This article has been accepted for publication in IEEE Internet Computing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MIC.2025.3590008

MEAL: Model of Empathy Augmented Logistics for Food Security

Seoyeong Park, *North Carolina State University, Raleigh, NC, USA*

Munindar P. Singh, *North Carolina State University, Raleigh, NC, USA*

Abstract—Millions globally lack reliable access to nutritious food. 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 in the network. However, they largely overlook consumer preferences and constraints. As a result, the resulting allocations lead to consumers either using foods they don't care for or discarding such foods, leading to food waste.

This paper presents a new model, evaluated via human study and agent-based simulation, that shows how to combine the consumer and provider perspectives. We find that persuasive messages that include individual circumstances and the social context can promote prosociality and empathy, leading to improved outcomes overall.

1. Introduction

Food insecurity is the condition of a household having poor access to adequate food and reduced quality of food intake[1]. One-eighth (approximately 17 million) of US households experience food insecurity[1].

The US foodbank system (Figure 1) is a nonprofit organization that reduces food waste and alleviates food insecurity by collecting, storing, and distributing food to those in need[2]. The federal government provides funding and capabilities to procure, store, transport, and distribute food[1]. Local foodbanks (providers) may receive donations from organizations, retailers, and individuals as well as allocations from regional foodbanks.

Ensuring equitable distribution is difficult when supplies are in shortage, 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. Current research addresses logistic efficiency [1] or concentrating on consumers' tastes [2], but not on both aspects together.

> In contrast, the objective of this paper is to produce equitable allocations that balance the logistic efficiency and consumer needs under consideration of supply constraints.

We propose **MEAL** for Model of Empathy Augmented Logistics for Food Security. MEAL optimizes food allocation and simulates consumer decisions. MEAL allocates food by considering both consumer needs and societal objectives such as reducing food waste and improving equity. MEAL incorporates *prosociality* by encouraging consumers to willingly accept less desirable allocations so that others in greater need (e.g., children or those who are sick) get more of the resources in shortage. MEAL naturally accounts for changing consumer preferences and resource availability.

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 than models that consider only one side, either consumers or providers. Additionally, a *human study* shows that persuasive messages, especially those that fit individual circumstances and the social context, can promote prosociality.

XXXX-XXX © 2021 IEEE
Digital Object Identifier 10.1109/XXX.0000.0000000

[1]USDA, Food security and nutrition assistance. https://www.ers.usda.gov/data-products/ag-and-food-statistics-charting-the-essentials/food-security-and-nutrition-assistance

[2]Feeding America. https://www.feedingamerica.org/our-work

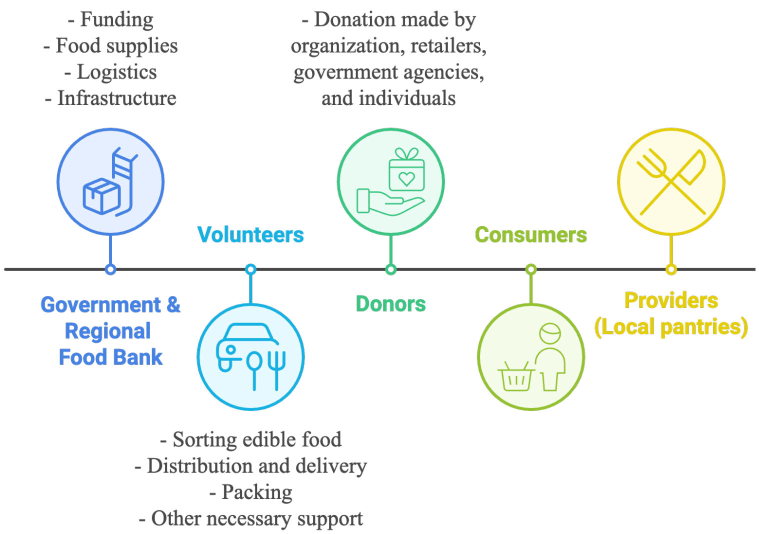

FIGURE 1: Food distribution system, based on the US setting.

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 foodbank, which has all the power and the food, decides how to allocate it to food-insecure households. 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 restructure the problem such that other acceptable allocations can be found. Our approach builds on key principles: social welfare, equity, prosociality, and empathy.

Social welfare reflects the greatest good of the greatest number [3]: it is the sum of the welfare of all participants, referred to as the *principle of utility*. Social welfare is the objective in traditional thinking on allocation. Here, a consumer's welfare is high when they get the foods they want; a provider's welfare reflects how low its operational costs are.

Equity is a property of a resource allocation that reflects not blind equality, which was famously mocked by Anatole France [4, p. 91][3]. In our setting, equity captures the intuition that all consumers have an equal opportunity to receive acceptable or desirable food allocations *in light of their specific needs*. For example, equity is enhanced if children are given more milk than adults. Equity also accounts for certain consumers being consistently favored over others, enhancing a balanced distribution and avoiding biases. Tracking and adjustment of the allocation patterns ensure that all consumers have fair access to preferred food items regardless of their circumstances.

Prosociality indicates an attitude or behavior directed to benefit others. Here, a prosocial consumer concedes their preferred foods to others and voluntarily accepts less preferred foods in short supply. MEAL promotes prosociality in consumers by helping bring out their empathy for others, leading to higher social welfare and equity.

Empathy involves understanding another person's situation and sharing their feelings [5]. Inequity-aversion theory [6] posits that people may be self-interested, but their decisions are affected by how relatively poorly others fare. Thus, empathy is a key driver of prosociality. MEAL is conceived on similar intuitions to promoting prosociality to system-level outcomes.

2.1. Stakeholders

We consider two main types of stakeholders: consumers and providers. *Consumers* are households served through MEAL agent. They aim to acquire food items that align with their preferences and needs. As consumers interact with the agent, 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,

[3] "The poor must work for this, in presence of the majestic quality of the law which prohibits the wealthy as well as the poor from sleeping under the bridges, from begging in the streets, and from stealing bread."

This article has been accepted for publication in IEEE Internet Computing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MIC.2025.3590008

household status, and willingness to make prosocial choices [7, 8]. 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[4]. 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. The provider has a measure of goodness for an allocation influenced by points such as:

- C_1 Perishable items have short shelf lives and items that demand refrigeration are expensive to store. Conversely, nonperishable goods can be held in stock at a lower cost.
- C_2 Limited-quantity foods may need to be set aside for consumers with priority needs, such as for infants.
- C_3 Foods with high quantities must be distributed quickly before they expire and to avoid taking up space.

Providers earn more benefits if they allocate food items with higher goodness.

2.2. 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.

Figure 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 receives the provider's inventory information to make accurate up-to-date recommendations.

[4]Client-Choice Food Pantry Guide. https://www.lsuagcenter.com/profiles/aiverson/articles/page1614284604730

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. 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.

2.3. Research Questions

Accordingly, this study investigates these research questions.

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

RQ$_{\textbf{persuasion}}$ Do persuasion and empathy influence human decisions about food and prosocial behavior?

3. Empirical Evaluation with Humans

Interactions with machines can largely influence human behavior. Therefore, MEAL's prosocial decision can shape the consumer's subsequent behaviors.

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.

3.1. Forming a Persuasive Message

Displaying empathy can build trust toward machines, which can lead to higher satisfaction and long-term engagement, particularly when invoking positive social behaviors. Paiva et al. [9] state that the prosocial decision taken by an agent influences the consumer's subsequent actions toward the agent, and so does sustaining prosocial communities. In our setting, fostering empathy towards both the foodbank system and other consumers can increase willingness to accept less preferred food items. Persuasive messages [10] can help consumers understand what is happening behind them.

3.2. Study Design

To conduct this study, we built a simple app that follows the streamlined flow of food requests and rec-

This article has been accepted for publication in IEEE Internet Computing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MIC.2025.3590008

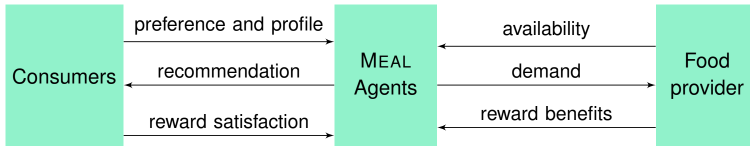

FIGURE 2: Model architecture.

ommends replacements. We recruited 49 (adult, US-based) volunteers without any restrictions, such as education level and socio-economic status, to ensure diverse representation. The details of the participant demographics are in the supplementary materials.

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

3.3. Data Collection

We collected the same profile elements (age, household size, whether having children, and whether having dietary restrictions) as we set in our simulation. Participants rated their preferences toward given food items. They chose a number between 0 (no experience) and 5 (extremely like) for each item. To diversify the data with a small number of participants and interaction, the participants were asked to select at least three items in each category (fruits, vegetables, and meats) in each requesting flow. At the end of the sessions, the participants were asked to rate satisfaction scores of accepted items (1 to 5). We considered rejected items as unsatisfactory.

After eliminating invalid responses, we had 515 responses from *No persuasion*, 463 responses from *Persuasion*, and 23 responses for the survey.

3.4. Interpretation

Finding 1. Persuasion

General persuasive messages have little influence on increasing satisfaction or acceptance rates.

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 [11] 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 agent 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 [11] 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 acceptance of substitutes may not be accompanied by enhancing the consumer's experience.

Finding 2. Empathy

Empathy is triggered when consumers clearly understand the need for social well-being and how their actions affect the community, unless customers have strong constraints.

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. For example, one participant stated *It is fine if recommendations have similar properties, such as I can replace chicken breast with duck breast and make similar food. Replacing a berry with a berry is fine too*. Additionally, the freshness of recommendations is another criterion consumers consider, with one participant noting *I would consider when recommendations have equal quality and similar freshness to me.*

This article has been accepted for publication in IEEE Internet Computing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MIC.2025.3590008

Finding 3. Personalized Persuasive Messages

Well-constructed persuasive messages explaining personal and social contexts can be much more effective.

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 likely change their decision (of refusing a recommended item) regardless of their personal situations if they knew their choice would promote social well-being, unless they had strong dietary restrictions. For instance, one participant answered that they would try a recommended item: *if I know someone who is in a similar situation to mine and by getting this item can benefit them then I would be open to a recommended item*. This suggests that the satisfaction and acceptance rates would become higher if a persuasive message described how their decision affects promoting overall social well-being.

Interestingly, some people expressed a general trust in the foodbank system. They believe that the system operates fairly and efficiently, prioritizing good outcomes. This finding suggests that building trust can be a critical means for nudging consumers to make prosocial decisions that benefit not just themselves but also society as a whole. Thus, our model should be transparent in its decision-making for the recommendations it generates and needs to be aware of the long-term consequences of recommendations.

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. Therefore, we define cumulative consumer satisfaction H and cumulative provider's benefit C as below:

$$H = \sum h_{u,d,t} \text{ and } C = \sum c_{u,d,t}$$

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 (Equation 2). Therefore, the agent's overall reward for the decision-making objective is a weighted combination of satisfaction and provider benefit as Equation 1.

$$r_\omega = \omega \cdot h + (1 - \omega) \cdot c \tag{1}$$

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

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.

We model this process using Q-learning. At each time step, the model recommends a food item $f \in F$ randomly or with the highest Q-value to consumer u. If f is accepted, the model collects the consumer satisfaction h and provider benefit c, and computes the reward r using Equation 1. 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. The details of the algorithm and notation we used are in the supplementary materials.

5. Experimental Setting

We evaluate our model through simulations to understand how prosocial decisions are made throughout interactions. The simulated environment comprises data

This article has been accepted for publication in IEEE Internet Computing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MIC.2025.3590008

consisting of three sets: consumer profiles, preference ratings, and food inventory. Since it is hard to acquire real-world food preference data and food availability, we arbitrarily approximated the values of food items in our simulation by seeding the survey results of food pantry needs [12]. Our approximation method provides a realistic representation of the effectiveness of our model and mimics the interactive environment and dynamic consumer behaviors. In each episode, the model recommends food items, and the consumer accepts or rejects them. An episode terminates when the inventory becomes empty or reaches a predefined number of steps.

5.1. 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 [12].

A consumer may accept or reject a recommendation. The probability of acceptance (Equation 3) 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. In equation 3, $p_{u,d_{u,f,t}}$ represents a numerical rating toward food $d_{u,f,t}$, where a food f for consumer u at time step t. e_u is the estimated prosociality of consumer u, and β is a weighting factor fixed to 0.9.

$$\text{Probability of acceptance} = \beta \cdot p_{u,d_{u,f,t}} + (1-\beta) \cdot e_u \quad (3)$$

In the real world, willingness to yield can be affected by many personality traits that are not fully observable. To quantify consumer prosociality, we illustratively break consumer profiles into three age groups, three family sizes, a binary indicator of whether they have dietary restrictions, and a binary indicator of whether they have one or more children, assuming that people who have fewer constraints are more likely to be able to think of others.

5.2. 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 USDA[5] (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 USDA[5] data lacks the specific attributes we need, we augmented attributes with feasible assumptions as close to demands mentioned in Caspi et al. [12]. For simplicity, we limit to considering quantity, expiration date, and perishability as key components of setting urgency of allocation. We arbitrarily formulate the values of each item representing the criteria mentioned in Section 2.1. If the quantity becomes zero, the associated value becomes zero, and the item should not be recommended. If the expiration date has passed while the item is still available, all the remaining becomes wasted.

6. Results

This section presents the experimental results of MEAL. 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. 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.

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

[5]USDA: Ag Data Commons User Guide. https://www.nal.usda.gov/services/agdatacommons

This article has been accepted for publication in IEEE Internet Computing. This is the author's version which has not been fully edited and content may change prior to final publication. Citation information: DOI 10.1109/MIC.2025.3590008

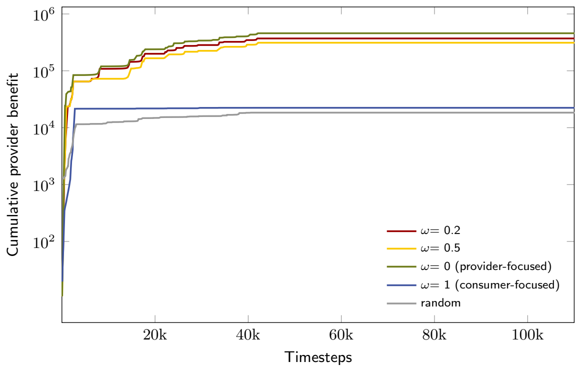

(a) Cumulative provider benefit. The provider-focused model gains the most while the consumer-focused model gains the least.

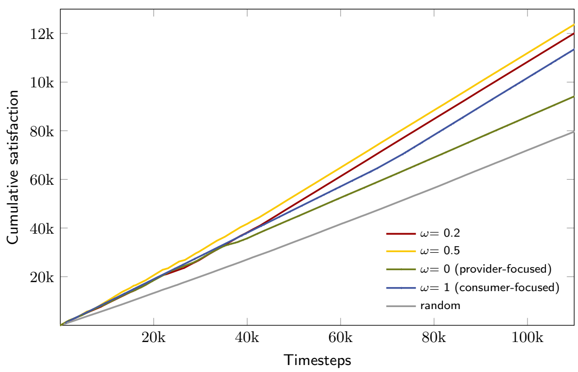

(b) Cumulative consumer satisfaction. Weighted models have the potential to achieve higher satisfaction.

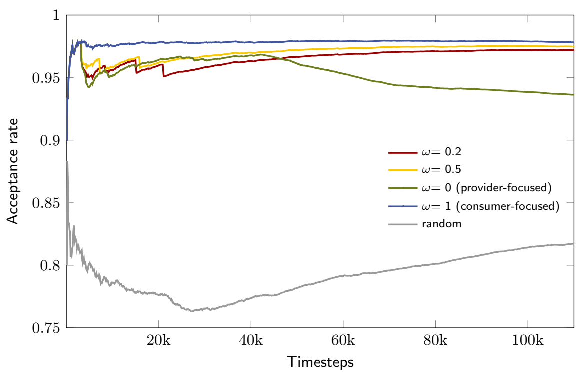

(c) Acceptance rates. Weighted models and the consumer-focused model achieve an increasing acceptance rate while the provider-focused model does not.

FIGURE 3: Simulation results with different ω values

of the weighting factor achieves our goal of satisfying both stakeholders' objectives.

6.1. Consumer Satisfaction

Finding 4. Satisfaction

Incorporating the provider's perspective alongside the customer preferences improves customer satisfaction beyond just focusing on the customers.

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.

As shown in Figure 3a, 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. However, Figure 3b shows high cumulative consumer satisfaction, underscoring the superiority of weighted models. Consumer satisfaction visibly improves, unlike our original expectation that both stakeholders would 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. This result indicates that a larger number of consumers get recommendations from MEAL recommends items that are closer to their preferences.

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. This outcome emphasizes that consumer satisfaction is enhanced when allocations are perceived as fair to both consumers and overall community needs.

6.2. Acceptance of Recommendations

Finding 5. Acceptance

The provider-focused model shows a drop in acceptance rate as time goes by while the weighted models achieve close to the consumer-focused model.

How much the model skews to consumer satisfaction affects the acceptance rate. The higher the weight on consumer preferences, the higher the acceptance rate. As shown in Figure 3c, 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. 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.

6.3. Potential in Food Waste Reduction

Finding 6. Waste Reduction

Achieving the right balance between the provider's benefit and consumer preferences leads to a reduction in waste.

Figure 4 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 (ω = 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.

7. Related Work

One prevalent technical approach to implementing recommendation systems is to employ reinforcement learning (RL) algorithms under Partially Observable Markov Decision Problems (POMDPs) [13]. Recent studies leverage RL to handle multiple objectives of recommendation systems in various domains [14]. Unlike traditional recommendation systems, multistakeholder recommendation systems address the objectives and needs of more than one party. These systems can help overcome limitations in incorporating system objectives that apply across stakeholders. For example, Kwak and Huettel [15] apply RL to establish a decision-making paradigm for oneself and charity and understand differences in prosocial tendencies. However, despite the diverse applications of multistakeholder RL in recommendation systems, it has not been applied to address the food insecurity problem.

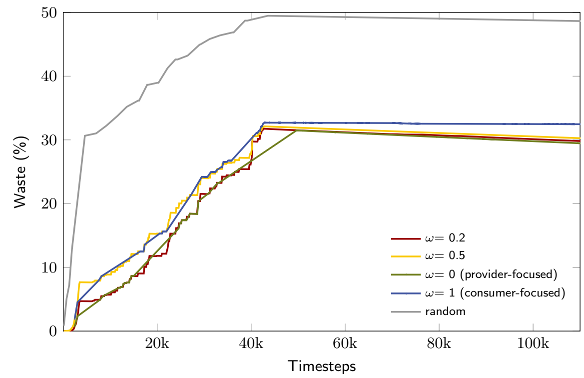

FIGURE 4: Food waste tendency

8. 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. However, considering nutrition, dietary needs, and health conditions could improve the suitability of suggested replacements for particular recipes. Likewise, attributes such as socioeconomic background, culture, religion, and other diversity across communities remain challenging for optimization.

MEAL omits human factors of decision-making such as cognitive biases, emotions, and personality traits. We may explore further psychological and social science approaches to capture more accurate measurements of human cognition and design persuasive strategies to nudge consumers to shift their preferences for the social good or to make healthier choices.

Incorporating additional stakeholder types (Figure 1) would provide a more holistic view but complicate ensuring well-being, fairness, and trust among the stakeholders.

One last challenge we encounter pertains to obtaining real-world data. This limitation may affect the realism of our simulation, especially in assessing the provider's benefit.

Acknowledgments

This work is supported by the NSF under grant 2125600.

REFERENCES

1. N. Zoha, T. Hasnain, and J. Ivy, "Tradeoff between geographic and demographic equity in food bank operations," in *IISE Annual Conference Proceedings*. Seattle, Washington: Institute of Industrial and Systems Engineers (IISE), may 2022, pp. 1–6. [Online]. Available: https://www.proquest.com/scholarly-journals/tradeoff-between-geographic-demographic-equity/docview/2715836589/se-2
2. L. Yang, C.-K. Hsieh, H. Yang, J. P. Pollak, N. Dell, S. Belongie, C. Cole, and D. Estrin, "Yum-Me: A personalized nutrient-based meal recommender system," *Transactions on Information Systems (TOIS)*, vol. 36, no. 1, pp. 1–31, 2017. [Online]. Available: https://doi.org/10.1145/3072614
3. R. Plant, "The greatest happiness." *Journal of Medical Ethics*, vol. 1, no. 2, p. 104, 1975. [Online]. Available: https://doi.org/10.1136/jme.1.2.104
4. A. France, *The Red Lily*. New York: Dodd-Mead, 1910. [Online]. Available: https://archive.org/details/theredlily00franiala/page/90/mode/2up?view=theater&q=majestic
5. M. L. Hoffman, *Empathy and Moral Development: Implications for Caring and Justice*. Cambridge: Cambridge University Press, 2000. [Online]. Available: https://doi.org/10.1017/CBO9780511805851
6. E. Fehr and K. M. Schmidt, "A theory of fairness, competition, and cooperation," *The Quarterly Journal of Economics*, vol. 114, no. 3, pp. 817–868, 1999. [Online]. Available: https://doi.org/10.1162/003355399556151
7. D. A. Ogundijo, A. A. Tas, and B. A. Onarinde, "Age, an important sociodemographic determinant of factors influencing consumers' food choices and purchasing habits: An english university setting," *Frontiers in Nutrition*, vol. 9, p. 858593, 2022. [Online]. Available: https://www.frontiersin.org/journals/nutrition/articles/10.3389/fnut.2022.858593
8. National Academies of Sciences, Engineering, and Medicine, "Understanding food waste, consumers, and the US food environment," in *A National Strategy to Reduce Food Waste at the Consumer Level*, B. O. Schneeman and M. Oria, Eds. Washington, DC: The National Academies Press, 2020. [Online]. Available: https://nap.nationalacademies.org/catalog/25876/a-national-strategy-to-reduce-food-waste-at-the-consumer-level
9. A. Paiva, F. Correia, R. Oliveira, F. P. Santos, and P. Arriaga, "Empathy and prosociality in social agents," in *The Handbook on Socially Interactive Agents: 20 Years of Research on Embodied Conversational Agents, Intelligent Virtual Agents, and Social Robotics Volume 1: Methods, Behavior, Cognition*, ser. ACM Books, B. Lugrin, C. Pelachaud, and D. R. Traum, Eds. New York: ACM / Morgan & Claypool, 2021, vol. 37, pp. 385–432. [Online]. Available: https://doi.org/10.1145/3477322.3477334
10. L. Shen, "Mitigating Psychological Reactance: The Role of Message-Induced Empathy in Persuasion," *Human Communication Research*, vol. 36, no. 3, pp. 397–422, 07 2010.
11. P. Sprent and N. C. Smeeton, *Methods for Two Independent Samples*. New York: CRC press, 2016, ch. 6, pp. 151–161.
12. C. E. Caspi, C. Davey, C. B. Barsness, N. Gordon, L. Bohen, M. Canterbury, H. Peterson, and R. Pratt, "Needs and preferences among food pantry clients," *Preventing Chronic Disease*, vol. 18, 2021. [Online]. Available: https://doi.org/10.5888/pcd18.200531
13. G. Shani, D. Heckerman, R. I. Brafman, and C. Boutilier, "An MDP-based recommender system," *Journal of Machine Learning Research*, vol. 6, no. 9, pp. 1265–1295, 2005.
14. M. M. Afsar, T. Crump, and B. H. Far, "Reinforcement learning based recommender systems: A survey," *ACM Computing Surveys*, vol. 55, no. 7, pp. 1–38, 2023. [Online]. Available: https://doi.org/10.1145/3543846
15. Y. Kwak and S. A. Huettel, "Prosocial reward learning in children and adolescents," *Frontiers in Psychology*, vol. 7, p. 1539, 2016. [Online]. Available: https://doi.org/10.3389/fpsyg.2016.01539

Seoyeong Park is a PhD student in Computer Science at NC State University. Her interests include persuasion, user experience, cognitive modeling, and multiagent systems. Contact her at spark43@ncsu.edu.

Munindar P. Singh is a Professor in Computer Science at NC State University. His interests include the engineering and governance of sociotechnical systems and AI ethics. Singh is a Fellow of AAAI, AAAS, ACM, and IEEE and a former Editor-in-Chief of *IEEE Internet Computing* and *ACM Transactions on Internet Technology*. Contact him at singh@ncsu.edu.