

MEAL: Model of Empathy Augmented Logistics for Food Security

Seoyeong Park¹[0009–0001–4402–6160] and Munindar P. Singh¹[0000–0003–3599–3893]

North Carolina State University, Raleigh, NC, USA
spark43@ncsu.edu, mpsingh@ncsu.edu

Abstract. Millions globally lack access to nutritious food, experiencing food insecurity. Efforts to address food insecurity seek to provide consumers food that may be *rescued* (i.e., what warehouses or grocers would otherwise soon discard as unusable), directly donated, or acquired using governmental funds.

Current approaches produce allocations that optimize global objectives to store and move food efficiently across food banks. However, they largely overlook consumer preferences and constraints. As a result, the resulting allocations lead to consumers either using foods they do not care for or discarding such foods, leading to food waste.

This paper presents a new model, studied via human study and agent-based simulation, that shows how incorporating the consumer perspective on par with the provider perspective can lead to better outcomes overall. We find that persuasive messages that include individual circumstances and the social context can promote prosociality and empathy.

Keywords: Food security · Multiagent system · Agent-based simulation

1 Introduction

Food insecurity is the condition of a household having poor access to adequate food and reduced quality of food intake [8]. One-eighth (approximately 17 million) of US households experience food insecurity [8], and it is a critical global concern [7].

The US food bank system is a nonprofit organization that reduces food waste and alleviates food insecurity by collecting, storing, and distributing food to those in need [1]. The federal government provides funding and capabilities to procure, store, transport, and distribute food [8]. Local food banks (providers) may receive donations from organizations, retailers, and individuals as well as allocations from regional food banks. Volunteers sort and distribute food to consumers and sometimes to smaller sites called food pantries. A consumer is a household experiencing food insecurity. Consumers deserve not only to satisfy their health-based or cultural dietary needs and to have a choice on what they eat, albeit limited by what is available.

Ensuring equitable distribution is difficult when supplies are in short supply, and preferences are diverse. Thus, a traditional approach may end up giving its limited supply of milk to a household without children while a household with children has to do without. Or, it might allocate starchy foods to a person with diabetes. Current research addresses logistic efficiency [2, 12] or concentrating on consumers’ tastes [11], but not on both aspects together.

We propose **MEAL** for Model of Empathy Augmented Logistics for Food Security. MEAL allocates food by considering both consumer needs and societal objectives such as reducing food waste and improving equity. MEAL’s **novelty** thus lies in combining prosociality with a multistakeholder model of food security. Through extensive simulation experiments, we find that MEAL reduces waste and increases satisfaction in distributing food items compared to models that consider only one side, either consumers or providers. Through a *human study*, we find that persuasive messages, especially those that fit individual circumstances and the social context, can promote prosociality.

2 Motivation for MEAL

In an ideal world, everyone would get the food items they most prefer. However, it is impossible to match everyone’s preferences with constraints. Previous approaches to promoting food security through sharing food with those in need have generally taken a rigid stance. In these approaches, an organization such as a food bank, which has all the power and the food, decides how to allocate it to food-insecure households. Besides the obvious challenges of not accommodating the wishes of the intended recipients of the food, this approach leads to greater food waste system-wide because foods that do not match the constraints and preferences of the recipients cannot be used by them. This top-down allocation inevitably ends up with some consumers not receiving their preferred items, which not only leaves them less satisfied but also worsens food waste. Therefore, we consider restructuring the problem such that other acceptable allocations can be found. Our approach builds on key principles: social welfare, equity, prosociality, and empathy.

2.1 Stakeholders

We consider two main types of stakeholders. *Consumers* are households served through our recommendation system. They aim to acquire food items that align with their preferences and needs. This consumer-centric perspective emphasizes the importance of enhancing consumer satisfaction and personalized experiences for food allocation [3]. As consumers interact with the system, their preferences for food items are constantly captured and refined. These preferences evolve over time and are shaped by factors such as age, health status, dietary constraints, household status, and willingness to make prosocial choices [4, 5]. The agent learns these dynamics by reflecting consumer feedback toward recommended

food items. This learning process allows the agent to provide recommendations matching a consumer’s tastes and current needs.

Providers seek to improve the effective distribution of the available food. This entails reducing food waste, maximizing the distribution of food, and meeting the needs of their community while providing food items that suit consumer preferences. The provider prioritizes not merely using in-stock items but also fulfilling consumer requests as closely as possible [3]. However, they might propose less-preferred alternatives when necessary. The provider intends to trigger empathy and gently nudge consumers to accept alternatives through social and psychological factors that influence decision-making.

2.2 Research Questions

Accordingly, this study investigates these research questions.

RQ_{prosociality} How can MEAL produce equitable allocations by incorporating a dynamic multistakeholder context (consumers and providers) and supporting prosocial behavior among consumers?

RQ_{persuasion} Do persuasion and empathy influence human decisions about food and prosocial behavior?

3 Empirical Evaluation with Humans

Even if MEAL recommends substitutes that are mostly consistent with preferences, simply offering those without any context or with a generic explanation is less effective and unhelpful for consumers. To validate our assumptions on human behavior and prosociality underlying our simulation, we conducted an IRB-approved human study on consumer decision-making. Our study shows personalized, context-rich persuasive messages may improve engagement compared to simple and generic ones.

We observed no significant difference in decision-making with or without a persuasive message. The acceptance rates were similar for *No persuasion* and *Persuasion*, 62.5% in the former and 63.7% in the latter. Applying the two-proportion Z-test [10] produced a p-value of 0.7, indicating no significant difference. This indicates that the consumers are highly likely not affected by persuasive situations when the system provides justification and context, implying that the persuasive message used in the study was too weak or generic to resonate with the participants’ priorities.

Similarly, we found no statistically significant difference in consumer satisfaction: the mean of 3.57 *No persuasion* and the mean of 3.43 *Persuasion*, with a Mann-Whitney U test [6] p-value of 0.193. The results show that consumer satisfaction was not greatly affected by the given persuasive message. This indicates that the observed increase in acceptance rate with persuasion may not necessarily translate to a corresponding increase in consumer satisfaction. In other words, simply encouraging to accept substitutes may not be accompanied to enhance the consumer’s experience.

Understanding what motivates consumers to accept recommendations is crucial. The survey given at the end of the study revealed that the participants are most likely to accept if the alternatives are what they like or similar to their original choices in terms of taste, type of food, or nutritional value; in other words, familiarity matters.

Consumers may measure their satisfaction not only with fulfilling personal desires but also by feeling rewarded for helping others. Almost all survey respondents answered that they would highly likely change their decision of refusing a recommended item regardless of their personal situations if they know their choice helps promote social well-being, unless they have strong dietary restrictions.

4 Model Design

Our goal is to simultaneously maximize consumer satisfaction and maximize the provider’s benefit. The agent understands stakeholders’ values, the future state of the world for each action it can perform, and the social experience its consumer will derive for each action it can perform. Then, since we cannot maximize both objectives, the agent moderates to achieve an optimal trade-off between two stakeholders.

We now formalize our problem setup. We have a set of consumers U and a set of food items F , where each consumer in U has profile information and unique food preferences toward each food item in F . Each item in F carries attributes that reflect its importance in consumption priority and benefits to the provider. These attributes include multiple factors, such as inventory capacity, expiration date, and perishability, shaping the provider benefits $c_{u,d,t}$ associated with each recommendation happening at time step t . Within this dynamic framework, $d_{u,f,t} \in D$ represents a recommendation for consumer u at a specific time step t . It contains two attributes: a recommended food item and a binary indicator of whether it is accepted. Subsequently, we define that consumer satisfaction $h_{u,d,t}$ comes as ratings at a time step t , ranging from 0 (no preference or experience) to 5 (extremely like). The provider’s benefit c is determined by the aggregate score of accepted food items, scaling to the same range as h . These scores are updated in real time as allocations are made.

The problem involves finding the optimal way to distribute the available food to consumers over time while considering their preferences and impact on the community, in other words, managing the trade-off between these two objectives. To balance these objectives, a weighted sum of consumer satisfaction H and provider benefit C is used with a weighting factor denoted as ω ($0 \leq \omega \leq 1$). We choose the optimal value of ω that maximizes both H and C . Therefore, the agent’s overall reward for the decision-making objective is a weighted combination of satisfaction and provider benefit.

By using Q-learning [9], our model effectively adapts to dynamic changes in consumers’ needs, food availability, and other factors and incorporates long-term interaction into their decision-making process.

5 Results

Our study considers three baselines: random recommendation, consumer-focused, and provider-focused approaches.

Random recommendation Recommends items randomly from in-stock inventory, regardless of consumer preferences or provider benefit. This baseline disregards fairness and trust.

Consumer-focused Solely prioritizes consumers' preferences based on their past interactions and preferences. This model is equivalent to assigning a weighting factor ω of 1 and completely ignores provider benefit.

Provider-focused Solely prioritizes the provider-side operation exclusively and disregards consumer preferences. It is equivalent to assigning a weighting factor ω of 0.

5.1 Consumer Satisfaction

Consumers find greater satisfaction with recommendations that consider both consumer preference and society's welfare. This trade-off indicates that MEAL fulfills the intended objectives even though it might sacrifice some provider benefits.

The provider-focused model delivers the highest cumulative provider benefit, and the consumer-focused model achieves the lowest provider benefit. The provider benefit decreases as the weight assigned to the provider decreases, in other words, it increases inversely related to ω . Consumer satisfaction visibly improves, unlike what we originally expected both stakeholders to sacrifice to some extent if we set a parameter for the reward. The evenly considered ($\omega = 0.5$) model and the optimal ($\omega = 0.2$) model outperform the consumer-focused model in terms of getting higher consumer satisfaction. It indicates that MEAL recommends items that consumers like more.

This implies that the weighted models distribute resources in a way that actually benefits both consumers and providers more. By incorporating the provider's perspective, MEAL achieves a more efficient and equitable allocation, meaning that a greater number of consumers are served or a greater number of consumers get better at matching their preferred items among the available inventory.

5.2 Acceptance of Recommendations

How much the model skews to consumer satisfaction affects the acceptance rate. The higher the weight on consumer preferences, the higher the acceptance rate. The gap in the acceptance rate between the consumer-focused and provider-focused models differs notably. The consumer-focused model dominates all other models, particularly the provider-focused and random recommendation. We could observe that the acceptance rate gradually drops in the provider-focused model, unlike increasing in other models. This result implies that when the provider recommends items that need to be sold quickly, without paying

much attention to whether they match the consumer’s preferences, consumers often find these recommendations less appealing. As a result, they are more likely to reject them.

Interestingly, models that incorporate some preference weighting tend to converge to acceptance rates that are similar to the consumer-focused model, with only slight differences of less than 1%. This observation indicates that while the consumer-focused model has the strongest alignment with consumer preferences and needs, weighted models still achieve comparable acceptance rates. It means that consumers are highly likely to accept substitutions even when recommendations are not perfectly tailored but reasonably close to their preferences, which eventually results in a better overall resource allocation.

5.3 Potential in Food Waste Reduction

Our result represents the estimated percentage of food wasted at each timestep. Waste after acceptance is excluded but all other expired food items are included. It shows that the percentage of food waste increases early stages but gently decreases after a certain point. The optimal model ($\omega = 0.2$) lowers the waste below the consumer-focused model and is close to the provider-focused model. That is, the optimal model shows only a small difference in food waste compared to the provider-focused model, even though the model considers the provider’s benefit less.

6 Limitations and Future Work

Our proposed model faces some limitations. First, MEAL elides nutritional factors and health considerations and recommends items solely relying on explicit preferences toward each food item given by consumers. Likewise, attributes such as socioeconomic background, culture, religion, and other diversity across communities remain challenging for optimization.

Incorporating additional stakeholder types would provide a more holistic view but complicate ensuring well-being, fairness, and trust among the stakeholders.

7 Conclusion

Achieving equitable food distribution requires a multifaceted endeavor that meets various goals. MEAL seeks to optimize the allocation strategy toward maximizing the rewards for consumer satisfaction and provider benefit, employing Q-learning. Our findings highlight that the right balance of the stakeholders’ objectives enhances consumer satisfaction while maximizing provider benefits. Our experiments simulate the society aligning with theoretical literature and other empirical findings in the relevant fields. Such alignment reinforces the robustness and applicability of our proposed method in real-world scenarios.

Acknowledgments. This work is supported by the NSF under grant 2125600.

References

1. America, F.: Our work (2024), <https://www.feedingamerica.org/our-work>, [Accessed 2024-10-06]
2. Iiyama, T., Kitakoshi, D., Suzuki, M.: Optimizing food allocation in food banks with multi-agent deep reinforcement learning. In: Proceedings of the International Conference on Technologies and Applications of Artificial Intelligence (TAAI). pp. 203–208. IEEE, Tainan, Taiwan (Dec 2022). <https://doi.org/10.1109/TAAI57707.2022.00045>, <https://doi.org/10.1109/TAAI57707.2022.00045>
3. Miller, M., Holston, D., Losavio, R.: Client-choice food pantry guide (2021), <https://www.lsuagcenter.com/profiles/aiverson/articles/page1614284604730>, [Accessed 2024-10-10]
4. National Academies of Sciences, Engineering, and Medicine: Understanding food waste, consumers, and the US food environment. In: Schneeman, B.O., Oria, M. (eds.) A National Strategy to Reduce Food Waste at the Consumer Level. The National Academies Press, Washington, DC (2020). <https://doi.org/10.17226/25876>, <https://nap.nationalacademies.org/catalog/25876/a-national-strategy-to-reduce-food-waste-at-the-consumer-level>
5. Ogundijo, D.A., Tas, A.A., Onarinde, B.A.: Age, an important sociodemographic determinant of factors influencing consumers’ food choices and purchasing habits: An english university setting. *Frontiers in Nutrition* **9**, 858593 (2022). <https://doi.org/10.3389/fnut.2022.858593>, <https://www.frontiersin.org/journals/nutrition/articles/10.3389/fnut.2022.858593>
6. Sprent, P., Smeeton, N.C.: Methods for Two Independent Samples, chap. 6, pp. 151–161. CRC press, New York (2016)
7. Tahtinen, L., Costa, N., Long, Y., Roberts-Harry, G.: SDG Good Practices: A Compilation of Success Stories and Lessons Learned in SDG Implementation. United Nations Department of Economic and Social Affairs, New York, 2 edn. (2022)
8. USDA: Food security and nutrition assistance (2023), <https://www.ers.usda.gov/data-products/ag-and-food-statistics-charting-the-essentials/food-security-and-nutrition-assistance/>, [Accessed 2024-10-06]
9. Watkins, C.J., Dayan, P.: Q-learning. *Machine Learning* **8**, 279–292 (1992). <https://doi.org/10.1007/BF00992698>, <https://doi.org/10.1007/BF00992698>
10. Webb, R.: Two Proportion Z-Test and Confidence Interval, chap. 9.3, p. 320. Portland State University Library, Portland, Oregon (2021), <https://pdxscholar.library.pdx.edu/pdxopen/36/>
11. Yang, L., Hsieh, C.K., Yang, H., Pollak, J.P., Dell, N., Belongie, S., Cole, C., Estrin, D.: Yum-Me: A personalized nutrient-based meal recommender system. *Transactions on Information Systems (TOIS)* **36**(1), 1–31 (2017). <https://doi.org/10.1145/3072614>, <https://doi.org/10.1145/3072614>
12. Zoha, N., Hasnain, T., Ivy, J.: Tradeoff between geographic and demographic equity in food bank operations. In: IISE Annual Conference Proceedings. pp. 1–6. Institute of Industrial and Systems Engineers (IISE), Seattle, Washington (may 2022), <https://www.proquest.com/scholarly-journals/tradeoff-between-geographic-demographic-equity/docview/2715836589/se-2>

Appendix

A.1 US Food Bank System

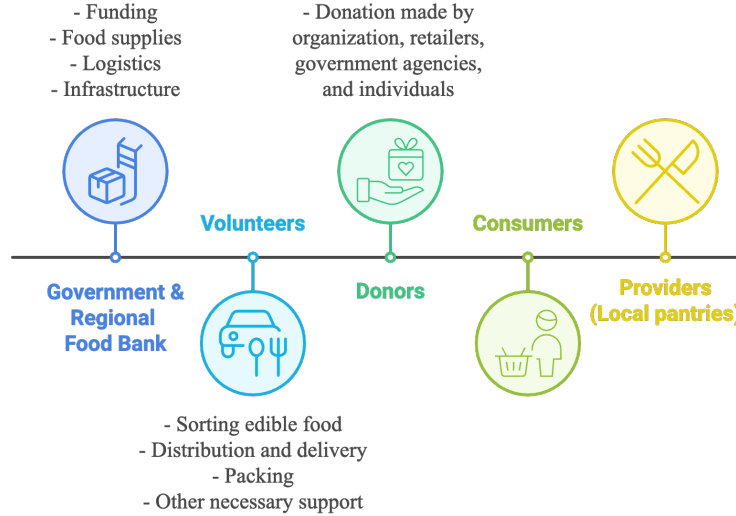


Fig. 1: Food distribution system, based on the US setting.

A.2 Concept of Operations in MEAL

We envision that consumers register with the food-sharing app by providing their profiles (e.g., household information). Consumers indicate preferences for some food items, e.g., fresh fruits and vegetables, milk, and whole grains. Then, they request food items as they need. Based on the inventory availability, community demands, and the consumer's profile and past selections, the app recommends alternative items from the same categories if one or more requested items are not available. The consumer can choose to accept or reject these substitutions and indicate their satisfaction with the accepted items, which the app uses to refine its suggestions.

Fig. 2 illustrates our conception. The agent serves as a mediator between consumers and a provider using consumer preferences and profiles to form the foundation for personalized recommendations. The agent received the provider's inventory information to make accurate up-to-date recommendations. Then, it aggregates demand and trends, estimates the level of prosociality of consumers and the goodness of food items, and processes interactions so that all parties benefit.




Fig. 2: Model architecture.

In general, the app cannot always recommend the consumer’s most preferred items. For instance, if apples have a higher demand than available stock, the app might suggest oranges. Doing so helps ensure as many people as possible get what they need and keeps the food bank running smoothly. Thus, consumers and providers have different perspectives. MEAL recognizes complexity by modeling consumers focusing on household needs and preferences, and a provider managing availability and community demand.

A.2.1 Model Formulation

Formally, we define the above problem as a Partially Observable Markov Decision Process (POMDP) where an agent (recommender) interacts with the environments (consumers and food provider) over time to maximize cumulative rewards of combined benefits. $\langle S, A, T, R, O, \Omega, \gamma \rangle$, where $s \in S$ is a finite set of states (i.e., consumer preferences and profiles, inventory status), $a \in A$ is a finite set of actions (i.e., the possible recommendations), T is a set of transition probabilities between states (i.e., the probability of acceptance), O is a set of observations (i.e., whether the recommendation is taken or not, consumers’ satisfaction feedback), Ω is a set of conditional observation probabilities of receiving an observation $o \in O$ after taking action $a \in A$ at state s , R is a reward function (i.e., a combination of consumer satisfaction and provider’s benefit controlled by the weighting factor ω , as defined in Equation 2), and $\gamma \in [0, 1)$ is the discount factor.

$$r_\omega = \omega \cdot h + (1 - \omega) \cdot c \quad (1)$$

$$\omega^* = \arg \max_{\omega \in [0, 1]} r(\omega) \quad (2)$$

A.3 Empirical Study Design

To conduct this study, we built a simple app that follows the streamlined flow of food requests and recommends replacements. We recruited 49 (adult, US-based) volunteers without any restrictions to ensure diverse representation.

The study involves two sessions of three food-requesting flows each. One session does not have persuasive messages when recommending replacements; the




Fig. 3: User study design.

other does. All 49 participants completed both the *Persuasion* and *No persuasion* sessions but in randomized order to consider potential dropout in the middle of the study. In each episode, the participants choose items from a list of fruits, vegetables, and meats. In the treatment, we replace two items in each food category after each request, and the participants can choose to accept or reject the replacements. At the end of the sessions, the participants rated how satisfied they were with the replacements they accepted.

Table 1: Data summary and results

	No persuasion	Persuasion
Total responses	515	463
Accepted	322	295
Rejected	193	168
No satisfaction response	91	54
Acceptance percentage	62.52	63.71
Mean satisfaction	3.57	3.43
Median satisfaction	4	4

A.4 Experimental Setting

We evaluate our model through simulations to understand how prosocial decisions are made throughout interactions. The simulated environment comprises data consisting of three sets: consumer profiles, preference ratings, and food inventory. Since it is hard to acquire real-world food preference data and food bank availability, we arbitrarily approximated the values of food items in our simulation by seeding the survey results of food pantry needs [1].

A.4.1 Consumer Profile and Prosociality

The main agents in our model are the consumers. We have crafted a consumer community with unique profiles. For simplicity, each consumer’s profile includes

age, whether they have one or more children, whether they have dietary restrictions or disease, family size, and ratings towards food items. We set 33% of consumers as aged over 65 and 45% of consumers as having a child. The family size distribution followed the statistics derived from a survey: the mean is three, and the standard deviation is two [1].

A consumer may accept or reject a recommendation. The probability of acceptance hinges on two factors: the consumer’s preference and inherent willingness to yield. Consumers don’t know how much the provider gains from their decisions. Ratings for particular items may be undefined. If undefined, we estimate satisfaction with the most similar consumer preferences using cosine similarity.

A.4.2 Food Inventory

Our simulation necessitates a comprehensive and realistic dataset that encompasses not just the items but also their attributes. We obtained a food list from [2] (169 different items) and classified it into six categories that people request every day, which are meat, fruits and vegetables, dairy, eggs, cooking items (like oils and seasoning), and others. However, since the [2] data lacks the specific attributes we need, we augmented attributes with feasible assumptions as close to demands mentioned in [1]. For simplicity, we limit to considering quantity, expiration date, and perishability as key components of setting urgency of allocation.

A.4.3 Trade-Offs: Provider versus Consumer

We evaluate various weightings to determine the optimal value of ω , as in Equation 2. We observe that the weighted models surpass the consumer-focused model in cumulative satisfaction, demonstrating the effectiveness of MEAL.

A.5 Visualizations of Results

To verify our model, we conduct simulations with 1,000 agents, each corresponding to consumers, one agent corresponding to the provider, and MEAL agent acting as a moderator between the consumers and the provider.

The parameter ω ranges between completely provider-focused valuation ($\omega = 0$) and completely consumer-focused ($\omega = 1$), with increment of 0.1. H and C are updated each time a particular recommendation is taken.

To evaluate our model’s performance, we consider two distinct values for the weighing factor (ω): 0.2, optimal in our setting determined by Equation 2, and 0.5, which evenly considers both sides. The results consistently show that our model with the optimal value of the weighting factor achieves our goal of satisfying both stakeholders’ objectives. The model is trained with a learning rate (α) of 0.1, a discount factor (γ) of 0.9, an exploration rate (ϵ) of 0.1, and a prosociality weight (β) of 0.1.




Fig. 4: Cumulative provider benefit. The provider-focused model gains the most while the consumer-focused model gains the least.




Fig. 5: Cumulative consumer satisfaction. Weighted models have the potential to achieve higher satisfaction.

References

1. Caspi, C.E., Davey, C., Barsness, C.B., Gordon, N., Bohen, L., Canterbury, M., Peterson, H., Pratt, R.: Needs and preferences among food pantry clients. *Preventing Chronic Disease* **18** (2021). <https://doi.org/10.5888/pcd18.200531>, <https://doi.org/10.5888/pcd18.200531>
2. USDA: What we eat in America (WWEIA) database | Ag Data Commons 2023 (Mar 2023), <https://data.nal.usda.gov/dataset/what-we-eat-america-wweia->




Fig. 6: Acceptance rates. Weighted models and the consumer-focused model achieve an increasing acceptance rate while the provider-focused model does not.




Fig. 7: Food waste tendency

database, [Accessed 2023-12-13]