Independent Study Final Report
Fair Allocation for Rachel's Table

Mason duBoef

December 2025

https://github.com/mduboef/food-rescue-fair-allocation

1 Introduction

Rachel's Table (RT) is a volunteer-run food rescue organization operating across Western Massachusetts. The organization collects excess food from approximately 100 donors, including grocery stores, restaurants, and farms, that would otherwise go to waste, and delivers it to 66 receiving agencies, mostly food pantries and soup kitchens. Through this network, Rachel's Table delivers over 50,000 meals per month to those in need throughout Western Massachusetts.

Currently, Rachel's Table employs a manual approach to routing and resource deployment. When a food donor like Whole Foods informs RT that they have excess food available for donation, a dispatcher must manually coordinate between the network of volunteer drivers and nearby receiving agencies before instructing a driver to complete the pickup and delivery. This process occurs daily with a rotating schedule of drivers and dispatch volunteers. While some donations follow regular schedules, many food items become available on an ad-hoc basis, requiring dispatchers to flexibly assign routes while working around numerous constraints.

RT dispatch must work around a number of real-world constraints. One notable constraint involves the Food Bank of Western Massachusetts (FBWM), the region's largest food bank. Various donors and agencies are partnered with FBWM, with RT facilitating deliveries between them. However, receiving agencies not partnered with FBWM are prohibited from receiving goods from FBWM-partnered donors. Additional constraints include refrigeration requirements for perishable items and the storage capacity of vehicles and receiving agencies. This system is challenging and time-consuming for dispatch volunteers, who must juggle the needs of donors, drivers, and receiving agencies in real time.

Beyond delivering as much food as possible, Rachel's Table seeks to distribute food fairly among their 66 partner agencies. In principle, RT would like to allocate food proportionally based on each agency's size, measured by the number of people they serve. However, RT has observed significant disparities in how well-served different agencies are. Some agencies receive food quantities nearly proportional to their size, while others, particularly the largest agencies, receive very little relative to their needs. This disparity motivated RT to reach out to the Fair and Explainable Decision Making lab at UMass Amherst.

Our goal is to provide an automated allocation solution that reduces dispatcher workload while ensuring fair outcomes among receiving agencies over time. This report describes the formalization of the fair allocation problem, the data processing pipeline I developed, and the integer linear program I implemented that computes fair allocations over the RT network. This project was a collaborative effort between myself, Prof. Yair Zick and Paula Navarrete Diaz, one of Prof. Zick's PhD students.

2 Measuring Fairness

To quantify how well-served each agency is, Rachel's Table uses a metric based on meals delivered relative to meals served. Let MD_i denote the average number of meals allocated to agency i by RT per week, and let MS_i denote the average number of meals served by agency i per week. The ratio

$$\mathrm{MDMS}_i = \frac{\mathrm{MD}_i}{\mathrm{MS}_i}$$

captures what fraction of an agency's food needs are met by Rachel's Table deliveries. After calculating MDMS values across all 66 agencies, RT observed dramatic disparities. Some agencies had MDMS values near 1.0, while other agencies, particularly the largest ones, had values as low as 0.0016. This motivated our work on fair allocation.

Note that RT lacks data on unmet demand (people who leave a pantry empty-handed because food ran out). The meals served statistic (MS_i) serves as a proxy for agency size and need, though it may underestimate true demand at under-resourced agencies.

3 Methods

3.1 Data Processing

A significant portion of my work involved processing and organizing RT's data. The original data was spread across multiple unorganized Excel spreadsheets with inconsistent formatting, missing fields, and qualitative descriptions that required manual interpretation.

I compiled agency data from these sources into a structured CSV format containing: agency name, address, city, state, ZIP code, FBWM partnership status (NFB, FBE, or FBNE), refrigerator count, freezer count, meals served per week (MS), and meals delivered per week (MD). For agencies missing MS or MD values, the system uses median imputation to avoid excluding them from the allocation.

Geographic coordinates (latitude and longitude) were needed for all donors and agencies. I manually geocoded each donor and agency location to enable future distance-based feasibility constraints. This involved looking up addresses and recording coordinates for all 96 donors and 65 agencies in the dataset.

To improve RT's ongoing data collection, I drafted a Google Form for drivers to complete as they make their routes.[1] While RT already uses a form for drivers to fill out, they expressed interest in updating it. This form captures pickup and drop-off locations, food weights by category, and timestamps. Responses to our new form automatically populate a CSV file that can be directly ingested by our allocation system, creating a sustainable data pipeline for RT's operations.

[1] https://docs.google.com/forms/d/e/1FAIpQLSd4FRxbO7dMhGO7DqUlKgOiUcbijTe2-yTep1doYvzB5EKdUg/viewform

3.2 Network Modeling

I built a model of the RT delivery network. It is modeled as a bipartite graph with donors on one side and agencies on the other. An adjacency matrix M encodes which donor-agency pairs represent feasible deliveries, where $M_{d,i} = 1$ indicates that a trip from donor d to agency i is permissible.

The feasibility constraints encoded in this matrix include:

- **FBWM Partnership Rules**: FBWM-partnered donors cannot deliver to non-FBWM agencies. This is a hard constraint imposed by the Food Bank of Western Massachusetts.
- **Geographic Proximity**: In practice, deliveries should be limited to reasonable distances. While the current implementation uses random connectivity for simulation purposes, the latitude/longitude data I collected will enable distance-based constraints in production.
- **Refrigeration Needs**: We are able to disconnect agencies without refrigeration from donors like butchers whose food largely requires refrigeration.

Individual food items are modeled with the following attributes: the donor providing the item, the timestep when the item becomes available, the item's weight in pounds, and a breakdown of the item's weight across seven food categories (dairy, meat, produce, grain, processed foods, non-perishables, and hygiene products). This food type information enables our fairness objective to consider not just total food quantity but also nutritional variety.

3.3 Fairness Objectives

We considered several approaches to formalizing fairness for this domain:

Envy-based criteria define fairness in terms of whether agents prefer their own allocation to others'. After speaking to RT we determined that this approach would be poorly suited to our setting, despite being a population option for fair allocation problems. In this setting agencies neither observe nor care about what other agencies receive, and with identical valuation functions (all agencies simply want as much food as possible with some preference for receiving a mix of food types), envy-based approaches are ill-suited to this problem.

Efficiency/Utilitarian Social Welfare (USW) maximizes total food delivered across all agencies. While this ensures efficiency it does little to ensure fair distribution among agencies. If we were to only maximize efficiency, food may concentrate at easily-accessible agencies while leaving others under-served.

Egalitarian Social Welfare (ESW) maximizes the minimum utility across all agents. This ensures that the worst-off agency is as well-served as possible, directly addressing RT's concern about disparities. We adopt an egalitarian objective as our primary fairness criterion.

To account for agency size, we weight each agency's utility by $W_i = \text{MS}_i$, the number of meals served per week. This means we maximize the minimum *per-capita* food allocation rather than absolute quantities, ensuring small and large agencies are treated equitably. A person in need of food is no more or less deserving because of the size of the agency they went to.

Additionally, we extend the egalitarian objective to individual food types. This prevents solutions that technically achieve high ESW but give certain agencies only bread while receive only protein. We want to incentivize the allocation of a mixture of foods to agencies. Our objective includes terms r_f for each food type f, representing the minimum per-capita allocation of that food type across all agencies.

3.4 Integer Linear Program Formulation

Paula and I iterated through a number of different ways to formalize the problem so that an integer linear program could find the optimal division of goods. Eventually, we settled on the following.

Let N be the set of agencies, G the set of food items, D the set of donors, F the set of food types, and T the set of timesteps. Each timestep represents a 2 to 3 hour shift during which drivers can make deliveries.

Parameters:

- W_i: Weight (meals served per week) of agency $i \in N$
- $q_{g,f}$: Quantity (pounds) of food type $f \in F$ in item $g \in G$
- $m_{i,d}^t$: Binary indicating if a trip from donor d to agency i is feasible at time t
- $h(g)$: Maps item g to its donor
- $z(g)$: Maps item g to its availability timestep

Decision Variables:

- $x_{i,g} \in \{0, 1\}$: Whether item g is assigned to agency i
- $y_{i,d}^t \in \{0, 1\}$: Whether a trip occurs from donor d to agency i at time t
- $r \geq 0$: Minimum weighted utility across all agencies (ESW)
- $r_f \geq 0$: Minimum weighted utility of food type f across all agencies

Objective:

$$\text{maximize} \quad \alpha \cdot r + \sum_{f \in F} \alpha_f \cdot r_f$$

This maximizes a weighted combination of overall egalitarian welfare and food-type-specific egalitarian welfare. The parameters α and α_f control the relative importance of total food quantity versus nutritional variety.

Constraints:

Egalitarian welfare definitions:

$$\sum_{g\in G}\sum_{f\in F}\frac{x_{i,g}\cdot q_{g,f}}{W_i}\geq r \qquad \forall i\in N \tag{1}$$

$$\sum_{g\in G}\frac{x_{i,g}\cdot q_{g,f}}{W_i}\geq r_f \qquad \forall i\in N, f\in F \tag{2}$$

Efficiency threshold (requiring allocation achieves fraction β of optimal USW):

$$\sum_{g\in G}\sum_{i\in N}\sum_{f\in F}x_{i,g}\cdot q_{g,f}\geq \beta\cdot \text{OPT} \tag{3}$$

Allocation and feasibility constraints:

$$\sum_{i\in N}x_{i,g}\leq 1 \qquad \forall g\in G \tag{4}$$

$$y_{i,d}^t\leq m_{i,d}^t \qquad \forall i\in N, d\in D, t\in T \tag{5}$$

$$x_{i,g}\leq y_{i,h(g)}^{z(g)} \qquad \forall i\in N, g\in G \tag{6}$$

Constraint (4) ensures each item is allocated to at most one agency. Constraint (5) restricts trips to feasible donor-agency pairs. Constraint (6) links item allocation to trip execution: an agency can only receive an item if a trip is made from that item's donor to that agency during the appropriate timestep.

The efficiency threshold (3) prevents the egalitarian objective from sacrificing too much total food delivery. We first solve a subproblem to find OPT, the maximum total food that can be allocated ignoring fairness, then require any fair solution to deliver at least $\beta\cdot$ OPT pounds of food.

We experimented with tuning several parameters:

- **USW fraction** (β): Higher values (e.g. 0.8) force more food to be delivered but may reduce fairness; lower values (e.g. 0.4) allow more flexibility for egalitarian optimization.
- **Timesteps**: We settled on 24 to 60 timesteps representing 2-3 hour shifts over a number of days, balancing temporal granularity with computational tractability.
- **Food type weights** (α_f): Equal weights ensure balanced consideration of all food categories; these could be adjusted based on nutritional priorities.

4 Results

We are able to find the optimal allocation for the full RT network in a reasonable time given typical hardware (my personal laptop). This includes 96 donors, 65 agencies, over 1,000 food items, and 24 timesteps. Phase 1 (efficiency optimization) completes in under 5 seconds, while Phase 2 (egalitarian optimization) completes within the 5-minute time limit, finding feasible solutions that achieve positive egalitarian welfare. Running the system on simulated data with RT's actual donor and agency lists produces the following results:

- **Network scale**: 96 donors, 65 agencies, 6,240 possible connections
- **Feasible connections**: 3,196 (after FBWM and random connectivity constraints)
- **Items generated**: 1,098 items totaling 32,813 lbs across 24 timesteps
- **Optimal USW**: 326.67 (weighted total across all agencies)
- **Achieved ESW**: 0.015 lbs per person served (minimum across agencies)
- **Food allocated**: 28,421 lbs (86.6% of available food)
- **Agencies served**: 65/65 (100%)
- **Connections used**: 176 (5.5% network utilization)

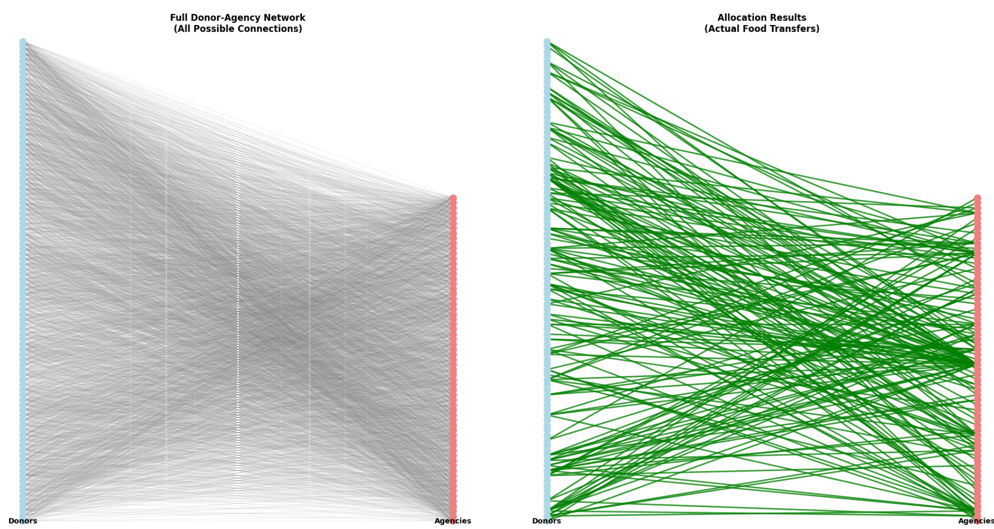

Figure 1: Visualization of allocation.

The key result is that our ILP formulation successfully computes allocations at the scale required for RT's operations, achieving positive egalitarian welfare while serving all agencies. The relatively sparse network utilization (5.5%) reflects the random connectivity in our simulation; real-world geographic constraints would likely produce denser local connectivity and higher ESW.

5 Next Steps

Several additions need to be made before we have a dispatch tool that is ready for RT to use.

Further Integration of Real-world data: The current results use simulated food donations with random connectivity. Integrating actual donation patterns and distance-based feasibility constraints will produce more realistic and actionable allocations.

User interface: For RT dispatchers to use this system, we need a user-friendly interface for inputting scheduled and ad-hoc donations and viewing the resulting shift schedules.

Assignment of drivers to deliveries: Our system determines which food goes to which agency but does not determine who will actually complete each delivery. Another ILP or matching algorithm could be used to assign trips to driver. For each timestep we would hope to divide the donors (and their associated drops offs) to drivers so that the shift assigned to the driver at that timestep can feasibly be done in 2 hours. We would aim to give all drivers a similar length route.

Driver routing: In addition to assigning drivers to deliveries, we need to optimal route drivers through all the stops assigned to them. This is essentially the Traveling Salesman Problem. We need to find the minimum spanning walk between the respective stops. This is a well-known subproblem. Many approaches exist for finding near optimal solutions to the Traveling Sales Problem in networks of this size.

Temporal fairness: Rather than optimizing fairness within a single planning period, we could track allocations over time and compensate agencies that received less in previous periods. If you received less food than you deserved last month, we can make it up to you this month.

I am considering different independent study options for the Spring, but will likely return to finish my work on this project. Regardless, Paula will continue to work on this project.

6 Conclusion

During this independent study I worked with Paula Navarrete Diaz to develop and implement a fair allocation system for Rachel's Table. We worked alongside Rachel's Table, taking time to understand their operation and seeking to best address their fairness concerns. While formalization work was largely collaborative effort between Paula, Prof. Zick and I, I performed most of the implementation and data processing.

By formalizing the problem as an integer linear program with egalitarian objectives, we compute allocations that maximize the welfare of the worst-off agency while ensuring efficient use of donated food. The implementation successfully handles the full scale of RT's network, demonstrating feasibility for real-world deployment. With continued development, this system can reduce dispatcher workload while ensuring fair food distribution throughout Western Massachusetts.