

Participation behavior and social welfare in repeated task allocations

Qing Chuan Ye

*Department of Econometrics
Erasmus University Rotterdam
Rotterdam, The Netherlands
Email: ye@ese.eur.nl*

Yingqian Zhang

*Information Systems Group
Eindhoven University of Technology
Eindhoven, The Netherlands
Email: yqzhang@tue.nl*

Abstract—Task allocation problems have focused on achieving one-shot optimality. In practice, many task allocation problems are of repeated nature, where the allocation outcome of previous rounds may influence the participation of agents in subsequent rounds, and consequently, the quality of the allocations in the long term. We investigate how allocation influences agents’ decision to participate using prospect theory, and simulate how agents’ participation affects the system’s long term social welfare. We compare two task allocation algorithms in this study, one only considering optimality in terms of costs and the other considering optimality in terms of primarily fairness and secondarily costs. The simulation results demonstrate that fairness incentivizes agents to keep participating and consequently leads to a higher social welfare.

Keywords—Repeated task allocation; Prospect theory; Fairness; Participation behavior.

I. INTRODUCTION

Task allocation problems have focused on achieving one-shot optimality, which typically aims at finding the minimum cost allocations ([1]). In [2] a fair task allocation problem has been studied, where we propose a fair task allocation algorithm that assigns companies jobs based not only on their costs but also tries to allocate jobs to all participants as fairly as possible. The motivation of developing fair task allocation algorithms was inspired by an actual transportation situation in the port of Rotterdam in the Netherlands, where many small inter-terminal tasks need to be assigned to companies who have trucks that are already present in the port. Those trucks that come from the hinterland to drop or pick up containers often have spare time in between tasks. Hence, terminals could take advantage of these idle trucks by providing them with jobs that they can perform within the port while waiting for their next scheduled job.

We hypothesize in [2] that due to psychological factors, using an allocation algorithm with fairness as a main criterion will encourage companies’ participation in the repeated task allocation game. More participants ensure more supplies in the system, which will eventually lead to a higher social welfare. The objective of this paper is to test this hypothesis. We study a repeated task allocation problem. Besides the inter-terminal transportation example of the port of Rotterdam, another example of repeated task allocation would be

private taxi services, in which any party is free to take up a job, different from traditional taxi services. An incoming job may be proposed to several drivers in the neighbourhood, but eventually only one of them will be assigned the job. In these settings, participants share their idle resources, and therefore, it is important to ensure some portion of the market share to the players to encourage their participation.

We consider agents to be not completely rational. Hence, the assumption that an agent will participate in the game as long as its expected utility is non-negative may not hold anymore. Instead, we take into account psychological factors of agents that may curb the decision to participate, and model agents’ participation based on prospect theory ([3]). Prospect theory has been widely studied in behavioral economics ([4], [5]). It is a behavioral model that shows how people handle decisions that involve risk and uncertainty. However, this behavioral theory is rarely used in task or resource allocation problems. We show how we extend it to model agents’ decisions on participating in each round of the games (Section III). This participation probability is derived based on the previous allocation outcomes, and particularly, on an agent’s perception on its received proportion in comparison to other agents. In Section IV, we investigate how the allocation influences agents’ decision to participate by using two task allocation algorithms, of which one only looks at optimality in terms of costs, and the other looks at optimality in terms of primarily fairness and secondarily costs ([2]). In turn, we also look at the effect of agents’ participation on the social welfare, which is measured by the allocation quality.

II. PROBLEM DEFINITION

We first introduce the task allocation setting that is studied in [2]. We assume that the set of available jobs to be distributed among agents is known in advance by the central planner. We assume a set of time periods \mathcal{T} , consisting of T time periods. The set of jobs, denoted by \mathcal{J} consisting of a total of J jobs, comes with an earliest available time and a latest completion time for each job. We assume that jobs are independent. We define for each job $j_i \in \mathcal{J}$ its possible starting time as a mapping: $\mathcal{J} \times \mathcal{T} \mapsto \{0, 1\}$. When it is clear from the context, we abuse the notation and use j_i^t to

Table I
THE BIDS OF THREE COMPANIES INCLUDE DESIRED JOBS IN EACH TIME PERIOD AND THEIR ASSOCIATED COSTS.

| | t_1 | t_2 | t_3 | t_4 | t_5 |
|-------|------------|--------------------------|--------------------------|--------------------------|------------|
| k_1 | $j_1 : 20$ | | | | |
| k_2 | $j_1 : 30$ | $j_2 : 40$ $j_3 : 25$ | | | |
| k_3 | $j_1 : 10$ | $j_2 : 20$ $j_3 : 20$ | $j_3 : 25$ $j_4 : 25$ | $j_2 : 30$ $j_4 : 20$ | $j_5 : 20$ |

denote that job j_i is available at time period $t \in \mathcal{T}$. Once the set of jobs \mathcal{J} together with their possible starting times has been made available, a set of companies \mathcal{K} , consisting of K companies, may bid on individual jobs. In addition to the selection of jobs that a company $k \in \mathcal{K}$ wishes to perform, the company also needs to provide their available capacity n_k^t in time period t in which it is able to perform the jobs. We assume that each job takes up one unit of capacity and can be completed within one time unit. Furthermore, the company k needs to provide its desired compensation (or cost), $c(j_i, k)$, for the bid job $j_i \in \mathcal{J}$. A bid, B_k , from a company k is thus a tuple: $(\mathbf{c}_k, \mathbf{n}_k)$, where \mathbf{c}_k is a set of costs $c(j_i^t, k)$, and \mathbf{n}_k is a set of capacities n_k^t . Once all bids from the bidding companies \mathcal{K} have been collected, the auctioneer determines a task allocation $\pi : \mathcal{J} \times \mathcal{T} \times \mathcal{K} \mapsto \{0, 1\}$. If a job j_i is allocated, it will result in a fixed value V . For all unallocated jobs, their values are set to 0. The social welfare U given an allocation π is defined as the difference between the total value of allocated tasks and the total costs of performing the allocated tasks. In a typical task allocation problem, the objective is to maximize the social welfare by choosing an optimal allocation. Note that for this one-shot task allocation problem, we can use any existing min-cost max-flow algorithm to find the optimal allocation ([2]).

In [2], however, a fair algorithm is developed to ensure a max-min fair allocation to agents. The max-min fairness principle means that given a total of Z jobs, the number of jobs for any agent cannot be increased by at the same time decreasing the number of jobs of other agents that have the same number of jobs or less. Intuitively speaking, we want to have an allocation that distributes the set of jobs among the agents as evenly as possible. To meet the fairness criterion a polynomial-time fair method is developed consisting of two novel algorithms: IMaxFlow, which computes a max-min fair vector, i.e., the most even distribution over agents given all bids, and FairMinCost, which finds the minimum cost allocation that satisfies the given max-min fair vector. The output of these two algorithms is a max-min fair task allocation with the least total cost. We now use the following example to illustrate the algorithms used in [2] to obtain the minimum-cost and fair allocations.

Example 1: Suppose we have 5 jobs, all having a value of $V = 100$. The jobs can be done in certain time periods and three companies submit their bids, as shown in Table I. Any min-cost max-flow algorithm results in π_{mincost}

assigning j_3^2 to k_2 and $j_1^1, j_2^2, j_4^4, j_5^5$ to k_3 , with a total cost of 95, and social welfare of 405. For the fair allocation, using the algorithms in [2] we obtain the max-min fair allocation vector $\phi = (1, 1, 3)$ using IMaxFlow. Thereafter, using FairMinCost, we obtain the fair allocation π_{fair} , which assigns j_1^1 to k_1 , j_3^2 to k_2 and j_2^2, j_4^4, j_5^5 to k_3 , with a total compensation of 105, and social welfare of 395. ■

In this paper, we will extend the one-shot task allocation problem in [2] to a multi-round task allocation problem. We introduce round $r \in \mathcal{R}$. We then have tasks $j_{i,r} \forall j_{i,r} \in \mathcal{J}_r$, with \mathcal{J}_r the set of jobs in round r , $\mathcal{J}_r \subseteq \mathcal{J}$. In the repeated task allocation game with R rounds, the objective is to maximize the social welfare over all rounds, that is, to maximize $\sum_{r=1}^R U_r = \sum_{r=1}^R \sum_{j_{i,r} \in \mathcal{J}_r, \pi_r(j_{i,r}, \cdot) = 1} |j_{i,r}| \cdot V - c(j_{i,r}, k)$, where π_r is an allocation in round r . In addition, we assume agents decide for themselves whether they would like to participate in a round, which will be modelled as a participation probability $p(\psi^l)$ dependent on earlier allocation outcomes. The repeated task allocation in a round r goes as follows. After the auctioneer announces the available tasks, agent k decides on whether to participate or not by computing its participation probability $p(\psi_{kr}^l)$, and if participating, it submits the bid $b_{r,k}$ based on its availability and costs. As we do not study agents' bidding strategies in this paper, we simply assume agents submit their bids based on their true values. The auctioneer then decides on the task allocation π_r and the payments $c(j_{i,r}^t, k)$. Finally, agent k observes all participants' bids and the outcome π_r .

III. MODELLING PARTICIPATION BEHAVIOR

Prospect theory ([3]) demonstrates that people view their expected utility not in absolute terms, what one would expect from expected utility theory, but rather relative to a reference point. In addition, it indicates that people are loss-averse, where they would be more willing to take risks in order to avoid a loss, and they would avoid taking risks if it concerns a gain. There is a distinction between two phases in the choice process: (1) the editing phase, where a simpler representation of the outcomes of alternatives is obtained, as they are coded as gains or losses relative to a reference point; and (2) the evaluation phase, where the edited prospects are evaluated and the prospect of highest value is chosen. The overall value of an edited prospect is expressed by a weighting function and a value function. In [3], the authors propose the form of the value function given by

$$v(x) = \begin{cases} x^\alpha & \text{if } x \geq 0 \\ -\lambda(-x^\beta) & \text{if } x < 0 \end{cases} \quad (1)$$

where $0 \leq \alpha \leq 1$ and $0 \leq \beta \leq 1$ are coefficients determining the concavity and convexity for gains and losses, respectively, and $\lambda > 1$ is the loss-aversion coefficient.

The participation probability $p(\psi_{kr}^l)$ of agent k in round r is dependent on their experience in previous rounds. To model the participation decision using prospect theory, we

use the average proportion in the previous round over all companies as the reference point in the editing phase, and a positive (or negative) difference between a company's proportion in the previous round and the average proportion as a gain (or loss). The intuition is that if an agent feels being treated worse in comparison with others, she might be more uncertain and will care less about participating again, because the time and effort put in the preparation when participating can then be seen as a loss. More formally, denote z_{rjtk} and x_{rjtk} as the binary variables that indicate whether in round r agent k has participated in bidding on job $j_{i,r}$ or is assigned $j_{i,r}$ in the allocation π_r , respectively. For ease of notation we use round r directly in the subscript of the variables. Denote $\mathcal{K}_r^+ \subseteq \mathcal{K}$ as the subset of agents where $k^+ \in \mathcal{K}_r^+$ has $\sum_{j,t} z_{rjtk^+} > 0$ in round r . For every agent k that has participated in the round prior to round r , $k^+ \in \mathcal{K}_{r-1}^+$, we can use the proportion of number of jobs won in round $r-1$ over the number of bids submitted in round $r-1$, ψ_{kr} , as a measurement of the possible gain and loss of an agent $k \in \mathcal{K}$ in round r , see Eq. (2). We set the reference point to be the average proportion over all companies $k' \in \mathcal{K}$, and we normalize the difference between the proportion of agent k and the average proportion, as in Eq. (3). For the evaluation phase, we use the value function in Eq. (4) to obtain the probability of agent k bidding in round r . In Eq. (5) the prospect probability $\hat{p}(\psi'_{kr})$ is scaled into the probability interval $[0, 1]$.

$$\psi_{kr} = \frac{\sum_{j,t} x_{(r-1)jtk}}{\sum_{j,t} z_{(r-1)jtk}} \quad (2)$$

$$\psi'_{kr} = \frac{\psi_{kr} - \frac{1}{K} \sum_{k' \in \mathcal{K}} \psi_{k'r}}{\max_{k^+ \in \mathcal{K}_{r-1}^+} (\psi_{k^+r} - \frac{1}{K} \sum_{k' \in \mathcal{K}} \psi_{k'r})} \quad (3)$$

$$\hat{p}(\psi'_{kr}) = \begin{cases} \psi'_{kr}^\alpha & \text{if } \psi'_{kr} \geq 0 \\ -\lambda(-\psi'_{kr})^\beta & \text{if } \psi'_{kr} < 0 \end{cases} \quad (4)$$

$$p'(\psi'_{kr}) = \max(p_{LB}, \frac{\hat{p}(\psi'_{kr})}{2} + 0.5) \quad (5)$$

When an agent's proportion is equal to the average proportion, we assume that every agent is indifferent on participating, therefore, $p'(\psi'_{kr}) = 0.5$. When an agent is allocated more tasks than others on average, then the participation probability increases, as the agent will feel more certain in participating. However, when an agent is allocated less tasks than others on average, the participation probability decreases. In order for the probability to not become negative we set a lower bound p_{LB} , so that agents will still be able to come back and participate in later rounds although with a very small probability. If an agent has not participated in bidding in the previous round, we set $p'(\psi'_{kr}) = p_{LB}$. In order to take into account the experience from previous rounds into the participation probability, we will apply simple exponential smoothing. Thus, the participation prob-

ability for agent k in round r is calculated as

$$p(\psi'_{kr}) = \gamma p'(\psi'_{kr}) + (1 - \gamma)p(\psi'_{k(r-1)}), \quad (6)$$


where γ is the smoothing factor. In the first round, $r = 1$, we assume that all agents will participate in bidding.

Example 2: In Example 1 we have seen three agents who bid on five jobs. Consider this as round 1 with participation probabilities $p(\psi'_{k_i1}) = 1.00$ for all three agents. We choose $\alpha = \beta = 0.88$ and $\lambda = 2$, as they are commonly adopted in prospect theory, and $p_{LB} = 0.01$ and $\gamma = 0.5$. In round 2, we obtain $\psi_{k_i2} = (\frac{1}{1}, \frac{1}{3}, \frac{2}{8})$, for agent k_1 , k_2 and k_3 , respectively. The reference point becomes 0.57, and the corresponding differences are $(0.43, -0.24, -0.19)$. The normalized proportions become $\psi'_{k_i2} = (1, -0.55, -0.45)$. Using (4) and (5), we obtain $p'(\psi'_{k_i2}) = (1.00, 0.01, 0.01)$. Eventually, the participation probability can be obtained through (6), resulting in $p(\psi'_{k_i2}) = (1.00, 0.505, 0.505)$. Agents k_1 and k_3 (despite a low $p(\psi')$) decide to participate in this round, bidding on 3 and 5 jobs, and are assigned 1 and 2 jobs, respectively. Using this information, we can repeat the calculations for round 3, which result in $p(\psi'_{k_i3}) = (0.81, 0.01, 1.00)$. Agent k_2 will be very unlikely to participate in the next round of allocation. ■

IV. EXPERIMENTS


We are interested in how social welfare develops over multiple rounds and how two different allocation algorithms: min-cost max-flow algorithm and fair algorithm ([2]), influence it. We will use the same test instances as described in [2]. We will use $T = 10$ time periods per round. we assume the tasks to be uniform with each having a value of $V = 100$. The tasks have a latest completion time, which we will set to 3 time periods after the earliest time the task become available. Tasks have a 25% chance of starting at t_2 and another 25% chance of starting at t_6 . If a task does not start at a peak hour, it has an equal chance to start at any time from t_1 to t_8 . A bidder k has a predetermined set of tasks she is interested in. In the simulation, each task in each time period has a chance of either 0.25 or 0.75 to be selected for this set, which are indicated as the *low* and *high competition* case, respectively. For each task in this set she has a bid. We will distinguish between two bid cases. The first bid case is where every agent draws their bid cost from an uniform distribution in the interval $[30, 60]$, based on the hourly wage of a driver and the fuel costs with an additional profit, which we call the *homogeneous cost* case. In the second bid case a predetermined half of the agents will draw their bid cost from an uniform distribution in the interval $[30, 50]$, and the other half of the agents will draw theirs from the interval $[40, 60]$. We call this the *heterogeneous cost* case. Combining scenarios with low/high competition and homogeneous/heterogeneous cost, we construct four scenarios: *L/hom*, *L/het*, *H/hom* and *H/het*. We set the parameters for the prospect function to

Average number of participants per round with low competition




(a) Low competition

Average social welfare per round for low competition




(a) Low competition

Average number of participants per round with high competition



(b) High competition

Average social welfare per round with high competition



(b) High competition

Figure 1. Average number of participants per round over 20 experiments over 50 rounds.

$\alpha = \beta = 0.88$ and $\lambda = 2$ ([3]), and we use $p_{LB} = 0.01$ and $\gamma = 0.5$. We conduct 20 experiments for each of the four scenarios, with 50 rounds, 50 agents and 250 jobs.

Figure 1 shows the average number of participants per round over the 20 experiments. In the low competition case the average number of participants does not differ much between the different cost cases and allocation algorithms over the rounds. This is due to a limited number of bids, resulting in similar allocations regardless of allocation algorithm. For the high competition case, however, the average number of participants when using the fair algorithm is substantially higher than when using the minimum-cost algorithm as the rounds progress. This is due to the jobs being allocated more evenly among agents, which results in higher participation probabilities over all rounds. The average social welfare over the 20 experiments is not very different between the different cases with low competition, as seen in Figure 2. In the high competition case, social welfare is substantially higher for the fair algorithm as rounds proceed, and even looks to be stagnant. This is due to the larger number of participants still present, which enables more bids for the algorithm to choose from. The stagnancy stems from the nature of the problem, where it is the maximum social welfare for the problem setting. For the *H/het* scenario, the fair allocation obtains a lower social welfare in the earlier rounds compared to the minimum-cost allocation, but eventually surpasses it, due to the larger number of participants still present.

Figure 2. Average social welfare per round over 20 experiments over 50 rounds.

V. CONCLUSION AND DISCUSSION

In a repeated task or resource allocation problem with sharing nature, the participation of agents is driven not only by their expected economic gain, but also by their willingness to put effort in participating in the face of uncertainty and risk. The main contribution of this paper is that we model agents' participation by prospect theory. In addition, we demonstrate when there are plenty of resources in the system, the fair allocations result in more participants than the minimum-cost allocations throughout the rounds, which eventually results in a higher social welfare.

REFERENCES

- [1] M. M. Weerdt, Y. Zhang, and T. Klos, "Multiagent task allocation in social networks," *Autonomous Agents and Multi-Agent Systems*, vol. 25, no. 1, pp. 46–86, Jul. 2012.
- [2] Q. C. Ye, Y. Zhang, and R. Dekker, "Fair task allocation in transportation," *Omega*, 2016. [Online]. Available: <http://dx.doi.org/10.1016/j.omega.2016.05.005>
- [3] A. Tversky and D. Kahneman, "Advances in prospect theory: Cumulative representation of uncertainty," *Journal of Risk and uncertainty*, vol. 5, no. 4, pp. 297–323, 1992.
- [4] C. F. Camerer, "Prospect theory in the wild: Evidence from the field," *Advances in behavioral economics*, pp. 148–161, 2004.
- [5] R. Lahdelma and P. Salminen, "Prospect theory and stochastic multicriteria acceptability analysis (smaa)," *Omega*, vol. 37, no. 5, pp. 961–971, 2009.