Scheduling and Allocation of Uncertain, Scarce, Equitable, and Perishable Supply in Humanitarian Contexts

Irem Sengul Orgut[1,*] and Meike Reusken[2,*]

June 26th, 2026

[1] Information Systems and Supply Chain Management, Quinlan School of Business, Loyola University Chicago
[2] Operations Research and Logistics, Wageningen University & Research
[*] Co-first authorship
Contact: iorgut@luc.edu, meike.reusken@wur.nl

Abstract

Problem definition: Humanitarian organizations often distribute scarce, perishable, and uncertain supplies over a finite planning horizon. In this context, managers face decisions on assigning the receivers of supply (e.g., collecting agencies or beneficiaries) to time periods for receiving their allotments before the supply is known, and on determining how much each receives once supply is revealed. These decisions are complex due to the trade-offs among equity, freshness, and waste, which are further exacerbated by the uncertainty in supply, i.e., donations. Distributing supply quickly improves freshness and responsiveness but may increase waste or inequity, whereas delaying distribution may improve utilization and fairness but reduce freshness. Decision support is needed to improve scheduling and allocation decisions. *Methodology/results:* We introduce, study, and solve a two-stage stochastic optimization model for the scheduling and allocation decisions faced by humanitarian organizations. In the first stage, the receivers of supply are scheduled across time periods before supply is known. In the second stage, realized supply is allocated, prioritizing earlier distribution while limiting inequity in fill rates. We characterize the allocation problem under perfect equity and show that it admits a closed-form optimal solution. We also establish structural properties of the full model, including bounds on collector fill rates. Exploiting these bounds, we introduce a practical and efficient heuristic for the perfect equity case. Using numerical experiments calibrated to American food pantry operations, we find that the timing of supply arrivals plays a central role: earlier supply arrivals improve timely distribution. We also find that strict equity requirements can restrict timely distribution, while modest equity relaxations substantially improve performance. Finally, better supply information is valuable, especially when supply arrivals vary substantially over time.

Keywords: Resource allocation, scheduling, equity, humanitarian decision problems, two-stage stochastic optimization, food bank operations.

1 Introduction

Humanitarian operations, which reach millions of people worldwide through disaster response, food aid, and health programs, face constant challenges in distributing limited resources. These challenges are increasing due to shrinking surplus supply, for example from reduced food waste in more efficient food systems, and growing funding gaps (Besiou et al., 2018). In such limited resources settings, planners must often decide when to serve a demand location or beneficiary—e.g., today, tomorrow, or the day after—without yet knowing how much supply will be available on those days. Given that resources are scarce, equitable allocation becomes an important objective. This raises considerations such as whether to reserve supply to ensure fair distribution over time (e.g., equal fill rates) and how to make robust decisions that avoid systematically disadvantaging certain beneficiaries. Therefore, humanitarian organizations must constantly make decisions regarding when to serve various beneficiaries and how much of their limited resources to allocate to each beneficiary over time.

Addressing these humanitarian decision problems, the initial motivation for our paper comes from our collaborations with food assistance networks, a type of humanitarian organization that delivers food to food-insecure populations. One author is an expert on food bank operations in Europe, particularly in the Netherlands, and the other author in the United States (US). In the Netherlands, uncertain food donations are received at distribution centers. Local food agencies are, in turn, responsible for collecting their allocation of the food donations from the distribution center. The distribution center determines a time schedule (that does not change in the short term) for the food agencies to coordinate their transportation and operations to collect their allocations during their allotted days/times. The distribution centers thus determine (1) which food agencies should be scheduled on which time periods, and (2) the allocation amounts to each agency. The similar dual-decision problem is observed at the food agency level in both the Netherlands and the US. Charitable agencies are smaller, locally managed nonprofits, such as food pantries, which receive donations from various sources such as food banks and local food serving establishments, and are responsible for providing food to the food-insecure clients and families in their neighborhoods. Some agencies use appointment-based systems in which clients are assigned specific days to visit, either to pick up pre-allocated grocery boxes (pre-pack pantry model) or to select items themselves (client-choice pantry model). These systems help avoid overcrowding and improve inventory management. In such settings, pantry managers must decide when to schedule different clients and how much food to allocate to each household based on factors such as household size and level of need.

Inspired by these problems, this paper introduces and studies *the Scheduling Allocation Problem where resources are Stochastic, Equitable, and Fast-moving* (SAP-SEF). This problem is illustrated in Figure 1, where a decision maker receives supply of resources in multiple time periods, and allocates this to collectors over the same time periods. We use the term "collectors" to denote all types of entities receiving supply. This may refer to individuals as well as organizations, irrespective of how resources are transferred from the decision maker to the collector. The term therefore covers a broader context than its wording may directly suggest. The SAP-SEF consists of two distinct decisions. First, a time-scheduling decision in which each collector is assigned to exactly one time period, and second, an allocation decision in which items are distributed equitably and as quickly as possible. There are three main problem characteristics concerning

the resources: (*i*) Resources are scarce, that is, supply is not sufficient to meet demand, which motivates the desire to allocate equitably. (*ii*) There is an interest in moving resources as quickly as possible, e.g., due to perishability, maintaining low stock levels, or, in the case of humanitarian crises, prioritizing helping people as quickly as possible. (*iii*) Supply is stochastic as each period's supply is not known when scheduling decisions are made. Schedules are considered strategic decisions and, once set, cannot easily be changed from one horizon to the next, e.g., to respect organizations' internal operational planning or individuals' personal, work, and school schedules. The supply is revealed at the beginning of the time horizon, after which allocation decisions are made. The need to time-schedule collectors before supply is known, together with the goal of allocating incoming uncertain supply equitably and quickly once it is revealed, creates a central challenge for humanitarian operations: *"How can humanitarian decision makers schedule collectors over time before supply is known and allocate realized supply in a way that balances equitable and timely distribution?"*

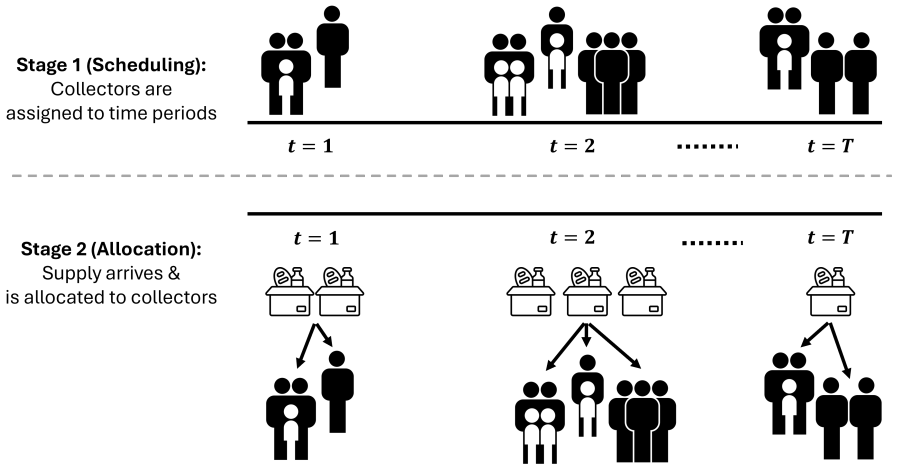

Figure 1: Illustration of the SAP-SEF decision structure.

Although our study is originally motivated by food assistance operations and humanitarian contexts, similar problems arise in many other settings where fast-moving/perishable, stochastic, and scarce supplies become available over time, and the decision maker must determine both which entities collect the supply at different time periods and how much each receives. Table 1 provides examples of such applications, specifying the supplied resource, the entity acting as the decision-maker, and the entity taking the role of the collector.

Supply	Decision-maker	Collector
Aid donations	Disaster relief sites	Disaster victims
COVID-19 vaccines	Health centers	Patients
Food donations	Distribution centers (the Netherlands context)	Food bank agencies
Food donations	Food bank agencies	Food-insecure clients
Blood donations	Blood banks	Hospitals and patients

Table 1: Example applications of the SAP-SEF.

In practice, humanitarian managers may rely on heuristic approaches to schedule beneficiaries or may not use an appointment system at all. This can lead to overcrowding in some periods and underutilization in others, resulting in both inequities and operational inefficiencies. We will now give a toy example to illustrate the main characteristics and trade-offs of the SAP-SEF.

Toy example. *Consider three time periods and four collectors with demands* $\boldsymbol{d} = [50, 80, 80, 70]$ *lbs. Initially, collectors must be assigned to time periods before supply* ξ_t *for each time period* t *is known. After collectors have been scheduled, the supply for each period in the planning horizon is revealed. Assume that the revealed supply is as follows:* $\hat{\boldsymbol{\xi}} = [40, 90, 50]$ *lbs. The realized supply is then distributed among collectors, following the pre-set schedule, to maximize total allocation, favor faster-moving items (e.g., due to perishability), and maintain equity among collectors. Table 2 compares three heuristic approaches that a humanitarian organization may use. Approach 1 schedules all collectors on the first day, prioritizing immediate access to aid. Approach 2 schedules all collectors on the last day, delaying distribution until all supplies over the planning horizon have been received. Approach 3 does not use a fixed schedule and allows collectors to arrive on their preferred day, subject to the restriction that each collector may visit only once. For illustration, we assume the following arrival pattern under Approach 3: collectors 1 and 2 arrive on the first day, while collectors 3 and 4 arrive on the last day. Table 2 reports the approach name and associated fill rates in the first two columns, followed by columns reporting three key performance metrics:* (i) *total waste, defined as the total supply that remains undistributed by the end of the planning horizon;* (ii) *an envy metric, which measures equity as the absolute difference between the maximum and minimum fill rates across collectors; and* (iii) *a freshness metric, defined as the average age of distributed items at the time of distribution, measured as the number of days between supply receipt and distribution. Calculations assume that the available supply is distributed equally among the collectors that visit on a given day.*

Table 2: Comparison of three heuristic approaches in the toy example.

Approach	Fill rates	Waste metric	Envy metric	Freshness metric
1: All first day	$\{\frac{40}{280}, \frac{40}{280}, \frac{40}{280}, \frac{40}{280}\}$	90 + 50 = 140 lbs ↓	0 ↑	0 days ↑
2: All last day	$\{\frac{180}{280}, \frac{180}{280}, \frac{180}{280}, \frac{180}{280}\}$	0 lbs ↑	0 ↑	$\frac{40\times2+90\times1+50\times0}{180} = 0.94$ days ↓
3: Random	$\{\frac{40}{130}, \frac{40}{130}, \frac{140}{150}, \frac{140}{150}\}$	0 lbs ↑	$\frac{140}{150} - \frac{40}{130} = 0.62$ ↓	$\frac{40\times0+90\times1+50\times0}{180} = 0.5$ days ↓

Note: Upward arrows denote relatively favorable outcomes and downward arrows denote unfavorable outcomes.

Table 2 highlights the trade-offs among waste, equity, and freshness. Approach 1 achieves perfect equity and distributes supplies immediately, but results in high waste because supplies received on days 2 and 3 are never distributed. Approach 2 eliminates waste and achieves perfect equity, but performs poorly in terms of freshness because supplies must be stored before distribution. Finally, the random approach eliminates waste, but performs poorly on equity because collectors' fill rates depend on when they arrive. It also leads to lower freshness, as some supplies must be stored before distribution. This toy example illustrates the inherent tension among these metrics, which becomes even more challenging in practice when supply is uncertain. We address these challenging decision problems in this paper.

We summarize the contributions of this paper. First, we introduce the SAP-SEF, a two-stage stochastic

optimization problem that schedules collectors over multiple time periods under supply uncertainty and then subsequently allocates realized supply while preserving equity in fill rates. Specifically, the objective of our model is to distribute as much supply as possible, as quickly as possible, subject to an envy-based equity constraint that limits deviations among collectors' fill rates. We derive a closed-form solution for the second-stage allocation problem under perfect equity. Furthermore, we establish several analytical properties of the SAP-SEF that provide structural insights into the problem and its optimal solutions. These results yield easily computable performance bounds for any equity level, and, based on the derived upper bounds, we develop a practical and computationally efficient heuristic for the perfect-equity case. Finally, we conduct a case study using pseudo-realistic data to evaluate the proposed model and heuristic. The results demonstrate their performance and provide managerial insights into the trade-offs between waste, equity, and freshness, and the impact of the degree of supply uncertainty.

The remainder of this research is organized as follows. Section 2 reviews related literature. In Section 3, we formally introduce a two-stage stochastic optimization model for the SAP-SEF. Section 4 examines the perfect equity variant of the SAP-SEP by deriving a closed-form solution. Section 5 presents structural properties of the proposed model. Building on some of the derived properties, Section 6 proposed a practical heuristic for the SAP-SEF under perfect equity. Section 7 describes the numerical experiments and discusses the results. Finally, Section 8 summarizes the managerial insights and concludes the paper.

2 Literature

This study falls within the broader field of humanitarian operations. Humanitarian organizations increasingly operate under conditions of resource scarcity, supply volatility, and uncertainty, which has motivated a substantial body of Operations Research (OR) and Operations Management (OM) work focused on decision support for humanitarian systems (Besiou et al., 2018; Swaminathan, 2025; Berenguer Shen, 2019). Within this stream, mathematical modeling and optimization have been widely used to support key operational decisions, including location planning, resource allocation, routing, and network design (Petropoulos et al., 2024).

While OR/OM has generated valuable insights and decision-support mechanisms for humanitarian practice (e.g., Peters et al., 2021; Reusken et al., 2023), concerns remain regarding the practical relevance and implementation of many proposed methods (Besiou van Wassenhove, 2020). Although the complex and pressing characteristics of humanitarian decision problems have been studied extensively (Vries Wassenhove, 2020), research outcomes often fail to translate into insights that are readily understandable, implementable, and accessible to practitioners (Besiou van Wassenhove, 2020). Surveys of OR/OM for humanitarian operations therefore emphasize the need for closer engagement with practice, simpler and more implementable solution methods, and clearer communication of actionable insights (Besiou van Wassenhove, 2020; Vries Wassenhove, 2020).

Our study responds directly to these calls by combining methodological development with practical relevance. First, the authors are connected to practitioners in humanitarian and food bank contexts, providing access to real-life operational challenges as well as a feedback loop for validating, refining, and returning

model outputs. Second, beyond its methodological contributions, the paper generates several practically actionable insights, including a comparison of the studied problem in the European and North American food-bank contexts (Section 1), analytical structural properties that can be derived without solving the optimization model (Section 5), a simple yet effective heuristic for operational decision making (Section 6), and results from solving our model with pseudo-realistic data (Section 7), along with the resulting managerial insights (Section 8). Third, to facilitate implementation and impact, all code will be made publicly available.

Our paper is connected to the literature on equitable resource allocation. Decision problems in which, like in the present paper, the resources are scarce and need to be equitably allocated typically arise in the humanitarian context in food assistance operations and disaster relief problems. In the food assistance context, the problem considers the allocation of a limited amount of donated food across multiple service regions or food agencies served by a food bank, while ensuring that allocations remain sufficiently fair (e.g., Sengul Orgut et al., 2016a; Lien et al., 2014; Sinclair et al., 2022). Within this stream, some studies do not consider uncertainty at all (e.g., Sengul Orgut et al., 2016b; Gómez-Pantoja et al., 2021; Sengul Orgut Lodree, 2023), others consider uncertainty in demand (e.g., Lien et al., 2014; Ma et al., 2023), in donations (Fianu Davis, 2018; Alkaabneh et al., 2021), in capacities of the receiving agencies (Sengul Orgut et al., 2017, 2018), in capacity and supply (Stauffer et al., 2022), and in the number of customers arriving each day (Sinclair et al., 2022). Similarly, and considering the same resource characteristics of scarcity and equity combined with uncertainty, but in the context of aid assistance in disaster response, Azizi et al. (2021) studies how a limited amount of humanitarian aid should be allocated across multiple refugee camps facing uncertain demand. In this setting, resources are scarce, and equity is incorporated through a priority structure that gives preference to camp-based refugees, rather than through an explicit fairness metric. Not focused on a particular application, but still considering the same resource characteristics (scarcity, equity, and uncertainty) from a more methodological perspective, is Hassanzadeh et al. (2023), who formulate a generic problem of allocating a limited resource to agents arriving over time under uncertain demand and use Nash Social Welfare as the fairness criterion. Related insights are provided by Gómez-Pantoja et al. (2021) and Tirkolaee et al. (2020), although these studies do not explicitly incorporate the equity characteristic. Moreover, none of these studies considers the problem of scheduling collectors or beneficiaries over time under supply uncertainty to generate equitable allocations.

The above shows that resource allocation models have been introduced for settings characterized by scarcity, equity, and uncertainty, with different combinations of these features considered across the literature. However, supply uncertainty remains relatively underexplored. We examine the three studies that explicitly consider supply uncertainty in this context. Fianu Davis (2018) develop a discrete-time, discrete-space Markov decision process model to determine optimal food donation allocation policies from food banks to counties under stochastic supply. Their model seeks to maximize equity across the network, where demand is measured based on the pounds of food needed per person in poverty in each county. Specifically, they evaluate three inventory policies commonly used in the inventory literature and identify the best-performing policy. Stauffer et al. (2022) also consider supply uncertainty in food bank operations, examining three strategies for improving food allocations: increasing food bank storage capacity, increasing agency storage capacity, and deploying mobile pantries to mitigate local capacity shortages. They formulate a two-stage

stochastic model that incorporates uncertainty in supply and agency capacity and minimizes food holding costs at warehouses and mobile pantry operating costs. While their model distinguishes among food items with different storage requirements, it represents perishability through fixed expiration dates rather than continuous deterioration. Further, their study focuses on strategic investment decisions, whereas our work combines strategic and tactical decisions related to scheduling and serving beneficiaries over time. Finally, Alkaabneh et al. (2021) model supply uncertainty through a multi-period Markov Decision Process in which food allocations are repeatedly updated as random donations arrive and inventories evolve over time, while our paper models supply uncertainty using a two-stage stochastic programming framework, where first-stage allocation decisions are made before the supply realization is known and second-stage recourse decisions are taken once the supply is revealed. Furthermore, in addition to scarcity, equity, and uncertainty, we explicitly consider perishability. In contrast to Alkaabneh et al. (2021), where perishability is implicitly captured through the evolving inventory state and dynamic allocation decisions over time, our model incorporates perishability directly through an objective-function penalty on carry-over inventory, thereby capturing the loss of value associated with delayed distribution. Thereby, our work contributes to the resource allocation literature by jointly modeling four key characteristics of resources: scarcity, equity, uncertainty, and perishability.

In addition, we combine resource allocation decisions with a strategic scheduling component over time. The studies that come closest to this idea are Lien et al. (2014) and Tirkolaee et al. (2020). Lien et al. (2014) determine the customer visitation sequence in addition to allocating resources. Specifically, the model sequentially decides which customer to serve next and how much resource to allocate after observing that customer's realized demand, while continuously updating the remaining inventory and fairness measure before proceeding to the next customer. In our setting, however, scheduling refers to the assignment of receiving entities to predefined time periods such as days or weeks, whereas Lien et al. (2014) focus on sequencing the order in which customers are served. Similarly, Tirkolaee et al. (2020) formulate a robust bi-objective disaster-response optimization model that allocates scarce rescue units to affected areas and schedules rescue operations over time. As in Lien et al. (2014), the model focuses on determining a sequence of operations, in this case the sequence of rescue tasks. Overall, these studies highlight the importance of integrating allocation decisions with temporal structure, but differ in the nature of the scheduling component. While these studies address operational routing decisions, such as determining which site to visit next, our paper focuses on a fundamentally different problem: developing a strategic schedule for collectors or beneficiaries over multiple time periods and then determining allocation quantities that preserve equity, while maximizing total allocation and freshness.

3 Two-stage stochastic optimization for the SAP-SEF

In this section, we introduce a two-stage stochastic optimization model for the Scheduling Allocation Problem where resources are Stochastic, Equitable, and Fast-moving (SAP-SEF). We first introduce notations:

$\mathcal{N}$	Set	Collectors $\mathcal{N} \equiv \{1, \ldots, n, \ldots, N\}$
$\mathcal{T}$	Set	Time periods $\mathcal{T} \equiv \{1, \ldots, t, \ldots, T\}$
D_n	Parameter	Demand of collector n, for all $n \in \mathcal{N}$
ξ_t	Random variable	Stochastic supply in time period t (nonnegative), for all $t \in \mathcal{T}$
$\hat{\xi}_t$	Realization	Realized supply of time period t (nonnegative), for all $t \in \mathcal{T}$
X_{nt}	Decision variable (First stage)	Binary; equals 1 if collector n is assigned to time period t, 0 otherwise, for all $n \in \mathcal{N}, t \in \mathcal{T}$
Y_{nt}	Decision variable (Second stage)	Allocation to collector n in time period t, for all $n \in \mathcal{N}, t \in \mathcal{T}$.

We use $\boldsymbol{\xi}$ to denote the vector of random variables, $\hat{\boldsymbol{\xi}}$ its realization, and $\mathbf{X}$ the matrix $(X_{nt})_{N \times T}$. In this problem, X_{nt} are the here-and-now decisions, fixed before uncertainty is known. Y_{nt} are the wait-and-see decisions, functions of the realizations $\hat{\xi}_t$. The following two assumptions are made. First, we assume that we operate in a supply scarcity environment where the realized total supply is less than the total demand, i.e., $\sum_{t \in \mathcal{T}} \hat{\xi}_t < \sum_{n \in \mathcal{N}} D_n$ (hereafter **Assumption 1**). Second, we assume that at the beginning of a week, after collection schedules have been determined, the decision-maker obtains complete knowledge of the donations that will occur during that week. Lastly, we assume that a positive supply amount is received in the first period of the planning horizon. This is without loss of generality, since any initial period with zero supply can be excluded from the planning horizon.

The first stage determines the collection schedules X_{nt} for all collectors n and all time periods t, maximizing the expected objective function from the second stage. The first-stage model can be represented as follows:

$$\max \quad \mathbb{E}_{\boldsymbol{\xi}}[Q(\mathbf{X}, \boldsymbol{\xi})] \tag{1}$$

$$\text{s.t.} \quad \sum_{t=1}^{T} X_{nt} = 1 \qquad n \in \mathcal{N} \tag{2}$$

$$X_{nt} \in \{0, 1\} \qquad n \in \mathcal{N}, t \in \mathcal{T}, \tag{3}$$

where constraint (2) ensures that each collector is scheduled once during the time horizon, reflecting the real-world operations observed in many humanitarian settings, and constraint (3) ensures that the first-stage decisions are binary.

After the collection schedules have been determined, the supply for the time horizon is observed. This could be considered as, e.g., donors providing their donation dates and quantities at the beginning of each week. The second-stage then determines the allocation amounts Y_{nt} for all collectors n and all time periods t, and targets the minimization of leftover resources at the end of each time period, where leftover resources at the end of period t is defined as $\sum_{i=1}^{t} \hat{\xi}_t - \sum_{i=1}^{t} \sum_{n=1}^{N} Y_{nt}$. Therefore, the total leftovers over the full time horizon is equal to:

$$\sum_{t=1}^{T} (\sum_{i=1}^{t} \hat{\xi}_t - \sum_{i=1}^{t} \sum_{n=1}^{N} Y_{ni}) = \sum_{t=1}^{T} \sum_{i=1}^{t} \hat{\xi}_t - \sum_{t=1}^{T} \sum_{i=1}^{t} \sum_{n=1}^{N} Y_{ni}. \tag{4}$$

Since the first term in (4) is a constant, minimizing (4) is equivalent to the following second-stage recourse

function:

$$Q(\mathbf{X}, \hat{\boldsymbol{\xi}}) = \max \sum_{t=1}^{T} \sum_{i=1}^{t} \sum_{n=1}^{N} Y_{ni}, \tag{5}$$

which maximizes the time-weighted amount of resources distributed. Note that $\sum_{t=1}^{T} \sum_{i=1}^{t} \sum_{n=1}^{N} Y_{ni}$ can be written as:

$$\begin{aligned} &(\sum_{n=1}^{N} Y_{n1}) + (\sum_{n=1}^{N} Y_{n1} + \sum_{n=1}^{N} Y_{n2}) + \cdots + \\ &\quad (\sum_{n=1}^{N} Y_{n1} + \sum_{n=1}^{N} Y_{n2} + \cdots + \sum_{n=1}^{N} Y_{nT}) && (6) \\ &= T \sum_{n=1}^{N} Y_{n1} + (T-1) \sum_{n=1}^{N} Y_{n2} + \cdots + \sum_{n=1}^{N} Y_{nt} && (7) \\ &= \sum_{t=1}^{T} (T - t + 1) \sum_{n=1}^{N} Y_{nt}, && (8) \end{aligned}$$

where (7) and (8) show that the weight assigned to the total allocation in time period t decreases as t increases. In other words, the coefficient for $\sum_{n=1}^{N} Y_{n1}$ is larger than that for $\sum_{n=1}^{N} Y_{n2}$, and so on. Since the objective function is maximized, this structure incentivizes earlier distribution rather than delaying it. Thus, distributing food in the period it is received yields a higher reward compared to postponing it to later periods.

The constraints of the second-stage problem are given as follows:

$$\left| \frac{\sum_{t=1}^{T} Y_{nt}}{D_n} - \frac{\sum_{t=1}^{T} Y_{mt}}{D_m} \right| \leq \theta \qquad n, m \in \mathcal{N} \tag{9}$$

$$D_n X_{nt} \geq Y_{nt} \qquad n \in \mathcal{N}, t \in \mathcal{T} \tag{10}$$

$$\sum_{i=1}^{t} \hat{\xi}_i \geq \sum_{i=1}^{t} \sum_{n=1}^{N} Y_{ni} \qquad t \in \mathcal{T} \tag{11}$$

$$Y_{nt} \in \mathbb{R}_+ \qquad n \in \mathcal{N}, t \in \mathcal{T}. \tag{12}$$

Constraint (9) is the equity constraint and is defined as an envy measure where the deviation between the absolute pairwise differences between each collector's fill rate (the amount of food allocated over the demand they serve) must be less than or equal to a parameter θ. Constraint (10) ensures that (i) allocation to collector n on time period t is positive if and only if collector n is scheduled for time period t, and (ii) that allocation to each collector does not exceed their demand. Constraint (11) limits the allocation in each period to the available supply. Finally, (12) ensures the nonnegativity of the second-stage allocation decision variables.

Since constraint (9) is nonlinear in its current form due to the absolute value, we can rewrite it using standard linearization reformulation as follows. Let r^{min} and r^{max} be auxiliary variables representing the minimum and maximum fill rates of collectors, respectively. We can then replace (9) with the linear constraints: $r^{max} \geq \frac{\sum_{t=1}^{T} Y_{nt}}{D_n}$ and $r^{min} \leq \frac{\sum_{t=1}^{T} Y_{nt}}{D_n}$ for $n \in \mathcal{N}$ and $r^{max} - r^{min} \leq \theta, r^{max}, r^{min} \in \mathbb{R}_+$.

Therefore, the extensive-form formulation for the SAP-SEF is as follows:

[SAP-SEF]

$$
\begin{aligned}
\max \quad & \mathbb{E}_{\boldsymbol{\xi}[\sum_{t=1}^{T}(T-t+1)\sum_{n=1}^{N} Y_{nt}]} & & \\
\text{s.t.} \quad & \sum_{t=1}^{T} X_{nt} = 1 & & n \in \mathcal{N} \\
& r^{max} \geq \frac{\sum_{t=1}^{T} Y_{nt}}{D_n} & & n \in \mathcal{N} \\
& r^{min} \leq \frac{\sum_{t=1}^{T} Y_{nt}}{D_n} & & n \in \mathcal{N} \\
& r^{max} - r^{min} \leq \theta & & \\
& D_n X_{nt} \geq Y_{nt} & & n \in \mathcal{N}, t \in \mathcal{T} \\
& \sum_{i=1}^{t} \xi_i \geq \sum_{i=1}^{t} \sum_{n=1}^{N} Y_{ni} & & t \in \mathcal{T} \\
& X_{nt} \in \{0,1\} & & n \in \mathcal{N}, t \in \mathcal{T} \\
& Y_{nt}, r^{max}, r^{min} \in \mathbb{R}_+ & & n \in \mathcal{N}, t \in \mathcal{T},
\end{aligned}
$$

which is a two-stage stochastic mixed-integer linear program containing NT binary variables and $NT + 2$ continuous variables.

The SAP-SEF is NP-hard since it contains the 0-1 Knapsack Problem as a special case. Consider a restricted version of the model with a single time period ($T = 1$) and without the equity constraints (or equivalently with $\theta \geq 1$). In this case, the model reduces to selecting binary allocation decisions subject to a single supply capacity constraint, which is equivalent to the 0-1 Knapsack Problem, known to be NP-hard.

Furthermore, SAP-SEF has complete recourse. This is straightforward to verify from the existence of a trivially feasible solution. In particular, setting $X_{nT} = 1$ for all $n \in \mathcal{N}$, $Y_{nt} = 0$ for all $n \in \mathcal{N}$ and $t \in \mathcal{T}$, $r^{\max} = 0$, and $r^{\min} = 0$ satisfies all constraints.

4 Special case: Perfect equity ($\theta = 0$)

Consider a special case of the SAP-SEF presented in Section 3 where $\theta = 0$, which represents perfect equity. Setting $\theta = 0$ implies that the fill rate of all collectors is exactly equal and hence represents an ideal case from an equity perspective. Let us focus on the second-stage problem, in which supply has been realized and the collection schedules have been fixed. We use the following definitions:

- $\tilde{Y}_n = \sum_{t=1}^{T} Y_{nt}$ represents the allocation to collector n over the total time horizon.
- $\tilde{\tilde{Y}} = \sum_{t=1}^{T} \sum_{n=1}^{N} Y_{nt}$ represents the total allocation.
- $\Delta = \sum_{n=1}^{N} D_n$ represents the total demand.
- $\{\mathcal{N}\}_t$ represent the subset of collectors to be served on time period t. Note two properties: (i) $\cup_{t\in\mathcal{T}}\{\mathcal{N}\}_t \equiv \mathcal{N}$, and (ii) $\{\mathcal{N}\}_t \cap \{\mathcal{N}\}_{t'} \equiv$ for all $t, t' \in \mathcal{T}$. These follow from constraints (2) and (3).
- $r_n = \frac{\tilde{Y}_n}{D_n}$ represents the fraction of demand satisfied for collector n, hereafter the fill rate.

Using these definitions, Lemma 1 shows that several natural second-stage objectives become equivalent, and Proposition 1 provides the closed-form optimal solution for the second-stage problem.

Lemma 1. *In the case of perfect equity ($\theta = 0$), the objective function for the second-stage problem in* (5) *is equivalent, with respect to optimality, to both (i) maximizing the total distribution, i.e., $\tilde{\tilde{Y}}$, and (ii) maximizing any fill rate, i.e., r_n for any $n \in \mathcal{N}$. That is, these objectives have the same optimal solution set.*

Proof of Lemma 1. Consider the equity constraint (9) under the perfect equity assumption:

$$\frac{\sum_{t=1}^{T} Y_{nt}}{D_n} - \frac{\sum_{t=1}^{T} Y_{mt}}{D_m} = 0 \quad n, m \in \mathcal{N},$$

which can be rewritten as $r_n = \frac{\tilde{Y}_n}{D_n} = \frac{\tilde{Y}_m}{D_m}$ for $n, m \in \mathcal{N}$. Setting this equal to a constant A, we get $r_n = \frac{\tilde{Y}_n}{D_n} = \frac{\tilde{Y}_m}{D_m} = A$ for $n, m \in \mathcal{N}$. Hence,

$$\tilde{Y}_n = AD_n \quad n \in \mathcal{N}. \tag{13}$$

Summing over all n, gives

$$\sum_{n=1}^{N} \tilde{Y}_n = \tilde{\tilde{Y}} = A \sum_{n=1}^{N} D_n = A\Delta, \tag{14}$$

and rewriting results in $A = \frac{\tilde{\tilde{Y}}}{\Delta}$, which, using (13), implies that

$$\tilde{Y}_n = D_n \frac{\tilde{\tilde{Y}}}{\Delta} \quad n \in \mathcal{N}. \tag{15}$$

Now, note that in the second-stage model, the first-stage decisions $\mathbf{X}$ are known. Further, due to constraint (2), for a given collector n, X_{nt} will be equal to one for only one time period t. This implies that due to (10), for a given collector n, Y_{nt} can be positive for only one period t and will be equal to zero for all other periods. Assume that for a given collector n, the period that they are assigned to receive donations is represented as t^*, i.e., $X_{nt^*} = 1$ and $Y_{nt^*} \geq 0$. Further, $X_{nt} = 0$ and $Y_{nt} = 0$ for all $t \neq t^*$. Considering this, we can write $Y_{nt} = \tilde{Y}_n X_{nt}$ for $n \in \mathcal{N}, t \in \mathcal{T}$. Then, using (8), we can write the objective function (5) as follows:

$$\sum_{t=1}^{T} (T - t + 1) \sum_{n=1}^{N} Y_{nt} = \sum_{t=1}^{T} (T - t + 1) \sum_{n=1}^{N} \tilde{Y}_n X_{nt}. \tag{16}$$

Now, inserting (15), this is equal to

$$\begin{aligned} &\sum_{t=1}^{T} (T - t + 1) \sum_{n=1}^{N} \frac{\tilde{\tilde{Y}}}{\Delta} D_n X_{nt} \\ &= \frac{\tilde{\tilde{Y}}}{\Delta} \sum_{t=1}^{T} (T - t + 1) \sum_{n=1}^{N} D_n X_{nt}. \end{aligned} \tag{17}$$

Since $X_{nt^*} = 1$ and $X_{nt} = 0$ for all $t \neq t^*$, and using the definition of Δ, this is equal to:

$$\frac{\tilde{\tilde{Y}}}{\Delta} \sum_{t=1}^{T} (T - t + 1)\Delta = \tilde{\tilde{Y}} \frac{T(T+1)}{2}. \tag{18}$$

This shows that the objective function (5) is equivalent to the last term in (18), which equals $\tilde{\tilde{Y}}$ multiplied by a positive constant (since $T > 0$). This positive constant only scales the objective values. Hence, maximizing $\tilde{\tilde{Y}} \frac{T(T+1)}{2}$ is equivalent to maximizing $\tilde{\tilde{Y}}$ in terms of which solution is optimal. This concludes the first part of the proof. Further, following from (14) and the definition of r_n, the last term in (18) is also equal to $\tilde{\tilde{Y}} \frac{T(T+1)}{2} = A\Delta \frac{T(T-1)}{2} = r_n \Delta \frac{T(T-1)}{2}$. Since $\Delta \frac{T(T-1)}{2}$ represents a positive constant, maximizing $\tilde{\tilde{Y}} \frac{T(T+1)}{2}$ is

also equivalent to maximizing r_n for any $n \in \mathcal{N}$. This completes the proof.

Proposition 1. *The optimal solution for the second-stage problem under perfect equity ($\theta = 0$) is represented as follows:*

$$\tilde{\tilde{Y}}^* = \Delta \min_{t \in \mathcal{T}} \frac{\sum_{i=1}^{t} \hat{\xi}_i}{\sum_{n \in \cup_{i=1}^{t} \{\mathcal{N}\}_i} D_n}. \tag{19}$$

Further, we also have:

$$Y_{nt}^* = \begin{cases} \tilde{\tilde{Y}}^* \frac{D_n}{\Delta} = D_n \min_{t \in \mathcal{T}} \frac{\sum_{i=1}^{t} \hat{\xi}_i}{\sum_{n \in \cup_{i=1}^{t} \{\mathcal{N}\}_i} D_n} & \text{if} \quad X_{nt} = 1, \\ 0 & \text{otherwise.} \end{cases} \tag{20}$$

Notice that the value indicated by the minimum function in (19) and (20) represents a random variable that has been realized in the second stage.

Definition 1. *We define the worst cumulative fill rate as the 'critical ratio', given by:*

$$R = \min_{t \in \mathcal{T}} \frac{\sum_{i=1}^{t} \hat{\xi}_i}{\sum_{n \in \cup_{i=1}^{t} \{\mathcal{N}\}_i} D_n}. \tag{21}$$

So, using this definition, we get

$$Y_{nt}^* = \begin{cases} RD_n & \text{if} \quad X_{nt} = 1, \\ 0 & \text{otherwise.} \end{cases} \tag{22}$$

Proposition 1 is helpful because it shows that we do not have to use a computer program to solve the second-stage model. Rather, we can implement the closed-form solution in a simple tool like Excel, which is usually preferred by humanitarian organizations, many of which are nonprofits, due to its simplicity and cost-effectiveness. In addition, it shows that the second stage is fully determined by a time-period bottleneck, which provides relevant information to the decision maker by highlighting the most urgent or problematic period and offering an opportunity to manage this bottleneck.

Proof of Proposition 1. We first start by showing that the optimal solution for the second-stage allocations Y_{nt} as given by (20) and (22) is feasible for the second-stage model by checking constraints (9)-(12).

We know from constraint (2) that for a given collector n, X_{nt} will be equal to one for only one time period t. Thus, due to (10), for a given collector n, Y_{nt} can be positive for only one period t. Therefore, we can insert (22) into (9) and we get $\left| \frac{\sum_{t=1}^{T} Y_{nt}}{D_n} - \frac{\sum_{t=1}^{T} Y_{mt}}{D_m} \right| = R - R = 0$. So, constraint (9) is satisfied.

Now, for any n and t such that $X_{nt} = 0$, inserting (22) into (10), we get $0 \geq 0$, which is satisfied.

For any n and t such that $X_{nt} = 1$, inserting (22) into (10), we get $D_n \geq D_n R$. Due to the supply scarcity assumption given above (Assumption 1), we must have $R < 1$, so this constraint is satisfied.

Let us rewrite the right-hand side of constraint (11) as follows:

$$\begin{aligned} \sum_{i=1}^{t} \sum_{n=1}^{N} Y_{ni} &= \sum_{n \in \cup_{i=1}^{t} \{\mathcal{N}\}_i} Y_{nt} \\ &= \sum_{n \in \cup_{i=1}^{t} \{\mathcal{N}\}_i} D_n R = R \sum_{n \in \cup_{i=1}^{t} \{\mathcal{N}\}_i} D_n. \end{aligned} \tag{23}$$

The first equality follows directly from the properties and notation of $\{\mathcal{N}\}_t$, the second equality follows from the insertion of (22), and the third follows from taking the constant out of the summation.

Since, by definition of the minimum of a function, for all t, we have $R \leq \frac{\sum_{i=1}^{t} \hat{\xi}_i}{\sum_{n \in \cup_{i=1}^{t} \{N\}_i} D_n}$, which is equivalent to $R \sum_{n \in \cup_{i=1}^{t} \{N\}_i} D_n \leq \sum_{i=1}^{t} \hat{\xi}_i$, and inserting this in the final expression in (23) proves that constraint (11) is satisfied. Finally, it is trivial to see that the nonnegativity constraints (12) are satisfied.

Therefore, the solution proposed by (20) and (22) is feasible for the second-stage model. Further, it is straightforward to see that this is also the optimal solution, as Lemma 1 allows the maximization of the total distribution, and any increase in any Y_{nt} value would render the equity constraint (9) infeasible. Therefore, the proposed solution is optimal. This concludes the proof.

Given these insights into the second-stage problem, Appendix A presents an alternative formulation of the SAP-SEF under perfect equity. This formulation is obtained by substituting the closed-form second stage into the first-stage problem and applying straightforward equivalent reformulations. By replacing the second-stage optimization with an explicit function of $\boldsymbol{\xi}$ and $\mathbf{X}$, the recourse problem is eliminated, reducing the model from a two-stage stochastic MILP to a single-stage stochastic integer program. The main benefits of this alternative formulation are the reduced number of decision variables and the alternative structure revealed.

5 Structural properties

This section presents analytical results to better understand the structure and characteristics of the SAP-SEF. First, we provide insights into the upper bound on the fill rate r_n, the fraction of demand satisfied for collector n. A simple and attainable upper bound is presented for the perfect equity case. For imperfect equity, the upper bound derived for the perfect equity case is used to construct an unattainable, yet ex ante derivable, upper bound, followed by a corollary that introduces a modification to ensure attainability. We show that the perfect equity upper bound is attained only when there is no early-time bottleneck. Furthermore, when the equity constraint is non-binding, the optimal total allocation equals the total available supply.

Proposition 2. *For $\theta = 0$, the maximum feasible r_n is*

$$\bar{R}_{\theta=0} := \frac{\sum_{t=1}^{T} \xi_t}{\sum_{n=1}^{N} D_n},$$

which is also an upper bound on R.

$\bar{R}_{\theta=0}$ represents the ideal fill rate in terms of supply usage, since all available supply is utilized and thus yields the maximum possible total allocation. Therefore, if collectors can be allocated over time periods such that $R = \bar{R}_{\theta=0}$, this yields the best attainable solution. In reality, this is usually not possible due to the collectors' demands being large and discrete quantities. Yet, knowing such an upper bound can be useful to the decision maker because it provides, in comparison to how close your fill rate r_n gets to this ideal fill rate, information on the quality of your result.

Proof of Proposition 2. The proposition contains two statements separately proven. First, that the maximum feasible r_n equals $\bar{R}_{\theta=0}$. Second, $\bar{R}_{\theta=0} \geq R$.

We start with this second statement and show that $\bar{R}_{\theta=0} \geq R$. Since R is defined as the minimum over

$t \in \mathcal{T}$, it holds for any fixed t that $R \leq \frac{\sum_{i=1}^{t} \xi_i}{\sum_{n=1}^{N} \sum_{i=1}^{t} D_n X_{ni}}$. Therefore, this also holds for $t = T$, i.e.:

$$R \leq \frac{\sum_{i=1}^{T} \xi_i}{\sum_{n=1}^{N} \sum_{i=1}^{T} D_n X_{ni}}. \tag{24}$$

Now, consider that, when multiplying both sides of (2) by $\sum_{n=1}^{N} D_n$ we obtain

$$\sum_{n=1}^{N} \sum_{i=1}^{T} D_n X_{ni} = \sum_{n=1}^{N} D_n. \tag{25}$$

Using (25) in (24) gives

$$R \leq \frac{\sum_{t=1}^{T} \xi_t}{\sum_{n=1}^{N} D_n} = \bar{R}_{\theta=0}, \tag{26}$$

which yields our result $\bar{R}_{\theta=0} \geq R$, i.e., that $\bar{R}_{\theta=0}$ is an upper bound for the critical ratio R, finalizing the proof for the second statement.

For the first statement, the proof consists of two parts: (i) $\bar{R}_{\theta=0} \geq r_n, n \in \mathcal{N}$ for all feasible solutions, and (ii) there exists a feasible solution such that $r_n = \bar{R}_{\theta=0}, n \in \mathcal{N}$.

Part (i) follows from Proposition 1 and finding (26). Proposition 1 yields that in perfect equity, every collector n receives the same fill rate R, and hence $r_n = R$ for all collectors $n \in \mathcal{N}$. Inserting this result in finding (26) gives $\bar{R}_{\theta=0} \geq R = r_n$ for $n \in \mathcal{N}$, which concludes (i).

Part (ii) follows from considering the allocation of all collectors to the last period T, i.e., $X_{nT} = 1, X_{n,t\neq T} = 0, n \in \mathcal{N}$. Let us show that this is a feasible solution. Constraints (2) and (3) are trivially satisfied. Since we consider $\theta = 0$, where $r_n = R, n \in \mathcal{N}$, (9) becomes $|R - R| \leq 0$, which is satisfied.

For $n \in \mathcal{N}$ and $t \neq T \in \mathcal{T}$, inserting $X_{n,t\neq T} = 0$ and (22) into (10), we get $0 \geq 0$, which is satisfied.

For $n \in \mathcal{N}$ and $t = T$, inserting $X_{nT} = 1$ and (22) into (10), we get $D_n \geq D_n R$. Due to the supply scarcity assumption (Assumption 1), we must have $R < 1$, so this constraint is also satisfied.

For $t \neq T \in \mathcal{T}$, inserting (22) into (11), we get $\sum_{i=1}^{t} \xi_i \geq 0$, which is satisfied.

For $t = T$, the right-hand side of (11) is:

$$\sum_{i=1}^{T} \sum_{n=1}^{N} Y_{ni} = R\Delta, \tag{27}$$

where the equality follows from applying Proposition 1 (use (19)) and Definition 1. Since, by definition of the minimum of a function, for all t, we have $R \leq \frac{\sum_{i=1}^{t} \xi_i}{\sum_{n \in \cup_{i=1}^{t} \{\mathcal{N}\}_i} D_n}$. Since this also holds for $t = T$, we have $R \sum_{n \in \cup_{i=1}^{T} \{\mathcal{N}\}_i} D_n \leq \sum_{i=1}^{T} \xi_i \Leftrightarrow R\Delta \leq \sum_{i=1}^{T} \xi_i$, and inserting this in (27) proves that constraint (11) is also satisfied for $t = T$. Finally, it is trivial to see that the nonnegativity constraints (12) are satisfied.

For showing attainability, consider the feasible solution $X_{nT} = 1, X_{n,t\neq T} = 0, n \in \mathcal{N}$. By Proposition 1 we have that $r_n = \bar{R}_{\theta=0}, n \in \mathcal{N}$, and the allocation is fully determined by R. For this feasible solution, all demand is pooled in the last period and R can be written as

$$R = \min_{t \in \mathcal{T}} \frac{\sum_{i=1}^{t} \xi_i}{\sum_{n \in \cup_{i=1}^{t} \{\mathcal{N}\}_i} D_n} = \frac{\sum_{i=1}^{T} \xi_i}{\sum_{n=1}^{N} D_n},$$

as for $t < T$, the denominator equals zero while the numerator is nonnegative, so these terms are not well-defined (or can be treated as $+\infty$) and do not affect the minimum. Hence, for the feasible solution

$X_{nT} = 1, X_{n,t\neq T} = 0, n \in \mathcal{N}$ we have

$$r_n = R = \frac{\sum_{i=1}^{T} \xi_i}{\sum_{n=1}^{N} D_n} = \bar{R}_{\theta=0} \quad n \in \mathcal{N},$$

which establishes attainability and completes the proof. □

Proposition 3. *In case of $\theta > 0$, the fill rate r_n of any collector n, has the upper bound:*

$$\bar{R}_{\theta>0} := \min\{1, \bar{R}_{\theta=0} + \theta\}. \tag{28}$$

The proposition presents a hard ceiling on the individual fill rates, defined only by the total supply by total demand, expressed by $\bar{R}_{\theta=0}$, and the equity parameter, which can therefore be derived without solving the SAP-SEF. Note that if supply is only limitedly scarce and θ is relatively large, $\bar{R}_{\theta=0} + \theta$ can be greater than one, but the allocation will be never more than the demand (Assumption 1), motivating the inclusion of the minimum function.

$\bar{R}_{\theta>0}$ is a structural upper bound, as it cannot be attained by any feasible solution. The true attainable maximum will be strictly smaller, depending on the demand distribution. To see why the bound cannot be reached, suppose one collector attains $\bar{R}_{\theta=0} + \theta$. The equity constraint then implies that every other collector must receive at least $\bar{R}_{\theta=0}$. Even if all others receive exactly this minimum level, the resulting total allocation would exceed the available supply, since allocating $\bar{R}_{\theta=0}$ to all collectors already exhausts the total supply.

Proof of Proposition 3. In this proof we combine a weighted average bound derived from (11) with a diameter bound derived from (9). We begin by dividing both sides of (11) by total demand, which yields:

$$\frac{\sum_{i=1}^{t} \xi_i}{\sum_{n=1}^{N} D_n} \geq \frac{\sum_{i=1}^{t} \sum_{n=1}^{N} Y_{ni}}{\sum_{n=1}^{N} D_n} \quad t \in \mathcal{T}.$$

Hence, for $t = T$ we have:

$$\frac{\sum_{i=1}^{T} \xi_i}{\sum_{n=1}^{N} D_n} \geq \frac{\sum_{i=1}^{T} \sum_{n=1}^{N} Y_{ni}}{\sum_{n=1}^{N} D_n} = \frac{\sum_{n=1}^{N} \sum_{i=1}^{T} Y_{ni}}{\sum_{n=1}^{N} D_n},$$

which can, when noting that the left-hand side equals $\bar{R}_{\theta=0}$ and using the definition of r_n (i.e., $r_n = \frac{\sum_{t=1}^{T} Y_{nt}}{D_n} \Leftrightarrow \sum_{t=1}^{T} Y_{nt} = r_n D_n$), be rewritten as:

$$\bar{R}_{\theta=0} \geq \frac{\sum_{n=1}^{N} r_n D_n}{\sum_{n=1}^{N} D_n}. \tag{29}$$

Note that the right-hand side is a weighted average of r_n with weights D_n. Hence, this inequality shows that the average fill rate (weighted by demand) cannot exceed the total supply by total demand, i.e., cannot exceed $\bar{R}_{\theta=0}$, and is thereby bounded by what total supply allows.

Now, consider (9) while inserting for the definition of r_n, which gives $|r_n - r_m| \leq \theta$ for $n, m \in \mathcal{N}$. As such, all fill rates lie inside an interval of width θ, and so $\max_n r_n - \min_n r_n \leq \theta$. Let $r_{\max} := \max_n r_n$, then we have for every $n \in \mathcal{N}$:

$$r_n \geq r_{\max} - \theta. \tag{30}$$

Now, combining the results (29) and (30) gives:

$$\bar{R}_{\theta=0} \geq \frac{\sum_{n=1}^{N} r_n D_n}{\sum_{n=1}^{N} D_n} \geq r_{\max} - \theta, \tag{31}$$

where the last inequality follows since all r_n are at least $r_{\max} - \theta$, and thus so is their weighted average. The

result (31) ensures that $r_{\max} \leq \bar{R}_{\theta=0} + \theta$. Provided that we have $r_n \leq 1$ (Assumption 1), we arrive at our final result that $r_n \leq \min\{1, \bar{R}_{\theta=0} + \theta\}$ for all $n \in \mathcal{N}$. □

Corollary 1. *The bound on r_n in Proposition 3 becomes tight if the term θ in* (28) *is replaced by $\theta\left(1 - \frac{D_k}{\sum_{n=1}^{N} D_n}\right)$ if $r_k \geq \theta$, and $\sum_{t=1}^{T} \xi_t \left(\frac{1}{D_k} - \frac{1}{\Delta}\right)$ if $r_k < \theta$, where k denotes the collector attaining the largest fill rate.*

In contrast to the bound proposed in Proposition 3, which is purely structural and cannot be attained, Corollary 1 provides a tight and attainable bound on the fill rate r_n. Proposition 3 establishes a simple ceiling that depends only on the total demand, the total supply, and the equity parameter. Corollary 1 refines this bound by incorporating the demand share of the collector that achieves the highest fill rate. As a result, the bound is more informative in terms of what can actually be achieved, but it depends on which collector receives the highest fill rate and is thus less directly available than the structural bound in Proposition 3. The proof of Corollary 1 can be found in Appendix B.

Proposition 4. *When $\theta = 0$, at optimality $r_n = \bar{R}_{\theta=0}$ holds if and only if*

$$\frac{\sum_{i=1}^{t} \hat{\xi}_i}{\sum_{n \in \cup_{i=1}^{t} \{\mathcal{N}\}_i} D_n} \geq \frac{\sum_{i=1}^{T} \hat{\xi}_i}{\sum_{n=1}^{N} D_n}, \tag{32}$$

for all t.

The left-hand side of inequality (32) represents the cumulative supply fraction at each period, which is also used in the critical ratio R (see Definition 1), while the right-hand side equals the upper bound $\bar{R}_{\theta=0}$. This proposition therefore states that the upper bound for perfect equity is attained if and only if there is no early-time bottleneck. If, for some $t < T$, the cumulative supply over cumulative assigned demand is strictly lower than the upper bound, the critical ratio R is determined by this bottleneck period. In other words, too much demand would be assigned before sufficient supply has arrived, preventing the upper bound from being reached.

The reason this can occur comes from the prioritization of early allocation in the objective. Suppose we were to remove this prioritization, i.e., replace the weight $T - t + 1$ in (7) with 1. Then, under perfect equity, allocating all demand to the last period would always be optimal, resulting in a uniform allocation where $r_n = R = \bar{R}_{\theta=0}$ for all collectors. In that case, the upper bound would always be attained. However, because the objective favors earlier allocations, achieving the upper bound requires that the cumulative supply relative to cumulative demand at each period be at least as large as $\bar{R}_{\theta=0}$, as expressed in (32). This ensures that no early period becomes a limiting bottleneck that would force the critical ratio R below the maximum possible.

Proof of Proposition 4. Proposition 1 states that for $\theta = 0$ the fill rates r_n for all collectors n are equal and take on the value derived by R.

$\Rightarrow$: Assume that $r_n = \bar{R}_{\theta=0}$ for all n. Suppose, for contradiction, that there exists a $t' \in \mathcal{T}$ such that $\frac{\sum_{i=1}^{t'} \hat{\xi}_i}{\sum_{n \in \cup_{i=1}^{t'} \{\mathcal{N}\}_i} D_n} < \frac{\sum_{t=1}^{T} \hat{\xi}_t}{\sum_{n=1}^{N} D_n}$. Then by definition of R, R selects the time period with the worst cumulative fill rate, and given the existence of this time period t', it must be that $R < \bar{R}_{\theta=0}$, which contradicts the assumption $r_n = \bar{R}_{\theta=0}$. Hence, (32) must hold for all t.

$\Leftarrow$: Assume that (32) holds for all t. Then, we have $R = \min_{t \in \mathcal{T}} \frac{\sum_{i=1}^{t} \hat{\xi}_i}{\sum_{n \in \cup_{i=1}^{t} \{\mathcal{N}\}_i} D_n} = \frac{\sum_{t=1}^{T} \hat{\xi}_t}{\sum_{n=1}^{N} D_n}$. Therefore, $r_n = \bar{R}_{\theta=0}$ for all n. □

Proposition 5. *Suppose that* $\xi_t \leq \sum_{m\in\{\mathcal{N}\}_t} D_m$ *for all* $t \in \mathcal{T}$. *When* $\theta \geq 1$, *the total allocation* $\tilde{\tilde{Y}}$ *is equal to the total available supply* $\sum_{t=1}^{T} \xi_t$.

Proposition 5 shows that relaxing the equity requirement expands the feasible allocation space and allows for a more efficient use of available supply. When $\theta \geq 1$, the equity constraint becomes permanently non-binding. As a result, all available supply can be allocated, and the total allocation reaches its maximum possible level. This highlights the trade-off between equity and efficiency: stricter equity requirements may prevent full utilization of supply, while relaxing them enables maximal allocation. Note that the extra assumption on which this proposition relies is stricter than the supply scarcity in Assumption 1 as each period's supply must be lower than the demand assigned to that period. However, in humanitarian operations such as food assistance, demand often substantially exceeds available supply, making this assumption realistic in many settings. The proof of Proposition 5 is deferred to Appendix C.

6 The critical-ratio balancing heuristic under perfect equity

We introduce the Critical-Ratio Balancing (CRB) Heuristic, which is an interpretable heuristic for the perfect-equity case of SAP-SEF. While SAP-SEF can be solved to optimality by using techniques such as Sample Average Approximation, the extensive-form formulation becomes computationally challenging as the number of scenarios, time periods, and collectors increases, due to the growth in both first-stage scheduling decisions and scenario-dependent second-stage allocation variables. Our computational experiments indicate that instances with more than 20 scenarios and 20 collectors can be difficult to solve to optimality within a practically reasonable time. This motivates the development of a scalable procedure that exploits the structural properties of the perfect-equity allocation problem.

The proposed heuristic builds on the insights from Propositions 1 and 2. Specifically, it uses the perfect-equity upper bound $R_{\theta=0}$ as a goal to construct first-stage scheduling decisions and then applies the closed-form allocation rule implied by the critical ratio R to determine second-stage allocations. Thus, the heuristic avoids solving the full stochastic mixed-integer program while preserving the key equity structure of the original model. In addition to being computationally efficient, the procedure is transparent and easy to implement, making it suitable for humanitarian organizations that must justify scheduling and allocation decisions to donors and other stakeholders. We first present the pseudocode for the heuristic and then describe its implementation and evaluation.

We now describe the logic of the heuristic. The heuristic is designed for the perfect-equity case, i.e., $\theta = 0$, where all collectors receive the same fill rate in each scenario. The key idea is to use the structural properties of the perfect-equity allocation problem to separate the scheduling and allocation decisions. Specifically, the heuristic first constructs a first-stage schedule using expected supply information for each period. Then, after each supply scenario is realized, it uses the closed-form optimal allocation rule from Proposition 1 to determine the second-stage allocation quantities.

The first-stage scheduling rule is based on balancing cumulative scheduled demand with cumulative expected supply. Let $\bar{\xi}_t$ denote the expected supply in period t, and let $\bar{R}_{\theta=0}$ denote the target fill rate under perfect equity, as defined in Proposition 2. If collectors are expected to receive approximately a

Algorithm 1 Critical-Ratio Balancing (CRB) Heuristic

Input: Demands D_n, scenarios $\mathcal{S}$, expected supplies $\bar{\xi}_t$, scenario supplies ξ_t^s $\forall n \in \mathcal{N}, s \in \mathcal{S}, t \in \mathcal{T}$, and target fill rate $\bar{R}_{\theta=0}$.
Output: Heuristic schedule $\mathbf{X}^H$, allocations $\mathbf{Y}^H$, and objective value Z^H.
Set $C_t = \min\left\{\left(\sum_{i=1}^{t} \bar{\xi}_i\right)/\bar{R}_{\theta=0}, \sum_{n\in\mathcal{N}} D_n\right\}$ for all $t \in \mathcal{T}$.
Sort collectors in nonincreasing order of D_n.
Initialize $X_{nt}^H = 0$ for all n, t, and $a_t = 0$ for all t, where a_t is demand scheduled in period t.
for each collector n in sorted order **do**
    Assign n to $t^* = \min\{t \in \arg\min_{t\in\mathcal{T}} \sum_{\tau=1}^{T} \left[\sum_{i=1}^{\tau} (a_i + D_n I\{i = t\}) - C_\tau\right]^2\}$ where $I\{i = t\} = 1$ is the indicator function and equals one if $i = t$ and zero otherwise.
    Set $X_{nt^*}^H = 1$ and $a_{t^*} \leftarrow a_{t^*} + D_n$.
end for
for each scenario $s \in \mathcal{S}$ **do**
    Compute $R_s(\mathbf{X}^H) = \min_{t\in\mathcal{T}:\sum_{i=1}^{t}\sum_{n\in\mathcal{N}} D_n X_{ni}^H > 0} \frac{\sum_{i=1}^{t} \xi_i^s}{\sum_{i=1}^{t}\sum_{n\in\mathcal{N}} D_n X_{ni}^H}$.
    Set $Y_{nts}^H = R_s(\mathbf{X}^H) D_n X_{nt}^H$ for all $n \in \mathcal{N}, t \in \mathcal{T}$.
end for
Set $Z^H = \frac{1}{|\mathcal{S}|} \sum_{s\in\mathcal{S}} \sum_{t\in\mathcal{T}} (T - t + 1) \sum_{n\in\mathcal{N}} Y_{nts}^H$.
Return $\mathbf{X}^H$, $\mathbf{Y}^H$, and Z^H.

$\bar{R}_{\theta=0}$ fraction of their demand, then cumulative expected supply through period t can support approximately $\frac{\sum_{i=1}^{t} \bar{\xi}_i}{\bar{R}_{\theta=0}}$ units of cumulative scheduled demand. Therefore, the heuristic defines this quantity as the cumulative scheduled-demand target for period t, capped above by total demand. These targets guide the scheduling decisions. Collectors are sorted in nonincreasing order of demand and assigned to periods sequentially. For each collector, the heuristic evaluates all possible periods and assigns the collector to the period that minimizes the squared deviation between cumulative scheduled demand and the cumulative targets across the planning horizon. If a tie exists, an earlier period is preferred. This step is intended to align the timing of scheduled demand with the timing of expected supply. For example, when expected supply arrives earlier in the horizon, the targets increase more quickly, encouraging the heuristic to schedule more demand earlier. Conversely, when expected supply arrives later, the targets increase more slowly, resulting in a more delayed schedule. Once the first-stage schedule $\mathbf{X}^H$ is constructed, the second-stage allocation is obtained in closed form for each scenario, using the result from Proposition 1. Finally, the heuristic objective value is computed by evaluating the resulting allocations over all scenarios. In Section 7, we evaluate the performance of the CRB heuristic by comparing it to the performance of exactly solving SAP-SEF.

7 Numerical experiments: US food pantries

The first goal of these numerical experiments is to improve understanding of the SAP-SEF, examining the effects of problem size, supply scarcity, temporal supply patterns, supply variability, and equity requirements. Additionally, these experiments test the problem's complexity and, using pseudo-real instances from American food pantries, offer managerial insights into food bank operations.

All computational experiments were conducted on a personal computer equipped with an Intel Core 7

165U CPU with a base speed of 1.70 GHz and 32 GB RAM running Microsoft Windows 11. The models were implemented in Python version 3.13.3 and solved using Gurobi Optimizer version 8.1.1. All experiments were performed using Gurobi's default multi-threading settings. The data and code used in this study are publicly available on the authors' GitHub repository (citation omitted for review).

7.1 US context and problem setting

In the United States, food banks provide food support to millions of food-insecure individuals and families. Food banks collect food donations from many sources such as the government, Feeding America, big box stores, and manufacturers. After collecting donations at their large warehouses, they distribute the food to their partner agencies, which are smaller nonprofit locations such as food pantries, homeless shelters, or faith-based organizations. These agencies, in turn, distribute the food to their food-insecure clients. In addition to receiving food donations from their food bank partner, agencies also typically receive donations from their local donors, such as supermarkets and farmers. Some agencies, such as soup kitchens, distribute prepared meals to their clients (meal-serving agency), whereas other agencies, such as food pantries, distribute groceries that clients pick up (grocery-distributing agency). For more details on US food bank and agency operations we refer to Sengul Orgut et al. (2025) and Davis et al. (2024).

We consider the context of a grocery-distributing food pantry that receives donations from various donors throughout the week and must assign their clients to each day of the week to come to the food pantry and pick up their allocated food donation, based on their demand, which is a function of their needs and household size and is fully observable by the food pantry. The food pantry does not know the exact donation amounts that they will receive over a given week when determining the clients' schedules. Further, schedules must be determined once and cannot be changed each week to respect the clients' work and school schedules. We assume that the donation amounts get revealed to the food pantry at the beginning of each week to help them determine the allocation quantities to each client.

7.2 Problem instances

The planning horizon consists of $T = 5$ periods, which can be interpreted as five collection or service days. To approximate the stochastic program, we use Sample Average Approximation (SAA) with a finite set of independently generated supply scenarios.

Demand. We consider a typical food pantry operating in the United States. Based on the United States Department of Agriculture (USDA, 2023), 62% of the food-insecure people were categorized as having "low food security" and 38% as "very low food security". The number of people served by food pantries varies considerably, so we consider four different values for N, $N \in \{20, 50, 100, 150\}$, where each index n represents a household (or collector) with heterogeneous demand. Household demand is generated to reflect variation in both household size and food insecurity severity. Specifically, approximately 62% of households are assigned to a lower-need group and the remaining 38% to a higher-need group, with counts rounded so that the totals match the target value of N. Household size is generated using the empirical proportions $0.29, 0.36, 0.15, 0.12, 0.05, 0.02$, and 0.01 for household sizes 1 through 7, respectively (Statista Research

Department, 2025). After assigning each household to a size and need category, its weekly demand D_n is computed as $D_n = 1.2 \times$ (household size) $\times$ (meals needed per person), where meals needed per person follow a discrete uniform distribution over $\{3, 4, 5, 6\}$ for the lower-need group and over $\{7, 8, 9, 10\}$ for the higher-need group, and using Feeding America's estimate of 1.2 pounds per meal (Feeding America, 2026). For each value of N, the household demand vector is generated once and then held fixed across all supply experiments so that differences in performance can be attributed solely to changes in the supply environment and equity parameter.

Supply. Supply is modeled as a stochastic daily process over the five periods. For each instance, total expected supply is set to be scarce relative to total demand. We consider two aggregate supply levels: a high-supply case with total mean supply equal to 0.75Δ, and a low-supply case with total mean supply equal to 0.25Δ. In both cases, expected supply is strictly below total demand, ensuring a scarcity setting.

To examine the impact of temporal supply structure, we consider six mean supply profiles across the five periods: (1) a *flat* profile, in which expected supply is identical across all days; (2) a *symmetric concave* profile, in which expected supply is relatively high on days 1 and 5 and lowest on day 3; (3) a *symmetric convex* profile, which is the reverse pattern: expected supply is lowest on days 1 and 5 and highest on day 3; (4) a monotonic *decreasing* pattern; (5) a monotonic *increasing* pattern; and (6) a *random* profile, in which day-specific mean weights are randomly generated and then normalized so that the five-period total still matches the required aggregate mean supply. Figure 2 provides an illustration of the six supply profiles.

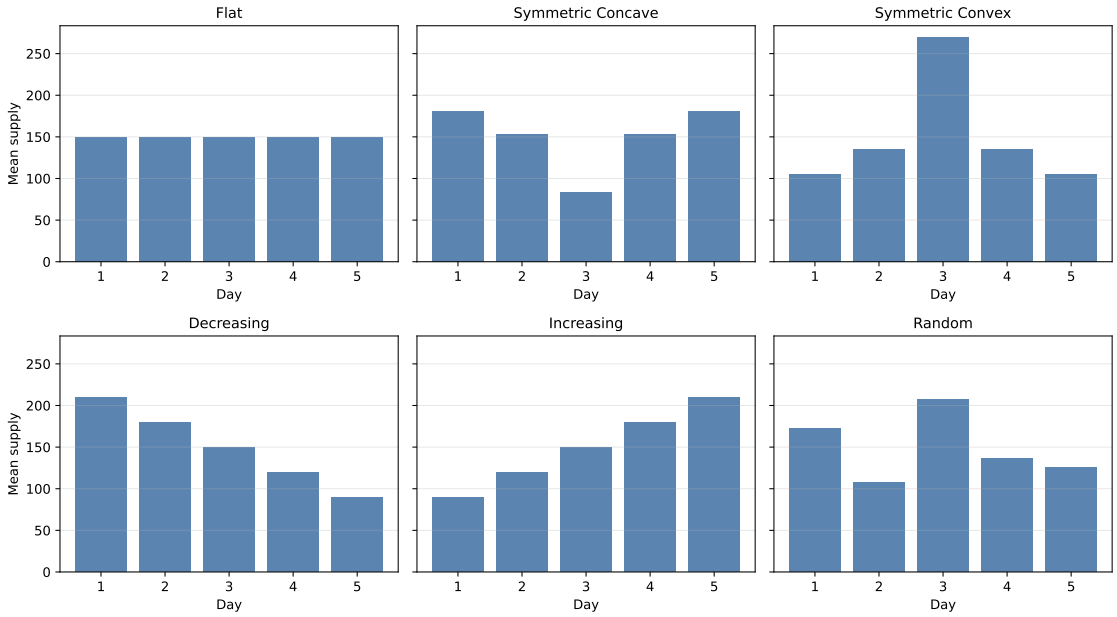

Figure 2: Supply profiles.

Around each mean profile, we generate stochastic daily supplies using a continuous probability distribution. Our base implementation uses a Lognormal distribution so that realized supply is always nonnegative. For each day, the standard deviation is specified as a percentage of that day's mean supply. We consider three variability settings: low variability, with standard deviation equal to 1% of the mean; middle variability, with

standard deviation equal to 10% of the mean; and high variability, with standard deviation equal to 25% of the mean. Daily supplies are sampled independently across days and across scenarios. For each combination of N, supply level, temporal pattern, and variability level, the same scenario set is used across all values of the equity parameter to ensure consistent comparisons.

Equity. The equity requirement is controlled by the parameter θ, which limits the disparity in fill rates across households. We vary θ from 0 to 1 in increments of 0.05.

Table 3 summarizes these problem instance settings. We use the 20 sample average scenarios, $N = 20$, middle supply variability of $\sigma = 0.10\mu$, and $\theta = 0.1$ as the base case instance, which is also highlighted in bold in the table.

Table 3: Summary of experimental design settings. Base case values are shown in bold.

Factor	**Settings**
Planning horizon	$T = 5$ periods
SAA sample size	$S = \{10, \mathbf{20}, 30, 40, 50, 60\}$ scenarios per instance
Service-region size	$N \in \{\mathbf{20}, 50, 100, 150\}$ households
Household need groups	Approximately 62% lower-need households and 38% higher-need households
Meals needed per person	Lower-need: discrete uniform on $\{3, 4, 5, 6\}$;
	Higher-need: discrete uniform on $\{7, 8, 9, 10\}$
Household size distribution	Number of people per household varies between 1 through 7 with proportions $(0.29, 0.36, 0.15, 0.12, 0.05, 0.02, 0.01)$
Demand calculation	$D_n = 1.2 \times$ (household size) $\times$ (meals per person) pounds
Aggregate supply level	High supply: total mean supply = 0.75Δ;
	Low supply: total mean supply = 0.25Δ
Temporal mean supply profile	Flat; Symmetric concave; Symmetric convex; Decreasing; Increasing; Random
Supply distribution	Lognormal (base implementation)
Supply variability	Low: $\sigma = 0.01\mu$; Middle: $\sigma = 0.10\mu$; High: $\sigma = 0.25\mu$
Equity parameter	$\theta \in [0.00 : 0.05 : 1.00]$, **Base:** $\theta = 0.1$

7.3 Sensitivity for the number of scenarios

The use of the Sample Average Approximation (SAA) method for solving a stochastic model requires the generation of random scenarios. In our case, a scenario corresponds to a sequence of supply values generated for each time period in the time horizon. Higher number of scenarios provides a better representation of the underlying uncertainty and hence increases solution quality. However, it comes at a cost of solution time. Figure 3 shows the change in average objective function value Z^* (left axis, solid line) and solution time (right axis, dashed line) for varying number of scenarios. We present results for the high average supply setting with a random supply pattern, as the other supply patterns exhibit similar behavior. For each scenario

size, we ran 50 replications and report the average objective value and solution time across those runs. We include only solutions that were within a 1% optimality gap.

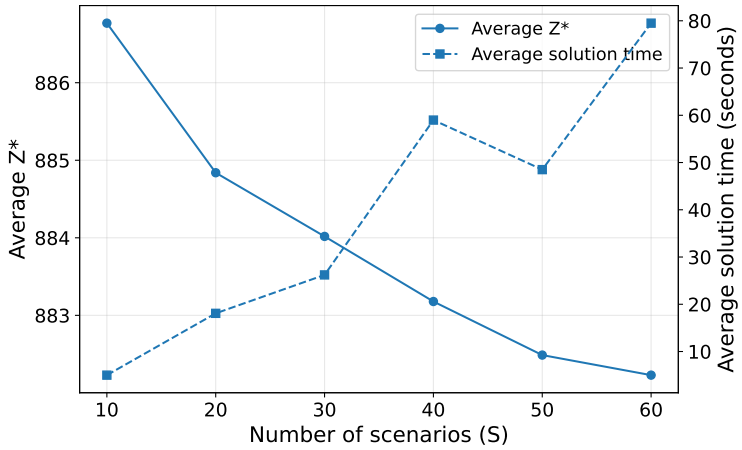

Figure 3: Average optimal solution value (Z^* - left axis) and the average solution time (right axis) for varying scenario sizes.

As expected, solution quality improves, that is, Z^* decreases, as the number of scenarios increases. However, solution time also rises substantially, from about 5 seconds with 10 scenarios to roughly 80 seconds with 60 scenarios. We also observe an elbow at 20 scenarios: the improvement in Z^* is pronounced between 10 and 20 scenarios but becomes much flatter beyond that point. Finally, although solution quality continues to improve, the overall reduction in Z^* is less than 0.5%. Based on this trade-off, we use 20 scenarios in the remainder of the analysis.

7.4 Cost of equity

In this section, we investigate how the equity restriction affects food allocation. Figure 4 illustrates how the objective function value changes with $\theta \in \{0, 0.05, 0.1, 0.15, 0.2, 0.25\}$. No meaningful change is observed for values of θ higher than 0.25 and hence are omitted here. The left and right panels report the results for the high- and low-supply settings, respectively, with different lines representing the supply profiles shown in the legend. Figure 5 similarly presents the effect of θ on solution time.

Figure 4 shows that, as θ increases, the objective function value also increases, as expected. This indicates that relaxing the equity requirement allows the system to distribute food earlier and in larger quantities, since the allocation decisions are less constrained by the need to equalize fill rates across collectors. A further result that emerges is the difference across supply profiles. For both the high- and low-supply settings, the decreasing supply pattern yields the highest Z^*, whereas the increasing supply pattern yields the lowest Z^*, despite the total mean supply being identical across profiles. In contrast, the remaining profiles—where supply is symmetric or random—produce very similar outcomes. This pattern is partly driven by the structure of the objective function, which rewards earlier distribution. However, this does not fully explain the result: even when total supply is held constant, the decreasing supply pattern consistently results in more food being distributed than the increasing supply pattern. The differences are modest, averaging 1.20% in the low-supply setting and 0.86% in the high-supply setting. Nevertheless, in humanitarian operations, even a

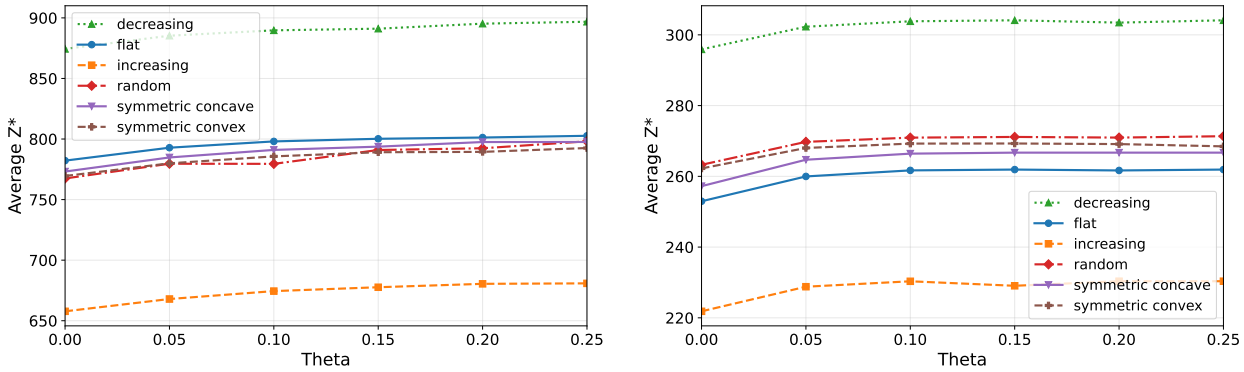

(a) High supply (b) Low supply

Figure 4: Average Z^* for varying equity limit Θ, for all supply profiles.

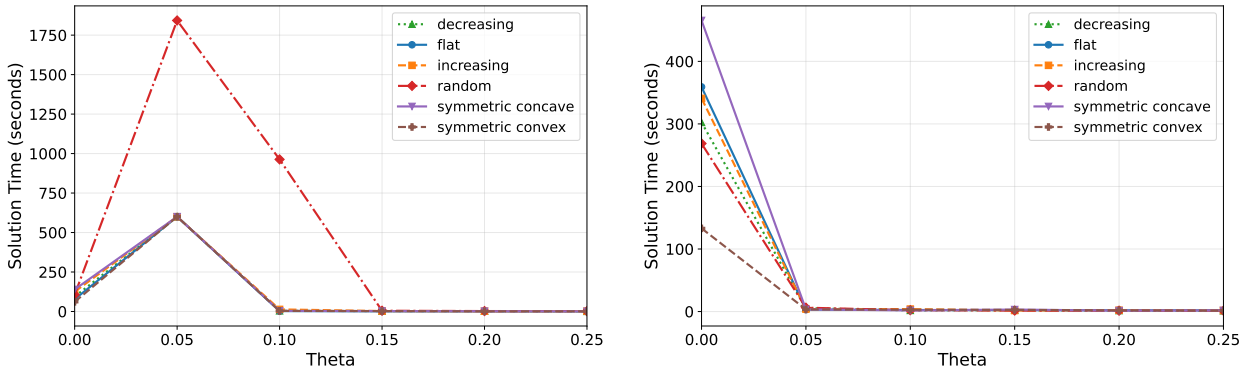

(a) High supply (b) Low supply

Figure 5: Average solution time for varying equity limit Θ, for all supply profiles.

1% improvement can translate into a meaningful increase in the amount of food distributed to beneficiaries. For example, one of our partner food pantries, Northwestern Settlement in Chicago, Illinois, distributes approximately 250,000 pounds of food annually, or about 4,800 pounds per week. Thus, even 1% of weekly supply corresponds to 48 pounds of food, enough to provide approximately 40 meals to individuals in need. This highlights the operational value of receiving donations earlier in the planning horizon, even when the total amount of supply remains unchanged.

Figure 5a reveals that, in the high-supply setting, solution times are moderate (around 200 seconds) under perfect equity but increase substantially at $\theta = 0.05$, after which they rapidly drop to nearly zero for larger values of θ. The longest solution times occur under the random supply pattern. A possible explanation is that under abundant supply, the model has many feasible allocations that satisfy demand while respecting the equity constraints. When θ is very small, the tight equity requirement restricts the feasible region, which can actually simplify the search. As θ increases slightly, however, the feasible region expands and introduces many additional candidate allocations with similar objective values, making it more difficult for the solver to identify and prove optimality. Once θ becomes sufficiently large, the equity constraint becomes largely nonbinding, significantly simplifying the problem and leading to very short solution times.

In contrast, when supply is limited (Figure 5b), solution times are highest under perfect equity and drop almost immediately to near zero as θ increases. In this case, the tight equity requirement forces the model to carefully balance allocations across collectors despite scarce supply, creating a more complex combinatorial structure. Relaxing the equity constraint allows the solver to find feasible and optimal solutions more easily.

7.5 Cost of supply variability

In stochastic optimization, the level of uncertainty can materially influence solution quality. In our setting, uncertainty stems from variability in stochastic supply and is captured through the variance of the lognormal distribution. To assess the cost of supply variability (CSV), we compare model performance under low- and high-variability supply settings while holding expected supply fixed. Specifically, supply is modeled using a Lognormal distribution, with the standard deviation set equal to 1% of the mean in the low-variability setting and 25% of the mean in the high-variability setting. We define the cost of supply variability as the relative loss in the objective value under high variability compared with low variability. We define the cost of supply variability as

$$\text{CSV} = \frac{Z^*_{\text{low}} - Z^*_{\text{high}}}{Z^*_{\text{high}}},$$

where Z^*_{low} and Z^*_{high} denote the optimal objective values under low- and high-variability supply, respectively. Hence, CSV measures how much solution quality deteriorates when variability in supply increases. It captures the relative gap between outcomes under low and high uncertainty, showing the performance loss caused purely by higher unpredictability in supply. Figure 6 reports how this measure varies across different supply levels and three supply patterns (flat, increasing, and decreasing) as others yield similar results.

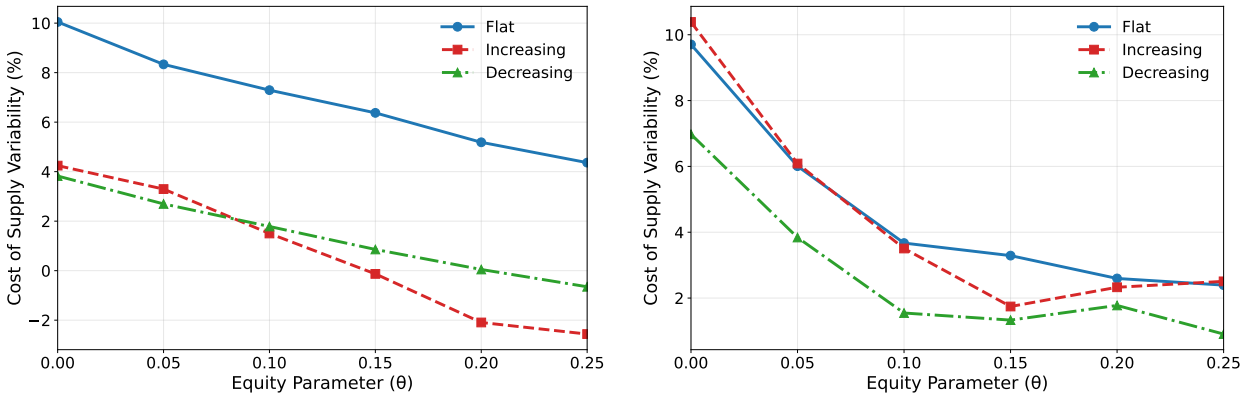

(a) High supply

(b) Low supply

Figure 6: Average solution time for varying number of collectors N, for all supply profiles.

From Figure 6a, we find that under high supply, the cost of supply variability declines as equity is relaxed. This indicates that variability in supply has the greatest effect when the model places a strong emphasis on equity. Put differently, when the objective is to achieve near-perfect equity, improved information about supply becomes particularly valuable. Figure 6b reveals a similar overall pattern in the low-supply setting. However, for the increasing and decreasing supply patterns, the relationship becomes slightly nonmonotone at higher values of θ. A possible explanation is that when θ is large, the model is no longer constrained by

equity and focuses on distributing all the supply as quickly as possible. In this regime, the supply realizations becomes more important, particularly when supply is concentrated earlier or later in the horizon. As a result, variability in supply can again have a noticeable impact on the allocation decisions, leading to a slight increase in the cost of uncertainty at high θ values.

7.6 Value of stochastic optimization and perfect information

We next assess the value of explicitly modeling supply uncertainty by comparing the stochastic program with two benchmarks (Birge Louveaux, 2011). First, we solve an *expected-value (EV) model* in which uncertain daily supply is replaced by its mean. The resulting first-stage schedule, denoted $\mathbf{X}^{EV}$, is then evaluated on the same stochastic scenarios used in the stochastic program, with second-stage allocations re-optimized for each scenario while keeping $\mathbf{X}^{EV}$ fixed. Let z^{EV} denote this expected performance. Second, we compute a *wait-and-see (WS) benchmark*, in which the realized supply vector is assumed to be known before any decisions are made. For each scenario, we solve the corresponding deterministic problem and average the resulting objective values across scenarios. Let this value be denoted by z^{WS}. The stochastic programming objective is denoted by z^{SP}. Since the WS solution has perfect information and the SP solution explicitly accounts for uncertainty, we expect $z^{WS} \geq z^{SP} \geq z^{EV}$. We report two standard performance measures. The value of the stochastic solution is VSS (%) $= \frac{z^{SP}-z^{EV}}{z^{EV}} \times 100$, which measures the percent improvement from using the stochastic model instead of the expected-value solution. The expected value of perfect information is EVPI (%) $= \frac{z^{WS}-z^{SP}}{z^{SP}} \times 100$, which measures the percent additional benefit of observing supply realizations before making scheduling decisions. Table 4 provides the results.

Table 4: Value of stochastic solution and expected value of perfect information.

Supply profile	θ	z^{EV}	z^{SP}	z^{WS}	VSS (%)	EVPI (%)
Flat	0.00	764.78	782.17	797.43	2.27%	1.95%
	0.10	794.35	798.03	801.06	0.46%	0.38%
Symmetric Concave	0.00	765.81	773.14	792.65	0.96%	2.52%
	0.10	785.01	791.00	795.88	0.76%	0.62%
Symmetric Convex	0.00	765.90	769.48	788.11	0.47%	2.42%
	0.10	782.58	785.67	790.93	0.39%	0.67%
Decreasing	0.00	872.78	874.39	890.70	0.18%	1.87%
	0.10	887.94	889.72	895.27	0.20%	0.62%
Increasing	0.00	648.14	657.77	675.72	1.49%	2.73%
	0.10	664.76	674.45	679.22	1.46%	0.71%
Random	0.00	762.59	767.45	792.71	0.64%	3.29%
	0.10	778.84	789.75	796.64	1.40%	0.87%

The results indicate that explicitly modeling supply uncertainty yields consistent, though moderate, improvements over the expected-value solution, with VSS values ranging from 0.18% to 2.27%. The value of perfect information is larger, reaching up to 3.29%, suggesting that improved information about supply realizations can provide greater benefits than stochastic optimization alone. Both VSS and EVPI generally decrease as the equity parameter increases. When equity constraints are tight, scheduling decisions are more

sensitive to supply uncertainty, increasing the value of anticipating uncertainty and observing realizations. As the constraint is relaxed, the system becomes less sensitive to uncertainty, and the performance of the expected-value, stochastic, and perfect-information solutions converges. Furthermore, the value of stochastic modeling and information varies significantly across supply patterns. In particular, increasing and random supply profiles exhibit the highest VSS and EVPI, highlighting the importance of temporal variability and uncertainty in determining the benefits of different approaches. Finally, although the percentage improvements appear modest, they can translate into substantial absolute gains in total distribution, particularly in large-scale systems.

7.7 Performance of CRB heuristic

Although we were able to solve SAP-SEF for instances with 20 collectors, the model becomes computationally challenging as the number of collectors increases and is often unable to reach optimality within a reasonable time. Therefore, Table 5 compares the SAP-SEF incumbent solutions with the Critical-Ratio Balancing heuristic (introduced in Section 6) under perfect equity. The column Z^{OPT} reports the incumbent objective value obtained by solving SAP-SEF, while OPT MIP Gap reports the final optimality gap of SAP-SEF at termination. OPT Time gives the solution time of SAP-SEF in seconds. The column Z^{H} reports the objective value obtained by the heuristic, evaluated over the same supply scenarios as SAP-SEF. H Gap is computed as $100 \times \frac{Z^{OPT} - Z^{H}}{Z^{OPT}}$. We do not report heuristic runtimes since they are negligible (< 0.1 seconds). Results are shown for the flat, increasing, and decreasing supply profiles only, as the other profiles yield similar patterns.

Table 5 shows that the CRB heuristic provides high-quality solutions in most instances and offers a computationally attractive alternative to solving SAP-SEF directly. Across all reported instances, the average and median gaps between the CRB heuristic and the SAP-SEF incumbent solutions are 4.77% and 1.59%, respectively. The heuristic performs worst under the increasing supply profile, with an average gap of 11.42%. In contrast, its performance is substantially stronger under the flat and decreasing supply profiles, with average gaps of 1.76% and 1.13%, respectively. These results are consistent with the structure of the heuristic. Because CRB constructs schedules by balancing cumulative scheduled demand against expected cumulative supply, it performs particularly well when supply is available earlier in the horizon or remains relatively stable over time. When supply is increasing, however, the optimal policy may benefit from more carefully delaying scheduled demand, making the first-stage scheduling decision more sensitive to the temporal supply pattern. We do not observe a statistically significant difference in heuristic performance between the high- and low-supply settings. The average gap also appears to decrease as the number of collectors increases, suggesting that the heuristic may benefit from greater scheduling flexibility in larger number of collectors.

8 Conclusion

This paper studies a two-stage humanitarian decision problem in which organizations must assign beneficiaries or collectors to pickup periods before supply is known and then determine allocation quantities after supply is realized. These schedules are often strategic and cannot be changed quickly, as they must

Table 5: Comparison of SAP-SEF incumbent solutions and the CRB heuristic under perfect equity.

N	Supply Level	Supply Profile	Z^{OPT}	OPT MIP Gap (%)	OPT Time (sec.)	Z^{H}	H Gap (%)
20	High	Flat	773.6	0.96	95	746.0	3.57
		Increasing	678.0	1.00	145	438.7	35.29
		Decreasing	890.2	0.95	53	876.7	1.53
	Low	Flat	257.3	0.99	182	249.2	3.17
		Increasing	226.2	0.99	204	141.1	37.60
		Decreasing	292.2	0.98	62	289.5	0.91
50	High	Flat	1886.6	1.41	600	1874.0	0.67
		Increasing	1653.1	1.35	600	1540.6	6.81
		Decreasing	2218.0	2.94	600	2177.7	1.82
	Low	Flat	626.0	1.93	600	621.5	0.72
		Increasing	561.9	1.28	600	539.0	4.08
		Decreasing	720.5	1.33	600	713.5	0.97
100	High	Flat	3784.5	3.17	600	3750.0	0.91
		Increasing	3291.8	2.92	600	3267.0	0.76
		Decreasing	4363.2	2.35	600	4333.5	0.68
	Low	Flat	1238.0	3.35	600	1217.6	1.65
		Increasing	1091.4	1.49	600	1067.2	2.22
		Decreasing	1444.4	3.00	600	1422.3	1.53
150	High	Flat	5811.1	3.00	600	5737.5	1.27
		Increasing	4973.1	2.80	600	4890.3	1.67
		Decreasing	6514.0	3.01	600	6449.6	0.99
	Low	Flat	1896.4	3.10	600	1855.6	2.15
		Increasing	1659.0	3.46	601	1610.3	2.94
		Decreasing	2155.3	3.05	600	2142.9	0.58
		Average					4.77
		Median					1.59
		Min					0.58
		Max					37.60
		St. Dev.					9.86

respect individuals' and families' work and school schedules. Therefore, schedule decisions must account for supply uncertainty, equity requirements, and the urgency of distributing items, while subsequent allocation decisions must balance timely distribution with fairness across beneficiaries. Although this decision problem is motivated by challenges faced by our food bank partners, it also arises in many other settings, including the distribution of aid donations by disaster relief organizations, the allocation of scarce vaccines by health centers, and the distribution of blood donations by blood banks.

To address this decision problem, we formulate a two-stage stochastic optimization model. The second-stage allocation problem under perfect equity is analyzed and we show that it admits a closed-form optimal solution. Hence, in settings where schedules are fixed, equitable allocation decisions can be implemented using simple decision-support tools such as spreadsheets, which are usually preferred by humanitarian organizations due to their simplicity and cost-effectiveness. We study structural properties of the full two-stage stochastic program, which is NP-hard, and derive, among other results, bounds on collector fill rates under both perfect and imperfect equity. These bounds provide humanitarian organizations with interpretable

best-case estimates on possible outcomes in terms of total supply usage, without solving an optimization problem. Building on the derived bounds, a simple heuristic for the perfect equity case is introduced which produces instantaneous (< 0.1sec) results for instances up to 150 collectors. Furthermore, the structural properties show that relaxing the equity requirement allows humanitarian organizations to become more efficient in terms of fully using available supply.

We conduct extensive numerical experiments using pseudo-realistic instances calibrated to US food pantry operations. Our results generate several managerial insights. First, the timing of supply arrivals plays a central role in allocation efficiency. When more supply arrives early in the planning horizon, more product can be distributed sooner, leading to higher objective values because earlier allocations receive greater weight and provide more opportunities for distribution. In contrast, when supply arrives later, distribution is delayed and the objective value declines. This finding highlights the operational value of early supply availability, even when total supply over the planning horizon is unchanged. Second, strict equity requirements can substantially restrict timely distribution. However, modest relaxations of equity, e.g., allowing approximately a 2% difference in collector fill rates, achieve the highest objective values in our experiments. This suggests that small deviations from perfect equity may provide meaningful gains in freshness and distribution efficiency while preserving a high degree of fairness. We examine the value of modeling uncertainty through the cost of supply variability, the expected value of perfect information (EVPI), and the value of the stochastic solution (VSS). The results show that supply uncertainty has the greatest impact when the model places strong emphasis on equity. Explicitly modeling uncertainty yields consistent, though moderate, improvements over the expected-value solution. The EVPI is larger, indicating that better information about future supply realizations can provide greater benefits than stochastic optimization alone. Increasing and random supply profiles exhibit the highest VSS and EVPI, underscoring the importance of temporal variability in determining the value of stochastic modeling and improved information. Finally, using our structural results, we develop a fast, interpretable and easy-to-implement heuristic that provide near-optimal solutions in most settings. We find that the heuristic performs best when supply arrives early in the horizon or is relatively flat, and as number of collectors increase.

There are several directions in which our work can be extended. First, in food bank settings, demand is often known to the distributing agencies because clients typically register to receive food assistance, and their needs may not change substantially over the short planning horizon considered in this paper. However, demand may be uncertain in other humanitarian contexts, such as disaster relief, or in settings where crises such as economic downturns or public health emergencies rapidly affect household needs. In these cases, incorporating demand uncertainty in addition to supply uncertainty would further improve the realism and applicability of the proposed models. Another potential avenue for future research is to explicitly account for beneficiary preferences in scheduling decisions, enabling better alignment between beneficiaries' needs and preferences and the food available at different points in time.

We would like to express our sincere gratitude to [acknowledge individuals, organizations, or institutions] for their invaluable contributions to this research. We are also grateful to [mention any additional acknowledgements, such as technical assistance, data providers, or colleagues] for their support and assistance throughout the course of this work.

References

Alkaabneh F, Diabat A, Gao HO (2021) A unified framework for efficient, effective, and fair resource allocation by food banks using an approximate fynamic programming approach. *Omega* 100:102300.

Azizi S, Bozkir CDC, Trapp AC, Kundakcioglu OE, Kurbanzade AK (2021) Aid allocation for camp-based and urban refugees with uncertain demand and replenishments. *Production and Operations Management* 30:4455–4474.

Berenguer G, Shen Z (2019) Om forum—challenges and strategies in managing nonprofit operations: An operations management perspective. *Manufacturing and Service Operations Management* 22(5):888–905.

Besiou M, Pedraza-Martinez AJ, Wassenhove LNV (2018) Or applied to humanitarian operations. *European Journal of Operational Research* 269:397–405.

Besiou M, van Wassenhove LN (2020) Humanitarian operations: A world of opportunity for relevant and impactful research. *Manufacturing and Service Operations Management* 22:135–145.

Birge JR, Louveaux F (2011) *Introduction to Stochastic Programming* (New York: Springer), 2 edition.

Davis LB, Sengul Orgut I, Jiang S, Aft E, Hale C, Morris L, Rykaczewski J (2024) Food bank operations: A us perspective on humanitarian food assistance. *Tutorials in Operations Research: Smarter Decisions for a Better World*, 322–357 (INFORMS).

Feeding America (2026) How feeding america turns $1 into at least 10 meals. `https://www.feedingamerica.org/ways-to-give/faq/about-our-claims`, Accessed 2025-05-18.

Fianu S, Davis LB (2018) A markov decision process model for equitable distribution of supplies under uncertainty. *European Journal of Operational Research* 264(3):1101–1115.

Gómez-Pantoja JÁ, Salazar-Aguilar MA, González-Velarde JL (2021) The food bank resource allocation problem. *Top* 29:266–286.

Hassanzadeh P, Kreacic E, Zeng S, Xiao Y, Ganesh S (2023) Sequential fair resource allocation under a markov decision process framework. *Proceedings of the Fourth ACM International Conference on AI in Finance*, 673–680.

Lien RW, Iravani SMR, Smilowitz KR (2014) Sequential resource allocation for nonprofit operations. *Operations Research* 62:301–317.

Ma Y, Wang T, Zheng H (2023) On fairness and efficiency in nonprofit operations: Dynamic resource allocations. *Production and Operations Management* 32(6):1778–1792.

Peters K, Silva S, Gonçalves R, Kavelj M, Fleuren H, den Hertog D, Ergun O, Freeman M (2021) The nutritious supply chain: Optimizing humanitarian food assistance. *INFORMS Journal on Optimization* 3:200–226.

Petropoulos F, et al. (2024) Operational research: methods and applications. *Journal of the Operational Research Society* 75:423–617.

Reusken M, Cruijssen F, Fleuren H (2023) A food bank supply chain model: Optimizing investments to maximize food assistance. *International Journal of Production Economics* 261:108886.

Sengul Orgut I, Brock III LG, Davis LB, Ivy JS, Jiang S, Morgan SD, Uzsoy R, Hale C, Middleton E (2016a) Achieving equity, effectiveness, and efficiency in food bank operations: Strategies for feeding america with implications for global hunger relief. *Advances in managing humanitarian operations* 229–256.

Sengul Orgut I, Davis LB, Ivy JS (2025) Food bank supply chain and operations in a data-driven world: objectives, challenges, and opportunities. *Nonprofit Operations and Supply Chain Management: Theory and Practice* 265–304.

Sengul Orgut I, Ivy J, Uzsoy R (2017) Modeling for the equitable and effective distribution of food donations under stochastic receiving capacities. *IISE Transactions* 49(6):567–578.

Sengul Orgut I, Ivy J, Uzsoy R, Wilson JR (2016b) Modeling for the equitable and effective distribution of donated food under capacity constraints. *IIE Transactions* 48(3):252–266.

Sengul Orgut I, Ivy JS, Uzsoy R, Hale C (2018) Robust optimization approaches for the equitable and effective distribution of donated food. *European Journal of Operational Research* 269(2):516–531.

Sengul Orgut I, Lodree EJ (2023) Equitable distribution of perishable items in a food bank supply chain. *Production and Operations Management* 32(10):3002–3021.

Sinclair SR, Banerjee S, Yu CL (2022) Sequential fair allocation. *Performance Evaluation Review*, volume 50, 95–96 (Association for Computing Machinery).

Statista Research Department (2025) Distribution of households in the united states from 1970 to 2024, by household size. `https://www.statista.com/statistics/242189/disitribution-of-households-in-the-us-by-household-size/`, Accessed 2025-04-22.

Stauffer JM, Vanajakumari M, Kumar S, Mangapora T (2022) Achieving equitable food security: How can food bank mobile pantries fill this humanitarian need. *Production and Operations Management* 31(4):1802–1821.

Swaminathan JM (2025) Om forum—operations management research: Relevance and impact. *Manufacturing and Service Operations Management* 27:1–7.

Tirkolaee EB, Aydın NS, Ranjbar-Bourani M, Weber GW (2020) A robust bi-objective mathematical model for disaster rescue units allocation and scheduling with learning effect. *Computers and Industrial Engineering* 149.

USDA (2023) Food security in the u.s. `https://www.ers.usda.gov/topics/food-nutrition-assistance/food-security-in-the-u-s/`, Accessed 2024-1-26.

Vries HD, Wassenhove LNV (2020) Do optimization models for humanitarian operations need a paradigm shift? *Production and Operations Management* 29:55–61.

Online Appendices for

"Scheduling and Allocation of Uncertain, Scarce, Equitable, and Perishable Supply in Humanitarian Contexts"

A Alternative SAP-SEF formulation with a closed-form second stage

Based on the findings in Lemma 1 and Proposition 1, we arrive at a possible alternative formulation for the SAP-SEF under perfect equity by taking the first stage and plugging in the closed-form solution for the second stage, resulting in:

$$\max \quad \mathbb{E}_{\boldsymbol{\xi}}\left[\sum_{t=1}^{T}(T-t+1)\sum_{n=1}^{N}Y_{nt}^{*}\right] \tag{A.1}$$

$$\text{s.t.} \quad Y_{nt}^{*}=\begin{cases} D_n R & \text{if} \quad X_{nt}=1 \\ 0 & \text{otherwise} \end{cases} \qquad n\in\mathcal{N}, t\in\mathcal{T} \tag{A.2}$$

$$\sum_{t=1}^{T}X_{nt}=1 \qquad n\in\mathcal{N}$$

$$X_{nt}\in\{0,1\} \qquad n\in\mathcal{N}, t\in\mathcal{T}.$$

We now study the stochastic objective and the tractability of this problem through rewriting (A.1) and (A.2). First, we use (16) and (17), to rewrite objective (A.1) as:

$$\mathbb{E}_{\boldsymbol{\xi}}\left[\sum_{t=1}^{T}(T-t+1)\sum_{n=1}^{N}\frac{\tilde{\tilde{Y}}^{*}}{\Delta}D_n X_{nt}\right]. \tag{A.3}$$

Since Y_{nt}^{*} and $\tilde{\tilde{Y}}^{*}$ depend on the random variable ξ_t, they are both random variables themselves and we can apply the sum property (the expectation of the sum of random variables is equal to the sum of their expectations) and the scale property (if you multiply a random variable by a constant, its expectation scales linearly), implying that (A.3) equals:

$$\begin{aligned} &\sum_{t=1}^{T}(T-t+1)\sum_{n=1}^{N}\mathbb{E}_{\boldsymbol{\xi}}\left[\frac{\tilde{\tilde{Y}}^{*}}{\Delta}D_n X_{nt}\right] \\ =&\sum_{t=1}^{T}(T-t+1)\sum_{n=1}^{N}\mathbb{E}_{\boldsymbol{\xi}}\left[\frac{1}{\Delta}D_n X_{nt}\Delta R\right], \end{aligned} \tag{A.4}$$

where the equality follows from applying Proposition 1. Then we use that $\sum_{n\in\cup_{i=1}^{t}\{N\}_i}D_n=\sum_{n=1}^{N}\sum_{i=1}^{t}D_n X_{nt}$, thus making (A.4) equivalent to:

$$\sum_{t=1}^{T}(T-t+1)\sum_{n=1}^{N}\mathbb{E}_{\boldsymbol{\xi}}\left[D_n X_{nt}\min_{t\in\mathcal{T}}\frac{\sum_{i=1}^{t}\xi_i}{\sum_{n=1}^{N}\sum_{i=1}^{t}D_n X_{ni}}\right]. \tag{A.5}$$

This formulation of the objective involves an expectation over a minimization, which is not tractable in its current form. Therefore, we reformulate the minimum operator through introducing auxiliary variable $Z\in\mathbb{R}$, allowing the following reformulation of (A.5):

$$\sum_{t=1}^{T}(T-t+1)\sum_{n=1}^{N}\mathbb{E}_{\boldsymbol{\xi}}\left[D_n X_{nt} Z\right],$$

with added constraints

$$Z \leq \frac{\sum_{i=1}^{t}\xi_i}{\sum_{n=1}^{N}\sum_{i=1}^{t} D_n X_{ni}} \quad t \in \mathcal{T}.$$

We thus present the following alternative SAP-SEP formulation, which is exactly equivalent to the original under perfect equity:

$$\begin{aligned}
\max \quad & \sum_{t=1}^{T}(T-t+1)\sum_{n=1}^{N}\mathbb{E}_{\boldsymbol{\xi}}[D_n X_{nt} Z] \\
\text{s.t.} \quad & Z \leq \frac{\sum_{i=1}^{t}\xi_i}{\sum_{n=1}^{N}\sum_{i=1}^{t} D_n X_{ni}} && t \in \mathcal{T} \\
& \sum_{t=1}^{T} X_{nt} = 1 && n \in \mathcal{N} \\
& X_{nt} \in \{0,1\} && n \in \mathcal{N}, t \in \mathcal{T}.
\end{aligned}$$

Compared to [SAP-SEF], this formulation replaces the second-stage optimization with an explicit function of $\boldsymbol{\xi}$ and $\mathbf{X}$, and thereby eliminating the recourse problem, converting the problem from a two-stage stochastic MILP into a single-stage stochastic integer program. This represents a simplification in terms of model structure, as it involves only one stage, fewer decision variables due to the elimination of the $\mathbf{Y}$ variables, and no nested optimization. However, note that this simplification comes at the expense of the bilinear term $X_{nt}Z$ in the objective, where $\mathbf{X}$ is binary and Z nonnegative and continuous. This term can be linearized exactly using standard techniques, such as McCormick-type linearization, yielding an equivalent MILP. Such linearizations are handled efficiently by modern solvers such as Gurobi, which apply them internally. In addition to potential computational advantages in specific cases, the alternative formulation is valuable as it provides conceptual insight by revealing an alternative structure. In particular, this more compact representation can improve interpretability by exposing a simpler and more direct relationship between decision stages.

B Proof of Corollary 1

Proof of Corollary 1. Let k denote the collector with the highest fill rate, i.e., $r_k = \max_{n\in\mathcal{N}} r_n$. The equity constraint implies $r_n \geq r_k - \theta$ for $n \neq k$. We consider two scenarios in our proof:

(i) Case 1: $r_k \geq \theta$: In this case, the minimum total supply required to support this r_k occurs when all other collectors receive the smallest fill rate allowed by the equity constraint, i.e., $r_n = r_k - \theta$ for all $n \neq k$. The resulting total allocation is therefore $D_k r_k + \sum_{n\neq k} D_n(r_k - \theta)$. Constraint (11) requires that the total allocation does not exceed the available supply, i.e., $D_k r_k + \sum_{n\neq k} D_n(r_k - \theta) \leq \sum_{t=1}^{T}\xi_t$. Rearranging gives $D_k r_k + \sum_{n\neq k} D_n r_k - \sum_{n\neq k} D_n\theta \leq \sum_{t=1}^{T}\xi_t \Leftrightarrow r_k \sum_{n=1}^{N} D_n \leq \sum_{t=1}^{T}\xi_t + \theta\sum_{n\neq k} D_n$. Solving for r_k yields

$$r_k \leq \frac{\sum_{t=1}^{T}\xi_t + \theta\sum_{n\neq k} D_n}{\sum_{n=1}^{N} D_n}.$$

Using the definition of $\bar{R}_{\theta=0}$ and $\sum_{n\neq k} D_n = \sum_{n=1}^{N} D_n - D_k$ yields

$$r_k \leq \bar{R}_{\theta=0} + \theta\left(1 - \frac{D_k}{\sum_{n=1}^{N} D_n}\right).$$

(i) Case 2: $r_k < \theta$: In this case, the smallest feasible fill rate for all the other collectors ($n \neq k$) is zero since we have $r_n \geq \max(0, r_k - \theta)$. Then, the resulting minimum total allocation becomes $D_k r_k$. From constraint (11), we get $D_k r_k \leq \sum_{t=1}^{T} \xi_t$, which yields $r_k \leq \frac{\sum_{t=1}^{T} \xi_t}{D_k}$. Using the definition of $\bar{R}_{\theta=0}$, we can write this as: $r_k \leq \bar{R}_{\theta=0} + \sum_{t=1}^{T} \xi_t(\frac{1}{D_k} - \frac{1}{\Delta})$. This completes the proof.

These bounds are attainable by allocating r_k to collector k and $r_k - \theta$ to all other collectors (Case 1) and zero to all the other collectors (Case 2), which satisfies both the equity constraint and the supply constraint. Hence the bounds are tight. □

C Proof of Proposition 5

Proof of Proposition 5. First, we note that the fill rate r_n for any n takes on values between zero and one, due to constraint (11) and Assumption 1. Since the absolute difference between two fractions that can only take on values between zero and one can at most be one, setting $\theta \geq 1$ results in the equity constraint (9) being no longer binding, and hence, can be dropped. Let us rewrite the second-stage model when $\theta \geq 1$:

$$\begin{aligned}
\max \quad & \sum_{t=1}^{T} (T - t + 1) \sum_{n=1}^{N} Y_{nt} \\
\text{s.t.} \quad & D_n X_{nt} \geq Y_{nt} && n \in \mathcal{N}, t \in \mathcal{T} && \text{(A.6)} \\
& \sum_{i=1}^{t} \hat{\xi}_i \geq \sum_{i=1}^{t} \sum_{n=1}^{N} Y_{ni} && t \in \mathcal{T} && \text{(A.7)} \\
& Y_{nt} \in \mathbb{R}_+ && n \in \mathcal{N}, t \in \mathcal{T}. && \text{(A.8)}
\end{aligned}$$

Consider the following solution for the second-stage variables:

$$Y_{nt} = \begin{cases} \frac{\xi_t D_n}{\sum_{m \in \{\mathcal{N}\}_t} D_m} & \text{if } X_{nt} = 1 \\ 0 & \text{otherwise.} \end{cases} \tag{A.9}$$

We will first show that this solution satisfies all the constraints. For constraint (A.6), if $X_{nt} = 0$ for any n and t, it follows from (A.9) that $Y_{nt} = 0$, satisfying constraint (A.6). If $X_{nt} = 1$ for any n and t, then, the left hand side of constraint (A.6) becomes D_n. Due to assumption $\xi_t < \sum_{m \in \{\mathcal{N}\}_t} D_m$, this constraint is satisfied. Next, let us consider constraint (A.7). For any time period t and all the collectors scheduled in that period $\{\mathcal{N}\}_t$ we get:

$$\begin{aligned}
\sum_{n \in \{\mathcal{N}\}_t} Y_{nt} &= \sum_{n \in \{\mathcal{N}\}_t} \frac{\xi_t D_n}{\sum_{m \in \{\mathcal{N}\}_t} D_m} \\
&= \xi_t \frac{\sum_{n \in \{\mathcal{N}\}_t} D_n}{\sum_{m \in \{\mathcal{N}\}_t} D_m} = \xi_t.
\end{aligned}$$

Hence, the proposed solution satisfies constraint (A.7). Furthermore, it is trivial to see that the nonnegativity constraint (A.8) is satisfied.

Finally, it is trivial to see that this is indeed an optimal solution since all the supply is distributed on the period on which it is received, which is the earliest possible moment. Hence, this shows that, when $\theta \geq 1$,

(A.9) represents the closed-form optimal solution for this problem, and that the distribution is equal to total available supply. This completes the proof. □