Gómez, and X. Fan. Improving food bank gleaning operations: An application in New York State. *American Journal of Agricultural Economics*, 98 (2):549–563, 2015.

[61] M. Sousa Alves Silva. Problema de recobrimento de rotas com coleta de prêmios. Master's thesis, Instituto de Computaçao, Universidade Federal Fluminense, Niterói, Rio de Janeiro, Brazil, 2009.

[62] H. Süral and J. H. Bookbinder. The single-vehicle routing problem with unrestricted backhauls. *Networks*, 41(3):127–136, 2003.

[63] R. G. van Anholt, L. C. Coehlo, G. Laporte, and I. F.A. Vis. An inventory-routing problem with pickups and deliveries arising in the replenishment of automated teller machines. Technical report, CIRRELT, 2013.

[64] C. van Winden and R. Dekker. Rationalisation of building maintenance by Markov decision models: A pilot case study. *Journal of the Operational Research Society*, 49 (9):928–935, 1998.

[65] P. Vansteenwegen, W. Souffriau, and D. Van Oudheusden. The orienteering problem: A survey. *European Journal of Operational Research*, 209(1):1–10, 2011.

[66] B. Verweij, S. Ahmed, A. J. Kleywegt, G. Nemhauser, and A. Shapiro. The sample average approximation method applied to stochastic routing problems: a computational study. *Computational Optimization and Applications*, 24(2-3):289–333, 2003.

[67] J. P. Vielma, S. Ahmed, and G. L. Nemhauser. Mixed integer linear programming formulations for probabilistic constraints. *Operations Research Letters*, 2012.

[68] M. Winne. *Closing the food gap: resetting the table in the land of plenty*. Beacon Press, 2008.

[69] H. Yıldız, M. P. Johnson, and S. Roehrig. Planning for meals-on-wheels: algorithms and application. *Journal of the Operational Research Society*, 64(10):1540–1550, 2013.

[70] A. Youn, M. Ollinger, and L.S. Kantor. Characteristics of Mid-Atlantic food banks and food rescue organizations. *FoodReview*, 22:45–51, 1999.

[71] Q. Zhao, S. Geirhofer, L. Tong, and B. Sadler. Opportunistic spectrum access via periodic channel sensing. *IEEE Transactions on Signal Processing*, 56(2):785–796, 2008.

APPENDIX A

Proofs

LEMMA 4.1.3: *Given a scenario $(d_1, d_2, \ldots, d_n)$ and initial load $\mathbf{S}_0$, no allocation policy allocates more in total than the MILB allocation policy $\mathbf{A}_i^l(\cdot)\ \forall\ i \in \{1, \ldots, n\}$.*

Proof. We first note that for a given load s at Agency i, no allocation policy could allocate more than the MILB allocation policy. This is a direct consequence of the structure of the MILB allocation policy, Equation (4.4): If $\mathbf{A}_i^{l^*}(s) = s - l_{i+1}^*$, allocating more would cause the intermediate load upon arrival to the next segment to be less than the minimum intermediate load; if $\mathbf{A}_i^{l^*}(s) = a_i^{\max}$, allocating more would violate the maximum allocation constraint at Agency i.

Since $\mathbb{S}_i^D$ and $\mathbb{S}_i^A$ are random variables, they are minimized in the sense of *stochastic dominance*. That is, for any scenario $(d_1, d_2, \ldots, d_n)$, the realizations of $\mathbb{S}_i^D$ and $\mathbb{S}_i^A$ are stochastically dominated by the realizations resulting from the application of any other allocation policy; therefore, the $\mathbb{S}_i^D$ and $\mathbb{S}_i^A$ resulting from the application of the MILB allocation policy are stochastically dominated by those resulting from the application of any other allocation policy.

For any allocation policy, the relationship between $\mathbb{S}_i^A$ and $\mathbb{S}_{i+1}^D$ from (4.1c) is given by:

$$\mathbb{S}_{i+1}^D = \mathbb{S}_i^A - \mathbb{A}_i$$

We solve the equation for the long-run allocation at Agency i:

$$\mathbf{A}_i = \mathbb{S}_i^A - \mathbb{S}_{i+1}^D$$

This relationship also holds in aggregate for the entire route:

$$\sum_{i=1}^{n} \mathbb{A}_i = \sum_{i=1}^{n} \left(\mathbb{S}_i^A - \mathbb{S}_{i+1}^D\right) = \sum_{i=1}^{n} \mathbb{S}_i^A - \sum_{i=1}^{n} \mathbb{S}_{i+1}^D$$

As a consequence of the above, the quantity $\sum_{i=1}^{n} \mathbb{S}_{i+1}^D$ is minimized by the application of the MILB allocation policy. An allocation policy could only allocate *more than* the MILB allocation policy by increasing at least one of the $\mathbb{S}_i^A$; but doing so would require increasing $\mathbb{S}_i^D$ by the same amount, leaving the value of $\sum_{i=1}^{n} \mathbb{S}_i^A - \sum_{i=1}^{n} \mathbb{S}_{i+1}^D$ unchanged. Therefore, no allocation policy can allocate more than the MILB allocation policy.

$\square$

THEOREM 4.1.4: *The MILB solution (consisting of initial load $\mathbf{S}_0^{l^*}$, allocation policy $\mathbf{A}_i^{l^*}(\cdot)\ \forall\ i \in \{1, \ldots, n\}$, and capacity $\mathbf{Q}^{l^*}$) is an optimal solution to the 1-PDA.*

Proof. We prove the claim by demonstrating that the MILB maximum intermediate load $u_i^{l^*}$ is a feasible lower bound on $\max \operatorname{supp}(\mathbb{S}_i^A)$ at each segment. By Constraints (4.1f), since $\mathbb{S}_i^A = \mathbb{S}_i^D + D_i$, the $\mathbb{S}_i^A$ determine the capacity, so demonstrating this statement about the $u_i^{l^*}$ proves the claim. We proceed by induction.

Initial step at Segment 1: For the first segment, $u_1^{l^*} = \mathbf{S}_0^{l^*} = l_1^*$, which is the minimum feasible value of $\mathbb{S}_1^A$ since it is the minimum intermediate load for the first segment.

Induction hypothesis for Segment $j-1$: We assume that u^{l*}_{j-1} is a feasible lower bound on $\max \operatorname{supp}(\mathbb{S}^A_{j-1})$.

Demonstration of claim for Segment j: We show that u^{l*}_j is a feasible lower bound on $\max \operatorname{supp}(\mathbb{S}^A_j)$.

We calculate u^{l*}_j by applying Equation (4.5):

$$u^{l*}_j = \max\left\{l^*_j, u^{l*}_{j-1} + d^{\max}_{j-1} - a^{\max}_{j-1}\right\}$$

Since none of the other values in the expression are decision variables, u^{l*}_j is a feasible lower bound on $\max \operatorname{supp}(\mathbb{S}^A_j)$ if u^{l*}_{j-1} is minimized. Due to the induction hypothesis, we know that u^{l*}_{j-1} is a feasible lower bound on $\max \operatorname{supp}(\mathbb{S}^A_{j-1})$, and hence that it is minimized, so u^{l*}_j is minimized. Therefore, the MILB solution minimizes $\max \operatorname{supp}(\mathbb{S}^A_j)$ at Segment j. By induction, this applies to all segments; therefore, since $\mathbf{Q} \geq \max_i \max \operatorname{supp}(\mathbb{S}^A_i)$, the MILB solution obtains the minimum capacity.

$\square$

THEOREM 4.2.1: *Given a solution to the implicit 1-PDA, the allocation policy:*

$$\mathbf{A}^l_i(s) = \min\{s - l_{i+1}, a^{\max}_i\} \ \forall\ s \in \operatorname{supp}(\mathbb{S}^A_i),\ \ \forall\ i \in \{1, \ldots, n\}$$

is feasible.

Proof. We prove the claim by demonstrating that $\mathbf{A}^l_i(\cdot)$ satisfies Constraints (4.1d) and (4.1e) of the 1-PDA.

Constraints (4.1d) require that $\mathbf{A}_i^l(s) \leq s$. By Equation (4.8), $\mathbf{A}_i^l(s) \leq s - l_{i+1} \leq s$, satisfying the constraint.

Constraints (4.1e) require that $\mathbf{A}_i^l(\cdot)$ allocate at least the minimum allocation and no more than the maximum allocation at each segment.

Regarding the minimum allocation: $\mathbf{A}_i^l(\cdot)$ allocates either $a_i^{\max}$ or $s - l_{i+1}$. If $a_i^{\max}$ is allocated, then clearly at least the minimum allocation is provided. If $s - l_{i+1}$ is allocated, then applying Constraints (4.7c):

$$s - l_{i+1} \geq l_i + d_i^{\min} - l_{i+1} > l_{i+1} - d_i^{\min} + a_i^{\min} + d_i^{\min} - l_{i+1} = a_i^{\min} \qquad \text{(A.1)}$$

The restriction regarding the maximum allocation is satisfied since one of the terms of the minimization in Equation (4.8) is $a_i^{\max}$.

□

THEOREM 4.2.2: *Inserting an agency on an edge adjacent to the depot cannot decrease the minimum capacity.*

Proof. Since we know that it produces an optimal solution (Theorem 4.1.4), we apply the MILB algorithm.

Restating Equation (4.6), the minimum required capacity for a route is given by:

$$\mathbf{Q}^{l^*} = \max\left\{u_1^{l^*} + d_1^{\max}, u_2^{l^*} + d_2^{\max}, \ldots, u_n^{l^*} + d_n^{\max}\right\} \qquad \text{(A.2)}$$

Consider the insertion of an agency at the beginning of the route. We refer to the agency considered for insertion as "Agency 0" with minimum allocation $a_0^{\min} > 0$. Recall that $u_1^{l^*} = \mathbf{S}_0^{l^*}$. For the route without the insertion of Agency 0, $\mathbf{S}_0^{l^*} = l_1^*$. If Agency 0 were inserted before Donor 1, the minimum allocation at Agency 0 would have to be provided from the initial load; that is, $\mathbf{S}_0^{l^*} = l_1^* + a_0^{\min}$. This would increase $u_1^{l^*}$, in turn increasing the first term of Equation (A.2), which can only increase $\mathbf{Q}^{l^*}$.

Consider the insertion of an agency at the end of the route. We refer to the agency considered for insertion as "Agency $(n+1)$" with minimum allocation $a_{n+1}^{\min} > 0$. Recall that $u_n^{l^*} = \max\left\{l_n^*, u_{n-1}^{l^*} + d_{n-1}^{\max} - a_{n-1}^{\max}\right\}$. For the route without the insertion of Agency $(n+1)$, $l_n^* = \left(a_n^{\min} - d_n^{\min}\right)^+$. If Agency $(n+1)$ were inserted after Agency n, the minimum allocation at Agency $(n+1)$ would have to be ensured by the minimum intermediate load at the prior segment; that is, $l_n^* = \left(a_{n+1}^{\min} + a_n^{\min} - d_n^{\min}\right)^+$. This can only increase $u_n^{l^*}$, in turn increasing the last term of Equation (A.2), which can only increase $\mathbf{Q}^{l^*}$.

$\square$

THEOREM 4.2.3: *For any route of the set of donors $\mathcal{D}$ and the set of agencies $\mathcal{A}$, a bound on the minimum required capacity is:*

$$\mathbf{Q}^* \geq \sum_{D \in \mathcal{D}} d_D^{\max} - \sum_{A \in \mathcal{A}} a_A^{\max}$$

Proof. We use $\mathbf{Q}^{lb}$ to denote the right side of Inequality (4.14). That is, $\mathbf{Q}^{lb} \equiv \sum_{D \in \mathcal{D}} d_D^{\max} - \sum_{A \in \mathcal{A}} a_A^{\max}$.

We prove the claim by contradiction. Assume that there exists a route σ with initial load $\mathbf{S}_0^{\sigma}$ such that the minimum capacity $\mathbf{Q}^{\sigma}$ is less than $\mathbf{Q}^{lb}$. We use $\mathbb{S}_1^{A\sigma}$ to denote the load after collecting the donation in Segment i on route σ under the maximum scenario; therefore $\max\operatorname{supp}(\mathbb{S}_1^{A\sigma})$ is the load upon arrival to Agency i under the maximum scenario.

For Segment 1,

$$\max\operatorname{supp}(\mathbb{S}_1^{A\sigma}) = \mathbf{S}_0^\sigma + d_1^{\max}$$

For Segment 2, the statement regarding $\max\operatorname{supp}(\mathbb{S}_2^{A\sigma})$ is an inequality because the allocation at Agency 1 could be less than the maximum allocation:

$$\max\operatorname{supp}(\mathbb{S}_2^{A\sigma}) \geq \mathbf{S}_0^\sigma + d_1^{\max} - a_1^{\max} + d_2^{\max}$$

In general, for Segment i:

$$\max\operatorname{supp}(\mathbb{S}_i^{A\sigma}) \geq \mathbf{S}_0^\sigma + \sum_{\iota=1}^{i} d_\iota^{\max} - \sum_{\iota=1}^{i-1} a_\iota^{\max}$$

Let i' denote the segment in which the final donor is visited. Therefore, $\sum_{\iota=1}^{i'} d_\iota^{\max} = \sum_{D\in\mathcal{D}} d_D^{\max}$. However, $\sum_{\iota=1}^{i'-1} a_\iota^{\max} \leq \sum_{A\in\mathcal{A}} a_A^{\max}$. Consequently:

$$\sum_{\iota=1}^{i'} d_\iota^{\max} - \sum_{\iota=1}^{i'-1} a_\iota^{\max} \geq \sum_{D\in\mathcal{D}} d_D^{\max} - \sum_{A\in\mathcal{A}} a_A^{\max}$$

Since $\mathbf{Q}^\sigma < \mathbf{Q}^{lb}$, it must be that $\Pr\left\{\mathbb{S}_{i'}^{A\sigma} < \mathbf{Q}^{lb}\right\} = 1$; however, upon calculating $\max\operatorname{supp}(\mathbb{S}_{i'}^{A\sigma})$ we obtain a contradiction:

$$\max\operatorname{supp}(\mathbb{S}_{i'}^{A\sigma}) \geq \mathbf{S}_0^\sigma + \sum_{\iota=1}^{i'} d_\iota^{\max} - \sum_{\iota=1}^{i'-1} a_\iota^{\max} \geq \sum_{D\in\mathcal{D}} d_D^{\max} - \sum_{A\in\mathcal{A}} a_A^{\max} = \mathbf{Q}^{lb}$$

□

Corollary 4.2.4: *If at most ρ agencies of the set $\mathcal{A}$ may be inserted into an FRP route of the set of donors $\mathcal{D}$, let $\mathcal{A}^\rho \subset \mathcal{A}$ represent the ρ agencies with the highest values of $a_j^{\max}$. Then,*

$$\mathbf{Q}^* \geq \mathbf{Q}^{lb}(\rho) = \sum_{D \in \mathcal{D}} d_D^{\max} - \sum_{j \in \mathcal{A}^\rho} a_j^{\max}$$

Proof. The right side of Inequality (4.15) results when Theorem 4.2.3 is applied to the set of donors $\mathcal{D}$ and the set of agencies $\mathcal{A}^\rho$. No other ρ-element subset of $\mathcal{A}$ could achieve a lower value for the bound.

$\square$

Theorem 4.3.3: *For a given route of donors and agencies, if the ZMIL condition is satisfied and $R_{\mathcal{A}_{i,i+k}} = \sum_{\iota=i}^{i+k} a_\iota^{\max}$ for all sets of consecutive agencies, then the node sequence in the route obtains the minimum vehicle capacity $\mathbf{Q}^*$ possible for that set of nodes.*

Proof. We prove the claim by demonstrating that, under the maximum scenario, the lower bound on vehicle capacity provided by Theorem 4.2.3 is achieved. Since the MILB solution is known to be optimal for a given route (Theorem 4.1.4), we use it throughout the proof. In particular, due to ZMIL and the use of the MILB solution, $\mathbf{S}_0 = 0$ and $\mathbf{A}_i(s) = \min\{s, a_i^{\max}\}$.

We first demonstrate that $\mathbf{A}_i^{\max} = a_i^{\max}$ at every agency. We proceed by induction. We denote sets of consecutive agencies by $\mathcal{A}_{i_j, i_j+k_j}$. That is, $\mathcal{A}_{i_1, i_1+k_1}$ is the first set of consecutive agencies to occur in the route, $\mathcal{A}_{i_2,i_2+k_2}$ is the second set of consecutive agencies, and so on.

Initial step for the first set of consecutive agencies $\mathcal{A}_{i_1,i_1+k_1}$: For the first set of consecutive agencies, since $R_{\mathcal{A}_{i_1,i_1+k_1}} = \sum_{\iota=i_1}^{i_1+k_1} a_\iota^{\max}$, it must be that $\sum_{\iota=1}^{i_1} d_\iota^{\max} \geq \sum_{\iota=i_1}^{i_1+k_1} a_\iota^{\max}$. Therefore, applying the structure of the MILB solution under the maximum scenario, $\sum_{\iota=i_1}^{i_1+k_1} \mathbf{A}_\iota^{\max} = \min\left\{\sum_{\iota=1}^{i_1} d_\iota^{\max}, \sum_{\iota=i_1}^{i_1+k_1} a_\iota^{\max}\right\} = \sum_{\iota=i_1}^{i_1+k_1} a_\iota^{\max}$, and hence $\mathbf{A}_i^{\max} = a_i^{\max}\ \forall\ i \in \{i_1, i_1+k_1\}$.

Induction hypothesis for segments prior to i_j: As the induction hypothesis, we assume that $\mathbf{A}_i^{\max} = a_i^{\max}\ \forall\ i < i_j$.

Proof for the jth set of consecutive agencies $\mathcal{A}_{i_j,i_j+k_j}$: The condition that $R_{\mathcal{A}_{i_j,i_j+k_j}} = \sum_{\iota=i_j}^{i_j+k_j} a_\iota^{\max}$ implies that:

$$\sum_{\iota=1}^{i_j} d_\iota^{\max} - \sum_{\iota=1}^{i_j-1} \mathbf{A}_\iota^{\max} \geq \sum_{\iota=i_j}^{i_j+k_j} a_\iota^{\max}$$

Applying the induction hypothesis, this becomes:

$$\sum_{\iota=1}^{i_j} d_\iota^{\max} - \sum_{\iota=1}^{i_j-1} a_\iota^{\max} \geq \sum_{\iota=i_j}^{i_j+k_j} a_\iota^{\max}$$

Therefore, under the maximum scenario:

$$\sum_{\iota=i_j}^{i_j+k_j} \mathbf{A}_\iota^{\max} = \min\left\{\sum_{\iota=1}^{i_j} d_\iota^{\max} - \sum_{\iota=1}^{i_j-1} a_\iota^{\max}, \sum_{\iota=i_j}^{i_j+k_j} a_\iota^{\max}\right\} = \sum_{\iota=i_j}^{i_j+k_j} a_\iota^{\max}$$

Therefore, $\mathbf{A}_i^{\max} = a_i^{\max}\ \forall i \in \{i_j, i_j+k_j\}$. By induction, $\mathbf{A}_i^{\max} = a_i^{\max}$ at every agency.

Applying $\mathbf{A}_i^{\max} = a_i^{\max}\ \forall\ i$ to Equation (4.16), the condition $R_{\mathcal{A}_{i,i+k}} = \sum_{\iota=i}^{i+k} a_\iota^{\max}$ for all sets of consecutive agencies implies that:

$$\sum_{\iota=1}^{i} d_\iota^{\max} \geq \sum_{\iota=1}^{i-1} a_\iota^{\max}\ \forall\ i \tag{A.3a}$$

$$\sum_{\iota=i+1}^{n} d_\iota^{\max} \geq \sum_{\iota=i}^{n} a_\iota^{\max}\ \forall\ i \tag{A.3b}$$

Inequality (A.3a) implies that, under the maximum scenario, the total donations already collected exceed the total allocations already made at every point in the route. Inequality (A.3b) implies that, under the maximum scenario, the total donations yet to be collected exceeds the total allocations yet to be made at every point in the route. Together, they imply that $\max \operatorname{supp}(\mathbb{S}_i^A)$ is monotone nondecreasing throughout the route, so the optimal capacity is $\max \operatorname{supp}(\mathbb{S}_n^A)$. Under the maximum scenario, $\mathbb{S}_n^A = \sum_{i=1}^n d_i^{\max} - \sum_{i=1}^n a_i^{\max}$. This is equal to the lower bound on $\mathbf{Q}^*$ established by Theorem 4.2.3.

□

Lemma A.0.1. *If* $b \leq c$, *then* $\min\{a, b\} \leq \min\{a, c\}$.

Proof. Consider the case that $a \leq b$. Since $b \leq c$, it must be that $a \leq c$. Therefore, $\min\{a, b\} = a$ and $\min\{a, c\} = a$, so $\min\{a, b\} \leq \min\{a, c\}$.

Alternately, consider the case that $a > b$. Therefore, $\min\{a, b\} = b$. Since $b < a$ and $b \leq c$, it must be that $b \leq \min\{a, c\}$, so $\min\{a, b\} \leq \min\{a, c\}$. □

Theorem 5.3.2: *For a one-supersegment route, any solution that minimizes the initial load* $\mathbf{S}_0$ *is optimal.*

Proof. The expected total collection for a one-supersegment route is the sum of the expected collection at each donor:

$$C = \sum_{i=1}^{|\mathcal{D}|} \mathbf{E}[\mathbb{C}_i] \tag{A.4}$$

The collection at Donor 1 is determined only by the donation D_1, the reserved capacity q_1, and the initial load $\mathbf{S}_0$. Upon departure from Donor 1, the load cannot be more than q_1; therefore, the collection cannot exceed $q_1 - \mathbf{S}_0$:

$$\mathbb{C}_1 = \min\{D_1, q_1 - \mathbf{S}_0\} \tag{A.5}$$

Of these, only $\mathbf{S}_0$ is a decision variable; D_i is a parameter of the instance, and q_i is a function of parameters (see Definition 5.2.1). Therefore, $\mathbb{C}_1$ is completely determined by the initial load $\mathbf{S}_0$.

We now consider the collection at Donor 2. In addition to D_2, q_2, and $\mathbf{S}_0$, the collection at Donor 2 is also determined by the collection at Donor 1, which is itself determined by D_1, q_1, and $\mathbf{S}_0$:

$$\mathbb{C}_2 = \min\{D_2, q_2 - \mathbb{C}_1 - \mathbf{S}_0\} = \min\{D_2, q_2 - \min\{D_1, q_1 - \mathbf{S}_0\} - \mathbf{S}_0\} \tag{A.6}$$

Therefore, $\mathbb{C}_2$, and indeed all the $\mathbb{C}_i$ for $i \in \{1, \ldots, |\mathcal{D}|\}$, are completely determined by $\mathbf{S}_0$.

Let $\mathbb{C}_i|\mathbf{S}_0$ denote the $\mathbb{C}_i$ resulting from initial load $\mathbf{S}_0$. Let $\mathbf{S}'_0$ denote a particular initial load value greater than the minimum feasible initial load. By applying induction, we demonstrate that $\mathbb{C}_i|\mathbf{S}'_0 \preccurlyeq \mathbb{C}_i|(\mathbf{S}'_0 - 1)$ (that is, that the collection with a lower initial load stochastically dominates the collection with a higher initial load) for all $i \in \{1, \ldots, |\mathcal{D}|\}$.

Initial step for Segment 1:

$$\mathbb{C}_1|\mathbf{S}'_0 = \min\{D_1, q_1 - \mathbf{S}'_0\} \tag{A.7a}$$

$$\mathbb{C}_1|(\mathbf{S}'_0 - 1) = \min\{D_1, q_1 - (\mathbf{S}'_0 - 1)\} = \min\{D_1, q_1 - \mathbf{S}'_0 + 1\} \tag{A.7b}$$

Since $q_1 - \mathbf{S}'_0 < q_1 - \mathbf{S}'_0 + 1$, we apply Lemma A.0.1 to conclude that $\mathbb{C}_1|\mathbf{S}'_0 \preccurlyeq \mathbb{C}_1|(\mathbf{S}'_0 - 1)$. As a consequence, we can obtain a similar statement regarding the amount of food in the vehicle upon arrival to Donor 2 in terms of the initial load:

$$\begin{aligned} \mathbb{S}^D_1|\mathbf{S}'_0 &= \mathbf{S}'_0 + \mathbb{C}_1|\mathbf{S}'_0 \\ &= \mathbf{S}'_0 + \min\{D_1, q_1 - \mathbf{S}'_0\} \\ &= \min\{\mathbf{S}'_0 + D_1, q_1\} \end{aligned} \tag{A.8a}$$

$$\begin{aligned} \mathbb{S}^D_1|(\mathbf{S}'_0 - 1) &= \mathbf{S}'_0 - 1 + \mathbb{C}_1|(\mathbf{S}'_0 - 1) \\ &= \mathbf{S}'_0 - 1 + \min\{D_1, q_1 - (\mathbf{S}'_0 - 1)\} \\ &= \min\{\mathbf{S}'_0 - 1 + D_1, q_1\} \end{aligned} \tag{A.8b}$$

Since $\mathbf{S}'_0 + D_1 > \mathbf{S}'_0 - 1 + D_1$, we apply Lemma A.0.1 to conclude that $\mathbb{S}^D_1|\mathbf{S}'_0 \succcurlyeq \mathbb{S}^D_1|(\mathbf{S}'_0 - 1)$.

Induction hypothesis: As the induction hypothesis, assume that $\mathbb{S}^D_k|\mathbf{S}'_0 \succcurlyeq \mathbb{S}^D_k|(\mathbf{S}'_0 - 1)$.

The induction hypothesis implies the stochastic dominance relationship we seek to demonstrate. We compare the collection at Donor k for initial load values $\mathbf{S}_0'$ and $(\mathbf{S}_0'-1)$:

$$\mathbb{C}_k|\mathbf{S}_0' = \min\{D_k, q_k - \mathbb{S}_k^D|\mathbf{S}_0'\} \tag{A.9a}$$

$$\mathbb{C}_k|(\mathbf{S}_0' - 1) = \min\{D_k, q_k - \mathbb{S}_k^D|(\mathbf{S}_0' - 1)\} \tag{A.9b}$$

Since $q_k - \mathbb{S}_k^D|\mathbf{S}_0' < q_k - \mathbb{S}_k^D|(\mathbf{S}_0' - 1)$, we apply Lemma A.0.1 to conclude that $\mathbb{C}_k|\mathbf{S}_0' \preccurlyeq \mathbb{C}_k|(\mathbf{S}_0' - 1)$.

Proof for next segment:

We demonstrate that the induction hypothesis holds for the next segment:

$$\begin{aligned}\mathbb{S}_{k+1}^D|\mathbf{S}_0' &= \mathbb{S}_k^D|\mathbf{S}_0' + \mathbb{C}_k|\mathbf{S}_0' \\ &= \mathbb{S}_k^D|\mathbf{S}_0' + \min\{D_k, q_k - \mathbb{S}_k^D|\mathbf{S}_0'\} \\ &= \min\{D_k + \mathbb{S}_k^D|\mathbf{S}_0', q_k\}\end{aligned} \tag{A.10a}$$

$$\begin{aligned}\mathbb{S}_{k+1}^D|(\mathbf{S}_0' - 1) &= \mathbb{S}_k^D|(\mathbf{S}_0' - 1) + \mathbb{C}_k|(\mathbf{S}_0' - 1) \\ &= \mathbb{S}_k^D|(\mathbf{S}_0' - 1) + \min\{D_k, q_k - \mathbb{S}_k^D|(\mathbf{S}_0' - 1)\} \\ &= \min\{D_k + \mathbb{S}_k^D|(\mathbf{S}_0' - 1), q_k\}\end{aligned} \tag{A.10b}$$

By the induction hypothesis, $D_k + \mathbb{S}_k^D|\mathbf{S}_0' > D_k + \mathbb{S}_k^D|(\mathbf{S}_0' - 1)$. We apply Lemma A.0.1 to conclude that $\mathbb{S}_{k+1}^D|\mathbf{S}_0' \succcurlyeq \mathbb{S}_{k+1}^D|(\mathbf{S}_0' - 1)$, and hence that $\mathbb{C}_{k+1}|\mathbf{S}_0' \preccurlyeq \mathbb{C}_{k+1}|(\mathbf{S}_0' - 1)$.

Therefore, by induction, $\mathbb{C}_i|\mathbf{S}_0' \preccurlyeq \mathbb{C}_i|(\mathbf{S}_0' - 1)$ for all $i \in \{1, \ldots, |\mathcal{D}|\}$.

Therefore, $\sum_{i=1}^{|\mathcal{D}|} \mathbf{E}[\mathbb{C}_i|\mathbf{S}_0'] \leq \sum_{i=1}^{|\mathcal{D}|} \mathbf{E}[\mathbb{C}_i|(\mathbf{S}_0' - 1)]$ for all $\mathbf{S}_0'$ greater than the minimum feasible initial load.

Therefore, total collection is maximized if the initial load is at its minimum feasible value.

□

THEOREM 5.3.3: *For a one-supersegment route with a single agency, the myopic allocation policy* $\mathbf{A}^m$ *is optimal.*

Proof. To prove the claim, we demonstrate that applying the myopic allocation policy minimizes the initial load. The minimum initial load when the myopic allocation policy $\mathbf{A}^m$ is applied to a one-supersegment route with one agency, denoted $\mathbf{S}_0^m$, is:

$$\mathbf{S}_0^m = (\max\{a_n^{\min} - \sum_{i=1}^{n} d_i^{\min}, a_n^{sust} - s_n^{1-\alpha}\})^+ \tag{A.11a}$$

$$= \max\{0, a_n^{\min} - \sum_{i=1}^{n} d_i^{\min}, a_n^{sust} - s_n^{1-\alpha}\} \tag{A.11b}$$

Assume that the initial load $\mathbf{S}_0' = \mathbf{S}_0^m - z$ for some $z \in \mathbb{Z}^+$ is feasible. We consider three cases, one for each of the terms of Equation (A.11b).

Case $\mathbf{S}_0^m = 0$: : $\mathbf{S}_0^m = 0$ implies that $\mathbf{S}_0' = -z$, violating Constraints (5.1j).

Case $\mathbf{S}_0^m = a_n^{\min} - \sum_{i=1}^n d_i^{\min}$:: We compute the minimum possible load upon arrival to the agency for initial load $\mathbf{S}_0'$:

$$\begin{aligned}
\min \operatorname{supp}(\mathbb{S}_n^A) &= \min \operatorname{supp}\left(\mathbf{S}_0' + \sum_{i=1}^n D_i\right) \\
&= \mathbf{S}_0' + \sum_{i=1}^n \min \operatorname{supp}(D_i) \\
&= \mathbf{S}_0' + \sum_{i=1}^n d_i^{\min} \\
&= a_n^{\min} - \sum_{i=1}^n d_i^{\min} - z + \sum_{i=1}^n d_i^{\min} \\
&= a_n^{\min} - z
\end{aligned}$$

Therefore, the minimum amount of food available upon arrival to the agency is less than the minimum allocation, violating Constraints (5.1h).

Case $\mathbf{S}_0^m = a_n^{sust} - s_n^{1-\alpha}$:: We compute the probability of the agency receiving at least its sustaining allocation for initial load $\mathbf{S}_0'$:

$$\begin{aligned}
\Pr\{\mathbb{S}_n^A \geq a_n^{sust}\} &= \Pr\left\{\mathbf{S}_0' + \sum_{i=1}^n D_i \geq a_n^{sust}\right\} \\
&= \Pr\left\{a_n^{sust} - s_n^{1-\alpha} - z + \sum_{i=1}^n D_i \geq a_n^{sust}\right\} \\
&= \Pr\left\{\sum_{i=1}^n D_i \geq s_n^{1-\alpha} + z\right\} \\
&\leq \Pr\left\{\sum_{i=1}^n D_i > s_n^{1-\alpha}\right\} \\
&< \alpha_n
\end{aligned}$$

Therefore, the probability of the agency receiving at least its sustaining allocation is less than α_n, violating Constraints (5.1g).

Each case violates a constraint of the 1-PDA-as. Therefore, $\mathbf{S}_0'$ is infeasible; $\mathbf{S}_0^m$ must be the minimum feasible initial load. Therefore, by Theorem 5.3.2, applying the myopic allocation policy yields an optimal solution. □

COROLLARY 5.3.4: *For a one-supersegment route with a single agency, if* $\Pr\{\mathbb{S}_n^A \geq s_n^{1-\alpha}\} = \alpha_n$ *and* $s_n^{1-\alpha} = a_n^{sust}$*, then the myopic allocation policy* $\mathbf{A}^m$ *is the unique optimal minimal allocation policy.*

Proof. Assume the existence of an optimal non-myopic allocation policy $\mathbf{A}'$ with associated long-run allocation $\mathbb{A}' = \mathbf{A}'(\mathbb{S}_n^A)$. Since $\mathbf{A}'$ is not myopic, there must be some $s' \in \text{supp}(\mathbb{S}_n^A)$, $s' \geq s_n^{1-\alpha}$, such that $\Pr\{\mathbf{A}'(s') = a_n^{sust}\} < \Pr\{\mathbb{S}_n^A = s'\}$.

We compute the probability of the agency receiving the sustaining allocation:

$$\begin{aligned}
\Pr\{\mathbb{A}' \geq a_n^{sust}\} &= \sum_{s \in \text{supp}(\mathbb{S}_n^A)} \Pr\{\mathbf{A}'(s) = a_n^{sust}\} \\
&= 0 + \sum_{s \in \text{supp}(\mathbb{S}_n^A): s \geq s_n^{1-\alpha}} \Pr\{\mathbf{A}'(s) = a_n^{sust}\} \\
&\leq \Pr\{\mathbf{A}'(s') = a_n^{sust}\} + \sum_{s \in \text{supp}(\mathbb{S}_n^A): s \geq s_n^{1-\alpha}, s \neq s'} \Pr\{\mathbb{S}_n^A = s\} \\
&= \Pr\{\mathbf{A}'(s') = a_n^{sust}\} + \alpha_n - \Pr\{\mathbb{S}_n^A = s'\} \\
&< \alpha_n
\end{aligned}$$

By contradiction, such $\mathbf{A}'$ does not exist; therefore, the myopic allocation policy is the unique optimal minimal allocation policy. $\square$

THEOREM 5.3.5: *If* $(a_k^{sust} - a_k^{\min}) \geq \sum_{i=k+1}^{n}(a_i^{sust} - a_i^{\min})$ *for all* $k \in \{|\mathcal{D}|, \ldots, n\}$, *then the myopic allocation policy* $\mathbf{A}^m$ *is optimal.*

Proof. Let $\mathbf{A}'$ denote an optimal solution, which comprises a sequence of allocation policies $\mathbf{A}'_i$ with associated supply distributions $\mathbb{S}'_i$ and long-run allocations $\mathbb{A}'_i$:

$$\mathbf{A}' = \mathbf{A}'_{|\mathcal{D}|}, \mathbf{A}'_{|\mathcal{D}|+1}, \ldots, \mathbf{A}'_n$$

Let the myopic solution be denoted $\mathbf{A}^m$, which comprises a sequence of allocation policies $\mathbf{A}_i^m$ with associated supply distributions $\mathbb{S}_i^m$ and long-run allocations $\mathbb{A}_i^m$:

$$\mathbf{A}^m = \mathbf{A}^m_{|\mathcal{D}|}, \mathbf{A}^m_{|\mathcal{D}|+1}, \ldots, \mathbf{A}^m_n$$

We assume that the optimal solution is not myopic; that is, that $\mathbf{A}' \neq \mathbf{A}^m$. Therefore, there must exist some $k \in \{|\mathcal{D}|, |\mathcal{D}|+1, \ldots, n\}$ such that:

- $\forall i < k$, $\mathbf{A}'_i = \mathbf{A}_i^m$, and
- $\mathbf{A}'_k \neq \mathbf{A}_k^m$.

Without loss of generality, we assume that:

- Agency k is the first agency in the supersegment; that is, $k = |\mathcal{D}|$.
- $\mathbf{A}'$ is a minimal allocation policy; that is, $\forall i \in \{|\mathcal{D}|, \ldots, n\}$, $\Pr\{\mathbb{A}'_i = a_i^{sust}\} = \alpha$ and $\Pr\{\mathbb{A}'_i = a_i^{\min}\} = 1 - \alpha$.

The first assumption is permissible because, for all agencies prior to Agency k, $\mathbf{A}'_i = \mathbf{A}^m_i$; therefore, $\mathbb{S}'_k = \mathbb{S}^m_k$.

The second assumption is permissible because allocating more than the minimum required cannot improve the objective in a one-supersegment problem.

Within the context of this proof, we introduce a special way to describe minimal allocation policies: The *Boolean allocation policy* $\hat{\mathbf{A}}_i : [0,1] \to \{0,1\}$ maps the cumulative probability t to 1 if the sustaining allocation is provided for the associated quantile $s^t_i \in \mathrm{supp}(\mathbb{S}^A_i)$, 0 otherwise. That is,

$$\hat{\mathbf{A}}_i(t) = \begin{cases} 0 & \text{if } \mathbf{A}_i(s^t_i) = a^{\min}_i \\ 1 & \text{if } \mathbf{A}_i(s^t_i) = a^{sust}_i \end{cases} \tag{A.15}$$

The *joint Boolean allocation policy* $\hat{\mathbf{A}} : [0,1] \to \{0,1\} \times \{0,1\} \times \cdots \times \{0,1\}$ is the composition of the $\hat{\mathbf{A}}_i$. A joint Boolean allocation policy meets the requirements for providing the sustaining allocation at each agency; that is, $\forall i \in \{|\mathcal{D}|, \ldots, n\}$:

$$\int_0^1 \hat{\mathbf{A}}_i(t)dt = \alpha_i \tag{A.16}$$

We denote the joint Boolean allocation policy for the optimal allocation policy by $\hat{\mathbf{A}}'$, the composition of the $\hat{\mathbf{A}}'_i$. By the second assumption above, it must be that, $\forall i \in \{|\mathcal{D}|, \ldots, n\}$:

$$\int_0^1 \hat{\mathbf{A}}'_i(t)dt = \alpha_i \tag{A.17}$$

We denote the joint Boolean allocation policy for the myopic allocation policy by $\hat{\mathbf{A}}^m$, the composition of the $\hat{\mathbf{A}}_i^m$, which are given by:

$$\hat{\mathbf{A}}_i(t) = \begin{cases} 0 & \text{if } t \geq 1 - \alpha_i \\ 1 & \text{if } t < 1 - \alpha_i \end{cases} \tag{A.18}$$

Therefore, $\forall i \in \{|\mathcal{D}|, \ldots, n\}$:

$$\int_0^1 \hat{\mathbf{A}}_i^m(t)dt = \alpha_i \tag{A.19}$$

We define the function $\mathbf{S}$ to compute the initial load required for the joint Boolean allocation policy $\hat{\mathbf{A}}$ given $t \in [0,1]$:

$$\mathbf{S}(\hat{\mathbf{A}}(t)) = \left(\sum_{i=|\mathcal{D}|}^{n} a_i^{\min} + \hat{\mathbf{A}}_i(t) \cdot (a_i^{sust} - a_i^{\min}) - s_{|\mathcal{D}|}^t \right)^+ \tag{A.20}$$

in which $s_{|\mathcal{D}|}^t$ is the t quantile of $\mathbb{S}_{|\mathcal{D}|}^A$.

For later steps of this proof, it is convenient to decompose one term of Equation (A.20) to obtain:

$$\mathbf{S}(\hat{\mathbf{A}}(t)) = \left(\sum_{i=|\mathcal{D}|}^{n} a_i^{\min} + \hat{\mathbf{A}}_{|\mathcal{D}|}(t) \cdot (a_{|\mathcal{D}|}^{sust} - a_{|\mathcal{D}|}^{\min}) + \sum_{i=|\mathcal{D}|+1}^{n} \hat{\mathbf{A}}_i(t) \cdot (a_i^{sust} - a_i^{\min}) - s_{|\mathcal{D}|}^t \right)^+ \tag{A.21}$$

Therefore, given a joint Boolean allocation policy, the required initial load is the maximum initial load over all $t \in [0,1]$:

$$\mathbf{S}_0(\hat{\mathbf{A}}) = \max_{t \in [0,1]} \mathbf{S}(\hat{\mathbf{A}}(t)) \tag{A.22}$$

We now consider the values $t \in [0,1]$ for which $\hat{\mathbf{A}}'_{|\mathcal{D}|}$ and $\hat{\mathbf{A}}^m_{|\mathcal{D}|}$ differ. Let $T_0, T_1 \subset [0,1]$ such that, $\forall t_0 \in T_0$ and $t_1 \in T_1$:

$$\hat{\mathbf{A}}^m_{|\mathcal{D}|}(t_0) = 0 \tag{A.23a}$$

$$\hat{\mathbf{A}}'_{|\mathcal{D}|}(t_0) = 1 \tag{A.23b}$$

$$\hat{\mathbf{A}}^m_{|\mathcal{D}|}(t_1) = 1 \tag{A.23c}$$

$$\hat{\mathbf{A}}'_{|\mathcal{D}|}(t_1) = 0 \tag{A.23d}$$

Since $\hat{\mathbf{A}}^m_{|\mathcal{D}|}(t_0) = 0$, it must be that $t_0 < 1 - \alpha_{|\mathcal{D}|}\ \forall t_0 \in T_0$. Since $\hat{\mathbf{A}}^m_{|\mathcal{D}|}(t_1) = 1$, it must be that $t_1 \geq 1 - \alpha_{|\mathcal{D}|}\ \forall t_1 \in T_1$. Therefore, $t_0 < t_1$, $\forall t_0 \in T_0, t_1 \in T_1$.

Let $T \subset T_0 \times T_1$ be a one-to-one mapping from T_0 onto T_1 (and hence from T_1 onto T_0). (Such a mapping must exist because, since both $\hat{\mathbf{A}}'$ and $\hat{\mathbf{A}}^m$ are minimal allocation policies, T_0 and T_1 have the same measure.)

Having described the differences between $\hat{\mathbf{A}}'$ and $\hat{\mathbf{A}}^m$ at Agency $|\mathcal{D}|$, let us define another joint Boolean allocation policy $\hat{\mathbf{A}}^{m\prime}$. In essence, $\hat{\mathbf{A}}^{m\prime}$ is $\hat{\mathbf{A}}'$, but with the element pairs in T exchanged so that the allocation policy is myopic at Agency $|\mathcal{D}|$. As the central argument of this proof, we demonstrate that applying $\hat{\mathbf{A}}^{m\prime}$ can only improve the objective with respect to applying $\hat{\mathbf{A}}'$. A rigorous definition of $\hat{\mathbf{A}}^{m\prime}$ is:

$$\hat{\mathbf{A}}^{m\prime}(t) = \begin{cases} \hat{\mathbf{A}}'(t) = \hat{\mathbf{A}}^m(t) & \text{if } t \notin T_0 \cup T_1 \\ \hat{\mathbf{A}}'(t_1), \text{ where } (t_0, t_1) \in T & t \in T_0 \\ \hat{\mathbf{A}}'(t_0), \text{ where } (t_0, t_1) \in T & t \in T_1 \end{cases} \tag{A.24}$$

We compare $\hat{\mathbf{A}}'$ and $\hat{\mathbf{A}}^{m\prime}$ by computing $\mathbf{S}(\hat{\mathbf{A}}')$ and $\mathbf{S}(\hat{\mathbf{A}}^{m\prime})$ for any $(t_0, t_1) \in T$. Applying Equation (A.21), we compute:

$$\mathbf{S}(\hat{\mathbf{A}}'(t_0)) = \left(\sum_{i=|\mathcal{D}|}^{n} a_i^{\min} + (a_{|\mathcal{D}|}^{sust} - a_{|\mathcal{D}|}^{\min}) + \sum_{i=|\mathcal{D}|+1}^{n} \hat{\mathbf{A}}_i'(t_0) \cdot (a_i^{sust} - a_i^{\min}) - s_{|\mathcal{D}|}^{t_0} \right)^+ \tag{A.25a}$$

$$\mathbf{S}(\hat{\mathbf{A}}'(t_1)) = \left(\sum_{i=|\mathcal{D}|}^{n} a_i^{\min} + \sum_{i=|\mathcal{D}|+1}^{n} \hat{\mathbf{A}}_i'(t_1) \cdot (a_i^{sust} - a_i^{\min}) - s_{|\mathcal{D}|}^{t_1} \right)^+ \tag{A.25b}$$

$$\mathbf{S}(\hat{\mathbf{A}}^{m\prime}(t_0)) = \left(\sum_{i=|\mathcal{D}|}^{n} a_i^{\min} + \sum_{i=|\mathcal{D}|+1}^{n} \hat{\mathbf{A}}_i'(t_1) \cdot (a_i^{sust} - a_i^{\min}) - s_{|\mathcal{D}|}^{t_0} \right)^+ \tag{A.25c}$$

$$\mathbf{S}(\hat{\mathbf{A}}^{m\prime}(t_1)) = \left(\sum_{i=|\mathcal{D}|}^{n} a_i^{\min} + (a_{|\mathcal{D}|}^{sust} - a_{|\mathcal{D}|}^{\min}) + \sum_{i=|\mathcal{D}|+1}^{n} \hat{\mathbf{A}}_i'(t_0) \cdot (a_i^{sust} - a_i^{\min}) - s_{|\mathcal{D}|}^{t_1} \right)^+ \tag{A.25d}$$

By Equation (A.22), $\mathbf{S}_0(\hat{\mathbf{A}})$ is determined by the maximum value of $\mathbf{S}(\hat{\mathbf{A}}(t))$; therefore, it suffices to compare $\mathbf{S}(\hat{\mathbf{A}}'(t_0))$ with $\mathbf{S}(\hat{\mathbf{A}}^{m\prime}(t_0))$ and $\mathbf{S}(\hat{\mathbf{A}}^{m\prime}(t_1))$.

(That is: If $\mathbf{S}(\hat{\mathbf{A}}'(t_1)) < \mathbf{S}(\hat{\mathbf{A}}'(t_0))$, then $\mathbf{S}_0(\hat{\mathbf{A}}')$ is at least $\mathbf{S}(\hat{\mathbf{A}}'(t_0))$, so the value of $\mathbf{S}(\hat{\mathbf{A}}'(t_1))$ is irrelevant. If $\mathbf{S}(\hat{\mathbf{A}}'(t_1)) > \mathbf{S}(\hat{\mathbf{A}}'(t_0))$, then demonstrating that $\mathbf{S}(\hat{\mathbf{A}}^{m\prime}(t_0))$ and $\mathbf{S}(\hat{\mathbf{A}}^{m\prime}(t_1))$ are less than $\mathbf{S}(\hat{\mathbf{A}}'(t_0))$ implies that they are also less than $\mathbf{S}(\hat{\mathbf{A}}'(t_1))$.)

Equations (A.25d) and (A.25a) differ only in the final term. Since $t_0 < t_1$, it must be that $s_{|\mathcal{D}|}^{t_0} \leq s_{|\mathcal{D}|}^{t_1}$. Therefore, $\mathbf{S}(\hat{\mathbf{A}}^{m\prime}(t_0)) \leq \mathbf{S}(\hat{\mathbf{A}}'(t_0))$.

By assumption, $(a_{|\mathcal{D}|}^{sust} - a_{|\mathcal{D}|}^{\min}) \geq \sum_{i=|\mathcal{D}|+1}^{n} (a_i^{sust} - a_i^{\min})$. Therefore, $(a_{|\mathcal{D}|}^{sust} - a_{|\mathcal{D}|}^{\min}) + \sum_{i=|\mathcal{D}|+1}^{n} \hat{\mathbf{A}}_i'(t_0)(a_i^{sust} - a_i^{\min}) \geq \sum_{i=|\mathcal{D}|+1}^{n} \hat{\mathbf{A}}_i'(t_1)(a_i^{sust} - a_i^{\min})$. Therefore, $\mathbf{S}(\hat{\mathbf{A}}^{m\prime}(t_1)) \leq \mathbf{S}(\hat{\mathbf{A}}'(t_1))$.

The relationship demonstrated among $\mathbf{S}(\hat{\mathbf{A}}'(t_0))$, $\mathbf{S}(\hat{\mathbf{A}}^{m\prime}(t_0))$, and $\mathbf{S}(\hat{\mathbf{A}}^{m\prime}(t_1))$ holds for any $(t_0, t_1) \in T$. For all $t \notin T_0 \cup T_1$, $\hat{\mathbf{A}}'(t) = \hat{\mathbf{A}}^m(t)$. Therefore, $\mathbf{S}_0(\hat{\mathbf{A}}') \geq \mathbf{S}_0(\hat{\mathbf{A}}^{m\prime})$. Since $\hat{\mathbf{A}}'$ is an optimal solution, by Theorem 5.3.2, $\hat{\mathbf{A}}^{m\prime}$ must be an optimal solution.

In general, any optimal solution that is myopic prior to Agency k and is not myopic at Agency k can be used to obtain an optimal solution that is myopic at Agency k. The procedure can be applied repeatedly at all agencies for which the solution is not myopic, always obtaining another optimal solution. The final result of this process is a solution that is myopic at all agencies; therefore, the myopic solution is an optimal solution.

$\square$

APPENDIX B

Dynamic programming formulation of the 1-PDA

The state space of the dynamic program (DP) in Segment i is the load upon arrival to Agency i. In terms of the stochastic program, $S_i = \text{supp}(\mathbb{S}_i^A)$. However, we do not know the support of the $\mathbb{S}_i^A$ a priori, so we must determine minimum and maximum values for the S_i.

The minimum value of S_i is determined by the necessity to provide the minimum allocation to Agency i and to all agencies later in the route. At Segment n, only the final agency must be considered; therefore, $\min(S_n) = a_n^{\min}$. At every other segment, the minimum value of S_i can be calculated recursively:

$$\min(S_i) = a_i^{\min} + \left(\min(S_{i+1}) - d_{i+1}^{\min}\right)^+ \;\; \forall i \in \{0, \ldots, n-1\} \tag{B.1}$$

Segment 0 represents the depot, so $\min(S_0)$ is the minimum feasible initial load.

For each segment, the amount by which the load increases cannot exceed the maximum donation of that segment. Therefore, we obtain a simple upper bound on the maximum value for S_i by summing the initial load and the maximum donations:

$$\max(S_i) = \min(S_0) + \sum_{i'=1}^{i} d_{i'}^{\min} \;\; \forall i \in \{1, \ldots, n\} \tag{B.2}$$

The action set $A_i(s_i)$ from state s_i in Segment i contains the possible allocations at Agency i:

$$A_i(s_i) = \{a_i^{\min}, ..., \min\{s_i, a_i^{\max}\}\} \tag{B.3}$$

The set of available actions from a state s_i is also limited by the state space of the next segment, since $\min(S_{i+1})$ may be greater than 0. Therefore, it must be that:

$$s_i - a_i + d_{i+1}^{\min} \in S_{i+1} \tag{B.4}$$

The transition probabilities from a state s_i to a state s_{i+1} under action (allocation) a are denoted $p_i(s_{i+1}|s_i, a)$. That is, $p_i(s_{i+1}|s_i, a) = \Pr\{S_{i+1} = s_{i+1}|S_i = s_i, A_i = a\}$. They are given by:

$$p_i(s_{i+1}|s_i, a) = \Pr\{D_{i+1} = s_{i+1} - s_i + a\} \tag{B.5}$$

The contribution function of the DP is q_i, the maximum capacity required to complete the route from the current state. For the final segment of the route, q_n is simply the load; that is, $q_n(s_n) = s_n$. For prior segments, q_i is the maximum possible load in the current segment or over the course of the rest of the route. This relationship is expressed through the optimality equation for the DP:

$$q_i(s_i) = \min_{a_i^{\min} \le a_i \le \min\{s_i, a_i^{\max}\}; s_i - a_i + d_{i+1}^{\min} \in S_{i+1}} \max\left\{s_i, q_{i+1}\left(s_i - a_i + d_{i+1}^{\max}\right)\right\} \tag{B.6}$$

The solution to the 1-PDA consists of the initial load $\mathbf{S}_0$, the capacity $\mathbf{Q}$, and the allocation policy $\mathbf{A}_i(\cdot)\ \forall i \in \{1, \ldots, n\}$. In terms of our DP formulation, $\mathbf{S}_0 = \min(S_0)$, $\mathbf{Q} = q_0(\mathbf{S}_0)$, and the allocation policy is determined by the optimality equation at each state.

Although it is possible to solve the 1-PDA through dynamic programming, the MILB algorithm is simpler and more efficient. Merely defining the minima of the state space of the DP would require the same effort as the first step of the MILB method. Obtaining the allocation policy from the DP would then require solving the optimality equation at every state, in stark contrast to the few simple calculations required to obtain the MILB allocation policy (Equation (4.4)).

APPENDIX C

Maps of NIFB regions

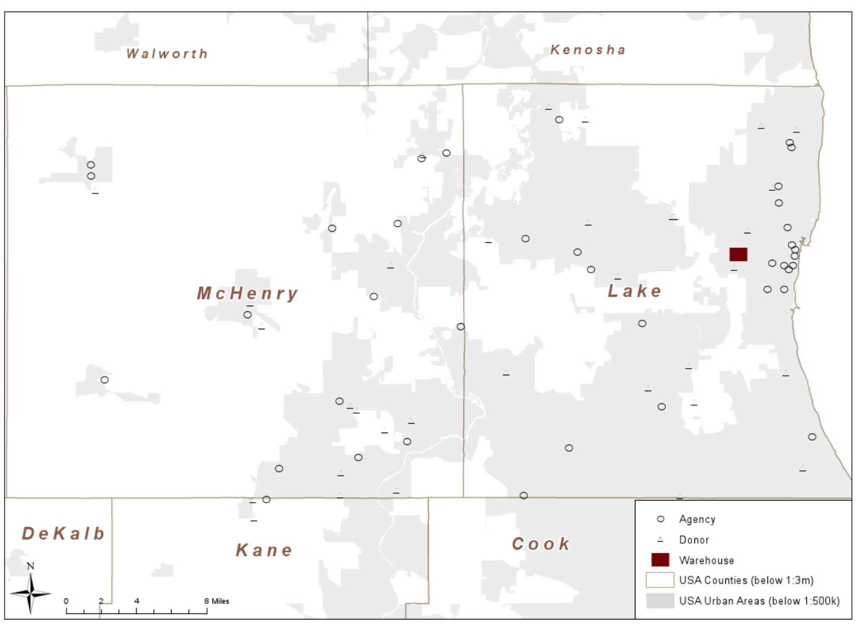

Figure C.1 North Suburban

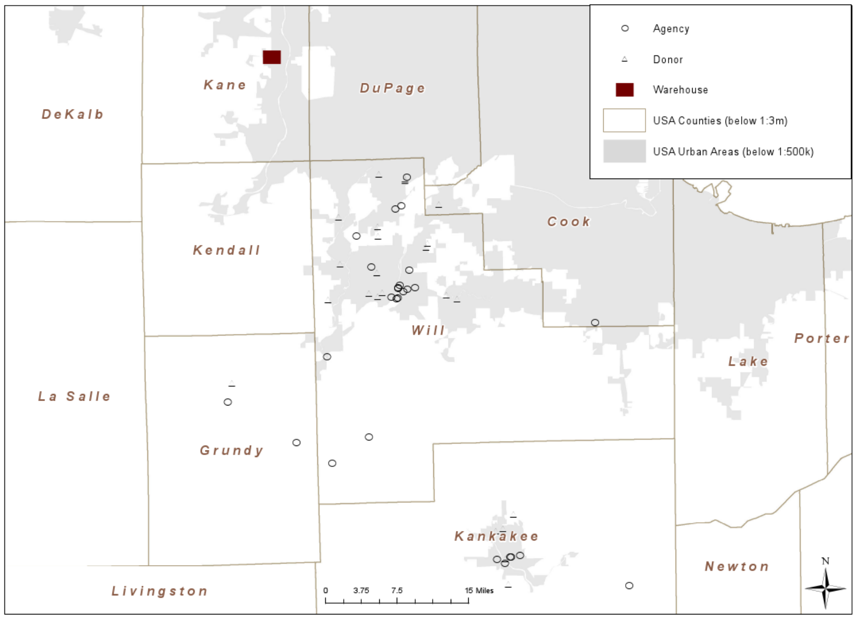

Figure C.2 South Suburban

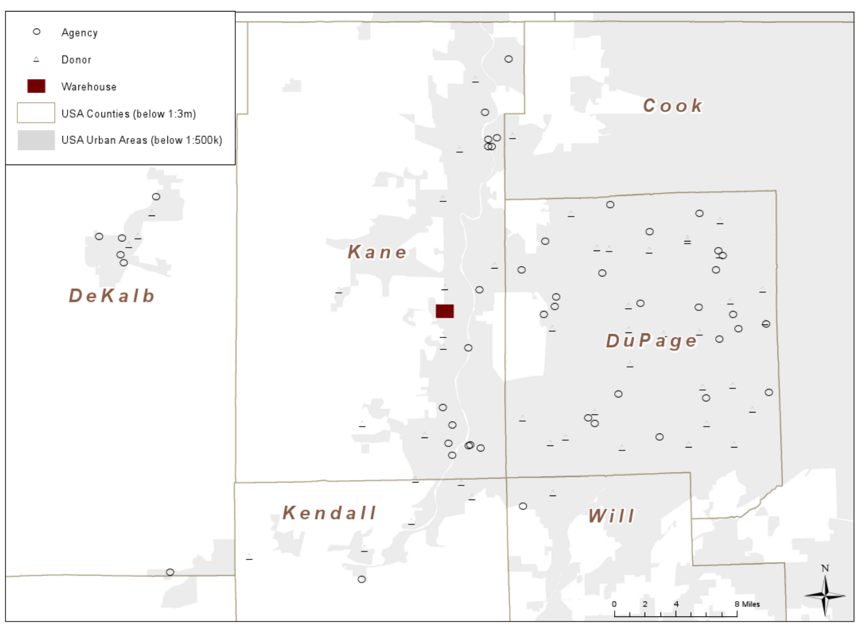

Figure C.3 West Suburban

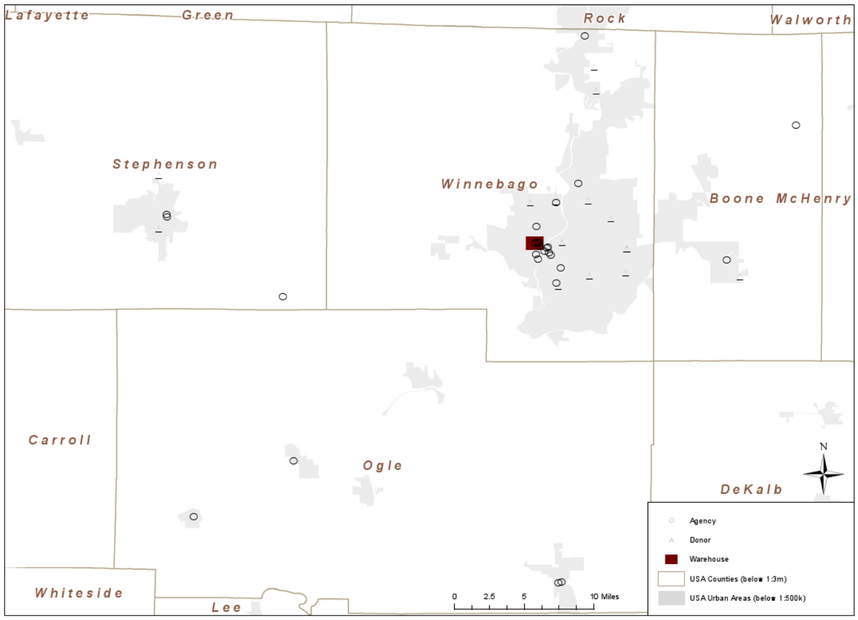

Figure C.4 Northwest

APPENDIX D

Alternate insertion algorithms

In this appendix, we describe the performance of several myopic agency insertion algorithms considered in addition to CRIH: nearest insertion (NI), insertion of the feasible agency with largest value of $a^{\max}$ ($MaxI$), and insertion of the feasible agency with the smallest value of $a^{\min}$ ($MinI$). In §D.1, we formally define the alternate agency insertion algorithms. In §D.2, we present summary statistics for the algorithms.

D.1. Description of algorithms

All of the alternate insertion algorithms restrict insertion to edges not adjacent to the depot, due to Theorem 4.2.2. We denote the tail and head of edge e, respectively, $u(e)$ and $v(e)$.

Nearest Insertion (NI). The nearest insertion algorithm chooses the agency nearest to the current route. As such, NI seeks to insert as many agencies as possible, without regard for the parameter values of those agencies.

Step 1: *Update:* Calculate C, the total travel time of the current route.

Step 2: *Calculate cost of nearest insertion for each agency:* For every agency k not on the current route, find an edge e_k not adjacent to the depot that minimizes the insertion cost $c_{ke} = t_{((u(e),k)} + t_{(k,v(e))} - t_e$, such that $c_{ke} \leq T - C$. If there are multiple such edges, choose one at random. If there are no such edges for any agency, return the current route as the NI solution; otherwise, go to Step 3.

Step 3: *Insert agency:* Let $k^* = \text{argmin}_k c_{k,e_k}$. Insert agency k^* on edge e_{k^*}. If there are multiple such agencies, insert one chosen among them at random. Go to Step 1.

Largest $a^{\max}$ Insertion ($MaxI$)

The largest $a^{\max}$ insertion algorithm chooses the agency with the largest maximum allocation that can be feasibly inserted into the current route. Like CRIH, $MaxI$ attempts to place agencies with large maximum allocation values in the route in order to reuse as much capacity as possible; unlike CRIH, $MaxI$ does not ensure that donations received before and after the insertion point are sufficient to allow capacity reuse to occur.

Step 1: *Update:* Calculate C, the total travel time of the current route.

Step 2: *Calculate cost of nearest insertion for each agency:* For every agency k not on the current route, find an edge e_k not adjacent to the depot that minimizes the insertion cost $c_{ke} = t_{((u(e),k)} + t_{(k,v(e))} - t_e$, such that $c_{ke} \leq T - C$. If there are multiple such edges, choose one at random. If there are no such edges for any agency, return the current route as the $MaxI$ solution; otherwise, go to Step 3.

Step 3: *Insert agency:* Let $k^* = \text{argmax}_k\{a_k^{\max}\}$. Insert agency k^* on edge e_{k^*}. If there are multiple such agencies, insert one with minimum insertion cost. Go to Step 1.

Smallest $a^{\min}$ Insertion ($MinI$)

The smallest $a^{\min}$ insertion algorithm chooses the agency with the smallest minimum allocation that can be feasibly inserted into the current route. As such, $MinI$ attempts to make the initial load and intermediate loads as small as possible.

Step 1: *Update:* Calculate C, the total travel time of the current route.

Step 2: *Calculate cost of nearest insertion for each agency:* For every agency k not on the current route, find an edge e_k not adjacent to the depot that minimizes the insertion cost $c_{ke} = t_{((u(e),k)}+t_{(k,v(e))}-t_e$, such that $c_{ke} \leq T-C$. If there are multiple such edges, choose one at random. If there are no such edges for any agency, return the current route as the $MinI$ solution; otherwise, go to Step 3.

Step 3: *Insert agency:* Let $k^* = \text{argmin}_k\{a_k^{\min}\}$. Insert agency k^* on edge e_{k^*}. If there are multiple such agencies, insert one with minimum insertion cost. Go to Step 1.

D.2. Performance of algorithms

As demonstrated in Tables D.1, D.2, and D.3, the optimality gaps for CRIH (see Table 4.11) are substantially smaller than those of all the other algorithms for every combination of maximum total travel time T and number of available agencies $|\mathcal{A}|$.

Table D.1. Average optimality gap of NI applied to heuristically-generated donor routes for varying maximum total travel time T and number of available agencies $|\mathcal{A}|$

	$|\mathcal{A}|$						
T	1	2	3	4	5	6	7
$T^D + 30$	6.0%	8.4%	9.4%	9.3%	9.5%	8.2%	7.5%
$T^D + 60$	1.2%	5.7%	10.7%	10.9%	15.3%	12.4%	15.8%
$T^D + 90$	0.4%	3.1%	9.0%	9.6%	18.8%	23.7%	28.9%
$T^D + 120$	0.4%	2.6%	6.5%	7.1%	22.7%	21.1%	36.0%

Table D.2. Average optimality gap of $MaxI$ applied to heuristically-generated donor routes for varying maximum total travel time T and number of available agencies $|\mathcal{A}|$

	$|\mathcal{A}|$						
T	1	2	3	4	5	6	7
$T^D + 30$	6.0%	6.5%	6.8%	4.7%	5.7%	3.2%	10.2%
$T^D + 60$	1.2%	4.2%	5.3%	4.4%	8.4%	2.7%	8.8%
$T^D + 90$	0.4%	1.9%	5.8%	4.9%	13.7%	9.0%	12.4%
$T^D + 120$	0.4%	1.9%	4.6%	4.4%	15.5%	10.4%	18.8%

Table D.3. Average optimality gap of $MinI$ applied to heuristically-generated donor routes for varying maximum total travel time T and number of available agencies $|\mathcal{A}|$

	$|\mathcal{A}|$						
T	1	2	3	4	5	6	7
$T^D + 30$	6.0%	8.6%	8.4%	7.4%	10.2%	6.6%	19.1%
$T^D + 60$	1.2%	6.8%	10.7%	11.6%	16.8%	12.9%	39.6%
$T^D + 90$	0.4%	2.3%	10.0%	12.9%	21.9%	22.1%	57.8%
$T^D + 120$	0.4%	1.9%	5.3%	7.2%	20.9%	22.4%	42.7%