Fair Allocation with Initial Utilities*

Extended Abstract

Niclas Boehmer
Hasso Plattner Institute, University of Potsdam
Potsdam, Germany
Niclas.Boehmer@hpi.de

Luca Kreisel
Hasso Plattner Institute, University of Potsdam
Potsdam, Germany
Luca.Kreisel@hpi.de

ABSTRACT

The problem of allocating indivisible resources to agents arises in a wide range of domains, including treatment distribution and social support programs. An important goal in algorithm design for this problem is fairness, where the focus in previous work has been on ensuring that the computed allocation provides equal treatment to everyone. However, this perspective disregards that agents may start from unequal initial positions, which is crucial to consider in settings where fairness is understood as *equality of outcome*. In such settings, the goal is to create an equal final outcome for everyone by leveling initial inequalities through the allocated resources. To close this gap, focusing on agents with additive utilities, we extend the classic model by assigning each agent an initial utility and study the existence and computational complexity of several new fairness notions following the principle of equality of outcome. Among others, we show that complete allocations satisfying a direct analog of envy-freeness up to one resource (EF1) may fail to exist and are computationally hard to find, forming a contrast to the classic setting without initial utilities. We propose a new, always satisfiable fairness notion, called minimum-EF1-init and design a polynomial-time algorithm based on an extended round-robin procedure to compute complete allocations satisfying this notion.

KEYWORDS

Fair Allocation; Indivisible Resources; Envy-Freeness; Equality of Outcome; Algorithmic Analysis.

ACM Reference Format:
Niclas Boehmer and Luca Kreisel. 2026. Fair Allocation with Initial Utilities: 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/DMYL9092

1 INTRODUCTION

Allocating resources to agents is a central problem in algorithmic decision-making, with applications ranging from task and house allocation to treatment distribution and scheduling [1, 6, 14, 15, 20, 21, 25]. Thus, the problem has been widely studied across artificial intelligence, machine learning, multi-agent systems, operations research, and computational social choice. Much of this literature focuses on the design and analysis of algorithms that compute allocations satisfying formal fairness criteria, for instance, envy-freeness [2, 7, 19]. Both within and outside computer science, the discourse on fairness distinguishes between two paradigms: On the one hand, *equality of opportunity* (also known as *formal equality* in law) demands equal treatment for everyone regardless of individual circumstances [13, 16]. On the other hand, *equality of outcome* (also known as *substantive equality* in law) acknowledges the existence of initial disparities and seeks to create an equal final outcome for everyone by leveling them [5, 20]. These two principles stand in an inherent conflict with each other, as countering initial disparities oftentimes necessitates the prioritization of those starting from an inferior position, contradicting the principle of equal treatment.[1]

In fair allocation, an overwhelming majority of works adhere to equality of opportunity: They assess fairness by how the allocation distributes resources—irrespective of agents' starting positions [2, 7, 19]. For example, under the classic notion of envy-freeness, an allocation is considered fair if no agent prefers the bundle allocated to another agent to their own. However, this ignores initial inequalities prevalent in many application areas of fair allocation. For example, in medical resource allocation programs, patients have different initial probabilities of recovery; educational interventions serve students with varying levels of prior knowledge or preparedness; and more generally, support programs often address individuals from diverse backgrounds and starting positions [10, 12, 14].

Motivated by the prevalence of such disparities, this paper initiates the study of equality of outcome in algorithmic fair allocation. In many real-world allocation problems, decision makers aim for equal post-allocation outcomes instead of equal treatment [18, 23]. One example is the Fairer Scotland Duty [23], which places a statutory responsibility on certain public bodies to consider how their decisions can help reduce inequalities of outcome. Notably, existing, intensively studied formal models of fairness reflecting equality of opportunity fail to capture this goal of equality of outcome, a gap that we address in this work.

Related Work. While fair allocation has been studied extensively, extending the model to account for heterogeneous starting positions has received surprisingly little attention. Most closely related are the works by Prakash HV et al. [22] and Deligkas et al. [11] on the following problem: given a partial initial allocation, can it be extended to a complete allocation satisfying a specified, traditional fairness criterion? Our model can be viewed as a special case of theirs, where all agents assign identical utilities to the initially allocated resources and each agent holds at most one resource in the initial allocation. As a result, some of their positive results carry over to our setting. However, our work fundamentally differs in that we design new fairness notions following the principle of

*The full version of our paper can be found at https://arxiv.org/abs/2602.14850.

[1] This conflict recently surfaced in an executive order of US President Trump [24].

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. https://doi.org/10.65109/DMYL9092

equality of outcome, tailored to the setting with initial utilities, rather than focusing on existing axioms. This enables stronger axiomatic guarantees more relevant to our setting. We refer to the full version for a detailed discussion and additional related work.

Preliminaries. We define $\mathbb{R}_{\geq 0} \coloneqq \{x \in \mathbb{R} \mid x \geq 0\}$. We consider the problem of allocating a set R of m *indivisible resources* among a set A of $n \geq 2$ *agents*. We refer to subsets $X \subseteq R$ of resources as *bundles* and denote by 2^R the set of all bundles. Each agent $i \in A$ has a *utility function* $u_i : 2^R \to \mathbb{R}_{\geq 0}$, where $u_i(\emptyset) = 0$.[2] We assume throughout that utility functions are *additive*, i.e., $u_i(X) = \sum_{r \in X} u_i(\{r\})$ for any $i \in A$ and $X \subseteq R$. In certain cases, we consider *identical resources*, where $u_i(\{r\}) = u_i(\{r'\})$ for all $i \in A$ and resources $r, r' \in R$. An *allocation* $\mathcal{X}$ is a tuple of n disjoint bundles $(X_1, \ldots, X_n)$ such that $X_i \cap X_j = \emptyset$ for all distinct $i, j \in A$, where X_i is *allocated* to agent $i \in A$. An allocation $\mathcal{X}$ is *complete* if $\bigcup_{i \in A} X_i = R$.

2 OUR CONTRIBUTIONS

We initiate the formal study of fair allocation for agents with initial disparities, focusing on fairness notions aligned with the principle of equality of outcome. We do so by equipping each agent i with an agent-specific *initial utility* value $b_i \in \mathbb{R}_{\geq 0}$, which we assume to be known to the allocation algorithm and comparable across agents. We restrict attention to additive utilities, which is a standard and well-established choice in the literature, particularly when introducing new settings or fairness notions [2–4, 8, 9, 17].

We adapt the classical fairness notions – envy-freeness (EF) and envy-freeness up to one resource (EF1) – to our setting and term the adaptations EF-init and EF1-init, respectively. Following the principle of equality of outcome, we assess the fairness of an allocation by the agents' total post-allocation utility. Concretely, we say that an allocation $\mathcal{X}$ is *EF1-init*, if for every pair of agents $i, j \in A$ either $X_j = \emptyset$[3] or there exists a resource $r \in X_j$ such that $b_i + u_i(X_i) \geq b_j + u_i(X_j \setminus \{r\})$. *EF-init* is defined analogously without the removal of one of j's resources. We find that introducing initial utilities fundamentally changes the nature of the problem: even in simple cases with identical resources, complete EF1-init allocations may fail to exist, since agents may disagree on how many resources are needed to bridge the differences in initial utility. Moreover, deciding whether such an allocation exists is NP-complete. This stands in stark contrast to the classical setting, where complete EF1 allocations always exist and are efficiently computable.

Nevertheless, we show that, for a constant number of agents (under a unary encoding of utility values), the existence of complete EF-init and EF1-init allocations can be decided efficiently using dynamic programming. Furthermore, for the special case of identical resources, we present a polynomial-time algorithm for deciding the existence of EF-init allocations. From a theoretical perspective, this result highlights the added complexity introduced by initial utilities: for identical resources, any envy-free allocation simply assigns the same number of resources to each agent in the classical setting without initial utilities. In contrast, in our model, ensuring fairness for identical resources requires a technically involved dynamic programming approach and a meticulous analysis of the interplay between initial utilities and allocated bundles.

The non-existence of complete EF1-init allocations limits the solution concept's applicability and highlights the need for fundamentally new approaches to achieve an always satisfiable envy-based notion. We propose minimum-EF1-init, whose core idea is to relax EF1-init when evaluating whether an agent i with higher initial utility envies an agent j with lower initial utility. To resolve the possible disagreement on the value of the resources needed to offset j's initial utility disadvantage discussed above, we depart from comparing the sums of the initial utilities and i's values for the bundles: Instead, we first allow j to offset their disadvantage with a subset X^* of j's resources, bounded in value by the initial utility difference. Then, we compare i's bundle to j's bundle without X^* under EF1, ensuring that after compensating for the initial utility difference, the remaining resources are distributed fairly from i's perspective. However, considering pairs of agents individually when restricting X^* proves insufficient to obtain a satisfiable notion, as the difference in initial utilities might be easier to offset for some agents than for others. To overcome this, we consider the minimum utility for each resource in X^* among agents with lower initial utility than i.

Definition 2.1. *An allocation $\mathcal{X}$ is* minimum-EF1-init *(min-EF1-init) if for every pair $i, j \in A$, we have $X_j = \emptyset$ or it holds that:*

(C1) If $b_i \leq b_j$, then $b_i + u_i(X_i) \geq b_j + u_i(X_j \setminus \{r\})$ for some $r \in X_j$.

(C2) If $b_i > b_j$, then there exists a resource $r \in X_j$ and a subset $X^* \subseteq X_j$ with

$$\sum_{r' \in X^*} \min_{j' \in A:\, b_{j'} < b_i} u_{j'}(\{r'\}) < b_i - b_j,$$

such that $u_i(X_i) \geq u_i(X_j \setminus (X^* \cup \{r\}))$ holds.

We show that minimum-EF1-init coincides with EF1-init in settings where resources' usefulness is diminishing in initial utility, i.e., when $b_i < b_j$ implies $u_i(\{r\}) \geq u_j(\{r\})$ for all $i, j \in A$ and $r \in R$. Such utilities capture practically relevant scenarios where utility is determined primarily by an agent's initial condition, rather than by individual preferences (e.g., aid-based applications).

We develop an adaptation of the classic round-robin algorithm for the setting with initial utilities that computes an allocation satisfying minimum-EF1-init. In our version, agents still select their most preferred available resource in rounds, but participation in each round is restricted to *active agents*, who pick sequentially according to a maintained *picking order* over all active agents. Initially, only agents with the lowest initial utility are active. An inactive agent i is activated as soon as all currently active agents have reached i's initial utility. Importantly, newly activated agents are inserted into the picking order after the agents that have already picked in the respective round, meaning that they get to pick directly after being activated. This ensures that each newly activated agent i prefers the resources picked by them over those picked, after i's activation, by any other agent that was active before them. While our algorithm is simple to state, the proof that it guarantees minimum-EF1-init is more involved, requiring a careful analysis of the position where newly added agents are inserted into the picking order.

[2] To rule out trivial edge cases, we assume that for every agent $i \in A$, there exists some resource $r \in R$ such that $u_i(\{r\}) > 0$, and for every resource $r \in R$, there exists some agent $i \in A$ such that $u_i(\{r\}) > 0$.

[3] This condition implies that an agent $i \in A$ with $b_i + u_i(X_i) < b_j$ does not envy an agent j under EF-init or EF1-init if $X_j = \emptyset$, as otherwise achieving an envy-free allocation would be impossible in settings with large initial utility disparities.

REFERENCES

[1] Martin Aleksandrov and Toby Walsh. 2020. Online fair division: A survey. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI '20)*. AAAI Press, 13557–13562.

[2] Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A. Voudouris, and Xiaowei Wu. 2023. Fair division of indivisible goods: Recent progress and open questions. *Artif. Intell.* 322 (2023), 103965.

[3] Haris Aziz, Bo Li, Hervé Moulin, and Xiaowei Wu. 2022. Algorithmic fair allocation of indivisible items: A survey and new questions. *SIGecom Exch.* 20, 1 (2022), 24–40.

[4] Siddharth Barman, Ganesh Ghalme, Shweta Jain, Pooja Kulkarni, and Shivika Narang. 2019. Fair division of indivisible goods among strategic agents. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems (AAMAS '19)*. IFAAMAS, 1811–1813.

[5] Catherine Barnard and Bob Hepple. 2000. Substantive equality. *The Cambridge Law Journal* 59, 3 (2000), 562–585.

[6] Hamsa Bastani, Kimon Drakopoulos, Vishal Gupta, Ioannis Vlachogiannis, Christos Hadjichristodoulou, Pagona Lagiou, Gkikas Magiorkinis, Dimitrios Paraskevis, and Sotirios Tsiodras. 2021. Efficient and targeted COVID-19 border testing via reinforcement learning. *Nature* 599, 7883 (2021), 108–113.

[7] Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. 2016. Fair allocation of indivisible goods. In *Handbook of Computational Social Choice*, Felix Brandt, Vincent Conitzer, Ulle Endriss, Jérôme Lang, and Ariel D. Procaccia (Eds.). Cambridge University Press, 284–310.

[8] Sylvain Bouveret and Michel Lemaître. 2016. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Auton. Agents Multi-Agent Syst.* 30, 2 (2016), 259–290.

[9] Mithun Chakraborty, Ayumi Igarashi, Warut Suksompong, and Yair Zick. 2021. Weighted envy-freeness in indivisible item allocation. *ACM Trans. Economics and Comput.* 9, 3 (2021), 18:1–18:39.

[10] Commission on Social Determinants of Health. 2008. *Closing the gap in a generation: Health equity through action on the social determinants of health.* World Health Organization, Geneva.

[11] Argyrios Deligkas, Eduard Eiben, Robert Ganian, Tiger-Lily Goldsmith, and Stavros D. Ioannidis. 2025. The complexity of extending fair allocations of indivisible goods. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI '25)*. AAAI Press, 13745–13753.

[12] Anmei Dong, Morris Siu-Yung Jong, and Ronnel B King. 2020. How does prior knowledge influence learning engagement? The mediating roles of cognitive load and help-seeking. *Frontiers in Psychology* 11 (2020), 591203.

[13] Gideon Elford. 2023. Equality of Opportunity. In *The Stanford Encyclopedia of Philosophy* (Fall 2023 ed.), Edward N. Zalta and Uri Nodelman (Eds.). Metaphysics Research Lab, Stanford University.

[14] Ezekiel J Emanuel and Govind Persad. 2023. The shared ethical framework to allocate scarce medical resources: A lesson from COVID-19. *The Lancet* 401, 10391 (2023), 1892–1902.

[15] Jonathan R. Goldman and Ariel D. Procaccia. 2014. Spliddit: Unleashing fair division algorithms. *SIGecom Exch.* 13, 2 (2014), 41–46.

[16] Moritz Hardt, Eric Price, and Nati Srebro. 2016. Equality of opportunity in supervised learning. In *Proceedings of the 30th Annual Conference on Neural Information Processing Systems (NIPS '16)*. 3315–3323.

[17] David Kurokawa, Ariel D. Procaccia, and Junxing Wang. 2018. Fair enough: Guaranteeing approximate Maximin shares. *J. ACM* 65, 2 (2018), 8:1–8:27.

[18] Haylee Lane, Mitchell Sarkies, Jennifer Martin, and Terry Haines. 2017. Equity in healthcare resource allocation decision making: A systematic review. *Social Science & Medicine* 175 (2017), 11–27.

[19] Shengxin Liu, Xinhang Lu, Mashbat Suzuki, and Toby Walsh. 2024. Mixed fair division: A survey. *J. Artif. Intell. Res.* 80 (2024), 1373–1406.

[20] Hervé Moulin. 2004. *Fair division and collective welfare.* MIT Press.

[21] Hervé Moulin. 2019. Fair division in the internet age. *Annual Review of Economics* 11, 1 (2019), 407–441.

[22] Vishwa Prakash HV, Ayumi Igarashi, and Rohit Vaish. 2025. Fair and efficient completion of indivisible goods. In *Proceedings of the 39th AAAI Conference on Artificial Intelligence (AAAI '25)*. AAAI Press, 14045–14053.

[23] Scottish Government. 2021. Fairer Scotland Duty: Guidance for public bodies.

[24] Donald J. Trump. 2025. Restoring equality of opportunity and meritocracy. Executive Order 14281.

[25] Shengwei Zhou, Rufan Bai, and Xiaowei Wu. 2023. Multi-agent online scheduling: MMS allocations for indivisible items. In *Proceedings of the 40th International Conference on Machine Learning (ICML '23) (Proceedings of Machine Learning Research, Vol. 202)*. PMLR, 42506–42516.