Decentralized School Choice

September 7, 2025

1 Abstract

We study two mechanisms for matching K–12 students with schools: the well-known Top Trading Cycles (TTC) mechanism and a new market-based mechanism inspired by Leshno and Lo (2021), which we refer to as the competitive equilibrium (CE) mechanism. While TTC is strategy-proof and Pareto efficient, students may fail to recognize that truthful reporting is optimal, resulting in welfare losses. Leshno and Lo (2021) develop a method to calculate priority cutoffs in a student-school market, ensuring that demand equals supply at each school, and show that the resulting allocation in a finite model coincides with the TTC outcome. Our CE mechanism builds on this approach, with a key modification in how priority cutoffs are determined. We conjecture that the CE mechanism is cognitively simpler than the TTC mechanism, as students only need to select one school from a set of schools of which they are eligible to apply, rather than providing a ranking of all the schools. We conduct an online experiment to assess whether the CE mechanism yields a higher rate of truth-telling (defined as students choosing their best eligible school), yields more social surplus, and more closely implements the allocation that would result if subjects were to report their preferences truthfully in the TTC mechanism. So far, only pilot data have been collected.

2 Experimental Details

2.1 Intervention

We compare two mechanisms, TTC and a new market-based mechanism inspired by Leshno and Lo (2021). We refer to the latter as the competitive equilibrium (CE) mechanism. In the comparison, we will vary the market size, i.e., the number of students and seats. A "small market" has 12 students and seats, whereas a "large market" has 120 students and seats. Both types of markets have 3 schools. The large market is closer to real-world applications, such as when K-12 students are matched to schools. We will also vary whether there is uncertainty about the students' aggregate preferences over the three schools. With three schools, there are 6 possible rankings of schools. In the "no aggregate uncertainty" treatment, exactly one-sixth of the students have each preference profile. In the "aggregate uncertainty" treatment, each student's preference profile is a random realization of one of the six preference profiles, with each preference profile being equally likely. The no aggregate uncertainty treatment mimics the feature of Leshno and Lo (2021)'s continuum model that there is no aggregate uncertainty over preferences, whereas the aggregate uncertainty treatment better matches the (finite) real world where students' aggregate preferences over schools are not known. In CE treatments, students' priorities at each school are random and independent, and hence characterized by aggregate uncertainty.

In sum, we have eight between-subject treatments. Table 1 summarizes the treatments. In every treatment, subjects read the instructions and completed comprehension check questions along the way. They then made school choice decisions over two rounds. In each round they had the same preferences, but their (random) priorities differed. The experiment concluded with a demographic survey and an open-ended question asking them to explain their decisions.

2.2 Primary outcomes

The main outcomes of interest are: (1) the probability at which subjects truthfully report their preferences under TTC, and the rate at which subjects choose their best eligible school in CE. (2) The payoffs realized under TTC and in CE. (3) The closeness of the TTC allocation under truth-telling to (a) the TTC allocation for reported preferences and to (b) the CE allocation. Here, we briefly explain these outcome measures. Thorough definitions and the analysis plan are explained in section 2.3 and section **??**.

The existing literature has shown that subjects largely do not report their preferences truthfully in the TTC mechanism, even though it is a weakly dominant strategy to do so (for example, Guillen and Hing, 2014; Guillen and Hakimov, 2017, 2018; Chen and Sönmez, 2006; Guillen and Veszteg, 2021). A likely explanation is that subjects do not recognize that truth-telling is optimal. This motivates our study of the CE mechanism inspired by Leshno and Lo (2021), in which a student's individual priorities, together with the CE priority prices, determine the set of schools to which the student is eligible to apply. While the CE priority prices clear markets in a continuum economy, they will typically fail to do so in a finite economy, and hence in our implementation of the market, students may be rationed. When choosing among eligible schools, a student is informed of the odds that they will be admitted to each school. In the CE treatment, we are interested in the rate at which students choose their best eligible school. Since the odds of admission are higher in the large market than in the small market, we expect students to be more likely to choose their best eligible school in the large market. For the TTC mechanism, we expect no difference in the rate at which subjects truthfully report their preferences in small and large markets.

Then, the natural next question would be how the payoffs would change between the two mechanisms. We calculate the average payoffs of each market and compare which mechanism subjects earn more. Since we expect there will be less difference in behaviors depending on the market size in the TTC treatments, the difference in average payoffs is

		Mechanisms	
		TTC	CE
Market size	Small	2 students/preference (6 preferences)	2 students/preference (6 preferences)
	Large	20 students/preference (6 preferences)	20 students/preference (6 preferences)

(a) No aggregate uncertainty

		Mechanisms	
		TTC	CE
Market size	Small	12 *i.i.d.* draw from unif dist	12 *i.i.d.* draw from unif dist
	Large	120 *i.i.d.* draw from unif dist	120 *i.i.d.* draw from unif dist

(b) Aggregate uncertainty

Table 1: Treatments

also expected only in the CE treatments. Especially, we expect to have an improvement in payoffs when the market size increase in the CE treatments.

Last, we compare which market produces the closer outcomes to the TTC equilibrium predictions, where everyone reports their preferences truthfully. For this purpose, we use the Borda distance metric proposed by Can et al. (2023).

2.3 Experimental Design

2.3.1 General Environment

There are three schools: A, B, and C. The number of seats in school A, B, C is $2k$, $4k$, and $6k$, respectively, where $k = 1$ and $k = 10$ depending on the size of the economy. Each student has a priority at each school. Let $r_i(s)$ denote the priority of student i at school s. We assume that $r_i(s) \sim U[0, 1]$ for each student i and school s. Each student's priority at each school is drawn from the same distribution to capture the idea that students are ex-ante symmetric from the perspective of each school.

For students, there are six possible preference orderings over the three schools, e.g., $A \succ B \succ C$, $A \succ C \succ B$, and so on. We consider an economy where the preference orderings are uniform. That each preference ordering is equally common captures the idea that schools are of the same quality. Schools differ only in the number of seats.

There are two possible ways to implement the uniform preference orderings in an experiment: (1) exogenously assign the same number of students to each preference ordering, and (2) for each student, independently draw their preference from the uniform distribution over the six preference orderings. We will do both cases. The first one is the *no aggregate uncertainty* treatment, and the second one is the *aggregate uncertainty* treatment.

The small market consists of 12 students. Thus, in the *small market* $\times$ *no aggregate uncertainty* treatments, there are two students for each of six preference orderings, while *small market* $\times$ *aggregate uncertainty* treatments require drawing 12 students' preferences independently with an equal chance for each preference ordering. In the case of the large market, there are 120 students. Similarly, there are 20 students for each preference ordering for *no aggregate* *uncertainty* treatments and we draw preferences 120 times independently for *aggregate uncertainty* treatments. Priority and preference profiles are the same between TTC and CE treatments.

For no aggregate uncertainty, we consider 10 large markets (LM1, LM2,..., LM10) and 10 small markets (SM1, SM2,...,SM10). The subjects' priorities in the large market are generated randomly.

To construct SM1, for each preference profile, take the first two students in LM1 with that preference profile, and then add them to SM1. This makes a market with 12 students. Repeat for SM2 and LM2, and so on.

For aggregate uncertainty, we again consider 10 large markets (LM1, LM2, ..., LM10) and 100 small markets (SM1, SM2, ..., SM10) The subjects' preferences and priorities in an uncertain large market (ULM) are generated randomly.

To construct USM1, take the first 12 students from ULM1 (regardless of preference profile). Repeat for USM2 and ULM2, and so on.

2.3.2 TTC treatment

In the TTC treatment, subjects are informed of (1) the number of seats in each school, (2) the number of students (size of the economy), (3) their earnings from each school (their own preferences), (4) their priority at each school, and (5) distribution of preferences (whether there's no aggregate uncertainty or not). Given the information, subjects submit a full preference ranking. Figure 1 is an example decision page of the TTC treatment.

They make choices twice (Round 1 and Round 2). All other things remain the same between the two rounds, except for priorities. Allocations follow the traditional TTC mechanism.

2.3.3 CE treatment

The key insight we take from Leshno and Lo (2021) is that the TTC allocation can be replicated using priority prices. In other words, each student is assigned a "budget set" of schools they can apply to based on their priorities. The TTC allocation can then be interpreted as assigning each student to their best school in their budget set. Because priorities determine the budget set, they also play a role in CE prices (or cutoffs) of schools in a decentralized economy. These cutoffs—and the resulting budget sets—can be computed directly from the distribution of preferences and priorities in a continuum economy.

We calculate the CE prices under the continuum economy where Schools A, B, and C have a measure $\frac{2}{12}$, $\frac{4}{12}$, and $\frac{6}{12}$, of seats respectively, and there is a measure $\frac{2}{12}$ of students with each preference profile. Then, we use these prices to form subjects' budget sets.

Subjects are informed of (1) the number of seats in each school, (2) the number of students (size of the economy), (3) their earnings from each school (their own preference), (4) their own budget set, (5) odds of getting into a school they choose conditioning on all others choose their highest-earning school from their budget set, and (6) distribution of preferences (whether there's no aggregate uncertainty or not). Given the information, subjects choose a school from

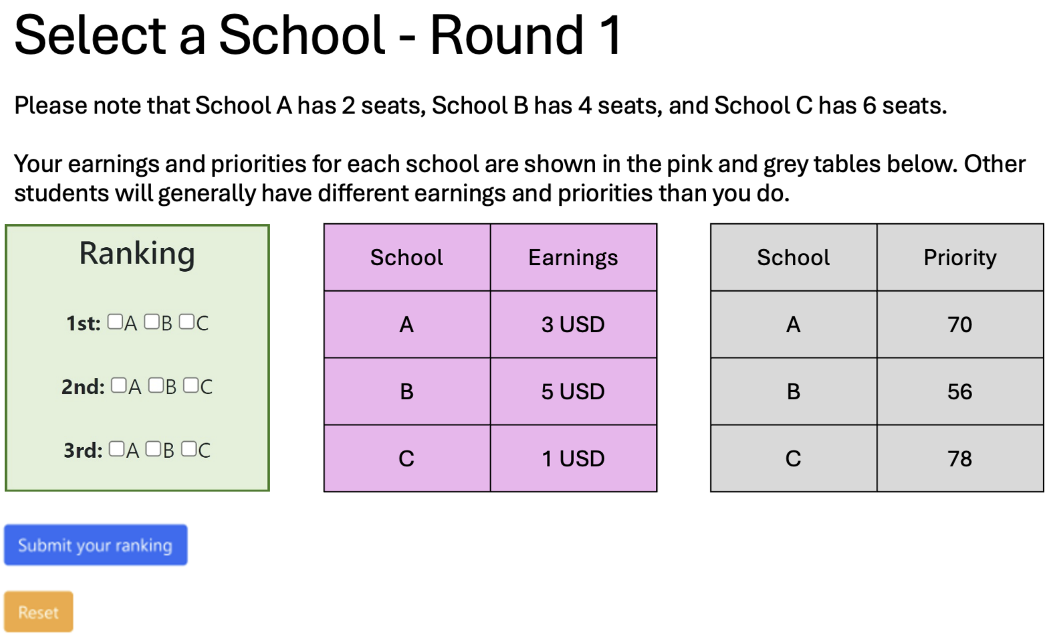

Figure 1: TTC example decision page

their budget sets. Figure 2 is an example decision page of the CE treatment. The gray panel is a budget set table. To avoid economic jargon and use a more comprehensible term, we say whether they are eligible for a school or not, instead of a budget set. Also, since in this example the subject is not eligible for School A, they are grayed out in the green panel.

Subjects make choices twice (Round 1 and Round 2). All other things remain the same between the two rounds, except for priorities. If a school has an excess demand, we randomly select which students will be assigned to it. Students who are not selected will be assigned to a school with available seats.

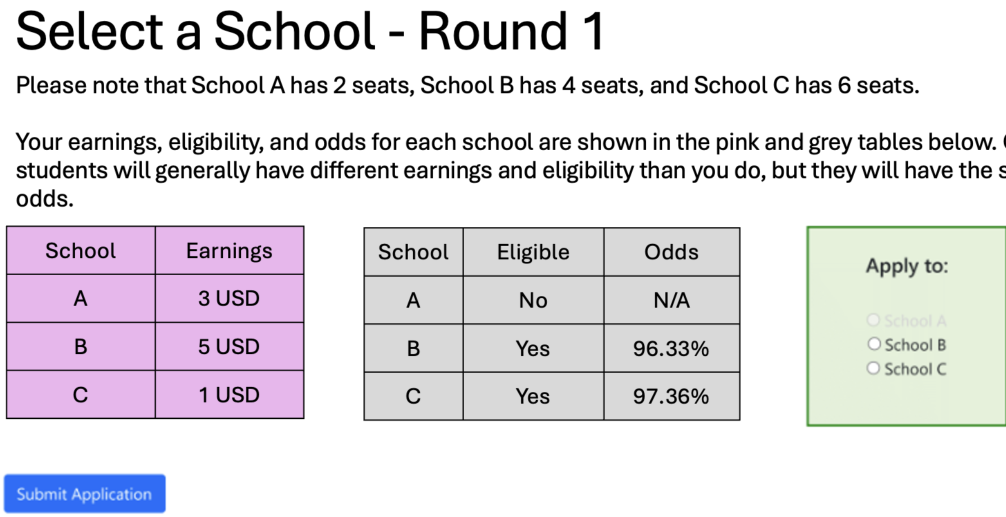

Figure 2: CE example decision page

References

Can, B., M. Pourpouneh, and T. Storcken (2023): "Distance on matchings: An axiomatic approach," *Theoretical Economics*, 18, 597–632.

Chen, Y. and T. Sönmez (2006): "School choice: an experimental study," *Journal of Economic theory*, 127, 202–231.

Guillen, P. and R. Hakimov (2017): "Not quite the best response: truth-telling, strategy-proof matching, and the manipulation of others," *Experimental Economics*, 20, 670–686.

——— (2018): "The effectiveness of top-down advice in strategy-proof mechanisms: A field experiment," *European Economic Review*, 101, 505–511.

Guillen, P. and A. Hing (2014): "Lying through their teeth: Third party advice and truth telling in a strategy proof mechanism," *European Economic Review*, 70, 178–185.

Guillen, P. and R. F. Veszteg (2021): "Strategy-proofness in experimental matching markets," *Experimental Economics*, 24, 650–668.

Leshno, J. D. and I. Lo (2021): "The cutoff structure of top trading cycles in school choice," *The Review of Economic Studies*, 88, 1582–1623.