
Workshop
“Mechanism Design Without Transfers”

June 15 - 19, 2026

organized by
Anna Bogomolnaia, Florian Brandl

Abstracts

Péter Biró (HUN-REN Centre for Economic and Regional Studies)

Smart Lotteries in School Choice: Ex-Ante Pareto-Improvement with Ex-Post Stability
(joint with Haris Aziz, Gergely Csáji, and Tom Demeulemeester)

Abstract: In a typical school choice application, the students have strict preferences over the schools while the schools have coarse priorities over the students based on their distance and their enrolled siblings. The outcome of a centralized admission mechanism is then usually obtained by the Deferred Acceptance (DA) algorithm with random tie-breaking. Therefore, every possible outcome of this mechanism is a stable solution for the coarse priorities that will arise with certain probability. This implies a probabilistic assignment, where the admission probability for each student-school pair is specified. In this paper, we propose a new efficiency-improving stable “smart lottery” mechanism. We aim to improve the probabilistic assignment ex-ante in a stochastic dominance sense, while ensuring that the improved random matching is still ex-post stable, meaning that it can be decomposed into stable matchings regarding the original coarse priorities. Therefore, this smart lottery mechanism can provide a clear Pareto-improvement in expectation for any cardinal utilities compared to the standard DA with lottery solution, without sacrificing the stability of the final outcome. We show that although the underlying computational problem is NP-hard, we can solve the problem by using advanced optimization techniques such as integer programming with column generation. We conduct computational experiments on generated and real instances. Our results show that the welfare gains by our mechanism are substantially larger than the expected gains by standard methods that realize efficiency improvements after ties have already been broken.

Anna Bogomolnaia (University of Glasgow)

Discrete Fair Division Under Objective Constraints

Abstract: We consider a discrete fair division model, where n agents are to divide a finite set of objects, under the additional assumption that each agent i is to receive a basket of pre-determined size q_i (her constraint, or “quota”). Objects are indivisible, and no transfers or randomization are allowed. Agents’ preferences are additive and publicly known. We are looking for an efficient and fair allocation.

We analyze both the ordinal case (only agents’ ordinal orderings of individual objects are known) and the cardinal case (full preferences are known). Full fairness is generally out of reach in discrete models,

so traditionally the literature considers approximate fairness (proportionality and envy-freeness) “up to one good”. We propose new notions of approximate fairness, more appropriate for the setting with quotas, “up to one upgrade”, which are stronger than the original ones. We also introduce a new notion of simultaneous fairness up to one circular exchange, or up to k circular exchanges.

For the ordinal case, we show that there exists an essentially unique ordinally efficient and ordinally fair allocation rule. Ordinal efficiency is weaker than full efficiency, but it is the only one that can be fully guaranteed under incomplete information on preferences. Ordinal fairness is much stronger than the standard one.

For the cardinal case, we show the existence of fully efficient and proportional up to one upgrade allocations, as well as the existence of efficient allocations that are simultaneously envy-free up to $n - 1$ circular exchanges.

Felix Brandt (Technical University of Munich)

Randomized Mechanisms Beyond Expected Utility

Abstract: This talk surveys results on three key questions in microeconomic theory when lotteries are permitted: Arrovian aggregation, strategyproof social choice, and fair assignment. Agents’ preferences are represented via skew-symmetric bilinear utility functions, a rich generalization of classic von Neumann–Morgenstern expected utility introduced by Fishburn.

Yeon-Koo Che (Columbia University)

Dynamic Market Design

Abstract: My lecture will feature a new framework for analyzing optimal resource allocation in stochastic dynamic environments which can be applied to both NTU and TU settings. In keeping with the week’s theme, I will illustrate two applications—(1) optimal queue design and (2) dynamic resource allocation with costly verification.

Jacob Coreno (University of Lausanne)

Axioms for Top Trading Cycles in Multi-Object Reallocation

(joint with Di Feng)

Abstract: This paper studies multi-object reallocation without monetary transfers, where agents initially own multiple indivisible objects and have strict preferences over bundles (e.g., shift exchange among workers at a firm). Focusing on marginal rules that elicit only rankings over individual objects, we provide axiomatic characterizations of the generalized Top Trading Cycles rule (TTC) on the lexicographic and responsive domains. On the lexicographic domain, TTC is characterized by balancedness, individual-good efficiency, the worst-endowment lower bound, and either truncation-proofness or drop strategy-proofness. On the responsive domain, TTC is the unique marginal rule satisfying individual-good efficiency, truncation-proofness, and either the worst-endowment lower bound or individual rationality. In the Shapley–Scarf housing market, TTC is characterized by Pareto efficiency, individual rationality, and truncation-proofness. Finally, on the conditionally lexicographic domain, the augmented Top Trading Cycles rule is characterized by balancedness, Pareto efficiency, the worst-endowment lower bound, and drop strategy-proofness. The conditionally lexicographic domain is a maximal domain on which Pareto efficiency coincides with individual-good efficiency.

Tangren Feng (Bocconi University)

Interim Strategy-Proof Voting

(joint with Qinggong Wu)

Abstract: Interim strategy-proofness (ISP) is a mechanism design criterion that extends strategy-proofness to environments with interdependent values. It requires that each agent have an interim dominant strategy—one that is optimal against any strategies of others, given her beliefs about their types. We characterize ISP mechanisms for binary collective choices without transfers: a mechanism is ISP if and only if it satisfies Ordinality, Monotonicity, and a novel condition, Limited Outcome Externality (LOE). LOE bounds how much an agent’s impact on the outcome can vary with others’ actions, with the permissible variation inversely related to the magnitude of informational externalities faced by her in the environment. We apply the characterization to the jury model and identify optimal ISP voting rules. Under ISP, a weak Condorcet Jury Theorem holds, whereas the strong version fails.

Ravi Jagadeesan (Stanford University)

Publication Design with Incentives in Mind

(joint with Davide Viviano)

Abstract: The publication process both determines which research receives the most attention, and influences the supply of research through its impact on researchers’ private incentives. We introduce a framework to study optimal publication decisions when researchers can choose (i) whether or how to conduct a study and (ii) whether or how to manipulate the research findings (e.g., via selective reporting or data manipulation). When manipulation is not possible, but research entails substantial private costs for the researchers, it may be optimal to incentivize cheaper research designs even if they are less accurate. When manipulation is possible, it is optimal to publish some manipulated results, as well as results that would have not received attention in the absence of manipulability. Even if it is possible to deter manipulation, such as by requiring pre-registered experiments instead of (potentially manipulable) observational studies, it is suboptimal to do so when experiments entail high research costs. We illustrate the implications of our model in an application to medical studies.

Deniz Kattwinkel (UCL)

TBA

Abstract: TBA

Patrick Lahr (ENS Paris-Saclay)

Extreme Points in Multi-Dimensional Screening

(joint with Axel Niemeyer)

Abstract: We characterize the extreme points of the set of incentive-compatible mechanisms for screening problems with linear utility. Our framework subsumes problems with and without transfers, such as monopoly pricing, principal-optimal bilateral trade and barter exchange, delegation and veto bargaining, or belief elicitation via proper scoring rules. In every problem with one-dimensional types, extreme points admit a tractable description. In every problem with multi-dimensional types, extreme points are dense in a rich subset of incentive-compatible mechanisms, which we call exhaustive

mechanisms. Building on these characterizations, we derive parallel conclusions for mechanisms that can be rationalized as (uniquely) optimal under a fixed objective. For example, in the multi-good monopoly problem, mechanisms that uniquely maximize revenue for some type distribution are dense among all incentive-compatible and individually rational mechanisms. The proofs exploit a novel connection between menus of extreme points and indecomposable convex bodies, first studied by Gale (1954).

Patrick Lederer (University of Amsterdam)

Approximate Strategyproofness in Approval-Based Budget Division

(joint with Jeremy Vollen and Haris Aziz)

Abstract: In approval-based budget division, the task is to allocate a divisible resource to the candidates based on the voters' approval preferences over the candidates. For this setting, Brandl et al. [2021] have shown that no distribution rule can be strategyproof, efficient, and fair at the same time. In this paper, we aim to circumvent this impossibility theorem by focusing on approximate strategyproofness. To this end, we analyze the incentive ratio of distribution rules, which quantifies the maximum multiplicative utility gain of a voter by manipulating. While it turns out that several classical rules have a large incentive ratio, we prove that the Nash product rule has an incentive ratio of 2, thereby demonstrating that we can bypass the impossibility of Brandl et al. by relaxing strategyproofness. Moreover, we show that an incentive ratio of 2 is optimal within three natural classes of rules and that the positive result for the Nash product rule even holds when voters may report arbitrary concave utility functions. Finally, we complement our results with an experimental analysis.

Hervé Moulin (University of Glasgow)

Compromises in Collective Choice: The Veto Solution

(joint with Anna Bogomolnaia)

Abstract: The tyranny of the majority disqualifies the Condorcet consistency principle in many contexts (typically not political elections) where individual agents (live persons or abstract criteria) seek a compromise between their different opinions. The Veto Core (Moulin 1981) finds such compromises by measuring the “objections” by coalitions of any size to each outcome, and eliminating those eliciting a “more than proportional” objection. Our first result generalises the Veto Core to choice problems with infinite sets of agents and/or outcomes, and rules that are neither Anonymous nor Neutral. As a choice rule the Veto Core enjoys very strong axiomatic properties: Maskin Monotonicity, Population Consistency and (the new) Clone Responsiveness. It typically selects a substantial short list rather than a deterministic choice. A natural selection from the Veto Core chooses the outcome(s) minimising the largest coalitional objection; it is Monotonic but loses the three strong properties above.

Axel Niemeier (Caltech)

Optimal Allocation with Peer Information

(joint with Justus Preusser)

Abstract: We study allocation problems without monetary transfers where agents have correlated types, i.e., hold private information about one another. Such peer information is relevant in various settings, including science funding, allocation of targeted aid, or intra-firm allocation. Incentive compatibility requires that agents cannot improve their own allocation by misrepresenting the merits of

allocating to others. We characterize optimal incentive-compatible mechanisms using techniques from the theory of perfect graphs. Optimal mechanisms improve on review panels commonly observed in practice by eliciting information directly from eligible agents and by using allocation lotteries to alleviate incentive constraints. Computational hardness results imply that exactly optimal mechanisms are impractically complex. We propose ranking-based mechanisms as a viable alternative and show that they are approximately optimal when agents are informationally small, i.e., when no single agent has information that is crucial for evaluating a large fraction of the other agents.

Szilvia Papai (Concordia University)

Pareto-Efficient Matching with Gradual Acceptances

(joint with Israa Hashem)

Abstract: We introduce the class of Gradual Acceptance rules, which assign objects to agents through a gradual procedure in which applicants are accepted selectively, rather than immediately accepting all current highest-priority applicants at each step, as in the Immediate Acceptance rule. The central idea is to slow down the acceptance process so that applicants can be accepted without creating priority violations whenever possible, while making permanent assignments at each step and hence maintaining Pareto-efficiency. We develop the secure-set concept and a priority-profile reduction procedure, leading to the Secure Gradual Acceptance rule, which is Pareto-efficient at every profile and selects a stable and Pareto-efficient matching whenever such a matching exists. Our main result uses these ideas to characterize, through a sequential procedure that imposes joint restrictions on preferences and priorities, the profiles at which Pareto-efficiency and stability are compatible. We also show that gradual acceptance can improve the stability and incentive properties of the Immediate Acceptance rule while preserving Pareto-efficiency.

William Phan (North Carolina State University)

Equal Opportunity in School Choice Lotteries

(joint with Acelya Altuntas and Yuki Tamura)

Abstract: In Berlin and North Carolina's Wake County school districts, all applying students are guaranteed a minimum chance of admission. This directly conflicts with the central property of respect for priorities in school choice. Our contribution is to formalize equal opportunity as a design property and study its implications. We relax other important properties to ensure compatibility, and show that existing mechanisms typically fail them. We propose probabilistic generalizations of three key mechanisms in the literature and evaluate their properties. Our generalization of Kesten (2010)'s mechanism performs best, lying on the efficient frontier given equal opportunity and priority constraints.

Marek Pycia (University of Zurich)

Evaluating with Statistics: Which Outcome Measures Differentiate Among Matching Mechanisms?

Abstract: The selection of mechanisms to allocate school seats in public school districts can be highly contentious. At the same time the standard statistics of student outcomes calculated from districts' data are very similar for many mechanisms. This paper contributes to the debate on mechanism selection by explaining the similarity puzzle as being driven by the invariance properties of the standard

outcome statistics: outcome measures are approximately similar if and only if they are approximately anonymous.

Joseph Root (University of Chicago)

How to Beat FCFS

(joint with Itai Ashlagi)

Abstract: We consider two homogeneous queues competing to attract agents who arrive stochastically. Agents choose which queue to join strategically to minimize their expected wait time. Both queues have identical service rates so can only compete by modifying their service rule, the order in which they serve the agents in their waitlist. We show that, when the queues can commit to their service rules, the most commonly used rule, first-come-first-served (FCFS) rule, is not an equilibrium. A simple alternative can be used to attract more than half of the arrivals. We bound the maximal arrival share for any non-preemptive rule competing against FCFS. Simulations illustrate the advantage of our rule, yet a substantial gap persists with the maximal bound. We also introduce a preemptive service rule that outperforms FCFS and simulations illustrate it gains a larger advantage. The picture changes when the queues cannot commit to their service rules. In this case, FCFS-FCFS is an equilibrium.

Jay Sethuraman (Columbia University)

TBA

Abstract: TBA

Olivier Tercieux (Paris School of Economics)

Dynamic Market Design in Large Economies: Foundations and Applications to the Design of Dynamic Incentives

(joint with César Barilla and Charlotte Martres)

Abstract: This paper develops a framework for dynamic matching markets without transfers. We study overlapping-generations environments in which, at each date, a finite set of institutions on one side of the market is matched to a continuum of agents on the other side, and a stable mechanism is operated in every period. We first characterize the stationary outcomes of these processes using stationary market-clearing cutoffs. We show that the stationary aggregate demand naturally violates the strict gross substitutes condition. As a result, structural properties that hold in static matching environments do not generally extend to the dynamic setting. Second, we provide foundations for stationary market-clearing cutoffs. Treating the distribution of agents' priorities and its evolution over time as design instruments, we show that any equilibrium outcome in a broad class of mechanisms can be implemented by the Deferred Acceptance (DA) mechanism. We also establish convergence results that clarify when the continuum model provides a good approximation to large but finite markets. Finally, we study dynamic incentive schemes—common in practice, for example in the assignment of civil servants. We characterize the trade-offs involved in using such dynamic incentives to achieve more equitable distributions of agents across institutions.

Alexander Westkamp (University of Cologne)

Marginal Mechanisms for Balanced Exchange

(joint with Vikram Manjunath)

Abstract: We study balanced exchange problems in which agents with responsive preferences are initially endowed with multiple indivisible objects and can trade without transfers, as in shift exchange or time-banking. In many such settings, eliciting or processing full preferences over bundles is infeasible and mechanisms instead often rely solely on marginal preferences, that is, rankings of individual objects.

We characterize when eliciting only marginal preferences is enough to identify allocations that are unambiguously efficient and individually rational in the sense that these properties hold with respect to any responsive preferences consistent with the elicited marginals. We parameterize domains of marginal preferences by which indifference classes can contain endowed and non-endowed objects. We show that the essentially unique maximal domain for which an unambiguously efficient and individually rational marginal mechanism exists is trichotomous: Agents rank objects in three tiers, with the bottom tier not containing any endowed objects.

We also consider agents' incentives for truthful preference revelation. The maximal domain for which an efficient, individually rational, and strategy-proof mechanism, marginal or not, exists is strongly trichotomous: Agents rank objects in three tiers, with the bottom tier not containing any endowed objects and the middle tier not containing any non-endowed objects. The canonical marginal mechanism that unambiguously achieves our three desiderata on that domain is a serial dictatorship over the set of individually rational allocations. Interestingly, when employed on the larger trichotomous domain, this serial dictatorship mechanism still has a weakly dominant strategy: Reveal the top preference tier truthfully and do not reveal any non-endowed objects in the middle tier. We propose a new family of gradual-revelation mechanisms that are also unambiguously efficient and individually rational on the trichotomous domain, while providing "better" incentives for the truthful revelation of all three preference tiers.