Switching Random Priority: A Draft Without Predictable Turns

Inácio Bó and Chunyu Zhao[1]

November 2026

Abstract

Course seats, work shifts, players in a league, and the contents of an estate are often divided by letting agents take turns choosing. Under a serial dictatorship, letting each agent pick all of her items at once protects truthful reporting but concentrates widely desired items in the hands of early pickers. Alternating turns, as in a draft, limits this advantage, but agents know when they will pick again and can gain by strategically postponing favorite items in order to secure contested ones. We propose a minimal change to the draft: after each pick, the turn passes to an agent drawn uniformly at random from all agents except the one who just picked. With at least three agents, the resulting *Switching Random Priority* mechanism is weakly SD-envy-free and satisfies weak SD-strategy-proofness, a standard incentive criterion in random assignment. Every realized allocation satisfies *bounded rank concentration*, an ex-post limit on the concentration of highly ranked items, which drafts also satisfy and serial dictatorships can violate. In simulations with Harvard Business School course-allocation data, SRP restores most of the individual welfare lost to manipulation in the observed draft and outperforms random serial dictatorship both with and without quotas, at the cost of higher utility inequality than the draft.

Keywords: market design, multi-object allocation, draft mechanisms, random assignment, fairness, strategy-proofness.
JEL classification: C78, D47, D63.

[1]Department of Economics, Faculty of Social Sciences, University of Macau. Emails: inaciobo@um.edu.mo; yc37336@um.edu.mo.

1 Introduction

Many allocation problems are solved by letting people take turns. Students choose elective courses in rounds (Budish and Cantillon, 2012a), sports and fantasy leagues distribute players through drafts, workers choose shifts and tasks from a common schedule, and heirs divide a family's belongings by picking items sequentially. In each case, participants receive several items, monetary transfers are unavailable, and participants must accept the procedure. Who ends up with the items everyone wants, and how participants decide what to pick, then depend on the rule that decides whose turn it is. A common way to organize such choices is a draft: participants are placed in an order and pick one item at a time, cycling through the order until the items run out. This paper studies a minimal change to the draft: instead of following a fixed order, the turn passes after each pick to a participant drawn at random from everyone except the one who just chose. The resulting mechanism prevents anyone from accumulating a run of top items, while simultaneously eliminating the profitable misreports that standard drafts allow.

A natural starting point is random serial dictatorship (RSD): participants are ordered randomly and pick sequentially from the remaining objects according to that order. When there are no restrictions on the number of objects a participant may receive and all objects are desirable, the first participant can take all the objects, leaving none for the others. While random order gives every participant the same chance of picking first, its realized allocations are maximally unequal in individual welfare. One response to this issue is to impose quotas. Under quota-RSD, participants are placed in a random order, and each consumes her entire quota by sequentially selecting her most preferred remaining objects in a single block. Quota-RSD is strong SD-strategy-proof[2] under additively separable preferences (Kojima, 2009). Yet while quotas eliminate quantity inequality, they do not equalize quality: earlier pickers sweep the most widely preferred objects before later agents have a turn, leaving the first-mover advantage of priority largely intact (Bouveret *et al.*, 2025).

Drafts address this concentration by spreading choices over rounds: a random order is repeated in a cyclic draft or reversed from one round to the next in a snake draft. This prevents the first picker from taking more than one item before everyone else picks. However, predictable future turns create opportunities for strategic manipulation (Brams and Straffin, 1979). Consider a student whose favorite course is one that few others want and whose second favorite is widely desired. In a draft, she knows that she will pick again before anyone else

[2]We distinguish three incentive properties. Strategy-proofness refers to standard dominant-strategy truthfulness, evaluated using expected utility for randomized outcomes or preferences over realized bundles for deterministic outcomes. Strong SD-strategy-proofness requires the truthful assignment to first-order stochastically dominate every deviation. Weak SD-strategy-proofness only rules out deviations that strictly stochastically dominate truthful reporting.

would take the unpopular course, so she does better by listing the popular course first and claiming her true favorite later (Bouveret and Lang, 2014; Coreno and Balbuzanov, 2026).

These concerns are not merely theoretical. In the course-allocation draft at Harvard Business School, Budish and Cantillon (2012a) find that students overstated their demand for popular courses and understated it for unpopular ones, and that this caused meaningful welfare losses. Yet they also find that the strong SD-strategy-proof alternative, random serial dictatorship, would have produced lower ex-ante welfare than the manipulated draft, because it lets students who draw an early position take popular courses without regard to how much those courses matter to others. When a procedure rewards strategic sophistication, participants who report truthfully, whether through inexperience or good faith, can be disadvantaged relative to those who strategize (Pathak and Sönmez, 2008).

This dilemma is not an accident of particular procedures. For deterministic mechanisms, strategy-proofness and Pareto efficiency, together with additional regularity conditions, single out serial dictatorships (Klaus and Miyagawa, 2002; Ehlers and Klaus, 2003; Hatfield, 2009), and the fairness properties that characterize draft rules are, in general, incompatible with strategy-proofness (Coreno and Balbuzanov, 2026). For randomized mechanisms in which each agent receives a fixed number of items, no mechanism combines efficiency, envy-freeness, and weak SD-strategy-proofness (Kojima, 2009).

We propose the *Switching Random Priority* (SRP) mechanism. Participants report rankings of individual items. The first picker is drawn at random; after each pick, the next picker is drawn with equal probability from all participants other than the one who just chose, and she receives her highest-ranked remaining item. Because no one ever picks twice in a row, no participant can accumulate a run of top items while others wait. And because no future turn is guaranteed, postponing a valued item is no longer riskless. The student in the example above can still list the popular course first, but with positive probability she does not pick again before her favorite course is gone, whereas by reporting truthfully she would have secured it.

We establish the following results for any number of items and at least three participants. First, SRP is *weakly SD-strategy-proof* (Theorem 5.2): no misreport yields a lottery over items that first-order stochastically dominates the one obtained by reporting truthfully. As formalized in Section 3.1, this is a standard incentive benchmark for random assignment without monetary transfers (Bogomolnaia and Moulin, 2001; Budish *et al.*, 2013; Katta and Sethuraman, 2006; Kojima, 2009). Second, the uniform randomization across eligible pickers ensures that SRP is anonymous and weakly SD-envy-free, so from each participant's perspective, no other participant's lottery stochastically dominates her own (Proposition 5.3). Third, every realized allocation satisfies a new ex-post fairness property, *bounded rank*

concentration: take any item a participant receives; the items in her bundle that she likes better than it cannot outnumber the items held by others that they like better than it (Proposition 5.6). Drafts satisfy this property; serial dictatorship, without quotas or with quotas of two or more items, do not. Finally, every realized allocation is *pairwise-dominance efficient*: no alternative allocation gives every participant weakly preferred replacements for all her items, with a strict improvement for some (Proposition 5.4).

Among the mechanisms we study, quota-RSD protects incentives but can violate bounded rank concentration, while standard drafts satisfy bounded rank concentration but violate even weak SD-strategy-proofness (Table 2). SRP is the only procedure combining both properties. It circumvents the impossibility results of Kojima (2009) by replacing ex-ante SD-efficiency with realized pairwise-dominance efficiency, and requiring weak rather than strong SD-envy-freeness.

We also evaluate SRP in counterfactual simulations using the elective course-allocation data of Budish and Cantillon (2012a) from Harvard Business School. When students report truthfully, SRP recovers most of the individual welfare lost to strategic play in the observed draft. Moreover, by forbidding consecutive picks, SRP prevents early pickers from exhausting popular courses before others can pick, achieving higher preference satisfaction and lower inequality than random serial dictatorship with or without quotas. However, it distributes utility less evenly than the draft.

The remainder of the paper is organized as follows. Section 2 reviews the related literature. Section 3 introduces the assignment problem and formalizes the criteria of incentives, fairness, and efficiency. Section 4 defines SRP and the benchmark mechanisms. Section 5 establishes our main theoretical results. Section 6 evaluates the mechanisms empirically using course-allocation data from Harvard Business School, and Section 7 concludes.

2 Literature Review

The literature on indivisible object allocation treats serial dictatorship as a central benchmark, where agents are processed in a fixed priority order, with each being assigned her most preferred set of remaining objects. For single-object allocation, Svensson (1999) shows that strategy-proofness, non-bossiness, and neutrality together characterize serial dictatorship. When generalized to multiple objects, the mechanism is characterized by combining strategy-proofness and Pareto efficiency with either population-monotonicity (Klaus and Miyagawa, 2002) or non-bossiness together with either resource-monotonicity (Ehlers and Klaus, 2003) or neutrality under quota constraints (Hatfield, 2009). While strategy-proof and Pareto optimal under the corresponding feasibility constraints, serial dictatorship can produce substantial ex-post inequality in multi-object settings. Although randomizing the priority order ensures

anonymity ex-ante and imposing quotas equalizes the number of objects received, quota-RSD does not eliminate the priority advantage of early pickers in accumulating top-tier objects (Kojima, 2009). Bouveret *et al.* (2025) further study how unequal quotas can compensate for this priority advantage.

Draft mechanisms, on the other hand, rotate the picking turn among agents rather than allowing each agent to select her entire quota in consecutive turns. In deterministic settings, standard drafts achieve pairwise efficiency under sincere play (which we formalize below as pairwise-dominance efficiency) but fail strategy-proofness (Brams and Straffin, 1979). Conversely, the booster draft (Caspari, 2026) ensures strategy-proofness by decomposing the allocation into independent picking sequences over disjoint subsets, but at the cost of pairwise efficiency. Coreno and Balbuzanov (2026) provide an axiomatic characterization of draft rules and establish impossibility results concerning the compatibility of their fairness properties with strategy-proofness. Empirically, Budish and Cantillon (2012a) study Harvard Business School's randomized course-allocation draft; they find that strategic manipulation in the draft caused significant welfare losses, yet the draft still outperformed random serial dictatorship in ex-ante welfare because serial dictatorship severely concentrates desirable courses among early pickers.

Several ex-post fairness concepts have been proposed for the allocation of indivisible objects. Amanatidis *et al.* (2017) demonstrate that no deterministic strategy-proof mechanism allocating all objects guarantees envy-freeness up to one good (EF1) for two agents with additive valuations when there are at least five objects. Other fairness notions include EFX (envy-free up to any item) (Caragiannis *et al.*, 2019), maximin share (MMS) fairness (Kurokawa *et al.*, 2018), and bounded envy under ordinal dominance relations (Caspari, 2026). Complementing these notions, we introduce *bounded rank concentration*, a property restricting the concentration of highly ranked objects within each realized allocation.

In probabilistic assignment, impossibility theorems demonstrate that SD-efficiency is incompatible with incentive compatibility and fairness. In the unit-demand setting, Bogomolnaia and Moulin (2001) document a fundamental trade-off: Random Serial Dictatorship is strong SD-strategy-proof and weakly envy-free but fails SD-efficiency, whereas the Probabilistic Serial mechanism achieves SD-efficiency and envy-freeness at the cost of strong SD-strategy-proofness while retaining weak SD-strategy-proofness. They prove that no mechanism can jointly satisfy SD-efficiency, strong SD-strategy-proofness, and equal treatment of equals for four or more agents. Subsequent work examines related incompatibilities in the presence of indifference and on restricted preference domains (Katta and Sethuraman, 2006; Kasajima, 2013). In settings where agents demand multiple objects, Kojima (2009) shows that SD-efficiency, SD-envy-freeness, and weak SD-strategy-proofness are incompatible under

fixed quotas and additively separable preferences. SRP satisfies weak SD-strategy-proofness and weak SD-envy-freeness, while ensuring pairwise-dominance efficiency and bounded rank concentration in every realized allocation ex-post.

3 Model

An assignment problem is a triple (N, O, P). $N = \{1, \ldots, n\}$ represents the set of agents, where $n \geq 3$. The set of goods is denoted as $O = \{o_1, o_2, \ldots, o_m\}$, comprising $m \geq 2$ indivisible objects. We use lowercase letters for individual objects and uppercase letters for bundles (subsets of O); in particular, the singleton bundle $\{o\}$ is distinct from the object o. Each agent $i \in N$ has preferences over bundles $S \subseteq O$, represented by a reflexive and transitive weak relation R_i and its strict part P_i, with a strict and complete ranking of individual objects. Our results use only this object ranking; additional assumptions on bundle preferences are stated where they are needed. We also use P_i and R_i for the induced ranking of individual objects: xP_iy means $\{x\}P_i\{y\}$, and similarly for R_i. For any object $o \in O$, let $\mathrm{Up}_i(o) = \{o' \in O : o'R_io\}$ denote agent i's upper-contour set at o, consisting of the objects weakly preferred to o.

A deterministic allocation $A = (A_i)_{i \in N}$ consists of pairwise disjoint bundles $A_i \subseteq O$ such that $\bigcup_{i \in N} A_i = O$, where A_i is the bundle assigned to agent i. A probabilistic allocation is a matrix $\pi = [\pi_i(o_j)]_{i \in N, o_j \in O}$, where $\pi_i(o_j) \in [0, 1]$ denotes the probability that object o_j is allocated to agent i. For every $o_j \in O$, $\sum_{i \in N} \pi_i(o_j) = 1$, so each object is allocated with probability one. Let Π be the set of all probabilistic allocations. For each agent i, her probabilistic assignment in $\pi \in \Pi$ is the vector $\pi_i = [\pi_i(o_j)]_{o_j \in O}$. For any $S \subseteq O$, write $\pi_i(S) = \sum_{o \in S} \pi_i(o)$. Let $\mathcal{P}_i$ denote the set of strict linear orders over O, and let $\mathcal{P} = \prod_{i \in N} \mathcal{P}_i$. A preference profile $P = (P_i)_{i \in N} \in \mathcal{P}$ specifies the agents' rankings of individual objects. These rankings are the only inputs to the mechanism. Preferences over bundles are evaluated in terms of these preferences over individual objects. A mechanism $\phi : \mathcal{P} \to \Pi$ maps each preference profile to a probabilistic allocation.

3.1 Desirable Properties

We evaluate probabilistic allocation mechanisms against properties of incentives, fairness, and efficiency under first-order stochastic dominance (SD). Consider two probabilistic allocations $\pi, \pi' \in \Pi$. We say π_i **weakly stochastically dominates** π'_i under preference P_i, denoted $\pi_i R_i^{\mathrm{sd}} \pi'_i$, if

$$\sum_{o' \in \mathrm{Up}_i(o)} \pi_i(o') \geq \sum_{o' \in \mathrm{Up}_i(o)} \pi'_i(o') \qquad \text{for every } o \in O.$$

We say π_i **strictly stochastically dominates** π_i', denoted $\pi_i P_i^{\text{sd}} \pi_i'$, if $\pi_i R_i^{\text{sd}} \pi_i'$ and the inequality is strict for at least one object $o \in O$. Since object rankings are strict, weak SD holds if and only if either the two assignment vectors are equal or strict SD holds.

Definition 3.1 (Weak SD-Strategy-Proofness). A mechanism ϕ is weakly SD-strategy-proof if for every agent $i \in N$, every preference profile $P \in \mathcal{P}$, and every misreport $P_i' \in \mathcal{P}_i$,

$$\begin{aligned} & \phi_i(P_i', P_{-i}) R_i^{\text{sd}} \phi_i(P_i, P_{-i}) \\ \implies \quad & \phi_i(P_i', P_{-i}) = \phi_i(P_i, P_{-i}). \end{aligned}$$

Equivalently, no misreport yields a strict SD improvement over truthful reporting.

Weak SD-strategy-proofness is a commonly used incentive criterion in the random-assignment literature. Bogomolnaia and Moulin (2001) use it to describe the incentive properties of the probabilistic serial mechanism, and Budish *et al.* (2013) show that their generalization of that mechanism to environments with constraints satisfies it. It is the incentive requirement in impossibility and characterization results for preferences with indifferences (Katta and Sethuraman, 2006; Aziz *et al.*, 2017), for multi-unit demands (Kojima, 2009), for fractional endowments (Athanassoglou and Sethuraman, 2011), for the probabilistic serial mechanism (Kesten *et al.*, 2017), and for randomized social choice (Brandl *et al.*, 2018); it is also the baseline against which intermediate incentive notions are defined (Balbuzanov, 2016; Mennle and Seuken, 2021).

Because object rankings are strict, the definition has a simple equivalent form: a mechanism is weakly SD-strategy-proof if and only if every misreport that changes agent i's probabilistic assignment strictly lowers $\pi_i(\text{Up}_i(o))$, the expected number of objects she receives among those she weakly prefers to o, for at least one object o. When preferences over bundles are additively separable, it follows that for each misreport that changes her assignment, some utility function consistent with her ranking over objects makes truthful reporting strictly better in expectation.

Definition 3.2 (Weak SD-Envy-Freeness). An allocation π is weakly SD-envy-free at preference profile P if, for any $i \in N$, there is no $j \in N$ such that $\pi_j P_i^{\text{sd}} \pi_i$. A mechanism ϕ is weakly SD-envy-free if $\phi(P)$ is weakly SD-envy-free for all $P \in \mathcal{P}$.

Definition 3.3 (Anonymity). A mechanism ϕ is anonymous if for every permutation ρ of N and profile P,

$$\phi_{\rho(i)}(P^\rho) = \phi_i(P) \quad \text{for all } i \in N,$$

where $P^\rho_{\rho(i)} = P_i$.

Definition 3.4 (Equal Treatment of Equals). A mechanism ϕ satisfies equal treatment of equals if for all agents $i, j \in N$ and every preference profile $P \in \mathcal{P}$, when $P_i = P_j$, then the allocations must satisfy: $\phi_i(P) = \phi_j(P)$.

Definition 3.5 (Pairwise-Dominance Efficiency). An allocation A is *pairwise-dominance efficient* if there is no other feasible allocation $B \neq A$ such that, for every agent i, there exists an injection $f_i : A_i \to B_i$ satisfying

$$f_i(x) R_i x \qquad \text{for every } x \in A_i.$$

In words, an allocation is pairwise-dominance efficient if there is no reallocation in which every agent can replace each object she receives with a distinct object she weakly prefers. The property requires only rankings of individual objects, without imposing preferences over bundles.

4 The Switching Random Priority Mechanism and Its Counterparts

We now introduce the SRP mechanism. It randomly selects an initial picker and then randomly selects an agent other than the current picker after each object has been allocated. This rule has two key features: no agent receives consecutive turns, and future picking turns are never predetermined. Let $\sigma_t \in N$ denote the agent assigned to round t. Given the reported rankings of individual objects in P, the mechanism operates in m rounds: at each round t, agent σ_t receives her most preferred available object according to P_{σ_t}. The turn sequence $\sigma = (\sigma_1, \ldots, \sigma_m)$ is generated by a Markov process:

- **Initial draw:** The initial agent σ_1 is drawn uniformly at random from N, so that $\Pr(\sigma_1 = i) = \frac{1}{n}$ for all $i \in N$.
- **Transition rule:** For each $t = 1, \ldots, m-1$, the next picker σ_{t+1} is drawn uniformly at random from the other $n-1$ agents:

$$\Pr(\sigma_{t+1} = j \mid \sigma_t = i) = \begin{cases} 0 & \text{if } j = i, \\ \dfrac{1}{n-1} & \text{if } j \neq i. \end{cases} \tag{4.1}$$

Let $\Omega = \{\sigma \in N^m : \sigma_{t+1} \neq \sigma_t \text{ for } t = 1, \ldots, m-1\}$ denote the set of all feasible picking sequences. Each $\sigma \in \Omega$ occurs with probability $1/[n(n-1)^{m-1}]$, independently of the reported preferences.

The following mechanisms will be treated as benchmarks to which we compare SRP.

Benchmark mechanisms. Each benchmark draws a permutation $\rho = (\rho_1, \ldots, \rho_n)$ of N uniformly at random, independently of P. At each round $t = 1, \ldots, m$, agent σ_t receives her most preferred available object according to P_{σ_t}, with the turn sequence defined as follows.

- **Unrestricted RSD.** Set $\sigma_t = \rho_1$ for every $t = 1, \ldots, m$. That is, the first agent selects all available objects before the next agent gets a turn.
- **Quota random serial dictatorship (quota-RSD).** Assume $m = kn$ for a positive integer k and set $\sigma_t = \rho_{\lceil t/k \rceil}$. Each agent thus receives k objects in consecutive turns.
- **Cyclic draft.** Set $\sigma_t = \rho_r$, where $r = 1 + ((t-1) \bmod n)$. This repeats the order ρ until all objects are allocated.
- **Snake draft.** Set σ equal to the sequence $(\rho_1, \ldots, \rho_n, \rho_n, \ldots, \rho_1)$, repeated until m turns have been assigned.

5 Incentives, Fairness, and Efficiency

We now examine the normative properties of the SRP mechanism under stochastic dominance, focusing on incentive compatibility, fairness, and efficiency.

We begin with strategy-proofness. A potential concern with draft mechanisms is that predictable future turns may allow agents to benefit from misreporting their preferences. The following example illustrates a profitable manipulation under cyclic draft.

Example 5.1 (Manipulating cyclic draft). There are three agents and six objects:

$$\begin{aligned} P_1 &: \quad a, b, c, d, e, f, \\ P_2 = P_3 &: \quad b, c, d, e, f, a. \end{aligned}$$

Suppose agent 1 swaps a and b in her report. If she is first in the randomly drawn order (probability 1/3), she receives $\{a, b\}$ instead of $\{a, d\}$; otherwise, her bundle is unchanged. The misreport therefore shifts marginal probability 1/3 from d to b, strictly improving her assignment under SD. The same misreport is profitable under the snake draft: when agent 1 is first in the drawn order, she picks in rounds 1 and 6 and receives $\{a, b\}$ instead of $\{a, f\}$; otherwise, her bundle is unchanged.

The example shows the key intuition behind manipulations in draft mechanisms. Agent 1 is guaranteed another turn before either of the other agents would pick a, which both rank last. Therefore, it is safe for her to postpone a and prioritize b instead, which is also desired by the other agents.

SRP removes this certainty. If agent 1 skips object a, there is a non-zero probability that she will not be able to get back to it later. This is the channel that makes SRP weakly SD-strategy-proof. The following theorem formalizes this intuition.

Theorem 5.2. *SRP is weakly SD-strategy-proof.*

See Appendix A.1 for the proof.

Beyond incentive compatibility, the SRP mechanism provides ex-ante fairness guarantees.

Proposition 5.3. *SRP satisfies anonymity, equal treatment of equals, and weak SD-envy-freeness.*

See Appendix A.2 for the proof.

Having established incentive compatibility and ex-ante fairness, we examine the efficiency of realized allocations.

Proposition 5.4. *Every realization of SRP, unrestricted RSD, quota-RSD, the cyclic draft, and the snake draft is pairwise-dominance efficient.*

This is a standard efficiency property of sincere sequential picking (Brams and Straffin, 1979; Coreno and Balbuzanov, 2026).

See Appendix A.3 for the proof.

5.1 Bounded Rank Concentration

The fairness guarantees in Proposition 5.3 concern the lottery induced by SRP. Draft mechanisms and the SRP, however, provide an intuitive ex-post fairness characteristic. Realized allocations are the result of agents taking turns picking objects. As a result, even for the SRP mechanism, in which the number of times that each agent will pick is stochastic, there is always an alternation. We formalize this by defining a new notion of ex-post fairness: *bounded rank concentration*.

Definition 5.5 (Bounded Rank Concentration). An allocation A satisfies *bounded rank concentration* if, for every agent i and every object $x \in A_i$,

$$|\{y \in A_i : yP_ix\}| \leq \sum_{j \neq i} |\{y \in A_j : yP_jx\}| .$$

In words, the condition can be visualized as a one-sided balance scale. Fix an object x received by agent i. On one pan are the objects in i's bundle that she ranks above x, and on the other are the objects held by everyone else that their recipients rank above x. Bounded rank concentration requires agent i's pan to never be the heavier one. It is thus a ceiling

on concentration, and a modest one, since each agent is compared with all other agents combined. Even so, random serial dictatorship with and without quotas violates it.

Proposition 5.6. *Every realization of SRP, the cyclic draft, and the snake draft satisfies bounded rank concentration. In contrast, unrestricted RSD and quota-RSD with a common quota $k \geq 2$ do not satisfy bounded rank concentration in general.*

See Appendix A.4 for the proof.

6 An Empirical Application: Course Allocation at Harvard Business School

We evaluate the performance of SRP using the course allocation environment and preference data originally assembled by Budish and Cantillon (2012a) for elective courses at Harvard Business School (HBS). We use the publicly available replication dataset (Budish and Cantillon, 2012b) to reconstruct the assignment problem, including course capacities and student preferences, and simulate allocations under alternative mechanisms. HBS fits our setting: students want several courses, no money is involved, and the school ran a sequential draft.

The data cover 456 students and 96 courses. Before the actual allocation, students answered a survey asking them to rank their five favorite courses truthfully. We construct truthful rankings by placing survey courses not flagged as likely preference changes first, using Budish and Cantillon's classification rules, and filling the remaining positions with courses in the order of the students' actual submitted rankings, without duplicates. Flagged courses can still enter through the submitted rankings, but receive no priority from the survey. These constructed rankings provide the common basis for evaluating outcomes under all mechanisms.

The empirical environment differs from our theoretical setting in some dimensions. First, courses have multiple seats, students can receive at most one seat in each course, and rankings are incomplete: we treat courses outside a student's constructed ranking as unacceptable.[3] Our objective is to compare the mechanisms under these common empirical constraints. To make the comparison as similar as possible across mechanisms, we remove the original HBS limit of ten courses per student and omit its aftermarket. Both HBS cases therefore use only the snake draft, without a limit on the number of courses a student can receive. Course seat capacities remain in place.

[3]Some of our theoretical results need additional assumptions to retain their analogous outcomes in this setting. In short, they require the number of students who deem a course acceptable to be sufficiently high. Bounded rank concentration holds with the modification that preference is weak instead of strict in the definition.

Mechanisms. We simulate six mechanisms over 10,000 independent draws, running all of them on the constructed truthful preferences except HBS Observed:

(i) *HBS Truthful*: the snake draft. Students pick in a random order in the first round, in the reverse order in the second, and so on, until no one can be assigned another course.

(ii) *HBS Observed*: the same snake draft, run with the rankings that students actually submitted for the HBS allocation. We use these reports as submitted, without an aftermarket.

(iii) *SRP Truthful*: the first student to pick is drawn uniformly at random, and each later picker is drawn uniformly from all students except the one who just picked. A student with no acceptable course left keeps her turns but takes nothing.

(iv) *RSD Truthful*: random serial dictatorship without quotas; each student takes all of her remaining acceptable courses before the next student picks.

(v) *Quota-RSD 12*: serial dictatorship with a cap of twelve courses per student, which is the number of seats that can be filled (5,419) divided by the number of students (456), rounded up.

(vi) *Quota-RSD 32*: serial dictatorship with the lowest cap which guarantees that no fillable seat is left unassigned.

Results. We compare the mechanisms on individual welfare, welfare inequality, and realized bundle sizes. Table 1 shows that using the observed reports lowers the share of students who receive their favorite course from 82.24% to 59.92%, and that SRP, under truthful reporting, recovers most of this loss, reaching 76.14%. The expected numbers of top-three and top-five courses show the same pattern. Making turns random therefore costs less than manipulation does.

SRP also outperforms the RSD benchmarks on the measures of preference satisfaction, especially on first-choice matches. Only 47.14% and 43.71% of students receive their favorite course under quota-RSD 12 and quota-RSD 16, compared with 40.04% under both unrestricted RSD and quota-RSD 32. Early choosers take popular courses in a block before anyone else chooses, which is the quality concentration that SRP prevents by never letting any student hold consecutive turns. Under quota-RSD 12, early choosers secure full 12-course bundles of high-ranked courses, while late choosers tend to receive worse ranks and fewer courses. Giving equal weight to each student's average assigned rank, Aggregate Mean Rank averages 10.099 under quota-RSD 12, well behind SRP's 9.535.

SRP distributes utility less evenly than the drafts (Gini 0.122 vs. 0.045–0.054), because drafts give every student one pick per round—ensuring both equal bundle sizes and balanced

Table 1: Mechanism Performance in the HBS Course Allocation Data

Mechanism	Aggregate Mean Rank	Favorite (%)	Top 3	Top 5	Gini	Mean Courses Received
HBS truthful	9.346	82.24	2.462	4.011	0.0453	11.884
HBS observed	9.987	59.92	1.806	3.098	0.0536	11.774
SRP truthful	9.535	76.14	2.358	3.832	0.1221	11.884
RSD truthful	13.818	40.04	1.353	2.355	0.4241	11.884
Quota-RSD 12	10.099	47.14	1.670	2.994	0.1479	10.949
Quota-RSD 32	13.818	40.04	1.353	2.355	0.4241	11.884

Notes: Simulation over 10,000 independent draws using the replication archive of Budish and Cantillon (2012a), with 456 students and 96 courses. All outcomes are evaluated against constructed truthful preferences. Lower values are better for Aggregate Mean Rank and the Gini coefficient;[a] higher values are better for the other measures. Aggregate Mean Rank is the average over students of the average rank of the courses each student receives, with rank 33 assigned to an empty bundle. Favorite (%) is the percentage of students receiving their first choice. Top 3 and Top 5 are the expected numbers of courses received among a student's top three and top five choices. Gini is the Gini coefficient of students' total utility, where a course of rank r gives utility $33-r$. Mean Courses Received is the average number of courses assigned per student. Quota-RSD 32 has a nonbinding cap and therefore matches unrestricted RSD on the same priority draws.

[a]For each draw, we calculate the population Gini as $G = \frac{\sum_{i,j}|U_i-U_j|}{2n\sum_i U_i}$, where $n = 456$ and $U_i = \sum_{c\in A_i}(33 - r_i(c))$ is student i's total utility, with ranks taken from the constructed truthful preferences. Students receiving no courses have zero utility and remain included. We report the average Gini across draws.

Table 2: Comparison of Sequential Allocation Mechanisms

Mechanism	Strategy-Proofness	Bounded Rank Concentration	Pairwise-Dominant Efficiency
Unrestricted RSD	Strong SD-SP	No	Yes
Quota-RSD	Strong SD-SP	No	Yes
Standard Drafts	Fails Weak SD-SP	Yes	Yes
SRP	Weak SD-SP	Yes	Yes

Notes: Bounded rank concentration requires that for every item an agent receives, she holds at most half of all items ranked higher by their respective recipients.

access across rounds—whereas under SRP the number and timing of turns are stochastic. Yet SRP distributes utility more evenly than quota-RSD and unrestricted RSD (0.148 and 0.424), where early priority blocks exhaust widely desired courses before others can pick. Moreover, quota-RSD assigns fewer courses per student (10.95, against 11.88 in all uncapped mechanisms) because students who reach the cap cannot take the remaining seats.

Moreover, quota-RSD with a quota of twelve leaves about 426 seats unassigned on average that could be filled without a quota, reducing the mean number of courses received from 11.88 to 10.95. In the profile of preferences and capacities in the data, preventing this for every priority ordering would require increasing the quota to thirty-two, at which point it never binds and the outcome coincides with unrestricted RSD. This illustrates another difficulty with quota-RSD when not every object is acceptable to every agent: choosing the quota is no longer straightforward. Dividing the number of seats by the number of agents ignores heterogeneity in acceptability, so sufficient aggregate quota capacity need not prevent unmatched seats. In this application, increasing the quota to eliminate those unmatched seats ultimately eliminates quota-RSD's inequality advantage over unrestricted RSD.

7 Concluding Remarks

Taking turns is one of the most common ways to divide indivisible items without money, and its familiar versions force a choice between protecting truthful reporting and limiting the early-mover advantage. Unrestricted RSD and quota-RSD protect incentives, but they allow early agents to accumulate highly ranked objects before others have a turn. Cyclic and snake drafts spread picking opportunities more evenly, yet their predictable turns create opportunities for manipulation. SRP changes only how the next picker is selected. It satisfies bounded rank concentration in every realization, and its uncertain future turns rule out misreports that strictly improve the SD assignment. Table 2 summarizes these sequential allocation mechanisms across their core theoretical properties.

A Proofs and Supplementary Derivations

A.1 Proof of Theorem 5.2

Proof. Fix an agent i, a true preference profile P, and an arbitrary misreport P_i'. Let $\pi_i = \phi_i(P_i, P_{-i})$ and $\pi_i' = \phi_i(P_i', P_{-i})$ denote agent i's allocation vectors under truthful reporting and under the misreport, respectively.

For each sequence σ, let $A_i(\sigma)$ and $A_i'(\sigma)$ denote the bundles assigned to agent i under P_i and P_i', respectively. Suppose, to the contrary, that $\pi_i' R_i^{\text{sd}} \pi_i$ but $\pi_i' \neq \pi_i$. Thus, $\pi_i'(\text{Up}_i(o)) \geq \pi_i(\text{Up}_i(o))$ for all $o \in O$. Because $\pi_i' \neq \pi_i$, allocations must diverge on some positive-probability sequence. On any such sequence σ, let $t^* = t^*(\sigma)$ denote the earliest round at which the object assigned to agent i differs across the two reports. Because a divergence requires at least two available objects, $t^* \leq m-1$. Let S denote the set of objects assigned to agent i at round t^* under P_i across all diverging sequences and let x^* be the most preferred object in S according to P_i.

Fix any $w \in O$ with wP_ix^*. On any sequence with a first divergence at round t^*, the object x allocated to agent i at round t^* under P_i satisfies x^*R_ix, so wP_ix. If w were available at round t^*, agent i would receive w under P_i, so w must have been allocated prior to round t^*. Since allocations prior to round t^* are identical under both reports, w is allocated to the same agent under P_i and P_i'. On sequences where allocations never diverge, w is likewise allocated identically. Hence $\pi_i'(w) = \pi_i(w)$ for all wP_ix^*.

Turning to x^*, we show that $\mathbf{1}\{x^* \in A_i'(\sigma)\} \leq \mathbf{1}\{x^* \in A_i(\sigma)\}$ on every sequence σ. On sequences where allocations never diverge, equality is immediate. On a diverging sequence, suppose $x^* \in A_i'(\sigma)$. Then either agent i was allocated x^* before round t^* under P_i', in which case she also receives it under P_i because all allocations prior to round t^* are identical, or x^* was not allocated before round t^*. In the latter case, x^* remains available at round t^* under P_i, so agent i receives an object weakly preferred to x^*, which by definition of x^* must be x^* itself. Thus $x^* \in A_i(\sigma)$ whenever $x^* \in A_i'(\sigma)$ for all $\sigma \in \Omega$, establishing $\pi_i'(x^*) \leq \pi_i(x^*)$.

Moreover, this inequality holds strictly on a set of sequences with positive probability. Since $x^* \in S$, there exists a positive-probability sequence σ with first divergence at round t^* where agent i receives x^* under P_i and some $z \neq x^*$ under P_i'. Since $n \geq 3$, choose two distinct agents $j, k \neq i$. A continuation that alternates between j and k from round t^*+1 onward is feasible and has conditional probability $(n-1)^{-(m-t^*)} > 0$. Agent i is excluded from all remaining turns on this continuation, so x^* is allocated to another agent under P_i'. On these sequences, $\mathbf{1}\{x^* \in A_i(\sigma)\} = 1 > 0 = \mathbf{1}\{x^* \in A_i'(\sigma)\}$, yielding $\pi_i'(x^*) < \pi_i(x^*)$.

Therefore,

$$\begin{aligned}\pi_i'(\mathrm{Up}_i(x^*)) &= \sum_{wP_ix^*} \pi_i'(w) + \pi_i'(x^*) \\ &< \sum_{wP_ix^*} \pi_i(w) + \pi_i(x^*) \\ &= \pi_i(\mathrm{Up}_i(x^*)),\end{aligned}$$

contradicting the assumption that π_i' weakly stochastically dominates π_i. □

A.2 Proof of Proposition 5.3

Proof. Fix any two distinct agents $i, j \in N$. Because the initial picker is chosen uniformly at random and the Markov transition probabilities are symmetric across all agents, the distribution over picking sequences is invariant under any permutation of agents: for any permutation ρ of N, we have $\Pr(\rho(\sigma)) = \Pr(\sigma)$. Under the permuted profile P^ρ and sequence $\rho(\sigma)$, induction over rounds shows that the same object is selected at each round, with its recipient relabeled by ρ. Hence the mechanism is anonymous.

If $P_i = P_j$, and ρ is the transposition swapping i and j, then $\rho(P) = P$. By anonymity, $\phi(P) = \phi(\rho(P)) = \rho(\phi(P))$, which implies $\pi_i = \pi_j$. This establishes equal treatment of equals. It remains to establish weak SD-envy-freeness.

If agents i and j have identical rankings of individual objects, swapping their identities in the picking sequence swaps their bundles, since the mechanism uses only these rankings. Thus $\pi_i = \pi_j$, so strict SD-envy is impossible.

Now suppose their rankings of individual objects differ. Let k^* be the first rank at which these rankings diverge; necessarily $k^* \leq m - 1$. Let S be their common top set of size $k^* - 1$, with $S = \emptyset$ if $k^* = 1$, and let x and y be their respective k^*-th ranked objects, so that xP_iy and yP_jx. To rule out $\pi_j P_i^{\mathrm{sd}} \pi_i$, it suffices to show $\pi_i(\mathrm{Up}_i(x)) > \pi_j(\mathrm{Up}_i(x))$.

For any $o \in S$, both agents rank S identically at the top of their preferences, so neither receives an object outside S while objects in S remain available. Let ρ_{ij} be the transposition that swaps i and j in the picking sequence, leaving all other agents unchanged. Before S is exhausted, every round allocates the identical object under σ and $\rho_{ij}(\sigma)$, because agents i and j share the same ranking on S, while any other agent faces the same available set and preferences. Consequently, for every $o \in S$, we have $o \in A_i(\sigma) \iff o \in A_j(\rho_{ij}(\sigma))$. Since ρ_{ij} preserves sequence probabilities, $\pi_i(o) = \pi_j(o)$ for all $o \in S$, and hence $\pi_i(S) = \pi_j(S)$. Because $\mathrm{Up}_i(x) = S \cup \{x\}$ is a disjoint union,

$$\pi_i(\mathrm{Up}_i(x)) - \pi_j(\mathrm{Up}_i(x)) = \pi_i(x) - \pi_j(x).$$

We first show that $\pi_i(x) \geq \pi_j(x)$. Take any sequence σ in which j receives x, and consider the transposed sequence $\rho_{ij}(\sigma)$. Before S is exhausted, the two agents receive identical

objects whenever assigned a turn. Under σ, agent i cannot hold any turn between the exhaustion of S and the round where j receives x, since i would then receive x. Hence, at the earliest round after S is exhausted at which j holds the turn, i holds the turn in the transposed sequence. Since j receives x at or after this round in σ, object x has not yet been allocated to anyone, so it remains available and i receives x in the transposed sequence. Thus, $\rho_{ij}\Big(\{x \in A_j\}\Big) \subseteq \{x \in A_i\}$, and since transposition preserves sequence probabilities, $\pi_i(x) \geq \pi_j(x)$.

To see that the inequality is strict, let E^* be the set of sequences whose first k^*+1 rounds strictly alternate between i and j, with round k^* belonging to i. This fixes a feasible prefix with probability $\Pr(E^*) = 1/[n(n-1)^{k^*}] > 0$.

On E^*, the first k^*-1 rounds exhaust S. At round k^*, agent i therefore receives x, her most-preferred remaining object. Under the transposed sequence $\rho_{ij}(\sigma)$, round k^* belongs to agent j. Because yP_jx and both x and y are available, agent j receives y. At round k^*+1, agent i holds the turn and receives x. Thus, j does not receive x under the transposed sequence. Hence $E^* \subseteq \{x \in A_i\} \setminus \rho_{ij}\Big(\{x \in A_j\}\Big)$. Because $\rho_{ij}\Big(\{x \in A_j\}\Big) \subseteq \{x \in A_i\}$ and $\Pr(E^*) > 0$, we obtain

$$\begin{aligned}\pi_i(x) &= \Pr(x \in A_i) \\ &> \Pr\Big(\rho_{ij}(\{x \in A_j\})\Big) \\ &= \pi_j(x).\end{aligned}$$

Therefore $\pi_i(\mathrm{Up}_i(x)) > \pi_j(\mathrm{Up}_i(x))$, ruling out $\pi_j P_i^{\mathrm{sd}} \pi_i$. By symmetry, $\pi_i P_j^{\mathrm{sd}} \pi_j$ cannot hold either, so no agent strictly SD-envies another, and thus the mechanism is weakly SD-envy-free. □

A.3 Proof of Proposition 5.4

Proof. Fix a realization of any of the five mechanisms, and let A be the resulting allocation. For unrestricted RSD, regard the first agent as selecting all objects in order of preference. Write $x_1, \ldots, x_m$ for the objects in their order of selection, and let i_t be the agent receiving x_t. In every case, each turn assigns the highest-ranked available object. The argument below uses only this rule.

Suppose, toward a contradiction, that some feasible allocation $B \neq A$ pairwise-dominates A. Let t be the first position at which x_t has a different recipient in B, and set $i = i_t$.

Since agent i selected her most-preferred available object at position t, every object she prefers to x_t was selected before t. By the choice of t, all these earlier objects have the same recipients in A and B.

Thus, using the upper-contour set $\mathrm{Up}_i(x_t)$, we have

$$|B_i \cap \mathrm{Up}_i(x_t)| = |A_i \cap \mathrm{Up}_i(x_t)| - 1,$$

because i receives x_t in A but not in B.

However, pairwise dominance requires an injection $f_i : A_i \to B_i$ that maps every object to a weakly preferred replacement. Such an injection must map the distinct objects in $A_i \cap \mathrm{Up}_i(x_t)$ to distinct objects in $B_i \cap \mathrm{Up}_i(x_t)$, contradicting the inequality above.

Therefore, no feasible allocation pairwise-dominates A. □

A.4 Proof of Proposition 5.6

Proof. Fix a realized picking sequence and its allocation A. Consider any agent i and object $x \in A_i$. Suppose that i receives x on her r-th turn, which occurs at position t in the picking sequence.

Since agents always select their highest-ranked available object, the objects in A_i ranked above x are exactly those selected by i before position t. Thus,

$$|\{y \in A_i : yP_ix\}| = r - 1.$$

Moreover, every object y selected before position t by another agent $j \neq i$ satisfies yP_jx, since x was still available when j selected y. There are $t - r$ such objects, so

$$\sum_{j \neq i} |\{y \in A_j : yP_jx\}| \geq t - r.$$

Consequently, bounded rank concentration holds whenever $t - r \geq r - 1$.

Under SRP, consecutive turns belong to different agents. Hence, between any two consecutive selections by i, at least one other agent selects an object. Therefore, $t - r \geq r - 1$.

Under both the cyclic and snake drafts, each complete round of n turns assigns exactly one turn to each agent. Before i makes her r-th selection, at least $r - 1$ complete rounds have elapsed, during which the other agents collectively made $(n - 1)(r - 1)$ selections. Thus,

$$t - r \geq (n - 1)(r - 1) \geq r - 1,$$

where the last inequality follows from $n \geq 2$. Therefore, all three mechanisms satisfy bounded rank concentration in every realization.

To establish the negative results, consider a profile in which all agents have the identical strict ranking

$$P_1 = \cdots = P_n : \quad o_1, o_2, \ldots, o_m.$$

Under unrestricted RSD, the first agent receives every object. Under quota-RSD with $k \geq 2$, the first agent receives the top k objects. In either case, letting i denote the first agent and $x = o_2$, we have

$$|\{y \in A_i : yP_ix\}| = 1, \qquad \sum_{j \neq i} |\{y \in A_j : yP_jx\}| = 0.$$

Hence bounded rank concentration is violated. □

References

Amanatidis, G., Birmpas, G., Christodoulou, G. and Markakis, E. (2017). Truthful allocation mechanisms without payments: Characterization and implications on fairness. In *Proceedings of the 2017 ACM Conference on Economics and Computation*, pp. 545–562.

Athanassoglou, S. and Sethuraman, J. (2011). House allocation with fractional endowments. *International Journal of Game Theory*, **40** (3), 481–513.

Aziz, H., Luo, P. and Rizkallah, C. (2017). Incompatibility of efficiency and strategyproofness in the random assignment setting with indifferences. *Economics Letters*, **160**, 46–49.

Balbuzanov, I. (2016). Convex strategyproofness with an application to the probabilistic serial mechanism. *Social Choice and Welfare*, **46** (3), 511–520.

Bogomolnaia, A. and Moulin, H. (2001). A new solution to the random assignment problem. *Journal of Economic theory*, **100** (2), 295–328.

Bouveret, S., Gilbert, H., Lang, J. and Méroué, G. (2025). Constrained serial dictatorships can be fair. In *Proceedings of the Thirty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-25)*, pp. 3762–3770.

— and Lang, J. (2014). Manipulating picking sequences. In *ECAI 2014: 21st European Conference on Artificial Intelligence*, IOS Press, *Frontiers in Artificial Intelligence and Applications*, vol. 263, pp. 141–146.

Brams, S. J. and Straffin, P. D. (1979). Prisoners' dilemma and professional sports drafts. *The American Mathematical Monthly*, **86** (2), 80–88.

Brandl, F., Brandt, F., Eberl, M. and Geist, C. (2018). Proving the incompatibility of efficiency and strategyproofness via SMT solving. *Journal of the ACM*, **65** (2), 6:1–6:28.

Budish, E. and Cantillon, E. (2012a). The multi-unit assignment problem: Theory and evidence from course allocation at harvard. *American Economic Review*, **102** (5), 2237–2271.

— and — (2012b). Replication data for: The multi-unit assignment problem: Theory and evidence from course allocation at harvard. American Economic Association [publisher], openICPSR [distributor], https://doi.org/10.3886/E112547V1.

—, Che, Y.-K., Kojima, F. and Milgrom, P. (2013). Designing random allocation mechanisms: Theory and applications. *American Economic Review*, **103** (2), 585–623.

Caragiannis, I., Gravin, N. and Huang, X. (2019). Envy-freeness up to any item with high nash welfare: The virtue of donating items. In *Proceedings of the 2019 ACM Conference on Economics and Computation*, pp. 527–545.

Caspari, G. (2026). Booster draft mechanisms for multi-object assignment. *International Journal of Game Theory*, **55** (1), 11.

Coreno, J. and Balbuzanov, I. (2026). Axiomatic characterizations of draft rules. *Journal of Economic Theory*, **234**, 106169.

Ehlers, L. and Klaus, B. (2003). Coalitional strategy-proof and resource-monotonic solutions for multiple assignment problems. *Social Choice and Welfare*, **21** (2), 265–280.

Hatfield, J. W. (2009). Strategy-proof, efficient, and nonbossy quota allocations. *Social Choice and Welfare*, **33** (3), 505–515.

Kasajima, Y. (2013). Probabilistic assignment of indivisible goods with single-peaked preferences. *Social Choice and Welfare*, **41** (1), 203–215.

Katta, A.-K. and Sethuraman, J. (2006). A solution to the random assignment problem on the full preference domain. *Journal of Economic theory*, **131** (1), 231–250.

Kesten, O., Kurino, M. and Ünver, M. U. (2017). On characterizations of the probabilistic serial mechanism involving incentive and invariance properties. *Mathematical Social Sciences*, **90**, 56–62.

Klaus, B. and Miyagawa, E. (2002). Strategy-proofness, solidarity, and consistency for multiple assignment problems. *International Journal of Game Theory*, **30** (3), 421–435.

Kojima, F. (2009). Random assignment of multiple indivisible objects. *Mathematical Social Sciences*, **57** (1), 134–142.

Kurokawa, D., Procaccia, A. D. and Wang, J. (2018). Fair enough: Guaranteeing approximate maximin shares. *Journal of the ACM (JACM)*, **65** (2), 1–27.

Mennle, T. and Seuken, S. (2021). Partial strategyproofness: Relaxing strategyproofness for the random assignment problem. *Journal of Economic Theory*, **191**, 105144.

Pathak, P. A. and Sönmez, T. (2008). Leveling the playing field: Sincere and sophisticated players in the Boston mechanism. *American Economic Review*, **98** (4), 1636–1652.

Svensson, L.-G. (1999). Strategy-proof allocation of indivisible goods. *Social Choice and Welfare*, **16** (4), 557–567.