Strategies in the Tallinn School Choice Mechanism

Andre Veski
Tallinn University of Technology, Estonia
Address: Akadeemia tee 3, 12618 Tallinn, Estonia
E-mail: andre.veski@ttu.ee

Kaire Põder
Estonian Business School, Estonia
Address: Lauteri 3, 10114 Tallinn, Estonia
E-mail: kaire.poder@ebs.ee

Abstract

In the first 20 years of the market economy in Estonia, the public school market was decentralised in Tallinn. Recently, a hybrid market was established by centralising the school allocations to comprehensive schools and also allowing some selective schools to autonomously select students for some groups. We contribute to mechanism design literature by studying the centralised clearing-house used in Tallinn – the Tallinn mechanism. By using genetic algorithms, we show that, the Tallinn mechanism incentivises families to manipulate their preference revelation by reporting only a few schools and not always from the top of their preference list. Also we see that the expected utility in the Tallinn mechanism is higher compared to the widely used Deferred-Acceptance mechanism, although the number of unassigned students is also higher.

JEL codes: D02, D04, D47, D82
Keywords: market design, school choice, preference revelation, stability

1. Introduction

In recent years, economists have gained significant experience (and fame) in practical market design. The applications of the theoretical principles of market design demonstrate that institutions matter at a level of details that economists have not often had to deal with. Much of this research is based on the seminal papers by Gale and Shapley (1962) that were initially used for entry-level job markets, such as the National-Resident Matching Program and others (Roth, 2008). The core of the curriculum is to apply a central mechanism that collects information from a market participant as ordered preferences and finds the allocation that has some merits. This approach has recently been very fruitful in many real-life resource allocation problems in public good provision (Milgrom, 2000; McAfee, Preston and Mcmillan, 1996) and also high school (college) allocation problems (Abdulkadiroğlu et al., 2005a, b; Abdulkadiroğlu and Sönmez, 2003; Balinski and Sönmez, 1999; Romero-Medina, 1998). Moreover, significant research has recently been carried out to explore the allocation of school seats to students in primary (e.g. Abdulkadiroğlu et al., 2006, 2011; Dur et al., 2013), secondary (Dur et al., 2013) as well as upper-secondary schools (Abdulkadiroğlu et al., 2015, 2009). In this agenda, two-sided matching markets are used as in the "marriage problem" to solve "the college admission problem" and some striking results concerning agent incentive schemes have recently been obtained (Abdulkadiroğlu et al., 2011; Abdulkadiroğlu and Sönmez, 2003; Pathak and Sönmez, 2013, 2008; Erdil and Kumano, 2013; Erdil and Ergin, 2008). In the case of two-sided markets it has been shown (Roth, 1982) that only a limited number of stable matching procedures exists to form a dominant strategy for families to reveal their true ordered preferences.

The existing matching mechanism literature is growing, not only in terms of new cases and designs, but also by adding new problematic design areas; that is, encouraging diversity with the use of quotas or priority classes that in many cases can fail to enforce social justice (Dur et al., 2013; Kominers and Sönmez, 2013; Fragiadakis and Troyan, 2013; Erdil and Kumano, 2013). However, to the best of our knowledge, there is no literature dealing with post-communist school allocation mechanisms. Our experience indicates that in the Soviet era, mechanisms were widely in use in many spheres; for example, the allocation of university graduates or university choice. One common characteristic of the communist mechanisms was the school-proposing nature while the submitted preferences were marginally considered. The latter has not diminished its prevalence – many applications in two sided markets are still initiated by the "stronger side" and have no welfare considerations.

We are contributing to the matching research agenda by studying the Tallinn school choice mechanism (Tallinn mechanism hereinafter). Notably, Soviet-style central matching was abandoned in the Tallinn school market during the liberal reforms after the 90s and substituted by decentralised or semi-centralised designs. Over the last few years, central matching has been reintroduced in the Tallinn school market for allocating children to primary schools. Through trial and error, local policy-designers established the Tallinn mechanism as a central marketplace in 2012. This mechanism has specific characteristics in addition to the school proposing nature. First, students are prioritised according to distance from the school. Second, families can submit three unordered preferences. Third, the mechanism uses immediate acceptance (Boston).

As with the shortcomings of the Boston mechanism, which has created a rule of thumb for submitting the preferences strategically (Ergin and Sönmez, 2006; Pathak and Sönmez, 2008; Pathak and Shi, 2013; Abdulkadiroğlu et al., 2011), we show there are similar rules of thumb for manipulation under the Tallinn mechanism. In the Boston mechanism there are different levels of sophistication among families who participate in the mechanism; that is, one strategy was to avoid ranking two over-demanded schools as their top choices or an unsubscribed school or popular school was recommended to be put as the first choice plus a "safe" second choice. Hence, as Pathak and Sönmez (2008) showed that the Boston mechanism is a coordination game among sophisticated families. Thereby "levelling the playing field" by diminishing the harm done to families, who do not strategise or do not strategise well, is emphasised as a condition for designing the new mechanism. Similarly, we introduce the Tallinn mechanism as a sophisticated game. We ask how many preferences it is rational to report under such a mechanism and whether families reveal their true top preferences or manipulate in both dimensions – report only a limited number of preferences which might not be at the top of their preference lists. In addition, we ask whether this behaviour is dependent on their preference structure – the functional form of their estimated cardinal utility function. The latter allows us to show whether the strategy of revealing preferences is dependent on the relative cardinal measure of utility from first, second etc. preference – which can be considered a measure of marginal utility. Moreover, we are interested in social inefficiencies defined as the difference between individual allocated ranks and unassigned families under the Tallinn mechanism compared to the Deferred-Acceptance mechanism.

Our research design is based on computational experiments. For descriptive analysis, we use data from the centralised database, e-school. The e-school database is an electronic register, where approximately 4000 7-year-old children with a known home address annually list their school preferences. The rest of our data is synthetic. Our research strategy is the following. After descriptive stylised facts, we use genetic algorithms to find the best strategies for revealing the preferences of families. We illustrate the results by indicating the extreme cases of utility function – utility over preferences might be uniform or extreme; in other words, might all be concentrated on the first preference. We use genetic algorithms to find good strategies for maximising utility for the families. Family agents optimise strategies by observing their allocation and the obtained utility.

We continue as follows. First, we describe the broader Tallinn school market, then the concrete mechanism used by the Tallinn education administration – the Tallinn mechanism. In section 3, we describe the preference generation, the utility function and genetic algorithms. In section 4, we describe the results of the parental strategies and the obtained allocation after revealing what and how much to report to the central marketplace. Finally, we conclude by highlighting the policy implications for Estonia and for other decentralised and centralised markets.

2. Background: Tallinn School Market

Over the years, some schools in Tallinn have become over-subscribed. These selective schools have inter-district admissions to primary school and have all introduced aptitude entrance tests (hereinafter exam schools). For intra-district comprehensive schools (hereinafter regular schools), the tradition has been a central or semi-central catchment-based allocation based on an application (single preference or multiple preferences) from the family. Rejected offers were not treated centrally – each school and student should find the match independently.

The admission process for the exam schools takes place between January and March. We note that it has been shifting from March (in 2012) to February (in 2013) and even to January (in 2014). The second stage (in the Tallinn mechanism) in regular schools starts on 1st of March with the submission of an electronic application to the e-school register. Central but manual entries are made by 25 May. By 10 June, parents must either accept or decline offers. There is a later decentralised round of applications for additional vacant positions after 15 June.

To make the entire school choice procedure more transparent, we highlight the following steps:

1. Students are assigned to exam schools based on the proposing Deferred-Acceptance (DA) mechanism of decentralised schools
2. The remaining students are centrally assigned to regular schools based on the *Tallinn mechanism*
3. Unassigned students are assigned to the closest schools potentially rejecting an already assigned student. Some students might be assigned to a school they did not apply to. This continues until all students are assigned.
4. Students can reject their assigned position. Once the rejection/acceptance deadline has passed, schools can autonomously accept students for any available positions.

Therefore, the hybrid structure of the Tallinn school market consists of exam schools (decentralised matching), the Tallinn mechanism (central matching) and the final decentralised round. We are only interested in the Tallinn mechanism.

2.1. Tallinn Mechanism

The Tallinn mechanism governs only the central admission procedure to all municipal primary schools. These schools rely on the following procedural steps. First, families submit an application where they list up to three schools. Then the seats are allocated based on the following procedure:

0. Look at the schools in a random order. Each student is only considered for the school to which the family applied.
1. Allocate students to the first school for which they have high (siblings and distance- based) priority until the quota is full.
2. Allocate students that were not allocated before to the second school for which they have high priority until the quota is full.

...

k. Allocate students that were not allocated before to the k-th school for which they have high priority until the quota is full.

It is important to stress that regular school applications are limited to three options; in other words, the parent has the right to list three schools, but these are not considered in any particular order. The application can also contain information about siblings and the school(s) they attend. Centralised school priorities are considered based on the student's distance from the school (in metres) from the officially registered address.

Table 1: Number of Reported Preferences under the Tallinn Mechanism

# of prefs	2011	2012	2013
1	52 %	74 %	76 %
2	18 %	15 %	14 %
3	11 %	11 %	9 %
> 3	18 %	0 %	0 %
Mean	2.2	1.4	1.3

Source: Authors' calculation

We use descriptive statistics to illustrate the micro-mechanism in Tallinn over three consecutive years – from 2011 to 2013. In 2011, the market was decentralised. However, applications were centrally collected without any upper limit on the submitted preferences. The Tallinn mechanism has been in use since 2012, limiting the amount of unordered preferences submitted to three (see Table 1).

This stylised fact illustrates the tendency to report a limited list. Most families submit only a single preference. However, there is no clear indication that parents do not manipulate as in the Boston mechanism and decide to reveal strategically lower preferences or "safe" choices. Therefore, we are interested in whether it is rational to report less than three preferences and what the rationality is of reporting truthful preferences.

2.2. Example of Deciding What to Report

We illustrate the choice set for parents using a simple extensive form game (Figure 1). In such a game, the parents in the starting node have three strategies – to report either 1, 2 or 3 preferences. In the following subgames, the designer randomly allocates the student to the reported school or an outside option. In the final nodes, the utilities are reported by indicating the preference – 1 stands for first preference and Ø indicates the utility of the outside option. In the illustration below, we assume risk neutral agents.

Figure 1: Extensive Form Reporting Game

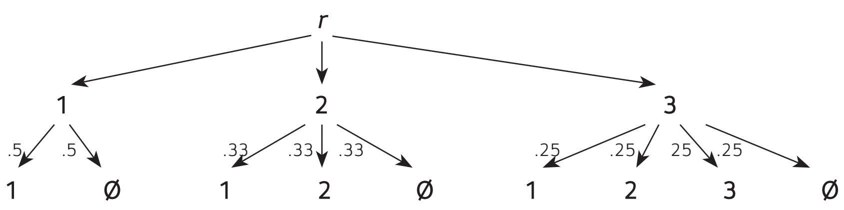

Source: Authors' illustration

Assume that we have two utility functions, where k indicates a position in a preference list:

- $u_1(k) = 0.358 - 0.025(k - 1)$
- $u_2(k) = 0.658 - 0.325(k - 1)$

Then we obtain cardinal utilities for $k \in \{1, 2, 3\}$ as in Table 2.

Table 2: Utilities

k	$u_1(k)$	$u_2(k)$
1	0.358	0.658
2	0.333	0.333
3	0.309	0.009

Source: Authors' calculation

Assuming the uniform probabilities of being unassigned or assigned to one of their preferences, as in Figure 1, we can compute the expected utilities for both utility functions and all cases of reported preferences. Notably, we do not take into account the demand for a school or the overall availability of places. Moreover, it is preferable to always report schools higher in the preference list, so we do not investigate cases where, for instance, only the second or third choice is reported, because the expected utility will definitely be lower. This might not be the case when the probabilities of being assigned to a particular school are not uniform.

We see that the probability of being left unassigned decreases as more preferences are reported, but so does the probability of getting a place in the most preferred school. The expected utilities for $u_1(k)$ are:

- for reporting one school $E_1[u_1(k)] = \frac{1}{2}(0.358 + 0) = .179$
- for reporting the first two schools $E_2[u_1(k)] = \frac{1}{3}(0.358 + 0.333) = .230$
- for reporting the first three schools $\mathrm{E}_3[u_1(k)] = \frac{1}{4}(0.358 + 0.333 + 0.309) = .250$

We see that reporting all three preferences maximised utility. With utility function $\mathrm{u}_2(\mathrm{k})$ the expected utilities are:

- for reporting one school $E_1[u_1(k)] = \frac{1}{2}(0.658 + 0) = .329$
- for reporting the first two schools $E_2[u_1(k)] = \frac{1}{3}(0.658 + 0.333) = .330$
- for reporting the first three schools $\mathrm{E}_3[u_1(k)] = \frac{1}{4}(0.658 + 0.333 + 0.009) = .250$

As Figure 1 illustrates the game, where under the expected utility maximisation assumptions, parents obtain higher utility by reporting only one or two schools with $u2(k)$.

We are interested in finding near-optimal strategies in large markets, where agents might have similar preferences or there are popular and over demanded schools. Additionally, the revealed demand also depends on the strategies of the agents and the revelation strategies depend on the revealed demand.

2.3. Deferred-Acceptance Mechanism

A widely used mechanism for school choice is the Deferred-Acceptance (DA) mechanism (Abdulkadiroğlu and Sönmez, 2003; Pathak and Sönmez, 2013). First introduced by Gale and Shapley (1962) it has been confirmed to be useful in many applications in matching residents to hospitals (Mullin and Stalnaker, 1952; Roth, 1984) and other labour market applications (e.g. Roth, 2008), students to schools (Abdulkadiroğlu et al., 2005a; Pathak and

Sönmez, 2013) and colleges (Abdulkadiroğlu et al., 2005a; Pathak and Sönmez, 2013) and probably more.

The DA mechanism has become so popular mainly due to two good properties it has: strategy-proofness and no justified envy (e.g. Abdulkadiroğlu and Sönmez, 2003). Strategy-proof mechanisms always make it safe and in the families' best interests to report their true preferences. In an allocation with no justified envy, families never have an option to a more preferred school, because another family with a higher priority has already been assigned to that school. When there is no justified envy, the allocation is also called stable. While there can be multiple stable matchings, we usually aim to obtain student-optimal stable-matching (e.g. Abdulkadiroğlu and Sönmez, 2003), as that is the only way to guarantee strategy-proofness. While there are other strategy-proof mechanisms, such as the Top Trading Cycles, this does not ensure that the final allocation is stable (e.g. Abdulkadiroğlu and Sönmez, 2003).

In general the Deferred-Acceptance works as follows:

1. All students are tentatively assigned to their first preference. If schools have more students assigned than places, they reject some students with lower priority.
2. All rejected students are tentatively assigned to their second preference. Again if schools have more students assigned than places, they reject some students with lower priority. Note that the school may reject students tentatively assigned in the previous round, if they have a lower priority than some new applicants.

...

k. In general, rejected students are tentatively assigned to their next preference. If a school has more students assigned than places then students with lower priority are rejected. The process continues until all students are assigned a place, or all preferences have been explored.

Due to its good properties, the DA is a good "ideal" model for our comparative welfare analysis. It allows us to show how final allocations differ under the Tallinn and DA mechanism.

3. Model

3.1. Environment

We are interested in understanding strategies in multiple environments. We characterise the environment with societal parameters (Tables 3 and 4) and the parameters of an individual. Societal parameters describe the number of schools, the number of exam (popular) schools, the correlation between ordered preferences, and so on. Exam schools exist because they are popular overall, so we consider them as a metaphor for globally popular schools. Moreover, in Tallinn, these schools are still allocated the most groups through the Tallinn mechanism.

We fix the number of schools, the number of places in a school and the number of students for all our experiments (Table 3). In addition, the maximum number of ordered preferences for each agent is fixed. We model families as agents. They are willing to apply to or can rank up to 15 schools at the most, although the utility from lower preferences is relatively small. This is partly driven by case specificities, as 15 was the maximum number of schools listed in the decentralised market in Tallinn in 2011. From those 15 ordered preferences, agents have to select three to report in the Tallinn mechanism.

We investigate societies, where agents can have random or spatially correlated preferences – the latter indicates that schools nearby are more desired (Table 4). We also look at the effect of having the same set of popular schools – exam schools. In these societies, all agents would prefer exam schools, even if they are further away than the nearest regular schools. In the case of spatial preferences among exam schools, agents would still prefer schools nearby, and no other criteria matters. In each computational experiment, all these parameters are fixed. The priorities for schools are always spatial, distance based. Agents closer to a school have a higher priority in that school.

Table 3: Fixed Societal Parameters

Parameter	Description
k = 15	Length of preference lists
n = 3000	Number of family agents
m = 50	Number of schools
qj = 60	Number of places in school j

Table 4: Variable Societal Parameters

Parameter	Description
$c \in 0{,}1$	Spatial correlation in preferences
$m_e \in \{0, 10\}$	Number of exam schools

For each agent looking for a place at the school, we only have one parameter: the functional form of the utility function described by the parameter (α). The latter indicates the slope of the utility function. In each experiment, our agents are heterogeneous, so they have different values for the slope of the utility function.

3.2. Preferences

We assume that agents have strict preferences for schools. In the simplest case, preferences are random; in other words, each agent has a totally idiosyncratic preference ordering. In general, we can think of more structured preferences in a society, parametrised by the length of the preference list (k) and the correlation between the preference lists (c). In our experiments, the preference lists are limited to k = 15. Correlated preferences stem from a spatial preference ordering, and can also be considered 2D-Euclidean preferences (Bogomolnaia and Laslier, 2007). The degree of correlation is also the same for all agents, but the preference ordering is not necessarily identical when comparing two agents due to the spatial nature of preferences.

We generate the preferences using the Algorithm 1 with parameters k, c and m. This algorithm is a modified version of a random permutation algorithm (Knuth, 1997, p. 145) to generate correlated preferences with parameter c. The algorithm starts with a master list of n numbers (agents). Then it iterates the list from beginning to end, each time at position j randomly selecting a position $q \in [j + 1, n]$ to exchange values with. The correlation parameter c illustrates how biased the randomly selected position is; higher values indicate that the exchange position is selected closer to the current position j. With $c = 0.0$ the selection is uniformly probable over all positions, until finally at $c = 1$ the exchange position is always the active position and all the generated lists are exactly the same. There is one global ordering of agents for each side of the market that is used for generating correlated preferences.

3.3. Utility Function

While agents have a preference ordering for schools, their behaviour might also be influenced by the cardinal utility they gain from assignment to the particular preference. A similar notion was illustrated by the reporting game in Section 2.2. In order to understand the behaviour with different cardinal valuations, we use exponentially declining utility over alternatives. When compared to consecutive schools i and $i + 1$, we assume that $\frac{u(i+1)}{u(i)} = 1 - \alpha$. Furthermore, we need to normalise the utility function such that $\Sigma_{i=1}^{k} u(i) = 1$. The resulting form of the utility function is in (1), where $i \in \{1, ..., k\}$ is a position in the preference ordering.

Algorithm 1 Correlated permutation

Require: $m,k = 15$, $c \in [0,1]$
Ensure: p is a permutation of unique numbers
$p \leftarrow 1, 2, 3, ..., n$, $j \leftarrow 0$
while $j < k$ **do**
  $r \leftarrow [0.0, 1.0]$ uniform random number between 0and 1
  $q \leftarrow [m - (m - j) \bullet \mathrm{r}^{1-c}]$
  $t \leftarrow p_q, p_q \leftarrow p_j, p_j \leftarrow t$
  $j \leftarrow j + 1$
end while
return $\{p_1, p_2, ..., p_k\}$

$$u(i) = \frac{\alpha(1-\alpha)^{i-1}}{1 - (1-\alpha)k} \tag{1}$$

When $\alpha \to 0$, then cardinal utilities for all alternatives are exactly the same $u(i) = \frac{1}{k} \forall i$. When $\alpha = 1$, then all utility is concentrated in the first preference, that is $u(1) = 1$. In Figure 2, we show the utility values using some examples of α.

Figure 2: Exponential Utility Function

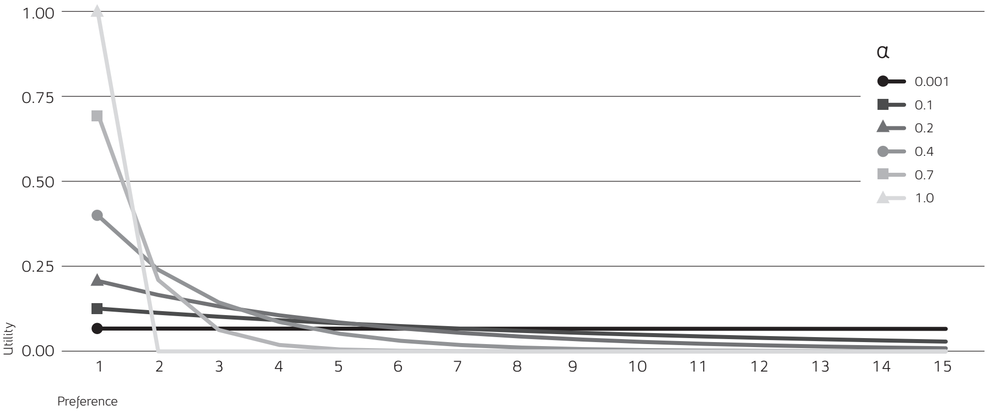

Source: Authors' illustration

The utility function can be compared to a linear utility function over perfect substitutes $u(x_1, ..., x_n) = \beta_1 x_1 + ... + \beta_n x_n$, where the consumer is allocated at most one good x_i. The β_i is the value of the allocated good x_i to the consumer (e.g. Varian, 2006, p. 61). Our utility function (1) states the shape of the decline in value β_i of the goods to consumers. We assume that agents are risk-neutral, i.e. they maximise their expected utility $E[u]$.

Using utility ratios $\frac{u(i+1)}{u(i)}$ to measure on preferences is also popular in decision theory, Saaty scale, and is supported by some psychological observations (e.g. Franek and Kresta, 2014, and references therein). Another reason is that differences are greater in geometrically declining function than linearly. So the effects we are investigating are more evident.

3.4. Genetic Algorithm Optimisation

We use genetic algorithms to find a near-optimal strategy for reporting in the Tallinn mechanism. The genetic algorithms adapt existing strategies to find better ones that would result in an increased utility. The result of a genetic algorithm after optimising is a steady state (e.g. Riechmann, 2001). While a steady state is also by definition a Nash equilibrium in a game, it could simply be one among many in multiple equilibria games. Our experiments are carried out with agents, not populations, as each individual might find a better strategy in every iteration, but for populations, a certain strategy would remain roughly constant.

There has been extensive use of genetic algorithms and programming in finance (e.g. Chen, 2002; Chen et al., 2011; Chen and Tai, 2010) and economics in general (e.g. Riechmann, 2001). Agents learn better trading strategies by observing the market. The main difference compared to our model is that agents do not have much to observe about the school market. Players do now know either the overall demand for schools or the preferences of other agents in the market. The only information source is their own allocation and the utility they gain from the market. With genetic algorithms, our approach is to find strategies that would maximise the utility of the agents.

Here we do not assume that the manner of genetic algorithms is in reality how humans learn. We only employ it for computational tractability, as exploring the entire strategy-space for 3,000 agents is resource consuming. However, there are studies that use a form of genetic algorithm as a model for learning (see e.g. Unver, 2001; Roth, 2002; Unver, 2005) and is also observed as exhibiting features with human subjects (e.g. Arifovic, 1994, 1996; Duffy, 2006).

Algorithm 2 Simple Genetic Algorithm – single iteration

Require: A set of agents, u agents utilities
Ensure: A is a set of agents
  $n \longleftarrow |A|$
  $s \longleftarrow \sum_{\alpha \in A}{}^{U\alpha}$
  $p \longleftarrow \{ \frac{U\alpha}{s}, \forall \alpha \in A\}$ {selection probablities}
  $i \longleftarrow 0$
  for all $r_1, r_2 \in$ Select (A, p, n) **do**
    {select with probability p, with replacement n pais of strategies}
    α_i $CrossOver(r_1, r_2)$ {assign new strategy to agent α_i
    if $RandomNumber$ () < 0.05 **then**
      $Mutate(\alpha_i)$
    end if
    $i \longleftarrow i + 1$
  end for
  return A

Genetic algorithms have two basic operations for finding an improved strategy (e.g. Simon, 2013): mutation and crossover. Mutation slightly tweaks an existing strategy and cross-over merges two successful strategies to find a better one. Finally, selection indicates an operation that eliminates the least successful strategies. Since agents in our model can have various utility functions, as specified by the α parameter, the strategy elimination and cross-over operations are contained in the α-population. Additionally, strategies for different α values might not be the same.

A strategy in the case of the Tallinn mechanism is simply a bit-string. A bit-string is a series of 1-s and 0-s, which respectively stand for reported and not reported preference. Since we limit our agent's preferences to k = 15, the length of the bit-string is 15 bits. Since the Tallinn mechanism is limited to just three preferences, the bit-string can contain at most three bits set to one. For example, a possible strategy for agent i might be a_i = 100110000000000; that is, the agents with this strategy would report their first, fourth and fifth preference.

We run our genetic algorithms for a fixed (2000) number of steps. In each step, an allocation is made based on the Tallinn mechanism and we get the utilities for each agent. Then based on the rules of the genetic algorithm the strategies evolve. In Algorithm 2, we present a simple genetic algorithm (e.g. Riechmann, 2001; Simon, 2013). It consists of three operations: selection, crossover and mutation. The selection operator selects strategies with replacement and probability proportional to the gained utility. The cross-over operation randomly selects the value from either strategy for each position. Finally, with a small 0.05 probability we mutate the new strategy.

We evaluate four versions of genetic algorithms:simple genetic algorithm; genetic algorithm with election; genetic algorithm with stud selection; and genetic algorithm with elitism. The last three are slight modifications of the simple genetic algorithm. In the election modification, the agents remember their previous strategy and the corresponding utility. Before the selection operation in the next allocation, each agent picks the strategy with a higher utility from the previously remembered and the newly evaluated strategies (e.g. Riechmann, 2001). In the stud selection, we pick the top 20% of strategies with higher utility and always set one of the strategies in the cross-over operator to be in the top 20% (Simon, 2013). In addition, we ignore the bottom 10% of strategies. In elitism, we keep the top 20% of strategies fixed and only use the remaining strategies in the crossover (Simon, 2013).

Figure 3: Mean Utility

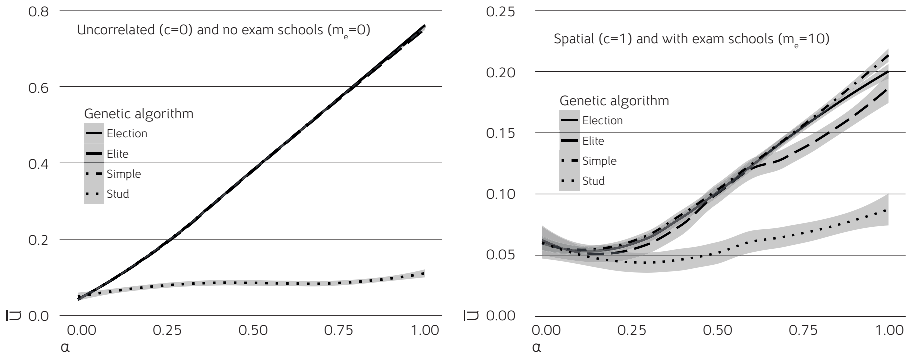

Source: Authors' calculation

Figures 3 and 4 show the results from the four variations of genetic algorithms. We see that the stud selection usually performs the worst, has the lowest utility and highest variation in utilities compared to the other variations. If preferences are spatially correlated and there is a large number of exam schools ($m_e = 10$), we can see that the simple genetic algorithm does slightly better with large values of α than with alternatives. For lower values of α, the simple model is statistically equivalent to the election and in some cases to elite selection. As the simple model does as good as others we further analyse the results from the simple optimisation method.

Figure 4: Ratio of Variance and Mean Utility

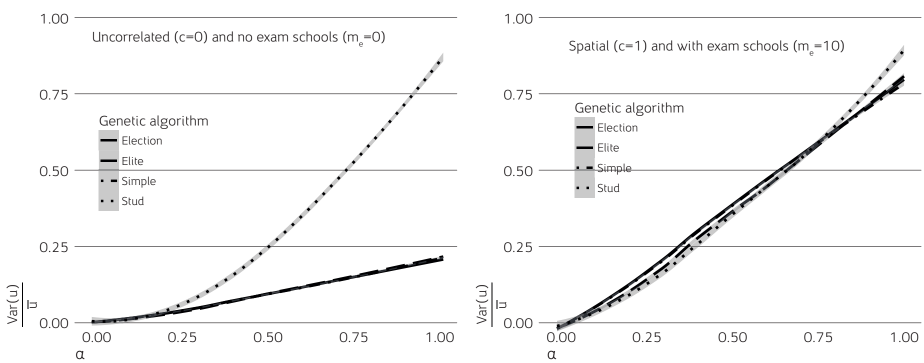

Source: Authors' calculation

4. Results

4.1. Expected Utility Maximising Strategies

The reported results are divided into four cases. In all of the figures illustrating the results, in the upper left corner the results with no correlation (random) preferences and no exam schools are indicated; in the upper right corner, the results with spatial (2D Euclidean) preferences and no exam schools; in the lower left corner, random preferences and ten exam schools; and in the lower right corner, correlated preference lists and 10 exam schools. Figures 5, 6 and 7 show a plot with the average and the standard deviation over multiple experiments. The standard deviation is often small so it is not always visible on the charts.

Figure 5: Reported Strategy Length

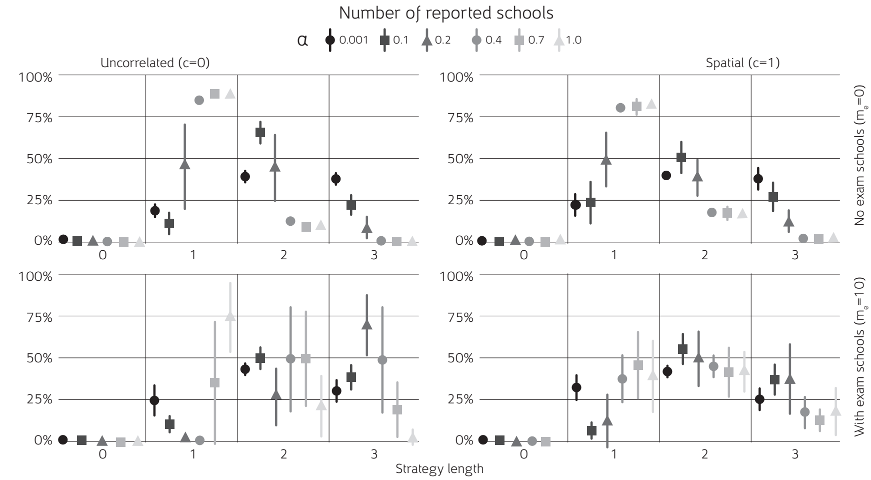

Source: Authors' calculation

Figure 6: Reported Preference by Utility Coefficient α

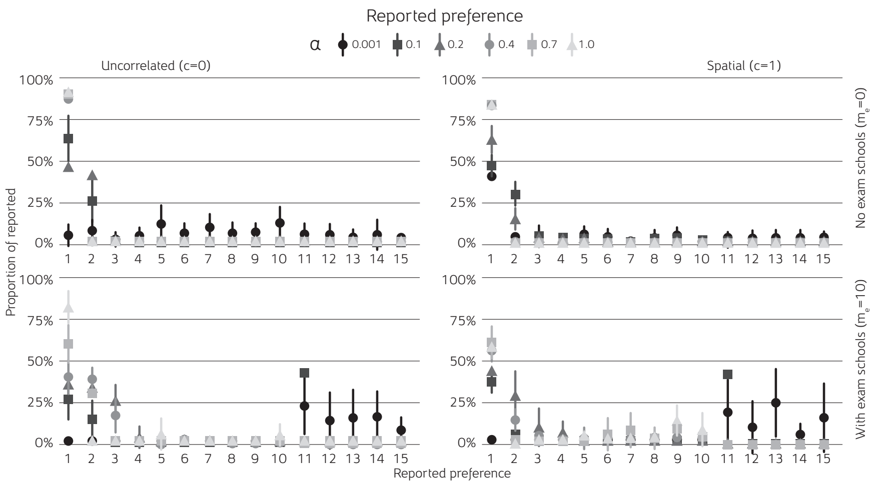

Source: Authors' calculation

Firstly, we are interested in the strategy length – the number of schools to be reported. In Figure 5, we show the strategy length by a proportion of the respective α-population. In general, it is elucidated that the decay in the utility function is a significant determinant of a good strategy. When $\alpha \approx 0.0$, it is best to randomly select the number of schools to report with roughly uniform probability. When $\alpha \geq 0.4$ and there are no exam schools, it would almost always be best to report only one school. With random preferences, when $\alpha \approx 0.2$, there is a phase transition in the number of schools to report, as the variation at this point is largest. Therefore, it is really difficult to pick a good strategy for how to report.

General trends show an increase in standard deviation in the strategy length when moving from spatial preferences to random preferences, or from having no exam schools to having 10 exam schools. Spatial preferences are aligned with school priorities, resulting is more predictable matches; therefore, the resulting strategies have a lower standard deviation. Standard deviation can also be interpreted as the uncertainty of the resulting match, when playing a certain strategy. In regard to exam schools, the uncertainty is greater than in the case of no exam schools, and even greater when the preferences are random in addition to exam schools.

Secondly, we are interested in how the mixed nature of the market – exam schools which are always preferred to regular neighbourhood schools – affect good strategies. We see that in the case of random preferences for high α, it is still often optimal to only report a single school. For medium α, the best strategy is to report 2 or 3, and only with low α (i.e. marginal utility is almost constant) is it best to randomly select the number of schools. If we assume that parents do not have a preference between the top three exam schools, they report the maximum number of preferences.

Figure 7: Reported Preference by Strategy Length

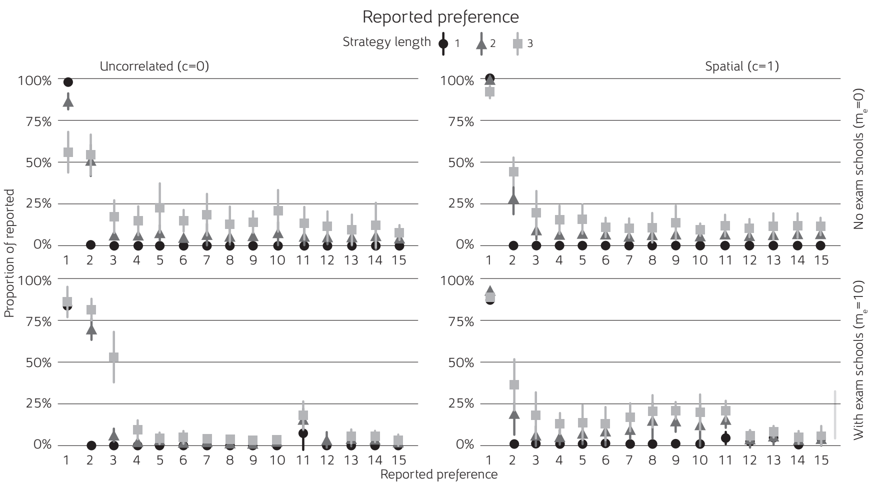

Source: Authors' calculation

We are also concerned with what to report. Figures 6 and 7 show the preferences reported by the agents' α and the strategy length. We see that without exam schools it is almost always (≈ 90%) optimal for $\alpha \geq 0.4$ to report from the top of the preference list, namely just their first preference. In regard to exam schools and random preferences, optimal reporting depends more on α, but generally the top three schools are reported. Figure 6 illustrates that in the case of exam schools and spatial preferences with high α, it would be better to report something from the higher and lower ends of exam schools, skipping the middle. Reporting schools lower on the preferences lists probably indicates that those agents would be otherwise unassigned, due to high demand, so they gain at least some utility. For medium α, the first three preferences are almost equally good. For indifferent agents, $\alpha \approx 0.0$, it would be best to randomly pick some schools from the list of regular schools. Also for agents with $\alpha = 0.1$, it would be beneficial to specify their most preferred exam school and most preferred regular school.

In Figure 7 the preferences are reported with different strategy lengths. The results show that it is always best to at least report one's most preferred school, as one might get lucky. If reporting more schools, it is useful to add the second most preferred school or with a small probability select something from even lower on the preference list. However, when reporting three choices, the selection of schools depends on the state of the school market. When preferences on the market in general are random with 50% probability, the first two preferences should be reported and the remaining options uniformly from the remainder of the preferences. In the case of spatially correlated lists or exam schools, the most preferred school should be almost always given. And when preferences are generally spatial, select the remaining options randomly. On the other hand, with exam schools and uncorrelated preferences when it is best to report three schools, it is usually best to report the top three.

4.2. Social Welfare

Previously we investigated the individual behaviour of agents, but now we consider how these behaviours influence the outcome for the entire society. For this, we compare the results of the Tallinn mechanism to the widely used Deferred-Acceptance (DA) mechanism (Gale and Shapley, 1962; Abdulkadiroğlu and Sönmez, 2003) as described in Section 2.3. Similar to the Tallinn mechanism, the priorities in the Deferred-Acceptance mechanism are also only based on distance.

Figure 8: Unassigned Agents

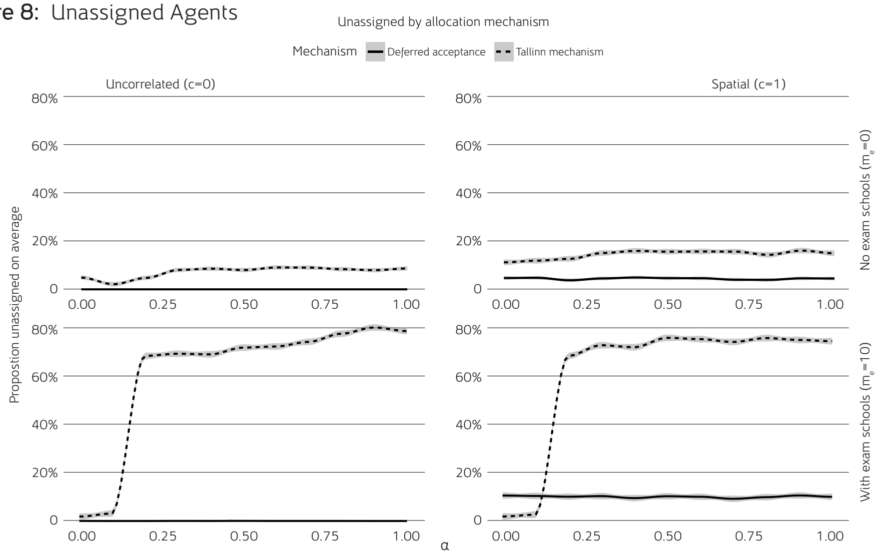

Source: Authors' calculation

We look at two measures of social welfare. First, the proportion of unassigned agents (Figure 8) and second the mean utility in the allocation (Figure 9). Usually, the measure used in matching problems are the allocated preferences, but this is mostly due to not having access to the utility. Since in our experiments, we know the agent's utility, we measure the mean utility over all the agents.

Figure 8 illustrates assignment probability based on the agents' α. We see that by using the DA mechanism and assuming random preferences (c = 0.0) and no exam schools (m_e = 0), there are no unassigned agents. When preferences are spatially correlated (c = 1.0), we can see that about 10% of students are unassigned, and this probability does not depend on agent type.

Figure 9: Mean Utility Comparison: Deferred-Acceptance and Tallinn Mechanism

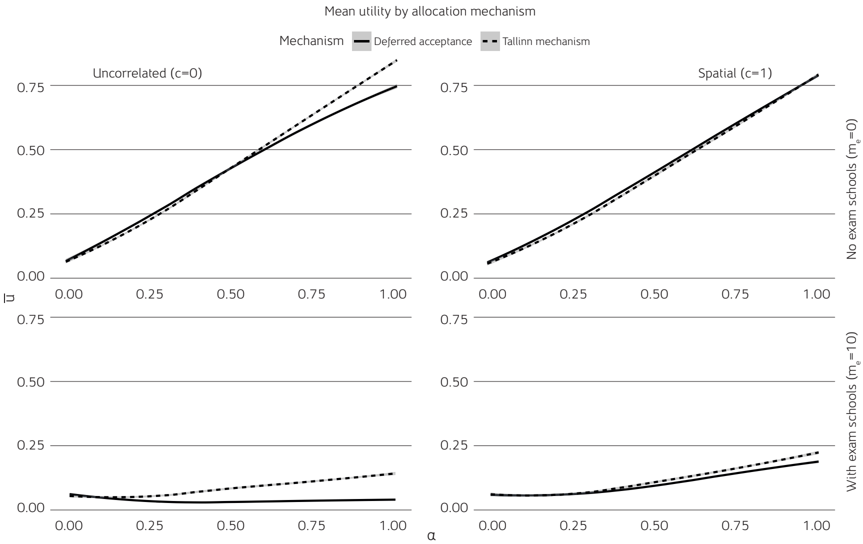

Source: Authors' calculation

As described in section 3.2., we have m_e = 10 exam schools that are always first on the agents' preference lists. Under such circumstances, only a small fraction of students receive a position in the top ten schools of their preference. Since there are fifty schools, exam schools account for 20% of places, so 20% of students receive a place in one of their top ten schools. Again, with uncorrelated preferences, DA can guarantee a place for all the students. Naturally, the students might receive a less preferred school. In the case of spatial preferences (c = 1.0, m_e = 10), even with DA, a significant number of students - about 10% - would be left unassigned. With the Tallinn mechanism, the number of unassigned students would be even higher – about 70% of students who have α > 0.2 would be unassigned. This is mainly due to agents maximising their expected utility and do not have a negative utility by being left unassigned.

In Figure 9 we show the expected utility under the two mechanisms. Expected utility is higher in the Tallinn mechanism compared to the DA results. This is because the strategies of the agents maximise their expected utility, which is greater under the Tallinn mechanism than under DA. A similar result was discovered in the manipulable Boston mechanism (Abdulkadiroğlu et al. 2011). This leads to the conjecture that manipulable mechanisms provide the option to maximise an agent's expected utility at the risk of being unassigned or assigned to a low ranked preference. Yet, as a result, a large number of agents are unassigned in the Tallinn mechanism.

We observe that agents with $\alpha = 0.999$ would maximise their expected utility by only reporting their first preference (Figure 5 and 6). This is due to the high utility value of their first preference, but because only a few preferences are reported, there is also a large probability of being unassigned under the Tallinn mechanism (Figure 8).

When agents are not particularly concerned with the school they are allocated to (α is small), the best strategy is to report randomly (see Figure 6). This also guarantees that students will not be unassigned, which is demonstrated in Figure 8. Other agents trade the probability of being unassigned with being assigned to a more preferred school. We see that for agents who have $\alpha \geq 0.3$, there is a high probability of being unassigned. However, there must be a considerable number of agents who are assigned to their top preferences on the condition of there being no exam schools, which increases the average utility from the allocation.

5. Conclusion and Discussion

Our aim was to contribute to the mechanism design literature about school choice by adding a description of the Tallinn mechanism, which is a centralised school-selecting assignment based on the student's distance from the school. Moreover, we wanted to indicate what the manipulative behaviour of agents is under such a mechanism; that is, how many preferences they report and how truthful their preference revelation is.

We used computational experiments to show the near-optimal strategies of the agents. For optimisation, we used a simple genetic algorithm, which outperformed the alternatives. Our model setup was the following: 50 schools (10 exam schools), 60 seats in each school and 3,000 agents. The agents (families) were heterogeneous, but their spatial preferences could have been correlated. Therefore, our emphasis in comparative static analysis has been on three parameters – the shape of the utility function of the agents, the number of exam schools and the correlation in the preferences of the agents. The first parameter space (α) illustrates the decreasing utility over alternatives and makes it possible to study cardinal preferences. The second parameter makes it possible to study case specificity – exam schools are popular schools at the centre of the city that are preferred by most families due to public information from league tables or from their reputation according to "hot knowledge". The third parameter makes it possible to indicate the effect of the homogeneity-heterogeneity of the agents. Homogeneity of agents can be interpreted as a post-Soviet tendency towards non-diversity of "good taste" – correlated preferences show that agents have similar preferences for schools. However, we used spatial preferences and we always put exam schools at the top of the list. This action is justified by empirical evidence (Põder and Lauri, 2014).

Our results show that in many circumstances under the Tallinn mechanism it is often best to report only one school, even if there is an option to report multiple schools. It is rarely beneficial to report three options (the maximum number). Nevertheless, it would benefit agents to report a school from the top of their preference lists. When reporting three schools, it is not always best to report the top schools and it seems to be advantageous to select the third option uniformly randomly from the remaining preferences. For agents with near-zero marginal utility, if they exist, it is best to report schools randomly. Additionally, the

Tallinn mechanism maximises the expected utility of the agents, if the agents learn what and how to report, but also runs a large risk of agents not being assigned to schools. The maximisation of expected utility seems similar to a similar phenomenon in the Boston mechanism (Abdulkadiroğlu et al., 2011) given that families know how to manipulate and might be a more general property of manipulable mechanisms.

Finally, we were interested in whether the Tallinn mechanism hurts families compared to a strategy-proof stable mechanism such as the Deferred-Acceptance mechanism. We saw that the number of unassigned students is much higher under the Tallinn mechanism. This can partially be interpreted as an inefficiency of behaviour due to the mechanism. However, there is no considerable mean welfare effect – agents optimise their utility maximising strategies under the Tallinn mechanism.

We see that we manage to find beneficial strategies under the Tallinn mechanism; however, due to the non-repetitive nature of the game, real-life learning can be relatively limited for most of the families. Nevertheless, as a stylised fact about the reporting of preferences indicated, agents learn not to report the maximum number of preferences, rather they limit their reported lists. In addition, in the case of exam schools, they tend to report schools from the top of the list, yet there remains a high probability of local regular schools also being reported. This could be the “learning effect” – the Tallinn mechanism prioritises neighbourhood kids by using the cardinal measure of distance.

In conclusion, it was demonstrated that post-Soviet school-proposing mechanisms use some properties of the central marketplace that are open to manipulation – such mechanisms force families to learn strategic behaviour by reporting non-truthful preferences. In this respect, the Tallinn mechanism is similar to the infamous Boston mechanism. Moreover, it was shown that both would result in a higher expected utility for the agents compared to the optimal, stable and strategy-proof Deferred-Acceptance mechanism, which might be the property of generally manipulable mechanisms.

References

Abdulkadiroğlu, A., Agarwal, N., and Pathak, P.A. 2015. The Welfare Effects of Coordinated Assignment: Evidence from the NYC HS Match. Technical Report May, National Bureau of Economic Research, Cambridge, MA.

Abdulkadiroğlu, A., Che, Y.-K., and Yasuda, Y. 2011. Resolving Conflicting Preferences in School Choice: The Boston Mechanism Reconsidered. *American Economic Review*, Vol. 101, No. 1, pp. 399–410.

Abdulkadiroğlu, A., Pathak, P. A., and Roth, A. E. 2005a. The New York City High School Match. *American Economic Review*, Vol. 95, No. 2, pp. 364–367.

Abdulkadiroğlu, A., Pathak, P. A., and Roth, A. E. 2009. Strategy-proofness versus Effi- ciency in Matching with Indifferences: Redesigning the NYC High School Match. *American Economic Review*, Vol. 99, No. 5, pp. 1954–1978.

Abdulkadiroğlu, A., Pathak, P. A., Roth, A. E., and Sönmez, T. 2005b. The Boston Public School Match. *American Economic Review*, Vol. 95, No. 2, pp. 368–371.

Abdulkadiroğlu, A., Pathak, P. A., Roth, A. E., and Sönmez, T. 2006. Changing the Boston School Choice Mechanism. *National Bureau of Economic Research, Working Paper Series*, 11965.

Abdulkadiroğlu, A. and Sönmez, T. 2003. School Choice: A Mechanism Design Approach. *American Economic Review*, Vol. 93, No. 3, pp. 729–747.

Arifovic, J. 1994. Genetic Algorithm Learning and the Cobweb Model. *Journal of Economic Dynamics and Control*, Vol. 18, No. 1, pp. 3–28.

Arifovic, J. 1996. The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies. *Journal of Political Economy*, Vol. 104, No. 3, pp. 510–541.

Aygün, O. and Bo, I. 2013. College Admission with Multidimensional Reserves: The Brazilian Affirmative Action Case. Available at: https://www2.bc.edu/inacio-bo/AygunBo2013.pdf (Accessed 02.08.2016).

Balinski, M. and Sönmez, T. 1999. A Tale of Two Mechanisms: Student Placement. *Journal of Economic Theory*, Vol. 84, No. 1, pp. 73–94.

Bogomolnaia, A. and Laslier, J.-F. 2007. Euclidean Preferences. *Journal of Mathematical Economics*, Vol. 43, No. 2, pp. 87–98.

Chen, S.-H. (Ed.). 2002. *Genetic Algorithms and Genetic Programming in Computational Finance*. Springer US, Boston, MA.

Chen, S.-H., Kampouridis, M. and Tsang, E. 2011. Microstructure Dynamics and Agent-Based Financial Markets. In: Bosse, T., Geller, A., and Jonker, C. M. (Eds.). *Multi- Agent-Based Simulation XI*, pp. 121–135. Springer: Berlin Heidelberg.

Chen, S.-H. and Tai, C.-C. 2010. The Agent-Based Double Auction Markets: 15 Years On. In: Takadama, K., Cioffi-Revilla, C., and Deffuant, G. (Eds.). *Simulating Interacting Agents and Social Phenomena*, pp. 119–136. Springer: Japan, Tokyo.

Duffy, J. 2006. Agent-Based Models and Human Subject Experiments. In: Tesfatsion, L. and Judd, K. L. (Eds.). *Handbook of Computational Economics Volume 2 - Agent-Based Computational Economics*, ch.19, pp. 948–1012. North-Holland.

Dur, U.M., Kominers, S.D., Pathak, P.A., and Sönmez, T. 2013. The Demise of Walk Zones in Boston: Priorities vs. Precedence in School Choice. *NBER Working Paper Series* 18981.

Erdil, A. and Ergin, H. 2008. What's the Matter with Tie-Breaking? Improving Efficiency in School Choice. *American Economic Review*, Vol. 98, No. 3, pp. 669–689.

Erdil, A. and Kumano, T. 2013. Prioritizing Diversity in School Choice. Available at: http://www. matching-in-practice.eu/wp-content/uploads/2013/09/Erdil-Prioritizing_ Diversity.pdf (Accessed 02.08.2016).

Ergin, H. and Sönmez, T. 2006. Games of School Choice under Boston Mechanism. *Journal of Public Economics*, Vol. 90, pp. 215–237.

Fragiadakis, D. and Troyan, P. 2013. Market Design under Distributional Constraints: Diversity in School Choice and Other Applications. Available at: http://tippie.uiowa.edu/ economics/ tow/papers/troyan-spring2014.pdf (Accessed 02.08.2016).

Franek, J. and Kresta, A. 2014. Judgment Scales and Consistency Measure in AHP. *Procedia Economics and Finance*, Vol. 12, pp. 164–173.

Gale, D. and Shapley, L.S. 1962. College Admissions and the Stability of Marriage. *The American Mathematical Monthly*, Vol. 69, No. 1, pp. 9–15.

Knuth, D.E. 1997. Seminumerical Algorithms. Addison-Wesley, Reading, MA, 3rd edition.

Kominers, S. D. and Sönmez, T. 2013. Designing for Diversity in Matching. Boston College Working Papers in Economics 806.

McAfee, Preston, R. and Mcmillan, J. 1996. Analyzing Airwaves Auction. *The Journal of Economic Perspectives*, Vol. 10, No. 1, pp. 159–175.

Milgrom, P.R. 2000. Putting Auction Theory to Work: The Simultaneous Ascending Auction. *Journal of Political Economy*, Vol. 108, No. 2, pp. 245–272.

Mullin, F.J. and Stalnaker, J.M. 1952. The Matching Plan for Internship Placement: A Report of the First Year's Experience. *Journal of Medical Education*, Vol. 27, No. 3, pp. 193–200.

Pathak, P.A. and Shi, P. 2013. Simulating Alternative School Choice Options in Boston. Technical report, MIT School Effectiveness and Inequality Initiative.

Pathak, P.A. and Sönmez, T. 2008. Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism. *American Economic Review*, Vol. 98, No. 4, pp. 1636–1652.

Pathak, P.A. and Sönmez, T. 2013. School Admissions Reform in Chicago and England: Comparing Mechanisms by their Vulnerability to Manipulation. *American Economic Review*, Vol. 103, No. 1, pp. 80–106.

Põder, K. and Lauri, T. 2014. When Public Acts like Private: The Failure of Estonia's School Choice Mechanism. *European Educational Research Journal*, Vol. 13, No. 2, pp. 220–234.

Riechmann, T. 2001. Learning in Economics. Springer-Verlag: New York.

Romero-Medina, A. 1998. Implementation of Stable Solutions in a Restricted Matching Market. *Review of Economic Design*, Vol. 3, No. 2, pp. 137–147.

Roth, A.E. 1982. The Economics of Matching: Stability and Incentives. *Mathematics of Operations Research*, Vol. 7, No. 4, pp. 617–628.

Roth, A.E. 1984. The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory. *The Journal of Political Economy*, Vol. 92, No. 6, pp. 991–1016.

Roth, A.E. 2002. The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics. *Econometrica*, Vol. 70, No. 4, pp. 1341–1378.

Roth, A.E. 2008. Deferred Acceptance Algorithms: History, Theory, Practice, and Open Questions. *International Journal of Game Theory*, Vol. 36, No. 3-4, pp. 537–569.

Simon, D. 2013. *Evolutionary Optimization Algorithms: Biologically-inspired and Population-based Approaches to Computer Intelligence*. Wiley.

Ünver, M.U. 2001. Backward Unraveling Over Time: The Evolution of Strategic Behavior in the Entry Level British Medical Labor Markets. *Journal of Economic Dynamics and Control*, Vol. 25, No. 6-7, pp.1039–1080.

Ünver, M.U. 2005. On the Survival of some Unstable Two-Sided Matching Mechanisms. *International Journal of Game Theory*, Vol. 33, No. 2, pp. 239–254.

Varian, H. 2006. *Intermediate Microeconomics*. W. W. Norton & Company, New York, seventh edition.

Veskioja, T. 2005. *Stable Marriage Problem and College Admission*. PhD thesis, Tallinn University of Technology.

Zhu, M. 2014. College Admissions in China: A Mechanism Design Perspective. *China Economic Review*, Vol. 30, pp. 618–631.