

Resource allocation with cooperative agents

Stefania Costantini¹, Giovanni De Gasperis¹, Pasquale De Meo², Francesco Gullo¹ and Alessandro Proveti³

¹University of L'Aquila, Italy

²University of Messina, Italy

³Birkbeck, University of London, UK, and University of Milan, Italy

Abstract

We study the problem of cooperative resource allocation in multi-agent systems, focusing on scenarios such as hospital networks. In our model, agents (e.g., hospitals) redistribute limited resources, such as medical personnel, in a way that satisfies both local constraints and global equity objectives. We devise ad-hoc optimization strategies for a static scenario, where resource needs are fixed over time. We empirically evaluate the proposed approaches through a set of experiments. Our results demonstrate that our approaches are highly effective.

Keywords

Multi-agent systems, Cooperating agents, Resource allocation, Reinforcement learning

1. Introduction

The fair and efficient allocation of resources in decentralised environments has long been a fundamental challenge in Artificial Intelligence and Economics [1]. In domains where autonomous agents pursue common goals, such as healthcare networks (e.g., the British National Health Service) [2], wireless networks [3], or cloud computing networks [4], the ability to enable cooperation without centralised authority is of both theoretical and practical importance.

In this work, we consider a network of agents managing local resources to achieve individual objectives while cooperating toward a collective goal through resource exchanges (e.g., lending). This paradigm aligns with the concept of *fairness*, which ensures equitable outcomes while maintaining system functionality. Our motivating example is an idealized model of the NHS healthcare network: hospitals manage their physician rosters but may temporarily lend doctors to others during localized emergencies (e.g., outbreaks of transmissible diseases). This scenario shares similarities with the inter-banking scenario [5, 6], though healthcare systems differ in their universal objective of avoiding hospital failures, unlike banking systems where central banks underwrite systemic stability.

A number of approaches have been proposed so far to allocating resources in multi-agent systems (MARA), including Distributed Constraint Optimization [7], Social Choice Theory [8], and Market-Based Coordination [9]. Among these, *Nash Welfare Optimization (NWO)* stands out as a principled method. NWO models societal welfare as the geometric mean of individual utilities, and balances efficiency and fairness by prioritizing improvements for agents with lower allocations and it comes with no surprise that NWO has been widely studied for its theoretical guarantees and practical effectiveness [10, 11].

However, we identify three critical limitations in its application to healthcare networks, namely:

(a) *Minimum vs. Target Staffing*: while NWO ensures hospitals meet *baseline staffing* thresholds (that is, we assume that each hospital has at least m_i doctors who ensure its functioning), it ignores *aspirational targets* (that is, we assume that each hospital wants at least t_i doctors) needed for optimal service delivery. In some cases, it is appropriate to concentrate more resources in highly specialised

Joint Proceedings of the Workshops and Doctoral Consortium of the 41st International Conference on Logic Programming, September 9–13, 2025, Rende, Italy

✉ stefania.costantini@univaq.it (S. Costantini); giovanni.degasperis@univaq.it (G. De Gasperis); pdemeo@unime.it (P. De Meo); francesco.gullo@univaq.it (F. Gullo); a.proveti@bbk.ac.uk (A. Proveti)

ORCID 0000-0002-5686-6124 (S. Costantini); 0000-0001-9521-4711 (G. De Gasperis); 0000-0001-7421-216X (P. De Meo); 0000-0002-7052-1114 (F. Gullo); 0000-0001-9542-4110 (A. Proveti)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).

medical centres (such as hospitals specialising in rare diseases or involved in innovative clinical trials), even if this may lead to slight inequalities in staff distribution. NWO does not meet this requirement.

(b) *Constraint Handling*: NWO lacks explicit mechanisms to enforce global constraints, such as the fact that the total number of doctors must be constant or hard lower bounds on m_i thresholds. Specifically, the constraint on the overall size of the labour force is a direct consequence of public expenditure management policies, which in some countries impose a freeze on recruiting new staff [12].

(c) *Dynamic Adaptability*: in emergencies (e.g., pandemics), staffing needs fluctuate rapidly. NWO (as well as any other method based on an optimisation algorithm) requires replanning the allocation of doctors from scratch after each change, but this is a computationally prohibitive task for large hospital networks. Moreover, the efficacy of reallocation is questionable if a new distribution of the workforce is necessitated after a brief period.

In this work, we propose novel approaches that address limitations (a)-(c). We focus on a *static scenario*, i.e., we assume that the staffing needs of each hospital are fixed over time. We defer the study of a dynamic scenario (in which the demand for staff vary over time) to future work. Our main contribution can be summarised as follows.

1. We introduce a new optimisation problem where the objective function is the sum of the squares of the differences between the number of doctors actually assigned to a hospital and the corresponding target. Our problem can be formulated as a *quadratic programming* (QP) problem [13], which can be solved with very accurate and efficient solvers.
2. In defining our QP problem, we explicitly introduce constraints on the minimum number of doctors to be allocated to each hospital, together with invariance of the total number of doctors in the system. We also reformulate the NWO method to incorporate these constraints and obtain a convex optimisation problem for which we have efficient solvers [13]. We compare the QP method with NWO and with a reallocation method called *Progressive Taxation*, which simulates a tax system in which the wealthiest individuals donate some of their resources to increase social justice [14]. We also devise a hybrid formulation that combines the objective function of the QP and NWO problems. Experimental results show that the QP model excels in reducing the difference between the number of doctors assigned to a hospital and the target, while the NWO model is more effective in ensuring a fair distribution of doctors, where fairness is measured by means of the *Gini Index* [15] of the staff size across all the various hospitals. We provide an empirical evaluation, whose results attest the high effectiveness of the proposed approaches, with QP and NWO mostly prevailing over Progressive Taxation.

While our approach is grounded in healthcare, its principles generalize to other domains requiring cooperative resource redistribution.

2. Multi-agent Resource Allocation: Proposed Approaches

Let us begin with an idealised multi-agent resource allocation framework that models an NHS-style healthcare domain. We are given a set of n hospitals, denoted as $\mathcal{H} = \{\vec{h}_1, \vec{h}_2, \dots, \vec{h}_n\}$ where each hospital \vec{h} has three dimensions: the current number of doctors available, c_i , the target rooster, t_i and the minimum number of doctors needed for the hospital to operate, m_i . In a static scenario, the total number of doctors available in \mathcal{H} is fixed, i.e., $\sum_{i=1}^n c_i = C$.

Clearly, if some hospital \vec{h}_i has critically-low levels of staff, i.e., $c_i \leq m_i$, the best option is to recruit more doctors. Yet, this may not be possible, even for relatively-long interim periods. Thus, we consider transferring doctors from other hospitals to improve the overall efficiency of the hospital system \mathcal{H} . The idea is that ‘wealthy’ hospitals (those operating at or near full rooster) could lend doctors to \vec{h}_i . This practice is common in NHS-style health systems, e.g., in Summer when population density in tourism areas shows huge alterations. Similarly, emergency situations, such as the outbreak of epidemics and

their containment in the geographical areas where they have occurred, may require the emergency re-assignment of medical staff to cope with the spike in hospital admissions.

The goal now is to model the emergency re-distribution scenario, so as to compute solutions that are (near-)optimal w.r.t. the global objective of maintaining all hospitals viable and operating, while respecting all constraints on the capacity of individual hospitals. In the remainder of this section, we describe the methods we devise for achieving this goal. In particular, Section 2.1 presents a *Quadratic Programming* [13] formulation. We then describe two approaches to redistributing resources that have been widely studied in economic theory, known as *Progressive Taxation* (Section 2.2) and *Nash Welfare Optimisation* (Section 2.3), along with a hybrid approach that combines them.

2.1. Quadratic Programming (QP)

We introduce the notion of *inter-hospital transfer matrix* (for short, *transfer matrix*):

Definition 1 (Transfer Matrix). *Given a set $\mathcal{H} = \{\vec{h}_1, \vec{h}_2, \dots, \vec{h}_n\}$ of hospitals, the transfer matrix $\mathbf{X} \in \mathbb{R}^{n \times n}$ is a matrix whose entries x_{ij} quantify the flow of staff from \vec{h}_i to \vec{h}_j . If $x_{ij} > 0$, then \vec{h}_i has lent staff to \vec{h}_j . If $x_{ij} = 0$, then there is no flow of doctors from \vec{h}_i to \vec{h}_j . If $x_{ij} < 0$, then \vec{h}_i has received doctors from hospital \vec{h}_j .*

Notic how Definition 1 assumes that the lease of doctors between hospitals may be *fractional*. This choice has important computational implications, as it enables effective and efficient algorithms. Also, this choice offers a higher level of flexibility in the sense that decision-makers receive suggestions on how to redistribute the available workforce but they have some leeway in the final choice.

We can now define y_i , the staff level at \vec{h}_i after redistribution, as the sum of the initial staff count at \vec{h}_i plus the number of incoming doctors and minus the number of those transferred elsewhere:

$$y_i = c_i + \sum_{j=1}^n x_{ji} - \sum_{j=1}^n x_{ij} \quad (1)$$

We collect all coefficients y_i into a vector $\mathbf{y} = [y_1, y_2, \dots, y_n]$. The following constraints apply:

a) *minimum satisfaction of individual hospital requests*: after the redistribution of doctors, each hospital \vec{h}_i must have a number of doctors at least equal to m_i , i.e., $y_i \geq m_i$;

b) *retention of staff*: the number of doctors working in each hospital after the redistribution must be equal to the number of doctors initially on duty, since no doctors have been hired/fired: $\sum_{i=1}^n y_i = C$. Now, a solution exists whenever the total number of available staff, C , is sufficient to cover the aggregated minimal demands from hospitals, i.e., $\sum_{i=1}^n m_i \leq C$.

The quantity $y_i - t_i$ expresses the difference between the number of doctors on duty in the hospital \vec{h}_i after the redistribution and the target number of doctors. A key desideratum for staff reallocation consists in making $y_i - t_i$ as small as possible (ideally, no staff should be asked to move to a new workplace). In practice, we may have the following two *sub-optimal scenarios*. First, an \vec{h}_i may have fewer doctors than its target (and therefore the difference $x_i - t_i$ is negative). Second, \vec{h}_i may have more doctors than it actually wants (and, thus, the difference $x_i - t_i$ is positive). To compensate the two errors above, we take the square of $x_i - t_i$ as an estimate of the error and we compute the sum of the errors across all the hospitals as the objective function to be minimized. This results in the following quadratic programming problem:

$$\min_{\mathbf{y}} \left\{ \sum_{i=1}^n (y_i - t_i)^2 \right\} \quad \text{s.t.} \quad \sum_{i=1}^n y_i = C \quad \text{and} \quad y_i \geq m_i \quad \forall i \in 1, \dots, n. \quad (2)$$

The above optimisation problem is a convex quadratic programming one. As such, it can be solved by off-the-shelf solvers¹.

¹E.g., the `qpsolvers` unified Python module for quadratic programming: <https://pypi.org/project/qpsolvers/>.

2.2. Progressive Taxation

The *progressive taxation* model of [14] is a fiscal framework wherein the effective tax rate escalates commensurately with increases in the taxable base (e.g., income or asset valuation). This mechanism facilitates wealth redistribution and mitigate the gap between rich and poor individuals in the society; consequently, progressive taxation aligns with the principle of *vertical equity*: those with greater resources have a bigger capacity to contribute to government funding.

We formalise progressive taxation in a multi-agent models with the following redistribution rule. Consider the top and bottom decile:

- R (Resource-abundant): top 10% hospitals by staff size c_i ;
- P (Resource-constrained): bottom 10% hospitals by staff size c_i .

The progressive-taxation-based reallocation procedure will iteratively apply the equations below (where the “ (k) ” superscript denotes the value at iteration k):

Taxation Phase:

$$\delta_{ij}^{(k)} = \begin{cases} -\frac{0.1y_i^{(k)}}{|P|} & \text{if } i \in R, j \in P \\ 0 & \text{otherwise} \end{cases} \quad (3)$$

$$\Delta_i^{(k)} = \begin{cases} \sum_{j \in P} \delta_{ij}^{(k)} = -0.1y_i^{(k)} & \text{if } i \in R \\ \sum_{j \in R} \delta_{ji}^{(k)} = \frac{1}{|P|} \sum_{j \in R} 0.1y_j^{(k)} & \text{if } i \in P \\ 0 & \text{otherwise} \end{cases} \quad (4)$$

Allocation Update:

$$y_i^{(k+1)} = \max \left\{ m_i, y_i^{(k)} + \Delta_i^{(k)} \right\} \quad (5)$$

The resulting entries $x_{ij}^{(k+1)}$ of the transfer matrix at the $(k+1)$ -th iteration are equal to $\delta_{ij}^{(k)}$ (resp. $-\delta_{ji}^{(k)}$) only if i (resp. j) belongs to R , j (resp. i) belongs to P , and $y_i^{(k)} + \Delta_i^{(k)}$ (resp. $y_j^{(k)} + \Delta_j^{(k)}$) is no less than m_i (resp. m_j). In all other cases, $x_{ij}^{(k+1)} = 0$.

The rationale of the above equations is as follows. Hospitals in class R donate 10% of their staff² to those in P . Thus, 10% of the total staff in R is equally distributed among hospitals in P . Conversely, hospitals in the middle deciles experience no change in staffing.

The above equations lead to a solution that satisfies the constraints of min satisfaction of individual hospital requests and staff retention (see Sect. 2.1). The final solution is found by either running the reallocation for a given number of steps K or setting convergence bounds (although no formal convergence is guaranteed for the general case).

2.3. Nash Welfare Optimization (NWO)

Nash Welfare Optimization (NWO) [16, 17] is a Pareto-efficient allocation mechanism whose maximization prevents extreme disparities in resource distribution. This makes NWO well-suited for our goal of redistributing doctors among hospitals, as it avoids scenarios where hospitals fall below minimum operational staffing levels ($y_i < m_i$) or become excessively overstaffed ($y_i \gg m_i$).

²The 10% threshold is arbitrary and can be adjusted by decision makers.

NWO solves the following constrained optimization problem, where the objective function is the product of agents' utilities:

$$\max_{\mathbf{y}} \{ \prod_{i=1}^n (y_i - m_i + 1) \} \quad \text{s.t.} \quad \sum_{i=1}^n y_i = C \quad \text{and} \quad y_i \geq m_i \quad \forall i \in \{1, \dots, n\}. \quad (6)$$

Next, the multiplicative objective in (6) can be transformed into a convex optimization problem:

$$\max_{\mathbf{y}} \{ \sum_{i=1}^n \log (y_i - m_i + 1) \},$$

which is computationally tractable and solvable with standard convex optimization solvers.

The choice of $y_i - m_i + 1$ (rather than $y_i - t_i$) in the utility function is critical. In fact, if we used $y_i - t_i$, the objective function to be maximised would be the product of terms of the type $y_i - t_i$, and each of these terms must be strictly positive (if this was not the case, the utility of some agents would be negative or equal to zero and the logarithm would make no sense). Thus, the condition $y_i - t_i > 0$ is verified if we assume that the number of available doctors is greater than the sum of the targets t_i , i.e. $\sum_{i=1}^n t_i \leq C$. In this case, each hospital could be assigned the target number of doctors and the surplus could be distributed randomly. In practice, we can assume that $\sum_{i=1}^n m_i \leq C$, but in general we expect $\sum_{i=1}^n m_i$ to be considerably less than $\sum_{i=1}^n t_i$. In this case, if the NWO algorithm tried to allocate each hospital a number of doctors equal to (or greater than) its target, the remaining doctors would not be sufficient to meet the minimum requirements of other hospitals, with the consequence that the constraint $y_i \geq m_i$ would be violated for some hospitals.

In what follows, we will also consider a hybrid optimisation strategy where the cost function is a linear combination of the quadratic programming and NWO ones:

$$\min_{\mathbf{y}} \{ \alpha \sum_{i=1}^n (y_i - t_i)^2 - (1 - \alpha) \sum_{i=1}^n \log (y_i - m_i + 1) \} \quad (7)$$

$$\text{s.t.} \quad \sum_{i=1}^n y_i = C \quad \text{and} \quad y_i \geq m_i \quad \forall i \in 1..n.$$

The objective function in (7) is convex as it is a linear combination of convex functions, thus efficient solvers are available. Note that the parameter α controls the contribution of the cost functions considered in the quadratic-programming and NWO cases. So, as $\alpha \rightarrow 0$ we see the cost function converge to NWO whereas for $\alpha \rightarrow 1$ it converge to the QP formulation instead.

3. Experiments

We evaluate the effectiveness of our approach through experiments designed to address how the proposed NWO, Progressive Taxation, QP and hybrid strategies perform in practice. Specifically, we are interested in assessing how effectively our approaches redistribute doctors to minimize inequality while ensuring each hospital meets its staffing target.

We adopt the following metrics to evaluate the performance of our approaches.

1. **Target Deviation** (or Mean Absolute error: MAE), which is defined as the average difference between actual staff levels and targets:

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |y_i - t_i| \quad (8)$$

Ideally, we would like each hospital to have as close to the target number of doctors as possible; thus, the lower MAE, the more effective the approach to redistributing doctors.

2. **Gini Inequality Index** (\mathcal{G}) [15], defined as

Table 1

Accuracy results (average and standard deviation) of the proposed approaches.

Method	Metric	
	Gini	Target Deviation
NWO	0.024 ± 0.001	35.013 ± 1.574
Progressive Taxation	0.023 ± 0.001	39.196 ± 2.355
QP	0.150 ± 0.010	1.132 ± 0.817

$$\mathcal{G} = \frac{1}{n} \left(n + 1 - 2 \sum_{i=1}^n \frac{\sum_{j=1}^i y_j}{\sum_{j=1}^n y_j} \right) \quad (9)$$

The Gini Index is one of the most commonly used inequality measures (e.g., income inequality or inequality in life expectancy). It ranges from 0 (perfect equality) to 1 (perfect inequality). In our case, a \mathcal{G} close to 1 indicates that the available doctors are concentrated in a few hospitals, which is undesirable. Thus, ideally, low values of \mathcal{G} should be sought.

3.1. Results

Our experiments aim to compare the effectiveness of medical staff reallocation strategies in a *static scenario*, i.e., where hospital staffing requirements do not change over time.

To do so, we generated a large test instance made up of 150 randomly-generated hospitals; i.e., for each hospital the minimum number of staff needed to function properly, the current number of doctors on the roster and the target number were assigned at random.

We ran the QP, Progressive Taxation, and NWO strategies of Sec. 2 on the random instance and computed the Gini Index and Target Deviation achieved by each method. To ensure statistical robustness, we repeated the experiment 20 times and calculated the mean and standard deviation of both measures. The results are shown in Table 1. We can observe that QP and NWO exhibit different behaviours. In fact, QP tends to redistribute resources to ensure that the number of doctors actually allocated to a hospital is as close to the target as possible. As a result, some hospitals may receive many more doctors than others, leading to a more unequal distribution of human resources and a corresponding increase in the Gini Index. NWO follows the opposite logic, as its redistribution aims to ensure that all hospitals end up with a comparable number of staff, which can have a negative impact on the Target Deviation.

Progressive Taxation behaves similarly to NWO and tends to smooth out inequalities. This is witnessed by a Gini Index value that gets close to that of NWO and by a drop in Target Deviation.

We then focus on the hybrid strategy and examine how the Target Deviation and the Gini Index vary as the parameter α changes. We have divided the range of variation of α (i.e., the segment from zero to one) into 35 intervals of equal size, and each division point corresponds to a value of α used to compute the objective function of the hybrid strategy. Figure 1 shows mean and standard deviation of Target Deviation and Gini Index obtained by the hybrid strategy. As α approaches zero, the objective function of the hybrid strategy coincides with that of NWO and this explains why we obtain increasingly higher values of the objective deviation, accompanied by smaller values of the Gini Index.

In the health sector, advanced therapies and innovative drugs are expensive and often require a large and well-trained workforce; consequently, some treatments should be concentrated in a few highly specialised hospitals, which should also have significant staffing levels. We therefore believe that minimising the Target Deviation is at least as important as minimising the Gini Index.

4. Related Work

Multiagent resource allocation (MARA) constitutes a fundamental challenge in multiagent systems, requiring autonomous agents to distribute limited resources in ways that balance efficiency and fairness.




Figure 1: Target Deviation and Gini Index of the hybrid strategy as a function of the α parameter.

Such problems arise ubiquitously [1]. As systems grow in scale and complexity, designing mechanisms that reconcile individual agent incentives with collective welfare becomes increasingly critical.

At its core, MARA involves agents negotiating resource distributions, often encountering dilemmas where self-interest conflicts with group optimality. Central here is the concept of *fairness*, a principle vital to both human societies and artificial systems [18, 19, 20].

Our work can be positioned under the broad umbrella of MARA. Specifically, we propose novel strategies for the re-allocation of resources in a multi-agent, cooperative setting.

Centralized vs. decentralized approaches. Traditional centralized methods, such as the Hungarian and Gale-Shapley algorithms, rely on a central authority with full system knowledge to compute optimal allocations [9]. The *Nash Welfare Optimization (NWO)* offers a principled framework for fair resource distribution by maximizing the geometric mean of agent utilities. NWO balances efficiency and equity, prioritizing improvements for agents with lower initial utility [8]. NWO occupies a middle ground between *utilitarian welfare* (maximizing total utility) and *egalitarian welfare* (maximizing minimum utility) and satisfies axiomatic properties such as scale invariance, Pareto efficiency, and independence of irrelevant alternatives [11, 10]. Due to these strengths, NWO has been applied successfully in collective decision-making, project funding allocation, and fair division problems [10].

In this work, we provide a novel contextualization of NWO to the multi-agent, cooperative resource re-allocation setting, based on quadratic programming.

5. Conclusions

Our work addresses cooperative resource re-allocation in a multi-agent environments, focusing on the important case of redistribution of hospital staff across a regional Health system such as the British NHS. We devise three principles approaches to this problem, which are based on quadratic programming (QP), Progressive Taxation, and Nash Welfare Optimization (NWO), respectively. We conduct experiments whose main results attest the high effectiveness of the proposed approaches. In the future, we plan to investigate the problem in a dynamic scenario, where the demand for staff vary over time.

Acknowledgments

Research partially supported by the PNRR Project CUP E13C24000430006 “Enhanced Network of intelligent Agents for Building Livable Environments - ENABLE”, and by PRIN 2022 CUP E53D23007850001 Project “TrustPACTX - Design of the Hybrid Society Humans-Autonomous Systems: Architecture, Trustworthiness, Trust, EthiCs, and EXplainability (the case of Patient Care)”.

Declaration on Generative AI

During the preparation of this work, the author(s) used Grammarly, solely to spell check and improve the grammar. The tool was not used to alter, generate, or influence the semantic contents of the paper. The authors retain full responsibility for the accuracy, originality, and integrity of the entire work.

References

- [1] N. Jong, P. Stone, M. Taylor, Multiagent resource allocation: A review of mechanisms and applications, *Autonomous Agents and Multi-Agent Systems* 22 (2008) 1–29.
- [2] Y. Zhao, N. Behari, E. Hughes, E. Zhang, D. Nagaraj, K. Tuyls, A. Taneja, M. Tambe, Towards a pretrained model for restless bandits via multi-arm generalization, in: *Proc. of the International Joint Conference on Artificial Intelligence, (IJCAI 2024)*, Jeju, South Korea, 2024, pp. 321–329.
- [3] J. Cui, Y. Liu, A. Nallanathan, Multi-agent reinforcement learning-based resource allocation for UAV networks, *IEEE Transactions on Wireless Communications* 19 (2020) 729–743.
- [4] Y. Zhao, Y. Liu, B. Jiang, T. Guo, CE-NAS: an end-to-end carbon-efficient neural architecture search framework, in: *Proc. of the International Conference on Advances in Neural Information Processing Systems* 38 (NIPS 2024), Vancouver, BC, Canada, 2024.
- [5] S. Battiston, M. Puliga, R. Kaushik, P. Tasca, G. Caldarelli, Debtrank: Too central to fail? financial networks, the fed and systemic risk, *Scientific Reports* 2 (2012) 541. URL: <https://doi.org/10.1038/srep00541>. doi:10.1038/srep00541.
- [6] J. Tong, B. de Keijzer, C. Ventre, Reducing systemic risk in financial networks through donations, in: *Proc. of the European Conference on Artificial Intelligence (ECAI 2024)*, volume 392, IOS Press, Santiago de Compostela, Spain, 2024, pp. 3405–3412.
- [7] S. de Jong, S. Uyttendaele, K. Tuyls, Learning to reach agreement in a continuous ultimatum game, *J. Artif. Intell. Res.* 33 (2008) 551–574. URL: <https://api.semanticscholar.org/CorpusID:13248455>.
- [8] L. Zhang, Y. Xu, F. Fang, Y. Yang, Online nash social welfare maximization in multi-agent systems, in: *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, 2024.
- [9] A. Nongailard, H. Sohler, V. Hilaire, Centralized and distributed approaches for resource allocation: A comparative study, *Journal of Intelligent Manufacturing* 27 (2016) 789–803.
- [10] T. Delemazure, F. Durand, F. Mathieu, Aggregating correlated estimations with (almost) no training, in: *Proc. of the European Conference on Artificial Intelligence, (ECAI 2023 - 26th)*, volume 372 of *Frontiers in Artificial Intelligence and Applications*, IOS Press, Krakow, Poland, 2023, pp. 541–548.
- [11] S. Hossain, E. Micha, N. Shah, Fair algorithms for multi-agent multi-armed bandits, *Advances in Neural Information Processing Systems* 34 (2021) 24005–24017.
- [12] I. Papanicolas, L. R. Woskie, A. K. Jha, Health care spending in the united states and other high-income countries, *Jama* 319 (2018) 1024–1039.
- [13] S. P. Boyd, L. Vandenberghe, *Convex optimization*, Cambridge university press, 2004.
- [14] J. E. Stiglitz, J. K. Rosengard, *Economics of the public sector: Fourth international student edition*, WW Norton & Company, 2015.
- [15] F. A. Farris, The gini index and measures of inequality, *The American Mathematical Monthly* 117 (2010) 851–864.
- [16] H. Moulin, *Fair division and collective welfare*, MIT press, 2004.
- [17] J. F. Nash, et al., The bargaining problem, *Econometrica* 18 (1950) 155–162.
- [18] D. Ceragioli, F. Rossi, B. Venable, Fairness-aware distributed planning for resource allocation in multiagent systems, in: *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI)*, 2024.
- [19] A. Jiang, K. Leyton-Brown, Fairness in multi-agent systems with reinforcement learning, *Artificial Intelligence* 275 (2019) 25–64.
- [20] Y. Chen, D. C. Parkes, Envy-free allocation in combinatorial auctions, *Games and Economic Behavior* 70 (2010) 1–14.