Learning to Allocate Fairly

Dawson Ren, Seyed Iravani
Department of Industrial Engineering & Management Sciences

Northwestern University

Northwestern
OFFICE OF UNDERGRADUATE RESEARCH

Motivations & Questions

The problem of fair Sequential Resource Allocation (SRA) arises when we distribute a resource among agents arriving one by one with uncertain demands.

Examples:
• Distributing medical supplies (medicine, ventilators, vaccines) during a pandemic [1].
• Equitable food distribution to food pantries [2].
• CPU/GPU resources for high-performance computing (HPC) centers [3].

Previous studies assume an attitude towards fairness (quantified by α), but α impacts important performance measures such as efficiency/waste [4]. This work generalizes previous studies by allowing decision-makers to choose their own level of *fairness (α)*.

Questions:

How should we handle uncertainty in demand as α changes?
What is the tradeoff between fairness (α) and efficiency?

Methods

Formulate the SRA problem as a Markov Decision Process, which is computationally infeasible to solve exactly as the number of agents grows.

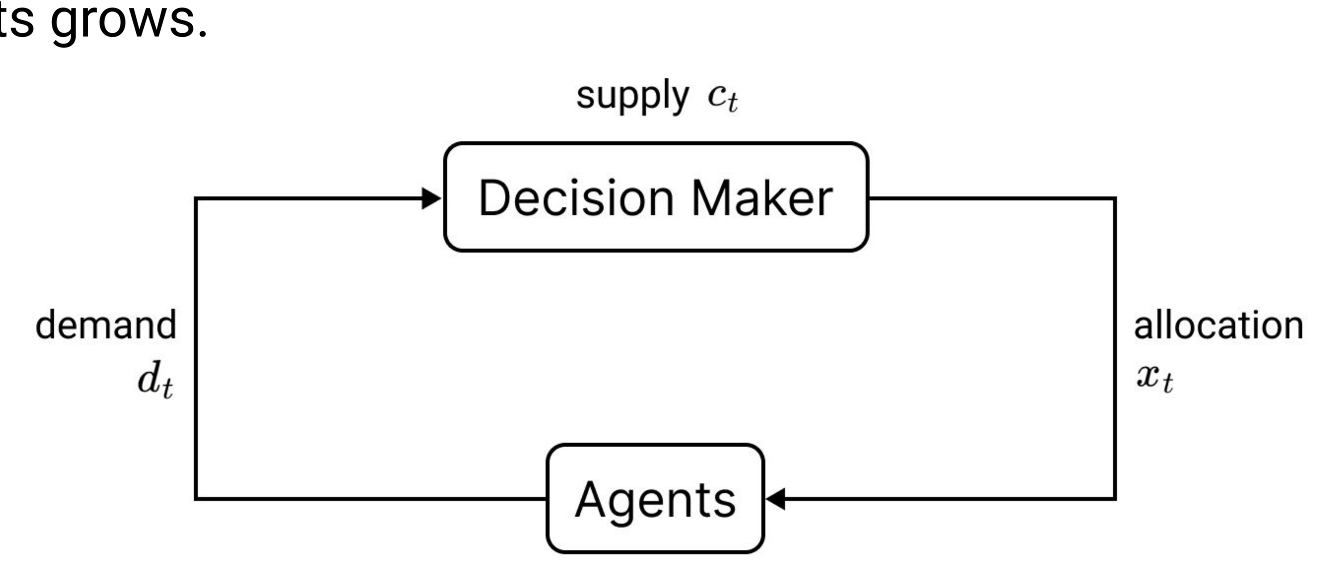

Instead of solving exactly, identify a *heuristic algorithm* and evaluate performance on real and synthetic datasets. The real dataset is a collection of 70 Mobile Food Pantries in New York.

Results

Theorem
When there is **no uncertainty in demand**, making sure that everyone receives the same proportion of their demand is optimal, no matter the fairness (α).

Corollary
The optimal allocation when there is **no uncertainty** is

$$x(c_t, d_t) = \min\left\{d_t, \frac{d_t}{d_t + \sum_{i=t+1}^{N} \mathbb{E}[d_i]} c_t\right\}$$

We now introduce the Projected Proportional Allocation – Alpha-Fair (PPA-AF) *heuristic algorithm*, which adds an **uncertainty parameter (λ)** that weighs the standard deviation (i.e. uncertainty) of the future demand.

$$x(c_t, d_t) = \min\left\{d_t, \frac{d_t}{d_t + \sum_{i=t+1}^{N} \mathbb{E}[d_i] + \lambda\sigma_i} c_t\right\} \quad \text{(PPA-AF)}$$

When **λ < 0**, we are **penalizing uncertainty in demand** and reserving more of the supply for the current agent being served. Conversely, when **λ > 0**, we are **rewarding uncertainty in demand** by reserving supply in case future demands are higher than expected.

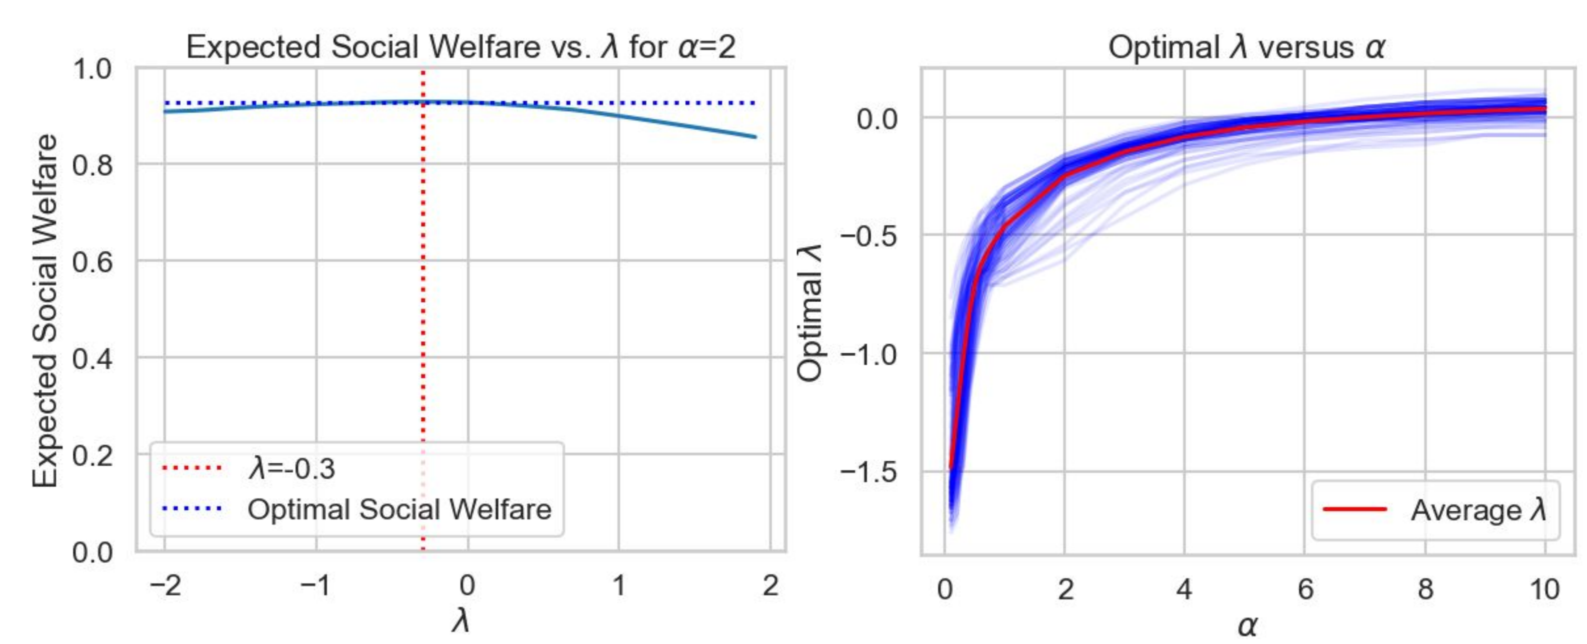

For N=3, Normal(μ=[15, 5, 10]), CV=0.4, 5 support points   For N=4, Normal(μ∈[10, 100]), CV=0.4, 5 support points

Given some **fairness (α)**, there is a value of the **uncertainty parameter (λ)** that maximizes the social welfare (left). Learning the value of **λ** gives us a relationship between fairness and uncertainty (right).

Observation
For the PPA-AF algorithm, there is a negative relationship between fairness and efficiency, and achieves lower waste than comparable algorithms, PPA (α=∞) and SAFFE (α=1).

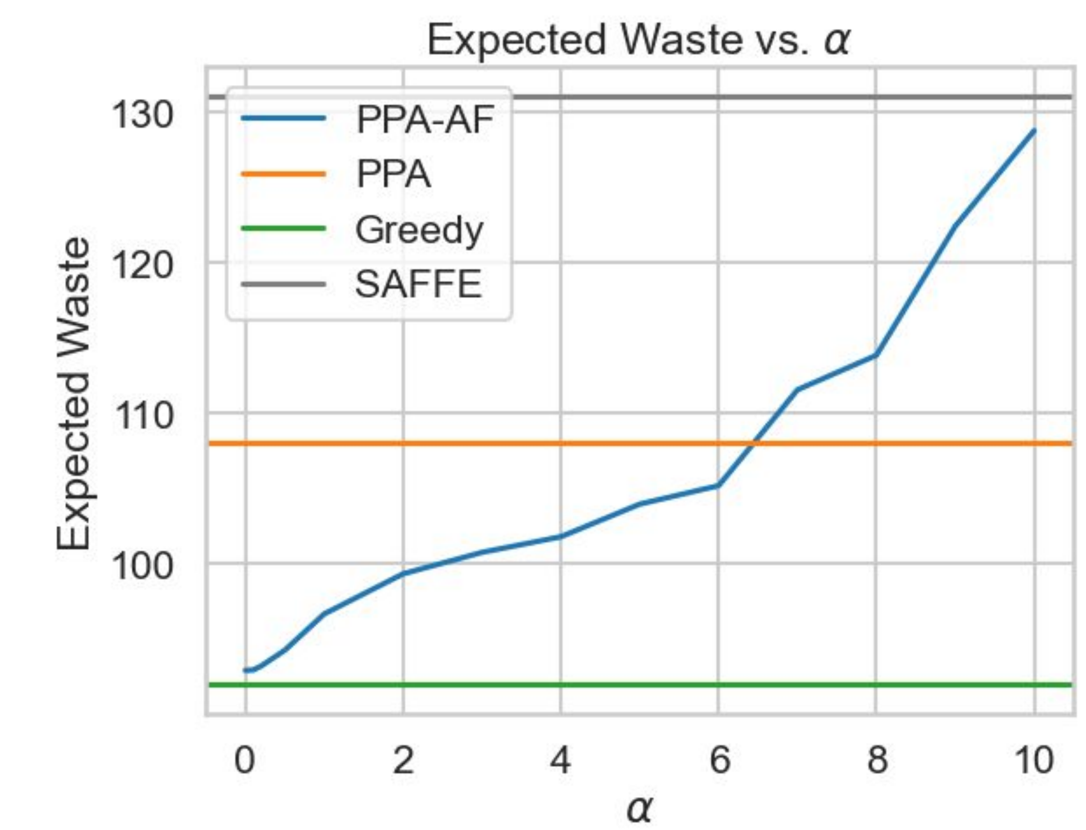

For N=70, Normal from FBST dataset, 10 support points

PPA-AF is able to outperform other algorithms that assume some level of fairness by adjusting how it responds to uncertainty in demand based on the desired fairness (α).

Acknowledgements

Thank you to Professor Seyed Iravani, Professor Sean Sinclair, and Haiqing Gao for their support and mentorship during this project. The study resulting in this presentation was assisted by a grant administered by Northwestern University's Office of Undergraduate Research. However, the conclusions, opinions, and other statements in this presentation are the author's and not necessarily those of the sponsoring institution.

Future Work

• Theoretically prove bounds on the efficiency and achievable social welfare of PPA-AF.
• Find a closed-form expression that relates the optimal uncertainty parameter (λ) and fairness (α). Currently, we use derivative-free optimization to find an optimal λ.
• Extend to multiple resource types.

Conclusions

• When we are **not very interested in fairness (α ≤ 1)**, we should **penalize uncertainty in demand**, and when **we are very interested in fairness (α >> 1)**, it *may be* optimal to **reward uncertainty in demand.** We can find the optimal
• **Efficiency decreases** as **fairness increases**, and a tradeoff curve can be constructed to help decision makers decide what level of fairness to pick.

[1] Vahideh Manshadi, Rad Niazadeh, and Scott Rodilitz. Fair dynamic rationing. Available at SSRN 3775895, 2021.
[2] Robert W Lien, Seyed MR Iravani, and Karen R Smilowitz. Sequential resource allocation for nonprofit operations. Operations Research, 62(2):301– 317, 2014.
[3] Parisa Hassanzadeh, Eleonora Kreacic, Sihan Zeng, Yuchen Xiao, and Sumitra Ganesh. Sequential Fair Resource Allocation under a Markov Decision Process Framework. Available at https://arxiv.org/abs/2301.03758
[4] Dimitris Bertsimas, Vivek F Farias, and Nikolaos Trichakis. On the efficiency-fairness trade-off. Management Science, 58(12):2234–2250, 2012.