



EFFICIENCY AND FAIR ACCESS IN KINDERGARTEN ALLOCATION POLICY DESIGN

André Veski

Tallinn University of Technology,
Estonia

andre.veski@ttu.ee

Péter Biró

Hungarian Academy of Sciences,
Hungary

biro.peter@krtk.mta.hu

Kaire Põder

Estonian Business School, Estonia

kaire.poder@ebs.ee

Triin Lauri

Tallinn University, Estonia

triin.lauri@tlu.ee

ABSTRACT

We examine Kindergarten allocation practices in an Estonian municipality, Harku. Based on our recommendations, the allocation mechanism in Harku was redesigned in 2016. The new mechanism produces a child-optimal stable matching, with priorities primarily based on siblings and distance. We evaluate seven policy designs based on 2016 admission data in order to understand efficiency and fairness trade-offs. In addition to the descriptive data analysis, we conduct a counter-factual policy comparison and sensitivity analysis using computational experiments with generated preferences. We fix the allocation mechanism to be the child-oriented Deferred-Acceptance algorithm, but we vary how the priorities are created by altering sibling and distance factors. Different lotteries are included for breaking ties. We find that different ways of considering the same priority factors can have a significant aggregate effect on the allocation. Additionally, we survey a dozen special features that can create significant challenges (both theoretical and practical) in redesigning the

This work was partially supported by Norway Grants V653. André Veski acknowledges support from European Cooperation in Science and Technology (COST) action IC1205 and Estonian IT Academy scholarship for PhD students. Péter Biró acknowledges support from the Hungarian Academy of Sciences under its Momentum Programme (LP2016-3/2016) and the Hungarian Scientific Research Fund, OTKA, Grant No. K108673. We are thankful for the comments of two anonymous reviewers, especially for the recommendation to include results with TTC algorithm. Any errors are our own.

allocation mechanism in Estonian Kindergartens, and potentially elsewhere as well.

Keywords: Kindergarten allocation; policy design; mechanism design

JEL Classification Numbers: C78, D47, D50, H75

1. INTRODUCTION

Families have become a much-debated issue in all developed countries and they form the focal point of debates about “new risks” and the much needed “new policies” for Western welfare states. The questions of who should care for children, to what extent and for how long, lie at the centre of conflicts about the values that shape not only policies and struggles around policies, but also individual and family choices (Saraceno, 2011). Moreover, in Eastern Europe, the Soviet legacy has paved the way for the dominance of publicly provided care, but in many countries, including the case examined here, there is a shortage of early childhood care places for children aged 18 months to three years. This shortage of places has forced municipalities, who are the main providers, to set priorities for the allocation of these places. Priorities are aimed not only at solving the problem of oversubscription, but also at implementing social goals. Thus, we conceptualise the process of implementing priorities accompanied with allocation principles (matching design) as policy design.

Policy design entails taking the approach of a matching mechanism design in order to propose a good way to allocate children to Kindergartens. There are process descriptions about the (re-)design of school choice mechanisms, e.g. in various cities in the US (Pathak & Sönmez, 2013; Pathak & Shi, 2013; Ergin & Sönmez, 2006) and in Amsterdam (de Haan et al., 2015). Nevertheless, to the best of our knowledge, our paper is the first to report such a redesign of a Kindergarten place allocation mechanism. However, our theoretical foundation relies on the mechanism design literature motivated by related applications, such as school choice (Abdulkadiroğlu & Sönmez, 2003; Abdulkadiroğlu et al., 2005a), college admissions (Biró et al., 2010b; S.-H. Chen et al., 2012) and job assignments (Roth, 2008). Mechanism design provides methods for allocation under given welfare criteria and selection priorities, but it does not prescribe the way in which these priorities should be applied. The general policy considerations for school choice are the allocation of siblings to the

same school and the proximity of the school. Some countries also use some affirmative action measures, e.g. prioritising children of low socio-economic status. Similar principles are applicable to our Kindergarten policy design case study while aiming for the clear-cut implementation and operationalisation of policies. The latter not only concerns a clear definition of proximity as a priority (i.e. defined as a walk-zone (Shi, 2015) or a continuous cardinal measure (West et al., 2004)) or the ordering of priority classes but also allows for the implementation of welfare considerations in policy evaluation.

Our welfare considerations aim at two social goals: efficiency and fairness. We define efficiency as the ability of a policy to meet predefined goals, the utility of families (high rank in their preferences and siblings in the same Kindergarten) accompanied with social goals such as minimising the travel distance or time to Kindergartens. Defining fairness is more problematic and entails more uncertainty. Our definition of fairness is based on the idea of equal access. It is operationalised by the probability that the child is assigned to her first preference.

Our study considers a local municipality Harku, in Estonia. Instead of implementing certain social goals by policy design, the most commonly used priority in Estonian municipalities is the date of application, while in limited cases, catchment areas are applied to ensure proximity. Children are ordered on the basis of the application date in a manner similar to a serial dictatorship mechanism, thus forcing one-sided matchings without enabling the implementation of affirmative action policies or social goals, such as fairness. In addition, parental preferences are either not considered or have been limited. In the Harku case, the number of preferences was bounded by three until 2015. The latter restriction implies that preferences are not revealed truthfully and moreover, the matching has been done manually.

Between 2014 and 2016, as part of an Estonian project we collaborated with representatives of the Harku municipality. We monitored their 2015 allocation practice and suggested a revision which led to a transitory system in 2016. In the 2016 allocation, the standard student-proposing Deferred-Acceptance mechanism was used under a special priority setting which is described in detail in Section 2.2. This mechanism is known to be strategy-proof, and the parents were encouraged to submit full preference lists, so we can expect the submitted applications to be truthful. We made a comparative assessment of policies using the 2016 data. As an input we used preference data collected from 152 families who have the right to a Kindergarten place.

In the assessment, we proposed seven different policies which consist of different metrics of indicating distance (as absolute, relative or binary measures), siblings, quotas; and their priority order. Ties are broken by assigning random numbers either with a single or with multiple lotteries. Our research methods are partially inspired by Shi (2015), but we investigated some novel policies as well. Perhaps the most interesting aspect of these policies is the way that distance is used in the priorities.

The classic way of creating proximity priorities is the *catchment area* system, where the city is partitioned into areas and the students living in an area have the highest priority in all schools in that area. This simple method can be seen as unfair, as one student can have a higher priority than another student, even though the actual distance of her location to the school is greater than for the other child. Therefore instead of catchment areas, most applications have switched to absolute or relative distance based priorities. The simplest absolute distance based policy is the walk-zone priority scheme, used in many US cities (e.g. New York (Abdulkadiroğlu et al., 2005a)), where the children living within a well-defined walking distance are in the high distance priority group for that school and the ties are broken by lottery. Strict priorities based on absolute distances are used in Sweden as well (Andersson, 2017). However, there were also discussions and court cases about the fairness of such absolute distance based priorities¹.

The absolute distance based priority schemes can be unfair for those living far from all (or most) of the (good) schools, therefore the so-called *relative distance* based methods are also commonly used in many applications (e.g. Calsamiglia & Güell (2014); Shi (2015)). The relative distance priority means that we give the highest priority to all children for their closest Kindergarten, no matter how far that is, and the children will be in the second priority group in their second closest Kindergarten, and so on. A rough version of this rule is to give high priority for all children in a given number of closest schools.

Barcelona changed its catchment area systems to a relative distance system in 2007. After the change, students have priority in at least six of their closest

¹ In the city of Lund parents have challenged allocation decisions in court based on an alternative option distance argument. The city used the absolute distance priority in their allocation, but some parents have found this policy unfair, as they would have to travel 1000m more to their second choice school than to their first choice school, whilst there was another student who would only need to travel 650m more if allocated to their second choice school rather than their first choice school. The court accepted this argument and gave a seat to the appealing student in their first choice school.

schools (Calsamiglia & Güell, 2014, Section 5.1)². In Boston, another relative distance policy was proposed recently by Shi (2015), mainly in order to reach the city's aim to cut down busing costs.

Note that there are also applications where the distance based priorities are considered unfair, as they can limit equal access to good schools. The Amsterdam school choice system (de Haan et al., 2015) does not use any distance based priority, only a pure lottery. In the Harku case, where Kindergartens are of more or less the same quality, the authority was in favour of using the distance based priorities in order to decrease the overall commuting costs and also to satisfy the preferences of the parents (typically for nearby Kindergartens). Based on the unfairness of the catchment area system described above, we only considered absolute and relative distance based priority approaches. We explain the distance-based priorities that we studied in more detail in Section 3 with examples.

Besides the distance, we also investigated different ways of taking the sibling priorities into account and also the way the lotteries are conducted in case of ties. The way the distance and sibling factors are considered has already been studied in the literature (Dur et al., 2013). The particular solution chosen for the 2016 transitory system is an interesting rotation priority scheme, which can lead to a well-balanced solution with respect to the two factors. Regarding the lotteries, we analysed the effects of using a single lottery for all Kindergartens compared to using multiple lotteries (one at each Kindergarten), and we have seen results similar to other research papers (Ashlagi & Nikzad, 2015).

As the second main contribution of our paper, we present a sensitivity analysis of various metrics of fairness and efficiency of policy designs based on counter-factual preference profiles. The policies that provide the best solutions for the current Harku data may not be ideal for other applications or robust for Harku, where the preferences of the parents are different. This can be the case in cities, or in other countries with different Kindergarten/school qualities, or for applications at different education levels (e.g. primary and

² “Before 2007, the city was divided into fixed neighbourhoods. The neighbourhoods varied in size for semi-public and public schools, but were conceptually the same. For semi-public schools, the neighbourhood coincided with the administrative district. For public schools, the neighbourhoods were smaller areas within the administrative district. The new neighbourhoods are based on distance between schools and family residences. An area (specifically, a minimum convex polygon) around every block of houses in the city was established to include at least the six closest schools (three public and three semi-public).”

secondary schools). Therefore, we found it important to investigate the effects of changes in priorities in the performance of different policies (i.e. different priority structures for the student-optimal Deferred-Acceptance mechanism). As a novel approach, we studied the fairness (or equal access) of the allocations measured by the probabilities of getting placed in the first choice schools.

In general our results indicate that preference structures, more precisely their endogeneity with respect to proximity, influence policy design. However, we advocate a relatively simple policy that prioritises siblings first and relative distance second. Relative distance gives all children priority in the closest Kindergarten independently of absolute distance from it. This policy is superior to others by our fairness criteria, especially when preferences of the families are aligned with policy priorities.

We structure our paper as follows. In Section 2 we review the practices and processes of Kindergarten choice of an Estonian municipality, Harku, before the process was redesigned on the basis of our recommendations in 2016. In Section 3 we define seven alternative policies and the descriptive statistics of our data, including our results from computational experiments. Finally, we discuss additional mechanism design challenges with some policy recommendations in Section 4 and give conclusions in Section 5.

2. MATCHING MECHANISM DESIGN

The design of an allocation mechanism is usually based on a two-sided matching market model, in this case between 1) families and 2) Kindergartens. Participants on both sides have linear orderings over the participants on the other side. Families have preferences over Kindergartens and they seek to get places at their most preferred Kindergartens. Kindergartens have a priority ranking over children. Priorities become important if there are fewer places available in a particular Kindergarten than the number of families who would like to be allocated to that Kindergarten. In those circumstances, Kindergartens accept children who are higher on their priority list, which in practice usually means children who live closer and/or who have a sibling in the Kindergarten. Kindergartens do not seek to admit higher priority children. This practice is different from some applications of two-sided markets. In college admissions for example (Gale & Shapley, 1962), both students and colleges seek to get more preferred matches, therefore they might act strategically in the allocation mechanism.

There are two prominent strategy-proof mechanisms for solving matching problems, the Deferred-Acceptance (DA) and the Top-Trading Cycles (TTC) mechanisms (Abdulkadiroğlu & Sönmez, 2003). The DA mechanism guarantees that no preferences and priorities (policies in our case) are violated, and there is no child who could get a place in a more preferred Kindergarten by priority, so there are no blocking pairs. A matching with no blocking pairs is called *stable*. A blocking pair can also be seen as a child having justified envy, since there is a family that would prefer a Kindergarten that either has free places or has accepted a child with lower priority. These kinds of justified envy situations are not tolerated in most applications (Pathak & Sönmez, 2013), and are sometimes even prohibited by law. Thus, stability is a crucial property for most applications.

While there are potentially a number of stable allocations (Knuth, 1997), the child-proposing DA mechanism that is usually implemented results in the best possible preference for all families among the stable solutions, and this option also makes it safe for the families to reveal their true preferences.

The theoretical properties and disadvantages of DA were studied by Haeringer & Klijn (2009), backed by evidence from laboratory experiments (Calsamiglia et al., 2010) and by practical applications across the world (Pathak & Sönmez, 2013). In addition to advocating DA, the main policy implications of these studies indicate that for an efficiency gain, it is advised to increase the bounds on the number of collected preferences or to abolish the limit on the number of submitted preferences.

Before its redesign, the application process of the Harku municipality had many design features, but it was not a transparent system. Families could submit up to three ordered choices. The application date and the home address were also collected. The application date was relevant for the allocation, as families with an earlier application date had higher priority. Therefore, families tended to submit their applications as early as possible, usually a few weeks after the birth of the child. The application data typically remained unchanged until the actual allocation occurred, which could make the originally true preferences out of date (e.g. it was possible that the family moved to a different place or their older sibling has received a place in a different Kindergarten during the waiting period). The address could be a factor, as some heads of Kindergartens considered it when assigning places. Secondly, a qualifying condition for a Kindergarten place is that the parents have to be registered residents in Harku, and residency is based on where local taxes are collected.

Moreover, the matching was done manually using the following procedural rules. First, the number of vacant places was settled by January of each year, when the allocation process started. Place offers were made to families by the heads of Kindergartens if their Kindergarten was the first choice of the family. Second, if there were more families than places, then priority was given to the applications with earlier registration dates, although proximity or siblings could also be occasionally relevant. Third, if an offer was accepted, the child became assigned to the Kindergarten, otherwise that place was offered to the subsequent family on the waiting list.

In the case of unassigned children, the procedural rules were complicated and discretionary. Generally the heads of the Kindergarten communicated with each other to find a place for the children who remained unassigned. In the case of families who ordered popular Kindergarten at the top of their list and remained unassigned in the first round, their second or third choice was considered, although these could already be full. If that was the case, the families with an earlier application date would be rejected from their second choice because the children already assigned there had listed that Kindergarten as their first preference, irrespective of their application dates. Thus, some children were allocated to a less preferred Kindergarten, simply because of how the family ordered their preferences. This is a well-known property of the Immediate-Acceptance mechanism (e.g. [Abdulkadiroğlu & Sönmez, 2003](#)) and the procedure that had been used in Harku until 2015 was very similar to this.

2.1. Building a mechanism for Kindergarten seat allocation

Our redesign of the Harku Kindergarten allocation mechanism inspired by the literature has four main areas as described in Table 1. The application procedure before 2016 which was initiated by collecting preferences had several drawbacks. First, since parents could get higher priority if they applied earlier, they tended to apply soon after the birth of their child. However, during the subsequent three years, the preferences of the families could change. This situation was usually not reflected in the application data, thus resulting in a high number of cancellations. Second, families could only list their top three choices. Limited preference not only created a large number of unassigned children, but also manipulation with the revelation of preferences.

Our design changed the data collection procedure and the number of

Table 1: Redesign of Harku mechanism

2015	2016
Application procedure	
Applications are collected after the birth of the child due to prioritising according to application dates	Applications are collected from 1 January until 1 February for allocating places from 1 September of the same year
Limited preference lists	
Limited to three Kindergartens	List all Kindergartens they are willing to attend (no limit)
Priorities (policies)	
Not clearly defined	See Section 3.2 for policy design alternatives
Matching mechanism	
Decentralised mechanism which has some properties of Serial dictatorship and Boston (Immediate-Acceptance)	Deferred-Acceptance

preferences collected. Families make application in the matching platform³ during monthly period six months before the service delivery (1. September) and list all their preferences. Giving up application date as a priority will be a necessary result of the procedural amendments.

Finally, the central allocation mechanism applied until 2016 was not transparent, the priorities were not clearly defined or adhered to by the heads of the Kindergartens. The first priority of the application date was sometimes violated. Children with siblings were usually considered to have higher priority, but not always. Our design introduces clearly defined priority metrics and a centralised allocation system that ensures that the criteria are always followed. Moreover, instead of an unstable and manipulable Immediate-Acceptance mechanism we propose the child-proposing DA. This is a standard method for school choice (Abdulkadiroğlu & Sönmez, 2003), which eliminates justified envy, and gives incentive for the families to state their true preferences.

³ <https://www.haldo.ee/>

2.2. Particularities of the 2016 system

Before the final implementation of our platform-based matching design, there was a transitory system in place in Harku in 2016 that partially applied our design recommendations, but experimented with priorities. Families were asked to rank all seven Kindergartens. Additionally, the home address, application date, status of siblings and the child's birth date were collected. The allocation process was designed on the basis of the DA mechanism with slots (Dur et al., 2013) while policy transformation regarding fixing priorities was more complex. There were four types of priorities that are defined per position as follows, in the order of precedence:

1. siblings, distance, age, application date
2. distance, age, application date, siblings
3. application date, siblings, distance, age
4. age, application date, siblings, distance
5. siblings, distance, age, application date
6. distance, age, application date, siblings
- ...

The positions are considered in order, with families first applying to the first position, then the second position, etc. This can also be thought of as each Kindergarten being split into a number of seats, with each seat potentially having a unique priority criteria. Then, the preferences of the families are modified so that within each Kindergarten, they rank the position with the higher precedence more highly. If the number of available places is not exactly divisible by four, then some type of priorities might have more positions available than others.

The main reason for the complicated policy design or for considering the four types of priorities rotationally was backed by the argument of equal treatment. Granting equal opportunity to all "types of families" (the ones that have siblings; those living nearby; early applicants; and families with an older child) was the preference of the local municipality. In future allocations, the application date will not be used. It was used here as some families still had the expectation of being allocated by the application date.

The precedence order of priority classes matters in the allocation procedure, as shown by [Dur et al. \(2013\)](#) by demonstrating that a simple priority scheme might be discriminating for some groups. For instance, let us assume there are five seats with siblings and distance priority and a further five seats with only distance priority. There are more than five children with a sibling and in total more than ten children. If for the first five positions we would consider children with siblings and then by distance, this would be disadvantageous for children with siblings compared to first only considering distance and then siblings as well as distance. In the latter case, some children with siblings might already be allocated by distance alone, so other children with siblings have lower competition and a better chance of getting a desired place. On the other hand, it might occur that some children living closer have an unfair disadvantage. The aim of the rotating scheme is to balance these two effects. That leads us to the equal treatment issues related to policy design.

3. POLICY DESIGN

3.1. Efficiency and fairness

In mechanism design the goals are usually related to designing an allocation method that maximises a form of efficiency, while not violating some constraint(s). In the matching domain, the usual criterion is selecting a Pareto optimal matching among a set of stable matchings. In a public resource two-sided matching setting, e.g. school seats, usually in fact two selections are made: first, the priorities of applicants and second, the mechanism. In a school choice setting, the priorities are often based on siblings and distance, although there are other alternatives ([MatchinginPractice, 2016](#)). However, in designing the allocation mechanism these priorities are usually treated as a given.

When evaluating the allocation methods we concentrate on two main criteria: efficiency and fairness. Efficiency characterises the level at which we, as designers, can satisfy the preferences of the applicants. Thus, we look at the average allocated preference. We also include the percentage of applicants receiving their first preference as this is often the case and the average might not always be a good indicator.

In addition to efficiency and stability (lack of envy), our policy design is driven by equality concerns. In the literature on distributive justice, discussion on fairness (fair access in our case) is often accompanied by discussion on

the principles of affirmative action, i.e. the Rawlsian difference principle (Rawls, 1971). In our case, fair access is defined as the chance for the family to access their most preferred Kindergarten. Moreover, we include in our design some positive discrimination, or controlled choice, through policies such as prioritising siblings.

Fair access is essentially different from the efficiency metrics for the priorities of local municipalities and the preferences of families. The goal of fair access is to provide an opportunity for every child to get a place at their most preferred Kindergarten. As some families might live far away from all Kindergartens (see Appendix C), they would always be low on the priority list for any Kindergarten. We measure fair access as the proportion of families placed in their most preferred Kindergarten on two levels, at least 10% chance and 50% chance. This is similar to access to quality in (Shi, 2015) where quality, in addition to being ranked high, contains an objective quality metric. Since there is no quality ranking for a Kindergarten in our case and only a small number of Kindergartens we look at the probability of being allocated to the first choice. Since not all policy designs use lotteries, some will be inherently unfair in terms of fair access.

The mechanism also allows the local authorities to have social objectives, which are usually, but not always aligned with the preferences of the parents. The two most prominent goals are

- having siblings in the same Kindergarten, and
- placing children in a Kindergarten near their home.

Prioritisation of proximity and siblings is also recommended by the regulations responsible for the allocation of Kindergarten places (“Preschool Child Care Institutions Act”, 2014). While prioritising proximity and siblings is common practice in the case of school and Kindergarten choice design, being favoured as the means to sustain community cohesion and avoid unreasonable transportation costs (see Shi, 2015, for instance), this practice may cause various concerns. The proximity principle may lead to problems in segregated areas, where it may result in the concentration of children from a similar socio-economic background into the same Kindergartens. Further social objectives could be the prioritisation of disadvantaged families or children with special needs, but there was no access to this kind of information in the data, so those goals were disregarded in this study. However, the main goal is still to provide families with a place at their most preferred Kindergartens.

3.2. Operationalisation of policy designs

A short list of social objectives indicated in the previous section does not mean that policy designs are limited to two alternatives, as the priority structures for siblings and proximity have many variants. Children with siblings might always have priority over others, or might only be prioritised over families living further away. Proximity can also be considered in many different ways, such as a walk-zone or a catchment area or a geographical distance.

A simple way to consider geographic aspects is to define catchment areas for each Kindergarten, and prioritise the children living in the catchment area where the Kindergarten is located. The drawback of this method is that these priorities may not reflect the personalised distances, as a Kindergarten might be relatively far from an address in the same area, whilst another Kindergarten in a different area can actually be nearby. Therefore, it may be more appropriate to use personalised distances. We can use continuous (real) distances or discretise them somehow, for instance giving priority to a Kindergarten within a 10-minute walking distance, or giving priority to the closest, or several closest Kindergartens. Another option is to give high priority to a child in a number of nearby Kindergartens. A special version of the latter so-called *menu* system has been evaluated and used in Boston school choice (Shi, 2015). Below we specify the distance-based priorities that we used in our policies.

- Absolute: Strict priorities based on the personalised absolute distances between the child's location and the school, measured in walk time or kilometres.
- Walk-zone: Coarse priorities based on the above-described absolute distance. A child is in the high priority group for a school if she lives within a 10-minute walking distance to this school.
- Relative: Every child is in the highest distance-based priority group in her closest school, she is in the second highest priority group in the second closest school, and so on.
- 3 closest: A binary variant of the above-defined relative distance policy, where every child is in the high priority group of a school, if this school is among the three closest schools for this child.

When we consider the children in walk-zones to have a higher priority, followed by children with siblings, the following priority groups are obtained:

1. siblings in walk-zones, 2. children in walk-zones, 3. siblings, 4. the rest. Siblings could also be considered to have a higher priority, which would result in the priority groups: 1. siblings in walk-zones, 2. siblings, 3. children in walk-zones, 4. the rest. This simple classification is used in many US cities, such as New York ([Abdulkadiroğlu et al., 2005a](#)) and Boston ([Abdulkadiroğlu et al., 2005b](#)), together with a randomised lottery for breaking ties. The lottery can also be conducted in two ways, either as a single lottery which is used in all Kindergartens, or as multiple lotteries, one for each Kindergarten. The typical choice, used in most US school choice programmes and also in Irish higher education admissions ([L. Chen, 2012](#)), is the single lottery. We will investigate both in our computational experiments. This question is discussed further by [Ashlagi & Nikzad \(2015\)](#) and [Pathak & Sethuraman \(2011\)](#).

If it is considered undesirable that a high proportion of children get admitted by sibling priority, then one option is to set a quota for siblings, for example 50% of the places. In this case, there is high priority for siblings for only some proportion of the places available, and the remaining places are prioritised by distance only. In such a setting, how the allocation is implemented is crucial. It can be done by allocating the places for siblings first and then the remaining seats or in reverse. [Dur et al. \(2013\)](#) showed that the reverse approach can benefit children with siblings, and [Hafalir et al. \(2013\)](#) showed that reserving places for a certain minority results in a better allocation for the minority than limiting the quota for the majority does. Under the latter policy, both groups (minority and majority) could be worse off. We evaluate policy design by the reservation of places for siblings or for families living nearby. In Harku, only about 20% of children have a sibling, so 20% of the places were set to have a sibling priority.

The Deferred-Acceptance algorithm can be slightly modified to accommodate for reserves and quotas. The priority quotas can be considered as separate Kindergartens. In this variant, the child is first placed in a quota group high in the precedence order and, if rejected, the child is then placed lower, etc. Thus, each child will be placed in the highest possible precedence quota group.

In this study, in order to explore the described aspects, we settled on seven priority policies (summarised in [Table 2](#)) for evaluation:

- DA1. Children with siblings always have the highest priority and children living closer have higher priority. Priority classes would be considered in the order: 1) siblings; 2) walking distance.

Table 2: Summary of policies (priority order in parentheses)

Policy	Distance (D)	Siblings (S)	Lottery	Quotas (Precedence)
DA1	absolute (2)	(1)	no	no
DA2	walk-zone (2)	(1)	(3)	no
DA3	walk-zone (1)	(2)	(3)	no
DA4	3 closest (2)	(1)	(3)	no
DA5	absolute (2)	(1)	no	[80%, 20%] ([D, S+D])
DA6	absolute (2)	(1)	no	[20%, 80%] ([S+D, D])
DA7	relative (2)	(1)	(3)	no

- DA2. Children with siblings always have the highest priority, then children in the walk-zone have higher priority. The walk-zone is defined as a 10-minute walking distance from home. Additional ties are ordered by a random lottery for all Kindergartens. The order of priority classes is: 1) siblings + walk-zone; 2) siblings; 3) walk-zone; 4) the remainder.
- DA3. Children in the walk-zone always have the highest priority, then children with siblings have higher priority. Additional ties are ordered by a random lottery for all Kindergartens. The order of priority classes is: 1) siblings + walk-zone; 2) walk-zone; 3) siblings; 4) the remainder.
- DA4. Children with siblings always have the highest priority, and children have higher priority for the three closest Kindergartens. Additional ties are ordered by a random lottery for all Kindergartens. Priority precedence order: 1) siblings + one-of-three-closest; 2) siblings; 3) one-of-three-closest; 4) the remainder.
- DA5. Children with siblings have the highest priority for the reserved 20% of places, otherwise priority is by distance. Precedence order: 1) by distance up to 80%; 2) children with siblings + distance up to 20%; 3) remaining places, if any, by distance.
- DA6. Children with siblings have the highest priority for the reserved 20% of places, otherwise priority is by distance. Precedence order: 1) children with siblings + distance up to 20%; 2) remaining places, if any, by distance.

DA7. Children with siblings always have the highest priority, and children have higher priority in the closest Kindergarten, second highest in the second-closest, etc. Additional ties are ordered by a random lottery for all Kindergartens. Priority precedence order: 1) siblings; 2) closest-number.

To demonstrate the effect of policies we construct a simple example. Let us assume we have four children $C = \{c_1, c_2, c_3, c_4\}$ and four Kindergartens $K = \{k_1, k_2, k_3, k_4\}$. In Table 3 we show the distances between homes and Kindergartens. We have no children with siblings in this example.

Table 3: Distances between homes and Kindergartens (km-s)

km	k_1	k_2	k_3	k_4
c_1	.7	1.2	1.0	1.7
c_2	.4	.6	.3	.7
c_3	.9	.5	.4	.3
c_4	.8	.3	.9	1.0

Assuming that walk-zone distance is $\leq .6$ km, the resulting priorities are in Table 4. We can observe that with absolute distance or walk-zone the child c_1 would not have a high priority in any Kindergarten. However with the 3-closest policy, there is at least some chance of having the highest priority in some Kindergarten, and with relative distance, each child has the highest priority in at least one Kindergarten. While this is not always guaranteed with relative distance, the lottery has lower impact compared to the 3-closest policy.

3.3. Data and initial policy design comparison

From a total of 152 families, 151 ranked all seven Kindergartens and only one family submitted a single Kindergarten as their preference. Table 5 shows the number of available places in each Kindergarten. Also 37 (about 24% of) children have a sibling in one of the Kindergartens.

Table 6 compares the allocations over all the policies with the submitted preferences. The listed Harku allocation does not exclude those few families who declined their assigned place. However, many (115, i.e. 76%) of the families were allocated to their most preferred Kindergarten. Since most

Table 4: Distance priorities

	absolute (DA1)	walk-zone (DA2, DA3)
k_1	$c_2 \prec c_1 \prec c_4 \prec c_3$	$c_2 \prec \{c_1, c_3, c_4\}$
k_2	$c_4 \prec c_3 \prec c_2 \prec c_1$	$\{c_2, c_3, c_4\} \prec c_1$
k_3	$c_2 \prec c_3 \prec c_4 \prec c_1$	$\{c_2, c_3\} \prec \{c_1, c_4\}$
k_4	$c_3 \prec c_2 \prec c_4 \prec c_1$	$c_3 \prec \{c_1, c_2, c_3\}$

	3-closest (DA4)	relative (DA7)
k_1	$\{c_1, c_2, c_4\} \prec c_3$	$c_1 \prec \{c_2, c_4\} \prec c_3$
k_2	$\{c_1, c_2, c_3, c_4\}$	$c_4 \prec \{c_1, c_2, c_3\}$
k_3	$\{c_1, c_2, c_3, c_4\}$	$c_2 \prec \{c_1, c_3\} \prec c_4$
k_4	$c_3 \prec \{c_1, c_2, c_4\}$	$c_3 \prec \{c_1, c_2, c_4\}$

families ranked all Kindergartens and there are more places than children, no children remained unassigned.

For policies that included lotteries, we computed averages over 20 lotteries. In the parentheses we show the standard error over the lotteries. In addition, we compared policies using a single (S) lottery for all Kindergartens or multiple (M) lotteries, one for each Kindergarten.

By using a simpler policy such as the DA1, we saw that there are fewer families receiving a place at their first choice Kindergarten⁴ than with the transitory Harku priority system. Moreover, two children (about 5%) are not allocated to the same Kindergarten as their siblings with the transitory rule, but with most other policies all siblings end up in the same Kindergarten. The only exception to this is DA3, which has siblings as a second priority over walk-zone, and on average also allocated 95% of siblings in the same Kindergarten, but fewer children to their first preferences.

It seems that the transitory policy of Harku invoked the so-called *vacancy chains* (Blum et al., 1997), where at the expense of one child with a sibling several others could obtain better places along an augmenting path. In particular, by denying places for two children in the same Kindergarten as their sibling, around seven more families could obtain their first choices. This leads

⁴ A more detailed allocated preference data is available in appendix B

Table 5: Harku allocation

Kindergarten	Number of places
A	20
B	20
C	34
D	18
E	20
F	38
G	5
Total	155

to an interesting trade-off between the goals of satisfying the sibling priority or granting the first choice of slightly more parents.

In 2016, the allocations based on policies DA5 and DA6 were exactly the same. This indicates that the gain in allocating more children to their first preference with Harku's policy is not due to allocating children to a closer Kindergarten, but due to application date and age priorities. Therefore, if these two criteria are not to be used in future policies, we expect that the rotation scheme based only on siblings and proximity will provide allocations similar to DA1, DA5 and DA6, assuming that the proportion of children and seats is similar.

3.4. Policy sensitivity to preferences

When comparing policies, one may wonder how sensitive the results are to changes in the preferences of parents. This can also be important when applying our policy recommendations in other applications. In Kindergarten allocation, and sometimes also in school choice, when the Kindergartens are more or less of the same quality, the most important factor influencing the preferences of parents is the location. Therefore, we conducted a comparative study wherein the intensities of this factor in the preferences of parents is varied. We evaluated the efficiency and fairness of the alternative policies accordingly. For the generation of preferences, we use the locations and the information on the siblings from the 2016 preference data. The detailed description of how

Table 6: Year 2016 comparison of policies using reported preferences

Policy	Mean preference	First	Mean distance (km)	With siblings
Harku	1.68	115	4.24	95 %
DA 1	1.76	110	4.26	100 %
DA 2 (M) ^a	1.85 (0.01)	98.75 (0.61)	4.59 (0.02)	100 % (0.0 %)
DA 2 (S)	1.72 (0.01)	108.05 (0.61)	4.44 (0.01)	100 % (0.0 %)
DA 3 (M)	1.83 (0.01)	98.30 (0.79)	4.51 (0.02)	95 % (0.25 %)
DA 3 (S)	1.72 (0.01)	107.75 (0.38)	4.45 (0.02)	96 % (0.3 %)
DA 4 (M)	1.91 (0.01)	89.25 (1.06)	4.53 (0.02)	100 % (0.0 %)
DA 4 (S)	1.75 (0.01)	104.85 (0.7)	4.49 (0.01)	100 % (0.0 %)
DA 5	1.76	110	4.26	100 %
DA 6	1.76	110	4.26	100 %
DA 7 (M)	1.78 (0.01)	107.60 (0.47)	4.30 (0.01)	100 % (0.0 %)
DA 7 (S)	1.76 (0.01)	107.75 (0.47)	4.31 (0.01)	100 % (0.0 %)

^a For policies with lotteries, (M) indicates multiple tie-breaking lotteries and (S) single. The standard errors over lotteries are in parentheses.

we generate the preferences of parents can be found in the Appendix A.

We characterise preference profiles by the conditional probability of a family ranking a closer Kindergarten higher ($\Pr(r_i \succ r_j \mid d_i < d_j), i \neq j$) and ranking a Kindergarten with a sibling higher ($\Pr(r_i \succ r_j \mid s_i > s_j), i \neq j$). Where r_i is rank of Kindergarten i , d_i is distance to Kindergarten i and s_i is one when there is a sibling and zero otherwise. In the collected 2016 preference data, the $\Pr(r_i \succ r_j \mid d_i < d_j) = 0.81$ and the $\Pr(r_i \succ r_j \mid s_i > s_j) = 1.0, i \neq j$.

The main dimensions of the evaluation are the preference rank achieved in an allocation as well as the effect of the average distance from Kindergartens and the share of siblings in the same Kindergarten.

For statistical comparison, we generated twenty preference profiles of each of the parameter values. A total of 200 preference profiles were generated. For each policy that has a lottery, we ran twenty different randomised lotteries for each instance. As we saw in Table 6 the standard errors over the twenty lotteries are small. All the figures of the results show the smoothed⁵ results of the ten allocations over policies with a 95% confidence bound. For policies with lotteries, there are results with a single (S) and multiple (M) lotteries over Kindergartens.

Each year the number of available Kindergarten positions varies. However, on average about 20 places should be available in each Kindergarten each year, as one group of children leaves for school. Occasionally, there might be more or fewer places. In our experiments, we set the number of available places at 20 in each Kindergarten. However, this creates additional competition and the resulting matched ranks will be lower (see Ashlagi et al., 2013a,b) in these experiments than in the actual data in Table 6. Additionally, in our interpretations we implicitly assume the effect of the competition will be similar for all the policies. We discuss here only the Deferred-Acceptance based results⁶. In addition we removed policies DA5 and DA6 from the chart, as these matchings were usually almost the same as DA1.

Figures 1a and 1b demonstrate the average preferences obtained and the proportion of families getting their first choices for all policies. Policy DA7 is the most sensitive to changes in the preferences of families. When preferences are strictly based on distance with conditional probability of $\Pr(r_i \succ r_j \mid d_i < d_j) \rightarrow 1.0$, one of the highest average rank scores is produced, one similar to

⁵ smoothed with local polynomial regression

⁶ In Appendix D we also provide for comparison results based on Top Trading Cycles algorithms, as defined by Abdulkadiroğlu & Sönmez (2003).

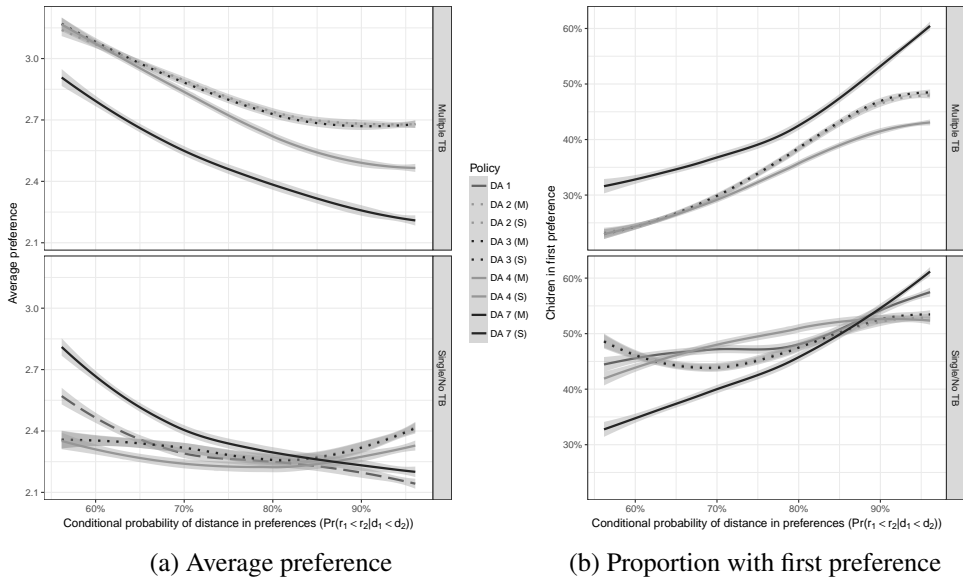


Figure 1: Conditional probability of distance

other policies such as DA1, DA5 and DA6. Surprisingly, when the preferences of families are close to random, with conditional probability of $\Pr(r_i \succ r_j | d_i < d_j) \rightarrow 0.5$, then DA7 (S) is the policy that has one of the lowest average ranks and the lowest number of families with a first preference. Policies that do worse are the ones using multiple lotteries, one per Kindergarten. In addition, the difference of having a single or multiple lotteries for Kindergartens is not very significant for DA7, most likely due to lower usage of tie-breaking in this policy compared to others with a lottery.

At face value, DA7 seems to be the most egalitarian policy as every family has the highest priority in at least one of the Kindergartens. However, it seems that families that do not prefer to be in the closest Kindergarten tend to be rejected more often from their preferred Kindergartens further away where they have a lower priority. As the matched rank drops more in DA7 than other policies, when $\Pr(r_i \succ r_j | d_i < d_j) \rightarrow 0.5$. Since the preferences and priorities are not aligned, the probability of the family being rejected in some round of the process is higher. The probability of being rejected at a certain point seems to be smaller for other policies.

In terms of average matched preference rank, the policies DA2 and DA3

are almost indistinguishable from each other, most likely because there are too few siblings in this data. Nevertheless, it is always better to use a single rather than multiple tie-breaking lotteries for both of these policies. The average preference achieved is always better with a single lottery and also there are more families with their first preference (Figure 1b). Policies with a single lottery, such as DA2 (S), DA3 (S) and DA4 (S) – with the exception of DA7 (S) – are significantly better for families in most situations. Only when $\Pr(r_i \succ r_j \mid d_i < d_j) > 0.9$, did policies DA1, DA5 and DA6, which use absolute distance, turn out to be better than the single lottery policies.

The policies DA1 and DA6 always produce exactly the same matching, DA5 is occasionally slightly different (for about 2-6 children), but the aggregate results are still very similar. This is most likely because the selected reserve of 20 % is close to the percentage of siblings in the data.

Interestingly, most policies, with the exception of DA7, are quite robust to changes in preferences. The same proportion of families almost always receive their first preferences, about 50% to 60% with DA2, DA3 and DA4 and 60% to 70% with DA1, DA5 and DA6. There is a slight increase in the average preference when preferences become determined by distance. With DA7, the proportion varies widely between 40% and 70%, and families fare better when preferences are aligned with distance.

Figure 2a shows the average distance between families and Kindergartens. The average distance is smaller for all policies when the preferences of families are determined more by distance. As might be expected, the smallest average distance is always with DA1 (including DA5 and DA6), as these policies are aimed to minimise distance. The average distance is the largest with DA2 and DA3, policies based on walk-zones, probably caused by the randomness in the priorities of Kindergartens. Furthermore, these policies have a slightly lower average distance with a single tie-breaking lottery, when preferences are correlated with distance. On the other hand, it is usually the case that if preferences are random, the multiple tie-breaking lotteries have a lower average distance than single lotteries. A small improvement in average distance in policies with lotteries is obtained by not using discretisation by walk-zones and instead having a higher priority for a fixed number of Kindergartens, as in DA4.

With random preferences, there is a trade-off between achieved preference and average distance in the results obtained by DA7 (M) and, DA2 (M) and DA3 (M), where DA4 (M) is at the middle point among these policies in this

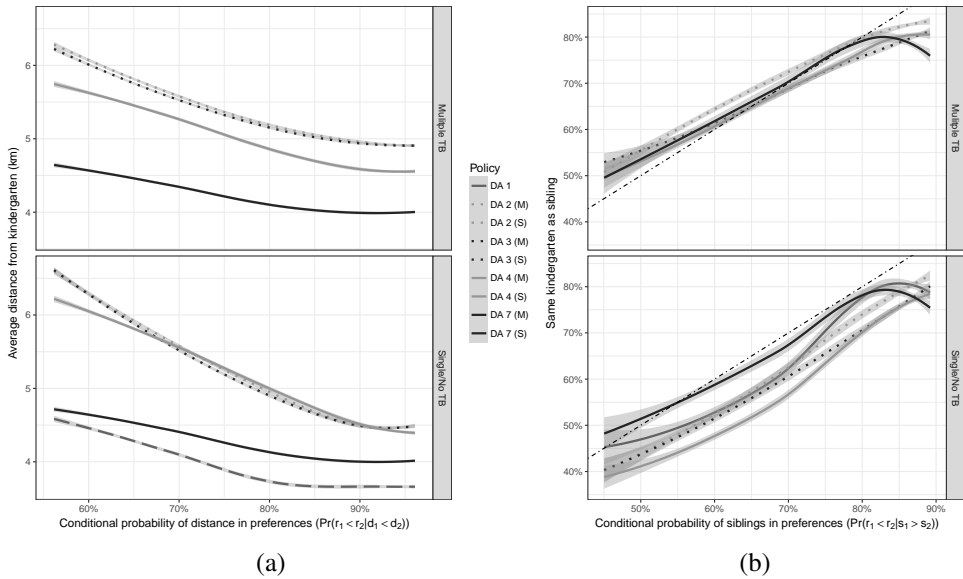


Figure 2: Average distance and conditional probability of distance in preferences

aspect. Policy DA7 always achieves the lowest average distance among the lottery policies, others produce better matched ranking. When preferences are more correlated with distance, DA7 performs better according to both average preference and distance.

Figure 2b depicts the probability of children being in the same Kindergarten as their siblings. When the preferences of families are random with respect to siblings, most policies place about 40% to 60% of children in the same Kindergarten as their siblings. When families prefer closer Kindergartens, then more siblings end up in the same place. This higher percentage is most likely due to siblings already being in a nearby Kindergarten. We have also added a 45 degree line, indicating that policies that are below this level have some children, who would prefer a Kindergarten with a sibling, result in children being assigned to a different Kindergarten. Multiple lottery policies seem to be better at placing children in the same Kindergarten with siblings.

In Figures 3a and 3b, the probability of a child being matched to the family’s first preference in at least one lottery is measured. This is a measure of fairness, or fair (equal) access to Kindergartens, which is similar to the

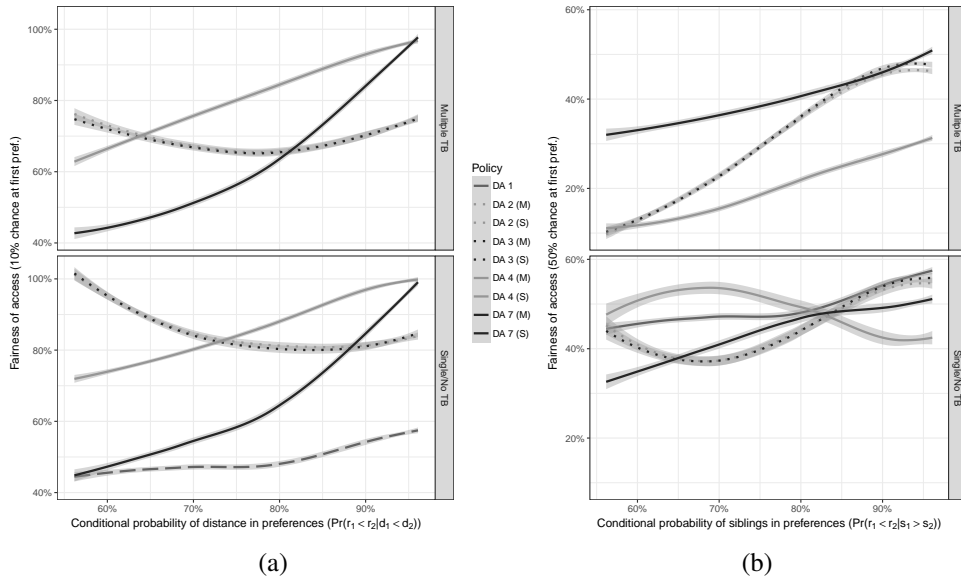


Figure 3: Fairness of access

measure of access to quality used by [Shi \(2015\)](#). We have plotted the fairness of access for policies DA1, DA5 and DA6, even though there is no sensible interpretation, since there are no lotteries. However, these policies are still useful for comparison.

With the lottery policies DA2, DA3 and DA4, with both single and multiple lotteries, about 60% to 95% of families have about a 10% chance of being granted a place in a Kindergarten that is their first preference. The DA4 (S) is the best performer when preferences are aligned with distance and DA2 (S) and DA3 (S) when preferences only have a Kindergarten effect. Policy DA7 (S) comes close to DA4 (S) only when preferences are almost perfectly aligned with distance.

However, when we make our fairness notion slightly stronger, i.e. when there has to be at least a 50% chance of a place in the family's first choice Kindergarten, the proportion of families achieving this drops to only about 40%. This is even lower than the case with deterministic policies like DA1. Therefore, it seems that with lotteries we can give some families a small 10%, chance of getting their first preference, but as a result, some families lose their first preferences. With a larger chance, 50%, there are more families losing

their first preference than those gaining.

In terms of trade-offs, the policy DA4 (S) is better on fairness and average matched preference, but worse on average matched distance. DA1 and similar policies do better on average matched rank and distance, however they fare worse on fairness, i.e. families living far away from all Kindergartens have a smaller chance of a preferred match. When preferences are not entirely determined by distance, then these two (DA1 and DA4) are the best options to choose from. However, with distance-based preferences, DA7 can prove to be an improvement. In the case of DA7, fairness is almost as good as with DA4, average distance was a significant improvement over DA4 and average allocated rank very close to DA1.

4. FURTHER ISSUES

We identify a dozen additional special features that should be further considered in the (re-)design of the mechanism in Harku. However, many of these features may pose significant challenges and require additional research. We describe these issues and give recommendations for possible adjustments in the allocation mechanism.

Children with special needs. In larger cities, there are schools for children with special needs, but in smaller municipalities these pupils are mixed with others. The standard practice is for Kindergartens to reserve places for children with special needs who require more attention and are thus considered to take up the space of three children. Usually, it is not known beforehand if there will be any such cases and the special needs may only become evident later. However, in most cases the extra places remain free and can be subsequently allocated to other children. Obviously, this has some effect on the fairness of the allocation.

A possible solution would be to have this data available before allocation and to take it into account in the allocation process. However, evaluating all of the applicants in advance could be very costly compared to the extra efforts needed for the reallocation process and the potential issues arising from the extended solution. It would be helpful if the parents of children that are likely to need special treatment were to register for evaluation. It should then be guaranteed that their chances of admission to their preferred Kindergartens would not be worsened, perhaps by giving priority for a number of places in each Kindergarten to such children.

Allocation in multiple rounds. Harku currently allocates students in multiple rounds, since two extra places could arise in each Kindergarten to which no student with special needs is admitted. The proportion of disadvantaged families is about 10% (PTVRR, 2015) and children with special needs make up about 3% (Paat et al., 2011). This question is similar to the question of the design of two-stage allocation mechanisms (Dur & Kesten, 2014) and also to the design of appeal processes (Dur & Kesten, 2015). The first option is to allocate the extra places exclusively among the unmatched children. This is a simple method with no reallocation of children, but it can be seen as unfair to families who are allocated seats in the main round and would prefer an extra place in a Kindergarten where they have higher priority than the unallocated children who get those extra places in the second round. The final solution could cause justified envy for the families. In addition, the parents might also act strategically in the main round, perhaps by not accepting an offer from the Kindergarten listed second, especially if they have information that they are first in the waiting list and the creation of extra places is very likely. Therefore, it appears reasonable to let everyone apply for the extra places, as is currently done in Harku. However, if the process is not centralised, then those who were assigned a place in the first round but now get a better match, would consequently create new available places. Even if this decentralised process could be continued until a stable solution was reached, this proposal-rejection chain would result in a stable matching that is the worst possible stable matching for the reallocated children, as proved by Blum et al. (1997). Therefore, this process would not be strategy-proof for the parents either. Hence, the only possible solution that is strategy-proof for the parents and avoids justified envy is a centralised second round, where parents can re-apply to all Kindergartens with the option of keeping their assignment if they wish to do so (technically this is achieved by putting the children already assigned to the Kindergarten at the top of the Kindergartens' rankings). Yet, this solution may affect a significant number of children, and in theory possibly all of them, which could result in high reallocation costs. These costs would be accepted by the parents, since they would always have the option of not changing their assignment, but could be seen as undesired by the local council and the Kindergartens.

Children with existing places. The parents of some children may request a transfer. This is especially relevant for children attending a class for 2-3-year olds who would like to go to a different Kindergarten for the 3-6-year period, since the classes for children aged 2-3 may not be available in the

Kindergartens preferred by the families. It is therefore a question of whether the reallocation of these children should be conducted as part of the yearly matching round. If so, these children should be guaranteed to get at least as good a seat in the reallocation, i.e. they should have the highest priority in their current Kindergarten. This question has been studied in the context of Danish daycare allocation (Kennes et al., 2014), and also for the reallocation of French teachers (Combe et al., 2015).

Overlapping admission processes. Some parents may be registered in more than one municipality, so they are able to apply for a place for their child in two systems, for example in Harku and in neighbouring Tallinn. This can lead to inefficiencies due to cancellations. Similar problems arise in some US cities where state schools and charter schools hold their admissions separately. Furthermore, the same phenomenon has also appeared in European college admission programmes, where an increasing number of students are applying for programmes in several countries disturbing the national matching schemes.

Outside options with subsidies. Somewhat related to the previous issue is the fact that private Kindergartens operate in Estonia, and some parents also consider the option of home schooling. However, if a municipality cannot provide enough Kindergarten places for its resident population, in some cases it may subsidise parents who choose an alternative option. In Harku, the local council financially supports parents who do not receive a place in a Kindergarten, but the council may withdraw their support if the parents do not accept a place that is offered. This conditional support can lead to strategic considerations, since some parents may find an alternative home or private option preferable to a local school, if and only if they receive the financial support, but this cannot be stated in the application. This special case can be modelled with the matching with contracts framework. A similar special feature is found in the Hungarian higher education matching scheme, where students can study on the same course under two different contracts, either free of charge or with a tuition fee. Furthermore, US cadets (Sönmez & Switzer, 2013) also face such a situation when they decide whether or not they are willing to take on some extra years of service in order to increase their chances of admission. The recommended solution is to let the parents list the option of not having a place in the Kindergarten but receiving financial support instead, when they give their applications to Kindergartens. Thus, all the listed options are considered preferable to the outside option with no financial support. In such a case, it is crucial that the parents-optimal stable solution is implemented

so as to make the parents reveal their true preferences for these outside options.

Lower quotas, opening of new groups. Sometimes Kindergartens are able to cancel groups or open new ones to fit with the applications. In particular, there are regulations determining the minimum number of children needed to start a new group. This feature is similar to the lower quotas used in the Hungarian higher education matching scheme (Biró et al., 2010a), where programmes may be cancelled if there is a lack of students. This is a natural requirement that makes the education service economical, but the theoretical model for college admissions with lower quotas is not always solvable. This means that a fair solution does not always exist and the problem of finding a fair solution is NP-hard. The problem becomes even more complicated if new groups can be created, since both the closures and the openings in a Kindergarten affect the number of students admitted elsewhere. However, clever heuristics and robust optimisation techniques, such as integer programming (Biró et al., 2014) can be used to tackle these generalised problems.

Homogeneous age groups and mixed groups. In Estonia, there are both homogeneous age groups and mixed groups. Having only same age groups can vary the number of groups opened in a Kindergarten, as a Kindergarten with five groups could open only one group every three years. This would be unsatisfactory for the local children in the years when no groups are opened. When mixed groups are created, the number of children admitted can be relatively stable if the available places are always filled. However, if there are some free places left in a year, then the age distribution of the children can be distorted.

Sharing places. In some Kindergartens, it is possible that some children only attend part of the week and the rest of the time is taken up by other children. This possibility again makes the underlying problem challenging to solve. Specifically, when there is a large number of part-time students then one might face the same problem as when allocating doctors and couples to hospitals, which is an NP-hard problem (McDermid & Manlove, 2010).

Historic dependence of preferences. In Harku, the applications of registered parents are listed on a public website. In Tallinn, the number of applications already submitted to the Kindergartens is also published. If the registration date is a criterion for priority and the parents can see the applications or the number of applications made before their turn, then this can affect their true as well as their submitted preferences. Potentially, if there are more applications than places, then parents will find it risky to apply. This can depend on the

birth date of the child, because if a child was born soon after 1 October, then the parents could have a good chance of obtaining a place everywhere, and so be more truthful. We did not find much evidence of significant changes in the preferences over time in the Harku data. However, in a similar study for Tallinn or other places where the registration date is important, attention should be paid to the potentially biased preferences caused by the published information about past applications.

Smooth transition to a new system. When designing the new mechanism, it may be important to consider how to engineer a smooth transition between the old and the final systems. This process is especially challenging in Harku, since the old priorities were based on registration date, and those parents who registered early may see it as unfair if this priority that they earned in the past is suddenly neglected. Therefore, in the 2016 transitory system, the priority of those who have already registered in the old regime is partly kept, as described above. Regarding the future years, how long these priorities should be kept, or whether they should be replaced with some age priority which is in correlation with the registration dates, is still to be debated.

The role of the heads of the Kindergartens. The heads of the Kindergartens were actively involved in the allocation system until 2015. The discussions among the heads and the personal communication with the parents were crucial in eliciting the true preferences of the parents and finding relatively good solutions through informal negotiations. In the centrally coordinated system, the head may fear losing their chance to influence the allocations, and the same could be true for the employees of the local municipality. It should be considered whether the heads of the Kindergartens could still have some power to adjust the priorities, or to make other decisions about their Kindergartens, for instance whether to open a new group or to create mixed groups.

The fairness of using proximity as a priority. Whether the use of proximity is fair may depend on the ease and/or cost of registering. Specially circumstances may vary, e.g., it is almost costless (as in Hungary); there may be some significant costs such as renting or having a flat in the area; or the family truly has to live there (for example in Barcelona, where somebody who is proved not to live at the stated address can lose their place). When it is easy to register at an address, then the parents may play a strategic game in which the first stage is to choose an address. When ownership and actual residency are required, and the priorities are important for the parents, this can affect the housing choices of the families, and influence house prices as well as the

socio-economic distribution of the population.

Restricting the choice of the parents. A simple restriction is to allow families to only apply to nearby Kindergartens. A more sophisticated method is to provide personalised choice menus, such as the system proposed in the Boston school choice mechanism (Shi, 2015). This would potentially provide parents with a choice of schools close to them where a child's siblings may have attended, with a limited number of further options. The advantage of this method over restricting the number of applications is that the mechanism remains strategy-proof, and the parents have a simpler task of ranking the available options. However, the disadvantage is the difficulty of estimating the preferences of the parents and therefore, there is a risk that some highly preferred Kindergartens could be missed out from some menus. In general, this type of restrictive policy can improve the overall quality of the allocation from the point of view of the municipality, perhaps by reducing the total travel distance. That was the main motivation in the Boston school choice redesign, as the bus costs had to be limited. However, the overall welfare of the children could be badly affected. We do not recommend this policy for Harku, due to the small size of the municipality, but it is suggested for consideration in larger cities, like Tallinn.

5. CONCLUSION AND DISCUSSION

We have reviewed the Kindergarten matching practices in one Estonian municipality, Harku. Until 2015, the collected preferences were unlikely to reflect the true preferences of the parents, since the data were out-of-date by the time of the allocation, the number of applications were limited and the allocation mechanism was not incentive-proof either. Therefore, the resulting allocation could create justified envy and it was also lacking transparency. In 2016, the municipality changed its allocation system mostly based on our recommendations.

In our study, we first listed well-known practices from matching mechanism design that present solutions to some of the problems and also provide policy tools for the local municipalities. These practices consist of:

- getting complete rather than limited preferences from families,
- using child-proposing stable matching for allocating places,

- defining clear policies for the local municipality based on a transparent priority system.

In assisting in the redesign of the allocation mechanism, it emerged that although the policy goals might be clear, the choice of exactly which implementation method to use can create significant differences in the results. In most cases, the goals of the local municipalities are to have siblings in the same Kindergarten and to provide a place in a Kindergarten close to home, in addition to the main consideration of providing a place in the most preferred Kindergartens of the families. We evaluated seven different policies for implementing the policy goals, first based on data from 2016, and then based on generated data. The 2016 transitory system that follows our main recommendations provides a child-optimal stable allocation under a rotational priority structure based on four factors, such as location, siblings, registration and birth dates. The limit on the number of applications was also removed, so the preferences of the families can be considered truthful. Our main findings regarding the seven policies evaluated on the real data and in the computational experiments are summarised below.

The simplest policy is to give higher priority to children with siblings and to families living nearby, which is policy DA1. This was also demonstrated to be one of the most effective policies. The resulting allocation had, on average, matched a lot of families with their most preferred Kindergarten, while also having one of the smallest average distances. This remained true when the preferences of families were agnostic about distance.

Policy DA1 might occasionally seem unfair, as small differences in distance might affect whether families are placed in their first preference or a lower one. Policies DA2, DA3 and DA4 group Kindergartens by distance within equal priority classes, DA2 and DA3 by defining a walk-zone and DA4 by having high priority in the three closest Kindergartens. Families in the walk-zone are treated equally and priorities are defined by lottery. It appears that the multiple tie-breaking rule might create a more egalitarian access to Kindergartens, however it is not without its cost. The average number of children who are placed in their most preferred Kindergarten is usually significantly lower and the average distance is greater. However, with a single tie-breaker over Kindergartens, families are on average allocated to their more preferred Kindergarten, even when compared to deterministic policies like DA1. Nevertheless, an allocation based on randomness might prove hard to justify to families. If having more egalitarian access is important, policy DA4

with a single tie-breaker would be the best of the three. The level of fair access is the same, satisfaction with average preferences is the best, and distance is the lowest.

Siblings always being given higher priority might prove another source of seemingly unfair treatment. If a family already has a child in a particular Kindergarten, they are almost guaranteed to get a place in the same Kindergarten for a sibling, even when there is another family living closer than them. We considered two policies, DA5 and DA6, which limit the number of places in a Kindergarten that consider having a sibling a priority at up to 20%. Even though the number of places reserved for siblings was low, most families still received a place in that Kindergarten if they preferred it. There is almost no difference from policy DA1 on any measure, nor between DA5 and DA6, although theoretically DA6 should provide more opportunity to nearby families, and DA5 to children with siblings.

A clear oddity is policy DA7, which was initially designed to deliver more equal access to Kindergartens for families who live far away from all Kindergartens. While policy DA1 would give such families low priority everywhere, DA7 would give them the highest priority in their closest Kindergarten. When most families have a high preference for nearby Kindergartens and for those where their siblings are, DA7 results in one of the best policy designs in all aspects. DA7 gives many families their first preference, it has the shortest average distance and one of the best results for equality of access. However, the result is radically different when family preferences are mostly idiosyncratic and are almost independent from distance. In this case, DA7 is the worst policy of all for families. On average less than 40% of children get matched to their first preferences, but the average distance is the one of lowest. Thus the lesson from policy DA7 seems to be that the policy designer needs to predict the preferences of the society fairly accurately to select a good trade-off. When preferences and priorities are aligned, both of the main goals can be met. A downside of this policy is that it is vulnerable when preferences and priorities are misaligned, and then the price paid is significant in terms of efficiency and fairness. If a local municipality aims to minimise the distances between homes and Kindergartens, then DA1 is the best option. The latter objective recently turned out to be crucial in Boston, where the local authority became concerned about the busing costs (Shi, 2015).

Finally, there remain several unsolved issues that we have not tried to address in the redesign. A dozen issues were listed along with a discussion

about possible solutions. For example, it would be reasonable to coordinate the allocation between neighbouring municipalities, but cooperation is usually hard to achieve. Similarly, it would be best to know about children with special needs before the allocation, but this is often not feasible.

A potential way to manage the shortage of Kindergarten places is to provide monetary incentives for parents to stay at home with their children or to seek a place in private childcare. The question of how to set this monetary compensation in an optimal manner is also interesting in terms of future research. Here, optimality could mean minimising the total cost of providing childcare services in the municipality.

There remain a few interesting aspects related to designing a more flexible mechanism which might improve the allocation outcome for families. Making decisions on the size and the age composition of the groups in Kindergartens and determining this in an optimal way based on the application data could give an additional boost to the number of families receiving a place at their most preferred Kindergarten. Some of this research has been done in terms of lower quotas for opening groups (Biró et al., 2010a).

A. GENERATING COUNTER-FACTUAL PREFERENCES

We use the 2016 data for counter-factual policy evaluation. To generate the counter-factual preferences only we use the distance between homes and Kindergartens and sibling status in a Kindergarten. The collected preference data is used to understand which features to use in the ranking function, the functional form of the utility function and the fixed effects of Kindergartens.

For each family and Kindergarten we know the geographical location from address lookup from google maps⁷ and Estonian Land Board (Maa-amet⁸) and distance calculations taken from Google maps distance⁹. We have a rich dataset for distance, as for each family-Kindergarten pair we know the driving and walking distances in kilometres and minutes. We also have the direct distance between the two points calculated with the haversine formula. The features are described in Table 7.

We fit a multinomial rank-ordered logit model (Croissant, 2011), which is similar to the model used by Shi (2015). The model assumes that families

⁷ <https://developers.google.com/maps/documentation/geocoding/intro>

⁸ <http://inaaddress.maaamet.ee/geocoder/bulk>

⁹ <https://developers.google.com/maps/documentation/distance-matrix/intro>

Table 7: Family's Kindergarten features

Feature	Description
preference_rank	Families rank of the Kindergarten, between 1-7
walking_distance_sec	walking time between family's home and Kindergarten, based on The Google Maps Distance Matrix API (2015)
walking_distance_m	walking distance between family's home and Kindergarten, based on The Google Maps Distance Matrix API (2015)
driving_distance_sec	driving time between family's home and Kindergarten, based on The Google Maps Distance Matrix API (2015)
driving_distance_m	driving distance between family's home and Kindergarten, based on The Google Maps Distance Matrix API (2015)
haversine_distance_m	direct distance between family's home and Kindergarten
walking_distance_rank	Kindergarten rank by walking distance
driving_distance_rank	Kindergarten rank by driving distance
haversine_distance_rank	Kindergarten rank by haversine distance
sibling	1 if Kindergarten has a sibling already attending, 0 otherwise
log_walking_distance_sec	$\log(walking_distance_sec)$
sqrt_walking_distance_sec	$\sqrt{walking_distance_sec}$
log_walking_distance_m	$\log(walking_distance_m)$
sqrt_walking_distance_m	$\sqrt{walking_distance_m}$
log_driving_distance_sec	$\log(driving_distance_sec)$
sqrt_driving_distance_sec	$\sqrt{driving_distance_sec}$
log_driving_distance_m	$\log(driving_distance_m)$
sqrt_driving_distance_m	$\sqrt{driving_distance_m}$
log_haversine_distance_m	$\log(haversine_distance_m)$
sqrt_haversine_distance_m	$\sqrt{haversine_distance_m}$

have an utility function of the form,

$$u_{ij} = \alpha_j + \sum_k \beta_k \cdot x_{kij} + \varepsilon_{ij} \tag{1}$$

where α_j are fixed effect of Kindergartens, β_k is the coefficient for feature k and ε_{ij} is the family’s personal unexplained preference. We further use the utilities to find a probability if a ranking. In a ranked-order logit model the probability of a ranking is a multiple of a Kindergarten begin is a particular position, which in our case is $Pr(\text{ranking}1, 2, \dots, 7) = Pr(\text{ranking} = 1) \cdot Pr(\text{ranking} = 2) \cdot \dots \cdot Pr(\text{ranking} = 7)$. The probability of family i ranking Kindergarten j at some position are,

$$\left\{ \begin{array}{l} Pr_{ij}(\text{ranking} = 1) = \frac{e^{u_{ij}}}{\sum_{r=1}^7 e^{u_{ir}}} \\ Pr_{ij}(\text{ranking} = 2) = \frac{e^{u_{ij}}}{\sum_{r=2}^7 e^{u_{ir}}} \\ \dots \\ Pr_{ij}(\text{ranking} = 6) = \frac{e^{u_{ij}}}{\sum_{r=6}^7 e^{u_{ir}}} \end{array} \right. \tag{2}$$

First our aim is to select one of the distance metrics from Table 7 to include in the utility model (1). For this we do 100 bootstrap runs with each metric. In Figure 4 we plot the resulting log-likelihood with its standard error. We see that the $\sqrt{\text{driving_distance_sec}}$ provides the best prediction on average. We also see that including the sibling status would improve the prediction accuracy, however the statistical significance of the coefficient is low (Table 8) in any combination of features. So we select the model (1) from Table 8 as our final model.

For policy comparison we generate the ranking over all Kindergartens. We do not model the cut-off levels for outside options, when the family would rather keep the child at home. We assume they would always rather have a place in any of Harku’s Kindergartens.

To obtain a full ranking of Kindergartens we use the probabilities from (2). For counter-factual preferences we vary the coefficient for distance. The parameter values are in (3). For each combination of parameters we generate several (7) different preference profiles and evaluate the policies on the average over all the preference profiles.

$$\beta_1 \in \{0.0, 0.05, 0.1, 0.23, 0.25, 0.5, 1, 2, 4, 10\} \tag{3}$$




Figure 4: Predictive features

Table 8: Rank-ordered logit coefficients


	preference rank		
	(1)	(2)	(3)
α_B	-0.690*** (0.150)	-0.685*** (0.150)	-0.560*** (0.143)
α_C	-0.565*** (0.173)	-0.540*** (0.176)	0.471*** (0.145)
α_D	0.157 (0.176)	0.185 (0.182)	1.479*** (0.154)
α_E	0.476*** (0.156)	0.500*** (0.159)	1.187*** (0.146)
α_F	0.275 (0.176)	0.351* (0.181)	1.608*** (0.153)
α_G	-1.769*** (0.193)	-1.789*** (0.195)	-1.580*** (0.179)
β_1 $\sqrt{\text{driving_distance_sec}}$	-0.229*** (0.015)	-0.220*** (0.015)	
β_2 sibling		20.750 (2,676.852)	20.812 (1,651.629)
Observations	906	906	906
Log Likelihood	-882.862	-840.256	-958.955
Note:	*p<0.1; **p<0.05; ***p<0.01		

To better interpret the results we look at the results by conditional probabilities of a parameter set. We look at two conditional effects: (a) probability of ranking Kindergarten higher given it is closer and (b) probability of ranking a Kindergarten higher given a Kindergarten has a sibling. Formally the conditional probability are defined in (4) and (5).

$$\Pr(r_1 < r_2 \mid d_1 < d_2) = \frac{\Pr(d_1 < d_2, r_1 < r_2)}{\Pr(d_1 < d_2)} \quad (4)$$

$$\Pr(r_1 < r_2 \mid s_1 > s_2) = \frac{\Pr(s_1 > s_2, r_1 < r_2)}{\Pr(s_1 > s_2)} \quad (5)$$

The mean conditional probability with fitted regression parameter, $\beta = 0.25$, is $\Pr(r_1 < r_2 \mid d_1 < d_2) \approx 0.79 \pm 0.02$ ¹⁰. This is similar to what we observe in the 2016 data, where $\Pr(r_i > r_j \mid d_i < d_j) = 0.81, i \neq j$. Figure 5a shows the relationship between the logistic parameters and the conditional probabilities.



(a) Conditional probability on distance (b) Conditional probability on siblings

Figure 5: Coefficients and conditional probabilities

¹⁰ 1.96 standard deviations, 95% probability

B. ALLOCATED PREFERENCES

Table 9: Year 2016 allocated preference comparison DA

Policy	1st	2nd	3rd	4th	5th	6th	7th
Harku	115	14	6	3	4	8	2
IA 1	122	8	2	5	8	3	4
DA 1	110	17	3	6	6	9	1
DA 2 (M) ^a	98.75 (0.61)	19.95 (0.61)	9.65 (0.49)	11.05 (0.78)	7.20 (0.43)	4.85 (0.36)	1.22 (0.15)
DA 2 (S)	108.05 (0.61)	19.90 (0.56)	4.65 (0.41)	4.95 (0.39)	7.70 (0.37)	5.85 (0.36)	1.20 (0.11)
DA 3 (M)	98.30 (0.79)	21.95 (1.09)	8.85 (0.6)	9.9 (0.49)	8.35 (0.3)	3.75 (0.24)	1.5 (0.29)
DA 3 (S)	107.75 (0.38)	20.65 (0.38)	4.79 (0.50)	4.95 (0.29)	6.60 (0.39)	6.2 (0.3)	1.86 (0.25)
DA 4 (M)	89.25 (1.06)	27.2 (0.84)	13.1 (0.88)	9.85 (0.76)	7.7 (0.53)	4.35 (0.43)	1.38 (0.26)
DA 4 (S)	104.85 (0.70)	19.15 (0.70)	7.05 (0.63)	8.15 (0.6)	8.05 (0.46)	4.10 (0.28)	1.63 (0.26)
DA 5	110	17	3	6	6	9	1
DA 6	110	17	3	6	6	9	1
DA 7 (M)	107.6 (0.47)	15.85 (0.52)	5.7 (0.36)	8.9 (0.42)	6.95 (0.33)	5.9 (0.32)	1.47 (0.17)
DA 7 (S)	107.75 (0.47)	16.90 (0.44)	5.20 (0.26)	7.40 (0.36)	8.40 (0.37)	5.15 (0.33)	1.33 (0.14)

^a For policies with lotteries in parenthesis are the standard errors

Table 10: Year 2016 allocated preference comparison TTC

Policy	1st	2nd	3rd	4th	5th	6th	7th
Harku	115	14	6	3	4	8	2
IA 1	122	8	2	5	8	3	4
TTC 1	112	16	1	5	9	8	1
TTC 2 (M)	110.25 (0.52)	17.85 (0.49)	4.00 (0.32)	6.55 (0.41)	7.75 (0.45)	4.70 (0.52)	1.38 (0.21)
TTC 2 (S)	110.10 (0.45)	17.50 (0.56)	3.85 (0.33)	6.40 (0.56)	7.95 (0.36)	4.80 (0.35)	2.00 (0.23)
TTC 3 (M)	109.50 (0.54)	18.95 (0.72)	4.55 (0.34)	5.75 (0.37)	7.00 (0.40)	5.20 (0.28)	1.50 (0.25)
TTC 3 (S)	110.55 (0.61)	17.05 (0.51)	4.65 (0.36)	6.20 (0.52)	7.05 (0.39)	5.75 (0.35)	1.25 (0.13)
TTC 4 (M)	109.25 (0.56)	18.65 (0.74)	3.78 (0.33)	7.40 (0.62)	8.80 (0.57)	3.60 (0.36)	1.80 (0.29)
TTC 4 (S)	109.35 (0.38)	18.75 (0.48)	4.47 (0.42)	7.30 (0.37)	7.45 (0.45)	4.05 (0.25)	1.31 (0.17)
TTC 7 (M)	109.40 (0.39)	15.45 (0.37)	3.75 (0.28)	7.45 (0.39)	8.55 (0.36)	5.70 (0.25)	1.89 (0.23)
TTC 7 (S)	109.25 (0.41)	16.00 (0.45)	3.63 (0.34)	7.25 (0.37)	9.15 (0.33)	5.60 (0.29)	1.53 (0.12)

C. MAP OF THE MUNICIPALITY




Figure 6: Locations of children and Kindergartens (with walk-zones) in 2016

D. RESULTS WITH TOP TRADING CYCLES (TTC)

Table 11: Year 2016 comparison of policies using reported preferences

Policy	Mean preference	First	Mean distance (km)	With siblings	Frac. children with JE ^a	BP per child with JE
DA 1	1.76	110	4.26	100 %		
DA 2 (M) ^b	1.85 (0.01)	98.75 (0.61)	4.59 (0.02)	100 % (0.0 %)		
DA 2 (S)	1.72 (0.01)	108.05 (0.61)	4.44 (0.01)	100 % (0.0 %)		
DA 3 (M)	1.83 (0.01)	98.30 (0.79)	4.51 (0.02)	95 % (0.25 %)		
DA 3 (S)	1.72 (0.01)	107.75 (0.38)	4.45 (0.02)	96 % (0.3 %)		
DA 4 (M)	1.91 (0.01)	89.25 (1.06)	4.53 (0.02)	100 % (0.0 %)		
DA 4 (S)	1.75 (0.01)	104.85 (0.7)	4.49 (0.01)	100 % (0.0 %)		
DA 7 (M)	1.78 (0.01)	107.60 (0.47)	4.30 (0.01)	100 % (0.0 %)		
DA 7 (S)	1.76 (0.01)	107.75 (0.47)	4.31 (0.01)	100 % (0.0 %)		

^a JE - Justified Envy, BP - Blocking Pairs

^b For policies with lotteries, (M) indicates multiple tie-breaking lotteries and (S) single. The standard errors over lotteries are in parentheses.

Table 12: Year 2016 comparison of policies using reported preferences

Policy	Mean preference	First	Mean distance (km)	With siblings	Frac. children with JE^a	BP per child with JE
TTC 1	1.76	112	4.39	100 %	11 %	1.06
TTC 2 (S)	1.71 (0.01)	110.1 (0.45)	4.49 (0.02)	100 % (0.0 %)	9 % (1.28 %)	1.08 (0.03)
TTC 2 (M)	1.69 (0.01)	110.25 (0.52)	4.51 (0.03)	100 % (0.0 %)	21 % (0.82 %)	2.16 (0.03)
TTC 3 (S)	1.7 (0.01)	110.55 (0.61)	4.47 (0.02)	96 % (0.37 %)	14 % (1.24 %)	1.27 (0.06)
TTC 3 (M)	1.69 (0.01)	109.5 (0.54)	4.46 (0.02)	95 % (0.33 %)	22 % (0.71 %)	2.16 (0.04)
TTC 4 (S)	1.69 (0.01)	109.35 (0.38)	4.6 (0.02)	100 % (0.0 %)	20 % (1.08 %)	1.86 (0.05)
TTC 4 (M)	1.7 (0.01)	109.25 (0.56)	4.58 (0.02)	100 % (0.0 %)	26 % (0.52 %)	2.26 (0.04)
TTC 7 (S)	1.77 (0.01)	109.25 (0.41)	4.48 (0.01)	100 % (0.0 %)	22 % (0.82 %)	1.70 (0.06)
TTC 7 (M)	1.78 (0.01)	109.4 (0.39)	4.51 (0.02)	100 % (0.0 %)	23 % (0.83 %)	1.90 (0.07)

^a JE - Justified Envy, BP - Blocking Pairs

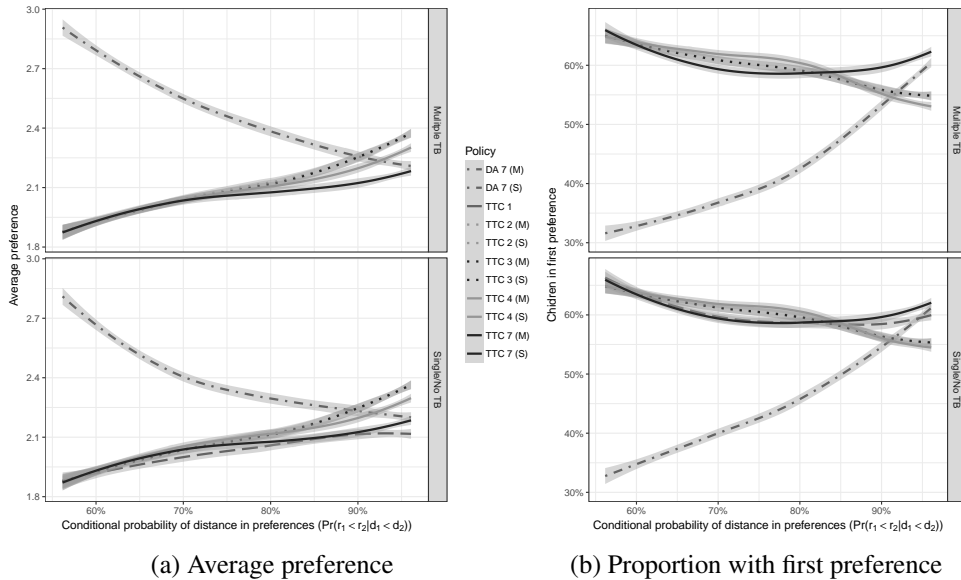


Figure 7: Conditional probability of distance (TTC)

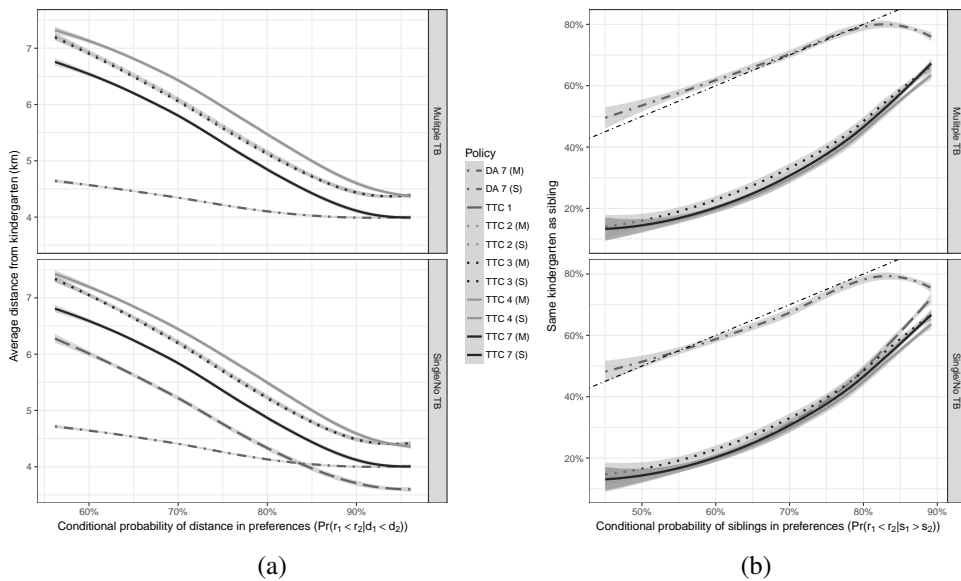


Figure 8: Average distance and conditional probability of distance in preferences (TTC)

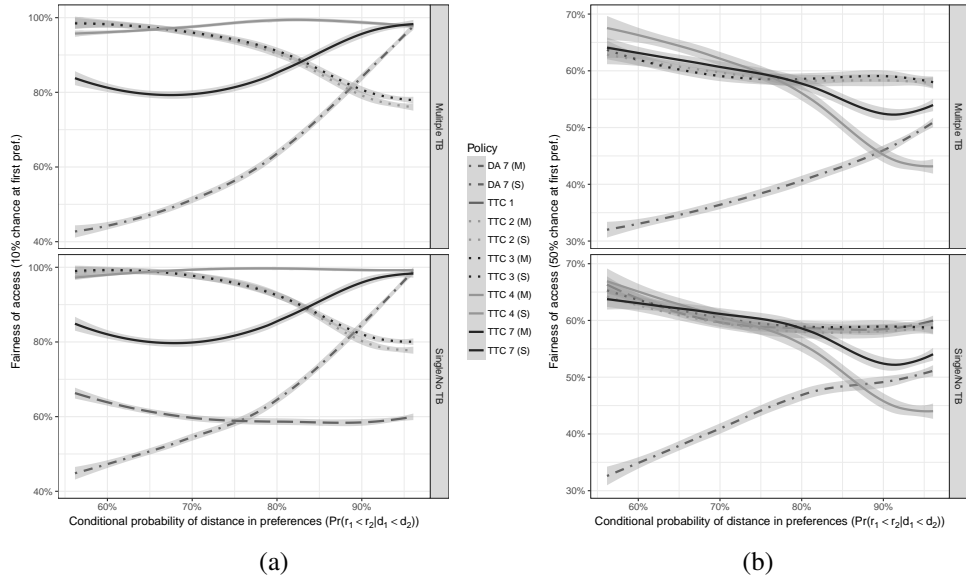


Figure 9: Fairness of access

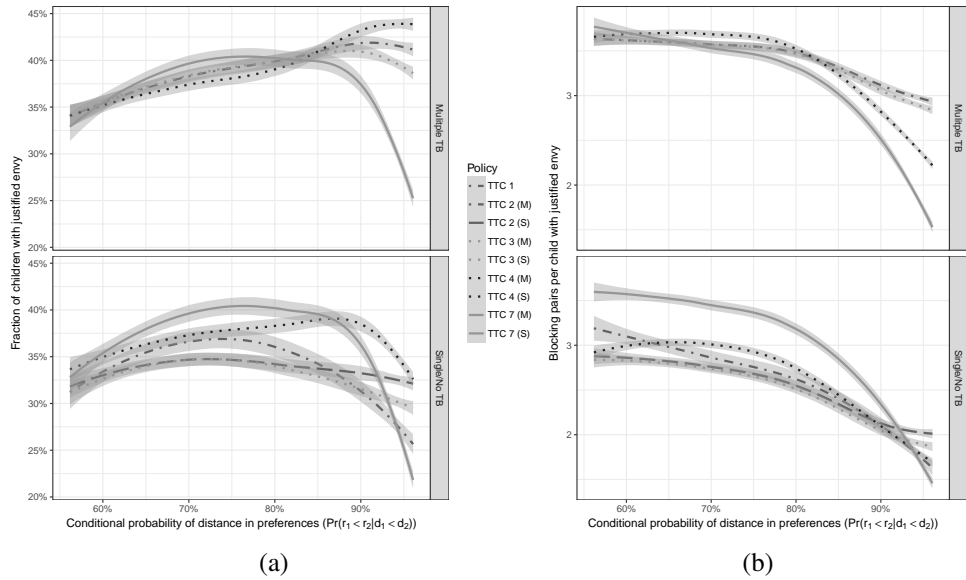


Figure 10: Justified envy and blocking pairs (TTC)

References

- Abdulkadiroğlu, A., & Sönmez, T. (2003). School choice: A mechanism design approach. *American Economic Review*, *93*(3), 729–747.
- Abdulkadiroğlu, A., Pathak, P. A., & Roth, A. E. (2005a). The New York city high school match. *American Economic Review*, *95*(2), 364–367.
- Abdulkadiroğlu, A., Pathak, P. A., Roth, A. E., & Sönmez, T. (2005b). The Boston public school match. *American Economic Review*, *95*(2), 368–371.
- Andersson, T. (2017). Matching practices for elementary schools - Sweden. *MiP Country Profile 24*. Retrieved from http://www.matching-in-practice.eu/wp-content/uploads/2017/01/MiP_Profile_No.24.pdf
- Ashlagi, I., Kanoria, Y., & Leshno, J. D. (2013a). Unbalanced random matching markets. In *Proceedings of the fourteenth acm conference on electronic commerce* (pp. 27–28). New York: ACM Press.
- Ashlagi, I., Kanoria, Y., & Leshno, J. D. (2013b). Unbalanced random matching markets: the stark effect of competition. *Unpublished working paper*. <http://web.mit.edu/iashlagi/www/papers/UnbalancedMatchingAKL.pdf> (Accessed 14.09.2016).
- Ashlagi, I., & Nikzad, A. (2015). What matters in tie-breaking rules? how competition guides design. *Unpublished working paper*.
- Biró, P., Fleiner, T., Irving, R. W., & Manlove, D. F. (2010a). The college admissions problem with lower and common quotas. *Theoretical Computer Science*, *411*(34–36), 3136–3153.
- Biró, P., Manlove, D. F., & McBride, I. (2014). The hospitals / residents problem with couples: Complexity and integer programming models. In J. Gudmundsson & J. Katajainen (Eds.), *Experimental algorithms* (pp. 10–21). Springer International Publishing.
- Biró, P., Manlove, D. F., & Mittal, S. (2010b). Size versus stability in the marriage problem. *Theoretical Computer Science*, *411*(16–18), 1828–1841.
- Blum, Y., Roth, A. E., & Rothblum, U. G. (1997). Vacancy chains and equilibration in senior-level labor markets. *Journal of Economic Theory*, *76*(2), 362–411.
- Calsamiglia, C., & Güell, M. (2014). The illusion of school choice: Empirical evidence from Barcelona. *Federal Reserve Bank of Minneapolis Working Paper*, 712. <https://www.mpls.frb.org/research/wp/wp712.pdf> (Accessed 22.12.2016).
- Calsamiglia, C., Haeringer, G., & Klijn, F. (2010). Constrained school choice: An experimental study. *American Economic Review*, *100*(4), 1860–1874.
- Chen, L. (2012). *University admission practices – Ireland*. <http://www.matching-in-practice.eu/higher-education-in-ireland/>

- (Accessed 18.01.2016).
- Chen, S.-H., Chang, C.-L., & Du, Y.-R. (2012). Agent-based economic models and econometrics. *The Knowledge Engineering Review*, *27*(2), 187–219.
- Combe, J., Tercieux, O., & Terrier, C. (2015). The design of teacher assignment: Theory and evidence. http://www.lse.ac.uk/economics/currentStudents/researchStudents/EDPjamboree/Terrier_PSE_EDPpaper.pdf (Accessed 04.08.2016).
- Croissant, Y. (2011). Estimation of multinomial logit models in R: The mlogit packages. *CRAN.R-project.org*. <https://cran.r-project.org/web/packages/mlogit/vignettes/mlogit.pdf> (Accessed 12.12.2006).
- de Haan, M., Gautier, P. A., Oosterbeek, H., & van der Klaauw, B. (2015). The performance of school assignment mechanisms in practice. *IZA Discussion Papers*, 9118.
- Dur, U. M., & Kesten, O. (2014). Sequential versus simultaneous assignment systems and two applications. <http://www.matching-in-practice.eu/wp-content/uploads/2014/01/Dur-Kesten.pdf> (Accessed 04.08.2016).
- Dur, U. M., & Kesten, O. (2015). The appeals process in NYC high school match. *Unpublished working paper*.
- Dur, U. M., Kominers, S. D., Pathak, P. A., & Sönmez, T. (2013). The demise of walk zones in Boston: Priorities vs. precedence in school choice. *NBER Working Paper Series*, 18981.
- Ergin, H., & Sönmez, T. (2006). Games of school choice under Boston mechanism. *Journal of Public Economics*, *90*, 215–237.
- Gale, D., & Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, *69*(1), 9–15.
- The google maps distance matrix api. (2015). <https://maps.googleapis.com/maps/api/distancematrix/> (Accessed 31.12.2015).
- Haeringer, G., & Klijn, F. (2009). Constrained school choice. *Journal of Economic Theory*, *144*(5), 1921–1947.
- Hafalir, I. E., Yenmez, M. B., & Yildirim, M. A. (2013). Effective affirmative action in school choice. *Theoretical Economics*, *8*(2), 325–363.
- Kennes, J., Monte, D., & Tumennasan, N. (2014). The day care assignment: A dynamic matching problem. *American Economic Journal: Microeconomics*, *6*(4), 362–406.
- Knuth, D. E. (1997). *Stable marriage and its relation to other combinatorial problems*. Providence: American Mathematical Society.
- MatchinginPractice. (2016). *Matching practices in Europe*. <http://www.matching-in-practice.eu/> (Accessed 27.12.2016).
- McDermid, E. J., & Manlove, D. F. (2010). Keeping partners together: algorithmic

- results for the hospitals/residents problem with couples. Journal of Combinatorial Optimization, 19(3), 279–303.
- Paat, G., Kaarna, R., & Aaviksoo, A. (2011). Meditsiiniliste erivajadustega laste üldhariduse korralduse analüüs (Tech. Rep.). Tallinn: Praxis Centre for Policy Studies.
- Pathak, P. A., & Sethuraman, J. (2011). Lotteries in student assignment: An equivalence result. Theoretical Economics, 6(1), 1–17.
- Pathak, P. A., & Shi, P. (2013). Simulating alternative school choice options in Boston (Tech. Rep.). MIT School Effectiveness and Inequality Initiative.
- Pathak, P. A., & Sönmez, T. (2013). School admissions reform in Chicago and England: Comparing mechanisms by their vulnerability to manipulation. American Economic Review, 103(1), 80–106.
- Preschool child care institutions act. (2014). Riigikogu RT I, 13.03.2014, 4. <https://www.riigiteataja.ee/en/eli/517062014005> (Accessed 04.08.2016).
- PTVRR. (2015). Peretoetuste, teenuste ja vanemapuhkuste roheline raamat. http://www.sm.ee/sites/default/files/content-editors/Lapsed_ja_pered/Perehuvitised/peretoetuste_teenuste_ja_vanemapuhkuste_roheline_raamat_2015.pdf.
- Rawls, J. (1971). A Theory of Justice. Cambridge: Harvard University Press.
- Roth, A. E. (2008). Deferred acceptance algorithms: history, theory, practice, and open questions. International Journal of Game Theory, 36(3-4), 537–569.
- Saraceno, C. (2011). Family policies. concepts , goals and instruments. Carlo Alberto Notebooks, 230, 20.
- Shi, P. (2015). Guiding school-choice reform through novel applications of operations research. Interfaces, 45(2), 117–132.
- Sönmez, T., & Switzer, T. B. (2013). Matching with (branch-of-choice) contracts at the United States military academy. Econometrica, 81(2), 451–488.
- West, A., Hind, A., & Pennell, H. (2004). School admissions and 'selection' in comprehensive schools: policy and practice. Oxford Review of Education, 30(3), 347–369.