Proceedings of the 2022 Winter Simulation Conference
B. Feng, G. Pedrielli, Y. Peng, S. Shashaani, E. Song, C.G. Corlu, L.H. Lee, E.P. Chew, T. Roeder, and P. Lendermann, eds.

FROM EFFICIENCY TO FAIRNESS: DESIGN OF ALLOCATION RULES FOR FOOD BANK OPERATIONS

Jinpeng Liang

School of Transportation Engineering
Dalian Maritime University
1 Linghai Rd, Ganjingzi District
Dalian, Liaoning 116026, P. R. CHINA

Guodong Lyu

School of Business and Management
Hong Kong University of Science and Technology
Clear Water Bay
Kowloon, HONG KONG

ABSTRACT

Food banks play an essential role in alleviating world hunger by allocating surplus food to eligible agencies or individuals. As non-profit organizations, food banks target agencies to achieve operational efficiency of food allocation (i.e., reduce food waste). However, this would result in inequitable service delivered to different agencies. In this work, we design real-time food allocation rules to serve the sequentially revealed demand of each targeted agency, and to ensure that adequate food is allocated to each agency (fairness) as much as possible. We measure allocation fairness by fill rate (i.e., the ratio of the allocated amount to revealed demand) and exploit online convex optimization tools to characterize the attainable fill rate of the agency. We use these insights to develop provably near-optimal allocation rules for food bank operations, and leverage on extensive numerical simulations to discuss the promising benefits of our allocation rules over the existing benchmark.

1 INTRODUCTION

Food is at the core of the United Nations' sustainable development agenda for the 21st century. As reported by the State of Food Security and Nutrition in the World 2021 report, despite some progress, over 800 million people across the globe suffered from hunger in 2020. Under the shadow of the COVID-19 pandemic, the situation continues to deteriorate. Hunger is not only the pain point of developing countries, millions of households in developed countries (e.g., America, Singapore, and Britain) also go to bed hungry. The paradox of scarcity in abundance is a consequence of food waste and unbalanced distribution. Estimates suggest that one-third of global food gets lost or wasted, while recycling and resharing just 25% of those leftovers would suffice to eradicate world hunger (The State of Global Food Banking 2018: Nourishing the World). To this end, the food bank model, which aims at matching surplus food with eligible agencies or individuals, provides an effective solution to hunger relief. Since the establishment of the world's first food bank in 1967, more than 1500 food banks have been set up across the world to prevent food waste and alleviate hunger. In 2019, the world's major food banks redirected 3.75 million metric tons of surplus food from landfills to more than 66.5 million beneficiaries worldwide (Advancing the Sustainable Development Goals: Roadmap to 2030).

In practice, the food banks operate a fleet of delivery trucks for food allocation and follow predetermined routes to visit the eligible agencies sequentially. However, due to the lack of personnel at different agencies and the sequential manner of food delivery, the demand information of each agency is revealed only upon the arrival of the truck. Consequently, the planner needs to make an irrevocable allocation decision to serve the demand required by each agency, while the random demand of unvisited agencies is unknown. To meet the demand of each agency as much as possible, the dynamic resource (food) allocation is at the heart of food bank operations. A great deal of research effort has been made to construct delicate resource allocation rules

978-1-6654-7661-4/22/$31.00 ©2022 IEEE

to maximize revenue or minimize cost (Talluri and Ryzin 2004). However, although operational efficiency (e.g., reducing food waste) is desirable for food banks, such efficiency-oriented allocation policies may possibly induce inequitable service experiences for different agencies. Notably, allocation fairness, which is measured by service equity, also plays a paramount role in food bank operations to deliver sustainable service (Lien et al. 2014). Indeed, fairness is one of the six critical issues faced by a nonprofit organization (Berenguer and Shen 2020), and the fairness concern becomes even more significant during this COVID-19 period when millions of vulnerable people are pushed into greater food insecurity (Manshadi et al. 2021; Ma et al. 2022).

In this paper, we aim to design real-time food allocation rules to serve the sequentially revealed demand of the targeted agencies, and to ensure that adequate food is allocated to each agency as much as possible (fairness). We use fill rate, which is defined as the ratio of the allocated amount to revealed demand, to measure allocation fairness, and our objective is to maximize the minimal expected fill rate attained across different agencies. We note that this fairness objective is in line with the spirit of Rawlsian justice that seeks to maximize the worst-off performance over all the agencies in a system (Manshadi et al. 2021). To achieve this goal, we first theoretically characterize the achievable region of fill rate for all agencies, and then numerically determine the maximal fill rate target. Next, we use these insights to design provably near-optimal allocation rules for food bank operations. Finally, we simulate extensive numerical scenarios to show that our allocation rule can effectively protect the most vulnerable agencies from food shortage, without scarifying too much operational efficiency.

The rest of this paper is organized as follows: In Section 2, we review the related works in the literature. We construct the allocation model and develop our allocation rules in Section 3. The numerical simulations are presented in Section 4. Section 5 concludes the paper. Some preliminary results to support the technical proof are relegated to Appendix A.

2 LITERATURE REVIEW

Our study is closely related to the research topics: food bank operations to alleviate hunger, and resource allocation to ensure fairness.

Food bank operations. The operations of food banks have piqued a surge of interest in academia. At a tactical level, the efficiency of food bank operations relies on the design of vehicle routing solutions to achieve cost-effective performance for food collection and delivery (Bartholdi et al. 1983; Gunes 2010; Lee et al. 2017). However, the aforementioned works did not address the issue of food allocation to each agency, which significantly affects the performance of food bank operations in terms of demand fulfillment. Since the food resource is in general scarce, the service experience delivered to the individual agency may vary considerably under different allocation policies. To this end, the problem of equitable and efficient food allocation through simultaneous routing and allocation decisions was studied by Nair et al. (2017) and Eisenhandler and Tzur (2019). We note that the food allocation problem to ensure equity and efficiency was also considered in a static setting (Orgut et al. 2016; Orgut et al. 2017; Orgut et al. 2018). Unfortunately, these static allocation solutions could not be used for real-time allocation decisions when the agencies' demand information reveals sequentially. To bridge the gap, as far as we know, Lien et al. (2014) was one of the earliest papers to study the fair resource allocation in a sequential setting where the food delivery truck visits the agencies under fixed routes and allocates a certain amount of food to each agency upon the visit to this agency. The goal was to maximize the expected minimum fill rate of all agencies, and this problem was extended to a multi-vehicle setting by Balcik et al. (2014). In this paper, we consider a similar setting as Lien et al. (2014) but focus on maximizing the minimal expected fill rate for all agencies.

Fair resource allocation. The notion of equity and fairness has been encountered in diverse contexts (Savas 1978; Brill 1979; Marsh and Schilling 1994), but the measurement of fairness varies under different allocation settings (Bertsimas et al. 2011). For example, a stream of research papers introduced the service level target to enforce operational fairness in supply chain management (Zhong et al. 2018; Lyu et al. 2019) and online platform (Lyu et al. 2019), etc. The aforementioned work addressed the resource allocation problem in a simultaneous allocation setting, in which the allocation decision is determined after knowing the realized demand of all agencies. In contrast, Zheng et al. (2022) designed static decision rules to solve the service level constrained cash-transfer problem without knowing the realized information a priori. Along this direction, various types of fair allocation problems have also been studied in a sequential setting (Lien et al. 2014; Manshadi et al. 2021; Ma et al. 2022). In particular, Manshadi et al. (2021) considered two types of fairness measures in terms of the expected fill rate, i.e., maximizing the expected minimum fill rate (ex-post fairness) and maximizing the minimum expected fill rate (ex-ante fairness). Similarly, we aim to ensure ex-ante fairness in the food allocation problem. The novelty of our work is to theoretically characterize the achievable region of ex-ante fill rate target to all agencies. Furthermore, we design provably near-optimal allocation policies to achieve the predetermined target, and design heuristic policies to boost the computational efficiency for large-scale problems. The numerical results show that our heuristic allocation policy outperforms the target fill rate policy proposed by Manshadi et al. (2021).

3 MODEL FORMULATION AND ALLOCATION POLICY

3.1 Model Formulation

We consider a food bank program in which one planner follows a fixed route to allocate a certain amount of food to N agencies with stochastic demand. The agencies are indexed as $\{1,\ldots,N\}$ according to the allocation sequence. We assume that the demands of agencies are drawn from a joint distribution, while their demands are sequentially revealed to the planner upon the visit to each agency. Let $\boldsymbol{d}$ denote the demand vector of all agencies with non-negative support set $\Omega_{\boldsymbol{d}}$, and d_i denote the demand of agency i. To enforce allocation fairness, we use fill rate (the ratio of allocated amount to realized demand) to quantify the service level delivered to each agency. To this end, the planner needs to sequentially determine the proportion of demand to be served for each agency so that the minimal expected fill rate among all agencies can be maximized.

Furthermore, we let s_i denote the remaining amount of food when the planner arrives at agency i and we have $s_1 = c$, where c represents the initial capacity level of food. Upon the visit to each agency i, the planner observes realized demand $d_i \in \mathbb{R}^+$ as well as available amount of food $s_i \in \mathbb{R}^+$, and then decides the proportion of demand fulfillment $x_i(s_i,\boldsymbol{d})$ for agency i to balance the trade-off between satisfying immediate demand d_i and reserving resources for the unvisited agencies. For the sake of feasibility, the allocation amount $x_i(s_i,\boldsymbol{d})\cdot d_i$ for agency i cannot exceed remaining quantity s_i nor can it exceed the realized demand d_i. More formally, we formulate the feasible region of $x_i(s_i,\boldsymbol{d})$ as follows:

$$X_i(s_i,\boldsymbol{d}) := \left\{ x_i \in \mathbb{R}^+ \;\middle|\; \begin{array}{l} 0 \le x_i(s_i,\boldsymbol{d}) \le 1, \\ x_i(s_i,\boldsymbol{d})\cdot d_i \le s_i \end{array} \right\}.$$

We say that the decision is feasible if $x_i(s_i,\boldsymbol{d}) \in X_i(s_i,\boldsymbol{d})$ for all $i = 1,\ldots,N$, and write $\boldsymbol{x} \in \boldsymbol{X}$ to denote that $\boldsymbol{x}$ is feasible in short. After determining the fill rate $x_i(s_i,\boldsymbol{d})$ for agency i, the amount of available food for agency $i+1$ is updated as:

$$s_{i+1} = s_i - x_i(s_i,\boldsymbol{d})\cdot d_i,\ \forall i = 1,\ldots,N-1.$$

It is evident that the straightforward first-come-first-serve policy that allocates all of the resources to satisfy the demand of early-visited agencies could achieve the highest efficiency (e.g., food waste minimization), but may induce unfair service experience among different agencies. This paper studies fair resource allocation policies for food bank operations that aim to maximize the minimal expected fill rate across all agencies. We formulate the fair resource allocation problem as follows:

$$
\begin{aligned}
\text{(P1)} \quad \max_{\boldsymbol{x}(s_i,\boldsymbol{d}),\theta} \quad & \theta \\
\text{s.t.} \quad & \mathbf{E}_{\boldsymbol{d}}\left[x_i(s_i,\boldsymbol{d})\right] \geq \theta,\ \forall i=1,\ldots,N \\
& x_i(s_i,\boldsymbol{d}) \in X(s_i,\boldsymbol{d}),\ \forall i=1,\ldots,N,\ \boldsymbol{d}\in\Omega_{\boldsymbol{d}} \\
& s_{i+1} = s_i - x_i(s_i,\boldsymbol{d})\cdot d_i,\ \forall i=1,\ldots,N-1,\ \boldsymbol{d}\in\Omega_{\boldsymbol{d}} \\
& x_i(s_i,\boldsymbol{d}) \text{ non-anticipatory},\ \forall i=1,\ldots,N.
\end{aligned}
$$

where the first set of constraints requires that the expected fill rate delivered to each and every agency is no less than θ, while our objective is to maximize the value of θ. For the last set of constraints, we ensure that the allocation decision is made in an online fashion, i.e., the decision for each agency i cannot use the unrevealed information from agency $i+1$ to N. This problem is challenging partly because the planner needs to sequentially build up the solution $x_i(s_i,\boldsymbol{d})$ for each agency i without full visibility of future demands. Furthermore, the allocation decision has to be jointly optimized with the minimal expected fill rate θ. To solve this problem, we first consider the following feasibility problem: given a fixed fill rate target θ, can we find a sequential allocation policy that can satisfy the fill rate constraints in Problem (P1)? We derive a set of necessary and sufficient conditions that only involve the demand fulfillment decision $\boldsymbol{x}\in\boldsymbol{X}$ and the fill rate target θ.

3.2 Necessary and Sufficient Conditions

In this subsection, we derive a set of necessary and sufficient conditions for a given fill rate target θ to be feasible. For ease of exposition, we use $x_i^{\pi}(s_i,\boldsymbol{d})$ to denote the fill rate decision for agency i under a feasible allocation policy π, while the collection of all feasible policies is denoted as Π. Next, we present the following Theorem 1 to describe the necessary and sufficient conditions.

Theorem 1 The expected fill rate target θ in the fair resource allocation Problem (P1) can be attained if and only if

$$
\max_{\boldsymbol{\lambda}\geq\mathbf{0}}\left\{\sum_{i=1}^{N}\lambda_i\theta - \max_{\pi\in\Pi}\mathbf{E}\left[\sum_{i=1}^{N}\lambda_i x_i^{\pi}(s_i,\boldsymbol{d})\right]\right\} \leq 0. \tag{1}
$$

Theorem 1 implies that the fill rate target θ is attainable if and only if $\max_{\pi\in\Pi}\mathbf{E}\left[\sum_{i=1}^{N}\lambda_i x_i^{\pi}(s_i,\boldsymbol{d})\right] \geq \sum_{i=1}^{N}\lambda_i\theta$ for any non-negative vector $\boldsymbol{\lambda}\geq\mathbf{0}$. Here the vector $\boldsymbol{\lambda}$ represents the priority weight attached to each agency when determining the allocation solution. Indeed, it is not easy to obtain any closed-form expressions for the above necessary and sufficient conditions, which concern complicated resource allocation decisions in food bank operations. Interestingly, we can leverage on the above results to evaluate the feasibility of a given θ and numerically characterize the largest feasible fill rate target for all agencies.

3.2.1 Necessary Conditions.

For any feasible sequential allocation policy $A\in\Pi$ that satisfies the fill rate requirement, we have the following conditions:

$$
\theta - \mathbf{E}_{\boldsymbol{d}}\left[x_i^{A}(s_i,\boldsymbol{d})\right] \leq 0, \quad \forall i=1,\ldots,N. \tag{2}
$$

By taking a linear combination with non-negative vector $\boldsymbol{\lambda}\geq\mathbf{0}$ among all agencies in Equation (2), we can conclude the necessity part of conditions (1) in Theorem 1.

3.2.2 Sufficient Conditions.

Note that Problem (P1) imposes the "expected" fill rate requirement in the single-period stochastic setting. It is straightforward to interpret the stochastic constraints by generating infinite i.i.d demand samples $\{\boldsymbol{d}_t\}_{t=1}^{T}$, where $\boldsymbol{d}(t)$ denotes the demand sample at each epoch t and the sample size is sufficiently large enough, i.e., $T \to \infty$. To this end, we explicitly construct a debt-associated-sequential (DAS) allocation policy for the multi-sample problem.

The key of our DAS allocation policy is the notion *debt*, which quantifies the gap between the attained average fill rate and the requirement θ for each agency. To be more specific, let $B_i(t)$ denote the debt for agency i at sampling epoch t, say the deficit owe to the fill rate target θ under the demand fulfillment decision $x_i(t)$:

$$B_i(t) := \theta - x_i(t).$$

With slight abuse of notation, we let $x_i(t)$ denote the allocation decision to agency i at epoch t, and $X_i(t)$ represents the corresponding feasible region. At the beginning of sampling epoch $(t+1)$, we let $\beta_i(t+1)$ denote the average debt for agency i, which is accumulated from sampling epoch 1 to t:

$$\beta_i(t+1) := \frac{1}{t}\sum_{s=1}^{t} B_i(s).$$

Intuitively speaking, the sign of average debt $\beta_i(t+1)$ implies whether the expected fill rate attained by agency i over the first t sampling epochs exceeds the target θ or not. More concretely, when $\beta_i(t+1) > 0$, the attained fill rate over the first t epochs has not achieved the target θ for agency i. However, when $\beta_i(t+1) \leq 0$, it means that the attained fill rate has already met the target θ for agency i. The key idea of our DAS policy is to prioritize agencies with positive average debt when making allocation decisions at sampling epoch $(t+1)$ so that the expected fill rate target θ can be attained for all agencies in the long run. Among the agencies with positive average debt, the agencies with larger positive average debt shall be given higher priority to be served at epoch $(t+1)$. In contrast, the agencies with negative debt shall not be considered unless there is idle resource after serving those with positive average debt. In this way, we set $\beta_i^+(t+1) = \max\{\beta_i(t+1), 0\}$ as the weight associated with agency i when determining the allocation decision $\boldsymbol{x}(t+1)$ at sampling epoch $(t+1)$. Altogether, our DAS policy at epoch $(t+1)$ aims to maximize the following weighted-sum stochastic dynamic programming (DP) problem:

$$f(\boldsymbol{\beta}^+(t+1)) = \mathbf{E}_{\boldsymbol{d}(t+1)}\left[\max_{\boldsymbol{x}(t+1)\in\boldsymbol{X}(t+1)} \sum_{i=1}^{N} \beta_i^+(t+1)\cdot x_i(t+1)\right]. \tag{3}$$

where $f(\boldsymbol{\beta}^+(t+1))$ denotes the optimal value of the DP model (3) at epoch $(t+1)$. To this end, our DAS policy decomposes the fill rate constrained problem into a series of DP problems. In what follows, we show that our DAS policy can deliver the expected fill rate target for each and every agency as long as conditions (1) hold. This result is motivated by tracking the evolution of debt vector under the DAS policy. Let $\boldsymbol{\beta}(0) = \mathbf{0}$ for notational convenience. For each sampling epoch $0 \leq t \leq T-1$, we have

$$\begin{aligned}
&\left\|\boldsymbol{\beta}^+(t+1)\right\|_2^2 \\
\leq& \left\|\boldsymbol{\beta}^+(t)\right\|_2^2 + 2\boldsymbol{\beta}^+(t)^\top[\boldsymbol{\beta}(t+1)-\boldsymbol{\beta}(t)] + \|\boldsymbol{\beta}(t+1)-\boldsymbol{\beta}(t)\|_2^2 && (4)\\
=& \left\|\boldsymbol{\beta}^+(t)\right\|_2^2 + 2\boldsymbol{\beta}^+(t)^\top\frac{\boldsymbol{B}(t)-\boldsymbol{\beta}(t)}{t} + \frac{\|\boldsymbol{B}(t)-\boldsymbol{\beta}(t)\|_2^2}{t^2} && (5)\\
\leq& \left\|\boldsymbol{\beta}^+(t)\right\|_2^2 + \frac{2\boldsymbol{\beta}^+(t)^\top\{\boldsymbol{B}(t)-\boldsymbol{\beta}(t)\}}{t} + \frac{N}{t^2} && (6)\\
=& \left\|\boldsymbol{\beta}^+(t)\right\|_2^2 + \frac{2\boldsymbol{\beta}^+(t)^\top\{\mathbf{E}[\boldsymbol{B}(t)|\boldsymbol{\beta}(t)-\boldsymbol{\beta}(t)]\}}{t} + \frac{2\boldsymbol{\beta}^+(t)^\top\{\boldsymbol{B}(t)-\mathbf{E}[\boldsymbol{B}(t)|\boldsymbol{\beta}(t)]\}}{t} + \frac{N}{t^2} && (7)
\end{aligned}$$

Step (4) is due to the inequality in the Claim 1 of Appendix A. According to the definition of $\boldsymbol{B}(t)$ and $\boldsymbol{\beta}(t)$, it is straightforward to see that $\boldsymbol{\beta}(t+1)-\boldsymbol{\beta}(t)=\frac{1}{t}\left[\boldsymbol{B}(t)-\boldsymbol{\beta}(t)\right]$, which justifies Step (5). In addition, we have

$$|B_i(t)-\beta_i(t)| = \left|\frac{1}{t-1}\sum_{s=1}^{t-1} x_i(s) - x_i(t)\right| \leq 1,\ \forall i = 1,\ldots,N$$

which facilitate us to bound the term $\|\boldsymbol{B}(t)-\boldsymbol{\beta}(t)\|_2^2$ with N in Step (6). Finally, we decompose the term $\{\boldsymbol{B}(t)-\boldsymbol{\beta}(t)\}$ by introducing the conditional expectation $\mathbf{E}\left[\boldsymbol{B}(t)|\boldsymbol{\beta}(t)\right]$ to arrive at Step (7). By definition, we have $\boldsymbol{\beta}^+(t)^\top\boldsymbol{\beta}(t) = \left\|\boldsymbol{\beta}^+(t)\right\|_2^2 \geq 0$. We then derive the following Step (8) and (9):

$$\begin{aligned}
&\left\|\boldsymbol{\beta}^+(t+1)\right\|_2^2 \\
\leq &\left\|\boldsymbol{\beta}^+(t)\right\|_2^2 + \frac{2\boldsymbol{\beta}^+(t)^\top\{\mathbf{E}\left[\boldsymbol{B}(t)|\boldsymbol{\beta}(t)\right]\}}{t} - \frac{\boldsymbol{\beta}^+(t)^\top\boldsymbol{\beta}(t)}{t} + \frac{2\boldsymbol{\beta}^+(t)^\top\{\boldsymbol{B}(t)-\mathbf{E}\left[\boldsymbol{B}(t)|\boldsymbol{\beta}(t)\right]\}}{t} + \frac{N}{t^2} \quad (8)\\
= &\frac{t-1}{t}\left\|\boldsymbol{\beta}^+(t)\right\|_2^2 + \frac{2\boldsymbol{\beta}^+(t)^\top\{\mathbf{E}\left[\boldsymbol{B}(t)|\boldsymbol{\beta}(t)\right]\}}{t} + \frac{2\boldsymbol{\beta}^+(t)^\top\{\boldsymbol{B}(t)-\mathbf{E}\left[\boldsymbol{B}(t)|\boldsymbol{\beta}(t)\right]\}}{t} + \frac{N}{t^2} \quad (9)\\
\leq &\frac{t-1}{t}\left\|\boldsymbol{\beta}^+(t)\right\|_2^2 + \frac{2\boldsymbol{\beta}^+(t)^\top\{\boldsymbol{B}(t)-\mathbf{E}\left[\boldsymbol{B}(t)|\boldsymbol{\beta}(t)\right]\}}{t} + \frac{N}{t^2} \quad (10)
\end{aligned}$$

where Step (10) is due to the inequality $\boldsymbol{\beta}^+(t)^\top\{\mathbf{E}\left[\boldsymbol{B}(t)|\boldsymbol{\beta}(t)\right]\} \leq 0$, which is proved in Lemma 2. This result is motivated by exploiting the convexity of $f(\boldsymbol{\beta}^+(t+1))$ (see Lemma 1) to derive its Fenchel dual. In what follows, we denote $W(t) := t\left\|\boldsymbol{\beta}^+(t+1)\right\|_2^2$ and restate Step (10) as follows:

$$W(t) \leq W(t-1) + 2\boldsymbol{\beta}^+(t)^\top\{\boldsymbol{B}(t)-\mathbf{E}\left[\boldsymbol{B}(t)|\boldsymbol{\beta}(t)\right]\} + \frac{N}{t}. \quad (11)$$

Apply the inequality (11) recursively from $t=1$ to $t=T$, we have

$$\begin{aligned}
W(T) &\leq W(0) + 2\sum_{t=1}^{T}\boldsymbol{\beta}^+(t)^\top\{\boldsymbol{B}(t)-\mathbf{E}\left[\boldsymbol{B}(t)|\boldsymbol{\beta}(t)\right]\} + \sum_{t=1}^{T}\frac{N}{t}\\
&\leq 0 + 2\sum_{t=1}^{T}\boldsymbol{\beta}^+(t)^\top\{\boldsymbol{B}(t)-\mathbf{E}\left[\boldsymbol{B}(t)|\boldsymbol{\beta}(t)\right]\} + N(1+\log(T))
\end{aligned}$$

Note that $\mathbf{E}\left[\boldsymbol{B}(t)|\boldsymbol{\beta}(t)\right] = \mathbf{E}\left[\boldsymbol{B}(t)\right]$ under i.i.d demand sampling process. Therefore, taking expectation on both sides of above inequality, we can obtain

$$\mathbf{E}\left[\left\|\boldsymbol{\beta}^+(T+1)\right\|_2^2\right] \leq \frac{N(1+\log(T))}{T}$$

which implies that $\mathbf{E}\left[\left\|\boldsymbol{\beta}^+(T+1)\right\|_2^2\right] \to 0$ as $T\to\infty$. In this way, the average debt $\beta_i(T+1)$ of agency i also satisfies the following non-asymptotic convergence guarantee in expectation:

$$\mathbf{E}\left[\left\|\beta_i^+(T+1)\right\|_2^2\right] \leq \frac{N(1+\log(T))}{T},\ \forall i = 1,\ldots,N.$$

To this end, we have shown that the expected fill rate target θ can be achieved under our DAS allocation policy if the set of conditions (1) holds. Altogether, we prove the sufficiency part. □

4 NUMERICAL EXPERIMENTS

In fact, it is challenging to characterize the optimal solution to the DP model (3) due to the curse of dimensionality, which may hinder its application in solving large-scale food allocation problems. For the ease of implementation, we design a heuristic DAS (abbreviated as H-DAS) allocation policy to sequentially determine the allocation quantity to each agency based on the realized demand of this agency and the mean demand of unvisited agencies. Next, we numerically evaluate the performance of this heuristic policy.

4.1 Design of Heuristic Allocation Policy

Let μ_i denote the mean demand of agency i. Upon the visit to agency i, we solve a linear programming (LP) model based on the revealed demand d_i and demand mean of unvisited agencies $\{\mu_{i+1},\dots,\mu_N\}$ to maximize the debt-weighted demand fulfillment objective, subject to the remaining capacity constraint. Note that there may exist multiple optimal solutions to the LP model when the resource is not used up after serving agencies with positive average debt. For tie-breaking and reduction of food waste, we modify $\beta_i^+(t+1) \leftarrow \max\{\varepsilon,\beta_i(t+1)\}$, where ε is a sufficiently small positive real number so that we can fully utilize the remaining resource. Altogether, we present the details of the H-DAS policy in Algorithm 1.

Algorithm 1 H-DAS Allocation Policy

INPUT: Capacity level c, demand samples $\{\boldsymbol{d}_t\}_{t=1}^T$, and expected fill rate target θ.
INITIALIZE: W.L.O.G., let the average debt vector $\boldsymbol{\beta}(1)=(1,\dots,1)$ for $t=1$.
for $t=1,\dots,T$ **do**
  for $i=1,\dots,N$ **do**
    Upon the visit to agency i, update the state variable:

$$s_i=\begin{cases} c & i=1\\ s_{i-1}-d_{i-1}(t)\cdot x_{i-1}(t) & i=2,\dots,N\end{cases}$$

    After observing the realized demand $d_i(t)$ of agency i, the demand fulfillment decision for agency i is determined by solving the following LP problem:

$$\begin{aligned} \max \quad & \sum_{k=i}^{N}\beta_k^+(t)\cdot x_k(t)\\ \text{s.t.} \quad & d_i(t)\cdot x_i(t)+\sum_{k=i+1}^{N}\mu_k\cdot x_k(t)\le s_i,\\ & 0\le x_k(t)\le 1,\ \forall k=i,\dots,N.\end{aligned}$$

  end for (i)
  Update the debt vector $\boldsymbol{B}(t)$ at epoch t and the sample-based average debt vector $\boldsymbol{\beta}(t+1)$.
end for (t)
OUTPUT: The collection of average debt vectors $\{\boldsymbol{\beta}(t)\}_{t=1}^T$.

4.2 Simulation Scenarios

In this subsection, we construct a simulation environment where a food bank sequentially serves $N=12$ agencies that are indexed by $\{1,\dots,12\}$. The stochastic demand of agency i follows a (truncated) normal distribution $d_i\sim\max\big(0,\text{Normal}(\mu_i,(\delta\mu_i)^2)\big)$, where δ denotes the degree of demand variation. We set the mean demand as $\boldsymbol{\mu}=[3,6,9,9,6,3,3,6,9,9,6,3]$ and use parameter α to denote the resource abundance level (i.e., the supply to mean demand ratio). In this way, the maximal fill rate to all agencies is $\min\{\alpha,1\}$ if the demand is deterministic. We solve the resource allocation problem under different demand variations, resource supply, and agency visit order. We compare the performance of our H-DAS policy with two benchmarks: (1) **Hindsight Optimal Solution**. This solution is obtained by the SAA method using the test data, which serves as an upper bound for any sequential policies. (2) **Target-fill-rate (TFR) policy proposed by Manshadi et al. (2021)**. Upon the visit to agency i, this policy either allocates the required amount of resource $d_i\theta^*$ or allocates all the remaining resource s_i to this agency.

In each simulation setting, we generate two sets of i.i.d demand samples (i.e., training and test data), each with sample size $T = 10^4$. We first exploit the classic sampling average approximation (SAA) method to estimate the max feasible value of θ^* with training data, then use it as the target to guide fair resource allocation with test data. We measure the performance of different policies in terms of (i) attained value of the minimal expected fill rate across all agencies, (ii) the average food waste, i.e., the ratio of remaining food after serving the last agency. We simulate the resource allocation environment using Python programming language and solve the optimization problem with Gurobi (9.5.1) solver. All the experiments are performed on a 2.8GHz i9-10900 CPU Windows PC with 64G RAM.

Service experience of each agency. Suppose the planner has $c = 43.2$ units of food for allocation (i.e., supply abundance level $\alpha = 0.6$), we depict the attained fill rate for each agency when the demand variation parameter δ takes value 0.1 and 0.5 respectively in Figure 1. First, it shows that our H-DAS policy delivers almost the same fill rate to each agency as the hindsight optimal solution when $\delta = 0.1$. However, the performance gap increases slightly when δ grows from 0.1 to 0.5. This is because the approximation error of our H-DAS policy increases as the demand exhibits larger deviations from the mean value. Second, since the TFR policy tries to meet the fill rate target for the early-visited agencies as much as possible, the remaining resources are insufficient to deliver the targeted fill rate to the downstream agencies. Third, we note that the largest feasible fill rate for all agencies is $\alpha = 0.6$ if the demand is deterministic (e.g., demand variation $\delta = 0$). However, the minimal expected fill rate under both the hindsight optimal solution and H-DAS policy is larger than 0.6 when $\delta > 0$. This implies that the planner could exploit demand fluctuations to achieve higher fill rate for all agencies under appropriately designed allocation rules.

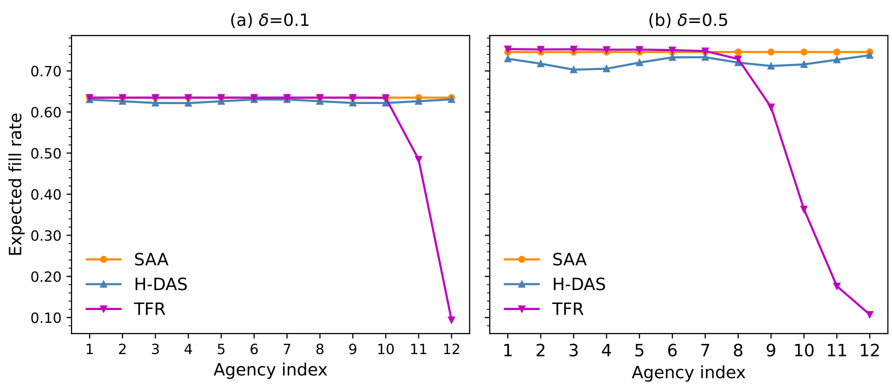

Figure 1: Expected fill rate of each agency under different policies

Impact of demand variations. To evaluate the impact of demand variation on each allocation policy, we set the supply abundance level as $\alpha = 0.6$ and evaluate the performance of each allocation policy under different demand variance δ in Figure 2. First, the minimal expected fill rate under both hindsight optimal solution and H-DAS policy increases as δ grows from 0.1 to 0.5. The reason is that larger demand fluctuation provides more potential for the planner to achieve higher fill rate by serving agencies with low demand. Second, the demand variation has marginal impacts on the TFR policy since it greedily meets the fill rate requirement for early-visited agencies. Third, the capacity waste under our H-DAS policy increases under larger demand variations. Recall that our H-DAS policy exploits revealed demand d_i and demand mean of unvisited agencies to make allocation decision at agency i, which results in a larger approximation gap when demand variation δ increases.

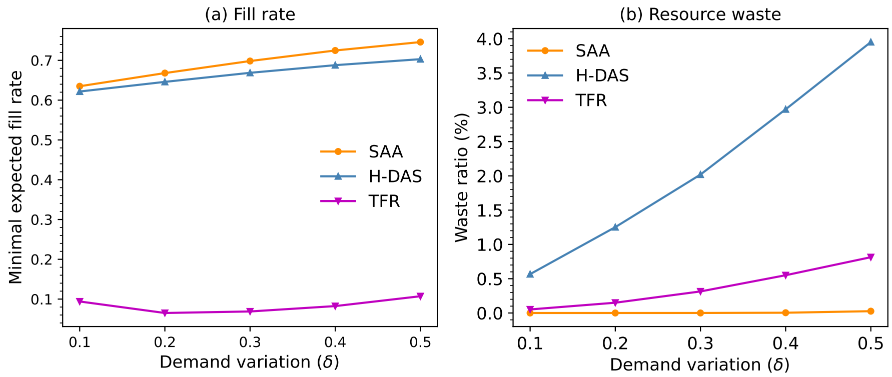

Figure 2: Performance of each allocation policy under various demand variation

Impact of supply level. Next, we evaluate the impact of supply level on the performance of each allocation policy. In this setting, we set the demand variation as $\delta = 0.3$, and depict each allocation policy's minimal expected fill rate and resource waste under various supply abundance levels in Figure 3. It is easy to see that the minimal expected fill rate under all policies would increase in the supply abundance level. However, compared to the TFR policy, the increase under our H-DAS policy appears to be more considerable. This means that our approach could exploit the resource more effectively. Third, the TFR policy wastes less resources since this policy allocates the resource to the early-visited agencies in a greedy manner while our H-DAS policy needs to strategically reserve the capacity for downstream agencies.

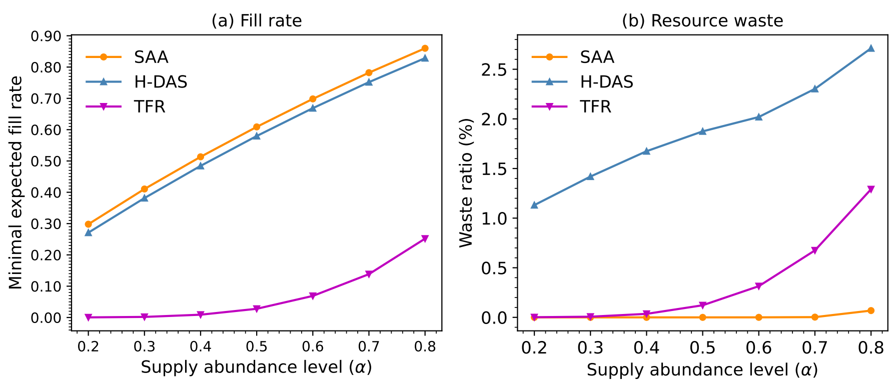

Figure 3: Performance of each allocation policy under various supply abundance level

Impact of agency visit order. Finally, we evaluate the impact of agency visit order on the performance of each allocation policy. In this setting, we fix the demand variation as $\delta = 0.3$ and supply abundance level as $\alpha = 0.6$. We generate 1000 resource allocation instances by randomly permuting the agency visit orders. We solve the resource allocation problem under each agency visit order and use boxplot to depict the performance of each policy in Figure 4. We can see that the minimal expected fill rate under our H-DAS policy exhibits negligible variation, which means that it is slightly affected by the order of visit. However, the TFR policy provides fluctuating fill rate guarantee under different visit orders since it endows higher priority to early-visited agencies. Moreover, the resource waste under our H-DAS policy is slightly higher and exhibits larger deviation compared with the TFR policy.

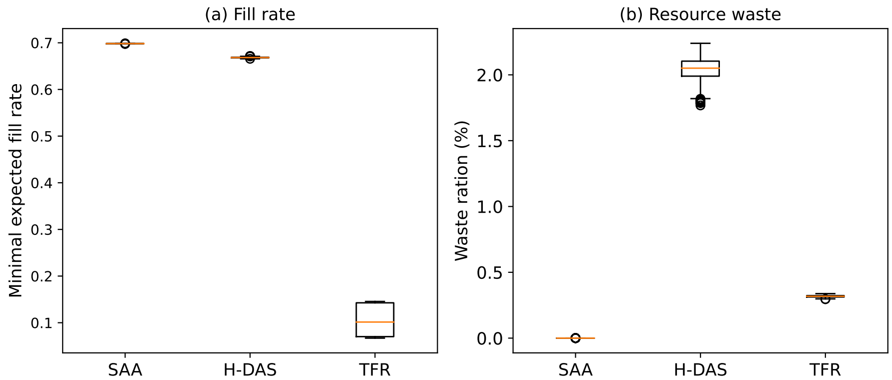

Figure 4: Performance of each allocation policy under different agency visit orders

5 CONCLUSION

Motivated by the allocation fairness challenges faced in food bank operations, this paper proposes an analytical framework to maximize the minimal expected fill rate across all eligible agencies. We first exploit the online convex optimization approach and Fenchel dual to characterize the attainable expected fill rate target, and then develop provable near-optimal allocation rules to facilitate the resource allocation in real-time. We show that our allocation rule can protect the most vulnerable agencies from food shortages without wasting too much resource. This insight holds regardless of the predetermined allocation sequence. Last but not least, we remark that this analytic framework can be easily adapted to other dynamic allocation or rationing problems with fairness guarantee.

ACKNOWLEDGMENTS

Jinpeng Liang is supported by the National Natural Foundation Science of China [Grant 72101042] and the Fundamental Research Funds for the Central Universities [Grant 3132022171]. Guodong Lyu is supported by the 2019 Academic Research Fund Tier 3, the Ministry of Education, Singapore [Grant MOE-2019-T3-1-010].

A SOME PRELIMINARY RESULTS FOR PROOF

Claim 1 Given any two vectors $\boldsymbol{p},\boldsymbol{q} \in \mathbb{R}^N$, the following inequality holds:

$$\left\|\boldsymbol{q}^+\right\|_2^2 \le \left\|\boldsymbol{p}^+\right\|_2^2 + 2\boldsymbol{p}^+\left[\boldsymbol{q}-\boldsymbol{p}\right] + \left\|\boldsymbol{q}-\boldsymbol{p}\right\|_2^2 \tag{12}$$

Proof. By definition, we can rewrite inequality (12) as

$$\sum_{i=1}^{N}(q_i^+)^2 \le \sum_{i=1}^{N}(p_i^+)^2 + \sum_{i=1}^{N}2p_i^+(q_i-p_i) + \sum_{i=1}^{N}(q_i-p_i)^2 \tag{13}$$

Therefore, it is sufficient to prove inequality (13) in component-wise. For any $i = 1,\dots,N$, we claim:

$$(q_i^+)^2 \le (p_i^+)^2 + 2p_i^+(q_i-p_i) + (q_i-p_i)^2 \tag{14}$$

We consider the following cases to prove inequality (14):

- $p_i \le 0$, $q_i \ge 0$, then inequality (14) reduces to $(q_i^+)^2 \le (q_i-p_i)^2$, which is true.
- $p_i \le 0$, $q_i \le 0$, then inequality (14) reduces to $0 \le (q_i-p_i)^2$, which is also true.
- $p_i \ge 0$, $q_i \ge 0$, then both sides of inequality (14) are equal to $(q_i)^2$.

- $p_i \geq 0$, $q_i \leq 0$, then both sides of inequality (14) are also equal to 0.

Hence, we complete the proof of inequality (12) by summing inequality (14) for all $i = 1, \dots, N$. □

Lemma 1. *For any sample epoch (t) and any average debt vector $\boldsymbol{\beta}(t)$, the optimal value of DP formulation $f_a(\boldsymbol{\beta}^+(t))$ is a convex function w.r.t. debt vector $\boldsymbol{\beta}^+(t)$.*

Proof. Note that formulation (3) is indeed a DP model with N stages. Let $J^*(s_i)$ denote the optimal value function under remaining capacity level s_i when the truck arrives at agency i. In this way, the optimal value function as stage N can be represented as

$$J_N^*(s_N) = \max_{x_N(s_N,\boldsymbol{d}) \in X_N(s_N,\boldsymbol{d})} \mathbf{E}\left[\beta_N \cdot x_N(s_N,\boldsymbol{d})\right] = \mathbf{E}\left[\beta_N \cdot \min\left(1, \frac{s_N}{d_N}\right)\right]$$

It is well known that $J_N^*(s_N)$ is convex w.r.t. debt vector $\boldsymbol{\beta}^+(t)$ since it appears as the coefficient in the objective function. W.L.O.G., we assume that $J_{i+1}^*(s_{i+1})$ at stage $i+1$ is convex w.r.t debt vector $\boldsymbol{\beta}^+(t)$. By the Bellman equation, $J_i^*(s_i)$ at stage i can be formulated as

$$J_i^*(s_i) = \max_{x_i(s_i,\boldsymbol{d}) \in X_i(s_i,\boldsymbol{d})} \mathbf{E}\left[\beta_i^+(t) \cdot x_i(s_i,\boldsymbol{d}) + J_{i+1}^*(s_{i+1})\right] \tag{15}$$

For any $x_i(s_i,\boldsymbol{d}) \in X_i(s_i,\boldsymbol{d})$, both terms in the RHS of equation (15) are convex w.r.t. $\boldsymbol{\beta}^+(t)$. Hence, we can claim that $J_i^*(s_i)$ is a convex function of $\boldsymbol{\beta}^+(t)$. As a result, the optimal value $f(\boldsymbol{\beta}^+(t))$, which is equivalent to $J_0^*(s_0)$, is also a convex function of $\boldsymbol{\beta}^+(t)$. □

Lemma 2. *For any sample epoch t and any average debt vector $\boldsymbol{\beta}(t)$, the debt vector $\boldsymbol{B}(t)$ obtained by solving DP problem* (3) *satisfies $\sum_{i=1}^N \beta_i^+(t)\mathbf{E}\left[B_i(t)|\boldsymbol{\beta}(t)\right] \leq 0$.*

Proof. The proof is motivated by exploiting the convexity of $f(\boldsymbol{\beta}^+(t))$ to derive its Fenchel dual. Given the ex-ante fairness target θ, we can define the Fenchel dual of $f(\boldsymbol{\beta}^+(t))$ as follows:

$$f^*(\theta) = \max_{\boldsymbol{\beta}^+(t+1) \geq \boldsymbol{0}} \left[\sum_{i=1}^N \beta_i^+(t+1)\theta - f(\boldsymbol{\beta}^+(t+1))\right].$$

According to Fenchel-Young inequality, we have:

$$f(\boldsymbol{\beta}^+(t+1)) + f^*(\theta) \geq \sum_{i=1}^N \beta_i^+(t+1)\theta.$$

We let $\boldsymbol{x}^*(\boldsymbol{s},\boldsymbol{d}_t)$ denote the optimal solution to DP model (3) given debt vector $\boldsymbol{\beta}(t)$, it is straightforward to obtain the following inequalities:

$$\begin{aligned} f^*(\theta) &\geq \sum_{i=1}^N \beta_i^+(t)\theta - f(\boldsymbol{\beta}^+(t)) \\ &= \sum_{i=1}^N \beta_i^+(t)\left\{\theta - \mathbf{E}\left[x_i^*(s_i,\boldsymbol{d}_t)|\boldsymbol{\beta}(t)\right]\right\} \\ &= \sum_{i=1}^N \beta_i^+(t)\mathbf{E}\left[B_i(t)|\boldsymbol{\beta}(t)\right]. \end{aligned}$$

Note that the conditions (1) imply that $f^*(\theta) \leq 0$, which means that

$$\sum_{i=1}^N \beta_i^+(t)\mathbf{E}\left[B_i(t)|\boldsymbol{\beta}(t)\right] \leq 0.$$

To this end, we complete the proof. □

REFERENCES

Balcik, B., S. Iravani, and K. Smilowitz. 2014. "Multi-vehicle Sequential Resource Allocation for a Nonprofit Distribution System". *IIE Transactions* 46(12):1279–1297.

Bartholdi, J. J., L. K. Platzman, R. L. Collins, and W. H. Warden. 1983. "A Minimal Technology Routing System for Meals on Wheels". *Interfaces* 13(3):1–8.

Berenguer, G., and Z.-J. M. Shen. 2020. "OM Forum—Challenges and Strategies in Managing Nonprofit Operations: An Operations Management Perspective". *Manufacturing & Service Operations Management* 22(5):888–905.

Bertsimas, D., V. F. Farias, and N. Trichakis. 2011. "The Price of Fairness". *Operations Research* 59(1):17–31.

Brill, E. D. 1979. "The Use of Optimization Models in Public-Sector Planning". *Management Science* 25(5):413–422.

Eisenhandler, O., and M. Tzur. 2019. "The Humanitarian Pickup and Distribution Problem". *Operations Research* 67(1):10–32.

Gunes, C. 2010. *Essays on Operations Management*. Ph.D. thesis, Tepper School of Business, Carnegie Mellon University, Pittsburgh, Pennsylvania.

Lee, D., E. Sönmez, M. I. Gómez, and X. Fan. 2017. "Combining Two Wrongs to Make Two Rights: Mitigating Food Insecurity and Food Waste through Gleaning Operations". *Food Policy* 68:40–52.

Lien, R. W., S. M. Iravani, and K. R. Smilowitz. 2014. "Sequential Resource Allocation for Nonprofit Operations". *Operations Research* 62(2):301–317.

Lyu, G., W.-C. Cheung, M. C. Chou, C.-P. Teo, Z. Zheng, and Y. Zhong. 2019. "Capacity Allocation in Flexible Production Networks: Theory and Applications". *Management Science* 65(11):5091–5109.

Lyu, G., W.-C. Cheung, C.-P. Teo, and H. Wang. 2019. "Multi-Objective Online Ride-Matching". SSRN 3538755, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=3538755, accessed 16th August 2022.

Ma, Y., T. Wang, and H. Zheng. 2022. "On the Fairness and Efficiency in Nonprofit Operations: Dynamic Resource Allocations". SSRN 4017058, https://papers.ssrn.com/sol3/papers.cfm?abstract_id=4017058, accessed 14th July 2022.

Manshadi, V., N. Rad, and R. Scott. 2021. "Fair Dynamic Rationing". Available at arXiv:2102.01240, https://arxiv.org/abs/2102.01240, accessed 14th July 2022.

Marsh, M. T., and D. A. Schilling. 1994. "Equity Measurement in Facility Location Analysis: A Review and Framework". *European Journal of Operational Research* 74(1):1–17.

Nair, D. J., D. Rey, and V. V. Dixit. 2017. "Fair Allocation and Cost-effective Routing Models for Food Rescue and Redistribution". *IISE Transactions* 49(12):1172–1188.

Orgut, I. S., J. Ivy, and R. Uzsoy. 2017. "Modeling for the Equitable and Effective Distribution of Food Donations under Stochastic Receiving Capacities". *IISE Transactions* 49(6):567–578.

Orgut, I. S., J. Ivy, R. Uzsoy, and J. R. Wilson. 2016. "Modeling for the Equitable and Effective Distribution of Donated Food under Capacity Constraints". *IIE Transactions* 48(3):252–266.

Orgut, I. S., J. S. Ivy, R. Uzsoy, and C. Hale. 2018. "Robust Optimization Approaches for the Equitable and Effective Distribution of Donated Food". *European Journal of Operational Research* 269(2):516–531.

Savas, E. S. 1978. "On Equity in Providing Public Services". *Management Science* 24(8):800–808.

Talluri, K. T., and G. J. V. Ryzin. 2004. *The Theory and Practice of Revenue Management*. New York: Springer.

Zheng, H., G. Lyu, J. Ke, and C.-P. Teo. 2022. "From Targeting to Transfer: Design of Allocation Rules in Cash Transfer Programs". *Manufacturing & Service Operations Management*. Advance online publications, https://pubsonline.informs.org/doi/abs/10.1287/msom.2022.1101, accessed 16th September 2022.

Zhong, Y., Z. Zheng, M. C. Chou, and C.-P. Teo. 2018. "Resource Pooling and Allocation Policies to Deliver Differentiated Service". *Management Science* 64(4):1555–1573.

AUTHOR BIOGRAPHIES

JINPENG LIANG is an Associate Professor in the School of Transportation Engineering at Dalian Maritime University. He received his Ph.D. degree in system science from Beijing Jiaotong University in 2020. His research interests include transportation system modeling and optimization.

GUODONG LYU is an Assistant Professor in operations management at the Hong Kong University of Science and Technology. He earned his Ph.D. in the Department of Analytics & Operations at National University of Singapore. His research interests include online optimization, robust optimization, and data analytics.