

Equity by Design in Task Allocation: Reverse Auctions with Group and Individual Fairness

Extended Abstract

Yunxiao Zhao

University of Electronic Science and Technology of China
Chengdu, China
yxzhao@std.uestc.edu.cn

Jiamou Liu

The University of Auckland
Auckland, New Zealand
jiamou.liu@auckland.ac.nz

Mengxiao Zhang*

The University of Auckland
Auckland, New Zealand
mengxiao.zhang@auckland.ac.nz

Bakh Khoussainov

University of Electronic Science and Technology of China
Chengdu, China
bmk@uestc.edu.cn

ABSTRACT

Reverse auctions are widely used for budgeted task allocation, but efficiency-only rules can systematically disadvantage higher-cost communities. We ask whether equity can be built into reverse-auction task allocation without sacrificing the incentive guarantees. We introduce two complementary fairness notions: aggregate group fairness and within-group Lipschitz individual fairness. We propose GIFTA, a two-stage mechanism that learns group-selection probabilities and then allocates within the chosen group using a bid-decreasing randomized rule with truthful payments. We prove that GIFTA is truthful in expectation, individually rational, budget-feasible, satisfies the Lipschitz guarantee, and achieves asymptotic ϵ -group fairness; experiments show large disparity reductions with modest social-cost increases.

KEYWORDS

Reverse auction; group fairness; individual fairness

ACM Reference Format:

Yunxiao Zhao, Mengxiao Zhang, Jiamou Liu, and Bakh Khoussainov. 2026. Equity by Design in Task Allocation: Reverse Auctions with Group and Individual Fairness: Extended Abstract. In *Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026)*, Paphos, Cyprus, May 25 – 29, 2026, IFAAMAS, 3 pages. <https://doi.org/10.65109/CCDY4126>

1 INTRODUCTION

Reverse auctions are widely used for task allocation and can lower expected expenditure and shorten turnaround time by awarding work to the lowest bidders [15]. However, when bidders face structural cost differences, this rule can systematically allocate far fewer tasks to higher-cost (often minority) communities, undermining equity. More broadly, systems optimized for a narrow objective can exhibit persistent group disparities unless fairness is made

*Corresponding author



This work is licensed under a Creative Commons Attribution International 4.0 License.

Proc. of the 25th International Conference on Autonomous Agents and Multiagent Systems (AAMAS 2026), C. Amato, L. Dennis, V. Mascardi, J. Thangarajah (eds.), May 25 – 29, 2026, Paphos, Cyprus. © 2026 International Foundation for Autonomous Agents and Multiagent Systems (www.ifaamas.org). <https://doi.org/10.65109/CCDY4126>

an explicit design constraint, as documented across multiple high-stakes domains [3, 11, 14]. Policy frameworks have therefore shifted toward proactive safeguards, often framed as “Equity by design”.

We study reverse-auction allocation with socially salient groups and ask whether equity can be embedded into the allocation rule without sacrificing the core incentive guarantees. Fairness in auctions has been widely studied, including envy-freeness (EF) and its variants [1, 5, 8, 16], but EF does not preclude the systematic exclusion of higher-cost groups. Related fairness principles have also been developed in AI at both the individual and group levels [4, 6, 7, 10]. Motivated by these insights, we will introduce two complementary fairness notions and design a mechanism that enforces them while preserving incentive guarantees.

2 PROBLEM FORMULATION

We study task allocation with a single *requester* and n *workers*, indexed by $\mathcal{N} = [n]$. The requester seeks to complete a set of tasks $\mathcal{K} = [\kappa]$. Each task $k \in \mathcal{K}$ has a publicly known value $v_k > 0$ to the requester; write $\vec{v} = (v_1, \dots, v_\kappa)$. Each worker $i \in \mathcal{N}$ has a privately known *cost* $c_{i,k} \in \Theta_k := [\underline{c}_k, v_k]$, where Θ_k is the feasible cost domain bounded by \underline{c}_k and v_k . We collect costs as $\vec{c}_i := (c_{i,1}, \dots, c_{i,\kappa})^\top$ and $C := (\vec{c}_1, \dots, \vec{c}_n)$. Each worker $i \in \mathcal{N}$ reports a *bid* $b_{i,k} \in \Theta_k$ for every task k . Collect worker i 's bids into the *bid vector* $\vec{b}_i := (b_{i,1}, \dots, b_{i,\kappa})^\top$; the *bid matrix* is $B := (\vec{b}_1, \dots, \vec{b}_n)$. Bids may deviate from true costs if profitable. We assume each task is assigned to at most one worker (or left unassigned), and each worker may be assigned any number of tasks.

By the revelation principle [13], it is without loss of generality to analyze *direct mechanisms* that map reports to (possibly randomized) allocations and payments. An *allocation outcome* is a function $x: \mathcal{N} \rightarrow \{0, 1\}^\kappa$ where $x(i)_k = 1$ iff worker $i \in \mathcal{N}$ is assigned task $k \in \mathcal{K}$. We usually write $x_{i,k}$ for $x(i)_k$. A *payment outcome* is a function $p: \mathcal{N} \rightarrow \mathbb{R}_{\geq 0}^\kappa$, where $p_{i,k} := p(i)_k$ is the payment to worker i contingent on completing task k . Let \mathcal{X} denote the feasible set of allocations (e.g., $\sum_i x_{i,k} \leq 1$ for all k). Given bids B , the *utility* of worker i is $u_i(x, p) := \sum_{k \in \mathcal{K}} x_{i,k} (p_{i,k} - c_{i,k})$, and the *social cost* is the total true cost of assigned workers, $sc(x) := \sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{N}} c_{i,k} x_{i,k}$.

DEFINITION 1. In a randomized mechanism M , the *allocation outcome* is a probability distribution over the set of feasible outcomes

\mathcal{X} , and the payment outcome is a probability distribution over the set of all possible payment functions $p: \mathcal{N} \rightarrow \mathbb{R}_{\geq 0}^{\kappa}$.

Let (X, P) denote the random variables of the allocation and payment outcomes, resp., which are induced by M at bids \mathbf{B} . For worker-task pair (i, k) , $X_{i,k}(\mathbf{B})$ and $P_{i,k}(\mathbf{B})$ are thus the random variables with realizations $x_{i,k}$ and $p_{i,k}$, resp. The *expected allocation probability* and *expected payment* for (i, k) are $\mathbb{E}[X_{i,k}(\mathbf{B})] \in [0, 1]$ and $\mathbb{E}[P_{i,k}(\mathbf{B})] \in \mathbb{R}_{\geq 0}$, resp. Worker i 's *expected utility* is $\mathbb{E}[U_i(\mathbf{B})] = \sum_{k \in \mathcal{K}} (\mathbb{E}[P_{i,k}(\mathbf{B})] - \mathbb{E}[X_{i,k}(\mathbf{B})] c_{i,k})$, and the *expected social cost* is $\mathbb{E}[SC(\mathbf{B})] = \sum_{k \in \mathcal{K}} \sum_{i \in \mathcal{N}} c_{i,k} \mathbb{E}[X_{i,k}(\mathbf{B})]$.

We define the mechanism desiderata. Let $C_{-i} := (\bar{c}_j)_{j \in \mathcal{N} \setminus \{i\}}$. A mechanism M is

- *incentive compatible* (IC) if for all $i \in \mathcal{N}$, all $\bar{c}_i, \bar{b}_i \in \Theta^{\kappa}$, and for all $C_{-i} \in \Pi_{k \in \mathcal{K}} \Theta^{n-1}$, $\mathbb{E}[U_i(\bar{c}_i, C_{-i})] \geq \mathbb{E}[U_i(\bar{b}_i, C_{-i})]$;
- *individually rational* (IR) if for all $i \in \mathcal{N}$, and for all $C_{-i} \in \Pi_{k \in \mathcal{K}} \Theta^{n-1}$, $\mathbb{E}[U_i(\bar{c}_i, C_{-i})] \geq 0$;
- *budget feasible* (BF) if for all $k \in \mathcal{K}$, $\sum_{i \in \mathcal{N}} p_{i,k} \leq v_k$.

Fairness notions. The set of workers \mathcal{N} is partitioned into τ disjoint, observable groups $\mathcal{N}_1, \dots, \mathcal{N}_{\tau}$ with $\bigcup_{t=1}^{\tau} \mathcal{N}_t = \mathcal{N}$ and $\mathcal{N}_s \cap \mathcal{N}_t = \emptyset$ for $s \neq t$. Group membership is exogenous and verified by the platform. We assume *structural cost heterogeneity* across groups: the distribution of \bar{c}_i may differ between groups. We now introduce fairness concepts for group-structured workers at both the group level and the individual level. At the group level, we adopt an aggregate equity lens: the total economic benefit accruing to each group should be comparable. Given an allocation-payment pair (x, p) , the *group utility* of group t is $au_t(x, p) := \sum_{i \in \mathcal{N}_t} u_i(x, p) = \sum_{i \in \mathcal{N}_t} \sum_{k \in \mathcal{K}} x_{i,k} (p_{i,k} - c_{i,k})$. For a mechanism M and bid profile \mathbf{B} , the *expected group utility* is $\mathbb{E}[AU_t(\mathbf{B})]$, where the expectation is over the mechanism's internal randomness.

DEFINITION 2. Fix $\epsilon \geq 0$. A mechanism M is ϵ -group fair if for all feasible bid profiles \mathbf{B} and all groups $s, t \in [\tau]$, $|\mathbb{E}[AU_s(\mathbf{B})] - \mathbb{E}[AU_t(\mathbf{B})]| \leq \epsilon$.

We also adopt an individual-level criterion: agents that are close under a task-relevant similarity should receive similar outcomes [4].

DEFINITION 3. A mechanism M satisfies Lipschitz fairness if, on every admissible bid profile \mathbf{B} , for any $k \in \mathcal{K}$, $t \in [\tau]$, and pair of workers $i, j \in \mathcal{N}_t$, $|\mathbb{E}[X_{i,k}(\mathbf{B})] - \mathbb{E}[X_{j,k}(\mathbf{B})]| \leq d(b_{i,k}, b_{j,k})$, where $d(b, b') := \frac{|b - b'|}{\max\{b, b'\}}$.

3 GIFTA MECHANISM

We present GIFTA (Group and Individual Fair Task Allocation), a reverse-auction mechanism for group-structured workers that enforces our fairness notions while preserving desirable properties. Given task values $\vec{v} = (v_1, \dots, v_{\kappa})$ and a bid matrix \mathbf{B} , GIFTA returns an allocation $x(\mathbf{B})$ and payments $p(\mathbf{B})$. It proceeds in two stages:

- (1) Learn group-selection probabilities** $\{\gamma_{t,k}\}_{t \in [\tau], k \in \mathcal{K}}$. (a) Randomly split workers into a learning set L and an allocation set A . (b) For each group $t \in [\tau]$ and task $k \in \mathcal{K}$, using the intra-group allocation and payment rule (in Stage 2), simulate allocation of task k among workers $\mathcal{N}_t \cap L$ with the roles of L and A swapped to estimate the expected social cost, $\widetilde{SC}_{t,k}$, and the expected aggregate utility, $\widetilde{AU}_{t,k}$ of workers in $\mathcal{N}_t \cap L$. (c) For each task k , choose a




Figure 1: Performance of R-SPA and GIFTA

probability simplex vector $(\gamma_{1,k}, \dots, \gamma_{\tau,k})$ by solving

$$\begin{aligned} \min_{\{\gamma_{t,k}\}} & \sum_{k \in \mathcal{K}} \sum_{t \in [\tau]} \gamma_{t,k} \widetilde{SC}_{t,k} \\ \text{s.t.} & \left| \sum_{k \in \mathcal{K}} \gamma_{s,k} \widetilde{AU}_{s,k} - \sum_{k \in \mathcal{K}} \gamma_{t,k} \widetilde{AU}_{t,k} \right| \leq \epsilon, \forall s, t \in [\tau] \quad (\text{OPT}) \\ & \sum_{t=1}^{\tau} \gamma_{t,k} = 1, \forall k \in \mathcal{K}, \gamma_{t,k} \geq 0, \forall 1 \leq t \leq \tau, \forall k \in \mathcal{K} \end{aligned}$$

We solve (OPT) over L via the Lagrangian dual method [12], rewriting each absolute-value group-fairness constraint into two linear inequalities and introducing multipliers $\{\lambda_{s,t}^+, \lambda_{s,t}^-\}_{s < t}$. The resulting Lagrangian is additively separable across tasks, so the dual decomposes into κ independent subproblems (one per task) over $\{\gamma_{t,k}\}_{t \in [\tau]}$, under the simplex and non-negativity constraints. We then iterate between solving these subproblems for $\{\gamma_{t,k}\}$ and updating Λ with Adam [9] until convergence.

(2) Allocate and pay. (a) For each $k \in \mathcal{K}$, draw $i^* \sim \text{Categorical}(\gamma_{1,k}, \dots, \gamma_{\tau,k})$ and restrict attention to eligible bidders in $\mathcal{N}_{i^*} \cap A$ for task k . (b) Apply an intra-group auction rule (X', P') to the eligible workers: assign each worker $i \in \mathcal{N}_{i^*} \cap A$ a probability $\Pr(i, k | \mathbf{B}) := (1/b_{i,k}) / \sum_{j \in \mathcal{N}_{i^*} \cap A} (1/b_{j,k})$, and draw a winner $w \sim \text{Categorical}(\Pr(1, k | \mathbf{B}), \dots, \Pr(n, k | \mathbf{B}))$. Pay i for task k as $P'_{i,k}(\mathbf{B}) = b_{i,k} + \int_{b_{i,k}}^{v_k} \mathbb{E}[X_{i,k}(\mathbf{B}[i, k/y])] dy / \mathbb{E}[X_{i,k}(\mathbf{B})]$, where $\mathbf{B}[i, k/y]$ is the bid matrix where the (i, k) th entry $b_{i,k}$ is replaced by y .

THEOREM 1. The mechanism GIFTA is IC, IR, and BF. Moreover, it satisfies within-group Lipschitz fairness and ϵ -group fairness asymptotically as the market size $n \rightarrow \infty$. \square

4 EXPERIMENTS

We evaluate GIFTA across a range of settings, focusing on its social cost, ϵ -group fairness on A using probabilities $\gamma_{t,k}$ learned from L , and within-group individual fairness. Experimental results (Figure 1) show consistent trends: relaxing the group-fairness parameter ϵ lowers social cost and improves the app. ratio (app. := $\mathbb{E}[SC(C)] / sc^*(C)$), reflecting the fairness–efficiency trade-off. Compared with the reverse second-price auction (R-SPA) [2], GIFTA yields substantially smaller group gaps and its realized fairness on A closely matches the target ϵ learned from L . Moreover, GIFTA keeps within-group allocation disparities near zero, whereas R-SPA has an individual-fairness gap of 1.0. These patterns persist across all parameter sweeps: GIFTA consistently meets the group-equity constraint with low individual disparity. Source code is available at <https://anonymous.4open.science/r/GIFTA-0B03>.

ACKNOWLEDGMENTS

This work is supported by NSFC No. 62172077 and No. 62350710215.

REFERENCES

- [1] Georgios Birmpas, Andrea Celli, Riccardo Colini-Baldeschi, and Stefano Leonardi. 2022. Fair equilibria in sponsored search auctions: The advertisers' perspective. In *Proceedings of the Thirty-First International Joint Conference on Artificial Intelligence, IJCAI-22*. 95–101.
- [2] Felix Brandt and Gerhard Weiß. 2002. Antisocial Agents and Vickrey Auctions. In *Intelligent Agents VIII. ATAL 2001 (Lecture Notes in Computer Science, Vol. 2333)*, J.-J. Ch. Meyer and Milind Tambe (Eds.). Springer, 335–347. https://doi.org/10.1007/3-540-45448-9_25
- [3] Julia Dressel and Hany Farid. 2018. The accuracy, fairness, and limits of predicting recidivism. *Science Advances* 4, 1 (2018), eaao5580.
- [4] Cynthia Dwork, Moritz Hardt, Toniann Pitassi, Omer Reingold, and Richard Zemel. 2012. Fairness through awareness. In *Proceedings of the 3rd innovations in theoretical computer science conference*. 214–226.
- [5] Alireza Fallah, Michael Jordan, and Annie Ulichney. 2024. Fair allocation in dynamic mechanism design. *Advances in Neural Information Processing Systems* 37 (2024), 125935–125966.
- [6] Nina Grgic-Hlaca, Muhammad Bilal Zafar, Krishna P Gummadi, and Adrian Weller. 2016. The case for process fairness in learning: Feature selection for fair decision making. In *NIPS symposium on machine learning and the law*, Vol. 1. Barcelona, Spain, 11.
- [7] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of opportunity in supervised learning. *Advances in neural information processing systems* 29 (2016).
- [8] Fengjuan Jia, Mengxiao Zhang, Jiamou Liu, and Bakh Khoussainov. 2024. Balancing Efficiency with Equality: Auction Design with Group Fairness Concerns. In *ECAI 2024*. IOS Press, 3252–3259.
- [9] Diederik P Kingma. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980* (2014).
- [10] Matt J Kusner, Joshua Loftus, Chris Russell, and Ricardo Silva. 2017. Counterfactual fairness. *Advances in neural information processing systems* 30 (2017).
- [11] Anja Lambrecht and Catherine Tucker. 2019. Algorithmic bias? An empirical study of apparent gender-based discrimination in the display of STEM career ads. *Management Science* 65, 7 (2019), 2966–2981.
- [12] D.G. Luenberger and Y. Ye. 2015. *Linear and Nonlinear Programming*. Springer International Publishing.
- [13] Roger B Myerson. 1986. Multistage games with communication. *Econometrica: Journal of the Econometric Society* (1986), 323–358.
- [14] Ziad Obermeyer, Brian Powers, Christine Vogeli, and Sendhil Mullainathan. 2019. Dissecting racial bias in an algorithm used to manage the health of populations. *Science* 366, 6464 (2019), 447–453.
- [15] Hai Tang, Mingjun Xiao, Guoju Gao, and Hui Zhao. 2020. Reverse-auction-based crowdsourced labeling for active learning. *World Wide Web* 23, 1 (2020), 671–689.
- [16] Taiki Todo, Runcong Li, Xuemei Hu, Takayuki Mouri, and Makoto Yokoo. 2011. Generalizing Envy-Freeness toward Group of Agents. In *IJCAI 2011, Proceedings of the 22nd International Joint Conference on Artificial Intelligence, Barcelona, Catalonia, Spain, July 16-22, 2011*.