On Balancing Fairness and Efficiency of Task Assignment in Agent Societies

Kendal Cousy[1,2], Marin Lujak[1(✉)], Alessio Salvatore[3], Alberto Fernández[1], and Stefano Giordani[2]

[1] CETINIA, University Rey Juan Carlos, Madrid, Spain
{kendal.cousy,marin.lujak,alberto.fernandez}@urjc.es
[2] Dipartimento di Ingegneria dell'Impresa, University of Rome "Tor Vergata", Rome, Italy
stefano.giordani@uniroma2.it
[3] Istituto per le Applicazioni del Calcolo "M. Picone", Consiglio Nazionale delle Ricerche, Rome, Italy
salvatore@iac.cnr.it

Abstract. Agent societies generally aim at collective provision of services (capabilities or resources) in a more efficient way than their agents could individually. In particular, some agents may be more efficient than the others in providing certain tasks. Thus, a task-agent assignment decision determines the overall performance of the society. The conventional linear sum assignment problem handles the assignment of tasks to a society of agents in a one-on-one manner. Such assignments typically only consider efficiency in terms of the overall cost or benefit for the system. However, an assignment strategy may be unfair if it does not explicitly consider fairness. Therefore, the conventional mathematical models for the task assignment problem should be rethought to explicitly consider fairness in the allocation of the tasks to the agents. In this paper, we study the utilitarian, egalitarian, and Nash social welfare in task assignment and propose two new assignment models that balance efficiency and fairness. Since fairness is a relatively abstract term that can be difficult to quantify, we propose three new fairness measures based on equity and equality and use them to compare the newly proposed models. Through functional examples, we show that a reasonable trade off between efficiency and fairness in task assignment can be found through the use of the proposed models.

Keywords: Task assignment · Multi-agent systems · Fairness · Efficiency · Resource allocation · Multi-agent coordination

1 Introduction

In this work, we focus on societies formed by self-concerned individually rational agents that share a common goal but have their own individual, possibly conflicting interests. They may share their capabilities and/or resources to carry out

A. González-Briones et al. (Eds.): PAAMS Workshops 2022, CCIS 1678, pp. 95–107, 2022.
https://doi.org/10.1007/978-3-031-18697-4_8

given tasks in a more efficient way and thus create synergies. Therefore, agents are expected to obtain higher performance by collective action. Examples of such societies are agriculture cooperatives, taxi, ride sharing, and hot meal delivery platforms.

In particular, we study linear sum task assignment problem where a set of tasks needs to be assigned to a set of agents in a one-on-one manner. We assume a centralised decision making process (or algorithm) that is in charge of deciding which agent is assigned to each task. An optimal solution, from a utilitarian point of view, would be the assignment that produces the lowest overall cost (or the highest benefit). However, this globally most efficient solution for the whole system may create large differences among the individual assigned costs of the participating agents (we refer to this as an "unfair" assignment). The perception of an unfair task assignment solution may motivate unsatisfied agents to leave the society, putting the survivability of the society at risk. Thus, assignment decisions in such societies should be made not only based on minimising overall assignment cost but should also consider social welfare and fairness.

The classical linear sum assignment problem is a largely studied, generally computationally easy problem, for which exact solutions can be produced relatively rapidly for even very large instances. However, to the best of our knowledge, related work on balancing fairness and efficiency in task assignment is scarce. Therefore, in this work, we explore the means of balancing the overall cost and fairness in task assignment in agent societies. These two aspects are generally opposed, i.e., solution approaches focusing on cost minimisation are likely to produce unfair assignments for some agents, while fair assignments may be far from the minimum cost solution. In this paper, we study trade-off between these two requirements and focus on finding task assignment solutions that are as fair as possible while not overly penalising the overall system's cost. This implies finding efficient and fair assignments considering the distribution of individual costs among agents.

The main contributions of this paper are twofold. First, we propose three new fairness measures for a multi-agent system composed of self-concerned individually rational agents: Egalitarian Fairness Measure (EFM), Relative All-to-all Fairness (RAF), and Overall Relative Opportunity Cost Fairness (OROCF) measure. Then, we present two new one-on-one task assignment models that maximise social welfare of the system while balancing efficiency and fairness: envy-free utilitarian model that uses the utilitarian social welfare function while constraining the differences of the costs between agents and the Nash model that optimises the Nash product of assigned tasks' benefits or costs of individual agents composing the system. We choose Nash social welfare due to its structure (being product of costs) that explicitly balances efficiency and fairness.

The rest of the paper is organised as follows. In Sect. 2, we give an overview of the state of the art. In Sect. 3, we give motivation for this work and define the general problem of one-on-one task assignment. We propose new equality and equity fairness measures in Sect. 4. The two new mathematical models for efficient and fair task assignment are presented in Sect. 5. Section 6 presents simple functional tests and discusses how the presented models differ based on the proposed fairness measures. In Sect. 7, we conclude the paper by giving an overview of the results and discuss the potential of the new proposed models and fairness measures to make a fairer task assignment. We also give lines of future work to improve the current models and fairness measures.

2 State of the Art

The assignment of resources, capabilities or tasks in a multi-agent society may vary when defining fairness and efficiency depending on mutual inter-dependencies among agents and their relation with the society's objectives (e.g., [4,7,21]).

Collaborative decision making considers a goal that is shared and owned among all agents in a society, while cooperative decision making considers working toward a shared goal even though its ownership is not shared [25]. Thus, cooperative decision making results are generally differentially beneficial to different agents [22], while collaboration is generally about equally sharing efforts, costs and benefits. Collaborative multi-agent task allocation problem is studied in, e.g., [16,19]. This problem has many different real world applications where fairness can be a challenge. For example, in Spatial Crowdsourcing [32], there is a need to minimise the payoff difference among workers while maximising the average worker payoff. Similarly, in Rideshare Platforms, it was shown in [24] that, during high-demand hours, lacking any consideration of fairness and seeking only an optimal number of trips could lead to increased societal biases in the choice of the clients. This problem is relevant for many other applications including manufacturing and scheduling, network routing and the fair and efficient exploitation of Earth Observation Satellites (e.g., [7]).

There exist in the literature many kinds of fairness measures for different contexts, e.g., machine learning (e.g., [10]), neural networks (e.g., [26]) and algorithm development (e.g., [13]). Of our interest are the fairness measures for the allocation of indivisible goods (e.g., [8]), and in more specific, the fairness measures for one-on-one assignment of tasks in collaborative or cooperative multi-agent systems. The most known fairness measures for allocation of indivisible goods are the maxmin of the utility of the agents in the system which maximises the utility of the agent that contributes the least to the global utility of the system; there is also proportionality which states that each agent should receive at least one n^{th} of the utility this agent would have received if it were alone. Max-min fairness is generalised in the case of resource allocation for systems with different resource types in [12] while max-min fairness, proportional fairness and balanced fairness are compared in the setting of a communication network of processor-sharing queues in [3].

The concept of fairness is studied as well in other contexts, like the multi-winner voting problem, machine learning and in recommender systems (e.g., [31]), but also more generally in decision making (e.g., [28]). The importance of the individual perception of fairness within a system in order to keep individual satisfaction high is emphasised in [29]. The potential contradiction between individual fairness and group fairness is studied in [2] in the Machine Learning context. Some works study more generally the concepts of distributive justice, equality and equity (e.g., [9]).

Related to the balance of efficiency and fairness are different social welfare concepts. Their modelling and importance in enhancing the quality of task allocation are studied in [7]. In this work, we study egalitarianism and utilitarianism in this regard. Egalitarianism is a trend of thought in political philosophy that favours equality among the individuals composing the society no matter what their circumstances are (e.g., [11]). Utilitarianism, on the other hand, is a theory of morality that advocates actions that maximise happiness or well-being for all individuals while opposing to the actions that cause their unhappiness or harm. When directed toward making social and economic decisions, a utilitarian philosophy aims at the improvement of the society as a whole (e.g., [23]).

The Nash social welfare combines efficiency and fairness considerations. This function, or variants of it, are studied in literature considering, e.g., fairness in the ambulance location problem [14], and in allocating indivisible goods [6]. The multiagent resource allocation problem considering Nash social welfare (the product of the utilities of the individual agents) is studied in [27].

In resource allocation, there can be agents desiring tasks (resources) more than others, or there can even be agents desiring tasks given to other agents, creating envy in the system (e.g., [7]). Envy-freeness criterion implies that an allocation should leave no agent envious of the other (e.g., [5]). However, it is not always enough to achieve envy-freeness for a fair solution (e.g., [1,15,20]).

3 Motivation and Problem Definition

Most of the State of the art literature on task assignment generally focuses on the efficiency of the assignment and does not consider fairness in the process, thus optimising only the system's overall general assignment cost or profit (e.g., time, distance, monetary value, etc.). This strategy is equivalent to optimising utilitarian social welfare function, a concept from welfare economy that sums the utility of each individual in order to obtain society's overall welfare (see, e.g., [7,30]). All agents are treated the same, regardless of their initial level of utility or cost distribution among the tasks. This strategy is admissible in case of a single decision maker, but might be unacceptable when multiple self-concerned and individually rational agents must decide on the assignment of tasks.

Let us introduce a simple example showing how unfair a task assignment optimising utilitarian social welfare function can be. Let us consider 3 self-concerned, individually rational agents (a_1, a_2, a_3) that need to be assigned to a set of 3 tasks (k_1, k_2, k_3) in a one-on-one manner and vice versa. The cost matrix containing the assignment costs for these agents and tasks is shown in Table 1a.

By applying the conventional (linear sum) task assignment model (i.e., the Utilitarian model) that optimises the overall cost of the system without considering fairness in the assignment (see, e.g., [17,19]), we might get the assignments (called solution s_1) marked in bold in Table 1b. The overall minimum assignment

Table 1. Example of a cost matrix and different one-on-one task assignment solutions with minimum overall cost (in bold)

(a) Cost Matrix

	k_1	k_2	k_3
a_1	50	60	70
a_2	30	40	50
a_3	10	50	30

(b) Solution s_1

	k_1	k_2	k_3
a_1	50	**60**	70
a_2	**30**	40	50
a_3	10	50	**30**

(c) Solution s_2

	k_1	k_2	k_3
a_1	**50**	60	70
a_2	30	**40**	50
a_3	10	50	**30**

(d) Solution s_3

	k_1	k_2	k_3
a_1	50	60	**70**
a_2	30	**40**	50
a_3	**10**	50	30

cost found by this model is 120. However, if we focus on its cost distribution on individual agents, we see large discrepancies. Indeed, the cost of agent a_1 is 60, while the cost of a_2 is only 30. Thus, a_1 is charged twice more than a_2. In Table 1d (i.e., solution s_3), this difference is even larger resulting in 7 times larger cost of the worst-off in respect to the best-off agent. Generally, an upper bound on the difference in the assignment cost is the maximum value in a given cost matrix. In centralised systems, where agents are owned and controlled by a single decision maker, this would not cause any problem. However, in the case of decentralised systems composed of self-concerned and individually rational agents, such an unfair solution might result in the worst-off agents leaving the system due to the lack of fairness in the solution.

Table 1c shows a fairer solution (called s_2) where the costs of the agents are as close as possible, thus minimising the envy of agents. This is an ideal situation in regard to fairness in this case where all agents are assigned tasks of similar costs. Notice that, in this case, we didn't have to sacrifice efficiency to achieve this situation. In case of repetitive task allocations, the assignments can be altered to further facilitate balance in the accumulated assignment costs.

Problem Definition. Given are a set of agents $a \in A$ and a set of tasks $k \in K$ that form a weighted complete bipartite graph $G = (A \bigcup K, E)$ with edge set $E = A \times K$ and with given edge weights c_{ak} on each edge $(a, k) \in E$, where c_{ak} is the cost of assigning task $k \in K$ to agent $a \in A$, $\forall a \in A, k \in K$. W.l.o.g, we assume that the cardinality of the sets is equal, i.e., $|A| = |K|$. In the case of unequal cardinality, we add a sufficient number of dummy vertices and assume that $c_{ak} = 0$ where $a \in A$ or $k \in K$ are dummy vertices. The objective is to assign agents $a \in A$ to tasks $k \in K$ in a one-on-one manner and, therefore, find a perfect matching among vertices in A and vertices in K considering both assignment efficiency and fairness. An edge (a, k) is matched if its (two) extreme vertices $a \in A$ and $k \in K$ are mutually matched, and a matching is perfect if every vertex of A is matched (assigned) exactly to one vertex of K, and vice versa. The following is the mathematical formulation of these constraints.

$$\sum_{k \in K} x_{ak} = 1, \forall a \in A \quad (1) \qquad \sum_{a \in A} x_{ak} = 1, \forall k \in K \quad (2)$$

$$x_{ak} \in \{0,1\}, \forall a \in A, \forall k \in K \tag{3}$$

where x_{ak} is a binary decision variable equal to 1 if agent $a \in A$ is assigned to task $k \in K$, and zero otherwise. Constraints (1) and (2) assure that there is one-on-one assignment for each agent $a \in A$ and task $k \in K$, respectively. Constraints (3) fix the ranges of the decision variables.

4 Proposed Fairness Measures

In this section, we introduce different fairness measures for quantifying the balance between fairness and efficiency in task assignment from the egalitarian and equity point of view. All the fairness measures are fractions ranging between 0 and 1. We avoided the division by 0 in some extreme cases by adding a very small number ϵ (e.g., $\epsilon = 1e^{-10}$) to both the numerator and the denominator of these fractions.

Egalitarian Fairness Measure. (EFM) focuses on the assignment cost faced by the worst-off agent (i.e., the agent with the highest assignment cost in a given feasible solution). Given the assignments x_{ak}^{sol}, with $a \in A$ and $k \in K$, of a feasible solution *sol*, EFM is computed as follows:

$$EFM(sol) = \frac{c_{max} - c_{sol}^{wo} + \epsilon}{c_{max} - c_{min}^{wo} + \epsilon} \tag{4}$$

where $c_{max} = max_{a \in A, k \in K}\{c_{ak}\}$ is the maximum value in the cost matrix, $c_{sol}^{wo} = \max_{a \in A}\{\sum_{k \in K} c_{ak} x_{ak}^{sol}\}$ is the cost paid by the worst-off agent in the given solution, and c_{min}^{wo} is the minimum cost that the worst-off agent could pay. In particular, c_{min}^{wo} is the optimal solution of the given mathematical problem:

$$c_{min}^{wo} = \min \lambda \tag{5}$$

s.t. (1)–(3) and

$$\sum_{k \in K} c_{ak} x_{ak} \le \lambda, \forall a \in A \tag{6}$$

$$\lambda \in \Re \tag{7}$$

where Constraints (6) impose that the cost (λ) paid by the worst-off agent must be not less than the cost paid by any agent, and Constraints (7) fix the range of the additional variable λ. When the worst-off assigned cost c_{sol}^{wo} is equal to c_{max}, $EFM(sol)$ will equal zero (ignoring ϵ). On the other hand, when c_{sol}^{wo} is equal to c_{min}^{wo}, $EFM(sol)$ will equal one; moreover, this also occurs when there exists an agent $a \in A$ such that $c_{ak} = c_{max}, \forall k \in K$.

For the cost matrix given in Table 1a, where $c_{max} = 70$ and $c_{min}^{wo} = 50$, we calculate the $EFM(sol)$ for each solution reported in Tables 1b–1d. All the solutions reported in Table 1, have minimum overall assignment cost equal to

120, while the values of c^{wo}_{sol} are $c^{wo}_{s_1} = 60$, $c^{wo}_{s_2} = 50$, $c^{wo}_{s_3} = 70$ for the solutions reported in Table 1b, Table 1c, and Table 1d, respectively. $EFM(sol)$ value for these solutions are: $EFM(s_1) = \frac{70-60}{70-50} = 0.5$, $EFM(s_2) = \frac{70-50}{70-50} = 1$, and $EFM(s_3) = \frac{70-70}{70-50} = 0$.

According to EFM measure, solution s_2 is the fairest one. Note that the increase in EFM value in solution s_2 corresponds to a distribution of the costs that leaves the worst-off agent better off than in s_1, and that solution s_3 leaves the worst-off agent with the worst possible cost. Note that, generally, there may be multiple such distributions.

Relative All-to-All Fairness. (RAF) evaluates fairness at a group level by taking into account each agent's perspective in comparison with the others. The measure is based on the sum between the squared differences of the assignment costs of each agent and the costs of the others, as seen in Eq. (8).

$$w_{sol} = \sum_{a \in A} \sum_{a' \in A | a' > a} (\sum_{k \in K} c_{ak} x^{sol}_{ak} - c_{a'k} x^{sol}_{a'k})^2, \tag{8}$$

Then, relative all-to-all fairness is computed as follows:

$$RAF(sol) = \frac{w_{max} - w_{sol} + \epsilon}{w_{max} - w_{min} + \epsilon}, \tag{9}$$

where w_{max} and w_{min} represent the maximum and the minimum value of Eq. (8) given Constraints (1)–(3).

For the cost matrix given in Table 1a, the two components of RAF that are independent of the assignment solution are $w_{max} = 5400$ and $w_{min} = 0$, related to solutions s_{max}, with $x^{s_{max}}_{13} = x^{s_{max}}_{22} = x^{s_{max}}_{31} = 1$, and s_{min}, with $x^{s_{min}}_{11} = x^{s_{min}}_{23} = x^{s_{min}}_{32} = 1$, respectively. The values w_{sol} for the solutions reported in Table 1b, Table 1c and 1d are $w_{s_1} = 1800$, $w_{s_2} = 600$, and $w_{s_3} = 5400$, respectively. Related RAF values are: $RAF(s_1) = \frac{5400-1800}{5400-0} = 0.67$, $RAF(s_2) = \frac{5400-600}{5400-0} = 0.89$, and $RAF(s_3) = \frac{5400-5400}{5400-0} = 0$.

Also according to the RAF measure, solution s_2 is the fairest one and the order of the three solutions is the same as for EFM. This is not surprising as both measures evaluate equality in a solution. However, s_2 is not the absolute fairest solution which, with respect to this indicator, is $x_{11} = x_{23} = x_{32} = 1$ where all the agents pay the same cost; in this case the RAF value is equal to 1. This is also not surprising as this particular measure considers not only the worst-off agent, but all of them, therefore making it less likely that one of the solutions with minimum cost also has the highest fairness value.

Overall Relative Opportunity Cost Fairness. (OROCF) focuses on achieving equity among the agents by taking into account the missed opportunities in terms of the assignment cost for each agent. The opportunity cost (e.g., [18]) is the concept in microeconomics of lost benefit that would have been derived by an agent from an option not chosen. As the reference value, we consider a task of the minimum cost and normalise the difference in the cost value between the assigned task and the best-off task (the task with minimum cost) over the amplitude of costs for each agent, as seen in Eq. (10).

$$y_{sol} = \sum_{a \in A} \frac{\sum_{k \in K} c_{ak} x_{ak}^{sol} - min_{k \in K} \{c_{ak}\} + \epsilon}{max_{k \in K} \{c_{ak}\} - min_{k \in K} \{c_{ak}\} + \epsilon} \tag{10}$$

$$OROCF(sol) = \frac{y_{max} - y_{sol} + \epsilon}{y_{max} - y_{min} + \epsilon}, \tag{11}$$

where y_{max}, y_{min} represents the maximum and the minimum value of Eq. (10) given Constraints (1)–(3).

The values y_{sol} for the solutions reported in Tables 1b, 1c and 1d are $y_{s_1} = 1$, $y_{s_2} = 1$, and $y_{s_3} = 1.5$, respectively. For the cost matrix given in Table 1a, the two components of OROCF that are independent of the assignment solution are $y_{max} = 2$ for $x_{13}^{s_{max}} = x_{21}^{s_{max}} = x_{32}^{s_{max}} = 1$ and $y_{min} = 1$ for $x_{11}^{s_{min}} = x_{22}^{s_{min}} = x_{33}^{s_{min}} = 1$. Related $OROCF$ values are: $OROCF(s_1) = \frac{2-1}{2-1} = 1$, $OROCF(s_2) = \frac{2-1}{2-1} = 1$, and $OROCF(s_3) = \frac{2-1.5}{2-1} = 0.5$.

Note that $OROCF$ value is the highest both for s_1 and s_2, meaning that these solutions offer the lowest highest opportunity cost for the sum of all agents. The reader can verify that the solution $x_{11} = x_{23} = x_{32} = 1$ would be the worst choice for agents a_2 and a_3 and would give a value of $OROCF$ equal to 0.

5 Proposed Models Considering Fairness and Efficiency

In this section, we propose new models that mitigate the equity issues posed by the classical linear sum assignment model (e.g., [4]) and achieve a solution that is as fair as possible while sacrificing as little as possible of the overall system's efficiency.

Nash Model. The proposed Nash Model is inspired by the Nash social welfare function, a well studied social welfare function in which the goal is to maximize the product of the utility functions of the agents composing the system. The proposed model is given next:

$$min \prod_{a \in A} \sum_{k \in K} c_{ak} x_{ak} \tag{12}$$

s.t. (1)–(3). Since Eq. (12) is a nonlinear objective function, solving the above problem is computationally expensive. Thus, we propose next its linearised equivalent, which is possible due to the one-on-one assignment constraints (1)–(3).

$$max \sum_{a \in A} \sum_{k \in K} log(M - c_{ak}) x_{ak} \tag{13}$$

s.t. (1)–(3), where $M > max_{k \in K, a \in A} \{c_{ak}\}$.

Envy-Free Utilitarian Model. This model focuses both on efficiency and fairness. We introduce the fairness variable f_u to ensure that all the costs for each agent are inside a certain interval that shrinks as f_u becomes smaller. The model is defined as follows:

$$\min\ \alpha f_u + (1-\alpha)\frac{\sum_{k\in K}\sum_{a\in A} c_{ak}x_{ak}}{|A|} \tag{14}$$

s.t. (1)–(3), and

$$\sum_{k\in K} c_{ak}x_{ak} - \frac{\sum_{k\in K}\sum_{a\in A} c_{ak}x_{ak}}{|A|} \leq f_u, \forall a \in A \tag{15}$$

$$f_u \in \Re \tag{16}$$

Constraints (15) guarantee that, for each agent, the difference between the cost of its assigned task and the average of the costs of the assigned tasks for all the agents is less than the value f_u. Constraint (16) fixes the range of variable f_u. Parameter α ranges between 0 and 1 and is used to weigh the fairness and the average cost paid by an agent in (14).

6 Functional Tests

To demonstrate the difference between the Nash model and the Envy-free Utilitarian model, we randomly generated three cost matrices (Table 2) with costs ranging from 1 to 1000. The models were solved for each matrix using IBM ILOG CPLEX Optimization Studio 20.0.1.

Table 2. Example cost matrices.

(a) Functional test 1

	k_1	k_2	k_3
a_1	382	816	366
a_2	846	544	175
a_3	578	824	526

(b) Functional test 2

	k_1	k_2	k_3
a_1	450	895	358
a_2	856	233	449
a_3	890	672	976

(c) Functional test 3

	k_1	k_2	k_3
a_1	683	170	699
a_2	943	364	894
a_3	557	741	127

To compare the efficiency of the solutions obtained using the models presented in Sect. 5, we calculate the following normalised efficiency indicator (Eff):

$$Eff(sol) = \frac{z_{max} - z_{sol} + \epsilon}{z_{max} - z_{min} + \epsilon} \tag{17}$$

where $z_{sol} = \sum_{k\in K}\sum_{a\in A} c_{ak}x_{ak}^{sol}$ with x_{ak}^{sol} being the solution returned by the considered model. The values z_{max} and z_{min} are, respectively, the maximum and the minimum values of $\sum_{k\in K}\sum_{a\in A} c_{ak}x_{ak}$ given Constraints (1)–(3).

Table 3. Results and comparison

Functional test 1				
Model	Eff	EFM	RAF	$OROCF$
Nash	0.91	1	1	1
Envy-free ($\alpha = 0$)	1	0.07	0	0.68
Envy-free ($\alpha \geq 0.5$)	0.91	1	1	1
Functional test 2				
Model	Eff	EFM	RAF	$OROCF$
Nash	0.93	1	0.91	1
Envy-free ($\alpha = 0$)	1	0.28	0.17	0.92
Envy-free ($\alpha = 0.5$)	0.93	1	0.91	1
Envy-free ($\alpha \geq 0.9$)	0	0	1	0
Functional test 3				
Model	Eff	EFM	RAF	$OROCF$
Nash	0.93	1	0.73	1
Envy-free ($\alpha = 0$)	1	0	0	0.99
Envy-free ($\alpha = 0.5$)	0.93	1	0.73	1
Envy-free ($\alpha = 0.7$)	0.59	0.94	0.93	0.64
Envy-free ($\alpha \geq 0.9$)	0.05	0.19	1	0.06

Table 3 shows the results of our experiments with all indicators and their values depending on the model used.

The case when $\alpha = 0$ corresponds to the case when we are optimising the global cost only (utilitarian social welfare function). We get very low values of fairness for this case according to our prior assumptions. It is interesting to notice similarities when we set α value to 0.5. Indeed, in that case, the Envy-free Utilitarian model and the Nash model have the same behaviour and give us the same solutions. These solutions for $\alpha = 0.5$ are ideal for the fairness indicators EFM and $OROCF$ in our three tests, while RAF also increases significantly. Moreover, the efficiency (Eff) is greater than 0.9. Equality and equity can be improved without significant decrease in efficiency. We notice in tests 2 and 3 that, generally, the higher the value of α, the lower the overall system's efficiency. This shows that striving for too much equality can be highly detrimental to the system's efficiency and even equity. The results for the cost matrix in Table 1a also support this claim in case $\alpha = 1$. Here, allocation $x_{11} = x_{23} = x_{32} = 1$ is an egalitarian allocation that decreases efficiency and equity simultaneously since agents a_2 and a_3 are allocated to their worst-off tasks and the overall allocation cost is 150 instead of the minimum cost of 120.

7 Conclusions

In this paper, we focused on balancing efficiency and fairness in one-on-one multi-agent task assignment. This problem is of high importance in agent societies composed of individually rational and self-concerned agents where an agent decides to participate only if it brings an individual benefit that is at least as good as when not participating. In this regard, we studied the utilitarian, egalitarian and Nash social welfare, the concepts from economics and philosophy that may be applied in such multi-agent societies to tackle this challenge. Since quantitative fairness measures for task assignment are scarce or missing, we proposed three new fairness measures: egalitarian fairness measure, all-to-all relative fairness measure, and overall relative opportunity cost fairness measure. We proposed the Nash model for task assignment that minimises the product of the costs of each agent, considering one-on-one assignment constraints, and the Envy-Free Utilitarian model which is a model combining the ideas of envy-freeness, equality and the utilitarian social welfare measure. We concluded with 3 functional tests showing that by using our proposed two models, we can achieve a better fairness with little sacrifice in the overall efficiency, and that our Envy-free Utilitarian model can be adjusted depending on the need for fairness.

The fairness measures presented should be computed only for non-dummy vertices to ensure that these measures can still reach either the value of 0 or 1 in practice. In the future, we will further study fairness measures, particularly one encompassing both equality and equity to better support decision-making in collaborative and cooperative open societies where agents can enter and leave as they wish. Moreover, we will focus on three-index assignment problem where each agent needs a tool to perform a task. The assignment here is also performed in a one-on-one manner. Similarly, crafting a multi-objective model which considers equality, equity and fairness for such a problem is a challenge worth facing henceforth.

Acknowledgements. This work has been partially supported by the Spanish Ministry of Science, Innovation, and Universities, co-funded by EU FEDER Funds, through grant number RTI2018-095390-B-C33 (MCIU/AEI/FEDER, UE) and the AGROBOTS Project funded by the Community of Madrid, Spain.

References

1. Alkan, A., Demange, G., Gale, D.: Fair allocation of indivisible goods and criteria of justice. Econometrica: J. Econ. Soc. 1023–1039 (1991)
2. Binns, R.: On the apparent conflict between individual and group fairness. In: Proceedings of the ACM FAccT 2020, pp. 514–524 (2020)
3. Bonald, T., Massoulié, L., Proutiere, A., Virtamo, J.: A queueing analysis of max-min fairness, proportional fairness and balanced fairness. QS **53**(1), 65–84 (2006)
4. Burkard, R., Dell'Amico, M., Martello, S.: Assignment problems: revised reprint. SIAM (2012)
5. Cappelen, A.W., Tungodden, B.: Fairness and the proportionality principle. Soc. Choice Welfare **49**(3), 709–719 (2016). https://doi.org/10.1007/s00355-016-1016-6

6. Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A.D., et al.: The unreasonable fairness of maximum Nash welfare. ACM-TEAC **7**(3), 1–32 (2019)
7. Chevaleyre, Y., Dunne, P.E., Endriss, U., et al.: Issues in multiagent resource allocation (2005)
8. Conitzer, V., Freeman, R., Shah, N., Vaughan, J.W.: Group fairness for the allocation of indivisible goods. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 33, pp. 1853–1860 (2019)
9. Cook, K.S., Hegtvedt, K.A.: Distributive justice, equity, and equality. Ann. Rev. Sociol. **9**(1), 217–241 (1983)
10. Corbett-Davies, S., Goel, S.: The measure and mismeasure of fairness: a critical review of fair machine learning. arXiv preprint arXiv:1808.00023 (2018)
11. Fleurbaey, M., Maniquet, F.: A Theory of Fairness and Social Welfare, vol. 48. University Press, Cambridge (2011)
12. Ghodsi, A., Zaharia, M., Hindman, B., et al.: Dominant resource fairness: fair allocation of multiple resource types. In: 8th USENIX Symposium NSDI 2011 (2011)
13. Hellman, D.: Measuring algorithmic fairness. VLR **106**(4), 811–866 (2020)
14. Jagtenberg, C., Mason, A.: Fairness in the ambulance location problem: maximizing the Bernoulli-Nash social welfare. SSRN 3536707 (2020)
15. Kranich, L.: Resource-envy-free and efficient allocations: a new solution for production economies with dedicated factors. J. Math. Econ. **89**, 1–7 (2020)
16. Kraus, S., Plotkin, T.: Algorithms of distributed task allocation for cooperative agents. Theoret. Comput. Sci. **242**(1–2), 1–27 (2000)
17. Kuhn, H.W.: The Hungarian method for the assignment problem. Naval Res. Logistics Q. **2**(1–2), 83–97 (1955)
18. Kurzban, R., Duckworth, A., Kable, J.W., Myers, J.: An opportunity cost model of subjective effort and task performance. Behav. Brain sci. **36**(6), 661–679 (2013)
19. Lujak, M., Giordani, S., Omicini, A., Ossowski, S.: Decentralizing coordination in open vehicle fleets for scalable and dynamic task allocation. Complexity **2020**(1047369) (2020)
20. Lujak, M., Giordani, S., Ossowski, S.: Route guidance: bridging system and user optimization in traffic assignment. Neurocomputing **151**, 449–460 (2015)
21. Moulin, H.: Fair Division and Collective Welfare. MIT Press, Cambridge (2004)
22. Moulin, H.: Axioms of Cooperative Decision Making, vol. 15. University Press, Cambridge (1991)
23. Mulgan, T.: Understanding Utilitarianism. Routledge, Milton Park (2014)
24. Nanda, V., Xu, P., et al.: Balancing the tradeoff between profit and fairness in rideshare platforms during high-demand hours. In: Proceedings of the AAAI Conference on Artificial Intelligence, vol. 34, pp. 2210–2217 (2020)
25. Nelson, L.M.: Collaborative problem solving. In: Instructional-Design Theories and Models: A New Paradigm of Instructional Theory, vol. 2, pp. 241–267 (1999)
26. Padala, M., Gujar, S.: FNNC: achieving fairness through neural networks. In: Proceedings of the Twenty-Ninth International Joint Conference on Artificial Intelligence, IJCAI 2020 (2020)
27. Ramezani, S., Endriss, U.: Nash social welfare in multiagent resource allocation. In: David, E., Gerding, E., Sarne, D., Shehory, O. (eds.) AMEC/TADA -2009. LNBIP, vol. 59, pp. 117–131. Springer, Heidelberg (2010). https://doi.org/10.1007/978-3-642-15117-0_9
28. Sampat, A.M., Zavala, V.M.: Fairness measures for decision-making and conflict resolution. Optim. Eng. **20**(4), 1249–1272 (2019). https://doi.org/10.1007/s11081-019-09452-3

29. Schappe, S.P.: Understanding employee job satisfaction: the importance of procedural and distributive justice. J. Bus. Psyc. **12**(4), 493–503 (1998)
30. Sen, A.: Collective Choice and Social Welfare. Harvard University Press, Cambridge (2017)
31. Shrestha, Y.R., Yang, Y.: Fairness in algorithmic decision-making: applications in multi-winner voting, machine learning, and recommender systems. Algorithms **12**(9), 199 (2019)
32. Zhao, Y., Zheng, K., Guo, J., Yang, B., Pedersen, T.B., Jensen, C.S.: Fairness-aware task assignment in spatial crowdsourcing: game-theoretic approaches. In: 2021 IEEE 37th ICDE, pp. 265–276. IEEE (2021)