# *Abstracts*

## Contents

Georgios Amanatidis
*University of Essex, UK*

## *Pushing the Frontier on Approximate EFX Allocations*

We study the problem of allocating a set of indivisible goods to a set of agents with additive valuation functions, aiming to achieve approximate envy-freeness up to any good (α-EFX). The state-of-the-art results on the problem include that (exact) EFX allocations exist when (a) there are at most three agents, or (b) the agents' valuation functions can take at most two values, or (c) the agents' valuation functions can be represented via a graph. For α-EFX, it is known that a 0.618-EFX allocation exists for any number of agents with additive valuation functions. In this paper, we show that 2/3-EFX allocations exist when (a) there are at most seven agents, (b) the agents' valuation functions can take at most three values, or (c) the agents' valuation functions can be represented via a multigraph. Our results can be interpreted in two ways. First, by relaxing the notion of EFX to 2/3-EFX, we obtain existence results for strict generalizations of the settings for which exact EFX allocations are known to exist. Secondly, by imposing restrictions on the setting, we manage to beat the barrier of 0.618 and achieve an approximation guarantee of 2/3. Therefore, our results push the frontier of existence and computation of approximate EFX allocations, and provide insights into the challenges of settling the existence of exact EFX allocations.

This is joint work with Aris Filos-Ratsikas and Alkmini Sgouritsa, and has appeared in the Proceedings of the 25th ACM Conference on Economics and Computation (EC'24).

Back to Contents Page →

Bhaskar Ray Chaudhury
*University of Illinois Urbana-Champaign, USA*

## *On the Theoretical Foundations of Data Exchange Economies*

The immense success of ML systems relies heavily on large-scale high-quality data. The high demand for data has led to several paradigms that involve selling, exchanging, and sharing data. This naturally motivates studying economic processes that involve data as an asset. However, data differs from classical economic assets in terms of (i) free duplication i.e., there is no concept of limited supply with data as it can be replicated at zero marginal cost, and (ii) ex-ante unverifiable, i.e., it is difficult to estimate the utility of the data to an agent apriori, without using it. These distinctions cause fundamental differences between economic processes involving data and those involving other assets.

We investigate the parallel of exchange markets (Arrow-Debreu markets) in settings where data is the asset, i.e., where agents in possession of datasets exchange data fairly and voluntarily for mutual benefit without any monetary compensation. This is relevant in settings involving non-profit organizations that are seeking to improve their ML models through data-exchange with other organizations and are not allowed to sell their data for profit. This work proposes a general framework for data-exchange from first principles. We investigate the existence and computation of a data-exchange satisfying the foregoing principles.

Back to Contents Page →

Jiarui Gan
*University of Oxford, UK*

---

*Envy-free policy teaching to multiple agents*

---

In this talk, I will present our work on envy-free policy teaching in multi-agent systems. We consider a scenario where multiple agents, each with their own reward function and discount rate, explore a shared environment. A teacher aims to guide them toward a target policy by modifying their rewards while ensuring envy-freeness (EF) so that the agents feel that the reward modification is fair. We explore how to design reward modifications to achieve this. I will discuss three main findings: 1) EF solutions always exist when both bonuses and penalties are allowed, 2) we can compute cost-efficient EF solutions using convex optimization, and 3) the cost of fairness—the "price of fairness" (PoF)—scales reasonably with problem size and agent complexity. Our results show that fairness can be incorporated into policy teaching without significant costs or computational difficulties.

Back to Contents Page →

Jugal Garg
*University of Illinois Urbana-Champaign, USA*

---

*Fair allocation of indivisible chores*

---

Fair division is an age-old problem of allocating a set of items among agents with preferences in a fair and efficient manner. It naturally arises in a wide range of real-life settings, from interpersonal to international conflicts. In my talk, I will present recent progress in the field, specifically focusing on the allocation of indivisible chores.

Back to Contents Page →

Erel Segal-Halevi
*Ariel University, Israel*

---

*Fairness in real-estate division*

---

I will present some of my attempts to apply fair division procedures to real-estate projects.

The first application involves "combination and redivision" projects, where an agricultural land becomes available for building. The agricultural land is jointly owned by hundreds of owners, each of whom owns a different share. After the construction, each owner should receive an asset with value proportional to their original share. Subject to this, it is required to minimize the amount of monetary transfers. The second application involves "evacuation and rebuilding" projects, where an old storeyhouse is destroyed and a new one is built instead. Each apartment owner in the original storeyhouse should receive an apartment in the new house, but the new houses are different, which might cause envy.

Based on joint work with Rica Gonen, Noga Klein and Eitan Lichtman.

Back to Contents Page →

Alexandros Hollender
*University of Oxford, UK*

## *Envy-Free Cake-Cutting for Four Agents*

In the envy-free cake-cutting problem we are given a resource, usually called a cake and represented as the [0,1] interval, and a set of n agents with heterogeneous preferences over pieces of the cake. The goal is to divide the cake among the n agents such that no agent is envious of any other agent. Even under a very general preferences model, this fundamental fair division problem is known to always admit an exact solution where each agent obtains a connected piece of the cake; we study the complexity of finding an approximate solution, i.e., a connected ε-envy-free allocation.

For monotone valuations of cake pieces, Deng, Qi, and Saberi (2012) gave an efficient algorithm for three agents, and it was conjectured by Brânzei and Nisan (2022) that the problem for four agents is hard. We present an efficient algorithm for the case of four agents. We also prove that as soon as the valuations are allowed to be non-monotone, the problem becomes hard, even in the communication model.

Based on joint work with Aviad Rubinstein.

Back to Contents Page →

Yasushi Kawase
*The University of Tokyo, Japan*

## *The Random Assignment Problem Under Constraints*

We investigate the problem of random assignment of indivisible goods, in which each agent has an ordinal preference and a constraint. Our goal is to characterize the conditions under which there always exists a random assignment that simultaneously satisfies efficiency and envy-freeness. The probabilistic serial mechanism ensures the existence of such an assignment for the unconstrained setting. In this paper, we consider a more general setting in which each agent can consume a set of items only if the set satisfies her feasibility constraint. Such constraints must be taken into account in student course placements, employee shift assignments, and so on. We demonstrate that an efficient and envy-free assignment may not exist even for the simple case of partition matroid constraints, where the items are categorized, and each agent demands one item from each category. We then identify special cases in which an efficient and envy-free assignment always exists. For these cases, the probabilistic serial cannot be naturally extended; therefore, we provide mechanisms to find the desired assignment using various approaches.

This is joint work with Hanna Sumita and Yu Yokoi.

Back to Contents Page →

Jérôme Lang
*Université Paris Dauphine, France*

## *How fair can strategy-proof fair division be?*

When allocating indivisible items to agents, it is known that the only strategy-proof 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). With these rules, 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 a particular emphasis on the Plackett-Luce model. We conclude with experimental results illustrating how the optimal sequence is impacted by various parameters of the problem.

This is joint work with Sylvain Bouveret, Hugo Gilbert and Guillaume Méroué.

Back to Contents Page →

Bo Li
*The Hong Kong Polytechnic University, Hong Kong*

## *MMS Allocation of Indivisible Chores with Subadditive Valuations and the Fair Surveillance Assignment Problem*

We study the maximin share (MMS) fair allocation of m indivisible chores to n agents who have costs for completing the assigned chores. It is known that exact MMS fairness cannot be guaranteed, and so far the best-known approximation for additive cost functions is 13/11 by Huang and Segal-Halevi [EC, 2023]; however, beyond additivity, very little is known. In this work, we first prove that no algorithm can ensure better than $\min\{n, \log m/\log\log m\}$-approximation if the cost functions are submodular. This result also shows a sharp contrast with the allocation of goods where constant approximations exist as shown by Barman and Krishnamurthy [TEAC, 2020] and Ghodsi et al. [AIJ, 2022]. We then prove that for subadditive costs, there always exists an allocation that is $\min\{n, \lceil \log m \rceil\}$-approximation, and thus the approximation ratio is asymptotically tight. Due to the hardness result for general subadditive costs, we turn to study specific subadditive costs, e.g., vertex cover, which is called the fair surveillance assignment problem, and more. For these settings, we show that constant approximate allocations exist.

Back to Contents Page →

Xinhang Lu
*University of New South Wales, Australia*

## *Best-of-Both-Worlds Fair Allocation of Indivisible and Mixed Goods*

We study the problem of fairly allocating either a set of indivisible goods or a set of mixed divisible and indivisible goods (i.e., mixed goods) to agents with additive utilities, taking the best-of-both-worlds perspective of guaranteeing fairness properties both ex ante and ex post. The ex-post fairness notions considered in this paper are relaxations of envy-freeness, specifically, EFX for indivisible-goods allocation, and EFM for mixed goods allocation. For two agents, we show that there is a polynomial-time randomized algorithm that achieves ex-ante envy-freeness and ex-post EFX / EFM. For n agents with bi-valued utilities, we show that there exist randomized allocations that are (i) ex-ante proportional and ex-post EFM, and (ii) ex-ante envy-free, ex-post EFX, and ex-post fractionally Pareto optimal, which can be sampled in polynomial time.

Joint work with Xiaolin Bu, Zihao Li, Shengxin Liu, and Biaoshuai Tao

Back to Contents Page →

Ryoga Mahara
*The University of Tokyo, Japan*

## *Proportionality up to the Least Valued Good on Average*

We study the problem of fairly allocating a set of indivisible goods to multiple agents and focus on proportionality, which is one of the classical fairness notions. Since proportional allocations do not always exist when goods are indivisible, approximate concepts of proportionality have been considered in the previous work. Among them,proportionality up to the maximin good (PROPm) has been the best approximate notion of proportionality that can be achieved for all instances. In this talk, we introduce the notion of proportionality up to the least valued good on average (PROPavg), which is a stronger notion than PROPm. We show that a PROPavg allocation always exists for all instances and can be computed in polynomial time. Our results establish PROPavg as a notable nontrivial fairness notion that can be achieved for all instances. Our proof is constructive, and based on a new technique that generalizes the cut-and-choose protocol and uses a recursive technique.

Back to Contents Page →

Pasin Manurangsi
*Google Research, Thailand*

---

*Fair Division for Random Utilities*

---

In recent years, several works have studied fair division in a random model where the agents' utilities for individual items are drawn at random from a probability distribution. A typical question is to determine when the allocation exists (with high probability); this question has been raised for many fairness notions and both in the individual setting--where each bundle is given to a single agent--and in the group setting--where each bundle is given to a group of agents. In this talk, we will survey the results and techniques from this line of work.

Based on joint works with Warut Suksompong

Back to Contents Page →

Hanna Sumita
*Tokyo Institute of Technology, Japan*

## *Fair Allocation with Binary Valuations for Mixed Divisible and Indivisible Goods via Hybrid Convex Optimization*

We study the fair allocation of mixed goods, consisting of both divisible and indivisible goods, to agents with binary additive valuations. This is a prominent topic of recent study in fair allocation. We define an allocation as fair if its utility vector minimizes a symmetric strictly convex function over utilitarian optimal allocations. This fairness criterion includes standard ones such as maximum egalitarian social welfare and maximum Nash social welfare. We address the problem of minimizing a given symmetric strictly convex function when agents have binary valuations. If only divisible goods or only indivisible goods exist, the problem is known to be solvable in polynomial time. In this talk, we demonstrate that the problem is NP-hard even when all indivisible goods are identical, while we provide a polynomial-time algorithm for the problem when all divisible goods are identical. The key technique is to view our problem from the perspective of convex optimization. Our problem is formulated as the problem of minimizing a given symmetric strictly convex function over the Minkowski sum of an integral base-polyhedron and an M-convex set. By exploiting the structures of this hybrid of continuous and discrete problems, we derive the computational results for mixed goods. We also discuss a connection to a relaxed envy-freeness for mixed goods.

This talk is based on joint work with Yasushi Kawase and Koichi Nishimura.

Back to Contents Page →

Biaoshuai Tao
*Shanghai Jiao Tong University, China*

## *Truuthful and Almost Envy-Free Mechanism of Allocating Indivisible Goods: the Power of Randomness*

We study the problem of *fairly* and *truthfully* allocating $m$ indivisible items to $n$ agents with additive preferences. Specifically, we consider truthful mechanisms outputting allocations that satisfy EF^{+u}_{-v}, where, in an EF^{+u}_{-v} allocation, for any pair of agents $i$ and $j$, agent $i$ will not envy agent $j$ if u items were added to $i$'s bundle and $v$ items were removed from $j$'s bundle. Previous work easily indicates that, when restricted to deterministic mechanisms, truthfulness will lead to a poor guarantee of fairness: even with two agents, for any $u$ and $v$, EF^{+u}_{-v} cannot be guaranteed by truthful mechanisms when the number of items is large enough. In this work, we focus on randomized mechanisms, where we consider *ex-ante* truthfulness and *ex-post* fairness. For two agents, we present a truthful mechanism that achieves EF^{+0}_{-1} (i.e., the well-studied fairness notion EF1). For three agents, we present a truthful mechanism that achieves EF^{+1}_{-1}. For $n$ agents in general, we show that there exist truthful mechanisms that achieve EF^{+u}_{-v} for some $u$ and $v$ that depend only on $n$ (not $m$).

We further consider fair and truthful mechanisms that also satisfy the standard efficiency guarantee: Pareto-optimality. We provide a mechanism that simultaneously achieves truthfulness, EF1, and Pareto-optimality for bi-valued utilities (where agents' valuation on each item is either $p$ or $q$ for some $p>q>0$). For tri-valued utilities (where agents' valuations on each item belong to $\{p,q,r\}$ for some $p>q>r>0$) and any $u,v$, we show that truthfulness is incompatible with EF^{+u}_{-v} and Pareto-optimality even for two agents.

This is a joint work with Xiaolin Bu. The link to our paper can be found here: https://arxiv.org/abs/2407.13634

Back to Contents Page →

Nicholas Teh
*University of Oxford, UK*

## *Temporal Fair Division of Indivisible Items*

We study a fair division model where indivisible items arrive sequentially, and must be allocated immediately and irrevocably. Previous work on online fair division has shown impossibility results in achieving approximate envy-freeness under these constraints. In contrast, we consider an informed setting where the algorithm has complete knowledge of future items, and aim to ensure that the cumulative allocation at each round satisfies approximate envy-freeness--which we define as temporal envy-freeness up to one item (TEF1). We focus on settings where items can be exclusively goods or exclusively chores. For goods, while TEF1 allocations may not always exist, we identify several special cases where they do---two agents, two item types, generalized binary valuations, unimodal preferences---and provide polynomial-time algorithms for these cases. We also prove that determining the existence of a TEF1 allocation is NP-hard. For chores, we establish analogous results for the special cases, but present a slightly weaker intractability result. We also establish the incompatibility between TEF1 and Pareto-optimality, with the implication that it is intractable to find a TEF1 allocation that maximizes any $p$-mean welfare, even for two agents.

(Joint work with Edith Elkind, Alexander Lam, Mohamad Latifian, and Tzeh Yuan Neoh)

Back to Contents Page →

Rohit Vaish
*Indian Institute of Technology Delhi, India*

---

*Fair Interval Scheduling of Indivisible Chores*

---

We will discuss the problem of fairly assigning a set of discrete tasks, or chores, among a set of agents. Each chore has a designated start and finish time, and each agent can perform at most one chore at any given time. We will explore the existence and computation of "fair" (specifically, envy-free up to one chore) and "efficient" (specifically, maximal or Pareto optimal) schedules under various settings. The presentation will cover novel technical ideas, including a color-switching technique and an application of the "cycle-plus-triangles" theorem (originally conjectured by Erdős) for achieving approximate envyfreeness. We will also highlight several open problems and directions for future work.

Joint work with Sarfaraz Equbal, Rohit Gurjar, Yatharth Kumar, Swaprava Nath, and Raghuvansh Saxena.

URL: https://arxiv.org/abs/2402.04353

Back to Contents Page →

Adrian Vetta
*McGill University, Canada*

---

*Six Candidates Suffice to Win a Voter Majority*

---

Given an election of n voters with preference lists over m candidates, Elkind, Lang, and Saffidine (2011) defined a Condocet winning set to be a collection of candidates that the majority of voters prefer over any individual candidate. Condocet winning sets of cardinality one (a Condorcet winner) or cardinality two need not exist. We prove however that a Condocet winning set of cardinality at most six exists in any election.

(Joint work with M. Charikar, P. Ramakrishnan, A. Lassota, A. Vetta and K. Wang)

Back to Contents Page →

Xiaowei Wu
*University of Macau, Macau*

---

*Fair Allocation of Chores with Subsidy*

---

The fair allocation problem has gained significant attention recently in the fields of theoretical computer science, artificial intelligence, and economics. In this presentation, I will discuss our latest research on ensuring fairness for the allocation of chores using subsidies. We consider the allocation of m indivisible chores among n agents with subsidies. Specifically, we focus on scenarios where agents have additive cost functions and assume that the maximum cost of an item to an agent can be offset by one dollar, we show that a total subsidy of n/4 dollars is sufficient to achieve a proportional allocation. Furthermore, we prove that n/4 is the minimum necessary subsidy, as there exists an instance with n agents where any proportional allocation requires at least n/4 dollars in subsidies. Additionally, we explore the weighted case and show that a total subsidy of n/3 dollars is sufficient to ensure weighted proportionality.

Back to Contents Page →

Sheung Man Yuen
*National University of Singapore, Singapore*

---

*Attaining EF1 Allocations by Exchanging Goods*

---

In the allocation of indivisible goods, a prominent fairness notion is envyfreeness up to one good (EF1). We initiate the study of reachability problems in fair division by investigating the problem of whether one EF1 allocation can be reached from another EF1 allocation via a sequence of exchanges such that every intermediate allocation is also EF1. In circumstances where this can be done, we investigate whether there is also an optimal sequence of such exchanges. Another problem that we study is the reformation of an unfair allocation into an EF1 allocation via such sequences. We investigate the complexity of deciding whether this reformation process is possible and the complexity of computing the number of exchanges needed whenever this is possible. We provide bounds to the number of exchanges required in the reformation process in the worst case.

Back to Contents Page →

William Zwicker
*Union College, USA*

---

*Cutsets and EF1 fair division of graphs, continued*

---

In fair division of a connected graph G = (V, E), each of n agents receives a share of G's vertex set V. The shares partition V, with each share required to induce a connected subgraph. By cutting the graph into "too many pieces," graph cutsets serve as forbidden substructures. For a graph that contains a cutset, divisions that are fair in the EF1 (envy-free up to one item) sense are blocked, for certain values of n. Two parameters – gap and valence – determine these blocked values. For some graphs G we can now, with help from some new positive results, pin down G's spectrum – the list of exactly which values of n do/ do not guarantee connected EF1 allocations. We also provide an example of a (non-traceable) graph H on eight vertices that has no cutsets, yet fails to guarantee connected EF1 allocations for three agents with additive preferences. Can the cutset notion be generalized to explain this result on H? Yes. Might there be some ultimate generalization, that pins down the spectrum of every graph? Good question!