



HAL
open science

Constrained Serial Dictatorships Can Be Fair

Sylvain Bouveret, Hugo Gilbert, Jérôme Lang, Guillaume Mérroué

► **To cite this version:**

Sylvain Bouveret, Hugo Gilbert, Jérôme Lang, Guillaume Mérroué. Constrained Serial Dictatorships Can Be Fair. IJCAI 2025 - Thirty-Fourth International Joint Conference on Artificial Intelligence, Aug 2025, Montreal, Canada. pp.3762-3770, <10.24963/ijcai.2025/418>. <hal-05393240>

HAL Id: hal-05393240

<https://hal.science/hal-05393240v1>

Submitted on 2 Dec 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire HAL, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons CC BY 4.0 - Attribution - International License

Constrained Serial Dictatorships Can Be Fair

Sylvain Bouveret¹, Hugo Gilbert², Jérôme Lang² and Guillaume Méréoué³

¹Univ. Grenoble Alpes, CNRS, Grenoble INP, LIG, 38000 Grenoble, France

²Université Paris-Dauphine, Université PSL, CNRS, LAMSADE, 75016 Paris, France

³Université Côte d’Azur, Inria, CNRS, I3S, Templier 1, 06410 Biot, France

Abstract

When allocating indivisible items to agents, it is known that the only strategyproof mechanisms that satisfy a set of rather mild conditions are *constrained serial dictatorships*: given a fixed order over agents, at each step the designated agent chooses a given number of items (depending on her position in the sequence). Agents who come earlier in the sequence have a larger choice of items; however, this advantage can be compensated by a higher number of items received by those who come later. How to balance priority in the sequence and number of items received is a nontrivial question. We use a previous model, parameterized by a mapping from ranks to scores, a social welfare functional, and a distribution over preference profiles. For several meaningful choices of parameters, we show that the optimal sequence can be computed exactly in polynomial time or approximated using sampling. Our results hold for several probabilistic models on preference profiles, with an emphasis on the Plackett-Luce model. We conclude with experimental results showing how the optimal sequence is impacted by various parameters.

1 Introduction

In an ideal world, a mechanism for dividing a set of indivisible goods (or items, we use both terms interchangeably) should be at the same time efficient, fair, and insensitive to strategic behaviour. Now, strategyproofness is a very strong requirement that severely limits the choice of mechanisms. The question we address in this paper is, *how can we design strategyproof mechanisms while retaining an acceptable level of fairness and/or efficiency?*

It is known that *under mild conditions, the only strategyproof mechanisms are within the family of serial dictatorships* (although the landscape is less dramatic when there are only two agents, see our related work section). A standard serial dictatorship is defined by a permutation of the set of agents; at each step, the designated agent chooses all the items she likes from those that are still available. A *constrained serial dictatorship (CSD)*, also called quota serial

dictatorship, is similar except that at each step, the designated agent chooses a predefined number of items.

(Constrained or unconstrained) serial dictatorships are strategyproof and elicitation-free: they do not require to know the preferences of the agents, which are only revealed through their picking choices. This is a major property, as in many contexts, it is not realistic to hope eliciting all the agent’s preferences, either because it would be too cumbersome, or for privacy reasons. However, are they acceptable on efficiency and fairness grounds? Unconstrained serial dictatorships are clearly not: if the first agent likes all items then she will pick them all. Constrained serial dictatorships do better, at the price of the loss of Pareto-efficiency; but still, agents appearing early in the sequence have a much larger choice than those appearing late. This is patent in the case where there are as many items as agents, each agent being entitled to only one item, CSDs cannot do better than this: the first agent will get her preferred item, and the last agent will have no choice and might receive her least preferred item.

However, when there are more items than agents, and agents can receive several items, things become better, because the advantage towards agents who come early in the sequence can be compensated by a higher number of items received by those who come later. Suppose, as a simple example, that three items have to be assigned to two agents, A(nn) and B(ob). Assuming that Ann picks first, there are three CSDs: AAA (Ann picks all items), AAB (Ann picks two, Bob one), and ABB (Ann picks one, Bob two). It is intuitively clear that ABB is optimal, but how can optimality be defined? With four items, things are less clear: AAAA and AAAB are clearly less desirable than AABB and ABBB, but which of these two should we choose? And what if we have five agents and seventeen items?

To sum up: strategyproofness leaves us almost no choice but (constrained) serial dictatorship; some are intuitively better than others. What remains to be done is to *define formal optimality criteria for choosing between CSDs*, and to *compute optimal ones*. Our paper addresses these questions.

The first question was addressed by Bouveret and Lang [2011] and further examined by Kalinowski *et al.* [2013a] in the broader context of *picking sequences* (which generalize CSDs by allowing non-consecutive picks).

First, since the cardinal values of agents for the items are not known, Bouveret and Lang [2011] propose to estimate

them using a scoring vector shared by all agents: for any agent i , the value assigned to the item ranked at position j is a fixed score s_j , independent of i .

Next, the efficiency and fairness of a CSD can be assessed by computing the resulting *social welfare*, according to a selected social welfare functional such as egalitarian, Nash, or utilitarian¹ [d’Aspremont and Gevers, 2002].

Finally, estimating the *expected* social welfare requires assuming a probability distribution over the ordinal preferences of agents. These can be generated in various models, such as impartial culture, Mallows [Mallows, 1957], or Plackett-Luce [Luce, 1959; Plackett, 1975].

Given these components, one can associate an expected social welfare with any CSD, and thereby define optimal CSDs for various scoring vectors, social welfare functionals, and preference distributions.

For egalitarian social welfare, we provide a simple algorithm which returns an optimal CSD given that one can compute the expected utility obtained by an agent when a CSD is used. This algorithm makes it possible to compute an optimal (respectively, close to optimal) CSD when this expected utility is polynomial-time computable (respectively, can be approximately evaluated, e.g., by sampling). We also provide a dynamic programming algorithm that computes an optimal CSD for utilitarian, Nash or egalitarian social welfare under a specific condition, which is met when preferences are fully correlated, or when they are fully independent and follow the impartial culture or more generally the Plackett-Luce model.

Sections 2 and 3 discuss related work and present our model. Section 4 presents our algorithms for computing an optimal CSD. These algorithms assume the existence of an oracle which can compute or estimate the expected utility of a picker given a CSD. Section 5 designs such oracles under various model assumptions. Section 6 gives results for small values of n , and depicts and comments on the evolution of the optimal sequences when all criteria except one are fixed.

2 Related Work

Strategyproof allocation of indivisible goods. Various characterization theorems state that, under mild additional conditions, strategyproof allocation mechanisms all have a serial dictatorship flavour: with strict preferences over subsets, only serial dictatorships are strategyproof, neutral, and nonbossy [Svensson, 1999], whereas only sequential dictatorships (a generalization of serial dictatorship where the identity of the agent picking in position k depends on the items assigned to the agents in positions 1 to $k - 1$) are strategyproof, Pareto-efficient, and nonbossy [Pápai, 2001]. If preferences are quantity-monotonic (a bundle of larger cardinality is always preferred to one of lower cardinality) then a mechanism is strategyproof, nonbossy, Pareto-efficient and neutral if and only if it is a CSD (also called a quota serial dictatorship) [Pápai, 2000]. Similar characterizations hold replacing quantity-monotonic by lexicographic preferences [Hosseini and Larson, 2019; Hosseini *et al.*, 2021]. With standard monotonicity, only quasi-dictatorships remain, where

¹If our main objective is *fairness*, utilitarian social welfare may not fit well. We will see further that it is the case indeed.

only the first agent in the sequence is allowed to pick more than one item [Pápai, 2000]. Variants of these characterizations have been established by Ehlers and Klaus [2003], Bogomolnaia *et al.* [2005] and Hatfield [2009]. Ignoring Pareto-efficiency or neutrality opens the door to more complex strategyproof mechanisms; a full characterization in the two-agent case is given by Amanatidis *et al.* [2017]. Amanatidis *et al.* [2016] show that the CSD where all agents except the last one pick only one item is a $1/\lfloor \frac{n-m+2}{2} \rfloor$ -approximation to maxmin fair share. Weakening strategyproofness into non obvious manipulability opens the door for more possibilities [Psomas and Verma, 2022].

Nguyen *et al.* [2018] show that when agents have preferences over sets of items defined from preferences over single items by an extension principle, some scoring rules are strategyproof for some extension principles. Allowing randomized mechanisms offers more possibilities, but not much [Bu and Tao, 2024; Garg and Psomas, 2022; Hosseini and Larson, 2019; Kojima, 2009]. CSDs are also considered in chore allocation [Aziz *et al.*, 2019].

Picking sequences. Sequential allocation of indivisible goods, also known as picking sequences, originates from Kohler and Chandrasekaran [1971], with a game-theoretic study of the alternating sequence for two agents. Still for two agents, Brams and Taylor [2000] consider other particular sequences. Bouveret and Lang [2011] define a more general class of sequences, for any number of agents, and argue that sequences can be compared with respect to their expected social welfare, using a scoring vector and a prior distribution over profiles. Kalinowski *et al.* [2013a] show that computing the expected utility of a sequence is polynomial under full independence, and that strict alternation is optimal for two agents, utilitarian social welfare and Borda scoring. The manipulation of picking sequences is studied by Bouveret and Lang [2014], Tominaga *et al.* [2016] and Aziz *et al.* [2017]. Flammini and Gilbert [2020] and Xiao and Ling [2020] study the parameterized complexity of computing an optimal manipulation. Game-theoretic aspects of picking sequences are addressed by Kalinowski *et al.* [2013b]. Chakraborty *et al.* [2021] study picking sequences for agents with different entitlements. While all these works are oblivious to agent identities, Caragiannis and Rathi [2023] try to find an approximately optimum order of agents in a serial dictatorship with a limited number of queries.

Maximizing social welfare in allocation of indivisible goods. A classic way of guaranteeing a level of fairness and/or efficiency consists in finding an allocation *maximizing social welfare*, under the assumption that the input contains, for each agent, her utility function over all bundles of goods (usually assumed additive). Egalitarian social welfare places fairness above all, utilitarian social welfare cares about efficiency only, and Nash social welfare is considered as a sweet spot in-between. See [Amanatidis *et al.*, 2023; Aziz *et al.*, 2022; Bouveret *et al.*, 2016; Lang and Rothe, 2024] for surveys. These mechanisms are not strategyproof.

3 Preliminaries : The Model

Given $n \in \mathbb{N}^*$, we use $[n]$ to denote $\{1, \dots, n\}$ and $[n]_0$ to denote $\{0, 1, \dots, n\}$. Bold symbols represent vectors.

Let $\mathcal{A} = \{a_1, \dots, a_n\}$ be a set of n agents with a_i the i^{th} agent to intervene in the allocation process and $\mathcal{G} = \{g_1, \dots, g_m\}$ a set of m goods. A preference profile $\mathbf{P} = (\succ_{a_1}, \dots, \succ_{a_n})$ describes the preferences of the agents: \succ_a is a complete ranking that specifies the preferences of agent a over the goods in \mathcal{G} . We denote by $\text{rk}_{\mathbf{P}}^a(g)$, the rank of item g in the ranking of a , given profile \mathbf{P} . *The preference profile is hidden, and therefore not part of the input*: we will assume that rankings are drawn independently according to some probabilistic model, that we denote by Ψ .

Two well-known probabilistic models are the *Mallows* and *Plackett-Luce models* [Mallows, 1957; Luce, 1959; Plackett, 1975]:

- The *Mallows model* is parameterized by a dispersion parameter $\phi \in [0, 1]$ and a ranking μ . We denote this model by $\text{M11}_{\mu, \phi}$. In this model, the probability of a ranking r is proportional to $\phi^{d_{\text{KT}}(r, \mu)}$, with $d_{\text{KT}}(r, \mu)$, the Kendall-Tau distance between rankings r and μ , i.e., the number of pairs of items that are in a different order in the two rankings.
- The *Plackett-Luce (PL) model* is parameterized by a value vector $\nu = (\nu_1, \dots, \nu_m)$. Intuitively, $\nu_i > 0$ represents the social value of good g_i . In this model, which we denote by PL_{ν} , the probability of a ranking $r = g_{i_1} \succ g_{i_2} \succ \dots \succ g_{i_m}$ is:

$$\prod_{j=1}^m \frac{\nu_{i_j}}{\sum_{l=j}^m \nu_{i_l}}.$$

The Plackett-Luce model has proven particularly good for learning a preference relation over a set of items (a.k.a. label ranking) [Cheng *et al.*, 2010] so it fits particularly well here.

These models generalize the two following sub-cases:

- *Impartial Culture*, denoted by IC, in which each preference ranking is drawn u.a.r. from the set of all possible rankings. Impartial culture is obtained when $\phi = 1$ for the Mallows model and when all values in ν are equal for the Plackett-Luce model.
- The *Full Correlation* case, denoted by FC stipulates that all agents have exactly the same preference ranking. Full correlation is obtained when $\phi = 0$ for the Mallows model (and also as the limit of Plackett-Luce models $\nu^M = (M^{m-1}, \dots, M, 1)$ when $M \rightarrow \infty$).

In the sequel, we obtain different results for $\Psi \in \{\text{FC}, \text{IC}, \text{M11}_{\mu, \phi}, \text{PL}_{\nu}\}$.

The items are allocated to the different agents according to a CSD: given a vector $\mathbf{k} = (k_1, \dots, k_n)$ of n non-negative integers, agent a_1 will first pick k_1 goods, then a_2 will pick k_2 goods within the remaining ones, and so on until a_n picks k_n items. In most cases, we will consider complete CSDs, in the sense that $\sum_{i=1}^n k_i = m$. However, we may also consider incomplete CSDs such that $\sum_{i=1}^n k_i < m$. We assume

that agents behave greedily by choosing their preferred goods within the remaining ones. This sequential process leads to an allocation that we denote by $\pi_{\mathbf{P}}^{\mathbf{k}}$. More formally, $\pi_{\mathbf{P}}^{\mathbf{k}}$ is a function such that $\pi_{\mathbf{P}}^{\mathbf{k}}(a)$ is the set of goods that agent a has obtained at the end of the sequential allocation process, given preference profile \mathbf{P} and vector \mathbf{k} .

The utility of an agent for obtaining an item i will be derived using a scoring vector. Stated otherwise, there is a vector $\mathbf{s} = (s_1, \dots, s_m) \in \mathbb{Q}^{+m}$ such that $s_i \geq s_{i+1}$ for all $i \in [m-1]$. The value received by an agent for obtaining her j^{th} preferred item is s_j . Different scoring vectors can be considered. An important example is the *Borda* scoring vector, where $s_i = m - i + 1$. Using scores as a proxy for utilities is classic in social choice: this is exactly how positional scoring voting rules (e.g., the Borda rule) are defined, and they are also used in fair division settings [Baumeister *et al.*, 2017; Brams *et al.*, 2003; Darmann and Klamler, 2016].

We denote by $U_{\mathbf{P}}^{\mathbf{k}}(a) = \sum_{g \in \pi_{\mathbf{P}}^{\mathbf{k}}(a)} s_{\text{rk}_{\mathbf{P}}^a(g)}$ the utility obtained by a when receiving $\pi_{\mathbf{P}}^{\mathbf{k}}(a)$ and by $EU_{\Psi}^{\mathbf{k}}(a) = \mathbb{E}_{\mathbf{P} \sim \Psi}[U_{\mathbf{P}}^{\mathbf{k}}(a)]$ her expected utility given model Ψ . This assumes that agents have *additive* preferences, which is very common in fair division. The utilitarian social welfare (USW) $W_{\Psi}^U(\mathbf{k})$, egalitarian social welfare (ESW) $SW_{\Psi}^E(\mathbf{k})$, and Nash social welfare (NSW) $SW_{\Psi}^N(\mathbf{k})$ are then defined by:

$$SW_{\Psi}^U(\mathbf{k}) = \sum_{a \in \mathcal{A}} EU_{\Psi}^{\mathbf{k}}(a), \quad SW_{\Psi}^E(\mathbf{k}) = \min_{a \in \mathcal{A}} EU_{\Psi}^{\mathbf{k}}(a),$$

$$SW_{\Psi}^N(\mathbf{k}) = \prod_{a \in \mathcal{A}} EU_{\Psi}^{\mathbf{k}}(a).$$

Note that our social welfare notions are meant *ex ante*, i.e., we define them on the expected utility values of the agents. This is different from the notion of *ex post* social welfare which considers the utility of the agents once the profile \mathbf{P} issued from Ψ is determined.

Our objective is to study the following class of optimization problems $\text{OptSD-}\Psi\text{-}x$ with $x \in \{U, E, N\}$.

OPTSD- Ψ - x
<p>Input: A number n of agents, a number m of goods, and a scoring vector \mathbf{s}.</p> <p>Find: A vector $\mathbf{k} = (k_1, \dots, k_n)$ of n non-negative integers with $\sum_{i=1}^n k_i = m$ maximizing $SW_{\Psi}^x(\mathbf{k})$.</p>

The following easy observation will be useful:

Observation 1. *For given n and m , the number of vectors $\mathbf{k} = (k_1, \dots, k_n)$ such that $\sum_{i=1}^n k_i = m$ equals $\binom{n+m-1}{n-1}$.*

From this observation, we can deduce that the number of potential vectors is lower-bounded by $\frac{m^{n-1}}{(n-1)!}$. This number does not take into account a natural further assumption that the optimal sequence is *non-decreasing*, that is, that $k_1 \leq k_2 \leq \dots \leq k_n$. We will see further that this assumption holds for ESW (under a mild condition), *but not* for USW. When the assumption holds, we can restrict the search to non-decreasing vectors; their number is the number of integer partitions m into n numbers; it is still exponentially large, but no closed form expression is known.

4 Computing an Optimal CSD

We now investigate the problem $\text{OptSD-}\Psi\text{-}x$ with $x \in \{U, E, N\}$. All algorithms in this Section assume access to an oracle algorithm $\mathcal{T}_\Psi(\mathbf{k}, i)$ computing $EU_\Psi^{\mathbf{k}}(a_i)$ in time $K(n, m, s)$. The computation of expected utilities of agents for various models will be addressed in Section 5.

We start by a positive result for Egalitarian Social Welfare: the optimal CSD can be computed by the greedy-like Algorithm 1. $\text{Completion}(\mathbf{k})$ denotes, for any partial CSD \mathbf{k} , the complete CSD such that $\text{Completion}(\mathbf{k})_i = k_i$ for $i \in [n-1]$ and $\text{Completion}(\mathbf{k})_n = m - \sum_{i \in [n-1]} k_i$. In informal terms, \mathbf{k} is completed by giving all remaining goods to the last agent.

Algorithm 1 GreedyESW

Require: the number of agents n , the number of goods m , the scoring vector s , the oracle algorithm \mathcal{T}_Ψ

- 1: $\mathbf{k} \leftarrow (0, \dots, 0)$ # empty CSD
- 2: $\text{max_k}, \text{max_esw} \leftarrow \mathbf{k}, 0$
- 3: **for** $t = 1$ to m **do**
- 4: $i \in \text{argmin}_{i \in [n]} EU_\Psi^{\mathbf{k}}(a_i)$
- 5: $k_i \leftarrow k_i + 1$
- 6: **if** $SW_\Psi^E(\mathbf{k}) > \text{max_esw}$ **then**
- 7: $\text{max_k}, \text{max_esw} \leftarrow \mathbf{k}, SW_\Psi^E(\mathbf{k})$
- 8: **end if**
- 9: **end for**
- 10: **return** $\text{Completion}(\text{max_k})$;

At line 1, we start with an empty CSD, that we will modify in a greedy fashion. In the **for** loop (lines 3-9), we identify an agent with minimal expected utility (line 4) and increment the number of goods that she gets (line 5). The CSD that is returned is not necessarily this CSD \mathbf{k} . During the algorithm, we keep in variables max_esw and max_k , the maximum ESW found so far and the corresponding (partial) CSD. The algorithm returns max_k completed by giving all remaining goods to the last agent (line 10). The completion step is not really necessary (the partial sequence obtained at line 9 already has maximum expected egalitarian social welfare); its role is to ensure that no good is left unallocated. The reason why one needs the test at line 6 is that letting the currently least happy agent pick one more good may decrease the ESW, as can be seen on the following example.

Example 1. Let $n = 2$, $m = 5$, $s = (50, 10, 4, 2, 1)$, and $\Psi = \text{IC}$. We show below the partial CSDs obtained in each iteration t together with the expected utilities of both agents (they can be computed easily, as we will see in Section 5) and the values of i_t and max_esw .

t	\mathbf{k}	max_k	$EU_\Psi^{\mathbf{k}}(a_1)$	$EU_\Psi^{\mathbf{k}}(a_2)$	i^t	max_esw
1	(0, 0)	(0, 0)	0	0	1	0
2	(1, 0)	(0, 0)	50	0	2	0
3	(1, 1)	(1, 1)	50	42	2	42
4	(1, 2)	(1, 2)	50	49.6	2	49.6
5	(1, 3)	(1, 3)	50	52.4	1	50
6	(2, 3)	(1, 3)	60	40.2	1	50

At iteration 5, the least happy agent is a_1 ; however, letting a_1 pick one more good, that is, $\mathbf{k} = (2, 3)$ gives $EU_\Psi^{\mathbf{k}}(a_1) =$

60 and $EU_\Psi^{\mathbf{k}}(a_2) = 40.2$ (iteration 6), decreasing the currently optimal expected ESW. Therefore, max_k is not replaced by $\mathbf{k} = (2, 3)$ at line 6 of the algorithm. The algorithm returns $\text{Completion}(\text{max_k}) = (1, 4)$ (with expected utilities 50 and 53.6) with the remaining good given to a_2 .

Proposition 1. Algorithm 1 returns a CSD \mathbf{k} maximizing $SW_\Psi^E(\mathbf{k})$, solving problem $\text{OptSD-}\Psi\text{-}E$, in time $O(nmK(n, m, s))$.

The proof is based on the following lemma:

Lemma 1. Let $\hat{\mathbf{k}}$ be a CSD. Let max_esw^t , \mathbf{k}^t and i^t denote max_esw , \mathbf{k} and i after line 4 of iteration t of the **for** loop in Algorithm 1. For all t , a necessary condition for $SW_\Psi^E(\hat{\mathbf{k}}) > \text{max_esw}^t$ is that $\hat{k}_j \geq k_j^t$ for all $j \in [n]$, and $\hat{k}_{i^t} > k_{i^t}^t$.

Proof. By induction. At iteration 1, the claim is obvious. Assume that the claim holds for iteration t , and let $\hat{\mathbf{k}}$ be a CSD such that $SW_\Psi^E(\hat{\mathbf{k}}) > \text{max_esw}^{t+1}$. Then obviously $SW_\Psi^E(\hat{\mathbf{k}}) > \text{max_esw}^t$ as $\text{max_esw}^{t+1} \geq \text{max_esw}^t$. Because the condition holds for iteration t and by construction of \mathbf{k}^{t+1} we have that $\hat{k}_j \geq k_j^{t+1}$ for all $j \in [n]$. Now suppose that $\hat{k}_{i^{t+1}} = k_{i^{t+1}}^{t+1}$. In that case, $SW_\Psi^E(\hat{\mathbf{k}}) \leq EU^{\mathbf{k}^{t+1}}(a_{i^{t+1}}) \leq \text{max_esw}^{t+1}$, a contradiction with the induction hypothesis. The first inequality is due to the fact that $a_{i^{t+1}}$ will get the same number of goods in $\hat{\mathbf{k}}$ and \mathbf{k}^{t+1} while the agents picking before her will get at least as many goods in $\hat{\mathbf{k}}$ than in \mathbf{k}^{t+1} . The second inequality is due to the definition of i^{t+1} . \square

Proof of Proposition 1. Suppose that there exists a CSD $\hat{\mathbf{k}}$ such that $SW_\Psi^E(\hat{\mathbf{k}}) > \text{max_esw}$. Lemma 1 applied at iteration $t = m$ implies that each agent receives more objects with $\hat{\mathbf{k}}$ than with the greedily constructed complete CSD \mathbf{k} obtained at the end of the **for** loop. As they both have m objects to allocate, they must be equal. This is a contradiction of the hypothesis as $\text{max_esw} \geq SW_\Psi^E(\mathbf{k})$. \square

We now go beyond ESW. For USW and NSW, we do not know of an efficient algorithm which would work for any distribution. A general approach could be to sample a large but hopefully reasonable number of preference profiles from Ψ and find a CSD with maximal social welfare considering the average utility of each agent. Yet, we prove in [Bouveret et al., 2025, Appendix B] that such an approach leads to an NP-hard problem for USW.

However, provided the distribution satisfies a natural condition, a CSD maximizing utilitarian and Nash social welfare can be computed by dynamic programming. This condition on Ψ states that $EU_\Psi^{\mathbf{k}}(a)$ only depends on the number of items picked by a , and the number of items that have been picked before a , but not on the number of agents who have picked before and how many items they have picked each.

Definition 1. A distribution Ψ satisfies **prefix independence** if for any sequence \mathbf{k} and $i \in [n]$, if a is the i^{th} picker in \mathbf{k} , then $EU_\Psi^{\mathbf{k}}(a)$ only depends on (1) $\kappa = k_i$, the number of goods that she picks, and (2) $\tau = \sum_{j=1}^{i-1} k_j$, the number of goods that have been picked before she starts picking.

Under prefix independence, the utility that agent a gets when picking κ goods while τ have already been picked, $\text{eu}(\kappa, \tau)$, is well-defined, and is exactly equal to $EU_{\Psi}^{\kappa}(a)$ when a is the i^{th} picker $\kappa = k_i$ and $\tau = \sum_{j=1}^{i-1} k_j$.

For pedagogical purposes, let us first focus on maximising USW. When prefix independence is met, one can use the following dynamic programming equations:

$$\begin{aligned} F(i, \tau) &= \max_{\kappa \in [m-\tau]_0} (\text{eu}(\kappa, \tau) + F(i+1, \tau + \kappa)), \\ &\quad \forall i, \tau \in [n-1] \times [m]_0, \quad (1) \\ F(n, \tau) &= \text{eu}(m - \tau, \tau), \forall \tau \in [m]_0, \end{aligned}$$

where $F(i, \tau)$ corresponds to the maximum USW that can be obtained by agents $\{a_i, a_{i+1}, \dots, a_n\}$ in the situation in which τ goods have already been allocated and we allocate the $m - \tau$ remaining goods to them. Of course the optimal value is given by $F(1, 0)$.

The other problems can be solved similarly. For problem OptSD- Ψ - E (resp. OptSD- Ψ - N), one should adapt Equation 1 by replacing the sum operation between $\text{eu}(\kappa, \tau)$ and $F(i+1, \tau + \kappa)$ by a min (resp. multiplication) operation.

Proposition 2. *If Ψ satisfies prefix independence, problems OptSD- Ψ - U , OptSD- Ψ - E and OptSD- Ψ - N can be solved in $O(nm^2 K(n, m, s))$ time.*

We conclude by giving a structural property satisfied by an optimal CSD for ESW when prefix independence holds. We will see that such property does not necessarily hold for USW (see [Bouvet et al., 2025, Appendix B] and Section 6).

Proposition 3. *Under prefix independence, there exists an optimal solution to OptSD- Ψ - E which is non-decreasing, i.e., in which the earlier an agent picks, the less goods she gets.*

5 Computing the Expected Utility of an Agent

In this section, we address the computation of $EU_{\Psi}^{\kappa}(a)$. Prefix independence again plays a crucial role: when it is satisfied, $EU_{\Psi}^{\kappa}(a)$ only depends on the number of items picked by a , and the number of items that have been picked before a , but not on the number of agents who have picked before and how many items they have picked each. We first investigate which of our different probabilistic models satisfy it.

Proposition 4. $\Psi \in \{\text{FC}, \text{IC}\}$ satisfy prefix independence.

Proof. Consider a situation where an agent starts picking while τ goods have previously been picked. When $\Psi = \text{FC}$ or $\Psi = \text{IC}$, the probability distribution on the set S of goods that have previously been picked only depends on τ : for $\Psi = \text{FC}$, this probability distribution assigns probability 1 to the set composed of the τ (unanimously) most preferred goods; for $\Psi = \text{IC}$, this probability distribution assigns equal probability to all sets of size τ and 0 to others. Note that, given the set S , the utility that the agents get is then determined by the number of goods she picks. \square

More interestingly, the PL_{ν} model, which generalizes FC and IC, also satisfies prefix independence.

Proposition 5. $\Psi = \text{PL}_{\nu}$ satisfies prefix independence.

To reason on the Plackett-Luce model, one can use the *vase model metaphor* [Silverberg, 1980]. Consider a vase filled with m types of balls, the proportion of balls of type j being $f(j) = \frac{\nu_j}{\sum_{i=1}^m \nu_i}$. The ranking is then generated by the following sequential process. At each stage, a ball is taken from the vase such that a ball of type j is chosen with probability $f(j)$. If the ball is of a different type than the ones previously picked, it yields the next good in the ranking. In either case, the ball is put back in the vase and the process continues. Using this metaphor, one can prove the following lemma (the formal proof can be found in [Bouvet et al., 2025, Appendix C]).

Lemma 2. *Let $I = (i_1, \dots, i_q)$ be a sequence of q different indices in $[m]$. Consider the following two cases:*

- i) *Agent a_1 picks q goods;*
- ii) *Agent a_1 picks q_1 goods and agent a_2 picks q_2 goods with $q_1 + q_2 = q$.*

For the PL model, the probability that for all $t \in [q]$, g_{i_t} is picked at timestep t is the same in cases i and ii.

Proof of Proposition 5. We recall that the preference rankings of the agents are drawn independently from PL_{ν} . Using Lemma 2 and a simple induction argument, we get that the probability of a specific sequence of q consecutive picks is the same regardless of whether they were picked by one, two or more agents. This entails that the probability distribution on the set S of goods that have been picked after τ timesteps only depends on the value of τ . Hence, the expected utility that an agent gets when choosing κ goods once τ have been picked only depends on the values of κ and τ . \square

Unfortunately, things are different for the Mallows model.

Proposition 6. *There exists $\phi \in (0, 1)$ and a ranking μ such that $\Psi = \text{M11}_{\phi, \mu}$ does not satisfy prefix independence.*

This holds even for 3 agents and 3 goods. See [Bouvet et al., 2025, Appendix C] for the proof.

Computation of $EU_{\Psi}^{\kappa}(a)$. Under prefix independence, we show how to compute $\text{eu}(\kappa, \tau)$ efficiently, starting by FC.

Proposition 7. *If $\Psi = \text{FC}$, $\text{eu}(\kappa, \tau) = \sum_{i=\tau+1}^{\tau+\kappa} s_i$. All values $\text{eu}(\kappa, \tau)$ can be computed in time $O(m^2)$ with the recursive formula $\text{eu}(\kappa, \tau) = \text{eu}(\kappa - 1, \tau) + s_{\kappa+\tau}$.*

We then turn to $\Psi = \text{IC}$, and show that the values $\text{eu}(\kappa, \tau)$ can be computed using a recursive formula. Let $T(j, \kappa, \tau)$ denote the utility that an agent can get if she can pick κ goods within the ones of rank in $\{j, \dots, m\}$, given that τ of these goods have been picked by preceding agents. Then, it is clear that we have:

$$\text{eu}(\kappa, \tau) = T(1, \kappa, \tau), \forall \kappa, \tau \in [m]_0 \times [m - \kappa]_0$$

The key point is that there is a probability $\frac{\tau}{m-j+1}$ that this good is one of the τ goods that have previously been picked, otherwise with a probability of $1 - \frac{\tau}{m-j+1}$ the good of rank j is free and the agent will pick this good. In both cases, we move to goods of rank in $\{j+1, \dots, m\}$. In the first case, we decrease τ by 1 as we have identified one of the goods

already picked within the ones of rank j to m . In the second case, we decrease κ by one as the agent has picked a good. Hence, $\text{eu}(\kappa, \tau)$ can be computed by the following formula:

$$T(j, \kappa, \tau) = \frac{\tau}{m-j+1} T(j+1, \kappa, \tau-1) + \left(1 - \frac{\tau}{m-j+1}\right) (s_j + T(j+1, \kappa-1, \tau))$$

$$\forall j, \kappa, \tau \in [m-1] \times [m-j+1] \times [m-j-\kappa+1], \quad (2)$$

with the following base cases:

$$T(j, 0, \tau) = 0, \forall j, \tau \in [m] \times [m-j+1]_0$$

$$T(j, \kappa, 0) = \sum_{j \leq i < j+\kappa} s_i, \forall j, \kappa \in [m] \times [m-j+1]_0.$$

By computing all values $T(j, \kappa, \tau)$ in $O(m^3)$ operations, we obtain the following result.

Proposition 8. *If $\Psi = \text{IC}$, then all values $\text{eu}(\kappa, \tau)$ can be computed in time $O(m^3)$ by using Equation 2.*

Propositions 2, 7, and 8 imply that $\text{OptSD-}\Psi$ - x for $x \in \{U, E, N\}$ can be solved in polynomial time for $\Psi = \text{FC}$ and $\Psi = \text{IC}$, in $O(nm^2)$ for $\Psi = \text{FC}$ and $O(m^2 \max(n, m))$ for $\Psi = \text{IC}$, by precomputing all values $\text{eu}(\kappa, \tau)$ before running the dynamic programming algorithm.

For $\Psi \notin \{\text{IC}, \text{FC}\}$, one can still use GreedyESW and the dynamic programming algorithm with values $EU_{\Psi}^{\kappa}(a)$ approximated by sampling, providing close-to optimal CSDs: the returned CSD is optimal with expected utility values replaced by their approximate values.²

For the general PL_{ν} model beyond FC and IC, we do not know whether values $\text{eu}(\kappa, \tau)$ can be computed exactly in polynomial time; however, they can be efficiently approximated by sampling preference profiles from Ψ and averaging the utility values obtained on the samples, with approximation guarantees from Hoeffding’s (1963) inequality.

To present this guarantee, let $u_{\kappa, \tau}(\mathbf{P}, \mathbf{s})$ denote the utility value obtained by the second picker when she picks her κ preferred (available) goods, while the first picker has picked her τ preferred ones, given the preference profile \mathbf{P} .

Proposition 9. *Let $\epsilon > 0$ and $\delta \in (0, 1)$ two fixed values, and Υ an upper bound on values $\text{eu}(\kappa, \tau)$ (e.g., $\sum_{i=1}^m s_i$).*

Let $\bar{\text{eu}}_{\kappa, \tau}$ be the value computed by averaging the values $u_{\kappa, \tau}(\mathbf{P}_i, \mathbf{s})$ over N preference profiles \mathbf{P}_i sampled independently from Ψ . If $N \geq (\Upsilon^2 \ln(2m^2/\delta))/2\epsilon^2$, then it holds with probability $1 - \delta$ that:

$$|\text{eu}(\kappa, \tau) - \bar{\text{eu}}_{\kappa, \tau}| \leq \epsilon, \forall \kappa, \tau \in [m] \times [m - \kappa].$$

Moreover, we show that these utility values can be computed exactly in time FPT (Fixed-Parameter Tractable) with respect to parameter m and XP (slice-wise polynomial) with respect to ρ , where ρ is the number of distinct values in ν . This seems particularly appealing as goods may often be partitioned in categories. When $\rho = 1$, all goods are in the same category and we obtain the IC model; when ρ equals 2 or 3 we obtain categories {high value, low value} or {high value, medium value, low value}.

²Some mild monotonicity conditions are required on the approximated $EU_{\Psi}^{\kappa}(a)$ values for the validity of Algorithm 1.

Proposition 10. *If $\Psi = \text{PL}_{\nu}$, then all values $\text{eu}(\kappa, \tau)$ can be computed in time $O(4^m \text{Poly}(m))$.*

Proposition 11. *If $\Psi = \text{PL}_{\nu}$, then all values $\text{eu}(\kappa, \tau)$ can be computed in time $O(m^{2\rho} \text{Poly}(m))$.*

6 Numerical Tests

We performed several experiments to explore the properties of the CSDs obtained by maximizing either USW, NSW or ESW. More precisely, we explored the impact of increasing one of the parameters, all other parameters being fixed.

Code and an interactive demo are available at <https://github.com/GuillaumeMeroue/CSD-can-be-Fair> and <https://guillaumemeroue.github.io/IJCAI25>.

Impact of the number of goods. Figure 1 displays the proportion of utility (left-hand side) and goods (right-hand side) obtained for $n = 5$ and increasing the number of goods m from 5 to 300 in steps of 5. To generate both figures, the IC model and the Borda scoring vector were used and we optimized either USW, ESW or NSW.³

Several comments can be made. First, as expected, in the egalitarian case (middle of Figure 1), we observe that as m increases, the distribution of utility received by each agent converges towards equal share.⁴ In order to achieve this, the agents who arrive later in the sequence receive more items.

Second, with Borda and utilitarianism, the first agent in the sequence may pick more items than others (plots on top of Figure 1). More generally, on this plot, the utility of an agent seems to decrease with the position in the sequence.

Finally, for the Borda scoring vector, egalitarian and Nash social welfare objectives tend to give similar results.

Impact of correlation. We explore the impact of correlation, through the parameters ϕ and ν of models PL_{ν} and $\text{M11}_{\phi, \mu}$. We use the Borda scoring vector and maximize ESW. To run Algorithm 1, we approximate the expected utility values of the agents by sampling 10000 preference profiles from PL_{ν} and from $\text{M11}_{\phi, \mu}$ with the PrefSampling library [Boehmer *et al.*, 2024]. Figure 2 displays the utility value (plots at the bottom) and the number of items (top) received by each of 5 agents with $m = 70$ goods, for models PL_{ν} (right) and $\text{M11}_{\phi, \mu}$ (left). In the former model, we use $\nu^x = (x^m, x^{m-1}, \dots, x^1)$ and decrease x from 1.5 (which already yields very correlated preference profiles similar to FC) to 1 (IC) in steps of 0.01. In the latter model, we increase ϕ from 0 (FC) to 1 (IC) in steps of 0.02.

Several comments are in order. First, as shown in Figure 2, the utility of all agents (and thus their total) increases as x decreases or ϕ increases. This is because preferences diverge when approaching IC, allowing even late agents to obtain some of their preferred items. Second, the number of goods received by the first agents in the CSD increases while it decreases for the last ones. Indeed, as these latter agents can

³In Figures 1, and 2, the 1st picker corresponds to the color *blue* (at the bottom of each plot) while the 5th and last agent to pick corresponds to the color *purple* (at the top of each plot). Moreover, note that the values plotted are in fact cumulative values.

⁴This observation is proven formally in [Bouveret *et al.*, 2025, Appendix D].

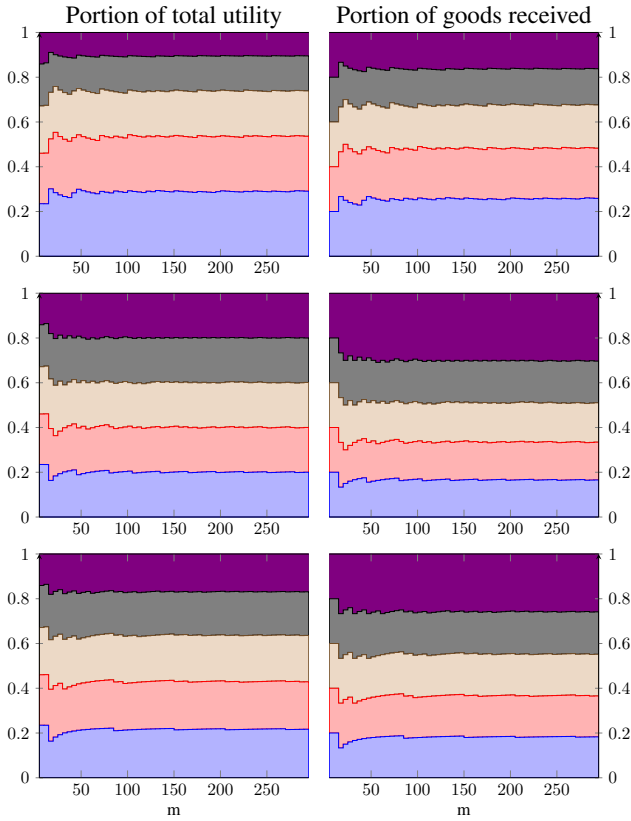


Figure 1: Portion of total utility (plots on the left) and of goods (right) received by each of 5 agents with m increasing from 5 to 300 in steps of 5. Maximizing USW (plots at the top), ESW (middle), or NSW (bottom), using Borda scoring vector and IC.

receive more preferred goods, the CSD needs less to compensate by giving them a high number of goods (recall that we optimize ESW). Third, we notice that both models PL_{ν} and $M11_{\phi,\mu}$ yield very similar plots as we decrease the level of correlation.

7 Discussion

The practical use of our setting raises a few questions.

First, we need to choose a distribution. The choice has to be tailored to the domain at hand, and distributions can be learnt using some preference learning models and techniques. If computation time is an important issue then it is wise to learn a Plackett-Luce model [Cheng *et al.*, 2010].

Second, we need to choose a scoring vector as a proxy for agents’ valuations over items. Again, this depends on the specific domain at hand. For each context, the scores can be estimated by an experiment where subjects are presented with a list of items to elicit their valuations; see [Bouveret *et al.*, 2025, Appendix E].

Third, we need to choose a social welfare functional. We have seen that, unsurprisingly, utilitarianism may lead to clearly unfair solutions and should be used only with care. As usual, egalitarianism may lead to a loss of efficiency, but is easier to compute or approximate; Nash is a good trade-off

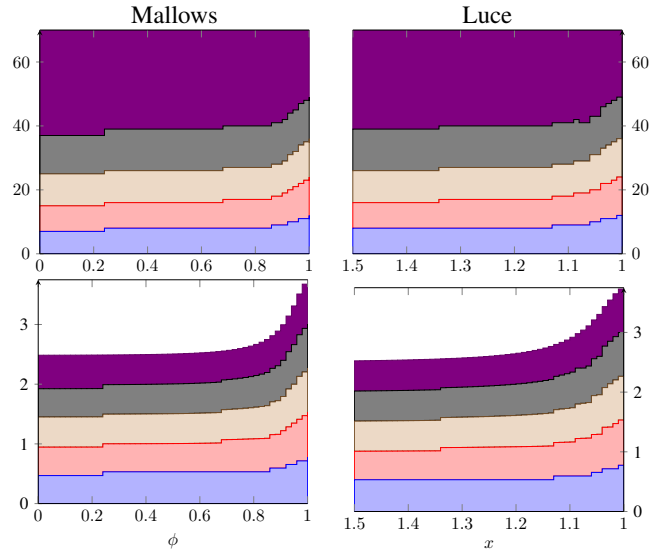


Figure 2: Number of goods received per agent (top); expected utility value (in thousands) per agent (bottom) as a function of ϕ for $M11_{\phi,\mu}$ (left) and x for PL_{ν^x} (right). Maximizing ESW, Borda scoring vector, $n = 5, m = 70$.

(see [Caragiannis *et al.*, 2019] for a manifesto towards using Nash social welfare in fair division) but is hard to compute if the distribution does not satisfy prefix independence.

Four, once a CSD is found, it is anonymous: for instance, with two agents, if the output is $(1, 2)$, it does not say who should start picking. Assigning agents to positions in the sequence has no impact on *ex ante* social welfare, but it may have an impact on *ex post* social welfare ([Bouveret *et al.*, 2025, Appendix F]).

8 Conclusion

Our main messages are: (1) imposing strategyproofness does not leave much choice beyond constrained serial dictatorships; (2) some constrained serial dictatorships are fairer than others; (3) their efficiency and fairness can be measured by expected social welfare, defined by a scoring vector, a distribution over profiles, and a social welfare functional; (4) depending on the social welfare functional and the distribution, the optimal sequence can be polynomial-time computable, efficiently approximated by sampling, or hard to approximate by sampling. The following table summarizes the results obtained. PI means that prefix independence is satisfied, poly means “polynomial-time computable”, and approx means “efficiently approximable by sampling”.

Ψ	PI	$EU_{\Psi}^k(a_i)$	Egal	Nash	Uti
FC	yes	poly	poly	poly	poly
IC	yes	poly	poly	poly	poly
PL_{ν}	yes	approx	approx	approx	approx
$M11_{\phi,\mu}$	no	approx	approx	?	?

If items were bads (e.g., chores) instead of goods, a similar methodology would work, with values in the scoring vector representing costs. Of course, agents coming first in the sequence should now take *more* items than those coming later.

Acknowledgements

This work has been supported by the French government, through the 3IA Côte d’Azur Investments project managed by the National Research Agency (ANR) with the reference number ANR-23-IACL-0001, through the France 2030 program, with the reference ANR-23-IACL-0008, and by the CITIZENS project, with the reference ANR-22-CE26-0019.

References

- [Amanatidis *et al.*, 2016] Georgios Amanatidis, Georgios Birmpas, and Evangelos Markakis. On truthful mechanisms for maximin share allocations. In Subbarao Kambhampati, editor, *Proceedings of the 25th International Joint Conference on Artificial Intelligence, IJCAI 2016, New York, NY, USA, 9-15 July 2016*, pages 31–37. IJCAI/AAAI Press, 2016.
- [Amanatidis *et al.*, 2017] Georgios Amanatidis, Georgios Birmpas, George Christodoulou, and Evangelos Markakis. Truthful allocation mechanisms without payments: Characterization and implications on fairness. In *Proceedings of the 2017 ACM Conference on Economics and Computation, EC ’17, Cambridge, MA, USA, June 26-30, 2017*, pages 545–562. ACM, 2017.
- [Amanatidis *et al.*, 2023] Georgios Amanatidis, Haris Aziz, Georgios Birmpas, Aris Filos-Ratsikas, Bo Li, Hervé Moulin, Alexandros A. Voudouris, and Xiaowei Wu. Fair division of indivisible goods: Recent progress and open questions. *Artif. Intell.*, 322:103965, 2023.
- [Aziz *et al.*, 2017] Haris Aziz, Sylvain Bouveret, Jérôme Lang, and Simon Mackenzie. Complexity of manipulating sequential allocation. In *Proceedings of the 31st AAAI Conference on Artificial Intelligence, San Francisco, California, USA, February 4-9, 2017*, pages 328–334. AAAI Press, 2017.
- [Aziz *et al.*, 2019] Haris Aziz, Bo Li, and Xiaowei Wu. Strategyproof and approximately maximin fair share allocation of chores. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence, IJCAI 2019, Macao, China, August 10-16, 2019*, pages 60–66. ijcai.org, 2019.
- [Aziz *et al.*, 2022] Haris Aziz, Bo Li, Hervé Moulin, and Xiaowei Wu. Algorithmic fair allocation of indivisible items: A survey and new questions. *CoRR*, abs/2202.08713, 2022.
- [Baumeister *et al.*, 2017] Dorothea Baumeister, Sylvain Bouveret, Jérôme Lang, Nhan-Tam Nguyen, Trung Thanh Nguyen, Jörg Rothe, and Abdallah Saffidine. Positional scoring-based allocation of indivisible goods. *Auton. Agents Multi Agent Syst.*, 31(3):628–655, 2017.
- [Boehmer *et al.*, 2024] Niclas Boehmer, Piotr Faliszewski, Łukasz Janeczko, Andrzej Kaczmarczyk, Grzegorz Lisowski, Grzegorz Pierczyński, Simon Rey, Dariusz Stolicki, Stanisław Szufa, and Tomasz Was. Guide to numerical experiments on elections in computational social choice. *arXiv preprint arXiv:2402.11765*, 2024.
- [Bogomolnaia *et al.*, 2005] Anna Bogomolnaia, Rajat Deb, and Lars Ehlers. Strategy-proof assignment on the full preference domain. *J. Econ. Theory*, 123(2):161–186, 2005.
- [Bouveret and Lang, 2011] Sylvain Bouveret and Jérôme Lang. A general elicitation-free protocol for allocating indivisible goods. In *Proceedings of the 22nd International Joint Conference on Artificial Intelligence, IJCAI 2011, Barcelona, Catalonia, Spain, July 16-22, 2011*, pages 73–78. IJCAI/AAAI Press, 2011.
- [Bouveret and Lang, 2014] Sylvain Bouveret and Jérôme Lang. Manipulating picking sequences. In *Proceedings of the 21st European Conference on Artificial Intelligence, ECAI 2014, Prague, Czech Republic, August 18-22, 2014*, volume 263 of *Frontiers in Artificial Intelligence and Applications*, pages 141–146. IOS Press, 2014.
- [Bouveret *et al.*, 2016] Sylvain Bouveret, Yann Chevaleyre, and Nicolas Maudet. Fair allocation of indivisible goods. In *Handbook of Computational Social Choice*, pages 284–310. Cambridge University Press, 2016.
- [Bouveret *et al.*, 2025] Sylvain Bouveret, Hugo Gilbert, Jérôme Lang, and Guillaume Méroué. Constrained serial dictatorships can be fair. *CoRR*, abs/2301.06086, 2025.
- [Brams and Taylor, 2000] Steven J. Brams and Alan D. Taylor. *The Win-win Solution. Guaranteeing Fair Shares to Everybody*. W. W. Norton & Company, 2000.
- [Brams *et al.*, 2003] Steven Brams, Paul Edelman, and Peter Fishburn. Fair division of indivisible items. *Theory and Decision*, 55:147–180, 02 2003.
- [Bu and Tao, 2024] Xiaolin Bu and Biaoshuai Tao. Truthful and almost envy-free mechanism of allocating indivisible goods: the power of randomness. *CoRR*, abs/2407.13634, 2024.
- [Caragiannis and Rathi, 2023] Ioannis Caragiannis and Nidhi Rathi. Optimizing over serial dictatorships. In *SAGT*, volume 14238 of *Lecture Notes in Computer Science*, pages 329–346. Springer, 2023.
- [Caragiannis *et al.*, 2019] Ioannis Caragiannis, David Kurokawa, Hervé Moulin, Ariel D. Procaccia, Nisarg Shah, and Junxing Wang. The unreasonable fairness of maximum nash welfare. *ACM Trans. Economics and Comput.*, 7(3):12:1–12:32, 2019.
- [Chakraborty *et al.*, 2021] Mithun Chakraborty, Ulrike Schmidt-Kraepelin, and Warut Suksompong. Picking sequences and monotonicity in weighted fair division. *Artif. Intell.*, 301:103578, 2021.
- [Cheng *et al.*, 2010] Weiwei Cheng, Eyke Hüllermeier, and Krzysztof J Dembczynski. Label ranking methods based on the plackett-luce model. In Johannes Fürnkranz and Thorsten Joachims, editors, *Proceedings of the 27th International Conference on Machine Learning (ICML-10), June 21-24, 2010, Haifa, Israel*, pages 215–222. Omnipress, 2010.

- [Darmann and Klamler, 2016] Andreas Darmann and Christian Klamler. Proportional borda allocations. *Soc. Choice Welf.*, 47(3):543–558, 2016.
- [d’Aspremont and Gevers, 2002] Claude d’Aspremont and Louis Gevers. Chapter 10 social welfare functionals and interpersonal comparability. In *Handbook of Social Choice and Welfare*, volume 1 of *Handbook of Social Choice and Welfare*, pages 459–541. Elsevier, 2002.
- [Ehlers and Klaus, 2003] Lars Ehlers and Bettina Klaus. Coalitional strategy-proof and resource-monotonic solutions for multiple assignment problems. *Soc. Choice Welf.*, 21(2):265–280, 2003.
- [Flammini and Gilbert, 2020] Michele Flammini and Hugo Gilbert. Parameterized complexity of manipulating sequential allocation. In *Proceedings of the 24th European Conference on Artificial Intelligence, ECAI 2020, Santiago de Compostela, Spain, August 29 - September 8, 2020*, volume 325 of *Frontiers in Artificial Intelligence and Applications*, pages 99–106. IOS Press, 2020.
- [Garg and Psomas, 2022] Rohan Garg and Alexandros Psomas. Efficient mechanisms without money: Randomization won’t let you escape from dictatorships. Manuscript, 2022.
- [Hatfield, 2009] John William Hatfield. Strategy-proof, efficient, and nonbossy quota allocations. *Soc. Choice Welf.*, 33(3):505–515, 2009.
- [Hoeffding, 1963] Wassily Hoeffding. Probability inequalities for sums of bounded random variables. *Journal of the American Statistical Association*, 1963.
- [Hosseini and Larson, 2019] Hadi Hosseini and Kate Larson. Multiple assignment problems under lexicographic preferences. In *Proceedings of the 18th International Conference on Autonomous Agents and MultiAgent Systems, AAMAS ’19, Montreal, QC, Canada, May 13-17, 2019*, pages 837–845. International Foundation for Autonomous Agents and Multiagent Systems, 2019.
- [Hosseini et al., 2021] Hadi Hosseini, Sujoy Sikdar, Rohit Vaish, and Lirong Xia. Fair and efficient allocations under lexicographic preferences. In *35th AAAI Conference on Artificial Intelligence, AAAI 2021, Virtual Event, February 2-9, 2021*, pages 5472–5480. AAAI Press, 2021.
- [Kalinowski et al., 2013a] Thomas Kalinowski, Nina Nardytka, and Toby Walsh. A social welfare optimal sequential allocation procedure. In *Proceedings of the 23rd International Joint Conference on Artificial Intelligence, IJCAI 2013, Beijing, China, August 3-9, 2013*, pages 227–233. IJCAI/AAAI Press, 2013.
- [Kalinowski et al., 2013b] Thomas Kalinowski, Nina Nardytka, Toby Walsh, and Lirong Xia. Strategic behavior when allocating indivisible goods sequentially. In *Proceedings of the 27th AAAI Conference on Artificial Intelligence, Bellevue, Washington, USA, July 14-18, 2013*. AAAI Press, 2013.
- [Kohler and Chandrasekaran, 1971] David A. Kohler and R. Chandrasekaran. A class of sequential games. *Operations Research*, 19(2):270–277, 1971.
- [Kojima, 2009] Fuhito Kojima. Random assignment of multiple indivisible objects. *Math. Soc. Sci.*, 57(1):134–142, 2009.
- [Lang and Rothe, 2024] Jérôme Lang and Jörg Rothe. Fair division of indivisible goods. In *Economics and Computation*, Springer texts in business and economics, pages 493–550. Springer, 2nd edition, 2024.
- [Luce, 1959] R Duncan Luce. *Individual choice behavior*, volume 4. Wiley New York, 1959.
- [Mallows, 1957] Colin L Mallows. Non-null ranking models. i. *Biometrika*, 44(1/2):114–130, 1957.
- [Nguyen et al., 2018] Nhan-Tam Nguyen, Dorothea Baumeister, and Jörg Rothe. Strategy-proofness of scoring allocation correspondences for indivisible goods. *Soc. Choice Welf.*, 50(1):101–122, 2018.
- [Pápai, 2000] Szilvia Pápai. original papers : Strategyproof multiple assignment using quotas. *Review of Economic Design*, 5(1):91–105, 2000.
- [Plackett, 1975] Robin L Plackett. The analysis of permutations. *Journal of the Royal Statistical Society Series C: Applied Statistics*, 24(2):193–202, 1975.
- [Psomas and Verma, 2022] Alexandros Psomas and Paritosh Verma. Fair and efficient allocations without obvious manipulations. In S. Koyejo, S. Mohamed, A. Agarwal, D. Belgrave, K. Cho, and A. Oh, editors, *Advances in Neural Information Processing Systems*, volume 35, pages 13342–13354. Curran Associates, Inc., 2022.
- [Pápai, 2001] Szilvia Pápai. Strategyproof and Nonbossy Multiple Assignments. *Journal of Public Economic Theory*, 3(3):257–271, July 2001.
- [Silverberg, 1980] Arthur Richard Silverberg. *Statistical models for q-permutations*. Princeton University, 1980.
- [Svensson, 1999] Lars-Gunnar Svensson. Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, 16(4):557–567, 1999.
- [Tominaga et al., 2016] Yuto Tominaga, Taiki Todo, and Makoto Yokoo. Manipulations in two-agent sequential allocation with random sequences. In *Proceedings of the 2016 International Conference on Autonomous Agents & Multiagent Systems, Singapore, May 9-13, 2016*, pages 141–149. ACM, 2016.
- [Xiao and Ling, 2020] Mingyu Xiao and Jiaying Ling. Algorithms for manipulating sequential allocation. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence, AAAI 2020, New York, NY, USA, February 7-12, 2020*, pages 2302–2309. AAAI Press, 2020.