

Fairness in Repeated Matching: A Maximin Perspective

Eugene Lim¹, Tzeh Yuan Neoh², Nicholas Teh³

¹National University of Singapore, Singapore

²Harvard University, USA

³University of Oxford, UK

elimwj@comp.nus.edu.sg, tzehyuan_neoh@g.harvard.edu, nicholas.teh@cs.ox.ac.uk

Abstract

We study a sequential decision-making model where a set of items is repeatedly matched to the same set of agents over multiple rounds. The objective is to determine a sequence of matchings that either maximizes the utility of the least advantaged agent at the end of all rounds (optimal) or at the end of every individual round (anytime optimal). We investigate the computational challenges associated with finding (anytime) optimal outcomes and demonstrate that these problems are generally computationally intractable. However, we provide approximation algorithms, fixed-parameter tractable algorithms, and identify several special cases whereby the problem(s) can be solved efficiently. Along the way, we also establish characterizations of Pareto-optimal/maximum matchings, which may be of independent interest to works in matching theory and house allocation.

1 Introduction

Traditional machine learning (ML) algorithms often focus on global objectives such as efficiency (e.g., maximizing accuracy or minimizing error rates in decision-making systems) or maximizing revenue/profit (e.g., maximizing click-through rates for recommendation systems), as they align closely with organizational goals and are more straightforward to quantify and optimize. However, modern approaches increasingly emphasize *fairness* as a key desideratum, as societal and regulatory demands push for more equitable and responsible ML systems.

We consider a multi-agent sequential decision-making scenario where a set of resources must be allocated among agents repeatedly over time, with the objective of achieving fairness in the assignment process. This framework encompasses applications such as dynamic spectrum allocation in wireless networks and energy distribution in smart grids (El-hachmi 2022; Jain et al. 2022; Rony, Lopez-Aguilera, and Garcia-Villegas 2021; Soares et al. 2024). In the case of spectrum allocation, communication channels must be repeatedly assigned to devices, with each device requiring exclusive access to one channel in each time slot. Persistent disparities in access can degrade system efficiency, reduce user satisfaction, and undermine trust. Similarly, in

many other ML-driven resource allocation systems, disparities in the distribution of resources—such as GPUs in distributed computing—can lead to unfair outcomes that compromise the perceived and actual effectiveness of the system. Numerous other applications where decisions are made dynamically—such as assigning tasks to workers in crowdsourcing platforms (Moayedikia, Ghaderi, and Yeoh 2020), or distributing compute resources in cloud systems (Belgacem 2022; Gupta, Samvatsar, and Singh 2017; Saraswathi, Kalaashri, and Padmavathi 2015)—call for central decision-makers to ensure that no agent is persistently disadvantaged, which is critical for both fairness and long-term trust in the system.

The scenarios described above can be captured using the *repeated matching* framework—a multi-agent sequential decision-making model in which a set of goods is repeatedly matched to agents over time, and each agent is assigned exactly one good at each round. This can also be viewed as a multi-round generalization of the *bottleneck assignment problem* (Ford and Fulkerson 1962) which is well-known in multi-agent task allocation: an application of this problem arises in *threat seduction*, where decoys are assigned to multiple incoming threats (Shames et al. 2017). Our problem can also be viewed as a sequential variant of the *Santa Claus problem* (Bansal and Sviridenko 2006), which is closely related to the classic scheduling problem of *makespan minimization* on unrelated parallel machines (Lenstra, Shmoys, and Tardos 1990; Bamas et al. 2024).

In particular, we focus on the *maximin* (or *egalitarian*) objective (Demko and Hill 1988; Thomson 1983), which aims to find a sequence of matchings that maximizes the minimum utility among agents. Maximin fairness serves as a principled trade-off between fairness and efficiency, as minimizing disparities often enhances overall system robustness and user satisfaction. Moreover, modern ML systems often involve iterative, data-driven decision-making, and maximin fairness integrates naturally with these systems by providing a fairness criterion that adapts dynamically, with its ability to handle both short-term and long-term outcomes.¹

¹This is in contrast to other *comparative* notions of fairness, such as *envy-freeness*, which has also been studied in the static matching (Aigner-Horev and Segal-Halevi 2022; Wu and Roth 2018; Yokoi 2020) and the two-sided repeated matching (Gollapudi, Kollias, and Plaut 2020) setting. Maximin fairness is also

1.1 Our Results

We study the *repeated matching* problem from the perspective of *maximin* (or *egalitarian*) fairness, a principle grounded in game theory, fair division, and matching problems. Leveraging techniques from classical matching algorithms, approximation methods, dynamic programming, and online decision-making, we analyze how to design fair repeated matching policies that ensure long-term fairness across multiple rounds.

In Section 2, we formally define the repeated matching problem and introduce the notion of (*anytime*) *optimality* in the egalitarian sense. We also introduce several tools that is central in proving some of our results.

In Section 3, we study the computation of optimal solutions. We begin by defining the decision variant of our matching problem and showing that it is NP-hard in general. Notably, this hardness holds even with only two timesteps and ternary agent valuations (i.e., when each agent’s utility for a good takes one of three possible values). Given these hardness results, we turn to the optimization variant of the problem and develop approximation algorithms that achieve an additive approximation bound independent of the number of rounds T . Crucially, this implies that as T increases—a scenario common in real-world applications—the solution produced by our algorithm converges to the optimal one. In addition, we also show that the problem is fixed-parameter tractable (FPT) with respect to the number of agents by providing a polynomial-time algorithm when the number of agents is a constant. Notably, in the process, we derive a characterization of *Pareto optimal* matchings in terms of the permutations of agents. This generalizes the previously-known result that *serial dictatorship* characterizes Pareto optimal matchings and may be of independent interest to communities working on the *house allocation* problem.

In Section 4, we shift our focus to *anytime optimal* solutions. We show that such solutions always exist for two agents, and we provide a polynomial-time algorithm for it. However, this does not extend to three or more agents—even with just two rounds, deciding if an instance admits an anytime optimal solution becomes coNP-hard. Nevertheless, we design an approximation algorithm that achieves anytime optimality with an additive bound independent of T . These results underscore the inherent difficulty of achieving anytime optimality in our setting.

In Section 5, we revisit optimality and identify three special cases admitting polynomial-time algorithms: (i) agents with binary valuations, (ii) two types of goods, and (iii) identical agent valuations. These special cases are well-motivated by the (temporal) fair division literature. For (i), we present an exact algorithm and a new characterization of Pareto optimal matchings under binary valuations. For (ii), we similarly provide an efficient exact algorithm. For (iii), despite NP-hardness in general, we show that optimal solutions can be computed in polynomial time when the number of rounds is a multiple of the number of agents. Finally, we extend our approximation approach to anytime optimality in these cases, giving us a stronger result than in the general

more *demonstrably* fair compared to an envy-based approach.

setting.

1.2 Related Work

We highlight several streams of research that are related to our work. We note that while there are many works on *online* matching and fair division, they are not directly relevant to our setting, as the underlying assumptions differ fundamentally. In our setting, the entire set of goods is made available in every round, whereas in online models, the set of goods may vary over time. Thus, we focus only on discussing works where meaningful implications can be drawn between their results and ours.

Repeated matching. Repeated matching was first studied by Hosseini, Larson, and Cohen (2015), which considered ordinal preferences that could change over time. They study strategyproofness and approximate envy-freeness. However, ordinal (their model) and cardinal (our model) preferences are vastly different, both in techniques and results. Gollapudi, Kollias, and Plaut (2020) subsequently looked at a two-sided repeated matching problem (i.e., each side have preferences over the other side). They also study approximate envy-freeness as the key desiderata, albeit under some strong assumptions. In contrast, our model is on one-sided repeated matching, which is fundamentally different. Our model is most aligned with that of Caragiannis and Narang (2024). However, they consider a slightly more general variant, whereby the value of an agent for a good in some round depends on the number of rounds in which the good has been given to the agent in the past. They study approximately envy-free notions, show an intractability result, and special cases where fairness can be guaranteed. Our model, while more specialized than theirs, has a few distinctions: (i) we have stronger negative and intractability results, (ii) the fairness concept we consider is not envy-based, and is therefore novel in this domain, and (iii) we consider a notion of fairness *at every round prefix*, something with prior work does not consider—they look at fairness at the end. Recently, Micheel and Wilczynski (2024) also studied essentially the same model (under a different name: *repeated house allocation*), but with ordinal preferences and other kinds of envy-based measures.

Repeated fair division. Igarashi et al. (2024) studied a model of repeated fair division, where a set of goods is available at each round, and every good must be allocated. This is in contrast to our model where each agent gets exactly one good. They consider the compatibility of envy-freeness and Pareto optimality, and show positive results in restricted cases. Balan, Richards, and Luke (2011) study a similar model, but with a focus on the *average* utility of goods received by the agents. Note that as with classical fair division, house allocation (where each agent gets exactly one good) is a special case and has considerably different results. Elkind et al. (2025) also consider a non-repeated (but also offline) variant of this model where a single good needs to be allocate at each round.

Multi-agent sequential decision-making. Several other works in multi-agent systems bear resemblance to our

model. For instance, Zhang and Shah (2014) also study the egalitarian objective multi-agent decision-making problems. However, they take a non-cooperative game-theoretic approach and do not study a matching problem. Lim, Tan, and Soh (2024) consider an assignment problem in the context of stochastic multi-armed bandits, with egalitarian fairness as the objective. In their setting, at each round, exactly one “arm” must be assigned to each user such that no two users are assigned to the same arm. However, the user’s utility (“reward”) in this case is stochastic, and therefore explores a different problem. Several other works (Cheng, Kellerer, and Kotov 2005; Kellerer et al. 1997) consider the problem of semi-online multiprocessor scheduling, with the objective of minimizing the *makespan* (i.e., minimize the maximum time taken by any any processor). This is analogous to the egalitarian objective. However, results in this setting only hold for identical valuations (since machines are identical), and primarily apply to a (semi-)online setting, where goods arrive one at a time (and so valuations over future goods are known not in advance), but the total valuation is known.

Santa Claus problem. Another related line of work is egalitarian fair division, also known as the *Santa Claus* problem. The standard model here is a single-shot fair division setting with an egalitarian objective, which was studied as far back as Thomson (1983), who axiomatically characterized the egalitarian solution using numerous desirable properties. Bansal and Sviridenko (2006) then initiated the study of approximation algorithms for this problem, by providing an $\mathcal{O}(\log \log m / \log \log \log m)$ approximation algorithm for the special case when agents have *restricted additive* valuations. Annamalai, Kalaitzis, and Svensson (2015) and Davies, Rothvoss, and Zhang (2020) subsequently provided a 12.33 - and $(4 + \varepsilon)$ -approximation algorithm for this restricted case, respectively. Numerous other works study *online* variants of this problem, but typically under various relaxations—since strong worst-case guarantees are impossible without additional assumptions. Some of these restrictions include allowing for some reordering in the allocation process (Epstein, Levin, and van Stee 2010) or restricting the number of agents (He and Jiang 2005; Tan and Cao 2005; Wu, Cheng, and Ji 2014), or allowing transfer of items after assignment (Chen and Qin 2011).

2 Preliminaries

Given a positive integer z , let $[z] = \{1, \dots, z\}$. We consider the problem of fairly matching a set of n agents $N = [n]$ to a set of $m \geq n$ goods $G = \{g_1, \dots, g_m\}$ over T rounds. We note that this is without loss of generality—to model the case of $m < n$, one can simply add zero-valued goods to arrive at the $m \geq n$ case and the results remain the same.

Matchings. A *matching* M is an injective map from N to G . We have $M(i) = g$ if and only if agent $i \in N$ is matched to good $g \in G$. In some instances, we also represent a matching either as a n -tuple $M = (M(1), \dots, M(n))$ or as an $n \times m$ matrix M , where $M_{ij} = 1$ if $M(i) = g_j$, and 0 otherwise. We denote the set of all sequences of matchings with length at least $t \in [T]$ as \mathbb{S}^t .

Valuations. Let $u_i(g)$ denote the non-negative value that agent $i \in N$ receives when matched to good $g \in G$. The *valuation profile* of a matching M is the n -tuple $(u_1(M(1)), \dots, u_n(M(n)))$. Given a sequence of T matchings $S = (M^1, \dots, M^T)$, the value that agent i receives under S up to round $t \in [T]$ is the sum of the values received up to that round, that is, $v_i^t(S) := \sum_{s=1}^t u_i(M^s(i))$.

Instances. An instance of the *egalitarian repeated matching* problem is a tuple $\mathcal{I} = (N, G, T, \{u_i\}_{i \in N})$. The egalitarian (or maximin) objective seeks to maximize the value received by the worst-off agents. Let $t \in [T]$. We define the *bottleneck agents* of a sequence $S \in \mathbb{S}^t$ at round t as the set of agents who received the lowest value under S up to that round. We further define the *bottleneck value* as the value received by the bottleneck agents, that is, $b^t(S) := \min_{i \in N} v_i^t(S)$.

Objective. Motivated by the egalitarian objective, we denote the maximum bottleneck value at round t as $\text{OPT}(t) := \max\{b^t(S) \mid S \in \mathbb{S}^t\}$. In this work, we consider two notions of optimality²: one that ensures the best outcome at a specific round, and another that ensures the best outcome at every round up to a given round. Both concepts of this nature (fairness at the end or at the end of each prefix) have been studied in temporal/repeated fair division (Elkind et al. 2025; Igarashi et al. 2024) and repeated matching (Caragianis and Narang 2024).

We first introduce the weaker notion of optimality,³ which is defined by mandating fairness at the end of a particular round $t \in [T]$. More formally, we say that a sequence $S \in \mathbb{S}^t$ is *optimal* at round $t \in [T]$ if $b^t(S) = \text{OPT}(t)$.

Note that this property does not require optimality to hold at any previous rounds s , for $s < t$. However, for any round $t \in [T]$, if we require optimality at every round $s \leq t$, then we get a stronger notion of optimality. More formally, we say that a sequence $S \in \mathbb{S}^t$ is *anytime optimal* up to round $t \in [T]$ if $b^s(S) = \text{OPT}(s)$ for all rounds $s \in [t]$.

Observe that while anytime optimality is significantly stronger than standard optimality, positive results for anytime optimality do not necessarily extend to the well-studied online setting. This is because, in the online setting, goods typically arrive one at a time, and valuations over these goods can be arbitrary—potentially over an unlimited set.

Efficiency. We also consider *Pareto optimality*, a notion of economic efficiency commonly studied in the social choice literature. Formally, a matching M is said to *weakly Pareto*

²For simplicity, we refer to *optimality* as shorthand for the egalitarian welfare-maximizing optimal solution.

³Note that our problem with optimality as an objective can be reformulated as a single-shot fair division problem with T copies of each good and an added constraint that each agent receives exactly T goods. While mathematically equivalent, this formulation is unintuitive in the classical setting, non-standard, and remains unexplored (with no known algorithms designed for it) in the literature. Furthermore, the *sequential* perspective is necessary for defining and motivating anytime-optimality and enabling potential extensions, neither of which can be naturally accommodated in a single-shot optimization framework.

dominates another matching M_0 if all agents $i \in N$ receive at least as much value under M as M_0 , that is, $u_i(M(i)) \geq u_i(M_0(i))$. A matching M is said to *strongly Pareto dominates* M_0 if M weakly Pareto dominates M_0 and there exist some agent $i \in N$ with $u_i(M(i)) > u_i(M_0(i))$. A matching M is *Pareto optimal* when no matching strongly Pareto dominates M .

2.1 Allocations and Bistochastic Matrices

Working with sequences of matchings can be challenging due to the constraints imposed by each matching. It would be helpful if we could ignore these constraints in our analysis and focus solely on the frequency with which each good is allocated to each agent. We refer to such an abstraction as an allocation. An *allocation* $A = (A_1, \dots, A_n)$ is a collection of multiset, where A_i is the multiset of goods that are allocated to agent $i \in N$. We can represent an allocation as a matrix A where A_{ij} is the number of times good $g_j \in G$ appears in A_i . The value that agent i receives under A is defined as

$$v_i(A) := \sum_{g \in A_i} u_i(g) = \sum_{g_j \in G} A_{ij} u_i(g_j).$$

Lemma 2.1 states that an allocation can be transformed into a polynomial-length sequence of unique matching. Hence, when a proof is phrased in terms of allocations instead of a sequence, no generality is lost. Accordingly, we will often reason with allocations in our proofs, invoking the lemma whenever an explicit sequence of matchings is required.

Lemma 2.1. *Suppose $A \in \mathbb{R}^{n \times m}$ is an allocation with*

$$\sum_{i \in N} A_{ij} \leq T \quad \text{and} \quad \sum_{g_j \in G} A_{ij} \leq T.$$

Then, there exist a sequence of matchings S consisting of $d \leq m^2 - m + 1$ unique matchings that satisfy $v_i^T(S) \geq v_i(A)$. This can be computed in polynomial time.

Several proofs of our results, including the preceding lemma, represent an allocation as a bistochastic matrix. A *bistochastic matrix* is a non-negative square matrix whose rows and columns each sum to 1, and a *scaled integer bistochastic matrix* is its integer counterpart, with non-negative integer entries and the sum of each row and column is a common integer. We defer an extended discussion of the mathematical preliminaries (along with all other omitted proofs in this paper) to the appendix.

3 Finding Optimal Sequences

We begin by focusing on optimality in this section. We first show that finding an optimal sequence of matchings is computationally intractable. We then show a relationship between a multiplicative approximation to our problem and the popular Santa Claus problem. Since computing exact solutions is intractable for large instances, we propose an approximation algorithm to find a near-optimal sequence efficiently. We also complement the hardness result by introducing a fixed-parameter tractable (FPT) algorithm that finds an

optimal sequence when n or m is a constant, thereby providing an efficient algorithm for practical applications.

We assume that the reader is familiar with basic notions of classic complexity theory (Papadimitriou 2007) and parameterized complexity (Flum and Grohe 2006; Niedermeier 2006).

3.1 Hardness Results

Consider the decision problem associated with the egalitarian repeated matching problem, as follows.

EGALITARIAN REPEATED MATCHING (ERM)

Input: An instance $(N, G, T, \{u_i\}_{i \in N})$ and a target κ .

Question: Is there a sequence $S \in \mathbb{S}^T$ with $b^T(S) \geq \kappa$?

We show that ERM is NP-complete by reducing from a known NP-hard problem, 3-OCC-3-SAT (defined in the proof). This result also implies that ERM is APX-hard—that is, there exists no polynomial-time approximation scheme (PTAS) for the problem. Our result is as follows.

Theorem 3.1. *ERM is NP-complete (and APX-hard) even when $u_i(g) \in \{0, 0.5, 1\}$ for all $i \in N$ and $g \in G$, for any $T \geq 2$.*

An implication of ERM not having a PTAS is that only constant-factor multiplicative approximations may be possible (though its existence is not guaranteed). We define this formally: for any $c \in [1, \infty)$, we say that an algorithm is *c-approximate* (or simply *c-approx*) if the sequence $S \in \mathbb{S}^t$ returned by the algorithm satisfy $b^t(S) \geq \text{OPT}/c$ for all $t \in [T]$. When $c = 1$, we have an exact algorithm. A natural question is whether ERM admits a *c-approx* algorithm, for some constant $c \in [1, \infty)$. Interestingly, we show that the existence of a *c-approx* algorithm for ERM would imply the existence of a *c-approx* algorithm for the single-shot egalitarian fair division problem (i.e., the Santa Claus problem with additive valuations⁴).

Proposition 3.2. *For any $c \in [1, \infty)$, there is a c-approx algorithm for ERM only if there is a c-approx algorithm for the Santa Claus problem with additive valuations.*

The result above implies that finding even a constant-factor multiplicative approximation algorithm for ERM is likely to be very challenging. This is because, despite the Santa Claus problem being a well-studied and long-standing problem, no constant-factor approximation is currently known for the version with general additive valuations. A constant-factor approximation is only known in the restricted additive case.⁵

⁴We specify “additive valuations” explicitly as some works (e.g., Davies, Rothvoss, and Zhang (2020)) consider a more restricted variant of the Santa Claus problem with restricted additive valuations.

⁵The current best known approximation factor is $(4 + \varepsilon)$, for a small $\varepsilon > 0$ in this restricted case (Davies, Rothvoss, and Zhang 2020).

Algorithm 1: Approximation algorithm for finding an optimal sequence of matchings

Input: An instance $\mathcal{I} = (N, G, T, \{u_i\}_{i \in N})$

- 1: let B be the solution to linear program (P1)
 - 2: decompose B into $\alpha_1 M_1 + \dots + \alpha_d M_d$ using Birkhoff's algorithm
 - 3: let S be an empty sequence
 - 4: add $\lfloor T\alpha_k \rfloor$ copies of M_k in S for each $k \in [d]$
 - 5: add any matchings into S so that $|S| = T$
 - 6: **return** S
-

3.2 Approximation Algorithm

Given the results above, we focus on whether we can achieve an *additive* approximation with respect to optimality instead. We now describe an approximation algorithm that achieves an additive approximation bound independent of the number of rounds T . Crucially, this implies that as T increases, the approximate solution converges rapidly to the optimal one. The setting when the number of rounds is large can be observed in applications where the matching process runs continuously over extended periods—such as dynamic spectrum allocation (where the system operates continuously, often measured in (milli)seconds), leading to an immense number of allocation rounds.

Without loss of generality, we can assume that $n = m$; otherwise, we can simply create $m - n$ dummy agents with $u_i(g_j) = \max_{i' \in N} \max_{g'_j \in G} u_{i'}(g'_j)$ for all dummy agents i and goods g . Then, consider the following linear program:

$$\begin{aligned}
 & \text{maximize}_{b, B} && b && \text{(P1)} \\
 & \text{subject to} && \sum_{g_j \in G} B_{ij} u_i(g_j) \geq b, && \forall i \in N, \\
 & && \sum_{g_j \in G} B_{ij} = 1, && \forall i \in N, \\
 & && \sum_{i \in N} B_{ij} = 1, && \forall g_j \in G, \\
 & && B_{ij} \geq 0, && \forall i \in N, \forall g_j \in G.
 \end{aligned}$$

Note that the solution to (P1) is a bistochastic matrix B . Our approximation algorithm uses Birkhoff's algorithm to decompose B into a convex combination of matchings. The number of times each matchings are included in the sequence is then determined by the convex coefficients (see Algorithm 1).

Then, we prove the following result.

Theorem 3.3. *Given an instance $(N, G, T, \{u_i\}_{i \in N})$, the sequence $S \in \mathbb{S}^T$ returned by Algorithm 1 satisfies*

$$b^T(S) \geq \text{OPT}(T) - m \cdot \max_{i \in N} \max_{g \in G} u_i(g).$$

Proof. Consider the allocation A in which $A_{ij} = \lfloor TB_{ij} \rfloor$ for all $i \in N$ and $g_j \in G$. Note that for each $g_j \in G$, we have

$$\sum_{i \in N} A_{ij} = \sum_{i \in N} \lfloor TB_{ij} \rfloor \leq \sum_{i \in N} TB_{ij} = T,$$

and similarly, for each $i \in N$, we have

$$\sum_{g_j \in G} A_{ij} = \sum_{g_j \in G} \lfloor TB_{ij} \rfloor \leq \sum_{g_j \in G} TB_{ij} = T.$$

By Lemma 2.1, there exist a sequence S over T rounds composed of at most $O(m^2)$ unique matchings such that $v_i^T(S) \geq v_i(A)$. Then, for any agent $i \in N$, we have

$$\begin{aligned}
 v_i^T(S) \geq v_i(A) &\geq \sum_{g_j \in G} u_i(g_j) \lfloor TB_{ij} \rfloor \\
 &\geq \sum_{g_j \in G} u_i(g_j) \cdot (TB_{ij} - 1) \\
 &= \sum_{g_j \in G} TB_{ij} u_i(g_j) - \sum_{g_j \in G} u_i(g_j) \\
 &\geq Tb - m \cdot \max_{g_j \in G} u_i(g_j) \\
 &\geq \text{OPT}(T) - m \cdot \max_{g_j \in G} u_i(g_j).
 \end{aligned}$$

Let $k \in N$ be a bottleneck agent of sequence S at round T so that $b^T(S) = v_k^T(S)$. Then, we have

$$\begin{aligned}
 b^T(S) &\geq \text{OPT}(T) - m \cdot \max_{g_j \in G} u_k(g_j) \\
 &\geq \text{OPT}(T) - m \cdot \max_{i \in N} \max_{g_j \in G} u_i(g_j). \quad \square
 \end{aligned}$$

Note that although the maximum valuation can be arbitrarily large, they are typically bounded in practice. Consequently, such a bound remains informative and relevant. Instance-dependent additive bounds of this type are well-established in the literature, particularly in the context of stochastic bandits (Lattimore and Szepesvári 2020; Lim, Tan, and Soh 2024) and online fair division (Benadè et al. 2018; Hajiaghayi et al. 2022).

3.3 Fixed-Parameter Tractable (FPT) Algorithm

Next, we consider another approach to dealing with computational intractability. We show that the problem is *fixed parameter tractable* (FPT) when the number of agents is a fixed parameter, i.e., there exists an algorithm that can compute an optimal sequence in polynomial-time when n is a constant. This provides a practical solution for small-group matching. Our result is as follows.

Theorem 3.4. *Given an instance $(N, G, T, \{u_i\}_{i \in N})$, ERM is FPT with respect to n .*

The proof of Theorem 3.4 relies on our newly established characterizations of Pareto-optimal and maximum matchings in terms of permutations of agents. These results may be of independent interest to researchers in matching and house allocation.

In particular, let $\pi : N \rightarrow [n]$ be a permutation of the agents. A matching M_* is said to be π -*optimal* if there exists no matching M such that

- Some agent $i \in N$ satisfies $u_i(M(i)) > u_i(M_*(i))$; and
- For every such agent i , it holds that for all agents $i' \in N$ with $\pi(i') < \pi(i)$, we have $u_{i'}(M(i')) \geq u_{i'}(M_*(i'))$.

Then, we obtain the following lemma.

Lemma 3.5. *A matching M is Pareto optimal if and only if it is π -optimal for some permutation π .*

In the context of house allocation without indifferences, it is well-established that *serial dictatorship* characterizes Pareto-optimal allocations (Abdulkadiroğlu and Sönmez 1998). However, when agents are allowed to express indifferences between houses, the allocations produced by serial dictatorship are not guaranteed to be Pareto optimal (Abraham et al. 2004). Therefore, our definition of π -optimal can be interpreted as an extension of serial dictatorship that ensures Pareto optimality even in the presence of indifferences.

We describe how this characterization leads to an FPT algorithm in Section 3.3.

4 Anytime Optimality

In this section, we consider the problem of *anytime optimality*, a stronger notion that requires optimality at every round prefix. We show that an anytime optimal sequence always exists when $n = 2$, but determining whether such a sequence exists for $n \geq 3$ is **coNP-hard**. The setting of $n = 2$ is a widely studied and is an important special case in related literature (Elkind et al. 2025; Gollapudi, Kollias, and Plaut 2020; Igarashi et al. 2024). Our results are as follows.

Theorem 4.1. *Given an instance $(N, G, T, \{u_i\}_{i \in N})$ with $n = 2$, there always exist an approximate anytime optimal sequence of matchings, and we can find it in polynomial time.*

However, we show that this positive result does not extend to the case when $n \geq 3$, for all $T \geq 2$, with the following impossibility result.

Proposition 4.2. *An anytime optimal sequence might not exist for any problem instance $(N, G, T, \{u_i\}_{i \in N})$ with $n \geq 3$ and $T \geq 2$.*

The above implies that we cannot hope for anytime optimality in most cases. However, given a problem instance, one may still wish to obtain an anytime optimal result *if it exists*. Unfortunately, we show that even determining whether an instance admits an anytime optimal solution is computationally intractable, with the following result.

Theorem 4.3. *Given instance $\mathcal{I} = (N, G, T, \{u_i\}_{i \in N})$, the problem of deciding if \mathcal{I} admits an anytime optimal sequence is **coNP-hard**.*

Finally, we complement the above hardness result with an approximation algorithm that achieves an additive approximation bound independent of the number rounds T . Again, this means that as T increases, the approximate solution converges rapidly to the optimal one.

Theorem 4.4. *Given an instance $(N, G, T, \{u_i\}_{i \in N})$, there always exist a sequence of matchings that is anytime optimal. Furthermore, Algorithm 2 outputs a sequence of matchings S , in polynomial time, that satisfy*

$$b^t(S) \geq \text{OPT}(t) - 5m \cdot \max_{i \in N} \max_{g \in G} u_i(g), \quad \forall t \in [T].$$

Algorithm 2: Approximate algorithm for anytime optimal sequence

Input: An instance $\mathcal{I} = (N, G, T, \{u_i\}_{i \in N})$

- 1: let B be the solution to (P1)
 - 2: decompose B into $\alpha_1 M_1 + \dots + \alpha_d M_d$ using Birkhoff’s algorithm
 - 3: initialize $n_k = 0$ for all $k \in [d]$
 - 4: **for** $t = 1, \dots, T$ **do**
 - 5: choose matching $M^t = \arg \min_{M_k} (n_k + 1) / \alpha_k$
 - 6: update $n_k \leftarrow n_k + 1$
 - 7: **end for**
 - 8: **return** $\{M_1, \dots, M_T\}$
-

Proof sketch. Let n_{kt} be the value of n_k after round t . After each round $t \in [T]$, we claim that our choice of matching M^t maintains the invariant $n_{kt} \geq \alpha_k \cdot t - 1$ for all $k \in [d]$. Intuitively, this says that by any round t , each matching M_k has been selected for roughly its intended α_k fraction of the rounds. Thus, we will get a result similar to that of Theorem 3.3. More specifically, we can show that $v_i^t(S) \geq \text{OPT}(t) - d \cdot \max_{g \in G} u_i(g)$ for all $i \in N$. Observe that since (P1) has $m^2 + 5m$ inequality constraints and $m^2 + 1$ variables, $m^2 + 1$ constraints will be tight at a vertex solution, meaning there are at most $5m$ non-zero entries in B , which implies that $d \leq 5m$. \square

5 Special Cases

In this section, we shift our focus back to optimality⁶ and consider three special cases: (1) when agents have binary valuations, (2) when there are only two types of goods, and (3) when agents share identical valuations. For each of the first two cases, we provide an algorithm that computes an optimal sequence of matchings in polynomial time. A key technique that we used here is to reduce the problem to one of *circulation with demand* and leveraging the Ford-Fulkerson algorithm to compute a feasible circulation. We then show for the third case that the problem is hard even for optimality, address a special case where it can be solved in polynomial time, and provide an approximate anytime optimal algorithm for it.

5.1 Binary Valuations

The first setting we consider is when agents have *binary valuations*, i.e. $u_i : G \rightarrow \{0, 1\}$ for all agents $i \in N$. This is an important and well-studied subclass of valuations (sometimes referred to as *binary additive* valuations). Numerous fair division (Aleksandrov et al. 2015; Amanatidis et al. 2021; Bouveret and Lemaître 2016; Freeman et al. 2019; Halpern et al. 2020; Hosseini et al. 2020; Suksompong and Teh 2022) and matching (Bogomolnaia and Moulin 2004;

⁶Unfortunately, anytime optimality is a strong condition with relatively strong negative results (as with many similar problems in the online setting). We leave the existence (or impossibility) of obtaining anytime optimality in special cases as an interesting direction for future work.

Gollapudi, Kollias, and Plaut 2020) papers consider this setting. Binary valuations can also be viewed as *approval votes*, which have long been studied in the voting literature (Brams and Fishburn 2007; Kilgour 2010), and permit very simple elicitation.

Notably, under binary valuations, maximizing egalitarian welfare is equivalent to maximizing *Nash welfare* (i.e., the geometric mean), which is an extremely popular concept in fair division, and has many desirable properties (Halpern et al. 2020; Suksompong and Teh 2022).

We first establish the following lemma.

Lemma 5.1. *Let G' be the goods in a maximum matching. Then, for any matching M , there is a matching M_* that weakly Pareto dominates M and that the goods matched by M_* is a subset of G' .*

The above lemma basically provides another characterization, this time, of maximum matchings under binary valuations in terms of Pareto optimality. To the best of our knowledge, this result is also novel in the context of house allocation, which may be of independent interest. This lemma is used to prove the following result.

Theorem 5.2. *Given an instance $(N, G, T, \{u_i\}_{i \in N})$ with binary valuations, we can find an optimal sequence of matchings in polynomial time.*

Note that the NP-hardness result of Theorem 3.1 implies that we cannot strengthen the positive results to the setting where agents have *ternary valuations* (or three-valued instances) (Fitzsimmons, Viswanathan, and Zick 2025).

5.2 Two Types of Goods

Next, we consider the setting with two *types* of goods: each good can be divided into two groups, and each agent values all goods in a particular group equally. This preference restriction is also commonly studied in (temporal) fair division (Aziz et al. 2023; Elkind et al. 2025; Garg, Murhekar, and Qin 2024). Formally, let $G_0, G_1 \subseteq G$ be a partition of the set of goods such that $G_0 \cap G_1 = \emptyset$, $G_0 \cup G_1 = G$, and for all agent $i \in N$ and all goods $g, g' \in G_r$ for some $r \in \{0, 1\}$, we have $u_i(g) = u_i(g')$. Then, our result is as follows.

Theorem 5.3. *Given an instance $(N, G_1 \cup G_2, T, \{u_i\}_{i \in N})$ with two types of goods, we can find an optimal sequence of matchings in polynomial time.*

5.3 Identical Valuations

The last special case we consider here is one where agents have *identical valuation* functions, i.e., $u_i = u_{i'}$ for all agents $i, i' \in N$. The setting with identical valuations is also well-studied in the repeated fair division/matching (Caragiannis and Narang 2024; Igarashi et al. 2024) and standard fair division (Barman and Sundaram 2020; Mutzari, Aumann, and Kraus 2023; Plaut and Roughgarden 2020) literature. Moreover, works on semi-online multiprocessor scheduling with the makespan minimization objective (analogous to the egalitarian objective) (Cheng, Kellerer, and Kovtsov 2005; Kellerer et al. 1997) focus on identical valuations as well (since machines are identical in that setting).

We show that even under this restricted setting of identical valuations, the problem of finding an optimal sequence is generally still NP-hard.

Theorem 5.4. *Given an instance $(N, G, T, \{u_i\}_{i \in N})$ with identical valuations, finding an optimal sequence of matchings is NP-complete.*

However, when T is a multiple of n , we shown that the problem can be solved in polynomial time. We note that the case when T is a multiple of n is also a popular special case studied in repeated matching/fair division (Caragiannis and Narang 2024; Igarashi et al. 2024)

Theorem 5.5. *Given an instance $(N, G, T, \{u_i\}_{i \in N})$ with identical valuations and $T = kn$ for some $k \in \mathbb{Z}$, we can find an optimal sequence of matchings in polynomial time.*

Finally, we complement the above with an approximation algorithm that achieves (even anytime) optimality up to an additive approximation factor of $\max_{g \in G} u(g)$.⁷ This gives us a stronger result compared to the general case, which is also only for optimality (as in Theorem 3.3).

Theorem 5.6. *Given an instance $(N, G, T, \{u_i\}_{i \in N})$ with identical valuations, we can find, in polynomial time, a sequence of matchings S that satisfy*

$$b^t(S) \geq OPT(t) - \Delta, \quad \forall t \in [T],$$

where Δ is the difference in value between the most valuable good and the n -th most valuable good.⁸

6 Conclusion

In this work, we introduced and studied a model of repeated matching with goal of obtaining egalitarian optimality. We investigated the computational complexity of achieving optimality and anytime optimality, and identified several settings where these problems can be solved efficiently, together with accompanying algorithms. Specifically, for optimality, we provided an approximation algorithm independent of T , and FPT algorithms with respect to n or m . For anytime optimality, we provided an approximation algorithm that complements the hardness and impossibility result even in simple cases. We also showed two special cases (binary valuations, two types of goods) where optimality can be achieved, and a final special case (identical valuations) where approximate anytime optimality can be achieved.

Directions for future work include considering other special cases that admit efficient optimal solutions, such as bi-valued utilities (where each agent values each good at either 1 or some integer $p > 1$) or identical rankings. It would also be interesting to study concepts that interpolate optimality and anytime optimality (e.g., optimality at every τ timesteps). In two of our special cases, we mentioned the equivalence between egalitarian and Nash welfare. It would be interesting to identify the conditions under which these two objectives are equivalent in this setting.

⁷We denote agents' identical utility function as u . Then, $v_i^t(S) := \sum_{s=1}^t u(M^s(i))$ for all $t \in [T]$.

⁸This is equivalent to the concept of *gap* in bandits literature.

Acknowledgments

This research/project is supported by the National Research Foundation, Singapore under its AI Singapore Programme (AISG Award No: AISG2-PhD/2021-08-011).

References

- Abdulkadiroğlu, A.; and Sönmez, T. 1998. Random serial dictatorship and the core from random endowments in house allocation problems. *Econometrica*, 66(3): 689–701.
- Abraham, D. J.; Cechlárová, K.; Manlove, D. F.; and Mehlhorn, K. 2004. Pareto optimality in house allocation problems. In *Proceedings of the 15th International Symposium on Algorithms and Computation (ISAAC)*, 3–15.
- Aigner-Horev, E.; and Segal-Halevi, E. 2022. Envy-free matchings in bipartite graphs and their applications to fair division. *Information Sciences*, 587: 164–187.
- Aleksandrov, M.; Aziz, H.; Gaspers, S.; and Walsh, T. 2015. Online fair division: Analysing a food bank problem. In *Proceedings of the 24th International Joint Conference on Artificial Intelligence (IJCAI)*, 2540–2546.
- Amanatidis, G.; Birmpas, G.; Filos-Ratsikas, A.; Hollender, A.; and Voudouris, A. A. 2021. Maximum Nash welfare and other stories about EFX. *Theoretical Computer Science*, 863: 69–85.
- Annamalai, C.; Kalaitzis, C.; and Svensson, O. 2015. Combinatorial algorithm for restricted max-min fair allocation. In *Proceedings of the 2015 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 1357–1372.
- Aziz, H.; Lindsay, J.; Ritossa, A.; and Suzuki, M. 2023. Fair Allocation of Two Types of Chores. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 143–151.
- Balan, G.; Richards, D.; and Luke, S. 2011. Long-term fairness with bounded worst-case losses. *Autonomous Agents and Multi-Agent Systems*, 22: 43–63.
- Bamas, E.; Lindermayr, A.; Megow, N.; Rohwedder, L.; and Schlöter, J. 2024. Santa Claus meets Makespan and Matroids: Algorithms and Reductions. In *Proceedings of the 2024 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2829–2860.
- Bansal, N.; and Sviridenko, M. 2006. The Santa Claus problem. In *Proceedings of the 38th ACM Symposium on Theory of Computing (STOC)*, 31–40.
- Barman, S.; and Sundaram, R. G. 2020. Uniform Welfare Guarantees Under Identical Subadditive Valuations. In *Proceedings of the 29th International Joint Conference on Artificial Intelligence (IJCAI)*, 46–52.
- Belgacem, A. 2022. Dynamic resource allocation in cloud computing: Analysis and taxonomies. *Computing*, 104(3): 681–710.
- Benadè, G.; Kazachkov, A. M.; Procaccia, A. D.; and Psoomas, C.-A. 2018. How to Make Envy Vanish Over Time. In *Proceedings of the 19th ACM Conference on Economics and Computation (EC)*, 593–610.
- Bogomolnaia, A.; and Moulin, H. 2004. Random matching under dichotomous preferences. *Econometrica*, 72(1): 257–279.
- Bouveret, S.; and Lemaître, M. 2016. Characterizing conflicts in fair division of indivisible goods using a scale of criteria. *Autonomous Agents and Multiagent Systems*, 30(2): 259–290.
- Brams, S. J.; and Fishburn, P. C. 2007. *Approval Voting*. Springer.
- Caragiannis, I.; and Narang, S. 2024. Repeatedly matching items to agents fairly and efficiently. *Theoretical Computer Science*, 981: 114246.
- Chen, X.; and Qin, S. 2011. On-line machine covering on two machines with local migration. *Computers & Mathematics with Applications*, 62(5): 2336–2341.
- Cheng, T. C. E.; Kellerer, H.; and Kotov, V. 2005. Semi-online multiprocessor scheduling with given total processing time. *Theoretical Computer Science*, 337(1): 134–146.
- Davies, S.; Rothvoss, T.; and Zhang, Y. 2020. A Tale of Santa Claus, Hypergraphs and Matroids. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms (SODA)*, 2748–2757.
- Demko, S.; and Hill, T. P. 1988. Equitable distribution of indivisible objects. *Mathematical Social Sciences*, 16(2): 145–158.
- Elhachmi, J. 2022. Distributed reinforcement learning for dynamic spectrum allocation in cognitive radio-based internet of things. *IET Networks*, 11(6): 207–220.
- Elkind, E.; Lam, A.; Latifian, M.; Neoh, T. Y.; and Teh, N. 2025. Temporal Fair Division of Indivisible Items. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*, 676–685.
- Epstein, L.; Levin, A.; and van Stee, R. 2010. Max-min Online Allocations with a Reordering Buffer. In *Proceedings of the 37th International Colloquium on Automata, Languages, and Programming (ICALP)*, 336–347.
- Fitzsimmons, Z.; Viswanathan, V.; and Zick, Y. 2025. On the Hardness of Fair Allocation under Ternary Valuations. In *Proceedings of the 24th International Conference on Autonomous Agents and Multiagent Systems (AAMAS)*.
- Flum, J.; and Grohe, M. 2006. *Parameterized Complexity Theory*. Texts in Theoretical Computer Science. An EATCS Series. Springer.
- Ford, L. R.; and Fulkerson, D. R. 1962. *Flows in Networks*. Princeton University Press.
- Freeman, R.; Sikdar, S.; Vaish, R.; and Xia, L. 2019. Equitable allocations of indivisible goods. In *Proceedings of the 28th International Joint Conference on Artificial Intelligence (IJCAI)*, 280–286.
- Garg, J.; Murhekar, A.; and Qin, J. 2024. Weighted EF1 and PO Allocations with Few Types of Agents or Chores. In *Proceedings of the 33rd International Joint Conference on Artificial Intelligence (IJCAI)*, 2799–2806.

- Gollapudi, S.; Kollias, K.; and Plaut, B. 2020. Almost Envy-Free Repeated Matching in Two-Sided Markets. In *Proceedings of the 16th International Conference on Web and Internet Economics (WINE)*, 3–16.
- Gupta, P.; Samvatsar, M.; and Singh, U. 2017. Cloud computing through dynamic resource allocation scheme. In *Proceedings of the 2017 International Conference on Electronics, Communication and Aerospace Technology (ICECA)*, 544–548.
- Hajiaghayi, M.; Khani, M.; Panigrahi, D.; and Springer, M. 2022. Online Algorithms for the Santa Claus Problem. In *Proceedings of the 36th International Conference on Neural Information Processing Systems (NeurIPS)*, 30732–30743.
- Halpern, D.; Procaccia, A. D.; Psomas, A.; and Shah, N. 2020. Fair division with binary valuations: One rule to rule them all. In *Proceedings of the 16th Conference on Web and Internet Economics (WINE)*, 370–383.
- He, Y.; and Jiang, Y. 2005. Optimal semi-online preemptive algorithms for machine covering on two uniform machines. *Theoretical Computer Science*, 339(2): 293–314.
- Hosseini, H.; Larson, K.; and Cohen, R. 2015. Matching with Dynamic Ordinal Preferences. In *Proceedings of the 29th AAAI Conference on Artificial Intelligence (AAAI)*, 936–943.
- Hosseini, H.; Sikdar, S.; Vaish, R.; Wang, H.; and Xia, L. 2020. Fair division through information withholding. In *Proceedings of the 34th AAAI Conference on Artificial Intelligence (AAAI)*, 2014–2021.
- Igarashi, A.; Lackner, M.; Nardi, O.; and Novaro, A. 2024. Repeated Fair Allocation of Indivisible Items. In *Proceedings of the 38th AAAI Conference on Artificial Intelligence (AAAI)*, 9781–9789.
- Jain, K.; Dhabu, M.; Kakde, O.; and Funde, N. 2022. Completely fair energy scheduling mechanism in a smart distributed multi-microgrid system. *Journal of King Saud University - Computer and Information Sciences*, 34(9): 7819–7829.
- Kellerer, H.; Kotov, V.; Speranza, M. G.; and Tuza, Z. 1997. Semi on-line algorithms for the partition problem. *Operations Research Letters*, 21(5): 235–242.
- Kilgour, D. M. 2010. Approval balloting for multi-winner elections. In Laslier, J.-F.; and Sanver, M. R., eds., *Handbook on Approval Voting*, chapter 6, 105–124. Springer.
- Lattimore, T.; and Szepesvári, C. 2020. *Bandit Algorithms*. Cambridge University Press.
- Lenstra, J. K.; Shmoys, D. B.; and Tardos, E. 1990. Approximation algorithms for scheduling unrelated parallel machines. *Mathematical Programming*, 46: 259–271.
- Lim, E.; Tan, V. Y. F.; and Soh, H. 2024. Stochastic Bandits for Egalitarian Assignment. *Transactions on Machine Learning Research*.
- Micheel, K. J.; and Wilczynski, A. 2024. Fairness in Repeated House Allocation. In *Proceedings of the 27th European Conference on Artificial Intelligence (ECAI)*, 3549–3556.
- Moayedikia, A.; Ghaderi, H.; and Yeoh, W. 2020. Optimizing microtask assignment on crowdsourcing platforms using Markov chain Monte Carlo. *Decision Support Systems*, 139: 113404.
- Mutzari, D.; Aumann, Y.; and Kraus, S. 2023. Resilient Fair Allocation of Indivisible Goods. In *Proceedings of the 22nd International Conference on Autonomous Agents and Multi-agent Systems (AAMAS)*, 2688–2690.
- Niedermeier, R. 2006. *Invitation to Fixed-Parameter Algorithms*. Oxford University Press.
- Papadimitriou, C. H. 2007. *Computational complexity*. Academic Internet Publ. ISBN 978-1-4288-1409-7.
- Plaut, B.; and Roughgarden, T. 2020. Almost envy-freeness with general valuations. *SIAM Journal on Discrete Mathematics*, 34(2): 1039–1068.
- Rony, R. I.; Lopez-Aguilera, E.; and Garcia-Villegas, E. 2021. Dynamic Spectrum Allocation Following Machine Learning-Based Traffic Predictions in 5G. *IEEE Access*, 9: 143458–143472.
- Saraswathi, A.; Kalaashri, Y.; and Padmavathi, S. 2015. Dynamic Resource Allocation Scheme in Cloud Computing. *Procedia Computer Science*, 47: 30–36.
- Shames, I.; Dostovalova, A.; Kim, J.; and Hmam, H. 2017. Task allocation and motion control for threat-seduction decoys. In *Proceedings of the 2017 IEEE 56th Annual Conference on Decision and Control (CDC)*, 4509–4514.
- Soares, J.; Lezama, F.; Faia, R.; Limmer, S.; Dietrich, M.; Rodemann, T.; Ramos, S.; and Vale, Z. 2024. Review on fairness in local energy systems. *Applied Energy*, 374: 123933.
- Suksompong, W.; and Teh, N. 2022. On maximum weighted Nash welfare for binary valuations. *Mathematical Social Sciences*, 117: 101–108.
- Tan, Z.; and Cao, S. 2005. Semi-online Machine Covering on Two Uniform Machines with Known Total Size. *Computing*, 78: 369–378.
- Thomson, W. 1983. Problems of fair division and the Egalitarian solution. *Journal of Economic Theory*, 31(2): 211–226.
- Wu, Q.; and Roth, A. E. 2018. The lattice of envy-free matchings. *Games and Economic Behavior*, 109: 201–211.
- Wu, Y.; Cheng, T.; and Ji, M. 2014. Optimal algorithms for semi-online machine covering on two hierarchical machines. *Theoretical Computer Science*, 531: 37–46.
- Yokoi, Y. 2020. Envy-Free Matchings with Lower Quotas. *Algorithmica*, 82(2): 188–211.
- Zhang, C.; and Shah, J. A. 2014. Fairness in Multi-Agent Sequential Decision-Making. In *Proceedings of the 28th International Conference on Neural Information Processing Systems (NeurIPS)*, 2636–2644.