

THESIS ON INFORMATICS AND SYSTEM ENGINEERING C125

Agent-Based Computational Experiments in Two-Sided Matching Markets

ANDRE VESKI



TALLINN UNIVERSITY OF TECHNOLOGY

School of Information Technologies

Department of Software Science

This dissertation was accepted for the defence of the degree of Doctor of Philosophy in Computer and Systems Engineering on 08.03.2017

Supervisors: Prof. Kaire Põder
Estonian Business School

Prof. Emer. Leo Võhandu
Department of Software Science
Tallinn University of Technology

Opponents: Alan Kirman, Ph.D., Aix-Marseille University, France
David Manlove, Ph.D., University of Glasgow, United Kingdom

Defence: 06.09.2017

Declaration:

I hereby declare that this doctoral thesis, my original investigation and achievement, submitted for the doctoral degree at Tallinn University of Technology, has not been submitted for any academic degree.

/Andre Veski/



Copyright: Andre Veski, 2017

ISSN 1406-4731

ISBN 978-9949-83-130-2 (publication)

ISBN 978-9949-83-131-9 (PDF)

INFORMAATIKA JA SÜSTEEMITEHNIKA C125

Agendipõhised arvutuslikud eksperimendid kahepoolsetel sobitusturgudel

ANDRE VESKI

Contents

List of Publications	9
Introduction	11
Motivation	15
The claims and contributions	16
Outline of the dissertation	17
1 Mechanism Design Background	19
1.1 Game theoretic foundations	19
1.2 Auctions	23
1.2.1 The first price auction	24
1.2.2 The second-price auction	25
1.2.3 The VCG mechanism	25
1.2.4 Optimal auctions	27
1.3 Fair division	29
1.3.1 Single divisible item	31
1.3.2 Multiple divisible items	32
1.3.3 Multiple indivisible items	35
1.4 Matching markets	35
1.4.1 Two-sided matching - stable marriage	39
1.4.2 Two-sided matching - school choice	41
1.4.3 Algorithms for on-line and decentralised matching	43
1.4.4 Other matching problems	44
1.5 Limits of the axiomatic approach	45
1.5.1 Combinatorial auctions and limits of theory	45
1.5.2 Double sided auctions and complex systems	51
2 Zero-Intelligence in Decentralised Matching	61
2.1 Introduction	61
2.2 Behaviour models	65
2.3 Preference generation	72
2.4 Computational experiments and convergence	75
2.5 Stability of a matching	75

2.6	Unassigned agents in a thick market	79
2.6.1	Analysis of convergence conditions	79
2.6.2	Structured preferences	84
2.7	Unassigned agents in a thin market – the Beveridge curve .	85
2.7.1	Beveridge curve and the movement along the curve .	85
2.7.2	Shifts in the Beveridge curve	86
2.7.3	Effect of a re-matching friction	88
2.8	Price of invisibility	90
2.8.1	Median rank in a thick market	90
2.8.2	Median rank in a thin market	93
2.8.3	Re-matching friction and median rank	96
2.9	Conclusion and discussion	97
3	Strategies in Tallinn School Choice Mechanism	101
3.1	Introduction	101
3.2	Background: Tallinn school market	103
3.2.1	Tallinn mechanism	104
3.2.2	Example of deciding what to report	105
3.3	Model	107
3.3.1	Environment	107
3.3.2	Preferences	108
3.3.3	Utility function	109
3.3.4	Genetic algorithms	110
3.4	Results	112
3.4.1	Expected utility maximising strategies	112
3.4.2	Social welfare	115
3.5	Conclusion and discussion	117
4	Policy Design for Kindergarten Allocation	121
4.1	Introduction	121
4.2	Matching mechanism design	125
4.2.1	Matching practices in Harku	126
4.2.2	Building a mechanism for kindergarten seat allocation	127
4.2.3	Particularities of the 2016 system	128
4.3	Policy design	130
4.3.1	Efficiency and fairness	130
4.3.2	Operationalisation of policy designs	131
4.3.3	Data and initial policy design comparison	135
4.3.4	Generating counter-factual preferences	137
4.3.5	Policy sensitivity to preferences	140
4.4	Further issues	147
4.5	Conclusion and discussion	148

5 Conclusions and Future Work	151
5.1 Discussion of the research methods	151
5.2 Answers to the research questions	153
5.3 Contributions	155
5.4 Further research	155
References	157
A Publication 1	181
B Publication 2	197
C Publication 3	211
D Publication 4	239
E Publication 5	261
Acknowledgements	303
Abstract	305
Kokkuvõte	307
Curriculum Vitae	309
Elulookirjeldus	313

List of Publications

1. Veski, A. and Vöhandu, L. (2011). Two Player Fair Division Problem with Uncertainty. In Barzdins, J. and Kirikova, M., editors, *Frontiers in Artificial Intelligence and Applications*, pages 394–407. IOS Press, Amsterdam
2. Veski, A. (2014). Price of Invisibility: Statistics of centralised and decentralised matching markets. In MacKerrow, E., Terano, T., Squazzon, F., and Sichman, J. S., editors, *Proceedings of the 5th. World Congress on Social Simulation*, pages 18–29, Sao Paulo
3. Veski, A. and Pöder, K. (2017). Zero-intelligence agents looking for a job. *Journal of Economic Interaction and Coordination*
4. Veski, A. and Pöder, K. (2016). Strategies in Tallinn school choice mechanism. *Research in Economics and Business: Central and Eastern Europe*, 8(1):5–24
5. Veski, A., Biro, P., Pöder, K., and Lauri, T. (2017). (forthcoming) Efficiency and fair access in kindergarten allocation policy design. *Journal of Mechanism and Institution Design*

Author's contribution to the publications

1. The author's contribution were the ideas of division methods, implementation of the described methods, performing the experiments, and writing the paper
2. The author's contribution were the ideas of decentralised matching models, implementation of the described methods, performing the experiments, and writing the paper
3. The author's contribution were the ideas decentralised matching models, implementation of all the described methods and performing the experiments. Writing the paper and interpreting results were with the co-author

4. The author's contribution was the idea of using genetic algorithms to learn equilibrium strategies, implementation of all the described methods and performing the experiments. Writing the paper and interpreting results were done with the co-author
5. The author's contribution was the idea of using counter-factual preferences to evaluate policy robustness, implementation of all the described methods and performing the experiments. Writing the paper and interpreting results were done with the co-authors

Introduction

Two-sided matching markets are everywhere around us, when booking a cab through a mobile app, applying for a school place for your child to school or even allocating landing slots to airplanes. Usually in these markets agents on both sides have certain preferences regarding each other. For instance, passengers have preferences about the price, quality and type of a taxi and similarly, the driver may prefer passengers with a longer travel distance. Families prefer certain schools over others, while due to local regulations, primary schools usually prioritise children on the basis of proximity. These matching markets have been extensively studied in a static centralised situation, where all participants share their information with a central clearing house, which then handles the allocation. A seminal paper by Gale and Shapley (1962) initiated this type of research by providing simple axioms and an algorithm to compute the allocation.

The inner workings of markets have long been at the centre of economics research. The (neo)classical assumptions are that sellers and buyers are rational, somehow find each other at the decentralised marketplace and agree upon a price. This commodity market usually has an equilibrium price, meaning that after all agents willing to sell or buy a good at a specific price have transacted, there are no more agents left willing to sell or buy at that price. If the price is too high, there might be some agents who are willing to sell. Alternatively, some might still be willing to buy at too low a price. Although the equilibrium price may exist, the question remains how self-interested agents would find that price. Walras (see e.g. Bowles, 2004, p. 216) proposed a process for such a market. This process involves an auctioneer, now known as a Walrasian auctioneer, who iteratively collects cost and value information from agents and proposes a price. However, this type of market tends to be considered as highly centralised as it precludes all out-of-equilibrium trading and thus, not realistic in real-world markets.

The questions of market and mechanism design are also prominent topics in theoretical economics (Klemperer, 2004b; Milgrom, 2004), also more recently in computer science (Roughgarden et al., 2007) and even physics (Smith et al., 2003; Farmer et al., 2012). Several Nobel prizes have been awarded on the study of economic design. To Leonid Hurwicz, Eric S. Maskin and Roger B. Myerson “for having laid the foundations of mech-

anism design theory” Nobelprize.org (2007). To Alvin E. Roth and Lloyd S. Shapley “for the theory of stable allocations and the practice of market design” (Nobelprize.org, 2012), which started by Gale and Shapley (1962). In fact William Vickery, joint with James A. Mirrlees, received the prize even sooner “for their fundamental contributions to the economic theory of incentives under asymmetric information” (Nobelprize.org, 1996). Vickery’s fundamental paper was on the second-price auction (Vickrey, 1961). Search friction, which is central in uncoordinated matching markets, has also received attention from the Nobel prize committee in the award to Peter A. Diamond, Dale T. Mortensen and Christopher A. Pissarides “for their analysis of markets with search frictions” (Nobelprize.org, 2010).

Computational agent-based economics studies the outcomes of economic processes of interacting agents (e.g. Tesfatsion and Judd, 2006; Sterling and Taveter, 2009; Leyton-Brown and Shoham, 2009). These agents might not behave in a perfectly rational manner, or possess enough information to do so and exhibit learning by doing. Neoclassical economics emphasises the study of equilibria as it is often feasible to solve the allocation market models by showing at least one good equilibrium with axiomatic analysis. This might be a reasonable approach, if operating under the assumption that economies tend to stay close to an equilibrium. Although there might exist an equilibrium, where the transactions take place under conditions of perfect knowledge and rationality, the equilibrium analysis does not specify a decentralised process by which to find such a competitive equilibrium (e.g. Bowles, 2004, p. 216).

A complex systems based approach to economics puts the agent-based interaction model centre stage (e.g. Kirman, 2016). The main argument is that the interaction model is an important factor in determining the outcome. Furthermore, it is no longer that easy to analyse these systems using fixed-point theorems, but as stated by Durlauf and Young (2001)

The hallmarks of this approach are, first, to explicitly model a socioeconomic system as a collection of heterogeneous individuals. Second, individuals interact directly as well as through prices generated by markets. Peer groups, social networks, role models, and the like have a prominent place when it comes to determining individual behaviour. Third, individual preferences, beliefs, and opportunities are themselves influenced by the interactions that characterize the system. Fourth, the analysis of such processes draws from methods in stochastic dynamical systems theory, supplemented by large-scale simulation techniques. (p. 11)

The double-auction market mechanism is often used in financial trading. This is similar to the Walrasian model, but more decentralised and detailed about how information is shared and trades actually occur. In terms of allocative efficiency, simple rules suffice to find the competitive equilibrium price even with randomised behaviour both in an agent-based model (Gode and Sunder, 1993b) and also in the context of human experiments (Gode and Sunder, 1993a). A significant part of research revolves around how to actually establish the required trading rules. The current state of the financial market demonstrates that setting such rules is not easy (e.g. Patterson, 2012; Budish et al., 2013; Lewis, 2014).

However, as regards matching markets, there is no market-clearing price, all the goods are substitutes and agents have to figure out their utility for each individual good. Conversely, in a commodity market, goods have a type and the only question is how to trade one type of good for another. The mediating institution is usually money, i.e. the goods have an associated price for which an agent can buy or sell it. In contrast, in a matching market all the goods are potentially different. As an example of a two-sided matching market, multiple jobs may be available, but agents have different preferences for them which are dependent on more than only the wage. Another example is multiple schools in which agents have a different level of desire to obtain a place. The main characteristic typically of a matching market is that there is no price associated with being matched to a job or a place in a school. People usually do not buy or sell jobs (see e.g. Bowles, 2004, p. 292 for discussion) or municipal school positions. In the latter case, allocation is usually restricted by the law and other considerations.

Although the motivations are different for job matching and school allocation, the models are similar. For example in the school market, there might be 50 or 1000 schools and agents have to form their preferences. In a job market, the positions, companies and locations are different and workers have to form their preferences, but it is hard to gain access to all of the information in order to form a preference list. Moreover, what further complicates the situation is the fact that agents need to determine how to act based on the (available/obtained) information. In some cases, it might be easiest to simply start from one's most preferred option and progress decreasingly. However, this method might not always be beneficial as it could potentially waste time while another agent could seize a feasible, but less preferred option. Furthermore, thinking strategically is not only relevant at decentralised marketplaces, but also regarding centralised clearing-houses. Sometimes even if school choices are centralised, it is beneficial for families to consider different strategies for revealing their preferences (e.g. Abdulkadiroğlu and Sönmez, 2003). However, designing safe and optimal matching markets is possible Gale and Shapley (1962). Still, there is a variety of de-

tails to be considered, such as application quotas (Dur et al., 2013; Aygün and Bo, 2013), couples (Roth and Peranson, 1999), contracts (Hatfield and Milgrom, 2005) and more. In these market the details matter. Nonetheless, as these papers show, there is success in designing these markets which has also been studied before in Estonia in the context of college admission (Veskioja, 2005).

Nevertheless, often the matching market may not be centralised. Roth (2008) points out three main aspects of failure for a matching market: unravelling, thickness (congestion) and strategic behaviour. A centralised clearing-house based on the deferred-acceptance algorithm by Gale and Shapley (1962) would solve the problems of unravelling, congestion and strategic behaviour to an extent. Nonetheless, there remains the problem of attracting participants to the market. Furthermore, establishing a centralised allocation may be difficult due to potential opposition from market participants. For instance, certain agents might benefit from a decentralised market and hence, not be interested in coordinating. Roth (2015) observes uncoordinated and decentralised matching markets becoming more relevant in the sharing economy and thus the issues mentioned more acute. Ackermann et al. (2008) studies the computational aspects of certain interaction models and shows that the time required to find a stable solution similar to the result of the deferred-acceptance algorithm is at least exponential regarding the number of agents.

However, the operation mechanisms of decentralised markets have not been extensively studied (Roth, 2008, p. 563). Similarly to a competitive equilibrium model, there is an analytical model for a decentralised job-market, namely the matching function (Petrongolo and Pissarides, 2001). This models the probability of a free agent being matched to a certain open position, but does not specify how the interaction and discovery process work. Rather the probability is attributed to characteristics such as education, experience, etc. and frictions that maybe attributable to an interaction model. Nevertheless, the matching function still neglects to specify the agent coordination model.

Moreover, with the growth of the internet economy, several matching markets have emerged (e.g. Evans and Schmalensee, 2016; Choudary et al., 2016). These matchmakers are often also two-sided, matching apps to users, taxis to passengers, etc. While in many situations the goods on one side can be considered unlimited, e.g. one user downloading an app does not limit others doing so as well, in other situations the goods are limited. For instance, if a taxi is booked, no other passenger can order the same taxi.

Currently, the matchmakers are concerned with attracting participants to their platform to solve the thickness problem. Thus, the transaction pricing model has been the main concern for these businesses (Rysman,

2009, p. 140). The main function of these platforms is to help market participants to find each other. Usually, there is no coordinating mechanism behind the platform, e.g. the Deferred-Acceptance algorithm. Rather, agents take simple myopic short-term greedy actions to select a match. The decentralised and dynamic nature of the markets compels us to ask if we can do better.

In the thesis, the progression is made from a decentralised to an optimal centralised two-sided matching market. We start the thesis by investigating behaviours in decentralised two-sided matching markets. We compare the matching size and allocated rank with different behaviours in a decentralised matching market, as well as to a centralised clearing-house based matching market. We proceed by studying the centralised two-sided clearing-house used in the Tallinn primary school choice mechanism. As this a manipulable mechanism, we investigate the equilibrium strategy and compare it to an optimal and strategy-proof mechanism. Finally, we use an optimal and strategy-proof mechanism to examine policies for constructing priorities for kindergarten allocation. We also compare the efficiency and fairness properties of the different policies.

Motivation

The main examples serving as our motivation are the different place allocation models in use in Estonia. The two cases we study are the primary school allocation in Tallinn and kindergarten allocation in the municipality of Harku. In brief, the procedure in Tallinn consists of two parts: first, a decentralised allocation of children to selective groups and schools; second, a centralised allocation of the remaining places. In Harku, the allocation is centralised, although it is decided by a committee of the heads of kindergartens.

It is unclear how effectively families manage to behave in these allocation settings. First, in a decentralised situation they might give up early rather than wait for an opening in a more preferred school, or they might limit the number of schools they apply to. This creates a kind of a lock-in, where several families stick to their position and thus, a favourable trade might not occur. The behaviour that would guarantee a place for a family at their most preferred school in the second centralised stage of the Tallinn process is also not clear. The collection procedure limits the choices of families to three, which moreover cannot even be submitted in a preferential order. Families have the motivation not to present their true preferences for various reasons and the data will show that families do try to consider how best to present their preferred schools.

The aim of the thesis is to investigate these matching markets. Firstly, a way of modelling the behaviour of agents in decentralised markets is proposed followed by a comparison of the aggregate outcomes with an optimal centralised mechanism. Second, a game-theoretic model of the centralised part of the Tallinn school choice mechanism is proposed and the equilibrium behaviour as well as the outcome of that mechanism is investigated.

The main way of improving the allocations in matching markets is to design a deferred-acceptance based centralised market. The municipality of Harku has operated a nearly centralised clearing-house for several years. However, families can only submit three alternatives and sometimes these were not enough to provide a place. Consequently, other alternatives were suggested to the families. This has the potential to create misallocations, as a family could have received a higher-ranked place, if it had stated initial preferences differently. In addition, the main criteria for selecting children is the application date. Harku is interested in making their allocation more flexible, based on distance from kindergarten and any siblings in the same kindergarten. As such considerations alter the ordering of children for kindergartens and the mechanism should take this into account. In order to have a simple and optimal allocation mechanism for families, a redesign of the clearing-house mechanism was required. The mechanism used thus far was widely known to be unsatisfactory for families, even more so when there is a limit on the number of preferences to be submitted.

The claims and contributions

The central claim of this dissertation is that designing allocation rules requires very detailed consideration, as minor elements in the design can have adverse effects on the final allocation. This can be due to the behaviour of the agents or because of rule design for a dynamic environment. In order to support this main claim we examine three claims in the dissertation:

- A In an uncoordinated and decentralised market, the behaviour of agents is the key determinant of the matching properties – size and rank. We aim to show that a decentralised market is significantly worse in terms of assigned agents and assigned rank than a deferred-acceptance based centralised matching.

- B Merely centralising the allocation process is insufficient. We aim to show that the centralised Tallinn school choice mechanism design incentivises agents to report insincerely. In addition these behaviour results in a less preferred match for some agents, while benefitting other agents compared to an optimal allocation.

C Even with a mechanism where agents are motivated to report truthfully and are guaranteed an optimal match, we aim to show that the allocation is sensitive to the implementation of the policy and to changes in the structure of family preferences. For example a policy of matching children to nearby kindergarten can be implemented either by absolute or relative distance. The comparative allocated rank is significantly different depending on the choice between the two.

To support these claims, we conduct three computational experiments in the matching market framework. The basic assumption of agents with preferences is the same in all models. We look at three two-sided matching environments: a decentralised market, a centralised manipulable and a centralised truthful mechanism design.

Outline of the dissertation

In chapter 1 we review the fundamentals of mechanism design and game theory, including the relevant concepts for this thesis. Examples are provided of axiomatic design of mechanisms for auctions, fair division and matching. Additionally, we discuss aspects of mechanism design wherein the analytic axiomatic approach is intractable and we also motivate a computational approach.

In chapter 2 we develop three simple behaviours in a uncoordinated and decentralised matching market. The behaviours range from simple one-shot matching choices to more sophisticated behaviour for producing a successful matching transaction. It emerges that the myopic behaviour proves superior to more informed behaviours in terms of matching size, but inferior to the utility of the average agent. The main dimension explored are market thickness, i.e. the balance between the numbers of agents on the two sides of the market. We manage to recreate a Beveridge curve similar to the one observed in job market literature. The second component is the structure of the preferences of agents and its impact on the Beveridge curve. We compare the decentralised mechanisms to a centralised clearing house based on the deferred-acceptance mechanism. We use a median rank to measure the efficiency of a matching. We found that the median rank can be as much as 20 times worse in a decentralised market than in a centralised deferred-acceptance based market, which occurs when correlation in preferences is small and the lists are long. Moreover, in the case of longer preference lists, the decentralised mechanism has many unassigned agents and the matching is unstable. On the other hand, a centralised mechanism computes a stable matching, usually with almost no unassigned agents.

In chapter 3 we describe a mechanism used for primary school allocations in the city of Tallinn. We observe that the mechanism is complicated and consists of multiple stages of decentralised and centralised matchings. We define and further explore the centralised mechanism in Tallinn school choice. By using genetic algorithms, we show how the Tallinn mechanism incentivises families to manipulate their preference revelation by reporting only a few schools, and not always from the top of their preference list. This claim appears to match the observed behaviour. In addition, it emerges that the expected utility in the Tallinn mechanism is higher compared to the widely used deferred-acceptance mechanism, although the number of unassigned students is also higher.

In chapter 4 we study kindergarten allocation practices in an Estonian municipality, Harku. We describe the allocation practice used until 2015, followed by an overview of the 2016 system, which was redesigned on the basis of our recommendations. The new mechanism provides a child-optimal stable matching, with priorities based on siblings, distance and other factors. We evaluate seven policy designs in order to understand efficiency and fairness trade-offs based on the 2015 and 2016 admission data. In addition to the real data analysis, we conduct a counter-factual policy comparison and sensitivity analysis using computational experiments with generated preferences. The findings show that different ways of considering the same priorities can have a significant aggregate effect on the allocation.

In chapter 5 we conclude the thesis. We discuss the assumptions in all of the conducted experiments. We view the types of markets as stages in moving towards an optimal centralised market allocation. Finally, we assess the claims and their validity in light of the obtained results.

1 Mechanism Design Background

1.1 Game theoretic foundations

Game theory analyses situations, where agents are seeking to maximise their utility in an interaction with other agents. We assume the agent's utility is known to the agent and he can assign values to different outcomes of the game. A simple game consists of agents and their strategies. In Table 1.1 we describe a well-known game of Prisoner's dilemma with two agents: Alice and Bob, with both having two possible strategies: cooperate and defect. The intersection of each strategy describe the payoff (utility) profile for the agents if they jointly choose to play some strategies. For example if both would select to cooperate, they payoff for both would be 4 for both.

The game can be thought of as a model of resource economy. If two agents would cooperate on managing the resource, by not over-consuming, they would both receive an utility of 4. However, if just one over-consumes then he would receive greater utility. And if you see the other agent over-consuming you might as well over-consume, because that will increase your utility from 0 to 2. Over-consuming will lead to unsustainable state, thus the utilities are lower.

Table 1.1: Tragedy of the commons (Prisoner's dilemma)

		Alice	
		Cooperate	Defect
Bob	Cooperate	4, 4	6, 0
	Defect	0, 6	2, 2

If Bob would cooperate we can easily observe that Alice would do better by playing defect, that is her utility u_a would increase from 4 to 6. Similarly we can reason, if Bob would play defect, Alice's payoff would also be greater if she played defect. Since the game is symmetric, we can change the names

of the agents and the payoffs would be the same, then the same reasoning applies to Bob. So the result of the game would be for both to play defect. This is the unique, dominant strategy, Nash equilibrium of the game.

Definition 1. *Best response for agent i to a strategy profile s_{-i} is a strategy s_i^* , such that $u_i(s_i^*, s_{-i}) \geq u_i(s_i, s_{-i})$ for all strategies s_i .*

We denote by s_i the strategy of agent i and with s_{-i} strategies for all other agents but i . In the Prisoner's dilemma example both agents play their best response strategies.

Definition 2. *Nash equilibrium is a strategy profile $s = (s_1, \dots, s_n)$ if for all agents i , s_i is a best response to s_{-i} .*

Depending on the game there might be multiple Nash equilibria (e.g. Leyton-Brown and Shoham, 2009). The game in Table 1.1 is also interesting that agent always play defect regardless of the other agent, so the defect strategy is the dominant strategy of the game.

Definition 3. *Dominant strategy s_i for an agent i is a strategy that is best response to all possible strategy profiles s_{-i}*

If we look at the game in Table 1.1, we see that both agents would be better off by cooperating, but this is not strategically viable. Ideally we would like the game's Nash equilibrium solution to be Pareto optimal.

Definition 4. *A strategy profile s is Pareto dominating if for any other strategy profile s' for all agents i , $u_i(s) \geq u_i(s')$ and for some agent j , $u_j(s) > u_j(s')$.*

Definition 5. *A Pareto optimal strategy profile s does not have any strategy profile s' that would Pareto dominate it.*

In the game described in Table 1.1 both agents cooperating is the Pareto optimal outcome of the game. Although there are other Pareto optimal outcomes, only both cooperating Pareto dominates the equilibrium solution. The question for mechanism designer is how to design the game such that when agents play their best strategy, the outcome would be good, e.g. Pareto optimal. In later sections we will see depending on the game model, there might be different optimality notions.

A central planner might provide incentives so the payoff structure of the games changed, like for example in Table 1.2. This might be also achieved using institution or social norm (Pöder, 2010), exactly how this is accomplished not important here, but the equilibrium solution should emerge for both agents to cooperate.

Table 1.2: Tragedy of the commons (Prisoner's dilemma)

		Alice	
		Cooperate	Defect
Bob	Cooperate	8, 8	6, 4
	Defect	4, 6	2, 2

So far we considered a game where the strategies of other agents we known to all agents. In a more general setting the payoffs for a particular strategy might be unknown, but there is some probability distribution over the payoff structure of the game.

There is a set $\theta \in \Theta$ of possible games each with a probability $p(\theta)$. The game still has a set A of actions. By $s_j(a_j|\theta_j)$ we denote the equilibrium strategy in a particular game, which is one when action a_j is played when game is θ_j . This might also be a probability, when the equilibrium strategy is not pure, but for simplicity we have not considered here.

Definition 6. *Agent's Ex ante expected utility is defined as*

$$E[u_i(s)] = \sum_{\theta \in \Theta} p(\theta) \sum_{a \in A} \left(\prod_j s_j(a_j|\theta_j) \right) u_i(a, \theta)$$

Definition 7. *Best response in a Bayesian game is*

$$BR_i(s_{-i}) = \arg \max_{s'_i \in S_i} E[u_i](s'_i, s_{-i})$$

Definition 8. *Bayes-Nash equilibrium is a strategy profile that $\forall i$ $s_i \in BR_i(s_{-i})$*

Assume now that with probability $\frac{1}{2}$ we play the game in Table 1.1 and with probability the game in Table 1.2. The expected utility of the game for Alice and Bob would be $E[u_A] = E[u_B] = \frac{1}{2}2 + \frac{1}{2}8 = 5$ as the equilibrium of the second game is both cooperating with a payoff of 8.

To find the equilibrium of the join game we have to compute the expected utilities for each combination of actions as in Table 1.3. The Nash equilibrium of the game is (DC, DC) as expected. In the first game Alice and Bob would both play defect and in the second game cooperate, as we also found when we reasoned individually. This might no always be the case, depending on the probabilities of being in one game or the other.


However real-world is more complex and interactions are usually not one-shot games as described above, rather interactions are repeated. So the strategies would have to specify longer term behaviour. One way to model these is to use so called folk theorems, which specify a strategy for next round as a response to observations from previous round(s).

Table 1.3: Joint game

		Alice			
		CC	CD	DC	DD
Bob	CC	6, 6	5, 4	7, 4	6, 2
	CD	4, 5	3, 3	5, 3	4, 1
	DC	4, 7	3, 5	5, 5	4, 3
	DD	2, 6	2, 2	3, 4	2, 2

In Figure 1.4 we present some folk strategies in repeated interactions. Here *Always defect* and *Always cooperate* ignore observations from previous rounds. However *Tit-for-Tat* always adjusts to the observations, when the other side defects (\cdot, D) it also defects and when the other side cooperates (\cdot, C) it also continues or changes to cooperation.

Table 1.4: Strategies in a repeated game



If we put these strategies to a similar matrix game, we would obtain the payoff structure as in Table 1.5. There are two equilibrium strategies here $s_1 = (\text{Always defect}, \text{Always defect})$ and $s_2 = (\text{Tit-for-Tat}, \text{Tit-for-Tat})$. So we see that once there is some adaptivity and retaliation to defection, cooperation is suddenly a feasible strategy as opposed to a myopic, one-shot, behaviour.

Table 1.5: Mean payoffs in an infinite game

		Alice		
		Always cooperate	Always defect	Tit-for-Tat
Bob	Always cooperate	4, 4	6, 0	4, 4
	Always defect	0, 6	2, 2	2, 2
	Tit-for-Tat	4, 4	2, 2	4, 4

Also the potential strategy-space in these games is very large and the best responses depend on the strategies in the population. Here computational experiments are the main tools applied. One potential implementation is by Wilensky (2002) in NetLogo (Wilensky, 1999). Table 1.6 shows

mean payoffs in three possible populations using Wilensky (2002) for about 200,000 iterations. In the experiment each agent remembers other agents' previous actions and acts according to its strategy.

Table 1.6: Mean payoffs (\bar{u}) in repeated Prisoner's dilemma

Strategies	\bar{u}	Strategies	\bar{u}	Strategies	\bar{u}
Tit-for-Tat	1.95	Tit-for-Tat	2.31	Tit-for-Tat	2.29
Always D	1.00	Always D	2.39	Always D	2.54
		Always C	1.96	Always C	1.85
				Random	2.25

In a population with no cooperators, Tit-for-Tat is better than always defecting as they have higher payoffs when playing against agents with the same strategy. However when cooperators are introduced defector can take advantage of them and gain higher payoffs. Even random agents do better than cooperators, in fact they do almost as good as Tit-for-Tat, although it is also dependent on the population. In longer and more complex repeated games it might be hard to figure out a best response as the strategy might be very complicated. For example to find a good strategy automata has been subject for research for quite a while (Axelrod, 1980; Nowak and Sigmund, 1993; Sigmund, 2010), and still is (Blake et al., 2015).

1.2 Auctions

In most well known auctions types there is one good for sale by one seller and there are multiple potential buyers for the good. Each buyer has some value v_i for the item. The question is how to design a system (auction) that have some good performance guarantees. We'll look at two types of performance measures: maximising social welfare and maximising expected revenue of the seller.

We usually concentrate on looking at sealed-bid auctions. In a sealed-bid auction buyers submit their bid to an auctioneer so that other buyers will not see it. Then the auctioneer makes the decision to whom to allocate the good and their payment. In most cases the good is allocated to the highest bidder and the payment amount will depend on the chosen auction format.

There is also another type of auctions, called open-cry auction. This is a more traditional format, where each bidder publicly announces his bid and the highest bidder wins. The price paid is usually the announced bid.

We continue by analysing the sealed-bid auctions. In some sense these auctions can be thought of as an equivalent to some open-cry auctions, although not in all aspects, like for example information about others valuations.

1.2.1 The first price auction

First price auction is probably the most well-known choice for auctioning. In the first price auction the good is allocated to the highest bidder and the bidder pays the amount he bid. The obvious question for the bidder is: What to bid?

Let us look at a situation where the bidders know the distribution of valuations of other bidders, but not the actual values. Assume it has a uniform distribution and is between $[0, 1]$. Let b_i denote the bid of agent i and as before v_i his true valuation. However a bidder i knows his own valuation v_i . If the bidder wins, his utility from the payoff will be $u_i = v_i - b_i$ and in the case of losing $u_i = 0$. This is known as the quasi-linear utility setting. That is agents only care about their own value and payment and it does not depend on the total amount of money he has (Leyton-Brown and Shoham, 2009, p. 268). Additionally, in uncertain settings, we assume the agents want to maximise his expected utility, i.e. they are risk-neutral.

An aspect we will not look at too closely is agent's risk attitude. In a first price auction it matters what is the agent's risk attitude to what he would bid, here we assume that all the bidders are risk-neutral. See Klemperer (2004a) for more references on risk-aversion.

If there are only two risk-neutral bidders with values drawn independently and uniformly from $[0, 1]$ then $(\frac{1}{2} \cdot v_1, \frac{1}{2} \cdot v_2)$ is a Bayes-Nash equilibrium strategy (Leyton-Brown and Shoham, 2009, p. 336). Bayes-Nash equilibrium is an equilibrium under uncertainty, we need to reason about the valuation of the other agent. Here we assume the value of the other bidder is drawn from a uniform distribution.

So far we looked at only two bidders. What happens when there are more? It would involve a lot of integrals. In general with n bidders the equilibrium strategy profile is $(\frac{n-1}{n}v_1, \dots, \frac{n-1}{n}v_n)$ (Leyton-Brown and Shoham, 2009, p. 337). Although we can observe that the strategy even for two bidders it is not dominant strategy to bid truthfully. Also the expected revenue from the auction is $\frac{n-1}{n} \cdot v_{max}$, where $v_{max} = \max\{v_1, \dots, v_n\}$.

Additionally, the winner of the auction paid the highest possible price for the item. Meaning there is nobody who would be willing to pay the same price. The bidder might have incorrectly valued the item and is now stuck with it, because nobody is willing to pay the same price. This is known as "winner's curse". So bidders might not have incentive to bid their true value.

1.2.2 The second-price auction

However there is a better way to run the auction, by using a second-price auction. In the sealed-bid second price auction the winner is still determined by the highest bidder, but the price for the good will be the price of the second highest bidder. Here the weakly dominant strategy is bid your true value (Leyton-Brown and Shoham, 2009, p. 334). Intuitively, when a bidder's value is the greatest, he will gain positive utility by bidding his true value although it does not matter how much above the second place value he will bid. Although when he bids lower his utility will be zero.

Both of these auction mechanisms guarantee that the item is allocated to the highest bidder. Although with the first price auction it is only in expectation, it does not necessarily happen in every instance. By allocating the item to the highest bidder we make sure that we maximise the total utility in the allocation – social welfare.

Since there exists a dominant strategy, agents do not need to reason how other agents bid. This is beneficial for bidders, but also for the seller. As revenue from Bayes-Nash equilibrium in first price auction is highly dependent on agents' behaviour, the second price auction is much less so.

If bidder valuations are drawn from a uniform distribution $[0, 1]$ and a winning bid is v_{max} there are $n - 1$ other valuations drawn from a uniform distribution $[0, v_{max}]$. To find the second-highest bid, we need the 1st order statistic of the uniform distribution $[0, v_{max}]$. The k th order statistic is $\frac{n+1-k}{n+k}v_{max}$, which is the expected k th largest value of the uniform distribution. So the 1st order statistic of $n - 1$ bidders is $\frac{(n-1)+1-1}{(n-1)+1}v_{max} = \frac{n-1}{n}v_{max}$. This is the same as in first-price auction. As it turns out these auctions are expected revenue equivalent.

Theorem 1. *Assume that each of n risk-neutral agents has an independent private valuation for a single good at auction, drawn from a common cumulative distribution $F(v)$ that is strictly increasing and atomless on $[\underline{v}, \bar{v}]$. Then any efficient auction mechanism in which any agent with valuation \underline{v} has an expected utility of zero yields the same expected revenue, and hence results in any bidder with valuation v_i making the same expected payment (Leyton-Brown and Shoham, 2009, p. 323).*

1.2.3 The VCG mechanism

The second price auction can be further generalised to multiple items and multiple bidders. The described second-price auction is also known as Vickrey auction, due to its author Vickrey (1961). But this auction applies only to individual items, what if we have a more general setting of multiple items? Or even more generally what might be the goals of a social planner

for selecting and designing a mechanism? Usually three goals are the most prominent (e.g. Narahari et al., 2009, p. 7; Leyton-Brown and Shoham, 2009, p. 273-274; Nisan, 2007, p. 225), although there could be others:

- a mechanism should be strategy-proof, it should always be in the participant's best interest to state their true valuations
- a mechanism should be efficient, the items should be allocated so that they would create the largest value in a society
- it should be feasible, in polynomial number of steps, to compute the solution of the mechanism

All of these properties were satisfied in the second-price auction. We saw it is easy to report ones true valuation; the items were allocated to the highest bidder, thus maximising efficient; last, it was easy to find the solution, it would require n steps, for n bidders.

To generalise the second-price auction, there is a class of mechanisms known as the Groves mechanisms (Leyton-Brown and Shoham, 2009, p. 273-274). It has two components, the allocation rule:

$$\chi(\hat{v}) = \arg \max_x \sum_i \hat{v}_i(x)$$

and the payment rule of the form

$$\wp_i(\hat{v}) = h_i(\hat{v}_{-i}) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))$$

First we see that the mechanism would always allocated to goods to highest bidders, as this would maximise the $\chi(\hat{v})$. It is important the payments does not depend on i -s valuation and the two components of $\wp_i(\hat{v})$ do not. If that is satisfied the payment rule guaranteed to be strategy-proof. There is also the freedom for the mechanism designer to choose h_i . Additionally, it turns out, a mechanism is strategy-proof only if it is a Groves mechanism. See **Theorem 10.4.3 (Green-Laffont)** in (Leyton-Brown and Shoham, 2009, p. 278).

The next obvious question is how to select h_i . There we have Clarke pivot rule (Leyton-Brown and Shoham, 2009, p. 280):

$$h_i(\hat{v}_{-i}) = \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}_{-i}))$$

This calculates the value of the allocation without the agent i . The χ is still the welfare maximising allocation function. As promised this function h_i does not depend on the valuation of agent i . The resulting mechanism would be:

$$\begin{aligned}\chi(\hat{v}) &= \arg \max_x \sum_i \hat{v}_i(x) \\ \wp_i(\hat{v}) &= \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}_{-i})) - \sum_{j \neq i} \hat{v}_j(\chi(\hat{v}))\end{aligned}\tag{1.1}$$

What is so special about the Clarke pivot (or payment) rule? First it guarantees that the payment $\chi_i(\hat{v})$ is always positive when all valuations are nonnegative (e.g. Roughgarden et al., 2007, p. 219). It is individually rational, meaning it does not create negative utility to be participating in the mechanism. For example an option would be to set $h_i = \infty$ or some other large number, but this would always create a huge payment for the participants and they would have negative utility and would not like to participate in the mechanism. And we would still like to make sure it is positive, so the seller can collect some revenue for the goods sold.

Although the VCG mechanism (1.1) satisfies the three properties above, the computational complexity being satisfied when finding an efficient allocation is doable in polynomial time, there still are some good properties it does not satisfy. For one the revenue for the auctioneer might be very low, compared to the winning valuation. There are many others reviewed in (Leyton-Brown and Shoham, 2009, p. 280-288).

1.2.4 Optimal auctions

So far we have concerned ourselves with maximising the utility of the society – social welfare. In many situations the auctioneer is instead interested in maximising its own revenue. This arises in many situations, for example Ad-Auctions (e.g. Edelman et al., 2007; Nisan et al., 2009) and selling goods on auction sites (eBay, 2016).

We compared the first- and second-price auctions and found find them to be equivalent in terms of expected revenue. If the second bidder’s value is low and the first bidders high, then the auctioneer would gain only low revenue from the auction. Furthermore the revenue can be arbitrarily low even in the VCG and other mechanisms with the same allocation rule. For example consider a situation with three bidders on two items (Ausubel and Milgrom, 2010). Bidder one is only interested in both of the items with a total value of 2 million €, whereas bidders 2 and 3 are both interested in just one item are willing to pay 2 million € for an individual item. The VCG would allocate the items to bidders 2 and 3. The payment in this situation would be zero for both bidders, due to the nature of the payment rule.

In a situation where the seller does not care about economic efficiency, but is just interested in maximising its revenue, revenue-maximising (or optimal) auctions are used (Leyton-Brown and Shoham, 2009, p. 329). The differentiating aspect is that the auctioneer would set a reserve price, below which the item would not be sold. This ensures, for example, if in the second-price auction the second bid is low and first high, the auctioneer would still extract at least his reserve price. Although it might happen that the reserve price is above the highest bid, in which case the item is not sold. This is the risk the seller would have to take in order to maximise the expected revenue.

Define for each bidder i a virtual valuation (1.2) (Myerson, 1981; Leyton-Brown and Shoham, 2009, p. 329).

$$\psi_i(v_i) = v_i - \frac{1 - F_i(v_i)}{f_i(v_i)} \quad (1.2)$$

To determine the virtual valuation we would need to know the distribution $f_i(v_i)$, from which the value is drawn and the distribution is such that ψ_i is increasing in v_i . Once the virtual valuations are known for each bidder, we use those to find the welfare maximising solution and compute the payments according to virtual valuations. When all the virtual valuations are all negative the item is not allocated. However the mechanism still remains strategy-proof.

So the point r_i^* where $\psi_i(r_i^*) = 0$ is when virtual value becomes positive and the item is allocated, this can be considered as a reserve price below which the item is not sold. When the distribution for all bidders is the same, this is simply the point r^* when $r^* - \frac{1 - F(r^*)}{f(r^*)} = 0$.

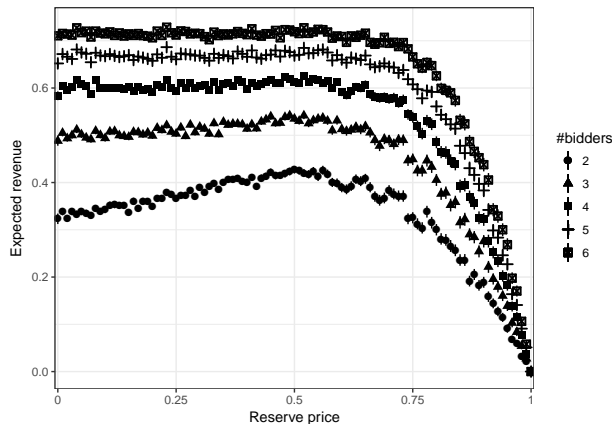


Figure 1.1: Revenue with uniform valuations

For illustration let's look at a small scale computational experiment with reserve prices in range $r^* \in [0, 1]$ and randomly drawn values for all the bidders from $\mathcal{U}(0, 1)$. The optimal reserve price r^* in this case is $r^* - \frac{1-r^*}{1} = 0 \implies r^* = 0.5$. In Figure 1.1 we see the optimal reserve price with between 2, ..., 10 bidders. We see that having slightly greater reserve price is slightly better than having a reserve price of zero. Also we can observe that overshooting, setting a too high reserve, can also be detrimental to revenue, much more so than not setting a reserve price at all.

Of course the assumption of uniform valuations might not hold and we would require more information to select revenue-maximising reserves, which is hard if the auction has not been run before. Although this has been successfully used in large scale ad-auctions (Ostrovsky and Schwarz, 2011) to increase revenue.

Additionally we see from Figure 1.1 that having more bidders increases revenue and even more than the reserve price. It has been shown that adding one more bidder to an auction is better than an optimal reserve in a single item auction (Klemperer and Bulow, 1996).

Theorem 2. *The expected revenue from an auction with $n + 1$ bidders and no reserve is at least as high as the revenue from the corresponding auction with n bidders using the optimal reserve price (Milgrom, 2004, p. 148).*

1.3 Fair division

A slightly different resource allocation problem is fair division. Here there is no seller, who is interested in revenue maximisation. Rather there is a pool of items that need to be allocated between agents, like for example inheritance, divorce settlement, a lot of land etc. A fair division problem has a set of agents and an item or a set of items to be shared. The problem comes in many flavours, it could be:

- one continuous divisible good to be divided – cake-cutting Robertson and Webb (1998); Brams and Taylor (1996)
- several divisible goods – fair division with divisible good Brams and Taylor (1996)
- several indivisible goods – fair division with indivisible goods Brams and Taylor (1996)

The problem of fair division has been present for a long time (e.g. Brams and Taylor, 1999), but is often cited to be first formalised as a mathematical problem by Steinhaus (1948). Initially the main goal was to find an allocation that would guarantee a proportionally $1/n$ fair share to all agents.

Over time other criteria have been introduced, like envy-freeness, utilitarian, egalitarian solutions, Pareto optimality and Nash’s Bargaining Solution (e.g. Chevaleyre et al., 2006; Veski, 2012). More recently with the advent of algorithmic game theory (Roughgarden et al., 2007) the equilibria and incentive properties of different allocation procedures have been studied (Han et al., 2011a; Van Essen, 2013; Brânzei et al., 2016; Aziz et al., 2015), but also computational complexity (Van Essen, 2013; Nguyen et al., 2013).

More formally we have a set of n agents $\mathcal{A} = \{a_1, \dots, a_n\}$ and a set m goods or resources $\mathcal{R} = \{r_1, \dots, r_m\}$. The goods, as mentioned, can be of different type: heterogeneous divisible, indivisible and a single continuous divisible good. Each agent a_i has some utility u_i for resources that is a mapping $u_i : \mathcal{R} \rightarrow [0, 1]$. It is usually assumed that $\sum_{r_j \in \mathcal{R}} u_i(r_j) = 1$. So the main difference with auctions is how the valuations are treated. While in auctions agents can have various levels of valuations, but eligibility for bidding might also be determined by their wealth, in fair division all agents are treated by ignoring their wealth or access to credit.

We now review some fundamental properties and solution concepts of fair division allocation procedures.

Definition 9. *Proportionality is satisfied if $u_i(X) \geq 1/n$ for every agent $a_i \in \mathcal{A}$*

Definition 10. *Utilitarian allocation is X if for any other allocation Y $\sum_{a_i \in \mathcal{A}} u_i(X) \geq \sum_{a_i \in \mathcal{A}} u_i(Y)$. Utility of an allocation is $\sum_{a_i \in \mathcal{A}} u_i(X)$.*

This is usually the preferred solution as it produces the highest overall welfare, similar to social welfare (efficiency) in auctions. However this might not always be desired, as some agent might have an unduly low access to a resource. Considering the properties above, utilitarian solution might not necessarily be proportional, although it is always Pareto optimal. Thus we introduce additional concepts that aim for a fairer allocation.

Definition 11. *Egalitarian allocation ensures that $u_i(X) = u_j(X)$ for any pair of agents i, j with $i \neq j$*

Definition 12. *Envy-free allocation ensures that $u_i(X(a_i)) \geq u_i(X(a_j))$ for any pair of agents i, j with $i \neq j$*

Definition 13. *Nash’s Bargaining Solution is a solution where the product of individual utilities is at its maximum, i.e. X is NBS if for any other allocation Y we always have $\prod_{a_i \in \mathcal{A}} u_i(X) \geq \prod_{a_i \in \mathcal{A}} u_i(Y)$*

In bargaining literature the latter is a predicted outcome of a bargaining game similar to the fair division problem. However, agents have to negotiate an allocation rather than having some procedure hand it to them.

Assuming that the agents have enough time and will to negotiate they will eventually reach the Nash's Bargaining Solution (Binmore, 2005; Osborne and Rubinstein, 2011).

1.3.1 Single divisible item

The simplest example of fair division is cake cutting between two agents – Alice and Bob. Suppose we have a cake with two flavourings (e.g. vanilla and chocolate). Alice has higher preference for vanilla and Bob for chocolate. How should the cake be divided? The cut-and-choose procedure is simple (e.g. Robertson and Webb, 1998):

- Randomly select an agent as a Cutter
- Cutter splits the cake in two
- Chooser select one of the pieces
- Cutter receives the remaining piece

In a situation where Alice prefers only vanilla and does not know anything about Bob's valuations, the best action for her would be to cut the vanilla part in half. This ensures that Alice will receive at least half of her value. However if Alice knows Bob's valuation they can do at least as good or better. When Bob prefers also only vanilla the result would be the same. When Bob is indifferent or prefers only chocolate, Alice would do better by exactly along the chocolate vanilla line. Bob will pick the chocolate half and the remaining vanilla will remain with Alice. This always assumes that agents are utility maximising.

The cut-and-choose procedure guarantees that each agent will receive at least $1/2$ of the total cake. Denote vanilla part of the cake with V and chocolate with C . When Alice values only vanilla, her utilities would be $u_a(V) = 1$ and $u_a(C) = 0$. If she cut the cake in half with $u_a(V_1) = u_a(V_2) = 0.5$, she guarantees herself at least half of her valuation. However, when she knows Bob's valuation, Alice would cut along the vanilla-chocolate line. Alice would get $u_a(V) = 1$ and Bob would pick $u_b(C) = 1$, thus achieving greater total social welfare of $u_a + u_b = 2$.

The 2-agent cut-and-choose could be generalised in many ways (see e.g. Brams and Taylor, 1996; Robertson and Webb, 1998). We will present one simpler, cut-your-own algorithm from (Steinhaus, 1969; Robertson and Webb, 1998).

This example of 3-agents can be easily extended to n -agents. Each agent divides the cake, or property, to 3 (or some n pieces) such that each piece j is worth at least $u_i(r_j) \geq 1/3$. Iteratively allocate to each agent a non-overlapping piece, there is always a way to do this. This already guarantees

$1/3$ to everyone. Although this procedure might not allocate all the cake and there usually is something left. This surplus could be allocated in a similar way, or sold for additional profit.




Figure 1.2: 3-agent example cuts

If we assume that each agent has drawn the lines with his proportional shares, we would obtain the cuts as overlaid in Figure 1.2. We can guarantee each agent a proportional share if we do the allocation as indicated by the colors. There is always a part remaining, unless all the lines overlapped. An extensions might be as in Figure 1.3. This would guarantee each agent more than a proportional share, although is might not necessarily be utility maximising, equitable or envy-free.




Figure 1.3: Extended division

In general finding an envy-free is harder than proportional. The above procedure can always find a proportional allocation for any number of agents, but envy-free procedures are only known up-to four agents (Rothe, 2016, e.g.). And if we are interested in finding a single connected piece for all agents, then it is impossible in general (Stromquist, 2008).

Regarding computational complexity, the above procedures require $n(n-1)$ -bits of information on cuts and only one-bit in case of two agents for proportional allocation. In general we could do better, $1 + nk - s^k = \mathcal{O}(n \log n)$ ($k = \lfloor \log_2 n \rfloor$) cuts, which is the best known bound (Robertson and Webb, 1998, p. 94). The lower bound for envy-free protocols is $\mathcal{O}(n^2)$ (Procaccia, 2009).

1.3.2 Multiple divisible items

These solution concepts are not necessarily aligned even if we had multiple *cakes*, e.g. by finding at utilitarian solution it might give unequal shares to agents. Binmore (2005) presents some simple examples. Caragiannis et al. (2012) has studied these trade-off in a more general setting, with a larger

Table 1.7: Price of X on utility for divisible goods

Price of	Lower bound	Upper bound
Proportionality	$\Omega(\sqrt{n})$	$\mathcal{O}(\sqrt{n})$
Envy-Freeness	$\Omega(\sqrt{n})$	$n - 1/2$
Equitability	$\frac{(n+1)2}{4n}$	n

Table 1.8: Adjusted Winner example (Brams and Taylor, 1996)

$v_i(\cdot)$	m_1	m_2	m_3
Agent 1	.06	.67	.27
Agent 2	.05	.34	.61

number of agents. In Table 1.7 (from Caragiannis et al., 2012; Veski, 2012), we present a few bounds on total social welfare in the allocation regarding other solution concepts. The Table 1.7 shows the ratio of utilities in the best overall utilitarian solution and best utilitarian solution under some restriction: proportionality, envy-freeness or equitability.

Brams and Taylor (1996) proposed a simple method for sharing a potentially divisible items among multiple agents. The procedure is called Adjusted Winner and works as follows:

1. There is a set \mathcal{K} of k goods to be divided and two agents
2. Collect valuations $v_i(k)$ from each agent i , such that $\sum_{j=1}^k v_i(m_j) = 1$
3. In an initial allocation X , assign each good to the agent who values it most
4. if $v_1(X) > v_2(X)$, for some two agents, adjust allocation X , by transferring items $m_j \in X(1)$ in order $\frac{v_1(m_1)}{v_2(m_1)} \leq \frac{v_1(m_2)}{v_2(m_2)} \leq \dots \leq \frac{v_1(m_k)}{v_2(m_k)}$ until $v_1(X) \leq v_2(X)$
5. if $v_1(X) < v_2(X)$, then select the last good transferred m_l and find α such that $v_1(X \setminus m_l) + \alpha \cdot v_1(m_l) = v_2(X \setminus m_l) + (1 - \alpha) \cdot v_2(m_l)$

For example, assume we have valuations as in Table 1.8. An initial allocation of goods would be $X = \{1 : \{m_1, m_2\}, 2 : \{m_3\}\}$, which would make $v_1(X) = 0.73$ and $v_2(X) = 0.61$. Clearly $v_1(X) > v_2(X)$. The order of items to transfer would be $m_1 \succ m_2$, as $\frac{v_1(m_1)}{v_2(m_1)} = \frac{0.06}{0.05} \leq \frac{v_1(m_2)}{v_2(m_2)} = \frac{0.67}{0.34}$. After transferring m_1 we would have $v_1(X) = 0.66 < v_2(X) = 0.67$. So we will have to split up m_2 by solving (1.3). This will result in $\alpha = 0.0099$

$$0.66 + \alpha \cdot 0.34 = 0.0 + (1 - \alpha) \cdot 0.67 \tag{1.3}$$

Table 1.9: Correlations

	Efficiency	Envy	Inequality	Nash
Efficiency	1			
Envy	-0.21	1		
Inequality	0.18	0.80	1	
Nash	0.65	-0.76	-0.59	1

Notice that both agents have an incentive to misrepresent their valuations. For example if Agent 1 would state $v_1(m_2) = 0.4$, then the Adjusted Winner would result in allocation $X = \{1 : \{m_1, m_2, \alpha \cdot m_3\}, 2 : (1 - \alpha) \cdot m_3\}$. Clearly Agent 1 gains for misreporting. In Nash equilibrium we assume all agents do their optimal misrepresentation and know all valuations. The resulting loss in efficiency is known as Price of Anarchy.

Definition 14. *Price of Anarchy is the rate between the optimal OPT solution and the equilibrium EQ solution, $PoA = \frac{OPT}{EQ}$*

It turns out the worst-case the PoA under some conditions in Adjusted Winner procedure is $4/3$ (Aziz et al., 2015). However, when agents are truthful, Adjusted Winner guarantees that the allocation is envy-free, equitable and Pareto optimal.

Additionally we see that the Adjusted-Winner procedure sacrifices some utility for a more egalitarian solution. The social welfare maximising solution would create the value of 134, however the Adjusted-Winner results in 132.67, which is lower. In Table 1.7 we saw what the costs of different properties are. Though the relevant solution concept for economist (e.g Binmore, 2005; Osborne and Rubinstein, 2011) is the Nash Bargaining Solution, mainly because it is produced by free bargaining and satisfies some reasonable axioms. Indeed it also appears that the Nash's solution also has some reasonable trade-off between efficiency, inequality and envy and has the desired correlation direction with these concepts (Veski and Vöhandu, 2010, 2011). In Table 1.9 we show the correlations among the criteria in some situations. We see that the Nash's Bargaining solution has positive correlation with efficiency and negative with envy and inequality. That Nash's Bargaining solution tends to be fair is also recently explored by Caragiannis et al. (2016).

Table 1.10: Price of X on total utility for indivisible goods

Price of	Lower bound	Upper bound
Proportionality	$n - 1 + 1/n$	$n - 1 + 1/n$
Envy-Freeness	$\frac{3n+7}{9} - O(1/n)$	$n - 1/2$
Equitability	∞	∞

1.3.3 Multiple indivisible items

While sometimes the good are divisible like cake or land, however in some situations the value of a good is destroyed when split-up. This procedure requires side payments to make the allocation equitable or envy-free. One of such procedures is the Knaster's procedure of sealed bids (Steinhaus, 1948; Brams and Taylor, 1996). The procedure works as follows:

1. In an initial allocation X allocate items to highest bidders
2. Compute side payments Y , such that

$$v_i(X) - \frac{v_i(\mathcal{K})}{n} + v_i(Y) = v_j(X) - \frac{v_j(\mathcal{K})}{n} + v_j(Y)$$

for any i, j

The main feature of the procedure is that the allocation can be considered equitable. Each agent receives the same amount of money in their valuation over what was their fair share. However, the total proportional share is not equal, as all agents value the entire bundle differently the monetary compensation is also different ratio from the total. In general it is not envy-free either. Haake et al. (2002) describes a compensation method that would be envy-free for any number of agents. Also, it can be easily seen that agents have incentive to misreport their valuations, so the allocation is not strategy-proof (Lyon, 1986). For example if an agent would under-report their value, but would still be highest bidder, they would increase their compensation $v_i(Y)$ amount and consequently their utility.

In Table 1.10 we see the trade-offs in solution concepts in case of indivisible items, similarly to the trade-offs in sharing divisible (Caragiannis et al., 2012; Veski, 2012). The main observation is that the bounds are more severe in the indivisible case. We lose more utility, if we need e.g. envy-freeness and can be infinite if equitable solution is desired.

1.4 Matching markets

A matching market consists of one or two sets of agents and/or items. At least one set has some preference relationship over the other set. For example considering one-sided market, with agents and items, only agents

have preferences over items. This is similar to the fair division model, except that agents do not have a cardinal value for items, but rather a preference ordering. Items higher in the order are more preferred and agents seek to gain items high on their list and they are strategic about maximising their utility.

We concentrate on two-sided matching, that is there is an ordering on both sides. Agents would have preferences over agents on the other side. This model arises in many situations, for example matching jobs to employees. Previously, we had valuations only on one, the agent's side. Also, when the other side is of items not of agents we talk about priorities that items have over to which agent it should be allocated to. In case of priorities we assume are non-manipulable, whereas preferences could be misreported for utility maximisation.

Table 1.11: Agents' preferences

Side \mathcal{A} preferences	Side \mathcal{B} preferences
$a_1 : b_2 \succ b_1 \succ b_3$	$b_1 : a_1 \succ a_2 \succ a_3$
$a_2 : b_1 \succ b_2$	$b_2 : a_3 \succ a_1 \succ a_2$
$a_3 : b_1 \succ b_2 \succ b_3$	$b_3 : a_1 \succ a_3$

In Table 1.11 we have an example of preferences on two-sides of the market. We could think of side \mathcal{A} and \mathcal{B} respectively as men or women, or jobs and employees depending on the situation. Ultimately we need to find a matching between \mathcal{A} and \mathcal{B} . If we are interested in maximising social welfare, we might well consider results from graph theory and find the maximum matching, using for example the Hungarian method $\mathcal{O}(n^2m)$ (Kuhn, 1955), (e.g. Lovasz and Plummer, 2009; Schrijver, 2003), which works with weighted and unweighted graphs, or more efficient Hopcroft-Karp algorithm $\mathcal{O}(n^{\frac{1}{2}}m)$ (Hopcroft and Karp, 1971, 1973) and (e.g. Lovasz and Plummer, 2009; Schrijver, 2003). Here n and m are respectively vertices and edges in a graph and an undirected graph can have at most n^2 unique edges.

However, the maximum matching might not be rational for all parties to be part of. It might be better for two agents, e.g. man and a woman, to form a match outside the matching procedure, as they would both result in a more preferred match. In case of preferences as in Table 1.11. The size of a maximum matching is 3 and in any maximum matching a_2 would have to be matched to either b_1 or b_2 . There are four possible matches of size 3 and in each matching there is a pair of agents who would prefer to match outside of the algorithm (Table 1.12).

Table 1.12: Maximum, but unstable matchings

Matching	Outside	Matching	Outside
$a_1 - b_1$		$a_1 - b_3$	$a_1 - b_2$
$a_2 - b_2$		$a_2 - b_2$	$a_2 - b_1$
$a_3 - b_3$	$a_3 - b_2$	$a_3 - b_1$	

Matching	Outside	Matching	Outside
$a_1 - b_2$		$a_1 - b_3$	$a_1 - b_1$
$a_2 - b_1$		$a_2 - b_1$	
$a_3 - b_3$	$a_3 - b_2$	$a_3 - b_2$	

For example in the first matching (upper left corner) agent pair $a_3 - b_2$ would rather be matched to each other than their current match. The agent a_3 prefers b_2 to b_3 and b_2 prefers a_3 to a_2 (Table 1.11). So we see that in all the maximum matchings there is no such matching that would not have this type of *blocking* pair. So we need to approach more carefully to find a matching.

Furthermore if the matching procedure depends on agents' preferences they might be strategic about how they present their preferences. There are only two known mechanisms (algorithms) that are strategy-proof for at least for one side of the market – the Deferred-Acceptance and the Top Trading Cycles (Abdulkadiroğlu and Sönmez, 2003, e.g.). Deferred-Acceptance is also known as Gale-Shapley algorithm (Gale and Shapley, 1962) for stable marriages. In a stable matching there are no *blocking* pairs as in the maximum matching example. Moving forward we concentrate on stable matchings and the Deferred-Acceptance algorithm and its properties.

More formally we employ a model similar to that used in modelling centralised two-sided matching markets (e.g. Roth, 2008). There is a set $\mathcal{A} = \{a_1, \dots, a_{n_A}\}$ of agents on one side and a set $\mathcal{B} = \{b_1, \dots, b_{n_B}\}$ of agents on the other side. The number of agents on both sides can differ ($n_A \neq n_B$) depending on market thickness. Each agent a_i from \mathcal{A} has a strict preference relation \succ_{a_i} over agents in \mathcal{B} , and similarly for $b_j \in \mathcal{B}$ there is a preference relation \succ_{b_j} over agents in \mathcal{A} . A matching μ is a mapping from $\mathcal{A} \cup \mathcal{B}$ to itself, so that for every $a_i \in \mathcal{A}$, is matched to $\mu(a_i) \in \mathcal{B} \cup \{a_i\}$, and similarly for $b_j \in \mathcal{B}$, $\mu(b_j) \in \mathcal{A} \cup \{b_j\}$. When an agent is matched to itself, $\mu(a_i) = a_i$ or $\mu(b_j) = b_j$ respectively indicates that they are in fact unmatched. Being matched to itself is the least preferred option for all the agents. In addition, for every $a_i, b_j \in \mathcal{A} \cup \mathcal{B}$, $\mu(a_i) = b_j$ implies $\mu(b_j) = a_i$.

Definition 15. A matching is unstable if there are at least two agents a_i and b_j from opposite sides of the market such that $b_j \succ_{a_i} \mu(a_i)$ and $a_i \succ_{b_j} \mu(b_j)$ – a blocking pair. A matching is stable, if it is not unstable.

A stable matching with preferences as in Table 1.11 would be: $\mu = \{a_1 - b_1, a_3 - b_2\}$, which is not maximum possible matching, as this was impossible. Making the resulting matching stable does not yield a pair of agents who would find a better matching outside the procedure, because by definition that opportunity does not exist. Still the question if it is beneficial for all agents to state their true preferences is open. For this we first need to define an optimality of a matching.

Definition 16. A stable matching μ is optimal for agents in \mathcal{A} if there is no stable matching ν for which $\nu(a_i) \succ_{a_i} \mu(a_i)$ or $\nu(a_i) = \mu(a_i)$ for all $a_i \in \mathcal{A}$ and $\nu(a_j) \succ_{a_j} \mu(a_j)$ for at least one $a_j \in \mathcal{A}$

Then it turns out that if the matching is stable and optimal for agents in \mathcal{A} then it is strategy-proof for those agents.

Theorem 3. In the matching procedure which always yields the optimal stable outcome for a given one of the two sets of agents (i.e., for \mathcal{A} or for \mathcal{B}), truthful revelation is a dominant strategy for all the agents in that set (Roth, 1982)

Table 1.13: Preferences with multiple stable matchings

Side a preferences	Side b preferences
$a_1 : b_1 \succ b_2$	$b_1 : a_2 \succ a_1$
$a_2 : b_2 \succ b_1$	$b_2 : a_1 \succ a_2$

In general the set of stable matchings can be greater than one (Knuth, 1997b) and not all are optimal matchings for all the agents in \mathcal{A} or \mathcal{B} . For example with preferences in Table 1.13 we can obtain two stable matchings: $\mu_1 = \{a_1 - b_1, a_2 - b_2\}$ and $\mu_2 = \{a_1 - b_2, a_2 - b_1\}$. In either case there are no blocking pairs and in one case agents in \mathcal{A} obtain their first preferences and in the second agents in \mathcal{B} obtain their first preferences. This turns out to be a general property of stable matchings (Knuth, 1997b).

Theorem 4. No stable matching procedure for the general matching problem exists for which truthful revelation of preferences is a dominant strategy for all agents. (Roth, 1982)

Also it turns out that the number of agents is the same in any stable matching, even more, in some conditions the same set of agents get matched.

Theorem 5. *When all preferences over individuals are strict, and hospitals have responsive preferences, the set of students employed and positions filled is the same at every stable matching. Furthermore, any hospital that has some empty positions at some stable matching is assigned precisely the same set of students at every stable matching (e.g. Roth, 2008; Manlove, 2013, “Rural Hospitals” Theorem)*

In cases, where the preferences are complete, as in Table 1.13, all agents can be matched in a stable matching. However, when preferences are incomplete, as in Table 1.11, the size of the maximum and stable matchings may be different. If we are allowed to have blocking pairs, the objective could be to find a maximum matching with a minimum number of blocking pairs, number of agents with blocking pairs (Eriksson and Häggström, 2007) or even number of instances of envied agents in blocking pairs (Abdulkadiroglu et al., 2017).

In general it turns out it is hard to find among maximum matchings, and matching with the minimum number of blocking pairs or agents.

Theorem 6. *Finding a matching with minimum number of blocking pairs or agents among maximum matchings is not approximable within $n^{1-\epsilon}$, where n is the number of agents in a given instance, for any $\epsilon > 0$, unless $P = NP$. (Biró et al., 2010b)*

Additional results showed that this holds when the length of the preference lists is limited to 3 (Hamada et al., 2009). However, when on at least one side preferences are limited to 2, finding a matching with minimum number of blocking pairs or agents is solvable in $\mathcal{O}(n^3)$ (Biró et al., 2010b).

1.4.1 Two-sided matching - stable marriage

The stable marriage problem (Knuth, 1997b) is the simplest of the two-sided matching problems. There are two sets of agents \mathcal{A} and \mathcal{B} or men and women. Each agent can be matched to at most one agent from the other set. In the simplest case all men and women have full list of preferences over respectively women and men. This ensures that all agents can be matched. Although this model can be easily generalised to a stable marriage with incomplete lists (Manlove, 2013, p. 22) by assuming that some agents can be left unmatched. The definition of stability remains the same if we assume that the last option on each agent’s preference list is to be matched to itself.

To find a man (or woman) optimal stable matching we can use the Deferred-Acceptance algorithms as proposed by Gale and Shapley (1962):

1. Each man proposes to his most preferred woman. Each woman who received more than one proposal rejects all, but her most favourite amount those who have proposed to her. However, she does not accept him yet, but keeps on a string to allow for the possibility of someone better may come along later
2. In the second stage those men who were rejected now propose to their second choices. Again each woman receiving proposals chooses her favourite from the group consisting of the new proposers and the man on her string, if any. She rejects all the rest and again keeps the favourite in suspense.
3. We proceed in the same manner. Those rejected at the second stage propose to their next choices, and the women again reject all but their best proposal they have had so far.

...

The algorithm continues until all the men are matched or have reached the end of their preference lists. Then the matching is declared final. The algorithms runs in $\mathcal{O}(nm)$ time, where n is the number of applicants and m the number of responders, or the average length of preference lists. Gale and Shapley (1962) showed that

Theorem 7. *Every applicant is at least as well off under the assignment given by the Deferred-Acceptance procedure as he would be under any other stable assignment*

The stable marriage problem can also be extended to a college admission problem or hospitals/residents problem (Gale and Shapley, 1962; Roth and Sotomayor, 1990; Manlove, 2013). On one side of the market we have students or residents, who can be accepted to at most one college or hospital. On the other side we have colleges or hospitals who have some capacity c_i on how many applicants they can accept. Both students and colleges still have preferences over each other.

The Deferred-Acceptance procedure would need to be slightly tweaked to work in this more general setting. To find a student optimal matching, student would still propose to colleges, but colleges only reject students who are above their capacity c_j and lower on their preference lists. Gale and Shapley (1962) show that Theorem 7 holds here as well.

The Deferred-Acceptance procedure can be reversed to find a woman optimal matching or a college optimal stable matching. It turns out that this matching is the worst possible stable matching for men or students respectively. Also depending on who is the proposer the matching procedure

is strategy-proof for the proposers. In equilibrium reporting in the Deferred-Acceptance allocation is still stable with respect to true preferences (Roth, 1984; Sotomayor, 2012). However this might not be the case even when the proposers are colleges, as in the college admission problem as colleges can accept multiple students and is thus a more general problem (Roth, 1985).

The earliest use of the Deferred-Acceptance algorithms was in US National Residency Matching Program that matched fresh doctors to their first jobs in hospitals. Initially a hospital proposing algorithm was used that created protest from applicants and was changed in 1998 (Roth, 1997) to be applicant proposing Deferred-Acceptance. However Roth (1997) found that this benefited a very small number of applicants, less than 1 in 1000. So the the difference in being a proposer is negligible in real world markets. Further results (Kojima and Pathak, 2009; Immorlica and Mahdian, 2015) show that in large markets the number of stable partners is small and hospital and student optimal stable matchings are very close.

Although the options for manipulations might be small agents are still motivated to find them if they exist. Teo et al.(2001) shows that if one agent from the accepting, \mathcal{B} -side, would know all the preferences of others participants, it can effectively use the DA algorithm to find an optimal truncation of preference list in $\mathcal{O}(n^2)$ time. They also show that probability of benefiting by cheating is small. In instances of size $|\mathcal{A}| = |\mathcal{B}| = 8$ on average about 5% of agents in \mathcal{B} benefit and in instances 100 and 500 agents on both sides on average about 10% of agents from \mathcal{B} benefit.

Matsui (2011) presents algorithms in stable marriage matching for agents in \mathcal{B} jointly to manipulate a matching, by permuting their preference lists. First he gave an $\mathcal{O}(n^2)$ to check, given a matching μ' , if there is a joint strategy for women that is an equilibrium and μ' is the outcome of \mathcal{A} -optimal matching. Second also gave an $\mathcal{O}(n^4)$ algorithm to check if a given joint strategy is an equilibrium or not.

Sönmez (1997) and Sönmez (1999) further show that colleges and hospitals can in general manipulate stable matchings by misreporting their capacities and pre-arranging matches.

1.4.2 Two-sided matching - school choice

In many situations the preferences for one side of the market are based on some objective non-manipulable criteria like an exam score, distance from home etc. In this situation the allocation mechanism could be considered strategy-proof overall. In the case when preferences are determined by objective criteria we call them priorities as they are not individual preferences of schools, but rather given by a central planner and cannot be manip-

ulated. It turns out that this arises in many situations of school choice (Abdulkadiroğlu and Sönmez, 2003; Pathak and Sönmez, 2013) and even college admission (Aygün and Bo, 2013).

In designing mechanisms for school choice another type of question develops – how to select the criteria and implement the priorities? In most situations the goal of the central planner is to use *fair* criteria for allocating available positions or affirmative action for more disadvantaged part of the society (Aygün and Bo, 2013; Kominers and Sönmez, 2013; Alcalde and Subiza, 2014). In some circumstances there is no objective criteria and a random lottery is used to allocate students to positions and experiments show (Abdulkadiroğlu et al., 2009) than one needs to be careful how the lotteries are implemented.

A critical part in ensuring strategy-proofness is not to limit the preferences revealed to the mechanism (Abdulkadiroğlu et al., 2009). The limitation is believed to help applicants state their preferences, as they do not have to evaluate so many schools. Actually it turns out this introduces another complexity to the problem, since participants also need to consider what to report (Haeringer and Klijn, 2009).

Table 1.14: Agents’ preferences

Side \mathcal{A} preferences	Side \mathcal{B} preferences
$a_1 : b_2 \succ b_1 \succ b_3$	$b_1 : a_1 \succ a_2 \succ a_3$
$a_2 : b_1 \succ b_2 \succ b_3$	$b_2 : a_3 \succ a_1 \succ a_2$
$a_3 : b_1 \succ b_2 \succ b_3$	$b_3 : a_1 \succ a_3 \succ b_2$

We have slightly augmented preferences in Table 1.14 from Table 1.11. An \mathcal{A} -optimal stable matching with original preferences would be: $\mu = \{a_1 - b_1, a_2 - b_3, a_3 - b_2\}$. Now if we limit the number of preferences to two a_2 would be unmatched, so it is better for him to report for example: $b_2 \succ b_3$. Of course this would depend on others’ preferences, information which might be hard to obtain, thus introduces uncertainty on the best course of action.

The best course for a designer is usually not to limit preference submission. However, in school choice families do not usually have resources to evaluate too many school. Some school district solve this problem by providing limited menus to students (Shi, 2015) and in many cases the preference lists are limited in practice (Pathak and Sönmez, 2013). Both approaches are not without their problems and there is currently no ideal solution.

1.4.3 Algorithms for on-line and decentralised matching

In some cases it might be hard to organise a centralised matching scheme. Roth and Vate (1990) observe that some of these decentralised matching markets are still in operation, thus might reach a stable outcome. They propose a procedure (RVV) by randomly satisfying blocking pairs and show that this random sequence converges to a stable matching with the probability of 1. In case of equal number of agents with full preference list the computational complexity of this algorithm is still polynomial – $\mathcal{O}(n^4)$ (see Manlove, 2013, p. 81-82). Another slightly different decentralised algorithm (ROM) is presented by Ma (1996). Many results from RVV carry over to ROM. In these decentralised mechanisms that agents arrive in sequence and are matched to already arrived agents. It turns out that it can be beneficial for an agent to arrive later (Blum and Rothblum, 2002). This is interesting when comparing to winning strategies in a continuous double auction market in section 1.5.2.

However, these algorithms rely on agents fully knowing their preferences and selecting the best possible blocking pair to satisfy in each step. In chapter 2 we explore a randomised setting, where agents do not know their preferences and cannot make proposals to their best blocking pair. Some of the randomised settings have been previously explored by Ackermann et al. (2008) and they give an exponential lower bound for the convergence time to a stable matching.

More recently on-line matching problems and their application in e-commerce have become more prominent. Mehta (2013a) survey matching algorithms, which include vertex-weighted and other more complex valuation functions. Furthermore, they do not necessarily consider stable matchings, rather they are interested in maximum matchings. As stable matchings might be smaller, then for an e-commerce maximum matchings can potentially produce more revenue. However, Mehta (2013a) does not consider manipulations, which might reduce the overall revenue.

Gu et al. (2015) describe a framework for wireless resource allocation using stable matching and also considers on-line (dynamic) versions. Other applications include taxi scheduling (Bai et al., 2014), allocating CPU resources (Wang et al., 2015) and more. Khuller et al. (1994) considers stable on-line matching and shows that randomised algorithms would result in at least $\Omega(n^2)$ blocking pairs. Lee (1999) experiment with windowing the incoming requests, so they are not immediately matched, but wait until window closes and potentially more agents or resources are available.

1.4.4 Other matching problems

There are many other matching settings, see books (Mehta, 2013b; Manlove, 2013) for a more extensive treatment. We will review two additional types and highlight some of their applications.

The stable roommates problem is a non-bipartite generalisation of the stable matching problem (Manlove, 2013). There is just one set of agents and each agent has a preference list over other agents in the same set. The goal is to find a pairing of agents that could for example share the same room in a college campus. The notion of stability can be defined in a similar manner as we did for two-sided markets. It turns out that stable roommates problem might not admit a stable matching (e.g. Manlove, 2013, p. 33), but if a stable matching exists we can find it in $\mathcal{O}(n^2)$ time Irving (1985).

The obvious application of stable roommates problem is allocating college students to share rooms on campus. Recently this has also been studied involving kidney exchange markets. In this market a patient in need of a kidney might have a donor, who is willing to donate, but is incompatible. The kidney exchange market can help find for a swap another patient, who also has an incompatible donor. Here preference lists would be constructed by compatibility between patients and donors. Allocation schemes exist in many countries, for example UK (Manlove and O'Malley, 2012), Netherlands Keizer et al. (2005) and US (Roth et al., 2007; Beard et al., 2012; Dickerson et al., 2012). There is also a similar organisation in Estonia (EKOTU, 2016), but how the actual matching is organised is unclear. (Manlove, 2013, p. 37-38) highlights few other applications of the stable roommates problem.

Another more prominent matching problem is the house allocation problem (Manlove, 2013; Schummer and Vohra, 2007). There is a set of agents and a set of houses. Each agent has a strict preference over houses and each agent owns a unique house. A goal is to reallocate the houses in a better way. Whereas in the two-sided market both sides had preferences, here houses do not have preferences over agents. A Pareto optimal matching of houses to agents is returned by the Top Trading Cycles (TTC) algorithm. The TTC algorithm is also strategy-proof for all the agents (e.g. Schummer and Vohra, 2007, p. 254). An important modification of the TTC algorithm (Abdulkadiroğlu and Sönmez, 2003) allows us to use this for school choice settings. The result of the TTC can be a better match for the applicant, however the resulting allocation might not be stable, but is still strategy-proof for the applicants. Potentially unstable allocation might explain why it is not widely in use. Although this has also been studied for kidney allocation (Roth et al., 2004).

In school choice setting we have preferences from one side and priorities from the other, which are not directly influenced for the other side, but are used for prioritising. The \mathcal{A} -optimal stable matching from Table 1.14 is $\mu = \{a_1 - b_1, a_2 - b_3, a_3 - b_2\}$. We see that agents a_1 and a_3 would be both better off if they exchanged their matches, they would both obtain their first preference instead of the second: $\mu = \{a_1 - b_2, a_2 - b_3, a_3 - b_1\}$. This of course assumes that we have priorities not preferences on side \mathcal{B} . But as a result the matching is no longer stable, as the pair $a_2 - b_1$ is a blocking pair and agent a_2 was not even involved in the exchange between a_1 and a_3 . So allowing these improving exchanges might create unstable matchings.

1.5 Limits of the axiomatic approach

We saw that auctions were mostly analysed in a homogeneous setting. For example analysing equilibrium strategies, agent were drawn from the same distribution and acted as in equilibrium. However, when agents are heterogeneous in their behaviour and preferences, it is harder to analytically understand potential outcomes. Often there are multiple equilibria and it is hard to predict which will prevail and in a dynamic setting there might be punctuated equilibria.

Computational agent-based models aim to make more general assumptions and study markets and economies in a less idealised situations. Having heterogeneous and interacting agents whose behaviour is based on heuristics rather than optimal rational choices. With the computational power these models have the potential to make predictions more relevant for human agent societies.

We survey two settings, combinatorial and double auctions, to show how agent-based experiments can augment the axiomatic results.

1.5.1 Combinatorial auctions and limits of theory

Combinatorial auctions consider complementaries or substitutes in agents' valuation. That is for some items S and T the value of obtaining these two items might be (Blumrosen and Nisan, 2007):

- lower than the total value $v(S \cup T) < v(S) + v(T)$ of just obtaining S or T , that is the items are substitutes
- higher than the total value $v(S \cup T) > v(S) + v(T)$ of just obtaining S or T , that is the items are complementary

This makes finding a solution much harder in many ways. First how to find the welfare maximising allocation, which is computationally NP-hard even for simple cases (Blumrosen and Nisan, 2007). Second collecting information on bidders' valuations for packages, which might be exponentially large. Third how to analyse strategic behaviour?

Lehmann et al. (2011) discusses the computational problem and find that in addition for the problem to be NP-hard it is also inapproximable in general. Intuitively we can think of having some set of k items there are 2^k subsets we would need to evaluate to determine the efficient solution. They even find it being NP-hard, when using a particular bidding language, like the OR and XOR.

How should bids be submitted to combinatorial auctions? In single-item auctions bidding was *easy*, the agents only needed to submit one bid. In combinatorial auctions each bidder should in a naive case submit 2^k bids for k items. In other words the bidder should evaluate each possible bundle of items. Even in an auction where the number of items is small this could amount to significant number of bids, eg when $n = 10$, $2^{10} = 1024$ bids, which is infeasible for the bidder to express or evaluate. A better way is required. We present a short overview from Nisan (2010).

Definition 17. *Simple atomic bids.* Each bidder submits a pair (S, b) , where $v(S) = b$ the price the bidder is willing to pay for a subset S of items. Meaning that for each $T \subseteq S$ $v(T) = p$ and otherwise $v(T) = 0$.

Definition 18. *OR bids.* Each bidder submits a number of atomic bids (S_i, b_i) , where $v(S_i) = b_i$. The bidder is willing to obtain any subset $T = S_i \cup S_j$ where $S_i \cap S_j = \emptyset$ for the combined price $p_T = v(S_i) + v(S_j)$. The bids/values for items can be computed by

$$v(S) = \max_{S_1, \dots, S_i \in S, S_1 \cap S_2 = \emptyset, \dots, S_1 \cap S_j = \emptyset, \dots, S_{j-1} \cap S_j = \emptyset} \sum_{i=1}^j v(S_i)$$

For example a bidder submits an OR set of atomic bids (1.4).

$$\{(s_1, 1), (s_2, 2), (s_3, 4), (\{s_2, s_3\}, 7)\} \quad (1.4)$$

Then the values for additional subsets would be:

- $v(\{s_1, s_2\}) = \max(v(s_1) + v(s_2)) = \max(1 + 2) = 3$
- $v(\{s_1, s_2, s_3\}) = \max((v(s_1) + v(s_2) + v(s_3)), (v(s_1) + v(\{s_2, s_3\}))) = \max(1 + 2 + 4, 1 + 7) = 8$

OR bidding language can represent additive valuations, for which $v(s \cup T) \geq v(S) + v(T)$, where $S \cap T = \emptyset$, and not substitutable valuations.

Definition 19. *XOR bids.* The bidder still submits an arbitrary number of atomic bids (S_i, b_i) , but he is willing to obtain just one of the subset of items. The bid for a set of items is calculated as

$$v(S) = \max_{S_i \subseteq S} b_i$$

XOR bids can represent all valuations Nisan (2010), although some of the additive values may require an exponential number (2^k) of bids.

For example a bidder submits a XOR set of atomic bids (1.5).

$$\{(s_1, 1), (s_2, 2), (s_3, 4), (\{s_2, s_3\}, 5)\} \tag{1.5}$$

Then the values for additional subsets would be:

- $v(\{s_1, s_2\}) = \max(v(s_1), v(s_2)) = \max(1, 2) = 2$
- $v(\{s_2, s_3\}) = \max(v(s_2), v(s_3), v(\{s_2, s_3\})) = \max(2, 4, 5) = 5$
- $v(\{s_1, s_2, s_3\}) = \max((v(s_1), v(s_2), v(s_3)v(\{s_2, s_3\}))) = \max(1, 2, 4, 5) = 5$

With XOR the value for $\{s_2, s_3\}$ is $v(\{s_2, s_3\}) = 5$, because the items are partially substitutable, whereas with the OR language the bid would had been $v(\{s_2, s_3\}) = \max(v(s_2) + v(s_3), v(\{s_2, s_3\})) = \max(2 + 4, 5) = 6$, because the bids would had been considered additive.

Obviously to make the bidding simpler we can consider a combination of OR and XOR bids to present additive and substitutable valuations. Nisan (2010) covers some simpler cases of combining OR and XOR bids.

From section 1.2.3 we know that the only way for strategy-proof value elicitation is to use Groves mechanisms. This takes us back to the computational complexity. We could need to compute n alternative efficient allocations to compute the payments and in combinatorial auctions already one of these was hard. However, alternative mechanisms would induce some sort of strategic behaviour.

Restricting the size of bundles is an option for limiting the complexity (Lehmann et al., 2011). Here computational experiments with various parameters are used to determine the auction format. An et al. (2005) defines some potential strategies for the bidders and vary the number of bidders submitting bundle bids and investigate the effect these have on the auctioneer's revenue. They find that more bundle bids increase the revenue of the auctioneer and allocative efficiency. Sureka and Wurman (2005) and Mochon et al. (2009) use meta-heuristics, genetic algorithms and tabu search, to find Nash equilibrium strategies in combinatorial auctions. This is similar to the evaluation problem – value queries to the bidding language. Given

some subset S of items determine $v(S)$. Value queries can be as hard as optimization problems for welfare maximising allocations. The intuition is that to compute the value in the OR bidding language, we would still need to evaluate many potential ways how to subset the items to calculate the sums, over which we would need to take the maximum value.

If bidders have some intuitive sense of the value functions it needs to be experimentally evaluated how to best transform it to actual bids. In some cases it might be easier when the bidder have some linear or quadratic additive value function, how to best transform it to a given bidding language. In some cases the bid intervals might be given as in energy markets (Contreras et al., 2001).

Calculating social welfare maximising allocation is hard in combinatorial auctions. If we consider some very general bidding language we usually need to evaluate n^k potential allocations, where n is the number of bidders and k the number of single items. Leyton-Brown and Shoham (2010) have developed a generator package to test winner determination in combinatorial auction. Holte (2001) run computational experiments for finding some expected, statistical, welfare guarantees. They use multidimensional knapsack problems to test for allocative efficiency. Schwind et al. (2003) uses similar meta-heuristic approach and also compares some existing deterministic formats. While others (Sandholm, 2010) have designed optimal allocations algorithms for some instances.

Some straightforward implementations for combinatorial auctions might turn out to be very bad. Sequential auctions are a bad idea. For example when there is a number of perfectly substitutable items for sale and we would allocate them in n sequential auctions. In the first second-price auction the price would be end highest, in the next auction thirds highest and so on. And at the last auction the price would be lowest, so bidders would always wait to bid in later auctions. So always bidding true value is not the best idea. See Swiss example by (Cramton, 2010, p. 102).

Additionally, sealed-bid auction on individual items is not a good idea as items have substitutes and complements and good package bidding is not available. It would be hard to figure out for which single items to bid, as winning any particular item is probabilistic, a bidder would never know if he would end up with a desirable package. And we already know why bidding for packages to not feasible.

In US and other countries have used the Simultaneous Ascending Auction (SAA) to allocate spectrum licenses to telecommunication providers (e.g. Klemperer, 2004b). This is not necessarily efficient or strategy-proof (Kagel et al., 2010). The SAA broadly works as follows:

- The seller (auctioneer) decides how to split the item for bidding. This is important decision as it will affect how bids are placed

- Buyers then bid for the items and may bid on multiple items (licenses) simultaneously. As for each bidder the aim is to cover some geographical area and usually is differs across bidders.
- The first round with no new bids ends the auction

Wellman et al. (2008) highlights bidding strategies that could be used in SAA, but also state that they are not optimal, which indicates the hardness of analysing the SAA mechanism. They also show that the near-optimal strategy is dependent on the structure of the environment, as characterized by complementary and substitutable valuations.

New types of *practical* problems arise in SAA. Bidders may not want a package of items that are missing some crucial items or package of items that are considered duplicates by the buyer. To avoid these problems usually some additional rules are defined in SAA:

- Activity rule: throughout the auctions the activity of the bidders can only decrease. As the prices will go up buyers are less willing to obtain the item.
- As some aggregations might fail, bidders are allowed to withdraw bids. This creates strategic behaviour as bidders are not responsible for their highest bid. Is usually regulated with fines to the high bidder after withdrawal
- Bidders often have a quantity cap to the number of items they can obtain
- Bidders usually have some initial down payment to ensure they are really interested in the auction. This is refunded depending on the final allocation
- Bids are incremented not freely placed by the bidders

The activity rule contributes to the *price discovery* process for all items. As bids increase buyers can decide on which items to continue bidding and for which items to stop. Additionally determining bids for all potential packages of items is hard, so price discovery will help in figuring out for which packages it might be profitable to bid and for which not. The specific packages of interest can potentially change throughout the auction.

Signalling other bidders to keep the prices down. If there are two items that are generally considered complements it is profitable for two buyers to agree among each other who will get which item, as they will not compete for the same item, thus the price will not increase. This has happened in the past. In one auction a winning bid was \$47,505,673, surely the buyer would not care up to a dollar how much he would pay. In another auction

a bidder made several bids ending with 378, outbidding another bidder several times. The bids seem to have been retaliation for bidding in area 378 (see footnote in Milgrom, 2004, p. 267).

Inefficiencies can arise in a SAA type auction. Suppose there are two parking spaces to be allocated. Bidder 1 has a trailer desires both places with $v_1(\{a, b\}) = 100$. The second bidder 2 values both spots at the same price $v_2(a) = v_2(b) = 75$. The efficiency maximising allocation would be to give both places to the first bidder. If we look at the SAA process, suppose the bids are \$50 for each of the spots and the first bidder is the high bidder. His utility in the case of winning would be 0. However for bidder 2 there is an incentive to increase any of the bids and would pay the bid. In case of package bidding and VCG the spots would be both allocated to bidder 1, for the price of \$75 (see example in Cramton, 2010, p. 102).

This has also been referred to as the exposure problem. To win both items bidder 1 would have to bid \$150, thus creating negative utility. When the current high bids for both items are \$40 and say the bid increment is \$10. The second bidder would increase the bid for one spot to \$50 and the second bidder does not have incentive to raise his bid anymore, because his utility would become 0. Nevertheless he is stuck with one of the spots for the price of \$40, which alone does not provide any utility. Maybe bidder 1 should not have bid for both spots at \$40, but given his limited knowledge of second bidders valuation (it might be \$40), he still has the potential to gain some positive utility.

Demand reduction involves not bidding on a second or a third substitute, so as not to increase bidding by other on the first item. In multi-unit auctions the demand reduction can be achieved by collusion, e.g. agreeing who will bid on what item, so the price will stay low. Suppose there are two identical items a and b and two bidders, with both valuing a single item at \$100. How would they bid? In a competitive environment both would increase their bids until it would be close to \$100 on both items, say \$99. Now if bidders would be able to communicate they would agree not to increase the bids on other items. Both would win the auction at the starting price (e.g. Cramton, 2010, p. 107).

Experimental papers (Kagel et al., 2010, 2014) compare two feasible mechanism for combinatorial auctions: combinatorial clock auction (CAA) and SAA. As theoretically analysing these mechanisms is hard, they use a combination of computational and human experiments to compare the two auction formats. Initially (Kagel et al., 2010) they find that SAA might provide more revenue than CAA, while in some situations having lower efficiency. Later (Kagel et al., 2014) that human bidders do not deal well with the *package selection* problem. Humans bias towards other signals, for example how the packages are named. This caused humans to behave

irrationally in CAA, they did not bid on more profitable packages (Kagel et al., 2014). When package names match human bidder valuations the outcome of CAA is efficient, but when not then SAA is more efficient.

1.5.2 Double sided auctions and complex systems

Single-item auctions are most popular, applicable when there is one unique item for sale, like a work of art. Combinatorial auctions arise when there are multiple different items for sale, which might be complementary or substitutable, like spectrum or bus routes etc. In a situation when we have multiple units of identical items the most often used mechanism type is the double auction. This is because in this situation there usually are also multiple sellers and buyers. In the case when there are only buyers, the VCG or some modified single-side auction type could be employed.

Double sided auctions arise when there are multiple identical items for sale and multiple sellers and buyers. The most frequent setting is buying and selling stocks, commodities like wheat, coffee etc. More recent applications include energy pricing (e.g. Nicolaisen et al., 2001; Faqiry and Das, 2016) or allocating cloud and communications resources (e.g. Ji and Ray Liu, 2006; Wang et al., 2010; Han et al., 2011b).

In a double auction there is a set of sellers \mathcal{S} and a set of buyers \mathcal{B} . To simplify we assume that all sellers are selling one unit of identical goods each and buyers are interested in buying one good each. All sellers have some cost, reserve price, c_i associated with good they are selling, below which they would not sell. Similarly each buyer has a valuation v_i for a unit of good, above which they would not obtain the good. The goal of the double auction mechanism is to find a price p that would realise all trading opportunities. The resulting utilities of the agents are $u_i = p - c_i$ for sellers and $u_j = v_j - p$ for buyers.

A seminal paper in designing a truthful static double auction is by McAfee (1992). They design a strategy-proof mechanism, however note that it is not necessarily efficient. McAfee (1992) mechanism assumes multiple buyers and sellers (≥ 2) otherwise the mechanism does not work. They analyse the efficiency of the mechanism via computational experiments, as the analytical form is uninformative in this regard.

A more general setting is the k-double auction, however there is no closed form solution to finding equilibria in this setting (Satterthwaite and Williams, 1993). Satterthwaite and Williams (1989a, 1993) analyse the k-double auction. The price in this mechanism is determined by $p = kb + (1 - k)a$, where $k \in [0, 1]$ and $[a, b]$ is the interval in which the market-clearing price is selected. After the price is determined, all agents with a

positive utility trade with the clearing price p . Satterthwaite and Williams (1989b) also conduct simulations to understand the efficiency of their auction format.

For example, assume we have four sellers and four buyers. Each agent is interested in buying or selling one unit of the good. Sellers have costs $\mathcal{C} = \{0.39, 0.43, 0.49, 0.64\}$ and buyers have valuations $\mathcal{V} = \{0.11, 0.32, 0.64, 0.71\}$. The respective supply-demand curves are presented in Figure 1.4.




Figure 1.4: Supply-demand curve (k-double auction mechanism)

In the k-double auction market clearing price is in the interval $[0.43, 0.49]$ and at most two trades can be realised. The interval is determined by ordering the costs and valuations, $s_{(i)} \in \mathcal{C} \cup \mathcal{V}$, in an ascending order $s_{(1)} \leq s_{(2)} \leq \dots \leq s_{(2m)}$ and setting the interval to $[s_{(m)}, s_{(m+1)}]$. McAfee (1992) uses a different method for selecting, but essentially yields a similar interval.

The cases when $k = 1$ and $k = 0$ are special in terms of incentives. In 1-double auction the price equals $s_{(m+1)}$, thus the sellers have no incentive to misreport their costs as the price is already at its maximum. If sellers currently below 0.49 would individually increase their reported cost above 0.49 they would not trade and would gain zero utility. Otherwise they would gain positive utility. Similarly in 0-double auction the price is $s_{(m)}$ and the buyers have no incentive to misreport. In case the price is in the interval $(0.43, 0.49)$ all agents have an incentive to misreport to some degree.

Satterthwaite and Williams (1993) show the form of symmetric Bayes-Nash equilibrium reporting strategies, assuming all agents know the distributions \mathcal{F} and \mathcal{G} from where respectively sellers and buyers valuations are drawn.


Theorem 8. (Satterthwaite and Williams, 1993, Theorem 1) Consider any equilibrium (f_S, f_B) in which trade occurs with positive probability, every seller always asks as much as his cost, and every buyer bids at most his value. A number κ exists, whose value is a function of \mathcal{F} and \mathcal{G} , but not of m or (S, B) , such that, for all $v \in (\underline{v}, 1]$ and $c \in [0, \bar{c}]$,

$$S(c) - c \leq \frac{\kappa}{m} \quad (1.6)$$

and

$$v - B(v) \leq \frac{\kappa}{m} \quad (1.7)$$

Also $\kappa \propto k$. There is also a certain region where trades do not occur due to misreporting, since sellers Bayes-Nash reporting strategy is above buyers' value and similarly buyers reporting strategy is below sellers' price. These values are \underline{v} and \bar{c} respectively. Figure 1.5 illustrates the equilibrium strategies and no-trade regions. The centre square has a positive probability for a trade, when the cost and value are suitable. On the borders trades never occur, even when the actual cost and valuations would be suitable, as the equilibrium reporting would exclude it. Even when the actual cost is close to zero, the seller would report \underline{v} , so if seller valuation is below this there is no trade. Similarly the seller would never report more than \bar{c} . From (1.6) and (1.7) we see that the misreporting decreases as the number of agents in the market increases in proportion to $\mathcal{O}(\frac{1}{m})$.



Note: Adapted with permission from Elsevier.

Figure 1.5: Equilibrium strategy pair (f_S, f_B) (Satterthwaite and Williams, 1993).

In reality the common assumptions required by the Bayes-Nash equilibrium concept might not be satisfied. The assumption was that all agents have same prior information about the distribution of private costs and values of each agent. Satterthwaite and Williams (1993) show that by slightly

modifying the the allocation mechanism the equilibrium strategies can vary widely, but Kagel and Vogt (1993) show that in human experiments there is no significant change in strategies.

For market efficiency we look at the total utility gain by the traders, agents who actually trade. Since the trade units have to be equal, the number of trading buyers and sellers is the same.

$$TU = \sum_{i \in \mathcal{S}} (p - c_i) + \sum_{i \in \mathcal{B}} (b_i - p) = \sum_{i \in \mathcal{B}} b_i - \sum_{i \in \mathcal{S}} c_i \quad (1.8)$$

Let's denote by $TU(f_S, f_B)$ the total utility gain, when agents play their equilibrium strategies and by $TU(f_S^*, f_B^*)$ the potential (optimal) gain when agents would reveal their true costs and valuations.

Definition 20. We define relative efficiency ρ in double auction as

$$\rho(\mathcal{S}, \mathcal{B}) = \frac{TU(f_S, f_B)}{TU(f_S^*, f_B^*)}$$

This definition is the inverse of a similar definition of Price of Anarchy (e.g. Roughgarden, 2005).

Theorem 9. (Satterthwaite and Williams, 1993, Theorem 2) Consider any equilibrium (f_S, f_B) in which trade occurs with positive probability, every seller always offers at least as much as his cost, and every buyer bids at most his value. A constant ξ exists, whose value is a function of \mathcal{F} and \mathcal{G} , but not of m or (f_S, f_B) , such that the relative efficiency in the k -double auction is at least

$$1 - \frac{\xi}{m^2}$$

So when the market is large, then the relative efficiency approaches one.

Conversely to the k -double auction, McAfee (1992) first discovered a strategy-proof static double auction mechanism. The McAfee-mechanism works as follows. There are m buyers and n sellers. Order the reported bids and asks:

$$b_{(1)} \geq b_{(2)} \geq \dots \geq b_{(m)} \quad (1.9)$$

and

$$c_{(1)} \leq c_{(2)} \leq \dots \leq c_{(n)} \quad (1.10)$$

Select $k \leq \min\{m, n\}$, such that $b_{(k)} \geq c_{(k)}$ and $b_{(k+1)} < c_{(k+1)}$. Define $p_0 = \frac{1}{2}(b_{(k+1)} + c_{(k+1)})$.

1. if $p_0 \in [c_{(k)}, b_{(k)}]$ the all buyers with $b_i \geq p_0$ and sellers with $c_i < p_0$ trade




Figure 1.6: Supply-demand curve (McAfee mechanism)

2. otherwise if $p_0 \notin [c_{(k)}, b_{(k)}]$ then only buyers and sellers (1) through (k-1) trade. Buyers pay $b_{(k)}$ and sellers receive $c_{(k)}$ and leaving a surplus of $(k - 1)(b_{(k)} - c_{(k)})$.

The example in Figure 1.4 would yield $p_0 = 0.405$ as in Figure 1.6.

Theorem 10. (McAfee, 1992, Theorem 1) *Honesty is a dominant strategy for the McAfee-mechanism*

Theorem 11. (McAfee, 1992, Theorem 3) *The expected efficiency loss due to a trade not being executed is $E[b_{(k)} - c_{(k)}] \propto \frac{1}{n}$*

In addition there is a utility loss due to earnings from the mechanism, which increases in the number of trades. McAfee (1992) also investigates the loss to the market due to the mechanism's earnings and finds it analytically difficult. With uniform distributions of costs and values McAfee (1992) finds using computational experiments, when $m = n$, that the mechanisms earnings also increase in the number of agents and already when $n = 1000$ the loss percentage due to trades not executed is only 0.0002%.

There are also modifications of the McAfee-mechanism that distribute the mechanism's earnings, from case 2 of p_0 , in the market, but still leaves some potential trades on the table (e.g. Segal-Halevi et al., 2016, and survey therein). There are also extensions to handle heterogeneous single-unit items Feng et al. (2012). There is even a continuous on-line truthful double auction design, but with an efficiency loss of no more than 20% (Wang et al., 2010).

A continuous double auction has an even richer strategy space and this game has proven extremely hard to analyse (Rust et al., 1993, p. 158). In a continuous double auctions buyers and sellers constantly submit their

offers to the mechanism and a trades occur immediately when the offers match. Some offers might stay in the market longer, if no suitable submission arrives from the other side. Rust et al. (1993) describes an experiment of continuous double auction strategies in the Santa Fe Institute in 1990. They collected some 30 human defined strategies that competed against each other. The strategy-set also included simple strategies like truth-telling and zero-intelligence¹, which is essentially playing randomly and is only constrained by its reservation price. They find that many of these strategies perform poorly and are easily exploited, although they guarantee market efficiency. A simple, Kaplan (Algorithm 1 for buying), strategy that always plays the current market price or reserve price, which ever produces positive utility, turned out to be the best. This strategy essentially waits until the market is about to close and then jumps in and makes its bid.

Algorithm 1 Kaplan buy strategy automata

Require: $p_{max}, p_{min}, bid_c, ask_c, p_{minc}, \pi$

Ensure: $bid \in \{\emptyset, \mathbb{R}^+\}$

$bid = \emptyset$

if $bid_c \neq \emptyset$ **then**

if $ask_c \neq \emptyset$ **then**

$m = \min(ask_c, \pi - 1)$

else

$m = \pi - 1$

end if

if $(m > bid_c)$ **and** $((ask_c < p_{max}$ **and** $\frac{u(.)}{\tau} > 2\%$ **and** $\frac{ask_c - bid_c}{ask_c} < 10\%)$
or $(ask_c \leq p_{min})$ **or** $(\delta t < \tau)$ **then**

$bid = \min(ask_c, m)$

end if

end if

return bid

Rust et al. (1993) also conducted an evolutionary game of the strategies, where survival of a strategy was proportional to their capital share of the market. They found that the Kaplan strategy dominated in this game as well. However, after being the only surviving strategy, the profits declined. Since Kaplan strategy awaits in the market until other strategies do the negotiating and when these strategies are extinct there is almost no trading. The Kaplan strategy activates at the end of the trading period and only few profitable are executed. The leave open the question: are there strategies

¹not to be confused with zero-knowledge, a concept from cryptography, a method to prove that a statement is true, without revealing anything other than the truth-value of the statement

and environments that are resistant to exploitation by Kaplan strategies (Rust et al., 1993)?. This seems even more relevant when looking at the exploitive nature of high-frequency traders (Budish et al., 2015).

Chen et al. (2009) and Chen and Tai (2010) compare the strategies evaluated by Rust et al. (1993) with evolutionary programming to find an even better strategy using learning with genetic programming. They train an even better strategy by randomly playing against strategies in the Santa Fe experiment (Rust et al., 1993, 1994). Chen et al. (2009) uses genetic algorithms to find a function based on trading statistics from previous trading day, current state of the order-book, timing and agent's values. Additionally using mathematical functions like addition, multiplication, absolute, maximum, minimum and if-then rules to combine the statistics. They find that using minimum bid from previous day is most often used in a profit maximising strategy. With larger *gene* population more complex strategies emerge combining the minimum bid from previous day with highest value the agent has for an item. The simplicity of this learned strategy is similar to the Kaplan strategy and can even outperform it.

Zero-intelligence models have been shown to be enough to guarantee almost 100% efficiency in continuous double auction markets (Gode and Sunder, 1993a,b). Gode and Sunder (1993b) investigate a *synchronized double auction*, where all bids and asks are collected. If highest bid and lowest ask cross, then they are traded in a binding transaction. Otherwise new asks and bids are asked from traders. Depending on the number of extra-marginal traders, traders who would not trade in an ideal truthful market, the efficiency is between 80%-100% and 95.7% on average as the ratio of extra-marginal traders goes to infinity. Some 50%-100% efficiency is achieved already in the first round. Gode and Sunder (1993a) run additional human and computational experiments with constrained and unconstrained zero-intelligence (ZI-U) trades. The ZI-U agents are not anymore constrained by their reservation value and can bid above or below it, depending if they are buyers or sellers. With the constrained ZI agents they obtain often 99% efficiency, while with unconstrained only 90%. Moreover with human experiments the efficiency is also close to 99%. Gode and Sunder (1993a,b) results suggest that individual rationality is not required to the extraction of market surplus.

As a results there has been more works on theoretical statistical models of continuous double auction (Chiarella and Iori, 2002; Smith et al., 2003; Daniels et al., 2003) and empirical (Madhavan, 2000; Farmer et al., 2005; Hasbrouck, 2007) to model aggregate properties like the bid-ask spread, price impact function, probability of order execution, price diffusion, volatility etc.


Almost all models of strategies assume agents know their own independent fundamental value v_i and based on this determine their pricing strategy. However many profiting strategies in double auctions do not require the knowledge of fundamental value. For example assuming that a price of a security is predicted to increase or decrease an agent can profit from arbitrage, by buying with a low and selling with a higher price. Thus the strategy-space of a continuous double auction is inconceivably large. Thus the price discovery or information dissemination is a critical part in understanding such a market mechanism. Therefore the human component is crucial part in understanding financial markets and computational experiments are a useful modern tool to investigate aggregate results from the behaviour (e.g. Friedman and Rust, 1993; Tesfatsion and Judd, 2006).

Both laboratory experiments and theoretical models show that agents' behaviour – and hence market outcomes – are highly sensitive to the assumed information structure (Madhavan, 2000, p. 207)

The financial market mechanism or market micro-structure design is more concerned with traders behaviour rather than selecting an *optimal* pricing rule. The design components include questions like: Should trades be mediated by a market maker? How to offer greatest liquidity? How to minimise trading costs? How traders react to price changes and how are prices disseminated? See (e.g. Madhavan, 2000) for further survey.

Hommes (2006) reviews some small heterogeneous agent-based models. The simplest model is with fundamentalist and chartist agents. First type of agents always have some individual fundamental valuation v_i for a good, as with previously described models, and make bids and asks according to that value. The other type of agents, chartist, behave by monitoring the market, they buy when the price increases and sell when the price falls. This behaviour composition already creates crashes and rapid increases and declines in stock prices. Additionally (Sornette, 2004, p. 132) the chartists are relatively successful free riders, not only matching the performance of fundamentalists in the long run, but outperforming them in the short run. Similarly to the Kaplan strategy mentioned previously.

Albin and Foley (1992) present a simple benchmark model for a decentralised exchange. In this model agents randomly broadcast messages indicating a willingness to trade. However, there is a cost associated with the signal. If the buyer's price is higher than the seller's asked price then an exchange occurs. In case of multiple potential trade partners, the agents meets sequentially in a random order with each partner. A subsequent trade occurs only if the current holdings permit. The price is determined by geometric average of the bid and ask prices. The agents know a lot



Note: Adapted with permission from Elsevier.

Figure 1.7: Exchange market efficiency (Albin and Foley, 1992).

about their potential trade partners: cost of advertising; meeting history; marginal rates of substitution of neighbours; distribution of goods among agents. Based on this information they estimate the expected gain from signalling and signal if there is an expected profit. They observe that with this model the market efficiency also converges to almost 100% with different messaging costs and number of targeted agents. On Figure 1.7 we have created the chart based on data from from (Albin and Foley, 1992) for two combinations of parameter values. In both cases the final allocation efficiency is very close to 100%. A similar model is, for example, employed in Beltratti and Margarita (1993).

LeBaron (2006) surveys additional centralised models and learning in a two-sided exchange. Most models concentrate on investigating some form of centralised allocation. Most experiments have some centralised market mechanism and adaptive agents mostly based on genetic algorithms and occasionally some other learning rule (LeBaron, 2006, p. 1223).

Marks (2006) reviews the main reasons to use computational agent-based modelling. First the analytical model is intractable. It is hard to obtain equilibrium solutions to some mechanisms, e.g. double auctions, due to the dimensionality of the task. Second the equilibrium behaviour is not relevant for modelling the real world and we are interested in out-of-equilibrium behaviour. In stock markets there is lot of dynamics in the price and the market almost never manages to converge. Third agents are not completely rational, but are bounded and make quick and myopic decisions. Additionally they learn market conditions and from past behaviour that guide their future decisions.

Marks (2006) also highlights a general analysis-design framework for mechanism design. Analysis as surveyed in previous section is important first step, but computational experiments, simulations, provide light were the axiomatic approach cast little or none. Marks (2006) concentrates on evolutionary and learning techniques. First which strategies would survive

in some environment or which strategies emerge as a result of adaptation? Both of these are important to understand the effect of picking a particular design and making sure the mechanism responds well to possible adaptation. In addition computational experiments are useful to propose behaviour model in some situation, e.g. in decentralised exchange (Albin and Foley, 1992) and zero-intelligence (Gode and Sunder, 1993a,b; Farmer et al., 2005).


2 Zero-Intelligence in Decentralised Matching

2.1 Introduction

Market economies in general experience large employment fluctuations and average unemployment rates differ between countries. The underlying job search and matching theory (the Diamond-Mortensen-Pissarides canonical framework) provides a conceptual explanation for some aspects of the relationship between vacancies and unemployment known as the Beveridge curve (Figure 2.1). The core of job-search and matching models is built on the assumption that the external rate of job creation and destruction, but also worker reallocation, determine the steady-state equilibrium of number of unmatched workers and jobs (unassigned agents in our model) (Mortensen and Pissarides, 1999). Because of search and recruiting costs, hiring and firing costs and other forms of matching-specific costs, decentralised markets create inefficiencies. This matching technology is implicitly characterised by its matching function, which summarises the trading technology between agents, their actions and choices that eventually bring them together into productive matches (Petrongolo and Pissarides, 2001). In the relevant “matching function literature”, it is stressed that such a theoretical tool is useful because it allows to reduce the complexity of information imperfection, heterogeneous agents and congestion into a tool-kit similar to the production function or money demand function. However, the interactions or matching technology are still rather treated as a black box.

We open this black box by providing a simple agent-based model of a decentralized market game in which agents (workers or job seekers) make proposals to the agents on the other side of the market (firms) in order to be matched to available positions. In our computational experiments, a market game is identified by three components: the preference structure of agents, market conditions, i.e. the relative number of positions and workers, and the behaviour of agents (workers or firms) based on the information they have about their own preferences and options in the market.

Search models in labour economics rely on three pillars: the decision of workers, the decision of firms and wage setting mechanism. We concentrate on the first two pillars. Thus, our model belongs to the literature



Source: Eurostat (online data code: JVS_Q_NACE2 and UNE_RT_Q)
 US Bureau of Labor Statistics (online data code: JOLTS and LAUS)

Figure 2.1: Beveridge curve based on empirical data

on agent-based partial labour market models (e.g. Neugart and Richiardi, 2012) which use microsimulation to explain stylised facts about the labour market. These agent based “micro-to-macro” models give insights into labour markets in the form of partial or general models. In the latter labour markets are embedded into larger economic models. We are mainly interested in literature aiming to reproduce the Beveridge curve with search and coordination in a partial agent-based model. Thus, models (Richiardi, 2004, 2006; Riccetti et al., 2015) aimed at explaining the job search on the basis of wages or more general market models that interact with the embedded labour market (Dawid et al., 2014; Deissenberg et al., 2008), that lie outside of our scope.

The partial agent-based models have been developed for replicating stylised facts from real labour markets, such as the negative-sloped Beveridge curve in the unemployment-vacancy (u, v)-space. Fagiolo et al. (2004), followed later by Silva et al. (2012), reproduce a Beveridge curve in a partial agent-based model and come up with a standard explanation that frictions and the institutional setting affect the position of the curve.

Moreover, giving up the assumption of rational expectations about the behaviour of agents has produced fruitful insights. Tassier and Menczer (2008) investigate job-hunting via social networks. They find that random social networks spread vacancy information better and thus achieve lower unemployment. Neugart (2004) uses an urn-ball matching model on a small scale (30-50 agents) and endogenous matching function workers send applications randomly, however the Beveridge curve is closer to origin than one would expect in large markets. Similarly Richiardi (2004) employs a similar model with wages and produces a Beveridge curve further from the

origin. While Richiardi (2006) models labour supply in a general setting by a non-equilibrium, adaptive agent-based model of heterogeneous workers and firms, with on-the-job searching, endogenous entrepreneurial decisions and endogenous wage and income determination. The latter is able to reproduce a number of stylized facts generally accepted in labour economics and industrial organization, including the negative-sloped Beveridge curve. Also, in this set-up, the matching process is based on random applications from job seekers for vacancies in the single labour market. Furthermore, this model allows for on-the-job searching, meaning that assigned agents can get job offers as well.

There are some other search-and-match models (e.g. Gabriele, 2002; Deissenberg et al., 2008; Boudreau, 2010). Some of them also produce a Beveridge curve, but usually investigate different aspect of matching. Often their aim is to study some other aspects, like stratification and use different underlying assumptions (e.g. centralised matching or perfect information).

Our modelled agents can be considered myopic – they make proposals and only accept better proposals without additional strategic thinking. In our base model, agents apply for a random position and a matching occurs only when both sides find their new partner preferable to their current match. This is similar to the Zero-Intelligence (ZI) model from financial markets (Gode and Sunder, 1993a; Chen and Tai, 2010). Our behavioural models are an extended version of the better response dynamic proposed by Knuth (1997b) and further analysed by Roth and Vate (1990) and Ackermann et al. (2008).

For comparative purposes, we include behaviours, where agents know more than in the ZI model, but less than in the better response dynamic. In our Better proposal model, we assume that agents, e.g. workers, know of a better match or a position that would also be suitable for them and thus do not make proposals in a wholly random manner. This can be considered similar to the Zero-Intelligence Plus model (Chen and Tai, 2010). In order to further extend the information pool available to the agent, we use the Blocking proposal model, where agents make proposals to their random blocking partner, which is equivalent to the better response dynamic proposed by Knuth (1997b).

Nevertheless, the agents will not be able to find an equilibrium (stable) matching in a micro sense. It is rather a steady-state equilibrium in a macro sense, as the size (number of matched agent-pairs) of the matching converges. In general, we show how the assumptions about the information available to agents, the structure of their preferences and market thickness determine the shape and placement of the Beveridge curve. Thus, our approach not only differs from the framework of Diamond-Mortensen-Pissarides, but also from recent discussions on job search and matching

efficiency (Veracierto, 2011; Shimer, 2013; Sahin et al., 2014) where the persistent empirically observed adverse shifts (outward shifts of the Beveridge curve) have been explained mostly by the deterioration in the efficiency of the matching technology. We use a partial agent-based model as classified by Neugart and Richiardi (2012) to develop aggregate regularities from the micro-behaviour of individual agents in order to illustrate the position of the Beveridge curve in the (u, v) -space. Our model is able to reproduce some well-known stylised facts from labour market literature like the negative sloped Beveridge curve. Moreover, we explain the shift of the Beveridge curve not only as related to information available to players, but also including a level of heterogeneity of the agents' preferences. We model the latter on the basis of the correlation between the preference lists of agents to indicate their similarity (or common understanding of a good job position). In addition, we allow for different lengths of the preference lists. Short lists indicate that there are only few acceptable positions in the market, i.e. geographical boundaries or asset specificity could determine the structural characteristics of the market. We show that short lists and high correlation shift the curve to the upper-right, further from the origin. In addition, we include a parameter for market thickness — the balance between market sides — indicating the ratio of positions to workers. Market thickness models the effect of interactions with other markets through job creation and destruction rates, i.e. the out-of-equilibrium state of a general market model. We see that this determines the position on the Beveridge curve.

In contrast to the decentralised matching in some situations we can employ a centralised mechanism. Then all the market participants report their preferences to a central clearing house that can then compute an optimal matching using for example Gale-Shapley deferred-acceptance algorithm (Gale and Shapley, 1962). Optimal usually means that the result is the best possible stable matching for one side of the market as the optimality can not be guaranteed for both sides. Roth (2008) observed that in some decentralised situations where there was not a central clearing house in place, market participants still executed a very similar algorithm as proposed by Gale and Shapley (1962). A major drawback of the execution was that usually it was time-capped, i.e. at some point the market had to be closed. This meant that the algorithm execution might not have finished and resulting matching may not be stable.

In addition to comparing the decentralised behaviours, we include an idealised deferred-acceptance based clearing house. This helps us to understand how far the decentralised market is from the optimal solution. There have been a few related works involving agent-based computational experiments in centralised matching markets. Oméro et al. (1997); Dzierzawa and

Oméro (2000) study the scaling behaviour in the obtained rankings in stable matchings. Others extend this to more general preferences, by allowing correlation (Caldarelli et al., 2001; Caldarelli and Capocci, 2001). Zhang (2001) studies the effect of having limited preference lists used in optimal stable matchings. This limitation is included to an extent in our decentralised model. Laureti and Zhang (2003) also investigate a decentralised model, however they assume optimal behaviour from agents compared to our model, namely that the agents always make the best possible proposals.

Another aspect more studied in centralised deferred-acceptance matching in the ranking of the match. We look at the cost of the decentralised matching game for agents in terms of median rank from matched agents. We discover similar results in decentralised markets as have been discovered for stable matchings – that proposers and agents on the smaller side of the market obtain better matches (Pittel, 1989; Ashlagi et al., 2013b) – but less extreme. Surprisingly the noise, ZI , behaviour results in best median rank among our decentralised behaviours.

We view our decentralised market game as an abstract model which in addition to studying agent-based interactions in job search and matching markets can be used in alternative settings, e.g. decentralised school or university choice. In all these cases, an institution or a central authority – a clearing house or similar – is missing and agents on both sides constantly have to react to new proposals and responses.

We continue as follows. In the second and third sections, we introduce the set-up of our model, concentrating on matching behaviour which includes search and commitment costs as well as the assumptions behind preference formation. In the fourth section, we describe our parameter selection and initial results of a steady-state. In the fifth section, we look at when the decentralised matching results in a stable matching. In the sixth and seventh section, we look at the number of unmatched agents in thick and thin markets, the latter shows how a Beveridge curve depends on the search behaviour and the structure of preferences. In the eighth section we compare the allocated ranks with all the decentralised behaviours to a centralised clearing-house based on Deferred-Acceptance. Finally, we conclude by discussing our contribution to the discourse on search and matching literature.

2.2 Behaviour models

In order to translate the neoclassical matching function into an agent-based version, we employ the framework offered by Guerrero and Axtell (2011). This relies on three orders of assumptions, which include rationality, agent homogeneity and “non-interactivness” in an agent-based model. The categorisation of the features is presented in Table 2.1.

Table 2.1: Matching model assumptions

	Matching function	Our agent-based model
First order assumptions		
Interaction mechanism	Implicit (black box)	Explicit (decentralised or centralised matching)
Second order assumptions		
Rationality	No explicit reference. Functional form captures information imperfections	Explicit, exact notion of how proposals are made and accepted
Equilibrium	Is inefficient due to negative externalities: congestion, skill mismatch and locational differences	Is inefficient because of assumptions about behaviour and information.
Agent types	Representative (homogeneous) agent(s), heterogeneous sectors	Heterogeneous (different preferences)
Third order assumptions		
Technological advances	Technological advances in matching shift “Beveridge curve” (search is less costly)	No technological advances. Preferences are static, do not adapt.
Supply shocks or business cycles	Affect job creation and destruction rates, but are generally treated as an empirical question	No explicit job creation and destruction, but the market thickness, is of interest
Contracts	A rate of matching and unmatching	A relationship can be broken whenever a better match is found (<i>and the contract has expired</i>)
Transaction costs	A limitation on matching rate	Once a better match is found, there is no additional cost for changing the match

Beginning with the first order assumptions about the nature of the interaction mechanism, we can see that the literature about the matching function often treats this procedure as a black box. In an agent-based model, the interaction is a central question for the investigation. In financial double-auction markets, Zero-Intelligence (ZI) interaction models have been fruitful in investigating aggregate market phenomena (e.g. Gode and Sunder, 1993a; Farmer et al., 2005; Ladley, 2012). We employ a similar approach for modelling search in the labour market.

Zero-intelligence (Ladley, 2012) is useful because it allows us to decouple the behaviour of an agent from the market structure. Moreover, we are interested in whether similarly to non-strategic agents (e.g. Farmer et al., 2005; Gode and Sunder, 1993a), interesting market phenomena can be produced in the current context. To our knowledge, these types of models have not been studied for job search. There are macro-level studies that concentrate on modelling unemployment and vacancies (e.g. Petrongolo and Pissarides, 2001; Mortensen and Pissarides, 1999). There also exist agent-based models of wage equilibrium (e.g. Guerrero and Axtell, 2011) and job search on social networks Guerrero and Axtell (2013); Zhou et al. (2014a); Hoefler and Wagner (2012), but there is no simple model for job search.

The labour market consists of two sets of agents – workers and firms (or positions within a firm). The main behavioural aspect is how the match is initiated, i.e. the worker-position pair selection. We study models where the proposing power is either only on one side of the market (A-proposing models) or where it is proportionally shared (random agent proposing). In other words, either workers always make proposals, proposals are made interchangeably, or firms always make proposals, depending on who is considered side A. We call the non-proposing agent the responding agent. In the centralised Deferred-Acceptance markets (Gale and Shapley, 1962), the matching is always optimal and stable for the proposing side, while it is the worst possible level for the responding side (e.g. Knuth, 1976; Roth and Sotomayor, 1990). However, in many practical (Roth and Peranson, 1999) and large markets (Immorlica and Mahdian, 2015), the difference seems to be small and the effect of market thickness is much greater (Ashlagi et al., 2013b). The matched rank structure may also affect the size of the matching as proposing probabilities are different under random agent proposing.

In fact, we investigate several models (Table 2.2), where a second-order assumption of Zero-Intelligence (ZI) is a characteristic of the base model. The ZI model is called the Noise proposal model wherein two random agents, one from each side, are selected, but a matching transaction occurs only when the new match is an improvement over their current matches. Similarly, in financial markets a deal is only accepted when the offered price is above the reserve price for both sides, i.e. the buyer and the seller, oth-

erwise the price is offered at random. Thus, in our mechanism, there are only pairwise interactions and a transaction occurs when the reserve offer is met on both sides.

Table 2.2: Explored behaviour models

Rationality (2nd order)	Interaction, proposer (1st order)	
	Random side	A-proposing
Noise proposal	Noise Proposal	Noise Proposal A
Better proposal	Better Proposal	Better Proposal A
Blocking proposal	Blocking Proposal	Blocking Proposal A

In the first Zero-Intelligence Plus (ZIP) behaviour, which is the Better proposal model, agents only make proposals to a better match, i.e. a position higher on their preference list than their current (reserve) match. Thus, the proposing side does not even consider non-acceptable matches. In contrast, in the Noise proposal model an offer is made to a random agent on the preference list, and might actually not be acceptable to the proposer. This is learned in the transaction. In other words, in the Better proposal model, even agents with a current high ranking match have a high probability of a new match, which only depends on the responding side finding it acceptable. On the other hand, in the Noise proposal model a transaction probability would also be lower for proposing agents, if their current match is high on their preference list.

The second ZIP behaviour is the Blocking proposal model. A blocking pair is formed by two agents from the opposite sides of the market who, if they met, would prefer to be matched to each other. Here agents only make a proposal to other agents when they know it would be accepted by the other party, i.e. the blocking pair, in the current state of the market. The match can still be broken in the future if either of the agents finds a more preferred partner.

Although we investigate multiple behaviour models of the proposing agents, we always assume that the responder is a ZI agent. He only and always accepts proposals made by agents higher than the current match on his preference list. In addition, the existence of an information aggregating institution is implicitly assumed in Better and Blocking proposal behaviour. For example finding a potential blocking pair in the Blocking proposal model can be thought of as being supported by an institution.

We are only interested in studying the aggregate results of the search behaviour, therefore we simplify most of the third-order assumptions. The preferences of the agents are fixed, so they do not adapt to market conditions during the search. There is also no creation or destruction rate of

new agents or positions, nor any external shocks that might trigger such destructions or creations. We do, however, explicitly model market thickness. The market is considered thick when there is exactly the same number of agents on both sides, so all the agents can potentially be matched. If there are more agents on one side, the market will not be thick and there will always be some agents unassigned. In addition, thickness has an impact on the search outcome, as agents from the smaller side have more options to choose from. Thickness is also an indicator of disequilibrium, as the number of jobs is not equal to the number of workers, characterising exogenous dynamics of job creation, destruction, etc.

Finally, in our main experiments there are no limitations on matching with a more preferred partner, i.e. no commitments to contracts or any transaction costs for changing a match. However, in Section 2.7.3 we show the effect of frictions of enforcing different types of obstacles, including contractual ones, on re-matching.

More formally, we employ a model similar to that used in modelling centralised two-sided matching markets (e.g. Roth, 2008). There is a set $\mathcal{A} = \{a_1, \dots, a_{n_A}\}$ of agents on one side and a set $\mathcal{B} = \{b_1, \dots, b_{n_B}\}$ of agents on the other side. The number of agents on both sides can differ ($n_A \neq n_B$) depending on market thickness. Each agent a_i from \mathcal{A} has a strict preference relation \succ_{a_i} over agents in \mathcal{B} , and similarly for $b_j \in \mathcal{B}$ there is a preference relation \succ_{b_j} over agents in \mathcal{A} . A matching μ is a mapping from $\mathcal{A} \cup \mathcal{B}$ to itself, so that every $a_i \in \mathcal{A}$, is matched to $\mu(a_i) \in \mathcal{B} \cup \{a_i\}$, and similarly for $b_j \in \mathcal{B}$, $\mu(b_j) \in \mathcal{A} \cup \{b_j\}$. When an agent is matched to itself, $\mu(a_i) = a_i$ or $\mu(b_j) = b_j$ respectively indicates that they are in fact unmatched. Being matched to oneself is the least preferred option for all the agents. Agents would still find only the acceptable positions in their preference list, which might not contain all positions. Similarly for the position, only some agents might be acceptable. In addition, for every $a_i, b_j \in \mathcal{A} \cup \mathcal{B}$, $\mu(a_i) = b_j$ implies $\mu(b_j) = a_i$.

A matching is unstable if there are at least two agents a_i and b_j from opposite sides of the market so that $b_j \succ_{a_i} \mu(a_i)$ and $a_i \succ_{b_j} \mu(b_j)$ – a blocking pair. A matching is stable, if it is not unstable. In Table 2.3 we have listed the notation for quick reference.

With the notation in Table 2.3 we can present a General Proposal Dynamic in Algorithm 2. The `SelectProposer()` and `SelectResponder()` procedures are distinct for each of the described models in Table 2.2. The `SelectProposer()` selects a random agent from set $\mathcal{A} \cup \mathcal{B}$ or \mathcal{A} depending on whether the behaviour model is Random side or A-proposing. The `SelectResponder()` selects an agent from the preferences of the proposer, and the actual selection depends on whether the behavioural model is the Random, Better or Blocking proposal.

Table 2.3: Notation

Symbol	Description
\mathcal{A}	Preferences of agents on side A
\mathcal{B}	Preferences of agents on side B
a_i	Preference profile for agent i , $a_i \in \mathcal{A}$
b_j	Preference profile for agent j , $b_j \in \mathcal{B}$
n_A	Number of agents on side A
n_B	Number of agents on side B
θ	Market thickness $\theta = \frac{n_B}{n_A}$
k	Length of preference lists
c	Correlation of preferences
τ	Re-matching friction
μ	Matching
s	Size of matching, counted in pairs of agents
u	Unassigned percentage on side A , $u = 1 - \frac{s}{n_A}$
v	Unassigned percentage on side B , $v = 1 - \frac{s}{n_B}$
$r(\mu(i))$	Matched rank of agent i in matching μ
\tilde{r}_a	Median matched rank of agents in \mathcal{A}
\tilde{r}_b	Median matched rank of agents in \mathcal{B}
ρ_i	Number of blocking pairs for agent i
$\bar{\rho}_i$	Number of blocking pairs with unmatched agents for agent i
$\tilde{\rho}_i$	Number of blocking pairs with matched agents for agent i

Algorithm 2 General Proposal Dynamic

Require: $\mathcal{A}, \mathcal{B}, \mu$

Ensure: μ is a matching

$p \leftarrow \text{SelectProposer}(\mathcal{A}, \mathcal{B})$

$m_p \leftarrow \mu(p)$

$r \leftarrow \text{SelectResponder}(p)$

$m_r \leftarrow \mu(r)$

if $p \succ_r m_r$ **and** $r \succ_p m_p$ **then**

$\mu(m_p) \leftarrow m_p, \mu(m_r) \leftarrow m_r$

$\mu(r) \leftarrow p, \mu(p) \leftarrow r$

end if

return μ

We study the macro-level convergence properties of the search behaviour. On an individual agent level, the market needs not to be in equilibrium. There have been studies on the equilibrium of decentralised matching processes. Niederle and Yariv (2009) study such applications where firms and workers have aligned preferences and show the conditions for having a stable matching in equilibrium. Haeringer and Wooders (2011) examine equilibrium behaviour with slightly different models, where agents cannot be re-matched to a previously rejected partner, but their model is otherwise similar to ours as agents have to respond immediately. Diamantoudi et al. (2015) look at stability when agents make a commitment to a partner, which can either be only for a certain period, or an infinite commitment so that participants exit the market, or only a one-sided commitment. They show that having a requirement for firms to commit to an employee can result in unstable matchings in equilibrium. In our models we mostly concentrate at the no-commitment scenario, except in the impediments scenario. Pais (2008) analyses the equilibrium with limited information about preferences. Eriksson and Häggström (2007) also study decentralised matching, but they do not make any underlying assumptions about how the matching is reached. Instead, they measure the degree of instability in some random matchings. However, if some decentralised matching model is assumed, the resulting matching would not be a uniform selection of all the possible matchings.

In addition there have been experimental studies (Echenique and Wilson, 2009; Echenique and Yariv, 2013) about decentralised matching markets with human subjects which show that stability tends to be a prevalent outcome, but is not always guaranteed. The interesting aspect in those cases is human behaviour, which usually also restricts the size of the experiments, which tend to be small – 10-20 participants. Zhou et al. (2014b) use real-world data from small and large on-line matching markets and study

the statistical regularities of those matchings, mainly how the size of the markets relates to the size of the matching. This is also what we are interested in. Unfortunately they do not count the size of the two sides of the markets, but only the overall size.

2.3 Preference generation

Nevertheless, we look at heterogeneous agents with various degrees of correlation in their preferences and the availability of matches. We model a situation where the preferences of the agents are all idiosyncratic (effectively random) and agents find all partners acceptable. Yet, we also look at some structural constraints. Firstly, we introduce limited preference lists indicating that an agent finds only a fraction of the partners acceptable. Second, preference lists are somewhat correlated, or in extreme cases, preferences are exactly the same, indicating common tastes.

In the real world, correlated preferences show "popular tastes", e.g. all agents have similar preferences for high paying jobs or are interested in simple assignments, etc. The length of the preference list, however, indicates the probability of an agent being found unacceptable, even if a certain agent would be the only candidate. So a person without a pilot's licence would never be employed as a pilot. Thus, shorter preference lists imply that not all agents are acceptable to a particular position.

We assume that agents have strict preferences for agents (workers or positions) from the opposite side of the market. In the simplest case preferences are random, i.e. each agent has a totally idiosyncratic preference ordering. In general we can think of more structured preferences in a society, parametrised by the length of the preference list (k) and the correlation between the preference lists (c). In our experiments, the preference list limit k is set to be the same for all agents. Correlated preferences are from a global preference ordering. The degree of correlation is also the same for all agents, but the preference ordering is not necessarily the same when comparing two agents.

We generate the preferences using algorithm 3 with parameters k , c and n . This algorithm is a modified version of a random permutation algorithm from Knuth (1997a) to generate correlated preferences with parameter c . The algorithm starts with a master list of n numbers (agents). Then it iterates the list from beginning to end, each time at position j randomly selecting a position $q \in [j + 1, n]$ to exchange values with. The correlation parameter c states how biased the randomly selected position is, higher values indicate that the exchange position is selected closer to the current position j . With $c = 0.0$ the selection is uniformly probable over all positions, until finally at $c = 1$ the exchange position is always the active

Algorithm 3 Correlated permutation

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




Figure 2.2: Preference probabilities with degrees of correlation

position and all the generated lists are exactly the same. There is one global ordering of agents for each side of the market that is used for generating correlated preferences.

The power of uniform distribution \mathcal{U}^{1-c} used to randomly select the exchange positions while generating the correlated preference list is proportional to the *Beta* distribution with parameters $Beta(\frac{1}{1-c}, 1) \sim \mathcal{U}^{1-c}$. In Figure 2.2, we see the probabilities of having a particular value at some position in a list of 10 values between 0 and 9. Each box is a position in a list and displays the probability of having a certain value in that position. We see that when $c = 1.0$ then all positions have a 100% probability of having the same value and when $c = 0.0$ then all values in any position are uniformly probable.

In Figure 2.3, we compute for comparison the mean Spearman ρ and Kendall τ correlation coefficients over all the preference lists. We compute two types of means over the correlation coefficients, first compared to the initial global ordering and then a mean over pairwise correlations among a random sample of preference lists. We see that the pairwise means are always below when compared to the correlation with the global ordering. This is because although all the preferences are a similar distance from the global list, the generated lists might still be far from each other, i.e. have a




Figure 2.3: Spearman ρ and Kendall τ of generated preferences

lower correlation. That is the case with small degrees of correlation. Still, when the correlation $c = 1.0$ all the lists are exactly the same and the ρ and τ values are also 1.0.

In reality, the limit of the preference list might be due to skill mismatch in position requirements and for agents based on utility. Similar limitations on the length of preferences have been studied in Zhang (2001) and Laureti and Zhang (2003). We consider the preferences to be “known” to the agents only in terms of the behaviour model employed. So, for example, in the Better proposal model, agents would select a random proposal that is an improvement over their current match, but it might not be their most preferred match.

We do not study societies where, in general, positions and workers might have aligned preferences as in Niederle and Yariv (2009). High correlations between preference lists are usually driven by people receiving similar information about alternatives and also due to similar value systems. It is observed by Roth and Peranson (1999) that high correlations limit the size of the core of stable matchings. Certain aspects of correlation have been investigated by Biró and Norman (2012) that looks at fully correlated preference lists by varying the length of preference lists and its effect on convergence to stability.

There have been additional studies on the effect of correlation. Generally, correlation is defined as the agent’s utility function in the form $u_{a_i}(b_j) = \beta \cdot \xi(b_j) + \xi_{a_i}(b_j)$ (Ashlagi et al., 2013a; Caldarelli and Capocci,




Figure 2.4: Convergence over time

2001; Boudreau and Knoblauch, 2010) and then sorted to obtain a preference ordering. The parameter β is the correlation parameter and in case of $\beta = 0$ we would recover the uncorrelated preferences. The $\xi(n_j)$ is the global popularity of the agent b_j and $\xi_{a_i}(b_j)$ is the specific utility of agent a_i for agent b_j . It should be noted that β can be arbitrarily large, thus it is hard to have fully correlated preference lists. (Boudreau and Knoblauch (2010)) define a similarity measure for preference lists after generation, but usually the results are still far from fully correlated ($c = 1.0$) preferences.

2.4 Computational experiments and convergence

All our experiments are carried out with $n_A = |\mathcal{A}| = 1000$ agents. We vary the market thickness $\theta = \frac{n_B}{n_A} \in [0.5, 2.0]$, which varies in the number of agents on side \mathcal{B} $n_B = |\mathcal{B}| \in [500, 2000]$. We do 300,000 matches. Figure 2.4 shows that with all the behaviour models and various values of market thickness, the matching size converges to a steady-state. This does not mean that there are no changes in the matching. Individual agents still change their matches whenever their behavioural mechanism conditions are met. In the experiments, the result was almost never a stable matching without blocking pairs. Therefore, small fluctuations always occur in the matching, but Figure 2.4 demonstrates that this does not have a significant effect on the macro-level of matching.

2.5 Stability of a matching

Ackermann et al. (2008) showed that the lower bound for convergence time to stability is exponential $2^{\Omega(n)}$ with respect to n being the number of agents on one side of a thick ($\theta = 1.0$) market and with full preference lists. This indicates that with large n stability becomes nearly impossible. Our decentralised behaviour models operate by satisfying blocking pairs in each

transaction, otherwise the proposal is rejected. These are not guaranteed to lead to stability, as some blocking pairs are satisfied some new blocking pairs are created. Furthermore it has been shown that by satisfying blocking pairs in our Blocking Proposal behaviour the matching may cycle (Knuth, 1997b; Ackermann et al., 2008).

The results of probability of stability with $n_A = n_B = 100$ and varying c and k is presented in Figure 2.5. Similar results are reported by Biró and Norman (2012), where they look at $k \cdot n \leq 8$. We observe that the behaviours do not always converge to a stable matching when $k \geq 0.20$, that is there are more than 19 candidates on agents' preference lists. This is in-line with the exponential convergence as $2^{19} \approx 500000$ expectation of potential proposals that are processed in our experiments.




Figure 2.5: Probability of a stable matching with $n_A = n_B = 100$

However, when we enlarge the number of agents to $n_A = n_B = 1000$ we see in Figure 2.6 that stability is very rare, as there are always blocking pairs. Stability arises only when preference lists are very short ($k < 0.10$) or when $c = 1$. The latter only occurs in Blocking proposal behaviours. Roth and Vate (1990) showed that by randomly satisfying blocking pairs, the matching eventually reaches stability with probability one. Although, there may be cycles in satisfying the blocking pairs, there is at least one path, sequence of such pairs, that always end up in a stable matching. Apparently when preference lists are short or very correlated, the possibility of a cycle in satisfying blocking pairs is small, thus we converge to a stable matching faster than with longer or less correlated preferences. Also the situation when preferences are fully correlated, $c = 1.0$, is special as it allows better coordination in a decentralised market. In Figure 2.5 we saw that fully correlated preference lists can be must longer as still have near 100% probability of a stable matching.

As stable matchings are rare with 1000 agents on both sides, we turn our attention the number of blocking pairs. In Figure 2.6 we show the number of blocking pairs by k and c . We see that the length of the preference list (k) has the greatest effect.




Figure 2.6: Average number of blocking pairs in a matching with $n_A = n_B = 1000$

If we look at the probability of having a random pair of agents a blocking pair, we see that this is almost constant in each of the mechanisms. The probability of a blocking pair is $Pr(\rho) = \frac{\rho}{n_A \cdot k \cdot n_B \cdot k}$. Fitting the probability to a linear model in the form (2.1), we obtain the results in Table 2.4.

$$Pr(\rho|c) = \beta_1 + \beta_2[c = 0.9] + \dots + \beta_{11}[c = 0.0] \quad (2.1)$$

We see that with Noise models ((1), (2)) the blocking pair probability is about $Pr(\rho) \approx 20\%$, slightly decreases with c , but is around $11\% - 12\%$ when $c = 0.0$. For Better proposal models ((3), (4)) the $Pr(\rho) \approx 30\%$ and slightly higher 33% only for $c = 0.9$ in Better Proposal (4). However in Better Proposal A behaviour the probability of a blocking pair decreases slightly with correlation decreasing, indicating some coordinating effect from just A-side proposing. For Best Proposal behaviour model (5) the fit when $c = 1.0$ is lowest compared to smaller correlation, but then increases significantly when $c \approx 0.9$ to 35% and finally settles to about 30% . Finally with Blocking Proposal A (6) blocking pairs probability settles to about 27% , that is lower compared to just Blocking Proposal, again indicating some coordinating effect from just A-side proposing.

To summarise, the noisy (ZI), models we have the lowest probability for blocking pairs compared to more sophisticated behaviours. Furthermore, the more sophisticated behaviour often benefit from having the proposing side fixed as the proportion of blocking pairs slightly is lower.




Figure 2.7: Blocking pairs in a matching with $c = 0.0$ and $k = 100\%$

Table 2.4: Probability of a blocking pair

	$Pr(\rho c) \cdot 100\%$					
	(1)	(2)	(3)	(4)	(5)	(6)
$c = 1.0$	19.690*** (0.042)	19.641*** (0.040)	30.022*** (0.067)	30.558*** (0.068)	24.683*** (0.102)	32.975*** (0.080)
$c = 0.9$	-5.759*** (0.062)	-5.648*** (0.060)	3.362*** (0.099)	0.651*** (0.101)	10.110*** (0.151)	-0.148 (0.118)
$c = 0.8$	-7.513*** (0.059)	-7.455*** (0.057)	1.915*** (0.095)	-1.172*** (0.096)	7.751*** (0.145)	-2.597*** (0.113)
$c = 0.7$	-8.488*** (0.060)	-8.266*** (0.058)	0.629*** (0.096)	-2.275*** (0.097)	6.292*** (0.146)	-4.205*** (0.115)
$c = 0.6$	-8.531*** (0.062)	-8.440*** (0.060)	0.211** (0.099)	-2.777*** (0.101)	5.789*** (0.151)	-4.748*** (0.118)
$c = 0.5$	-8.640*** (0.062)	-8.561*** (0.059)	-0.288*** (0.099)	-3.062*** (0.100)	5.244*** (0.150)	-5.214*** (0.118)
$c = 0.4$	-8.445*** (0.060)	-8.454*** (0.058)	-0.162* (0.096)	-3.272*** (0.098)	5.293*** (0.147)	-5.443*** (0.115)
$c = 0.3$	-8.289*** (0.059)	-8.421*** (0.057)	-0.300*** (0.094)	-3.337*** (0.096)	5.260*** (0.143)	-5.592*** (0.113)
$c = 0.2$	-8.330*** (0.061)	-8.264*** (0.059)	-0.225** (0.098)	-3.611*** (0.100)	4.954*** (0.149)	-5.884*** (0.117)
$c = 0.1$	-8.319*** (0.059)	-8.234*** (0.057)	-0.207** (0.094)	-3.526*** (0.096)	5.183*** (0.143)	-5.792*** (0.113)
$c = 0.0$	-8.348*** (0.060)	-8.260*** (0.057)	-0.238** (0.095)	-3.412*** (0.097)	5.086*** (0.146)	-5.904*** (0.114)

Note: *p<0.1; **p<0.05; ***p<0.01

However, the market thickness has some effect on the probability of blocking pairs. Figure 2.7 shows how the probability of a blocking pair is impacted by the thickness parameter θ . First we observe that the probability of a blocking pair is always highest when the market is thickest, that is $\theta = 1$. Also the effect is symmetrical with respect to thickness as $Pr(\rho|\theta = 0.5) \approx Pr(\rho|\theta = 2.0)$, which is expected and the same holds true for models where the agents from A-side are always in the proposers role. Moreover, the noisy model is again superior to other models, as it has always a lower probability for a blocking pair.

The Figure 2.7 indicates that the number of blocking is bounded by the number of agent pairs, as $\frac{\rho}{n_A \cdot n_B} \rightarrow const..$ Khuller et al. (1994) show that there is a lower bound of $\Omega(n^2)$ for on-line randomised matchings. However, we see that the constant is important, as the Noise Proposal mechanisms result about three time lower number of blocking pairs.

2.6 Unassigned agents in a thick market

2.6.1 Analysis of convergence conditions

Our probabilistic analysis considers the simpler situation where the market is thick ($\theta = 1.0$), all agents have full ($k = 1.0$) and uncorrelated preference lists ($c = 0.0$). With limited preference lists, the analysis would not hold and would need to be augmented with probabilities of having certain agents on a preference list. Similarly, with correlation, we would need to assume some probability of having certain agents higher on the preference lists. Calculations are much more simplified, when we can assume this to be of uniform probability for all the relevant agents.

There are four types of events that can occur in all of the decentralised matching behaviour models:

- e_1 Two previously unmatched agents are matched. The size of the matched population increases by one on both sides and nobody becomes unmatched. The net change in the size of the matching will be one.
- e_2, e_3 One unmatched agent (either from \mathcal{A} or \mathcal{B}) is matched to another matched agent. The matched population increases by one, but one previously matched agent now becomes unmatched due to the divorce of the already matched agent. The net change in the matching size will be zero.
- e_4 Two already matched agents are matched to each other and consequently two divorces occur. The net change in the size of the matching will be minus one.

We are interested in understanding the convergence of the size of the matching. Since for events e_2 and e_3 the net change in the matching is zero, we are not interested in those events. The size of the matching changes only with the events e_1 and e_4 and has converged when $\Delta s = P(e_1) - P(e_4) \rightarrow 0$. We analyse the probabilities of the events e_1 and e_4 for all of the six decentralised behaviour models. This is similar to the model in (Mortensen and Pissarides, 1999, p. 1185). However, Mortensen and Pissarides (1999) analyse the model on a macroscopic level with transition probabilities on a Markov chain. Yet, we analyse the model on an agent level, where the transition probabilities depend on the states of the agents.

Noise Proposal and Noise Proposal A Whenever two unmatched agents meet, they always prefer to be matched rather than unmatched, given our assumptions about preferences. Hence, the probability for event e_1 is the probability for two unmatched agents to meet as in (2.2).

$$\widehat{P}(e_1) = \left(1 - \frac{s}{n_A}\right) \left(1 - \frac{s}{n_B}\right) \quad (2.2)$$

The probability of event e_4 is selecting two matched agents that prefer to be matched. This firstly depends on selecting two matched agents, one from \mathcal{A} and the other from \mathcal{B} . Secondly, the selected agents would both have to be higher on each other's preference lists than their current match. The latter is an average over all the matched agents. This is summarised in (2.3).

$$\begin{aligned} \widehat{P}(e_4) = & \left(\frac{s}{n_A}\right) \left(\frac{s}{n_B}\right) \left(\frac{1}{s} \sum_{i \in \mathcal{A}, i \in \mathcal{B}} P(r(\mu(i)) > x) P(X = x)\right) \\ & \cdot \left(\frac{1}{s} \sum_{i \in \mathcal{A}, i \in \mathcal{B}} P(r(\mu(i)) > x) P(X = x)\right) \end{aligned} \quad (2.3)$$

Since the probability of selecting an agent in a particular position is uniform, $P(X = x) = \frac{1}{n}$, and the number of agents n is either n_A or n_B , depending on which side we are looking at, we can simplify (2.3), which results in (2.4).

$$\begin{aligned} \widehat{P}(e_4) = & \left(\frac{s}{n_A^2}\right) \left(\frac{s}{n_B^2}\right) \left(\frac{1}{s} \sum_{i \in \mathcal{A}, i \in \mathcal{B}} P(r(\mu(i)) > x)\right) \\ & \cdot \left(\frac{1}{s} \sum_{i \in \mathcal{A}, i \in \mathcal{B}} P(r(\mu(i)) > x)\right) \end{aligned} \quad (2.4)$$

Better proposal A The probability of event e_1 is the same as with Noise behaviour. The probability that an unmatched agent a_i is selected from \mathcal{A} and then the probability that the agent will select an unmatched agent is $b_i \in \mathcal{B}$. Since agent a_i has a full preference list, the selection is made from the entire set \mathcal{B} . This results in the probability of two unmatched agents being selected, as expressed in Equation (2.5).

$$\widehat{P}(e_1) = \left(1 - \frac{s}{n_A}\right) \left(1 - \frac{s}{n_B}\right) \quad (2.5)$$

To find the probability of event e_4 of Better proposal A, we first have to take the probability of selecting a matched agent from \mathcal{A} . Then the selected agent a_i will randomly select an agent from the set of better matches on its preference list. The matching is successful only if the selected agent from \mathcal{B} side finds a_i acceptable as well. This means, by definition, that the two agents have to form a blocking pair. With $\tilde{\rho}_i$ we count the number of blocking pairs with another matched agent from \mathcal{B} for agent a_i . This results in probability for event e_4 as in Equation 2.6.

$$\widehat{P}(e_4) = \left(\frac{s}{n_A}\right) \left(\frac{1}{s} \sum_{i \in \mu, i \in A} \frac{\tilde{\rho}_i}{r(\mu(i))}\right) \quad (2.6)$$

Better proposal When an agent from either side can act as a proposer, we only need to weigh the proposer's selection probabilities by the size of the respective agent-sets. For the probability of event e_1 , this would result in Equation (2.7), which simplifies to the same result as (2.5).

$$\begin{aligned} \widehat{P}(e_1) &= \left(\frac{n_A}{n_A + n_B} + \frac{n_B}{n_A + n_B}\right) \left(1 - \frac{s}{n_A}\right) \left(1 - \frac{s}{n_B}\right) = \\ &= \left(1 - \frac{s}{n_A}\right) \left(1 - \frac{s}{n_B}\right) \end{aligned} \quad (2.7)$$

The probability of event e_4 for Better proposal behaviour is again similar to the Better proposal A. We first take the probability of selecting a matched agent from either from \mathcal{A} or \mathcal{B} . If the proposing agent is selected, then selecting an accepting responder has the same probability as (2.6), but over all of the agents in $\mathcal{A} \cup \mathcal{B}$, which results in total probability as in

Equation (2.8).

$$\begin{aligned}
\widehat{P}(e_4) &= \frac{n_A}{n_A + n_B} \binom{s}{n_A} \left(\frac{1}{s} \sum_{i \in \mu, i \in A} \frac{\tilde{\rho}_i}{r(\mu(i))} \right) + \\
&+ \frac{n_B}{n_A + n_B} \binom{s}{n_B} \left(\frac{1}{s} \sum_{i \in \mu, i \in B} \frac{\tilde{\rho}_i}{r(\mu(i))} \right) = \\
&= \frac{1}{n_A + n_B} \sum_{i \in \mu} \frac{\tilde{\rho}_i}{r(\mu(i))} \tag{2.8}
\end{aligned}$$

Blocking proposal A The probability of event e_1 depends on selecting an unmatched agent a_i from \mathcal{A} and then agent a_i selecting an unmatched agent from among its blocking pairs. On average, this results in probability as in Equation (2.9). When all agents have full preference lists, we could simplify even further with $\bar{\rho}_i = n_B - s$ as all unmatched B agents would be blocking pairs for any unmatched A agent.

$$\widehat{P}(e_1) = \left(1 - \frac{s}{n_A} \right) \left(\frac{1}{n_A - s} \sum_{i \notin \mu, i \in A} \frac{\bar{\rho}_i}{\rho_i} \right) = \frac{1}{n_A} \sum_{i \notin \mu, i \in A} \frac{\bar{\rho}_i}{\rho_i} \tag{2.9}$$

Similarly the probability of event e_4 depends on selecting a matched agent $a_i \in \mathcal{A}$ and this agent a_i selecting a blocking pair with a matched agent from among all blocking pairs, including the ones with and unmatched agent. This results in probability as in Equation (2.10).

$$\widehat{P}(e_4) = \binom{s}{n_A} \left(\frac{1}{s} \sum_{i \in \mu, i \in A} \frac{\tilde{\rho}_i}{\rho_i} \right) = \frac{1}{n_A} \sum_{i \in \mu, i \in A} \frac{\tilde{\rho}_i}{\rho_i} \tag{2.10}$$

Blocking proposal Similarly to the Better proposal behaviour, we need to weigh the probabilities in Equations (2.9) and (2.10) against the probabilities of selecting an agent either from \mathcal{A} or \mathcal{B} . This results in probabilities as in Equations (2.11) and (2.12) for e_1 and e_4 respectively.

$$\begin{aligned}
\widehat{P}(e_1) &= \frac{n_A}{n_A + n_B} \left(1 - \frac{s}{n_A} \right) \frac{1}{n_A - s} \sum_{i \notin \mu, i \in A} \frac{\bar{\rho}_i}{\rho_i} + \\
&+ \frac{n_B}{n_A + n_B} \left(1 - \frac{s}{n_B} \right) \frac{1}{n_B - s} \sum_{i \notin \mu, i \in B} \frac{\bar{\rho}_i}{\rho_i} = \\
&= \frac{1}{n_A + n_B} \sum_{i \notin \mu} \frac{\bar{\rho}_i}{\rho_i} \tag{2.11}
\end{aligned}$$




Figure 2.8: Comparison of expected and experimental probabilities of event e_1




Figure 2.9: Comparison of expected and experimental probabilities of event e_4

$$\begin{aligned}
 \widehat{P}(e_4) &= \frac{n_A}{n_A + n_B} \frac{s}{n_A s} \frac{1}{s} \sum_{i \in \mu, i \in A} \frac{\tilde{\rho}_i}{\rho_i} + \frac{n_B}{n_A + n_B} \frac{s}{n_B s} \frac{1}{s} \sum_{i \in \mu, i \in B} \frac{\tilde{\rho}_i}{\rho_i} = \\
 &= \frac{1}{n_A + n_B} \sum_{i \in \mu} \frac{\tilde{\rho}_i}{\rho_i} \tag{2.12}
 \end{aligned}$$

In Figures 2.8 and 2.9, we compare the results of the probabilistic matching estimations from the structural properties of the matchings $\widehat{P}(\cdot)$ from the specified equations and actual observed probabilities $P(\cdot)$ from computational experiments. The figures show the average difference in these probabilities with 99% confidence bounds on normal distributions. We see that with all the behaviours, the statistical difference between the estimated and observed probabilities is close to zero and is always within the 99% bound (Figure 2.8 and 2.9). This indicates that the structural properties of the matchings are as expected.




Figure 2.10: Statistical difference in $P(e_1)$ and $P(e_4)$

In Figure 2.10 we investigate the convergence of the matching process. The process converges when $\Delta s = P(e_1) - P(e_4) \rightarrow 0$. This figure demonstrates that the difference $P(e_1) - P(e_4) \approx 0$ is statistically within the 99% confidence bound. This is not to say that the matching freezes. There are still new matches made as well as broken. Rather the statistical properties of the matching, in terms of size, distribution of obtained rank, and the distributions of blocking pairs, converge and become stationary.

2.6.2 Structured preferences

Now we vary the correlation and the length of preferences, but still keep the market thick, that is $\theta = 1.0$. As previously in Figure 2.6 we saw that only with very short preference lists, $k = 0.02$, the decentralised behaviours lead to stable matchings. In Figure 2.11 we show the unassigned agents in those matchings. Due to small k there are many unmatched agents, about 70% to 90%, even when these preferences are uncorrelated ($c = 0.0$).




Figure 2.11: Unassigned agents in a stable matching $\theta = 1.0$

However, when k is increased, the probability of being unassigned decreases dramatically (Figure 2.12), to around 20%, when preferences are uncorrelated. Then again these matchings are unstable and there is potentially a significant amount of re-matches.

Additionally, in Figure 2.11, we see that the size of all the stable matchings is the same, even with the optimal Deferred-Acceptance. This shown always be the case in the “Rural Hospitals” theorem reviewed previously.




Figure 2.12: Unassigned agents in an unstable matching $\theta = 1.0$

However, in the case when the decentralised behaviour find unstable matchings, we see that the optimal matching would often be an improvement, at least when the preferences are moderately correlated, $c \leq 0.5$. Noise proposal does slightly better, than other decentralised behaviours, but mostly only in the moderate correlation region. Once the preference lists are short and highly correlated, there is always large percentage of agent unassigned, even in the optimal stable matching.

2.7 Unassigned agents in a thin market – the Beveridge curve

2.7.1 Beveridge curve and the movement along the curve

We look at the Beveridge curve without any structure in preferences, that is $c = 0.0$ and $k = 1.0$. Figure 2.13 contains data points for all six of the behaviours.

The thickness (θ) of the market sides determines where on the Beveridge curve the steady-state of the matching is situated. In Figure 2.13 the lines represent some examples of market thickness. When the market is thick, i.e. we have equal number of agents on both sides ($\theta = 1.00$), the result will lie on the 45 degree line. When the market is biased toward one or the other side, we move along the Beveridge curve to the upper left or lower right. Different values for thickness can be considered the effects of outside influences, e.g. economic state that influence the job destruction and creation rates. So the curve is the result of out-of-equilibrium state of the wider market.

It appears that the best and largest matching outcome is obtained when agents behave randomly in the market, as in the two Noise proposal behaviours. Moreover, it is not relevant how the proposing power is distributed, the results are the same on average. Strangely enough, the size of the matching is much smaller when agents exhibit more intelligent behaviour, by proposing only to more preferred agents (Better proposing behaviours) or only to agents they know will accept (Blocking proposing behaviours). This is most likely the result of a matching transaction being




Figure 2.13: Beveridge curves

much more likely in the latter two cases than with noisy proposals. Thus a lower steady-state is reached due to agents making many more swaps in their partner.

It is also clear that in any of the mechanisms, which side proposes selection does not affect the size of the match. This might be an indication of the fact that being a proposer is not that relevant in large markets, as has been also discovered in large centralised matching markets (Roth and Peranson, 1999; Immorlica and Mahdian, 2015).

Depending on market thickness, the ratio of free agents is higher on the larger side on the market. In case of Better and Blocking response behaviours, the change in free agents depends linearly on market thickness. In the case of the zero-intelligence Noise proposing behaviour, the relation of the number of free agents with thickness is not linear, but is akin to a square root function. So the size of the matching increases faster with Noise proposing behaviour when the market is becoming thicker ($\theta \rightarrow 1.00$) compared to other behaviours.

2.7.2 Shifts in the Beveridge curve

The Beveridge curve is concerned with the size of the matching. Most empirical curves show a relationship between unemployment and vacancies, and it is never the case that neither of them is zero. This is usually attributed to the structural properties of preferences – some workers are not suitable for some jobs. We observe a similar effect of having structure in preferences. In addition, we show that the shift can also be the effect of search behaviour.

It has been long assumed that shifts of the Beveridge curve are due to the structure of preferences in the labour market (Abraham and Katz, 1986; Blanchard et al., 1989; Mortensen and Pissarides, 1999; Sahin et al., 2014): namely, where workers can and would like to work, and similarly who the employers would like to hire. If agents are low on the preference lists for every position or not on the list at all, it is very hard for them to find a match.

We model the preferences of the agents in terms of correlation c in a society and the length k of the preference lists. Both these factors play a significant role in how good, large, the match is. In the experiments in Figures 2.14 and 2.15, we vary c and k simultaneously to understand their interaction effects. If preference lists are very correlated (Figure 2.14), the matching tends to be small, which is also the case when the lists are short (Figure 2.15). However, it also appears that either of these features can determine the location of the Beveridge curve on its own. Conversely to the trend even when the lists are short, but with low correlation, large matches can result. This can also occur when the lists are long and correlation is high. Naturally, with long lists and low correlation, the matching is the largest. So the relationship in preference list parameters (c and k) and the size of the matching is not straightforward.




Figure 2.14: Beveridge curve and correlation in preferences c



Figure 2.15: Beveridge curve and length of preferences k

To simplify thinking about the structural properties of the preferences of agents, we use the maximum potential matching to determine the effect of correlations and limited lists on a matching. The maximum potential matching is computed using the Hopcroft-Karp algorithm (e.g. Cormen

et al., 2004, p. 696) in networkX library (Hagberg et al., 2008). This matching is then compared to the maximum matching with no correlations or limits on preferences and this comes down to the number of agents in the smaller side of the market. Thus, when $n_A = 1000$ and $n_B = 500$, the maximum matching can be $s_m = 500$. However, this can never be obtained because the preferences are somewhat structured. The maximum matching returns the size that can be obtained given the preference structure — the potential size p_m . Figure 2.16 looks at the effect of $\varphi = \frac{p_m}{s_m}$ on the Beveridge curve. We see that φ is close to 1 when we find a large matching with ZI and close to 0 when the matching with ZI is small. The latter may reflect certain skill mismatches in the market similarly to the established stylised facts in macroeconomic literature.




Figure 2.16: Beveridge curve and maximum potential matching

2.7.3 Effect of a re-matching friction

The main effect of the behaviours on the size of the matching originates from the differences in probabilities that a transaction ends with a successful re-matching. In the Noise proposal behaviour, a pair of agents is randomly selected, whereas in Better and Blocking proposal behaviours, only the agent on one side is randomly selected. In the latter two, the randomly selected agent only makes proposals that they already find acceptable. ZI-agents make proposals to random agents from the other side of the market and learn its ranking during the transaction. Thus, they might eventually reject the match.

Therefore, to disentangle the effect of the Noise proposal behaviour in a comparative context, we tweak our model slightly by adding a friction timer τ to an individual match. This would lower the probability of a transaction to result in a successful rematch. We investigate the friction effect with preferences, where $c = 0.0$ and $k = 100\%$. A new matching is only accepted, when the timer condition is satisfied for both agents forming a match. We implement three types of friction timers:

1. After making a new match *lazy* agents wait for τ iterations before accepting a new offer;

2. After making a new match *patient* agents wait for τ proposals before accepting a new offer;
3. After making a new match *greedy* agents accept only matches that are τ positions higher than their current match.

In Figures 2.17, 2.18 and 2.19 we show the results of lazy, patient and greedy agents respectively. In our experiments the friction is fixed for all agents. We see that by introducing frictions to agents before allowing them to be re-matched, the resulting match becomes larger and the Beveridge curve shifts closer to the origin. This is true for all our modelled behaviours. However, the effect is significantly greater for Better and Blocking proposal behaviours with lazy and patient agents (Figures 2.17 and 2.18). This is caused by the initial lower re-matching probability already present in frictionless Noise proposal behaviour. Moreover, the types of frictions cause some overlap. In frictionless Better and Blocking proposal behaviour, once a pair of agents was selected, the re-matching probability was higher compared to the Noise proposal, so the effect of the friction is also greater.




Figure 2.17: Beveridge curve of lazy agents




Figure 2.18: Beveridge curve of patient agents




Figure 2.19: Beveridge curve of greedy agents

For lazy and patient agents, the effects of the friction are similar. For a patient agent to re-match, they would have to be selected on τ occasions, whereas a lazy agent would have to wait for τ iterations. If the selection probabilities of an agent are the same in both cases, it should be straightforward to scale the results of lazy agents to the matching size of patient agents.

Furthermore, regarding patient (Figure 2.18) and greedy (Figure 2.19) agents with large $\tau \approx 100$, the resulting Beveridge curve is close to the origin for all behaviours. With a slightly unbalanced market, the number of free agents on the smaller side becomes effectively zero.

The greedy agent type of friction with smaller values for τ does not significantly improve the size of the matching (Figure 2.19). Greedy agents accept a re-match only when it improves their position by at least τ ranks. They would still accept any match if they were unmatched. Similarly to behaviours without the frictions, the Better and Blocking proposal behaviours still result in more unmatched agents than the Noise proposal, as agents would tend to accept proposals more often. Also, the probability of selecting a match for a greedy agent that is a τ -improvement over their current match appears high for the selected τ as the size of the matching does not increase significantly (Figure 2.19).

2.8 Price of invisibility

2.8.1 Median rank in a thick market

Unassigned agent is one important aspect of a matching and as a designer we might always aim to minimise the number of unmatched agents. However agents themselves are usually more interested in being matched to a higher ranked position or partner.

Pittel (1989) shows that in the case of random preferences the average rank for side A should be $\ln n$ and for side B $\frac{n}{\ln n}$. From our data we obtain average rank for side A to be 7.57 and side B 134, which is close to the expectation from Pittel (1989), as $\ln 1000 \approx 6.9$ and $\frac{1000}{\ln 1000} \approx 145$. As the distribution of matched rank is not necessarily normal the average is slightly different from the median, the median ranks are ≈ 5.4 for side A and ≈ 94 for side B.

However, we use the median rank of a matching as a descriptive statistic. The main reason is that the distribution of matched ranks is exponential, most receive their first some their second and then the number of agents decays by rank, and median is much better statistic for an exponential distribution than the mean rank. Secondly median has a much better interpretation to it as half of the agents received a better and half a worse rank than the median, but there is hard to find an agent who received the

average rank. Another alternative would be the the rate parameter of the exponential distribution, but the parameter describes more the skewness of the distribution than the outcome. We denote the median rank by \tilde{r}_a and \tilde{r}_b for agents on A- and B-side respectively.

We start by looking at thick, balanced, markets. In Figures 2.20, 2.21, 2.22 and 2.23¹ we have plotted the median rank as a function of correlation (c) and length (k) of preference lists. Interesting observations are that in an A-proposing deferred-acceptance matching the median rank for proposers is usually very high (lower number indicate ranked higher), except in a highly correlated markets. However for the responding B-side the median rank, while in many situations it is similar to the proposing side, when preferences are uncorrelated ($c \approx 0.0$) the median rank is much higher when compared to the proposers. This can be explained by the fact that proposers have idiosyncratic preferences, thus make proposals to different agents and face less competition and consequently the responders receive only a few proposals and are sat a disadvantage.

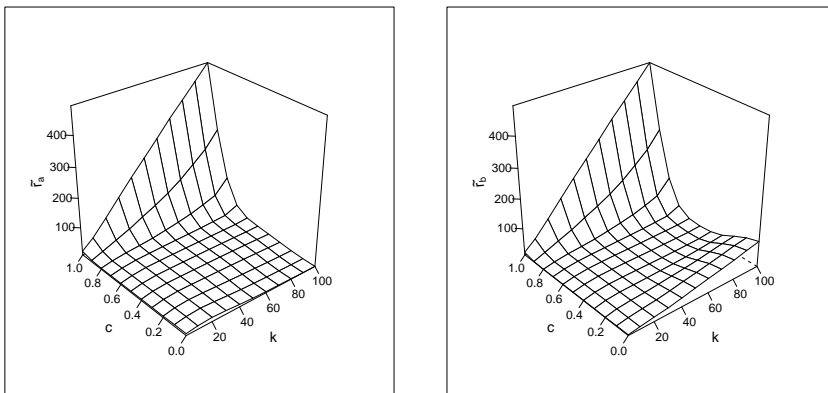


Figure 2.20: Median rank of a deferred-acceptance matching in a thick market ($\theta = 1$, $n_A = n_B = 1000$)

This is the result of the stable matching lattice, reviewed in section 1.4, as visible in the example in Table 1.13 on page 38. When preferences become more correlated, responders have more choice and can be matched to a higher ranked agent. These observations are also confirmed by other papers – Immorlica and Mahdian (2005) and Kojima and Pathak (2009) show that when the preference lists are short, even on one side, the set of stable matchings is likely to be small, and the difference in ranks is also small, which we observe when $k \leq 0.4$. Roth and Peranson (1999) also

¹Surfaces are smoothed with local linear regression

observed empirically that the size of the core is small when preference lists are short. However, in the opposite situation, the difference in matched ranks is greater, as is visible in Figure 2.20.

In Figure 2.21 we show the results from the decentralised behaviours. The Noise Proposal behaviour results in very similar median ranks for both sides, A and B , which can be expected as neither side has a definite advantage over the other. The Better and Best Response behaviours show inferior median rank compared to the Noise Proposal. In Noise Proposal model, even the proposer can later reject their own proposed match, the probability that a matching pair will change is much lower than, for example, in Better response. This creates some delay in changing a match so more agents are able to find an initial potential match, thus making future changes more robust. Many changes in a matching means that there are more free agents (as seen in Figure 2.12) willing to accept any match, which can cause the median rank to be lower.




Figure 2.21: Median rank of a matching in a thick market ($\theta = 1$, $n_A = n_B = 1000$)

In Figures 2.22 and 2.23 are the results when the proposing power is concentrated on A-side. The median rank for agents in \mathcal{A} and \mathcal{B} are shown respectively. We see that with the Noise Proposal A model the median ranks for the two sides do not differ by much and are also very similar to the Noise Proposal model. This indicates that when the behaviour of agents has a significant amount of uncertainty, market power is not really important. However, with more rational behaviour, in the Better Proposal A and Best Proposal A models, the proposing side is able to obtain a better rank. However, the median rank for the B side remains similar to the shared proposing model (Figure 2.23), which suggests that agents on B-side have nothing to lose by A-side becoming a proposer, but has potential regret by not being proposers themselves.




Figure 2.22: Median rank for A in a matching in a thick market ($\theta = 1$, $n_A = n_B = 1000$)




Figure 2.23: Median rank for B in a matching in a thick market ($\theta = 1$, $n_A = n_B = 1000$)

2.8.2 Median rank in a thin market

In Table 2.5 we have summarised the median ranks from decentralised matchings by pairwise comparisons. We compare the decentralised behaviour models to each other and in addition to centralised deferred-acceptance results. We fit a regression model of the form (2.13) to the data. We are only interested in γ coefficients that are statistically significant on $p < 0.05$ level and the effect of the coefficient itself is also significant $e^\gamma - 1 > 0.05$. This shows that when comparing the two mechanism, the difference in a median rank is significant over the experiments. In the table same or higher indicates that either one behaviour model resulted in same, that there was no statistically or effectively significant differences, or better rankings, the model resulted in statistically and effectively significant differences.

$$\ln(\bar{r}) = \beta_1[c = 0.0, k = 0.0, m = I] + \dots + \beta_n[c = 1.0, k = 1.0, m = I] + \gamma_1[c = 0.0, k = 0.0, m = II] + \dots + \gamma_n[c = 1.0, k = 1.0, m = II] \quad (2.13)$$

The Table 2.5 shows that almost always Noise Proposal model has higher median rank than any other decentralised behavioural model. Some exceptions occur when the preference lists are really short ($k = 0.02$) or highly correlated ($c = 1.0$). Exceptions occur for agents on the B-side, they achieve better median rank in A proposing models, when preferences become more correlated. This is partially explained again by the competition example mentioned before. So it appears that responding side agents are, in some cases, also better off giving having only A-side proposing.

Table 2.5: Comparison of median rank

I	II	A-side	B-side
Noise	Better	I same or higher	
Noise	Best	I same or higher, except when $k > 0.6$ and $c = 1.0$	
Noise	Noise A	I same or higher	II often same, but higher when $0.4 \leq c \leq 0.9$
Better	Better A	II same or higher	I same or higher, except when $0.7 \leq c \leq 0.9$ and $0.2 \leq k \leq 0.4$
Best	Best A		I higher, except when $c \geq 0.7$ and $0.2 \leq k \leq 0.5$
Noise A	Better A	I same or higher, except when $k \leq 0.6$ and $c < 0.6$	I same or higher
Noise A	Best A	I same or higher, except when $k \leq 0.5$ and $c \leq 0.5$	
DA	Noise	I higher, except when $c = 1.0$	
DA	Noise A		

Although, both sides would be better off by having just one side making the proposals, this requires coordination, which is not always easy. As with the Best or Better Proposal behaviours it is beneficial to be the proposing side, as their median rank is almost always higher than for the responders (Figures 2.22 and 2.23). In these situations the responders, in more correlated cases, would not prefer a shared market power either.

Even though the noise models achieved higher median rank than other decentralised behaviour models, the fully centralised deferred-acceptance matching is still an improvement. Only in extreme situations ($k = 0.02$ or $c = 1.0$) can the Noise models do better. In many cases the noise behaviours are two to three times worse than the Deferred-Acceptance, for example when $k = 1.0$ and $c = 0.0$ the Noise Proposal behaviour results is 20 times worse median rank, $\tilde{r}_a \approx 110$, than the deferred-acceptance, $\tilde{r}_a \approx 5.4$. This is the Price of Indivisibility of agents being merely guided by an *invisible hand*, as proposed by Smith (1776) for individual behaviour in general, in a matching market.

However, there can also be a different kind of market power. Figure 2.24 shows the results of the median rank (r_a, r_b) over matched weighted by the length of the preference list (n_b, n_a). First we see that with deferred-acceptance mechanism the median rank abruptly changes when market thickness shifts away from $\theta = 1.0$. When $\theta = 1.0$ then agents on both sides obtain a high rank. However when θ is slightly more or less than one then there is a significant drop in matched median rank for agents on the larger side of the market. Similar observation is made by Ashlagi et al. (2013a,b).




Figure 2.24: Median rank in a decentralised thin market with $c = 0.0$ and $k = 100\%$

In decentralised behaviour models the effect of θ is not so stark. Still agents on the smaller side usually benefit as they have greater opportunity to choose from a larger pool. In most models there is an intersection at $\theta = 1.0$, so when the market is thick both sides are matched similarly ranked agents. However, when A-side has proposing power, then in Better Proposal A and Blocking Proposal A behaviours result in better rankings for A-side. Agents from B-side usually end up with slight lower median rank, however, as mentioned in Table 2.5, occasionally can benefit.

The Noise Proposal models still produce a higher median rank for both sides among decentralised models. The Noise Proposal even results in a higher median rank for the larger side in a thin market compared to the centralised deferred-acceptance matching. This also comes with a cost for the smaller side, as they would have much higher median rank in a centralised market. This can also explain why in some situations it is hard to

agree among market participants to organise a centralised clearing house as some might benefit by deviating. See for example job-market for lawyers, where some firms, despite fixed rules, managed to circumvent and make early offers (Roth, 2015, p. 68).

2.8.3 Re-matching friction and median rank

Previously, in section 2.7.3, we saw that re-matching frictions can potentially increase the size of the match in a decentralised market. However, the effect of these frictions is not all positive. The re-matching friction increases the size of the matching, but decreases the median rank of the matched agents. In Figures 2.25, 2.26 and 2.27 we show, similarly to previous section, the results of the median rank (\tilde{r}_a, \tilde{r}_b) over matched agents from the lazy agent experiments, weighted by the length of the preference list (n_b, n_a). Lower values indicate that the median agent has a more preferred match.




Figure 2.25: Median rank for A-side lazy agents

First, with lazy agents, we observe that with the Noise Proposal behaviour and minimal impediment the median rank is in about top 20% position. And surprisingly slightly worse with Blocking Proposal. However, when only A-side has proposing power this side achieves better median rank. However, if A-side is smaller from the two, their median rank is in top 10% regardless of the friction. The re-matching friction also has significant interaction effect with market thickness (θ). With noisy behaviour the agents on the smaller side have the power to get matched to more preferred agents, regardless of the friction. However, with Blocking Proposal behaviour the effect depends on who is on which side. For agents on the larger side, the friction has an adverse effect, i.e. the median agent has a less preferred match. Conversely, for agents on the smaller A-side ($\theta > 1.0$), the longer waiting time will result in more preferred matches.

Considering patient agents the friction effects on median rank are similar to lazy agents. However, the median rank from different behaviour models are much more similar. The Noise Proposal models are not significantly superior to Best or Blocking proposal models. In addition, with lazy agents we saw that the friction does not affect much the smaller side of the market, however, patients agents can be negatively affected if case the friction is too




Figure 2.26: Median rank for A-side patient agents

severe and result in a lower median rank. We see similar effect, in different degrees, for greedy agents (Figure 2.27). As the friction increases, initially the median rank improves as well, however when passing a certain threshold (here about $\tau \approx 20$) the median rank starts decreasing.




Figure 2.27: Median rank for A-side greedy agents

As friction waiting times decrease the expected matched rank of an agent, it might not be rational for agents to participate in such a market. So market participants would advocate for lowering friction, for the reward of being better matched. However, they would also be taking an additional risk of being left unmatched.

2.9 Conclusion and discussion

Recent contributions to the economists' understanding of the micro- foundations of the Beveridge curve have enriched the early work of Blanchard et al. (1989). However, substantial gaps remain in our understanding of both the impact of the matching technology as well as the process of including mechanisms which affect the Beveridge curve. Thus, we contribute to this research gap by studying the micro-foundations underlying the Beveridge curve.

We translated the framework of Diamond-Mortensen-Pissarides to an agent-based model, with the intention of explaining both the movements along the Beveridge curve and the shifts (location) of the curve itself. Our simple model shows a two-sided decentralized market game with three key determinants – preferences, information and market conditions. Thus, it may be argued that instead of explicitly modelling labour market institutions, we implicitly include features of institutions by modelling the var-

ious behaviours of agents. Our agents can have degrees of heterogeneous or completely homogeneous preferences. The structure of the preferences indicates a notion associated with the possible mismatch of the skills of workers across jobs. There might be high demand for the same jobs and same workers, which form the source of the mismatch. We have multiple approaches to model preferences. Firstly, agents can be heterogeneous with random preferences and full-length preference lists. Secondly, preferences can be correlated to some degree which is common to all agents. Thirdly, the length of the preference list of the agents can vary, which indicates that not all positions are acceptable or not all agents are suitable for certain positions. This allows us to model the limitations of structural unemployment.

The cornerstone of our analysis is our assumption about information. Information determines how the market game is played. Generally our agents are myopic – at each stage of the game, they make random decisions and accept better proposals without any alternative strategic thinking. Agents do not learn. However, we studied different behavioural models. In our initial Noise proposal (zero-intelligence) model, agents make random proposals. For comparative purposes, we constructed two alternative decision models — the Better proposal and the Blocking proposal model. In the Better proposal model, agents randomly make proposals only to a more preferred agent than their current match. In the Blocking proposal model, agents only make proposals to a random blocking pair, indicating that the proposal is always accepted.

Through the computational experiments, we found the aggregate number of vacancies and unmatched agents which constitute the Beveridge curve. We have three relevant agent related dimensions that explain the position of the curve and/or the current position along the curve — the correlation of preferences, the length of the preference lists, and the assumptions about the decision-making mechanisms of the agents. For comparative statics, we first showed that low correlation (heterogeneous agents) will shift the Beveridge curve downward and long lists of preferences have a similar effect. We also observed that the assumptions about the decision-making behaviour affect the location of Beveridge curve considerably. Noise proposing models shift the Beveridge curve toward the origin compared to the Better or Blocking proposal models. This insight can be interpreted in light of the search and wait unemployment concept – zero-intelligence agents make random proposals that are not always accepted, while more advanced players make better proposals, thus resulting in a better match for the agent, but smaller matching overall.

In addition, we were interested in the effect of market thickness. This is the indicator for measuring the balance between market sides, i.e. equal number of jobs (agents) and worker agents indicates a thicker market. We demonstrated that thickness affects movement along the Beveridge curve. For instance, in the case of random preferences, we move right-down along the curve if there is an decreasing number of positions (job offerors) compared to agents (job seekers). This shows that the Beveridge curve is mostly the result of out-of-equilibrium dynamics in interrelated markets, affecting job creation and destruction rates. It appeared that regarding the Better and Blocking proposal mechanisms, changing market thickness simply means shifting the number of free agents or positions from one side of the market to the other. On the other hand, when agents make proposals randomly and market thickness becomes closer to one, the decrease in the rate of free agents is not linear, but a square root of free agents from the other side. Therefore, each additional position has a larger effect than one additional match, meaning that it creates opportunities for more agents to be matched. As the Better and Blocking proposals implicitly model search institutions, e.g. job hunters, it seems that these have a decreasing effect on employment.

The investigation of Noise behaviour revealed that the decreasing effect on unemployment and vacancies is related to limiting the probability of re-matches. Additional experiments showed that by enforcing some obstacle, friction, on the termination of the contract brings the Beveridge curve closer to the origin. These frictions might also have a basis from human psychology, as a sense of duty might limit an agent's willingness to terminate an existing contract. The longer the obstacle lasts, the closer to the origin the Beveridge curve locates. However, we also saw that frictions affect the matched rankings, i.e. with stronger friction the expected matched rankings decrease.

Finding a stable solution in a decentralised market is hard as the market becomes large. Even when there are 1000 agents on both sides, in our experiments, we rarely found a stable matching. In a very extreme situations with short and correlated lists we find stability. In decentralised markets the stability notion is not very useful, rather the existence of blocking pairs results in some dynamic in the market. Furthermore, Best and Better Proposal models tend to perform worse than the Noise Proposal models. In a sequentially matched results they tend to result in more blocking pairs, more unassigned agents and agents having a less preferred match.

The Price of Invisibility showed that in a thick market all the participants, from both sides, would prefer the centralised clearing house to the decentralised market. In a centralised case more agents are matched and the median rank either remain the same or is in fact better, when preference

list are longer or more correlated. However, all decentralised behaviours were better in terms on median rank for the larger side in a thin, unbalanced, market. So in a thin market, agents on the larger side would prefer a decentralised market to a centralised matching.

The effect of market re-matching friction on median rank also depend on the thickness of the market. More severe frictions have an adverse effect, the median rank increases, on the larger side of the market, while beneficial for the smaller side. Moreover, too severe frictions can be damaging to all participants, so when regulating a market setting this parameter need to be carefully considered. With potentially unobservable behaviours and dynamic market conditions, this might become a Sisyphean task.

Our approach had several simplifying assumptions: no transaction costs, no search and matching costs, no agency, homogeneous behaviour, and no dynamics (behaviour learning, new agents or change in preferences). Despite this, we open a path of research in agent-based modelling in order to contribute to the search and matching literature. Modelling matching technology by including some kind of a job board or alternative agency to the agent-based model remains a challenge for the future research.

3 Strategies in Tallinn School Choice Mechanism

3.1 Introduction

Significant research has been recently carried out to explore the allocation of school seats to students in primary (e.g. Abdulkadiroğlu et al., 2006, 2011; Dur et al., 2013), secondary (Dur et al., 2013) as well as upper-secondary schools (Abdulkadiroğlu et al., 2015, 2009). In this agenda, two-sided matching markets are used as in the “marriage problem” to solve “the college admission problem”. Some unexpected results concerning agent incentive schemes have been obtained (Abdulkadiroğlu et al., 2011; Abdulkadiroğlu and Sönmez, 2003; Pathak and Sönmez, 2013, 2008; Erdil and Kumano, 2013; Erdil and Ergin, 2008).

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

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

students are prioritised according to distance from the school. Second, families can submit three unordered preferences. Third, the mechanism uses immediate acceptance (Boston).

As with the shortcomings of the Boston mechanism, which has created a rule of thumb for submitting the preferences strategically (Ergin and Sönmez, 2006; Pathak and Sönmez, 2008; Pathak and Shi, 2013; Abdulkadiroğlu et al., 2011), we show there are similar rules of thumb for manipulation under the Tallinn mechanism. In the Boston mechanism there are different levels of sophistication among families who participate in the mechanism; that is, one strategy was to avoid ranking two over-demanded schools as their top choices or an unsubscribed school or popular school was recommended to be put as the first choice plus a “safe” second choice. Hence, as Pathak and Sönmez (2008) showed that the Boston mechanism is a coordination game among sophisticated families. Thereby “levelling the playing field” by diminishing the harm done to families, who do not strategise or do not strategise well, is emphasised as a condition for designing the new mechanism. Similarly, we introduce the Tallinn mechanism as a sophisticated game. We ask how many preferences it is rational to report under such a mechanism and whether families reveal their true top preferences or manipulate in both dimensions – report only a limited number of preferences which might not be at the top of their preference lists. In addition, we ask whether this behaviour is dependent on their preference structure – the functional form of their estimated cardinal utility function. The latter allows us to show whether the strategy of revealing preferences is dependent on the relative cardinal measure of utility from first, second etc. preference – which can be considered a measure of marginal utility. Moreover, we are interested in social inefficiencies defined as the difference between individual allocated ranks and unassigned families under the Tallinn mechanism compared to the optimal deferred-acceptance mechanism (Gale and Shapley, 1962).

Our research design is based on computational experiments. For descriptive analysis, we use data from the centralised database, e-school. The e-school database is an electronic register, where approximately 4000 7-year-old children with a known home address annually list their school selections. The rest of our data is synthetic. Our research strategy is the following. After descriptive stylised facts, we use genetic algorithms to find the best strategies for revealing the preferences of families. We illustrate the results by indicating some cases of utility functions. Family agents optimise strategies by observing their allocation and the obtained utility and adapt according to the rules of the genetic algorithm. We do not argue, that ge-

netic algorithms is necessarily the way how families learn, however we use it as a tool to find an approximate Nash equilibrium strategies (Riechmann, 2001b,a).

We continue as follows. First, we describe the broader Tallinn school market, then the concrete mechanism used by the Tallinn education administration – the Tallinn mechanism. In section 3.3, we describe the preference generation, the utility function and genetic algorithms. In section 3.4, we describe the results of the parental strategies and the obtained allocation after revealing what and how much to report to the central marketplace. Finally, we conclude by highlighting the policy implications for Estonia and for other decentralised and centralised markets.

3.2 Background: Tallinn school market

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

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

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

1. Students are assigned to exam schools based on an uncoordinated school proposing Deferred-Acceptance (DA) mechanism of decentralised schools
2. The remaining students are centrally assigned to regular schools based on the *Tallinn mechanism*
3. Unassigned students are assigned to the closest schools potentially rejecting an already assigned student. Some students might be assigned to a school they did not apply to. This continues until all students are assigned.

4. Students can reject their assigned position. Once the response deadline has passed, schools can autonomously accept students for any available positions.

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

3.2.1 Tallinn mechanism

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

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

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

We use descriptive statistics to illustrate the micro-mechanism in Tallinn over three consecutive years – from 2011 to 2013. In 2011, the market was decentralised. However, applications were centrally collected without any upper limit on the submitted preferences. The Tallinn mechanism has been in use since 2012, limiting the amount of unordered preferences submitted to three (see Table 3.1).

This stylised fact illustrates the tendency to report a limited list. Most families submit only a single preference. However, there is no clear indication that parents do not manipulate as in the Boston mechanism and decide

Table 3.1: Number of reported preferences under the Tallinn mechanism

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

to reveal strategically lower preferences or “safe” choices. Therefore, we are interested in whether it is rational to report less than three preferences and what the rationality is of reporting truthful preferences.

While there Tallinn school market in more complicated and families may report less because they are *guaranteed* a position in the 3rd round. We argue that the differentiation is hard to determine and aim to show that even when the 3rd round is not present, families may prefer to report less than possible.

3.2.2 Example of deciding what to report

We illustrate the choice set for parents using a simple extensive form game (Figure 3.1). In such a game, the parents in the starting node have three strategies – to report either 1, 2 or 3 preferences. In the following subgames, the designer randomly allocates the student to the reported school or an outside option. In the final nodes, the utilities are reported by indicating the preference – 1 stands for first preference and \emptyset indicates the utility of the outside option. In the illustration below, we assume risk neutral agents.




Figure 3.1: Extensive form reporting game

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

- $u_1(k) = 0.358 - 0.025(k - 1)$

- $u_2(k) = 0.658 - 0.325(k - 1)$

Then we obtain cardinal utilities for $k \in \{1, 2, 3\}$ as in Table 3.2.

Table 3.2: Utilities

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

Assuming the uniform probabilities of being unassigned or assigned to one of their preferences, as in Figure 3.1, we can compute the expected utilities for both utility functions and all cases of reported preferences. Notably, we do not take into account the demand for a school or the overall availability of places. Moreover, it is preferable to always report schools higher in the preference list, so we do not investigate cases where, for instance, only the second or third choice is reported, because the expected utility will definitely be lower. This might not be the case when the probabilities of being assigned to a particular school are not uniform.

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

- for reporting one school $E_1[u_1(k)] = \frac{1}{2}(0.358 + 0) = .179$
- for reporting the first two schools $E_2[u_1(k)] = \frac{1}{3}(0.358 + 0.333) = .230$
- for reporting the first three schools $E_3[u_1(k)] = \frac{1}{4}(0.358 + 0.333 + 0.309) = .250$

We see that reporting all three preferences maximised utility. With utility function $u_2(k)$ the expected utilities are:

- for reporting one school $E_1[u_2(k)] = \frac{1}{2}(0.658 + 0) = .329$
- for reporting the first two schools $E_2[u_2(k)] = \frac{1}{3}(0.658 + 0.333) = .330$
- for reporting the first three schools $E_3[u_2(k)] = \frac{1}{4}(0.658 + 0.333 + 0.009) = .250$

As Figure 3.1 illustrates the game, where under the expected utility maximisation assumptions, parents obtain higher utility by reporting only one or two schools with $u_2(k)$.

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

3.3 Model

3.3.1 Environment

We are interested in understanding strategies in multiple environments. We characterise the environment with societal parameters (Tables 3.3 and 3.4) and the parameters of an individual. Societal parameters describe the number of schools, the number of exam (popular) schools, the correlation between ordered preferences, and so on. Exam schools exist because they are popular overall, so we consider them as a metaphor for globally popular schools. Moreover, in Tallinn, these schools are still allocated the most groups through the Tallinn mechanism.

We fix the number of schools, the number of places in a school and the number of students for all our experiments (Table 3.3). In addition, the maximum number of ordered preferences for each agent is fixed. We model families as agents. They are willing to apply to or can rank up to 15 schools at the most, although the utility from lower preferences is relatively small. This is partly driven by case specificities, as 15 was the maximum number of schools listed in the decentralised market in Tallinn in 2011. From those 15 ordered preferences, agents have to select three to report in the Tallinn mechanism.

We investigate societies, where agents can have random or spatially correlated preferences – the latter indicates that schools nearby are more desired (Table 3.4). We also look at the effect of having the same set of popular schools – exam schools. In these societies, all agents would prefer exam schools, even if they are further away than the nearest regular schools. In the case of spatial preferences among exam schools, agents would still prefer schools nearby, and no other criteria matters. In each computational experiment, all these parameters are fixed. The priorities for schools are always spatial, distance based. Agents closer to a school have a higher priority in that school.

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

Table 3.3: Fixed societal parameters

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

Table 3.4: Variable societal parameters

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

3.3.2 Preferences

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

We generate the preferences using the Algorithm 3 with parameters k , c and m . This algorithm is a modified version of a random permutation algorithm (Knuth, 1997a, p. 145) to generate correlated preferences with parameter c . The algorithm starts with a master list of n numbers (agents). Then it iterates the list from beginning to end, each time at position j randomly selecting a position $q \in [j + 1, n]$ to exchange values with. The correlation parameter c illustrates how biased the randomly selected position is; higher values indicate that the exchange position is selected closer to the current position j . With $c = 0.0$ the selection is uniformly probable over all positions, until finally at $c = 1$ the exchange position is always the active position and all the generated lists are exactly the same. There is one global ordering of agents for each side of the market that is used for generating correlated preferences.

3.3.3 Utility function

While agents have a preference ordering for schools, their behaviour might also be influenced by the cardinal utility they gain from assignment to the particular preference. A similar notion was illustrated by the reporting game in Section 3.2.2. In order to understand the behaviour with different cardinal valuations, we use exponentially declining utility over alternatives. When compared to consecutive schools i and $i+1$, we assume that $\frac{u(i+1)}{u(i)} = 1 - \alpha$. Furthermore, we need to normalise the utility function such that $\sum_{i=1}^k u(i) = 1$. The resulting form of the utility function is in (3.1), where $i \in \{1, \dots, k\}$ is a position in the preference ordering.

$$u(i) = \frac{\alpha(1 - \alpha)^{i-1}}{1 - (1 - \alpha)^k} \quad (3.1)$$

When $\alpha \rightarrow 0$, then cardinal utilities for all alternatives are exactly the same $u(i) = \frac{1}{k} \forall i$. When $\alpha = 1$, then all utility is concentrated in the first preference, that is $u(1) = 1$. In Figure 3.2, we show the utility values using some examples of α .




Figure 3.2: Exponential utility function

The utility function can be compared to a linear utility function over perfect substitutes $u(x_1, \dots, x_n) = \beta_1 x_1 + \dots + \beta_n x_n$, where the consumer is allocated at most one good x_i . The β_i is the value of the allocated good x_i to the consumer (e.g. Varian, 2006, p. 61). Our utility function (3.1) states the shape of the decline in value β_i of the goods to consumers. We assume that agents are risk-neutral, i.e. they maximise their expected utility $E[u]$.

Using utility ratios $\frac{u(i+1)}{u(i)}$ to measure on preferences is also popular in decision theory, Saaty scale (Saaty, 1978), is supported by some psychological observations (e.g. Franek and Kresta, 2014, and references therein) and is often used in human decision making (Herrera et al., 2001; Chen et al., 2013; Gavalec et al., 2015). Another reason is that differences are greater in geometrically declining function than linearly. So the effects we are investigating are more evident.

3.3.4 Genetic algorithms

We use genetic algorithms to find a near-equilibrium strategy for reporting in the Tallinn mechanism. The genetic algorithms adapt existing strategies to find better ones that would result in an increased utility. The result of a genetic algorithm after optimising is a near-equilibrium steady state (e.g. Riechmann, 2001a,b). While a steady state is also by definition a Nash equilibrium in a game, it could simply be one among many in multiple equilibria games. Additionally, there is always a random mutation in genetic algorithms, which keeps the state near, but not exactly at equilibrium. Our experiments are carried out with populations of agents, where each type of agent population, defined by the utility function, learns a distribution of strategies. This is also known as one-population social-learning (Vriend, 2000) as opposed to individual multi-population learning (Chen and Tai, 2010; Chen et al., 2012).

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

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

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

A strategy in the case of the Tallinn mechanism is simply a bit-string. A bit-string is a series of 1-s and 0-s, which respectively stand for reported and not reported preference. Since we limit our agent's preferences to

Algorithm 4 Simple Genetic Algorithm - single iteration

Require: \mathcal{A} set of agents, u agents utilities

Ensure: \mathcal{A} is a set of agents

$n \leftarrow |\mathcal{A}|$

$s \leftarrow \sum_{a \in \mathcal{A}} u_a$

$p \leftarrow \{\frac{u_a}{s}, \forall a \in \mathcal{A}\}$ {selection probabilities}

$i \leftarrow 0$

for all $r_1, r_2 \in \text{Select}(\mathcal{A}, p, n)$ **do**

 {select with probability p , with replacement n pairs of strategies}

$a_i \leftarrow \text{CrossOver}(r_1, r_2)$ {assign new strategy to agent a_i }

if $\text{RandomNumber}() < 0.05$ **then**

$\text{Mutate}(a_i)$

end if

$i \leftarrow i + 1$

end for

return \mathcal{A}

$k = 15$, the length of the bit-string is 15 bits. Since the Tallinn mechanism is limited to just three preferences, the bit-string can contain at most three bits set to one. For example, a possible strategy for agent i might be $a_i = 100110000000000$; that is, the agents with this strategy would report their first, fourth and fifth preference.

We run our genetic algorithms for a fixed (2000) number of steps. In each step, an allocation is made based on the Tallinn mechanism and we get the utilities for each agent. Then based on the rules of the genetic algorithm the strategies evolve. In Algorithm 4, we present a simple genetic algorithm (e.g. Riechmann, 2001b; Simon, 2013). It consists of three operations: selection, crossover and mutation. The selection operator selects strategies with replacement and probability proportional to the expected utility. The cross-over operation randomly selects the value from either strategy for each position. Finally, with a small 0.05 probability we mutate the new strategy.

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

2013). In addition, we ignore the bottom 10% of strategies. In elitism, we keep the top 20% of strategies fixed and only use the remaining strategies in the crossover (Simon, 2013).




Figure 3.3: Mean utility




Figure 3.4: Ratio of variance and mean utility

Figures 3.3 and 3.4 show the results from the four variations of genetic algorithms. We see that the stud selection usually performs the worst, has the lowest utility and highest variation in utilities compared to the other variations. If preferences are spatially correlated and there is a large number of exam schools ($m_e = 10$), we can see that the simple genetic algorithm does slightly better with large values of α than with alternatives. For lower values of α , the simple model is statistically equivalent to the election and in some cases to elite selection. As the simple model does as good as others we further analyse the results from the simple optimisation method.

3.4 Results

3.4.1 Expected utility maximising strategies

The reported results are divided into four cases. In all of the figures illustrating the results, in the upper left corner the results with no correlation (random) preferences and no exam schools are indicated; in the upper right corner, the results with spatial (2D Euclidean) preferences and no exam schools; in the lower left corner, random preferences and ten exam

schools; and in the lower right corner, correlated preference lists and 10 exam schools. Figures 3.5, 3.6 and 3.7 show a plot with the average of the population playing a type of strategy and the standard deviation of the population over 400 experiments. The standard deviation is often small so it is not always visible on the charts.

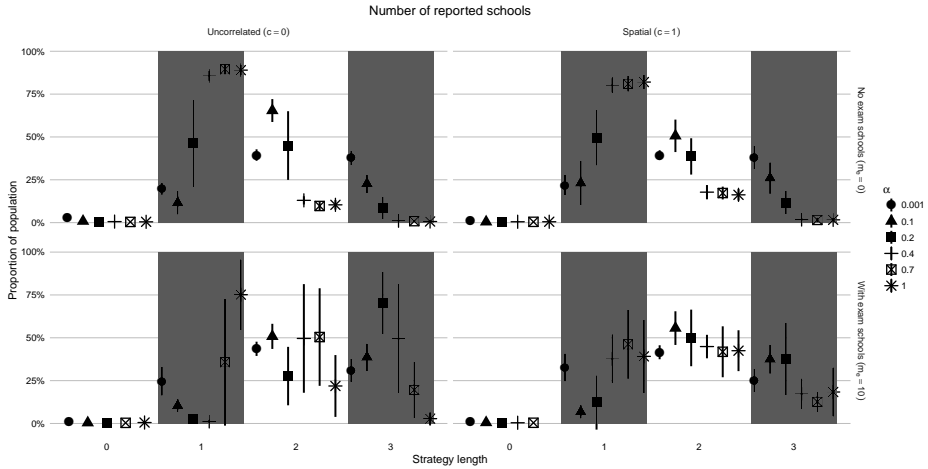


Figure 3.5: Reported strategy length




Figure 3.6: Reported preference by utility coefficient α

Firstly, we are interested in the strategy length – the number of schools to be reported. In Figure 3.5, we show the strategy length by a proportion of the respective α -population. In general, it is elucidated that the decay in the utility function is a significant determinant of a good strategy. When $\alpha \approx 0.0$, it is best to randomly select the number of schools to report with




Figure 3.7: Reported preference by strategy length

roughly uniform probability. When $\alpha \geq 0.4$ and there are no exam schools, it would almost always be best to report only one school. With random preferences, when $\alpha \approx 0.2$, there is a phase transition in the number of schools to report, as the variation at this point is largest. Therefore, it is really difficult to pick a good strategy for how to report.

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

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

We are also concerned with what to report. Figures 3.6 and 3.7 show the preferences reported by the agents' α and the strategy length. We see that without exam schools it is almost always ($\approx 90\%$) optimal for $\alpha \geq 0.4$ to report from the top of the preference list, namely just their first prefer-

ence. In regard to exam schools and random preferences, optimal reporting depends more on α , but generally the top three schools are reported. Figure 3.6 illustrates that in the case of exam schools and spatial preferences with high α , it would be better to report something from the higher and lower ends of exam schools, skipping the middle. Reporting schools lower on the preferences lists probably indicates that those agents would be otherwise unassigned, due to high demand, so they gain at least some utility. For medium α , the first three preferences are almost equally good. For indifferent agents, $\alpha \approx 0.0$, it would be best to randomly pick some schools from the list of regular schools. Also for agents with $\alpha = 0.1$, it would be beneficial to specify their most preferred exam school and most preferred regular school.

In Figure 3.7 the preferences are reported with different strategy lengths. The results show that it is always best to at least report one's most preferred school, as one might get lucky. If reporting more schools, it is useful to add the second most preferred school or with a small probability select something from even lower on the preference list. However, when reporting three choices, the selection of schools depends on the state of the school market. When preferences on the market in general are random with 50% probability, the first two preferences should be reported and the remaining options uniformly from the remainder of the preferences. In the case of spatially correlated lists or exam schools, the most preferred school should be almost always given. And when preferences are generally spatial, select the remaining options randomly. On the other hand, with exam schools and uncorrelated preferences when it is best to report three schools, it is usually best to report the top three.

3.4.2 Social welfare

Previously we investigated the individual behaviour of agents, but now we consider how these behaviours influence the outcome for the entire society. For this, we compare the results of the Tallinn mechanism to the widely used Deferred-Acceptance (DA) mechanism (Gale and Shapley, 1962; Abdulkadiroğlu and Sönmez, 2003) as described in Section 1.4. Similar to the Tallinn mechanism, the priorities in the Deferred-Acceptance mechanism are also only based on distance.

We look at two measures of social welfare. First, the proportion of unassigned agents (Figure 3.8) and second the mean utility in the allocation (Figure 3.9). Usually, the measure used in matching problems are the allocated preferences, but this is mostly due to not having access to the utility. Since in our experiments, we know the agent's utility, we measure the mean utility over all the agents.




Figure 3.8: Unassigned agents

Figure 3.8 illustrates assignment probability based on the agents’ α . We see that by using the DA mechanism and assuming random preferences ($c = 0.0$) and no exam schools ($m_e = 0$), there are no unassigned agents. When preferences are spatially correlated ($c = 1.0$), we can see that about 10% of students are unassigned, and this probability does not depend on agent type.

As described in section 3.3.2, we have $m_e = 10$ exam schools that are always first on the agents’ preference lists. Under such circumstances, only a small fraction of students receive a position in the top ten schools of their preference. Since there are fifty schools, exam schools account for 20% of places, so 20% of students receive a place in one of their top ten schools. Again, with uncorrelated preferences, DA can guarantee a place for all the students. Naturally, the students might receive a less preferred school. In the case of spatial preferences ($c = 1.0, m_e = 10$), even with DA, a significant number of students - about 10% - would be left unassigned. With the Tallinn mechanism, the number of unassigned students would be even higher – about 70% of students who have $\alpha > 0.2$ would be unassigned. This is mainly due to agents maximising their expected utility and do not have a negative utility by being left unassigned.

In Figure 3.9 we show the expected utility under the two mechanisms. Expected utility is often higher in the Tallinn mechanism compared to the DA results. A similar result was discovered in the manipulable Boston mechanism Abdulkadiroğlu et al. (2011). This leads to the conjecture that manipulable mechanisms provide the option to maximise an agent’s ex-




Figure 3.9: Mean utility comparison: Deferred-Acceptance and Tallinn mechanism

pected utility at the risk of being unassigned or assigned to a low ranked preference. Yet, as a result, a large number of agents are unassigned in the Tallinn mechanism.

We observe that agents with $\alpha = 0.999$ would maximise their expected utility by only reporting their first preference (Figure 3.5 and 3.6). This is due to the high utility value of their first preference, but because only a few preferences are reported, there is also a large probability of being unassigned under the Tallinn mechanism (Figure 3.8).

When agents are not particularly concerned with the school they are allocated to (α is small), the best strategy is to report randomly (see Figure 3.6). This also guarantees that students will not be unassigned, which is demonstrated in Figure 3.8. Other agents trade the probability of being unassigned with being assigned to a more preferred school. We see that for agents who have $\alpha \geq 0.3$, there is a high probability of being unassigned. However, there must be a considerable number of agents who are assigned to their top preferences on the condition of there being no exam schools, which increases the average utility from the allocation.

3.5 Conclusion and discussion

Our aim was to contribute to the mechanism design literature about school choice by adding a description of the Tallinn mechanism, which is a centralised school-selecting assignment based on the student's distance from

the school. Moreover, we wanted to indicate what the manipulative behaviour of agents is under such a mechanism; that is, how many preferences they report and how truthful their preference revelation is.

We used computational experiments to show the near-optimal strategies of the agents. For optimisation, we used a simple genetic algorithm, which outperformed the alternatives.

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

Our results show that in many circumstances under the Tallinn mechanism it is often best to report only one school, even if there is an option to report multiple schools. It is rarely beneficial to report three options (the maximum number). Nevertheless, it would benefit agents to report a school from the top of their preference lists. When reporting three schools, it is not always best to report the top schools and it seems to be advantageous to select the third option uniformly randomly from the remaining preferences. For agents with near-zero marginal utility, if they exist, it is best to report schools randomly. Additionally, the Tallinn mechanism maximises the expected utility of the agents, if the agents learn what and how to report, but also runs a large risk of agents not being assigned to schools. The maximisation of expected utility seems similar to a similar phenomenon in the Boston mechanism (Abdulkadiroğlu et al., 2011) given that families know how to manipulate and might be a more general property of manipulable mechanisms.

We were interested in a situation, when agent have significant utility differences over preferences, we used a multiplicative form utility function to model the change in utilities over alternatives. To some extent our results

are influenced by this assumption. However, the multiplicative form has been shown to be useful and applicable in many situation with human decision making. Alternative form utility function remain for future research. In addition, the derived equilibrium is sensitive to the distribution of the degrees of the utility function in the population. As population changes, other equilibria might emerge.

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

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

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

4 Policy Design for Kindergarten Allocation

4.1 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 and Sönmez, 2013; Pathak and Shi, 2013; Ergin and 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 allocation mechanism. However, our theoretical founding relies on the mechanism design literature motivated by related applications, such as school choice (Abdulkadiroğlu and Sönmez, 2003; Abdulkadiroğlu et al., 2005a), college admissions (Biró et al., 2010b; 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.

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 not considered or these are 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 the 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 4.2.3. 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 the distance is used in the priorities.

The classical 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 and 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.

¹In 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.

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 schools (Calsamiglia and Güell, 2014, Section 5.1)². In Boston, another relative distance policy was proposed recently by Shi (2015), mainly in order to reach the goal of the city 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 that were 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 4.3 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 and 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 secondary schools). Therefore, we found it important to investigate the effects of the changes in priorities in the performance of

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).”

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 in the probabilities of getting placed in the first choice schools. The results indicate that preference structures, more precisely their endogeneity on proximity, influence the “optimal” policy design. However, in general we can advocate for a relatively simple policy that prioritises siblings first and relative distance second.

We structure the chapter as follows. In Section 4.2 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 4.3 we define seven alternative policies and descriptive statistics of our data, including our results from computational experiments. Finally, we conclude in Section 4.5.

4.2 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 allocated to 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, which is different from some applications of two-sided markets. In college admissions for example (Gale and 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 and 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 and Sönmez, 2013), and are sometimes even prohibited by law. Thus, stability is a crucial property of any mechanism.

While there is potentially a number of stable allocations (Knuth, 1997b), 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 and Klijn (2009), backed by evidence from laboratory experiments (Calsamiglia et al., 2010) and by practical applications across the world (Pathak and Sönmez, 2013). In addition to advocating for DA, the main policy implications of these studies indicate that for 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.

4.2.1 Matching practices in Harku

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 child-birth. 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 the 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 on the top of their list and remained unassigned in the first round, 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 and Sönmez, 2003) and the procedure that had been used in Harku until 2015 was very similar to this.

4.2.2 Building a mechanism for kindergarten seat allocation

Our redesign of Harku kindergarten allocation mechanism inspired by literature has four main areas as described in Table 4.1. 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 the child. However, during the subsequent three years, the preferences of the families could have changed. That 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 preferences collected. Families use the application platform³ during a limited period (one month) six months before the service delivery (1. September) and listed all their preferences. Giving up application date as a priority will be an imminent 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 introduced clearly defined priority metrics and a centralised allocation system that ensures that the criteria are always followed. Moreover, instead of unstable and manipulable Immediate Acceptance mechanism we proposed the child-proposing DA. This is

³<https://www.haldo.ee/>

Table 4.1: 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 4.3.2 for policy design alternatives
Matching mechanism	
Decentralised mechanism which has some properties of Serial dictatorship and Boston (Immediate acceptance)	Deferred-Acceptance

a standard method for school choice (Abdulkadiroğlu and Sönmez, 2003), which eliminates justified envy, and gives incentive for the families to state their true preferences.

4.2.3 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 higher. 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 anymore. It was used as 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.

4.3 Policy design

4.3.1 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 (Matching in Practice, 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 a designer, 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 everyone to get into their most preferred kindergarten. As some families might live far away from all kindergartens, 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 families having a chance to be allocated to their 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 proximity and siblings are common practice in the case of school and kindergarten choice design, often favoured as the means to sustain community cohesively 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 in their most preferred kindergartens.

4.3.2 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 multiple 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 (Chen, 2012), is the single lottery. We will investigate both in our computational experiments. This question is discussed further by Ashlagi and Nikzad (2015) and Pathak and 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.

Table 4.2: 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

In this study, in order to explore the described aspects, we settled on seven priority policies (summarised in Table 4.2) 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.
- 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 4.3 we show the distances between homes and kindergartens. We have no children with siblings in this example.

Table 4.3: 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 4.4. 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.

Table 4.4: 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	$c_3 \prec \{c_1, c_2, c_4\}$	$c_3 \prec \{c_1, c_2, c_4\}$

4.3.3 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 4.5 shows the number of available places in each kindergarten. Also 37, about 24% of, children have a sibling in one of the kindergartens.

Table 4.5: Harku allocation

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

Table 4.6 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 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.

Table 4.6: Year 2016 comparison of policies using reported preferences

Policy	Mean preference	First	Unassigned	Mean distance (km)	With siblings
Harku	1.68	115	0	4.24	95 %
DA 1	1.76	110	0	4.26	100 %
DA 2 (M) ^a	1.85 (0.01)	98.75 (0.61)	0	4.59 (0.02)	100 % (0.0 %)
DA 2 (S)	1.72 (0.01)	108.05 (0.61)	0	4.44 (0.01)	100 % (0.0 %)
DA 3 (M)	1.83 (0.01)	98.30 (0.79)	0	4.51 (0.02)	95 % (0.25 %)
DA 3 (S)	1.72 (0.01)	107.75 (0.38)	0	4.45 (0.02)	96 % (0.3 %)
DA 4 (M)	1.91 (0.01)	89.25 (1.06)	0	4.53 (0.02)	100 % (0.0 %)
DA 4 (S)	1.75 (0.01)	104.85 (0.7)	0	4.49 (0.01)	100 % (0.0 %)
DA 5	1.76	110	0	4.26	100 %
DA 6	1.76	110	0	4.26	100 %
DA 7 (M)	1.78 (0.01)	107.60 (0.47)	0	4.30 (0.01)	100 % (0.0 %)
DA 7 (S)	1.76 (0.01)	107.75 (0.47)	0	4.31 (0.01)	100 % (0.0 %)

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

By using a simpler policy such as the DA1, we saw that there are fewer families receiving a place in their first choice kindergarten 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 kinder-

garten. 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 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 will not 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.

4.3.4 Generating counter-factual preferences

We use the 2016 data for counter-factual policy evaluation. To generate the counter-factual preferences we only 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 maps⁴ and Estonian Land Board (Maaamet⁵) and distance calculations taken from Google maps distance⁶. 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 4.7.

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 have an utility function of the form,

$$u_{ij} = \alpha_j + \sum_k \beta_k \cdot x_{kij} + \epsilon_{ij} \quad (4.1)$$

⁴<https://developers.google.com/maps/documentation/geocoding/intro>

⁵<http://inaaddress.maaamet.ee/geocoder/bulk>

⁶<https://developers.google.com/maps/documentation/distance-matrix/intro>

Table 4.7: 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 Google (2015)
walking_distance_m	walking distance between family’s home and kindergarten, based on Google (2015)
driving_distance_sec	driving time between family’s home and kindergarten, based on Google (2015)
driving_distance_m	driving distance between family’s home and kindergarten, based on Google (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	$\sqrt{haversine_distance_m}$

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 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,

$$\begin{cases} Pr_{ij}(\text{ranking} = 1) = \frac{e^{u_{ij}}}{\sum_{r=1}^7 e^{u_{ir}}} \\ Pr_{ij}(\text{ranking} = 2) = \frac{e^{u_{ij}}}{\sum_{r=2}^7 e^{u_{ir}}} \\ \dots \\ Pr_{ij}(\text{ranking} = 6) = \frac{e^{u_{ij}}}{\sum_{r=6}^7 e^{u_{ir}}} \end{cases} \quad (4.2)$$

First our aim is to select one of the distance metrics from Table 4.7 to include in the utility model (4.1). For this we do 100 bootstrap runs with each metric. In Figure 4.1 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 4.8) in any combination of features. So we select the model (3) from Table 4.8 as our final model.




Figure 4.1: Predictive features

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 (4.2). For counter-factual preferences we vary the coefficient for distance. The parameter values are in (4.3). For each combination of parameters we generate several (7) different preference profiles and evaluate the policies on the average over all the preference profiles.

$$\beta_1 \in \{0.0, 0.05, 0.1, 0.23, 0.25, 0.5, 1, 2, 4, 10\} \quad (4.3)$$

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 (4.4) and (4.5).

$$\Pr(r_1 < r_2 \mid d_1 < d_2) = \frac{\Pr(d_1 < d_2, r_1 < r_2)}{\Pr(d_1 < d_2)} \quad (4.4)$$

$$\Pr(r_1 < r_2 \mid s_1 > s_2) = \frac{\Pr(s_1 > s_2, r_1 < r_2)}{\Pr(s_1 > s_2)} \quad (4.5)$$

The mean conditional probability with fitted regression parameter, $\beta = 0.25$, is $\Pr(r_1 < r_2 \mid d_1 < d_2) \approx 0.79 \pm 0.02^7$. This is similar to what we observe in the 2016 data, where $\Pr(r_i > r_j \mid d_i < d_j) = 0.81$, $i \neq j$. In Figure 4.2a shows the relationship between the logistic parameters and the conditional probabilities.

4.3.5 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.


⁷1.96 standard deviations, 95% probability

Table 4.8: Rank-ordered logit coefficients

	preference rank		
	(1)	(2)	(3)
α_B	-0.643*** (0.150)	-0.646*** (0.151)	-0.532*** (0.143)
α_C	-0.642*** (0.179)	-0.637*** (0.181)	0.556*** (0.146)
α_D	0.223 (0.181)	0.206 (0.187)	1.690*** (0.155)
α_E	0.250 (0.153)	0.232 (0.156)	0.994*** (0.142)
α_F	0.223 (0.178)	0.247 (0.184)	1.662*** (0.153)
α_G	-1.868*** (0.193)	-1.916*** (0.197)	-1.563*** (0.177)
β_1 $\sqrt{\text{driving_distance_sec}}$	-0.247*** (0.016)	-0.247*** (0.017)	
β_1 sibling		20.910 (2,616.377)	21.560 (2,702.194)
Observations	906	906	906
Log Likelihood	-871.479	-825.165	-955.807

Note:

*p<0.1; **p<0.05; ***p<0.01



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

Figure 4.2: Coefficients and conditional probabilities

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 run twenty different randomised lotteries for each instance. As we saw in Table 4.6 the standard errors over the twenty lotteries are small. All the figures of the results show the smoothed⁸ 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 to 20 in each kindergarten. However, this creates additional competition and the resulting matched ranks will be lower (see

⁸smoothed with local polynomial regression

Ashlagi et al., 2013a,b) in these experiments than in the actual data in Table 4.6. Additionally, in our interpretations we implicitly assume the effect of the competition will be similar for all the policies.

Figures 4.3a and 4.3b 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 $\Pr(r_i \succ r_j \mid d_i < d_j) \rightarrow 1.0$, it produces one of the highest average rank score, one similar to other policies such as DA1, DA5 and DA6. Surprisingly, when the preferences of families are close to random, with conditional probability of $\Pr(r_i \succ r_j \mid d_i < d_j) \rightarrow 0.5$, 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 $\Pr(r_i \succ r_j \mid d_i < d_j) \rightarrow 0.5$. 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 4.3b). 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.

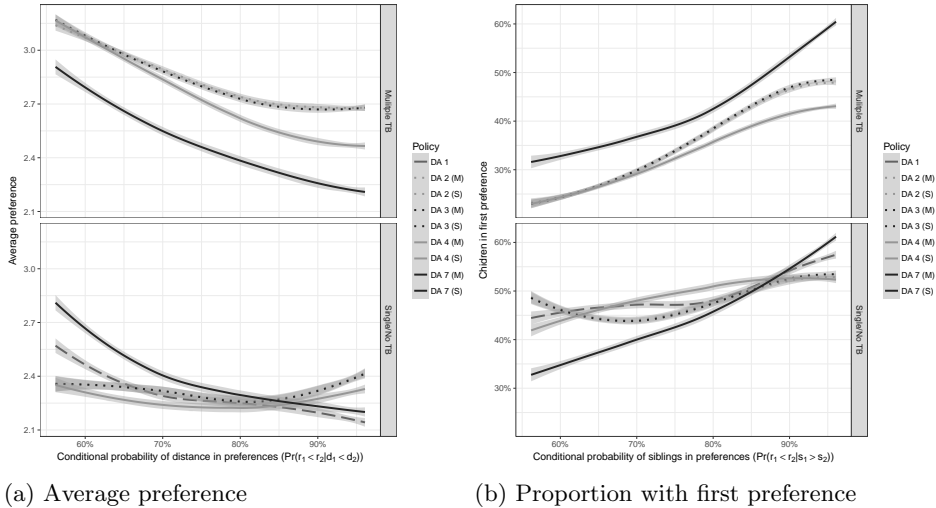


Figure 4.3: Conditional probability of distance

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 4.4a 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. Expectedly, 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 often, if preferences are random, the multiple tie-breaking lotteries have a lower average distance than single lotteries. A small improvement of 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 poli-

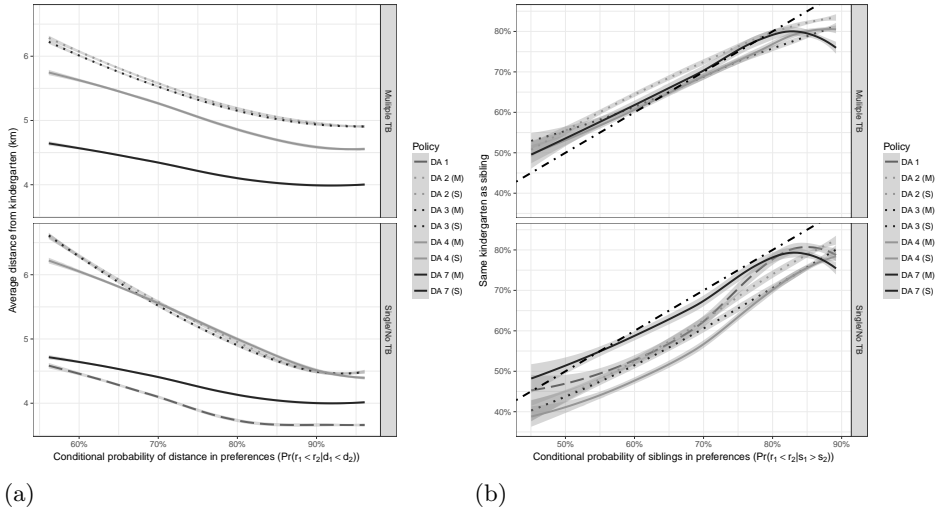


Figure 4.4: Average distance and conditional probability of distance in preferences

cies in this 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, then DA7 is better by both average preference and distance.

Figure 4.4b 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 siblings 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 the 45 degree line, indicating that policies that are below this have some children, who would prefer a kindergarten with sibling, assigned to a different kindergarten. Multiple lottery policies seem to be better at placing children in the same kindergarten with siblings.

In Figures 4.5a and 4.5b, the probability of a child being matched to the family’s first preference in at least one lottery is measured. This is a measure for fairness, or fair (equal) access to kindergartens, which is similar to the 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.




Figure 4.5: Fairness of access

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 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. 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 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 this case with DA7, the fairness is almost as good as with DA4, average distance was a significant improvement over DA4, and average allocated rank very close to DA1.

4.4 Further issues

There are some additional aspect of kindergarten allocation to consider. These were not included in the current design mainly because it would require a more profound change to the data collection during the application procedure or were very small scale, involving potentially one or two children a year. We present here a summarised overview of Further issues section from (Veski et al., 2017).

One of the issues with potentially the biggest impact is the procedure of considering children with special needs. Currently some two to four places are held back due to knowing beforehand if there will be any such cases as children with special needs are considered to take up-to three places. If no such cases arise the parental committee might make these places available. This creates another round of applications, where potentially also already allocated children might be interested to get a place in a more preferred kindergarten. There are two potential mitigation to this reallocation: first, evaluate early on the need for special needs; second, make the reallocation available to all families to avoid potential blocking pairs in the resulting match.

Some parents may works in a neighbouring municipality and might also consider getting a kindergarten place there. If these allocation procedure are not coordinated, the final selection of the parents creates an additional free place in the other municipality. After the allocation is already made it would be hard to ensure stability in the new allocation.

In many municipalities the shortage of places is acute so private childcare provides additional flexibility. However this is usually compensated for the family. Some families for some price might prefer a place in the private childcare, but for the local municipality it is often cheaper to provide a place their kindergarten. How this compensation should work to allow for optimal use of resources is for further research.

Often kindergarten have the option to open multiple groups or mixed age groups. Usually the age and mixing of the groups is decided before the allocation. It might be beneficial for the families if these decision were made during the allocation procedure to ensure that more preferred kindergartens to more families. However in general this problem is computationally hard to solve and might require heuristics to be tractable (Biró et al., 2010a, 2014).

4.5 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 appeared 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 still 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, the result of DA7 is one of the best policy designs in all aspects. DA7 gives many families their first preference, it has the shortest average distance and even 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).

A few interesting aspects of designing a more flexible mechanism might improve the allocation 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 in their most preferred kindergarten. Some of this research has been done in terms of lower quotas for opening groups (Biró et al., 2010a).

5 Conclusions and Future Work

5.1 Discussion of the research methods

In this thesis, we have applied three agent-based computational techniques to model the behaviours and outcomes in two-sided matching markets. First, we employed a very general agent-based model for behaviour in decentralised matching markets. Second, we used genetic algorithms in an agent-based model with a centralised clearing-house to study equilibrium behaviour in that market. Third, we proposed and, using computational experiments, compared policy designs for kindergarten allocation.

The decentralised agent-based models were partially motivated by the complicated school market in Tallinn, but also by the goal of understanding the outcomes of decentralised two-sided matching markets in general. We proposed three behavioural models, some of which were already studied before in the literature, while others were inspired by similar studies in the double-auction markets. The included behaviours were a noisy behaviour or zero-intelligence agents, which is also used as a benchmark model in double-auction markets. Two additional behavioural models were included that incorporated some knowledge of the market in order to make a more rational proposal choice.

While decentralised markets require sophisticated behaviour, the centralised allocation component used for school choice in Tallinn is also vulnerable to manipulation. Families are incentivised not to report their true preferences, but rather have to reason on the basis of which schools they are the most likely to get admitted to. We used genetic algorithms to investigate potential beneficial near-equilibrium strategies in the Tallinn school choice mechanism. We implemented four variations of genetic algorithms based on literature and select the results of the best performing, expected utility maximising, variation for interpretation. We ran computational experiments in four idealised environments in order to model the effect of the preference structure on the size and average utility of the matching for the agents.

In the final chapter, we commenced designing an centralised allocation mechanism. Kindergartens in Harku have long been centrally allocated by a board of the heads of kindergartens. In most cases, the allocations were

settled using a mix of immediate-acceptance mechanisms and negotiations. Harku was interested in changing the working of the allocations by minimising the effort required by parents as well as ensuring that families have an option to send siblings to the same kindergarten, which would also be close to home. We based the design on a well-known strategy-proof and optimal deferred-acceptance mechanism. This mechanism requires two-sided preferences. On the one side, there are families with preferences and on the other, kindergartens with local priorities. In the case of Harku, the priorities would be siblings and distance. Our question was how these priorities should be implemented. Multiple options can be found both in the literature and in practice. We compared seven alternatives based on actual and counter-factual preference profiles.

The computational experiments are all agent-based, as in each case, there was a set of agents with heterogeneous preferences that exhibited some type of behaviour. In a decentralised marketplace, the significant aspects were the behaviour of the agent and the confounding effect of market thickness. While in societies behaviour is rarely so clear-cut, we could nevertheless see the effect on allocation. Nevertheless, a mix of behaviours might have unexpected effects on the macro steady-state and the dynamics. However, we saw that in some real-world markets, the aggregate outcome of unassigned agents is in-between our described behaviours. A large aspect that we neglected was the market dynamics and adaptability of the behaviour of agents. The dynamics might include external shocks or interactions with other markets or even adapting by agents based on feedback from the market or due to learning from other agents through social networks. All these are fruitful avenues for further research, as our current contribution can be considered as a baseline model.

Another major application of the agent-based model is to study the effect of learning and adaptability. Using genetic algorithms, we found a distribution of near-equilibrium strategies in the Tallinn school choice mechanism. However, this automatically assumes that all agents act in a near-optimal manner, but this might not always be the case. It is still unclear how to adapt to a distribution of strategies. Also, the four modelled environments are idealised, but show some stylised facts similar to the data. As the true preferences of the families are unknown, we assumed a distribution of multiplicative utility functions. With more realistic preference structures, other types of equilibrium strategies might emerge.

We conducted several computational experiments to compare the policies based on the deferred-acceptance algorithm for kindergarten allocation in Harku. Since the mechanism is strategy-proof, we can assume truthful behaviour by the agents and we do not have to model their behaviour. First and foremost, we compared the allocations to the actual reported

preferences of the families. However, these might be misreported due to vulnerabilities in the mechanism or might change from year-to-year. In addition, we were interested in understanding how robust the policy designs are to changes in preferences. Due to these two aspects, we built a family utility function based on siblings in and the driving distance to kindergartens. Having varying degrees of utility for allocation to a nearby kindergarten or to one with a sibling, we ran the experiments using existing data. We found that the degrees of utility for siblings and distance can have a significant effect on allocation for some policies.

While we assumed that the preferences of families are based on sibling allocation and distance from home, there might be other factors that influence the utility of a particular kindergarten that are currently unobserved. Another simplification was to assume that families would prefer to be matched to any of the available kindergartens, rather than to be unmatched. That is to say they did not have an outside option. Also, as we observed in the 2016 allocation, families tend to rank all alternatives. However, some later declined a position, clearly indicating some outside option. As families and local municipalities learn how the mechanism works, it would make sense not to rank all the alternatives. This in turn might require a re-evaluation of the policies. Additionally, some kindergartens hold back positions to be allocated later and this has an effect on the stability of the outcome. This is hopefully also remedied by learning from experience by kindergartens, as was the case with schools in the New York High School matching (Roth, 2015).

5.2 Answers to the research questions

Claim A stated that in an uncoordinated and decentralised market, the behaviour of the agents is the key determinant of the matching properties – size and rank. We aimed to show that a decentralised market is significantly worse in terms of assigned agents and assigned rank than a deferred-acceptance based centralised matching.

The results showed that the noisy behaviour in a decentralised is often better on aggregate, resulting in fewer unassigned agents and a higher median rank compared to other decentralised models. Nevertheless, the results using a centralised deferred-acceptance based clearing-house would almost always assign more agents and to a higher ranked alternative. In extreme situations with highly correlated preferences, a decentralised model would have a higher median rank, but at a cost to the stability of the market. A key determinant of both decentralised and centralised is the state of market thickness. A decentralised market seems to benefit, in terms of median rank, the agents of the larger side, thus explains why in some situation centralised matching is hard to implement.

Claim B stated that merely centralising the allocation process is insufficient. We aim to show that the centralised Tallinn school choice mechanism design incentivises agents to report insincerely. In addition these behaviour results in a less preferred match for some agents, while benefitting other agents compared to an optimal allocation.

We investigated the Tallinn school choice mechanism and discovered that with some utility functions, agents are better off only reporting one or two schools. Although often these choices should be the top preferences for the families, it is not immediately clear if it should be only the first, second or third or some combination of three. In a less competitive environment, it would be optimal to report just the most preferred school. However, if the environment is more competitive, reporting schools from the very end of the preference lists would be the best option, depending on the utility functions, as well as the risk a family is willing to take. The observations from empirical data show that about 75% of families only report one school, which is a good strategy in the case of there is significantly utility in highly ranked schools.

Claim C stated that even with a mechanism wherein agents are motivated to report truthfully and are guaranteed an optimal match, the allocation properties are sensitive to policy implementation and to changes in the structure of family preferences. The policy of matching children to a nearby kindergarten can be implemented either by absolute or relative distance. The comparative allocated rank is significantly different depending on which of the two is used.

We experimented with several policy designs for allocating children to kindergartens in Harku. The policies were based on two essential metrics: the distance between home and kindergarten and the placement of any siblings in the same kindergarten. With the current preferences, we found that the main difference in allocation is determined by using a lottery or a deterministic policy. When evaluating with counter-factual preferences, we found that when preferences are close to random, families do not necessarily prefer nearby kindergartens. Thus, a policy using relative distance as a priority creates adverse competition, so eventually most families are allocated to less preferred kindergartens than with a policy with absolute geographic distance as a priority. Moreover, we observed trade-offs in policy design. Policies with randomised lotteries gave more families a chance to obtain a place in their most preferred kindergarten, while those kindergartens were on average further away from home and often a less preferred alternative. Deterministic policies minimised the average distance and average matched rank, but many families were unfairly treated due to their location. A policy design using relative distance can be fair for some preference profiles,

on average with minimal distance and matched to a high-ranking alternative. However, this policy becomes less desirable when families preferences are more random. Thus, we showed that the aggregate properties of the matching are determined by policy design and not all criteria can always be optimised. Details matter in selecting the sweet-spot of trade-offs.

5.3 Contributions

The main contributions of this thesis:

1. We proposed a simple zero-intelligence behavioural model for decentralised two-sided matching markets, which reproduce some well-known stylised fact from labour markets
2. We quantified several loss metrics in decentralised two-sided matching markets compared to a fully centralised model. The metrics included the probability of having a stable match, the probability of being unassigned and the median rank achieved in the market. The main aspects investigated are the thickness of the market as well as the structure of preferences in terms of the available alternatives and correlation among alternatives.
3. We specified the mechanism used in school choice in Tallinn and showed that it might be beneficial for families to manipulate their revealed preferences. Furthermore, we demonstrated the effective reporting strategies in this centralised admission procedure, assuming that all agents behave strategically. We further observed similar behaviour by the families in the Tallinn primary school allocation data compared to the behaviours in our experiments.
4. We showed that the Tallinn school choice mechanism might increase the expected utility of families with only a few preferences compared to widely used and strategy-proof deferred-acceptance mechanism.
5. We proposed seven alternatives for re-designing kindergarten allocation in the municipality of Harku. We compared the allocation criteria in terms of efficiency and fairness.
6. We proposed an approach based on counter-factual preference estimation for comparing the alternatives.

5.4 Further research

The experiments regarding the proposed decentralised model are the first of their kind and provide a simple and robust benchmark model. These models could easily be extended to include dynamics, e.g. a process of agents

entering and leaving the market, and/or adapting to market conditions based on feedback signals. Also, the market composition was very simple for these experiments. A more diverse mix of behavioural strategies might have a significant effect on the outcome, dynamics and convergence.

The policy space in the kindergarten allocation experiments was relatively strict, just seven alternatives. It would be interesting to have a larger policy space from where to search for policies that would satisfy some social goals. Much of the research on matching mechanism design treats policy criteria as given, but it is unclear how the criteria affect the goal. Is the matching reasonably close to minimal distance? Policy makers usually have these social goals in mind rather than specific criteria. It seems to be worthwhile to consider the criteria and goals together for a better optimisation of the policy goals.

The growing sharing economy could be seen as a significant application of research in decentralised matching markets. Currently, the most researched issue is pricing (e.g. Choudary et al., 2016; Evans and Schmalensee, 2016) for attracting supply and demand to the app-based marketplace. Once the marketplace is established, designing the matching criteria becomes more relevant. These marketplaces seek to maximise revenue, and we observed that this goal might sometimes conflict with efficiency maximisation, as sometimes the goods are not allocated. Similar issues are likely to arise when matching apartments to tenants or taxis to passengers, as participants have specific preferences and the allocated goods cannot be considered a commodity. Roth (2015) already observes some aspects of matching markets to be relevant in designing the allocation practices in these new electronic marketplaces.

References

- Abdulkadiroglu, A., Che, Y.-K., Pathak, P., Roth, A. E., and Tercieux, O. (2017). Minimizing Justified Envy in School Choice: The Design of New Orleans' OneApp. NBER Working Paper No. 23265.
- Abdulkadiroğlu, A., Agarwal, N., and Pathak, P. A. (2015). The Welfare Effects of Coordinated Assignment: Evidence from the NYC HS Match. Technical Report May, National Bureau of Economic Research, Cambridge, MA.
- Abdulkadiroğlu, A., Che, Y.-K., and Yasuda, Y. (2011). Resolving Conflicting Preferences in School Choice: The “Boston Mechanism” Reconsidered. *American Economic Review*, 101(1):399–410.
- Abdulkadiroğlu, A., Pathak, P. A., and Roth, A. E. (2005a). The New York City High School Match. *American Economic Review*, 95(2):364–367.
- Abdulkadiroğlu, A., Pathak, P. A., and Roth, A. E. (2009). Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match. *American Economic Review*, 99(5):1954–1978.
- Abdulkadiroğlu, A., Pathak, P. A., Roth, A. E., and Sönmez, T. (2005b). The Boston Public School Match. *American Economic Review*, 95(2):368–371.
- Abdulkadiroğlu, A., Pathak, P. A., Roth, A. E., and Sönmez, T. (2006). Changing the Boston School Choice Mechanism. National Bureau of Economic Research, Working Paper Series 11965.
- Abdulkadiroğlu, A. and Sönmez, T. (2003). School Choice: A Mechanism Design Approach. *American Economic Review*, 93(3):729–747.
- Abraham, K. and Katz, L. F. (1986). Cyclical unemployment: sectoral shifts or aggregate disturbances? *Journal of Political Economy*, 94(3):507–522.

- Ackermann, H., Goldberg, P. W., Mirrokni, V. S., Röglin, H., and Vöcking, B. (2008). Uncoordinated two-sided matching markets. In *Proceedings of the 9th ACM conference on Electronic commerce*, pages 256–263, New York. ACM Press.
- Albin, P. and Foley, D. K. (1992). Decentralized, dispersed exchange without an auctioneer. A Simulation Study. *Journal of Economic Behavior & Organization*, 18(1):27–51.
- Alcalde, J. and Subiza, B. (2014). Affirmative action and school choice. *International Journal of Economic Theory*, 10(3):295–312.
- An, N., Elmaghraby, W., and Keskinocak, P. (2005). Bidding strategies and their impact on revenues in combinatorial auctions. *Journal of Revenue and Pricing Management*, 3(4):337–357.
- Andersson, T. (2017). Matching Practices for Elementary Schools - Sweden. MiP Country Profile 24 http://www.matching-in-practice.eu/wp-content/uploads/2017/01/MiP_-Profile_No.24.pdf.
- Arifovic, J. (1994). Genetic algorithm learning and the cobweb model. *Journal of Economic Dynamics and Control*, 18(1):3–28.
- Arifovic, J. (1996). The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies. *Journal of Political Economy*, 104(3):510–541.
- Ashlagi, I., Kanoria, Y., and Leshno, J. D. (2013a). Unbalanced random matching markets. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 27–28, New York. ACM Press.
- Ashlagi, I., Kanoria, Y., and Leshno, J. D. (2013b). Unbalanced random matching markets: the stark effect of competition. <http://web.mit.edu/iashlagi/www/papers/UnbalancedMatchingAKL.pdf> (Accessed 14.09.2016).
- Ashlagi, I. and Nikzad, A. (2015). What matters in tie-breaking rules? How competition guides design. Unpublished working paper.
- Ausubel, L. M. and Milgrom, P. (2010). The Lovely but Lonely Vickrey Auction. In Cramton, P., Shoham, Y., and Steinberg, R., editors, *Combinatorial Auctions*, chapter 1, pages 1–37. MIT Press.
- Axelrod, R. (1980). Effective Choice in the Prisoner’s Dilemma. *The Journal of Conflict Resolution*, 24(1):3–25.

- Aygün, O. and Bo, I. (2013). College Admission with Multidimensional Reserves: The Brazilian Affirmative Action Case. <https://www2.bc.edu/inacio-bo/AygunBo2013.pdf> (Accessed 02.08.2016).
- Aziz, H., Brânzei, S., Filos-Ratsikas, A., and Frederiksen, S. K. S. (2015). The Adjusted Winner Procedure: Characterizations and Equilibria. In Yang, Q. and Wooldridge, M., editors, *Proceedings of the 24th International Conference on Artificial Intelligence*, pages 454–460.
- Bai, R., Li, J., Atkin, J. A. D., and Kendall, G. (2014). A novel approach to independent taxi scheduling problem based on stable matching. *Journal of the Operational Research Society*, 65(10):1501–1510.
- Beard, T. R., Jackson, J. D., Kaserman, D., and Kim, H. (2012). A time-series analysis of the U.S. kidney transplantation and the waiting list: donor substitution effects. *Empirical Economics*, 42(1):261–277.
- Beltratti, A. and Margarita, S. (1993). Evolution of trading strategies among heterogeneous artificial economic agents. In Meyer, J.-A., Roitblat, H. L., and Wilson, S. W., editors, *Proceedings of the second international conference on From animals to animats 2 : simulation of adaptive behavior: simulation of adaptive behavior*, pages 494–501, Cambridge. MIT Press.
- Binmore, K. (2005). *Natural Justice*. Oxford University Press, New York.
- Biró, P., Fleiner, T., Irving, R. W., and 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., and McBride, I. (2014). The Hospitals / Residents Problem with Couples: Complexity and Integer Programming Models. In Gudmundsson, J. and Katajainen, J., editors, *Experimental Algorithms*, pages 10–21. Springer International Publishing.
- Biró, P., Manlove, D. F., and Mittal, S. (2010b). Size versus stability in the marriage problem. *Theoretical Computer Science*, 411(16-18):1828–1841.
- Biró, P. and Norman, G. (2012). Analysis of stochastic matching markets. *International Journal of Game Theory*, 42(4):1021–1040.
- Blake, P. R., Rand, D. G., Tingley, D., and Warneken, F. (2015). The shadow of the future promotes cooperation in a repeated prisoner’s dilemma for children. *Scientific Reports*, 5:14559.
- Blanchard, O. J., Diamond, P., Hall, R. E., and Yellen, J. (1989). The Beveridge curve. *Brookings Papers on Economic Activity*, 1(1989):1–76.

- Blum, Y., Roth, A. E., and Rothblum, U. G. (1997). Vacancy Chains and Equilibration in Senior-Level Labor Markets. *Journal of Economic Theory*, 76(2):362–411.
- Blum, Y. and Rothblum, U. G. (2002). “Timing Is Everything” and Marital Bliss. *Journal of Economic Theory*, 103(2):429–443.
- Blumrosen, L. and Nisan, N. (2007). Combinatorial Auctions. In Roughgarden, T., Nisan, N., Tardos, E., and Vazirani, V. V., editors, *Algorithmic Game Theory*, chapter 11, pages 267–299. Cambridge University Press, New York.
- Bogomolnaia, A. and Laslier, J.-F. (2007). Euclidean preferences. *Journal of Mathematical Economics*, 43(2):87–98.
- Boudreau, J. W. (2010). Stratification and growth in agent-based matching markets. *Journal of Economic Behavior and Organization*, 75(2):168–179.
- Boudreau, J. W. and Knoblauch, V. (2010). Marriage matching and inter-correlation of preferences. *Journal of Public Economic Theory*, 12(3):587–602.
- Bowles, S. (2004). *Microeconomics: Behavior, Institutions and Evolution*. Princeton University Press, New Jersey.
- Brams, S. J. and Taylor, A. D. (1996). *Fair Division: From Cake-Cutting to Dispute Resolution*. Cambridge University Press, New York.
- Brams, S. J. and Taylor, A. D. (1999). *The win-win solution: guaranteeing fair shares for everybody*. W. W. Norton & Company, New York.
- Brânzei, S., Caragiannis, I., Kurokawa, D., and Procaccia, A. D. (2016). Equilibria of Generalized Cut and Choose Protocols. <http://arxiv.org/abs/1307.2225> (Accessed 30.06.2016).
- Budish, E., Cramton, P., and Shim, J. (2015). The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response. *The Quarterly Journal of Economics*, 130(4):1547–1621.
- Budish, E. B., Cramton, P., and Shim, J. J. (2013). The High-Frequency Trading Arms Race: Frequent Batch Auctions as a Market Design Response. Chicago Booth Research Paper No. 14-03.
- Caldarelli, G. and Capocci, A. (2001). Beauty and distance in the stable marriage problem. *Physica A: Statistical Mechanics and its Applications*, 300(1-2):325–331.

- Caldarelli, G., Capocci, A., and Laureti, P. (2001). Sex-oriented stable matchings of the marriage problem with correlated and incomplete information. *Physica A: Statistical Mechanics and its Applications*, 299(1-2):268–272.
- Calsamiglia, C. and 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., and Klijn, F. (2010). Constrained School Choice: An Experimental Study. *American Economic Review*, 100(4):1860–1874.
- Caragiannis, I., Kaklamanis, C., Kanellopoulos, P., and Kyropoulou, M. (2012). The Efficiency of Fair Division. *Theory of Computing Systems*, 50(4):589–610.
- Caragiannis, I., Kurokawa, D., Moulin, H., Procaccia, A. D., Shah, N., and Wang, J. (2016). The Unreasonable Fairness of Maximum Nash Welfare. In *Proceedings of the 17th ACM Conference on Economics and Computation*, pages 305–322, New York. ACM Press.
- Chen, L. (2012). University admission practices – Ireland. <http://www.matching-in-practice.eu/higher-education-in-ireland/>. accessed 2016-01-18.
- Chen, S., Liu, J., Wang, H., and Augusto, J. C. (2013). Ordering based decision making – A survey. *Information Fusion*, 14(4):521–531.
- Chen, S.-H., editor (2002). *Genetic Algorithms and Genetic Programming in Computational Finance*. Springer US, Boston, MA.
- Chen, S.-H., Chang, C.-L., and Du, Y.-R. (2012). Agent-based economic models and econometrics. *The Knowledge Engineering Review*, 27(2):187–219.
- Chen, S.-H., Kämpouridis, M., and Tsang, E. (2011). Microstructure Dynamics and Agent-Based Financial Markets. In Bosse, T., Geller, A., and Jonker, C. M., editors, *Multi-Agent-Based Simulation XI*, chapter 9, pages 121–135. Springer, Berlin.
- Chen, S.-H. and Tai, C.-C. (2010). The agent-based double auction markets: 15 years on. In Takadama, K., Cioffi-Revilla, C., and Deffuant, G., editors, *Simulating Interacting Agents and Social Phenomena*, chapter 9, pages 119–136. Springer Japan, Tokyo.

- Chen, S.-H., Zeng, R.-J., and Yu, T. (2009). Co-Evolving Trading Strategies to Analyze Bounded Rationality in Double Auction Markets. In Worzel, B., Soule, T., and Riolo, R., editors, *Genetic Programming Theory and Practice VI*, chapter 13, pages 1–19. Springer, Boston.
- Chevalere, Y., Dunne, P. E., Endriss, U., Lang, J., Lemaitre, M., Maudet, N., Padget, J., Phelps, S., Rodriguez-Aguilar, J. A., and Sousa, P. (2006). Issues in Multiagent Resource Allocation. *Informatica*, 30(1):3–31.
- Chiarella, C. and Iori, G. (2002). A simulation analysis of the microstructure of double auction markets. *Quantitative Finance*, 2(5):346–353.
- Choudary, S. P., Alstyne, M. W. V., and Parker, G. G. (2016). *Platform Revolution: How Networked Markets Are Transforming the Economy – and How to Make Them Work for You*. W. W. Norton & Company, New York.
- Contreras, J., Candiles, O., de la Fuente, J., and Gomez, T. (2001). Auction design in day-ahead electricity markets. *IEEE Transactions on Power Systems*, 16(3):409–417.
- Cormen, T. H., Leiserson, C. E., Rivest, R. L., and Stein, C. (2004). *Introduction to algorithms*. MIT Press, Boston, 2nd edition.
- Cramton, P. (2010). Simultaneous Ascending Auctions. In Cramton, P., Shoham, Y., and Steinberg, R., editors, *Combinatorial Auctions*, chapter 4, pages 99–114. MIT Press.
- Croissant, Y. (2011). Estimation of multinomial logit models in R: The mlogit Packages. <https://cran.r-project.org/web/packages/mlogit/vignettes/mlogit.pdf> (Accessed 12.12.2006).
- Daniels, M. G., Farmer, J. D., Gillemot, L., Iori, G., and Smith, E. (2003). Quantitative Model of Price Diffusion and Market Friction Based on Trading as a Mechanistic Random Process. *Physical Review Letters*, 90(10):108102.
- Dawid, H., Gemkow, S., Harting, P., Hoog, S. V. D., and Neugart, M. (2014). An agent-based macroeconomic model for economic policy analysis: The Eurace@ unibi model. Working Papers in Economics and Management 01-2014.
- de Haan, M., Gautier, P. A., Oosterbeek, H., and van der Klaauw, B. (2015). The performance of school assignment mechanisms in practice. IZA Discussion Papers 9118.

- Deissenberg, C., van der Hoog, S., and Dawid, H. (2008). EURACE: A massively parallel agent-based model of the European economy. *Applied Mathematics and Computation*, 204(2):541–552.
- Diamantoudi, E., Miyagawa, E., and Xue, L. (2015). Decentralized matching: The role of commitment. *Games and Economic Behavior*, 92:1–17.
- Dickerson, J. P., Procaccia, A. D., and Sandholm, T. (2012). Optimizing kidney exchange with transplant chains: theory and reality. In *Proceedings of the 11th International Conference on Autonomous Agents and Multiagent Systems - Volume 2*, pages 711–718, Richland. IFAAMAS.
- Duffy, J. (2006). Agent-Based Models and Human Subject Experiments. In Tesfatsion, L. and Judd, K. L., editors, *Handbook of Computational Economics Volume 2 - Agent-Based Computational Economics*, chapter 19, pages 948–1012. North-Holland, Amsterdam.
- Dur, U. M., Kominers, S. D., Pathak, P. A., and Sönmez, T. (2013). The Demise of Walk Zones in Boston: Priorities vs. Precedence in School Choice. NBER Working Paper Series 18981.
- Durlauf, S. N. and Young, H. P. (2001). The New Social Economics. In Durlauf, S. N. and Young, H. P., editors, *Social Dynamics*, chapter 1, pages 1–14. MIT Press, Cambridge.
- Dzierzawa, M. and Oméro, M.-J. (2000). Statistics of stable marriages. *Physica A: Statistical Mechanics and its Applications*, 287(1-2):321–333.
- eBay (2016). eBay. <http://pages.ebay.com/help/sell/reserve.html>. accessed 2016-06-27.
- Echenique, F. and Wilson, A. J. (2009). Clearinghouses for two-sided matching: an experimental study. Social Science Working Paper 1315. California Institute of Technology.
- Echenique, F. and Yariv, L. (2013). An experimental study of decentralized matching. <http://people.hss.caltech.edu/%7Elyariv/papers/ExpDecentralizedMatching.pdf> (Accessed 25.05.2016).
- Edelman, B., Ostrovsky, M., and Schwarz, M. (2007). Internet Advertising and the Generalized Second-Price Auction: Selling Billions of Dollars Worth of Keywords. *American Economic Review*, 97(1):242–259.
- EKOTU (2016). Eesti Kudede ja Organite Transplantatsiooni Ühing. <http://www.elundidoonorlus.ee/uhingust/>. accessed 2016-06-17.

- Erdil, A. and Ergin, H. (2008). What's the Matter with Tie-Breaking? Improving Efficiency in School Choice. *American Economic Review*, 98(3):669–689.
- Erdil, A. and Kumano, T. (2013). Prioritizing Diversity in School Choice. http://www.matching-in-practice.eu/wp-content/uploads/2013/09/Erdil-Prioritizing_Diversity.pdf (Accessed 02.08.2016).
- Ergin, H. and Sönmez, T. (2006). Games of school choice under Boston Mechanism. *Journal of Public Economics*, 90:215–237.
- Eriksson, K. and Häggström, O. (2007). Instability of matchings in decentralized markets with various preference structures. *International Journal of Game Theory*, 36(3-4):409–420.
- Evans, D. S. and Schmalensee, R. (2016). *Matchmakers: The New Economics of Multisided Platforms*. Harvard Business Review Press, Boston.
- Fagiolo, G., Dosi, G., and Gabriele, R. (2004). Matching, bargaining, and wage setting in an evolutionary model of labor market and output dynamics. *Advances in Complex Systems*, 07(02):157–186.
- Faqiry, M. N. and Das, S. (2016). Double-Sided Energy Auction Equilibrium Under Price Anticipation. <http://arxiv.org/abs/1605.06564> (Accessed 11.08.2016).
- Farmer, J. D., Gallegati, M., Hommes, C., Kirman, A., Ormerod, P., Cincotti, S., Sanchez, A., and Helbing, D. (2012). A complex systems approach to constructing better models for managing financial markets and the economy. *The European Physical Journal Special Topics*, 214(1):295–324.
- Farmer, J. D., Patelli, P., and Zovko, I. I. (2005). The predictive power of zero intelligence in financial markets. *Proceedings of the National Academy of Sciences*, 102(6):2254–2259.
- Feng, X., Chen, Y., Zhang, J., Zhang, Q., and Li, B. (2012). TAHES: A Truthful Double Auction Mechanism for Heterogeneous Spectrums. *IEEE Transactions on Wireless Communications*, 11(11):4038–4047.
- Fragiadakis, D. and Troyan, P. (2013). Market Design under Distributional Constraints: Diversity in School Choice and Other Applications. <http://tippie.uiowa.edu/economics/tow/papers/troyan-spring2014.pdf> (Accessed 02.08.2016).
- Franek, J. and Kresta, A. (2014). Judgment Scales and Consistency Measure in AHP. *Procedia Economics and Finance*, 12:164–173.

- Friedman, D. and Rust, J. (1993). *The Double Auction Market: Institutions, Theories, and Evidence*. Perseus Publishing, Cambridge, Massachusetts.
- Gabriele, R. (2002). Labor market dynamics and institutions: an evolutionary approach. LEM Working Paper Series 2002/07.
- Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15.
- Gavalec, M., Ramík, J., and Zimmermann, K. (2015). Pairwise Comparison Matrices in Decision Making. In *Decision Making and Optimization*, chapter 2, pages 29–90. Springer.
- Gode, D. K. and Sunder, S. (1993a). Allocative efficiency of markets with zero-intelligence traders: market as a partial substitute for individual rationality. *Journal of Political Economy*, 101(1):119–137.
- Gode, D. K. and Sunder, S. (1993b). Lower Bound for Efficiency of Surplus Extraction in Double Auctions. In Friedman, D. and Rust, J., editors, *The Double Auction Market: Institutions, Theories, and Evidence*, chapter 7, pages 199–219. Perseus Publishing, Cambridge.
- Google (2015). The Google Maps Distance Matrix API. <https://maps.googleapis.com/maps/api/distancematrix/>. accessed 2015-12-31.
- Gu, Y., Saad, W., Bennis, M., Debbah, M., and Han, Z. (2015). Matching theory for future wireless networks: fundamentals and applications. *IEEE Communications Magazine*, 53(5):52–59.
- Guerrero, O. A. and Axtell, R. L. (2011). Using agentization for exploring firm and labor dynamics. In *Emergent Results of Artificial Economics*, chapter 12, pages 139–150. Springer, Berlin.
- Guerrero, O. A. and Axtell, R. L. (2013). Employment growth through labor flow networks. *PLoS ONE*, 8(5):e60808.
- Haake, C.-J., Raith, M. G., and Su, F. E. (2002). Bidding for envy-freeness: A procedural approach to n-player fair-division problems. *Social Choice and Welfare*, 19(4):723–749.
- Haeringer, G. and Klijn, F. (2009). Constrained school choice. *Journal of Economic Theory*, 144(5):1921–1947.
- Haeringer, G. and Wooders, M. (2011). Decentralized job matching. *International Journal of Game Theory*, 40(1):1–28.

- Hafalir, I. E., Yenmez, M. B., and Yildirim, M. A. (2013). Effective affirmative action in school choice. *Theoretical Economics*, 8(2):325–363.
- Hagberg, A. A., Schult, D. A., and Swart, P. J. (2008). Exploring network structure, dynamics, and function using NetworkX. In Varoquaux, G., Vaught, T., and Millman, J., editors, *Proceedings of the 7th Python in Science Conference*, pages 11–15, Pasadena.
- Hamada, K., Iwama, K., and Miyazaki, S. (2009). An improved approximation lower bound for finding almost stable maximum matchings. *Information Processing Letters*, 109(18):1036–1040.
- Han, L., Su, C., Tang, L., and Zhang, H. (2011a). On Strategy-Proof Allocation without Payments or Priors. In Chen, N., Elkind, E., and Koutsoupias, E., editors, *Proceedings of 7th International Workshop on Internet and Network Economics*, pages 182–193, Berlin. Springer.
- Han, Z., Niyato, D., Saad, W., Başar, T., and Hjørungnes, A. (2011b). *Game Theory in Wireless and Communication Networks*. Cambridge University Press, New York.
- Hasbrouck, J. (2007). *Empirical Market Microstructure: The Institutions, Economics, and Econometrics of Securities Trading*. Oxford University Press.
- Hatfield, J. W. and Milgrom, P. R. (2005). Matching with Contracts. *The American Economic Review*, 95(4):913–935.
- Herrera, F., Herrera-Viedma, E., and Chiclana, F. (2001). Multiperson decision-making based on multiplicative preference relations. *European Journal of Operational Research*, 129(2):372–385.
- Hoefler, M. and Wagner, L. (2012). Locally stable matching with general preferences. <http://arxiv.org/abs/1207.1265> (Accessed 22.08.2016).
- Holte, R. C. (2001). Combinatorial Auctions, Knapsack Problems, and Hill-Climbing Search. In *14th Biennial Conference of the Canadian Society for Computational Studies of Intelligence*, pages 57–66.
- Hommes, C. H. (2006). Heterogeneous Agent Models in Economics and Finance. In Tesfatsion, L. and Judd, K. L., editors, *Handbook of Computational Economics Volume 2 - Agent-Based Computational Economics*, chapter 23, pages 1109–1186. North-Holland, Amsterdam.
- Hopcroft, J. E. and Karp, R. M. (1971). An $n^{5/2}$ Algorithm for Maximum Matchings in Bipartite Graphs. In *Proceedings of the 12th Annual Symposium on Switching and Automata Theory*, pages 122–125, New York. IEEE.

- Hopcroft, J. E. and Karp, R. M. (1973). An $n^{5/2}$ Algorithm for Maximum Matchings in Bipartite Graphs. *SIAM Journal on Computing*, 2(4):225–231.
- Immorlica, N. and Mahdian, M. (2005). Marriage, Honesty, and Stability. In *Proceedings of the 16th annual ACM-SIAM symposium on Discrete algorithms*, pages 53–62, Philadelphia. SIAM.
- Immorlica, N. and Mahdian, M. (2015). Incentives in large random two-sided markets. *ACM Transactions on Economics and Computation*, 3(3):1–25.
- Irving, R. W. (1985). An efficient algorithm for the “stable roommates” problem. *Journal of Algorithms*, 6(4):577–595.
- Ji, Z. and Ray Liu, K. (2006). Belief-Assisted Pricing for Dynamic Spectrum Allocation in Wireless Networks with Selfish Users. In *Proceedings of 3rd Annual IEEE Communications Society on Sensor and Ad Hoc Communications and Networks*, pages 119–127. IEEE.
- Kagel, J. H., Lien, Y., and Milgrom, P. (2010). Ascending Prices and Package Bidding: A Theoretical and Experimental Analysis. *American Economic Journal: Microeconomics*, 2(3):160–185.
- Kagel, J. H., Lien, Y., and Milgrom, P. (2014). Ascending prices and package bidding: Further experimental analysis. *Games and Economic Behavior*, 85:210–231.
- Kagel, J. H. and Vogt, W. (1993). Buyer’s Bid Double Auctions: Preliminary Experimental Results. In Friedman, D. and Rust, J., editors, *The Double Auction Market: Institutions, Theories, and Evidence*, chapter 10, pages 285–306. Perseus Publishing, Cambridge.
- Keizer, K., de Klerk, M., Haase-Kromwijk, B., and Weimar, W. (2005). The Dutch Algorithm for Allocation in Living Donor Kidney Exchange. *Transplantation Proceedings*, 37(2):589–591.
- Khuller, S., Mitchell, S. G., and Vazirani, V. V. (1994). On-line algorithms for weighted bipartite matching and stable marriages. *Theoretical Computer Science*, 127(2):255–267.
- Kirman, A. (2016). Complexity and Economic Policy: A Paradigm Shift or a Change in Perspective? A Review Essay on David Colander and Roland Kupers’s Complexity and the Art of Public Policy. *Journal of Economic Literature*, 54(2):534–572.

- Klemperer, P. (2004a). A Survey of Auction Theory. In *Auctions : Theory and Practice*, chapter 1, pages 9–61. Princenton University Press, Princeton.
- Klemperer, P. (2004b). *Auctions: Theory and Practice*. Princeton University Press, Princeton.
- Klemperer, P. and Bulow, J. (1996). Auctions Versus Negotiations. *American Economic Review*, 86(1):180–194.
- Knuth, D. E. (1976). *Mariages stables*. Les Presses de l’Université de Montréal, Montréal.
- Knuth, D. E. (1997a). *Seminumerical Algorithms*. Addison-Wesley, Reading, MA, 3rd edition.
- Knuth, D. E. (1997b). *Stable marriage and its relation to other combinatorial problems*. American Mathematical Society, Providence.
- Kojima, F. and Pathak, P. A. (2009). Incentives and Stability in Large Two-Sided Matching Markets. *American Economic Review*, 99(3):608–627.
- Kominers, S. D. and Sönmez, T. (2013). Designing for Diversity in Matching. Boston College Working Papers in Economics 806.
- Kuhn, H. W. (1955). The Hungarian method for the assignment problem. *Naval Research Logistics Quarterly*, 2(1-2):83–97.
- Ladley, D. (2012). Zero intelligence in economics and finance. *The Knowledge Engineering Review*, 27(02):273–286.
- Laureti, P. and Zhang, Y.-C. (2003). Matching games with partial information. *Physica A: Statistical Mechanics and its Applications*, 324(1-2):49–65.
- LeBaron, B. (2006). Agent-based Computational Finance. In Tesfatsion, L. and Judd, K. L., editors, *Handbook of Computational Economics Volume 2 - Agent-Based Computational Economics*, chapter 24, pages 1187–1233. North-Holland, Amsterdam.
- Lee, H. (1999). Online stable matching as a means of allocating distributed resources. *Journal of Systems Architecture*, 45(15):1345–1355.
- Lehmann, D., Müller, R., and Sandholm, T. (2011). The Winner Determination Problem. In Cramton, P., Shoham, Y., and Steinberg, R., editors, *Combinatorial Auctions*, chapter 12, pages 297–317. MIT Press, Boston.

- Lewis, M. (2014). *Flash Boys: A Wall Street Revolt*. Simon & Schuster Audio.
- Leyton-Brown, K. and Shoham, Y. (2009). *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*. Cambridge University Press, New York.
- Leyton-Brown, K. and Shoham, Y. (2010). A Test Suite for Combinatorial Auctions. In Cramton, P., Shoham, Y., and Steinberg, R., editors, *Combinatorial Auctions*, chapter 18, pages 451–478. MIT Press, Cambridge.
- Lovasz, L. and Plummer, M. D. (2009). *Matching Theory*. AMS Chelsea Publishing, New York.
- Lyon, R. M. (1986). Equilibrium properties of auctions and alternative procedures for allocating transferable permits. *Journal of Environmental Economics and Management*, 13(2):129–152.
- Ma, J. (1996). On randomized matching mechanisms. *Economic Theory*, 8:377–381.
- Madhavan, A. (2000). Market microstructure: A survey. *Journal of Financial Markets*, 3(3):205–258.
- Manlove, D. F. (2013). *Algorithmics of Matching Under Preferences*, volume 2 of *Series on Theoretical Computer Science*. World Scientific Publishing, Singapore.
- Manlove, D. F. and O’Malley, G. (2012). Paired and Altruistic Kidney Donation in the UK: Algorithms and Experimentation. In *Lecture Notes in Computer Science*, volume 7276, pages 271–282. Springer, Berlin.
- Marks, R. E. (2006). Market Design Using Agent-Based Models. In Tesfatsion, L. and Judd, K. L., editors, *Handbook of Computational Economics Volume 2 - Agent-Based Computational Economics*, chapter 27, pages 1339–1380. North-Holland, Amsterdam.
- Matching in Practice (2016). Matching Practices in Europe. <http://www.matching-in-practice.eu/> (Accessed 27.12.2016).
- Matsui, T. (2011). Algorithmic Aspects of Equilibria of Stable Marriage Model with Complete Preference Lists. In Hu, B., Morasch, K., Pickl, S., and Siegle, M., editors, *Operations Research Proceedings 2010*, pages 47–52. Springer, Berlin, Heidelberg.
- McAfee, R. P. (1992). A dominant strategy double auction. *Journal of Economic Theory*, 56(2):434–450.

- Mehta, A. (2013a). Online Matching and Ad Allocation. *Foundations and Trends® in Theoretical Computer Science*, 8(4):265–368.
- Mehta, A. (2013b). Online Matching and Ad Allocation. *Foundations and Trends® in Theoretical Computer Science*, 8(4):265–368.
- Milgrom, P. (2004). *Putting Auction Theory to Work*. Cambridge University Press, Cambridge.
- Mochon, A., Saez, Y., Isasi, P., and Gomez-Barroso, J. (2009). Testing bidding strategies in the clock-proxy auction for selling radio spectrum: A Genetic Algorithm approach. In *Proceedings of the IEEE Congress on Evolutionary Computation*, pages 2348–2353. IEEE.
- Mortensen, D. T. and Pissarides, C. A. (1999). Job reallocation, employment fluctuations and unemployment. In Taylor, J. B. and Woodford, M., editors, *Handbook of Macroeconomics*, volume 1, chapter 18, pages 1171–1228. North Holland, Amsterdam.
- Myerson, R. B. (1981). Optimal Auction Design. *Mathematics of Operations Research*, 6(1):58–73.
- Narahari, Y., Garg, D., Narayanam, R., and Prakash, H. (2009). *Game Theoretic Problems in Network Economics and Mechanism Design Solutions*. Springer, London.
- Neugart, M. (2004). Endogeneous matching functions: and agent-based computational approach. *Advances in Complex Systems*, 07(02):187–201.
- Neugart, M. and Richiardi, M. (2012). Agent-based models of the labor market. Laboratorio Riccardo Revelli Working Paper no. 125.
- Nguyen, N.-T., Nguyen, T. T., Roos, M., and Rothe, J. (2013). Computational complexity and approximability of social welfare optimization in multiagent resource allocation. *Autonomous Agents and Multi-Agent Systems*, 28(2):256–289.
- Nicolaisen, J., Petrov, V., and Tesfatsion, L. (2001). Market power and efficiency in a computational electricity market with discriminatory double-auction pricing. *IEEE Transactions on Evolutionary Computation*, 5(5):504–523.
- Niederle, M. and Yariv, L. (2009). Decentralized matching with aligned preferences. NBER Working Paper Series 14840.

- Nisan, N. (2007). Introduction to Mechanism Design (for Computer Scientists). In Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V., editors, *Algorithmic Game Theory*, chapter 9, pages 209–241. Cambridge University Press, Cambridge.
- Nisan, N. (2010). Bidding Languages for Combinatorial Auctions. In Cramton, P., Shoham, Y., and Steinberg, R., editors, *Combinatorial Auctions*, chapter 9, pages 215–231. MIT Press, Boston.
- Nisan, N., Bayer, J., Chandra, D., Franji, T., Gardner, R., Matias, Y., Rhodes, N., Seltzer, M., Tom, D., Varian, H. R., and Zigmond, D. (2009). Google’s Auction for TV Ads. In Albers, S., Marchetti-Spaccamela, A., Matias, Y., Nikolettseas, S., and Thomas, W., editors, *Automata, Languages and Programming*, chapter 26, pages 309–327. Springer, Berlin.
- Nobelprize.org (1996). The Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel. Nobel Media AB 2014. http://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/1996/ (Accessed 20.12.2016).
- Nobelprize.org (2007). The Sveriges Riksbank Prize in Economic Sciences in Memory of Alfred Nobel. Nobel Media AB 2014. http://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/2007/ (Accessed 08.12.2016).
- Nobelprize.org (2010). The Prize in Economic Sciences. Nobel Media AB 2014. http://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/2010/ (Accessed 20.12.2016).
- Nobelprize.org (2012). The Prize in Economic Sciences. Nobel Media AB 2014. http://www.nobelprize.org/nobel_prizes/economic-sciences/laureates/2012/ (Accessed 08.12.2016).
- Nowak, M. A. and Sigmund, K. (1993). A strategy of win-stay, lose-shift that outperforms tit-for-tat in the Prisoner’s Dilemma game. *Nature*, 364:56–58.
- Oméro, M.-J., Dzierzawa, M., Marsili, M., and Zhang, Y.-C. (1997). Scaling Behavior in the Stable Marriage Problem. *Journal de Physique I*, 7(12):1723–1732.
- Osborne, M. J. and Rubinstein, A. (2011). *A Course in Game Theory*. MIT Press, Boston.
- Ostrovsky, M. and Schwarz, M. (2011). Reserve prices in internet advertising auctions. In *Proceedings of the 12th ACM conference on Electronic commerce*, pages 59–60, New York. ACM Press.

- Pais, J. (2008). Incentives in decentralized random matching markets. *Games and Economic Behavior*, 64(2):632–649.
- Pathak, P. A. and Sethuraman, J. (2011). Lotteries in student assignment: An equivalence result. *Theoretical Economics*, 6(1):1–17.
- Pathak, P. A. and Shi, P. (2013). Simulating Alternative School Choice Options in Boston. Technical report, MIT School Effectiveness and Inequality Initiative.
- Pathak, P. A. and Sönmez, T. (2008). Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism. *American Economic Review*, 98(4):1636–1652.
- Pathak, P. A. and Sönmez, T. (2013). School Admissions Reform in Chicago and England: Comparing Mechanisms by their Vulnerability to Manipulation. *American Economic Review*, 103(1):80–106.
- Patterson, S. (2012). *Dark Pools: The Rise of the Machine Traders and the Rigging of the U.S. Stock Market*. Random House Business Books, London.
- Petrongolo, B. and Pissarides, C. A. (2001). Looking into the black box: a survey of the matching function. *Journal of Economic Literature*, 39(2):390–431.
- Pittel, B. (1989). The Average Number of Stable Matchings. *SIAM Journal on Discrete Mathematics*, 2(4):530–549.
- Pöder, K. (2010). *Structural Solutions to Social Traps: Formal and Informal Institutions*. PhD thesis, Tallinn University of Technology.
- Pöder, K. and Lauri, T. (2014). When Public Acts like Private: the failure of Estonia’s school choice mechanism. *European Educational Research Journal*, 13(2):220–234.
- Pöder, K., Lauri, T., Karmo, K., Veski, A., Roosalu, T., and Simm, K. (2015). *Lasteaiakohtade jagamine Soovitused kohalikele omavalitsustele*. Gutenbergi Pojad, Tallinn.
- Pöder, K., Lauri, T., and Veski, A. (2016). Does School Admission by Zoning Affect Educational Inequality? A Study of Family Background Effect in Estonia, Finland, and Sweden. *Scandinavian Journal of Educational Research*, pages 1–21.
- Pöder, K., Veski, A., and Lauri, T. (2014). Eesti põhikooli- ja gümnaasiumivõrgu analüüs aastaks 2020. Technical report, PRAXIS Poliitikauuringute keskus.

- Preschool Child Care Institutions Act (2014). Riigikogu RT I, 13.03.2014, 4. <https://www.riigiteataja.ee/en/eli/517062014005> (Accessed 04.08.2016).
- Procaccia, A. D. (2009). Thou shalt covet thy neighbor's cake. In *Proceedings of the 21st International Joint Conference on Artificial Intelligence*, pages 239–244, Palo Alto. AAAI Press.
- Rawls, J. (1971). *A Theory of Justice*. Harvard University Press, Cambridge.
- Riccetti, L., Russo, A., and Gallegati, M. (2015). An agent based decentralized matching macroeconomic model. *Journal of Economic Interaction and Coordination*, 10(2):305–332.
- Richiardi, M. (2004). A search model of unemployment and firm dynamics. *Advances in Complex Systems*, 07(02):203–221.
- Richiardi, M. (2006). Toward a non-equilibrium unemployment theory. *Computational Economics*, 27(1):135–160.
- Riechmann, T. (2001a). Genetic algorithm learning and evolutionary games. *Journal of Economic Dynamics and Control*, 25(6-7):1019–1037.
- Riechmann, T. (2001b). *Learning in Economics: Analysis and Application of Genetic Algorithms*. Springer, Berlin.
- Robertson, J. and Webb, W. A. (1998). *Cake-Cutting Algorithms: Be Fair if You Can*. A. K. Peters, Natick.
- Roth, A. E. (1982). The Economics of Matching: Stability and Incentives. *Mathematics of Operations Research*, 7(4):617–628.
- Roth, A. E. (1984). Misrepresentation and stability in the marriage problem. *Journal of Economic Theory*, 34(2):383–387.
- Roth, A. E. (1985). The college admissions problem is not equivalent to the marriage problem. *Journal of Economic Theory*, 36(2):277–288.
- Roth, A. E. (1997). The Effects of the Change in the NRMP Matching Algorithm. *JAMA: The Journal of the American Medical Association*, 278(9):729.
- Roth, A. E. (2002). The Economist as Engineer: Game Theory, experimentation, and Computation as Tools for Design Economics. *Econometrica*, 70(4):1341–1378.

- Roth, A. E. (2008). Deferred acceptance algorithms: history, theory, practice, and open questions. *International Journal of Game Theory*, 36(3-4):537–569.
- Roth, A. E. (2015). *Who Gets What - and Why? Understand the Choices You Have, Improve the Choices You Make*. William Collins, London.
- Roth, A. E. and Peranson, E. (1999). The redesign of the matching market for American physicians: some engineering aspects of economic design. *American Economic Review*, 89(4):748–780.
- Roth, A. E., Sönmez, T., and Ünver, M. U. (2004). Kidney Exchange. *The Quarterly Journal of Economics*, 119(2):457–488.
- Roth, A. E., Sönmez, T., and Ünver, M. U. (2007). Efficient Kidney Exchange: Coincidence of Wants in a Structured Market with Compatibility-Based Preferences. *The American Economic Review*, 97(3):828–851.
- Roth, A. E. and Sotomayor, M. A. O. (1990). *Two-sided matching: a study in game-theoretic modeling and analysis*. Cambridge University Press, Cambridge.
- Roth, A. E. and Vate, J. H. V. (1990). Random paths to stability in two-sided matching. *Econometrica*, 58(6):1475–1480.
- Rothe, J., editor (2016). *Economics and Computation*. Springer Texts in Business and Economics. Springer, Berlin.
- Roughgarden, T. (2005). *Selfish Routing and the Price of Anarchy*. MIT Press, Cambridge, Massachusetts.
- Roughgarden, T., Nisan, N., Tardos, E., and Vazirani, V. V. (2007). *Algorithmic Game Theory*. Cambridge University Press, New York.
- Rust, J., Miller, J. H., and Palmer, R. (1993). Behavior of Trading Automata in a Computerized Double Auction Market. In Friedman, D. and Rust, J., editors, *The Double Auction Market: Institutions, Theories, and Evidence*, chapter 6, pages 155–198. Perseus Publishing, Cambridge.
- Rust, J., Miller, J. H., and Palmer, R. (1994). Characterizing effective trading strategies. *Journal of Economic Dynamics and Control*, 18(1):61–96.
- Rysman, M. (2009). The Economics of Two-Sided Markets. *Journal of Economic Perspectives*, 23(3):125–143.

- Saaty, T. L. (1978). Exploring the interface between hierarchies, multiple objectives and fuzzy sets. *Fuzzy Sets and Systems*, 1(1):57–68.
- Sahin, A., Song, J., Topa, G., and Violante, G. L. (2014). Mismatch unemployment. *American Economic Review*, 104(11):3529–3564.
- Sandholm, T. (2010). Optimal Winner Determination Algorithms. In Cramton, P., Shoham, Y., and Steinberg, R., editors, *Combinatorial Auctions*, chapter 14, pages 337–368. MIT, Boston.
- Saraceno, C. (2011). Family policies. Concepts , goals and instruments. Carlo Alberto Notebooks, (230).
- Satterthwaite, M. A. and Williams, S. R. (1989a). Bilateral trade with the sealed bid k-double auction: Existence and efficiency. *Journal of Economic Theory*, 48(1):107–133.
- Satterthwaite, M. A. and Williams, S. R. (1989b). The Rate of Convergence to Efficiency in the Buyer’s Bid Double Auction as the Market Becomes Large. *The Review of Economic Studies*, 56(4):477.
- Satterthwaite, M. A. and Williams, S. R. (1993). The Bayesian Theory of k-Double Auction. In Friedman, D. and Rust, J., editors, *The Double Auction Market: Institutions, Theories, and Evidence*, pages 99–124. Perseus Publishing, Cambridge.
- Schrijver, A. (2003). *Combinatorial Optimization: Polyhedra and Efficiency*. Springer, Berlin.
- Schummer, J. and Vohra, R. V. (2007). Mechanism Design without Money. In Nisan, N., Roughgarden, T., Tardos, E., and Vazirani, V. V., editors, *Algorithmic Game Theory*, chapter 10, pages 243–266. Cambridge University Press, New York.
- Schwind, M., Stockheim, T., and Rothlauf, F. (2003). Optimization heuristics for the combinatorial auction problem. In *Proceedings of the 2003 Congress on Evolutionary Computation*, volume 3, pages 1588–1595.
- Segal-Halevi, E., Hassidim, A., and Aumann, Y. (2016). SBBA: a Strongly-Budget-Balanced Double-Auction Mechanism. <http://arxiv.org/abs/1607.05139> (Accessed 16.08.2016).
- Shi, P. (2015). Guiding School-Choice Reform through Novel Applications of Operations Research. *Interfaces*, 45(2):117–132.

- Shimer, R. (2013). Job search, labour force participation, and wage rigidities. In Acemoglu, D., Arellano, M., and Dekel, E., editors, *Advances in Economics and Econometrics: Theory and Applications: Tenth World Congress*, chapter 5, pages 197–234. Cambridge University Press, New York.
- Sigmund, K. (2010). *The Calculus of Selfishness*. Princeton University Press, Princeton.
- Silva, S. T., Valente, J. M. S., and Teixeira, A. A. C. (2012). An evolutionary model of industry dynamics and firms’ institutional behavior with job search, bargaining and matching. *Journal of Economic Interaction and Coordination*, 7(1):23–61.
- Simon, D. (2013). *Evolutionary optimization algorithms: biologically-inspired and population-based approaches to computer intelligence*. Wiley, Hoboken.
- Smith, A. (1776). *An Inquiry into the Nature and Causes of the Wealth of Nations*. Metalibri, London, 2007 edition.
- Smith, E., Farmer, J. D., Gillemot, L., and Krishnamurthy, S. (2003). Statistical theory of the continuous double auction. *Quantitative Finance*, 3(6):481–514.
- Sönmez, T. (1997). Manipulation via Capacities in Two-Sided Matching Markets. *Journal of Economic Theory*, 77(1):197–204.
- Sönmez, T. (1999). Can Pre-arranged Matches Be Avoided in Two-Sided Matching Markets? *Journal of Economic Theory*, 86(1):148–156.
- Sornette, D. (2004). *Why Stock Markets Crash: Critical Events in Complex Financial Systems*. Princeton University Press, Princeton.
- Sotomayor, M. (2012). A further note on the college admission game. *International Journal of Game Theory*, 41(1):179–193.
- Steinhaus, H. (1948). The Problem of Fair Division. *Econometrica*, 16(1):101–104.
- Steinhaus, H. (1969). *Mathematical Snapshots*. Oxford University Press, New York, 3rd edition.
- Sterling, L. S. and Taveter, K. (2009). *The Art of Agent-Oriented Modeling*. MIT Press, Cambridge, MA.
- Stromquist, W. (2008). Envy-free cake divisions cannot be found by finite protocols. *The Electronic Journal of Combinatorics*, 15:R11.

- Sureka, A. and Wurman, P. R. (2005). Applying metaheuristic techniques to search the space of bidding strategies in combinatorial auctions. In Beyer, H.-G., editor, *Proceedings of the 7th annual conference on Genetic and evolutionary computation*, pages 2097–2103, New York. ACM Press.
- Tassier, T. and Menczer, F. (2008). Social network structure, segregation, and equality in a labor market with referral hiring. *Journal of Economic Behavior and Organization*, 66(3-4):514–528.
- Teo, C.-P., Sethuraman, J., and Tan, W.-P. (2001). Gale-Shapley Stable Marriage Problem Revisited: Strategic Issues and Applications. *Management Science*, 47(9):1252–1267.
- Tesfatsion, L. and Judd, K. L. (2006). *Handbook of Computational Economics, Volume 2: Agent-Based Computational Economics*. North-Holland, Amsterdam.
- Ünver, M. U. (2001). Backward unraveling over time: The evolution of strategic behavior in the entry level British medical labor markets. *Journal of Economic Dynamics and Control*, 25(6-7):1039–1080.
- Ünver, M. U. (2005). On the survival of some unstable two-sided matching mechanisms. *International Journal of Game Theory*, 33(2):239–254.
- Van Essen, M. (2013). An Equilibrium Analysis of Knaster’s Fair Division Procedure. *Games*, 4(1):21–37.
- Varian, H. R. (2006). *Intermediate Microeconomics*. W. W. Norton & Company, New York, 7th edition.
- Veracierto, M. (2011). Worker flows and matching efficiency. *Economic Perspectives*, 35(4):147–169.
- Veski, A. (2012). Some issues in Multi-Agent Resource Allocation. In *Proceedings of the 6th Annual Conference of the Estonian National Doctoral School in Information and Communication Technologies*, pages 101–104, Tallinn. TUT Press.
- Veski, A. (2014). Price of Invisibility: Statistics of centralised and decentralised matching markets. In MacKerrow, E., Terano, T., Squazzon, F., and Sichman, J. S., editors, *Proceedings of the 5th. World Congress on Social Simulation*, pages 18–29, Sao Paulo.
- Veski, A., Biro, P., Pöder, K., and Lauri, T. (2017). (forthcoming) Efficiency and fair access in kindergarten allocation policy design. *Journal of Mechanism and Institution Design*.

- Veski, A. and Pöder, K. (2015). Primary School Choice in Tallinn: Data and Simulations. TUTECON Working Paper No. WP-2015/1.
- Veski, A. and Pöder, K. (2016). Strategies in Tallinn school choice mechanism. *Research in Economics and Business: Central and Eastern Europe*, 8(1):5–24.
- Veski, A. and Pöder, K. (2017). Zero-intelligence agents looking for a job. *Journal of Economic Interaction and Coordination*.
- Veski, A. and Vöhandu, L. (2010). Another View on Territory Fair Division. In Barzdins, J. and Kirikova, M., editors, *Databases and information systems : proceedings of the Ninth International Baltic Conference*, pages 261–276, Riga. University of Latvia Press.
- Veski, A. and Vöhandu, L. (2011). Two Player Fair Division Problem with Uncertainty. In Barzdins, J. and Kirikova, M., editors, *Frontiers in Artificial Intelligence and Applications*, pages 394–407. IOS Press, Amsterdam.
- Veskioja, T. (2005). *Stable Marriage Problem and College Admission*. PhD thesis, Tallinn University of Technology.
- Vickrey, W. (1961). Counterspeculation, auctions, and competitive sealed tenders. *The Journal of Finance*, 16(1):8–37.
- Vriend, N. J. (2000). An illustration of the essential difference between individual and social learning, and its consequences for computational analyses. *Journal of Economic Dynamics and Control*, 24(1):1–19.
- Wang, L., Liu, S., Lu, C., Zhang, L., Xiao, J., and Wang, J. (2015). Stable Matching Scheduler for Single-ISA Heterogeneous Multi-core Processors. In Chen, Y., Ienne, P., and Ji, Q., editors, *Advanced Parallel Processing Technologies. APPT 2015. Lecture Notes in Computer Science, vol 9231*, pages 45–59. Springer, Cham.
- Wang, S., Xu, P., Xu, X., Tang, S., Li, X., and Liu, X. (2010). TODA: Truthful Online Double Auction for Spectrum Allocation in Wireless Networks. In *Proceedings of IEEE Symposium on New Frontiers in Dynamic Spectrum*, pages 1–10. IEEE.
- Wellman, M. P., Osepayshvili, A., MacKie-Mason, J. K., and Reeves, D. (2008). Bidding Strategies for Simultaneous Ascending Auctions. *The B.E. Journal of Theoretical Economics*, 8(1).
- West, A., Hind, A., and Pennell, H. (2004). School admissions and 'selection' in comprehensive schools: policy and practice. *Oxford Review of Education*, 30(3):347–369.

- Wilensky, U. (1999). NetLogo. <http://ccl.northwestern.edu/netlogo/>. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston.
- Wilensky, U. (2002). NetLogo PD N-Person Iterated model. <http://ccl.northwestern.edu/netlogo/models/PDN-PersonIterated>. Center for Connected Learning and Computer-Based Modeling, Northwestern University, Evanston.
- Zhang, Y.-C. (2001). Happier world with more information. *Physica A: Statistical Mechanics and its Applications*, 299(1-2):104–120.
- Zhou, B., He, Z., Jiang, L.-L., Wang, N.-X., and Wang, B.-H. (2014a). Bidirectional selection between two classes in complex social networks. *Scientific Reports*, 4:7577.
- Zhou, B., Qin, S., Han, X.-P., He, Z., Xie, J.-R., and Wang, B.-H. (2014b). A model of two-way selection system for human behavior. *PLoS ONE*, 9(1):e81424.

A Publication 1

Veski, A. and Võhandu, L. (2011). Two Player Fair Division Problem with Uncertainty. In Barzdins, J. and Kirikova, M., editors, *Frontiers in Artificial Intelligence and Applications*, pages 394–407. IOS Press, Amsterdam

Two Player Fair Division Problem with Uncertainty¹

Andre VESKI^{a,2}, Leo VOHANDU^{a,3}

^a *Ehitajate tee 5, Tallinn, Estonia*
Technical University of Tallinn, Institute of Informatics

Abstract. This paper analyses the territory fair division, problem initially posed by Hugo Steinhaus [1], by studying the solutions given by different algorithms on a large generated set of inputs for two players. Main algorithm used is Adjusted Winner, developed by S. Brams and A. Taylor [2]. We compare it to combinatorial enumeration and some algorithms proposed for experimentation by authors. Additionally we define measures to characterize the initial task and game theoretic measures to select the best solution. Moreover we extend the problem by allowing uncertainties in the players' value representation of items to be divided, based on the example of territorial division. For uncertainty management we use the belief system from Dempster-Shafer Theory [3].

Keywords. Fair division, Dempster-Shafer Theory

Introduction

Division problem is a more general task of the well known partitioning problem. Several different techniques have been developed to provide a fair division of goods or a single homogeneous good (e.g. cake). In general there are two classes of techniques: turn based information sharing and decision making; or providing full information to a mediator who then proposes a solution. In this paper we concentrate on the later type of algorithms, where we have all the information available from all players.

At first it may seem that a fair division is unreachable, because each player has their own subjective opinion about the goods being distributed. Ultimately just because of those subjective estimations, a fair division is possible. Unfortunately most of the world relies on experts' *objective* opinions and not on their own attitudes.

Initially this problem was described by Steinhaus [1], [4] as the problem of fair division. In his example two heirs have a territory to divide and both expect to get half of it. In Steinhaus' solution they both draw a vertical dividing line that would split the territory into two subjectively equal parts [1] (Figure 1). At first both players receive a piece of the territory they value more per area unit. In our example I_1 would be attributed to bidder B and I_3 to A . The remaining piece, I_2 , can be divided similarly or split randomly, since each of the players has already gained a half.

¹The final publication is available at IOS Press through <http://dx.doi.org/10.3233/978-1-60750-688-1-394>

²E-mail: andre.veski@gmail.com

³E-mail: leov@staff.ttu.ee




Figure 1. Division

To characterize fitness of a fair distribution, we use three measures proposed by Brams and Taylor [2]: *envy*, *equality* and *efficiency*. Envy shows how much a participant desires some one else's gains. Efficiency illustrates the final results of all players. A division is efficient when every participant got at least what he bargained for – with two player half of his initial values. Equality shows similarity in each participant's total gain. We will give a formal definition of these measures later in the paper. **In order for the division to be fair it has to be as envy-free, efficient and equal as possible.** There are cases where these measures contradict each-other and one can not drive all of them to the maximum. We also look how to fuse these three measures and use a single measure to evaluate the result such as Nash's Bargaining Solution.




Figure 2. Conflicting division

In our paper we also add a level of generalization to the fair division problematique, based on an example of territory division. Instead of having non-crossing division lines as in Figure 1 we introduce crossing lines as on Figure 2. Hence we get a total of four items to be divided, but valuations only for item sets of two. This means we have some level of uncertainty. While we know how participants value part of their division, we don't know how it translates to each of the four individual sections as depicted on Figure 2.

1. More Formal Description of the Problem

We have a set S of n bidders S_1, \dots, S_n and a set I of m items I_1, \dots, I_m . Each bidder has a real-valued valuation function v_i that for each item $I_j \in I$ gives the value $v_i(I_j)$ that bidder S_i obtains if he receives I_j . A division of items among the bidders is a matching mutually exclusive subsets of items to bidders. For every bidder S_i, S_j where $i \neq j$ $D_i, D_j \subseteq I$ is an allocation such that $D_i \cap D_j = \emptyset$. For all bidders $\bigcup_{S_i \in S} D_i = I$. The total utility obtained is defined by $\sum_{S_i \in S} v_i(D_i)$.

As already mentioned, the problem of fair division is similar to the known NPC problem of partitioning, where items of different value need to be partitioned into n

Table 1. Valuations

v	I_1	I_2	I_3
v_A	0.44	0.6	0.50
v_B	0.50	0.5	0.45

distinct sets, all with equal total value. In fair division case, each bidder can be considered as a partition. And for each bidder $S_i \in S$ the value functions are the same, meaning that for any i, j and l the $v_i(I_l) = v_j(I_l)$. Often total equality between partitions is not achievable then we need to use a measure of fitness – equality between partitions. Having m items to be partitioned into two sets D_1 and D_2 , the goal is to minimize the difference

$$Equality(D_1, D_2) = \left| \sum_{I_i \in D_1} v_1(I_i) - \sum_{I_i \in D_2} v_2(I_i) \right| \quad (1)$$

For example having items with values $\{1, 2, 3\}$ then these can be easily divided into two sets of equal value $\{1, 2\}$ and $\{3\}$. The simplest algorithm to solve this task is Algorithm 1.1. Start from the largest element, assign it to a set, then second largest element to the set with a lower value, until there are no more items left. Partitioning problem can be generalized by adding some degrees of freedom. For example in fair division we have to take into account differences of opinion, i.e. the same item can have different values depending on the set they are assigned to – A or B . For simplicity, in the rest of the paper we will look at a problem with two bidders denoted $A = S_A$ or $B = S_B$.

Algorithm 1.1: MFP(I)

```

[Initialize]  $A, B \leftarrow \{\}$ 
while not empty  $I$ 
  do  $\left\{ \begin{array}{l} \text{[Find maximum] } \max\{v(I_1), \dots, v(I_m)\}, \text{ and } I_l \text{ is the maximal} \\ \text{[Select lower value set] if } Total(v(A)) \leq Total(v(B)) \\ \quad \text{then } A[\text{elements} + 1] \leftarrow I_l \\ \quad \text{else } B[\text{elements} + 1] \leftarrow I_l \end{array} \right.$ 
return  $(A, B)$ 

```

Let us take an example of fair division as presented on Figure 1. Assume that A and B have the value functions for items $\{I_1, I_2, I_3\}$ as in Table 1. The valuations are normalized so that for both A and B the total value would be 1. A good algorithm to solve this has been developed by Brams and Taylor [2], called Adjusted Winner (Algorithm 1.2), which has some similarities with Algorithm 1.1. In the first round the main goal is to return maximal efficiency by assigning each item to the highest bidder. In the second, adjustment step the result is equalized by giving most similarly valued items to the worst-off player. The algorithm tries to optimize all three criteria mentioned above.




Figure 3. Valuations

Alternatively we may aim for highest utility as in Algorithm 1.3, which, as we shall see later, will in some cases yield a better result.

Algorithm 1.2: AWT(I)

```

[Initialize]  $A, B \leftarrow \{\}, m \leftarrow |I|, v \leftarrow 0$ 
for  $i \leftarrow 1$  to  $m$ 
  do  $\left\{ \begin{array}{l} \text{[Select largest and add]} \\ \text{if } v_A(I_i) > v_B(I_i) \\ \quad \text{then } A[\text{elements} + 1] \leftarrow I_i \\ \quad \text{else } B[\text{elements} + 1] \leftarrow I_i \end{array} \right.$ 
[Adjust A and B to be equal valued] if  $Total(v_A(A)) > Total(v_B(B))$ 
  then repeat
    for  $i \leftarrow 1$  to  $n$ 
      do  $\left\{ \begin{array}{l} c \leftarrow 2/(v_A(I_i)/v_B(I_i)) \\ \text{[Select best] if } v < c \\ \quad \text{then } v \leftarrow c, s \leftarrow i \\ B[\text{elements} + 1] = I_s \end{array} \right.$ 
  until  $Total(v_A(A)) \leq Total(v_B(B))$ 
  else if  $Total(v_A(A)) < Total(v_B(B))$ 
    then [Exchange]  $A \leftrightarrow B, \text{ancontinueatstep[Adjust]}$ 
return  $(A, B)$ 
  
```

Algorithm 1.3: MFT(I)

```

[Initialize]  $A, B \leftarrow \{\}, J \leftarrow \{I_i | i = \{1, \dots, m\}, |I_i| = 1\}, o \leftarrow |J|$ 
for  $i \leftarrow 1$  to  $o$ 
  do  $\left\{ \begin{array}{l} \text{[Select largest and add] if } v_A(J_i) > v_B(J_i) \\ \quad \text{then } A[\text{elements} + 1] \leftarrow J_i \\ \quad \text{else } B[\text{elements} + 1] \leftarrow J_i \end{array} \right.$ 
return  $(A, B)$ 
  
```

Let's look at the case where participant's valuations are uncertain, example on Figure 3. We have values for player A , with the value function $v_A(\{I_1, I_2\}) = v_A(I_1 \cup I_2) =$

0.5 and for individual items $v_A(\{I_1\}) = 0.42$ and $v_A(\{I_2\}) = 0.6$. The leftover value $v_A(I_1 \cap I_2) = v_A(\{I_1, I_2\}) - v_A(\{I_1\}) - v_A(\{I_2\}) = 0.5 - 0.42 - 0.06 = 0.02$ is the value that comes from owning both of these items together. In other words, it is the value of the connection between those two items. In Dempster-Shafer Theory (henceforth DST) [3] we have upper and lower bound values for all items. Going forward, we describe briefly the necessary part of DST [3]. On a valuation lattice (Figure 3) there are at least three ways of looking at each item set:

1. Total value of items which an item contains - belief
2. Total value of items in which an item is part of - plausibility
3. Total value of items in which an item is not part of and does not contain itself - doubt

The base values for each item or set of items is called mass or basic probability assignment [3] and is the value only for that particular item. The mass values for all item sets in the power set add up to 1 and value of an empty set is 0. Meaning that for single items the mass is the value of that item and in larger item sets the value of the connection (intersection). In our example (Figure 3) the mass $m_A(I_1, I_2) = v_A(I_1 \cup I_2) = 0.02$ and the belief value $b_A(I_1, I_2) = v_A(I_1, I_2) = 0.5$. The belief function from DST corresponds to our idea of a value function so that $b_A(I_i) = v_A(I_i)$. Here $\sum_{I_i \in I} m_A(I_i) = 1$, whereas $\sum_{I_i \in I} v_A(I_i) \geq 1$. The definitions for item set value functions are given below.

Definition 1. Belief is a sum of masses from all the sets where observed element I_i is a part of or in other words a lower bound for an item set A that contains the certain knowledge about the item set. On Figure 4 we have belief for $I_6 = \{I_2, I_3\}$ presented with gray background.

$$Belief(I_i) = b(I_i) = \sum_{I_j \subset I_i} v(I_j) \quad (2)$$

Definition 2. Plausibility is a sum of masses from all the item sets where the union with the observed element I_i is not an empty set or in other words an upper bound for a set I_i that I_i would have if it got assigned all the uncertain values. On Figure 5 there is plausibility for I_1 with a gray background.

$$Plausibility(I_i) = p(I_i) = \sum_{I_j \cap I_i \neq \emptyset} v(I_j) \quad (3)$$

Definition 3. Doubt is a sum of masses from all the item sets where the union with the observer element I_i is an empty set or value that I_i can never have even if it got assigned all the uncertain values. On Figure 4 there is doubt for I_1 presented with a gray background.

$$Doubt(I_i) = d(I_i) = \sum_{I_j \cap I_i = \emptyset} v(I_j) \quad (4)$$




Figure 4. Belief I_6 , Doubt I_1




Figure 5. Plausibility I_1

2. Measuring Solutions and Tasks

With each additional degree of freedom there are more possibilities to evaluate the fitness of a solution. On the simple partitioning task there is a measure of equality between sets. As the number of partitions grows, new measures come up such as a total difference of values in pairwise comparison etc.

A fair division is usually described by three measures according to Brams and Taylor [2]: efficiency, envy and equality and we also use total product. The latter is also used by Nguyen and Kreinovich [5] and is known in the game theory as Nash's Bargaining Solution [6].

Definition 4. (Pareto) Efficiency is the total value of the solution. This is actually a condition where no player can be made better off by making someone else worse off. But if an item can change owners and create more value from that, the initial owner can be compensated by some uniform measure – e.g. money.

$$Efficiency(A, B) = ef(A, B) = \sum_{I_i \in A} m_A(I_i) + \sum_{I_i \in B} m_B(I_i) \quad (5)$$

Definition 5. Envy is the amount by which in one player's valuations other players result was larger than his. The total envy is the total sum on pairwise comparisons and consists of two parts:

$$Envy(A, B) = en(A, B) = \max \left(\left(\sum_{I_i \in B} m_A(I_i) - \sum_{I_i \in A} m_A(I_i) \right), 0 \right) \quad (6)$$

$$Envy(B, A) = en(B, A) = \max \left(\left(\sum_{I_i \in A} m_B(I_i) - \sum_{I_i \in B} m_B(I_i) \right), 0 \right) \quad (7)$$

$$TotalEnvy(A, B) = Envy(A, B) + Envy(B, A) \quad (8)$$

Definition 6. Equality is the amount by which end results differ for each player and is calculated as a sum of pairwise differences

$$Equality(A, B) = eq(A, B) = \left| \sum_{I_i \in A} m_A(I_i) - \sum_{I_i \in B} m_B(I_i) \right| \quad (9)$$

Definition 7. Product is the total product of all players end results

$$Product(A, B) = pr(A, B) = \sum_{I_i \in A} m_A(I_i) \cdot \sum_{I_i \in B} m_B(I_i) \quad (10)$$

Envy and equality are somewhat similar. In partitioning task we have equality as our measure of fitness. Since each subset of items has now different values based on every players individual preferences, we also need the envy measure to assess fitness from all players viewpoints.

Efficiency is a new kind of measure. With simple partitioning this measure is equal to the total value of items and is always the same, regardless of the actual solution. In fair division, since the total values of partitioned subsets are different in each solution, it is important to make sure that we would get the maximum possible total value.

Similarly measuring the goodness of a solution, there are measures to characterize the initial players valuations. We will use these to generalize the input in order to create a relationship with the output.

Definition 8. Conflict is defined as a conflict measure in the DST [3]

$$Conflict(A, B) = c(A, B) = 1 - \sum_{i=1}^n \sum_{j=i, I_i \cap I_j = \emptyset}^n m_A(I_i) \cdot m_B(I_j) \quad (11)$$

Definition 9. Difference is a pairwise difference in participants' valuations. This is recommended by authors and as we see later it has a good correlation with the solution.

$$Difference(A, B) = d(A, B) = \frac{\sum_{i=1}^n |v_A(I_i) - v_B(I_i)|}{2} \quad (12)$$

Difference describes the total difference in players' value functions. Looking at the Algorithm 1.2, the more different the players' value functions are, the more efficient the solution should be. In extreme cases where valuations are completely opposite, the resulting gain would be double of that the players initially subjectively expected.

Definition 10. Uncertainty is an amount of uncertainty in the valuations as total sum of mass on item sets with greater volume than 2.

$$Uncertainty() = un() = \sum_{i=1, |I_i| > 1}^n m_A(I_i) + m_B(I_i) \quad (13)$$

In Table 2 we have presented an example calculation of definitions given above. It is based on problem from Table 1 where the result is for $A = \{I_1\}$ and for $B = \{I_2, I_3\}$.

Table 2. Example calculations

	Player A	Player B
Result	0.50, 0.6	0.50
Efficiency	$0.6 + 0.50 + 0.50 = 1.06$	
Envy	$0.67 - (0.27 + 0.6) = 0.34$	$0.61 + 0.5 - 0.34 = 0.33$
Equality	$0.67 - 0.61 - 0.05 = 0.01$	
Product	$0.56 \cdot 0.50 = 0.28$	
Conflict	$1 - 0.44 \cdot 0.05 - 0.44 \cdot 0.45 - 0.06 \cdot 0.5 - 0.06 \cdot 0.45 - 0.5 \cdot 0.5 - 0.5 \cdot 0.05 = 0.448$	
Difference	$\frac{1}{2} \cdot (0.44 - 0.50 + 0.6 - 0.5 + 0.50 - 0.45) = 0.06$	
Uncertainty	0.0	

3. Algorithms Modifications for Uncertainty

So far we have examined two algorithms for fair division: Maximal First and Adjusted Winner for two bidders. As a next step we need to figure out how to handle uncertain valuations in these algorithms. In authors' views handling uncertainty could be done in at least two ways.

1. Using different value functions from DST to determine a value for item comparison, e.g. belief or plausibility
2. Start the division process with different volumes sets of items, e.g. on the second level of the valuation lattice where all item sets have a cardinality of two. Without uncertainty, all item sets have a cardinality of one.

In the paper we will explore only the first. More precisely we shall compare results from three algorithms.

1. Enumerating all possible combinations (Algorithm CT and Algorithm C)
2. 2-step Adjusted Winner (Algorithm AWT and Algorithm AW)
3. Maximal valued First (Algorithm MFT and Algorithm MFL)

In other words, all algorithms will be used in two models.

1. Model with uncertainty (Algorithms C, AW and MFL)
2. Model without uncertainty (Algorithms CT, AWT and MFT)

Algorithms AW (Algorithm 3.1) and MFL (Algorithm 3.2) are modifications of their certain world counterparts AWT (Algorithm 1.2) and MFT (Algorithm 1.3). With AW we still loop only over single items, but when comparing the item sets we use the plausibility function. We also add a step to check if we can add some item sets already covered by single items. In generalizing MFT to MFL we use the belief measure, not just mass, but

again we loop through single items and add larger item sets at the final stage.

Algorithm 3.1: AW(I)

```

[Initialize]  $A, B \leftarrow \cdot, J \leftarrow \{I_i | i = \{1, \dots, m\}, |I_i| = 1\}, o \leftarrow |J|$ 
for  $i \leftarrow 1$  to  $o$ 
  do  $\left\{ \begin{array}{l} \text{[Select largest and add] if } p_A(J_i) > p_B(J_i) \\ \text{then } A[\text{elements} + 1] \leftarrow J_i \\ \text{else } A[\text{elements} + 1] \leftarrow J_i \end{array} \right.$ 
[Adjust]  $v \leftarrow 0$ 
if  $Total(v_A(A)) > Total(v_B(B))$ 
  then repeat
    for  $i \leftarrow 1$  to  $o$ 
       $\left\{ \begin{array}{l} a \leftarrow b_A(J_i) b \leftarrow b_B(J_i) c \leftarrow 2/(a/b) \\ \text{do } \left\{ \begin{array}{l} \text{[Keep best matching item] if } v < c \\ \text{then } v \leftarrow c, s \leftarrow i \end{array} \right. \\ A[i] \leftarrow nil, B[i] = J_i \end{array} \right.$ 
    until  $Total(v_A(A)) \leq Total(v_B(B))$ 
    else if  $Total(v_A(A)) < Total(v_B(B))$ 
      then [Exchange]  $A \leftrightarrow B$ , and continue at step[Adjust]
 $K \leftarrow \{I_i | i = \{1, \dots, m\}, |I_i| > 1\}, q \leftarrow |J|$ 
for  $i \leftarrow 1$  to  $q$ 
   $\left\{ \begin{array}{l} \text{if } K_i \subset A \\ \text{do } \left\{ \begin{array}{l} \text{then } A[\text{elements} + 1] \leftarrow K_i \\ \text{else if } K_i \subset B \\ \text{then } B[\text{elements} + 1] \leftarrow K_i \end{array} \right. \end{array} \right.$ 
return  $(A, B)$ 

```

Algorithm 3.2: MFL(I)

```

[Initialize]  $A, B \leftarrow \cdot, J \leftarrow \{I_i | i = \{1, \dots, m\}, |I_i| = 1\}, o \leftarrow |J|$ 
for  $i \leftarrow 1$  to  $o$ 
  do  $\left\{ \begin{array}{l} \text{[Select largest and add] if } b_A(J_i) > b_B(J_i) \\ \text{then } A[\text{elements} + 1] \leftarrow I_i \\ \text{else } B[\text{elements} + 1] \leftarrow I_i \end{array} \right.$ 
 $K \leftarrow \{I_i | i = \{1, \dots, m\}, |I_i| > 1\}, q \leftarrow |J|$ 
for  $i \leftarrow 1$  to  $q$ 
   $\left\{ \begin{array}{l} \text{if } K_i \subset A \\ \text{do } \left\{ \begin{array}{l} \text{then } A[\text{elements} + 1] \leftarrow K_i \\ \text{else if } K_i \subset B \\ \text{then } B[\text{elements} + 1] \leftarrow K_i \end{array} \right. \end{array} \right.$ 
return  $(A, B)$ 

```

Table 3. Valuation example

	I_1	I_2	I_3	I_4	I_1, I_2	I_3, I_4	I_1, I_3	I_2, I_4
Player A	0	0	0	0	0.5	0.5	0	0
Player B	0	0	0	0	0	0	0.5	0.5

Algorithm 3.3: $C/CT(I)$

```

[Initialize]  $l \leftarrow 0, A, B \leftarrow \{\}$ 
repeat
   $A \leftarrow \text{Combination of Items}(I)$ 
   $B \leftarrow I - A$ 
  [Compare the new solution with the current solution] if  $pr(A, B) > l$ 
    then  $l \leftarrow pr(A, B)$ 
[All combinations tested]
until  $\text{OutOfCombinations}()$ 
return  $(A, B)$ 

```

4. Algorithm Comparison Setup

To compare the results for the algorithms we look at a large set of generated inputs, basically taking assumptions for both player on their values for single items (I_1, I_2, I_3 and I_4). Our goal is to look at many different settings. Some of them have uncertainty, some of them are closely valued and some are very different. To make sure we have all the possible settings we generate a full set of combinations. Since $v_A(\{I_1, I_2\}) = v_A(\{I_3, I_4\}) = 0.5$ will hold and respectively $v_B(\{I_1, I_3\}) = v_B(\{I_2, I_4\}) = 0.5$. If we know that $v_A(I_1) = 0.1$ then we also know that $v_A(I_2)$ can be at most 0.4). Meaning that for the valued item set $H = \{H_1, H_2\}$ we have $6 + 5 + 4 + 3 + 2 + 1 = 21$ ways to share the $v_A(H) = 0.5$ between the individual items H_1 and H_2 . Since there are two such item sets for both players, we get $21^{2*2} = 194481$ examples. We don't need examine all the examples. To reduce the amount of computation we skip mirrored valuations, such that $v_{B_2}(I_1) = v_{A_1}(I_1), v_{B_2}(I_3) = v_{A_1}(I_2), v_{B_2}(I_2) = v_{A_1}(I_3), v_{B_2}(I_4) = v_{A_1}(I_4), v_{A_2}(I_1) = v_{B_1}(I_1), v_{A_2}(I_2) = v_{B_1}(I_3), v_{A_2}(I_3) = v_{B_1}(I_2)$ and $v_{A_2}(I_4) = v_{B_1}(I_4)$. Assuming that Player A is A_1 and A_2 from example 1 and example 2 respectively and B_1 and B_2 Player B. This leaves us with 97461 examples.

1. Table 3 example has high uncertainty and high conflict.
2. Table 4 example has no uncertainty, has high conflict and high difference.
3. Table 5 has some conflict, some uncertainty and by our definition no difference.

Additionally to compare algorithms that can't handle uncertainty, we need to transform the input. For this we use the plausibility measure. For all single items we calculate their plausibility and then normalize that to add up to 1. This calculation basically distributes the uncertainty uniformly between the items.

Table 4. Valuation example

	I_1	I_2	I_3	I_4	I_1, I_2	I_3, I_4	I_1, I_3	I_2, I_4
Player A	0.5	0	0.5	0	0	0	0	0
Player B	0	0.5	0	0.5	0	0	0	0

Table 5. Valuation example

	I_1	I_2	I_3	I_4	I_1, I_2	I_3, I_4	I_1, I_3	I_2, I_4
Player A	0.2	0.2	0.2	0.2	0.1	0.1	0	0
Player B	0.2	0.2	0.2	0.2	0	0	0.1	0.1

Table 6. Correlations

	Efficiency	Envy	Inequality	Product
Efficiency	1			
Envy	-0.21	1		
Inequality	0.18	0.80	1	
Product	0.65	-0.76	-0.59	1

5. Assessing Metrics

Based on the metrics presented, which would be the best to maximize? We have presented Envy, Efficiency, Equality and Nash's Bargaining Solution. In Brams and Taylor [2] Adjusted Winner algorithm efficiency is preferred to reach at least 1 and after that the aim is to equalize the solution. In general we wish to go in all directions at once: minimize envy, maximize efficiency and equality. It appears that Nash's Bargaining Solution does just that.

On Figure 6 we can see the relation between all the four metrics. These values have been generated from all possible solutions (16) to our set of 97461 examples using Algorithm 3.3. This results in total about 1.6 million cases. In a rough summary the relationships between the metrics are (based on Figure 6):

1. Product has a good correlation with Efficiency (Table 6)
2. The highest Product is almost always with low Envy. There are always tasks where Envy is unavoidable, the extreme case being the example in Table 3 in the previous section
3. With a higher inequality the product approaches 0, as can be seen from the chart (Figure 6) starting from bottom left and moving to upper right

Since all the correlations with the product are in the right direction (Table 6) as stated: maximizing efficiency and equality and minimizing envy and all correlation coefficients with Nash's Bargaining Solution are above 0.5 which makes it a good metric to assess the fitness of our solutions. Moreover there is more justification on using the Nash's Bargaining solution in [6] and [7].




Figure 6. Solution metrics

6. Initial Results

Solving tasks in the model without uncertainty, the resulting solution is usually well predicted by the initial difference of opinions between the players. The greater the difference the better the result will be as on Figure 7.

As expected, adding uncertainty creates additional level of complexity. The differences of opinion don't have such strong impact on the result as before. On Figure 8 we see that the difference itself does not give as good an explanation as before. Difference still has a similar direction as in the model without uncertainty, the best result is achieved only with the biggest difference, but not the other way around i.e. larger difference does not guarantee a better solution. Moreover from Figure 8 we see that there is no definite impact by uncertainty and conflict either. Although there is visible direction with uncertainty, i.e. the greater uncertainty implies lower product value, but with decreasing uncertainty the variability of the Nash's Bargaining Solution increases.

Next let us compare results from different algorithms, with (Figure 10) and without (Figure 9) uncertainty. On the certain world (Figure 9) we see that all the algorithms produce quite similar results. Most of the results from algorithms AWT and MFT produce results quite near to the ideal, both have some deviation from CT, but not very much. In general there is room for improvement for both algorithms, although AWT manages always to beat MFT.

Adding uncertainty to algorithms the best algorithm is not that obvious anymore. When comparing algorithms AW and MFL it is visible on Figure 10 that MFL is in some




Figure 7. Difference without Uncertainty




Figure 8. Difference with Uncertainty

cases able to achieve better results than AW. This means that second stage in Algorithm AW, adjustment for equality, has some negative impact on other metrics in terms of Nash’s Bargaining Solution. This is definitely one of the places to start creating an even better algorithm.




Figure 9. Results without Uncertainty




Figure 10. Results with Uncertainty

7. Conclusion

It has been confirmed once more that Nash’s Bargaining Solution contains all the necessary properties of the bargaining game. While maximizing utility value it also tries to balance mutual equality and minimize envy. It is also important to understand the limita-

tion we have here; we just have looked at a limited, but a large set of examples. It might be that examples included here belong to a separate class of behavior.

Looking at the various charts it is apparent that the quality of the solution is determined by initial limits we have set on item values as seen on Figure 10. Having certain valuation it is clear that the outcome is defined by the initial differences in players valuations. Although in the case of uncertainty the determination is not that clear. But in general, as would be expected by intuition, adding uncertainty also degrades the possible solution. Uncertainties impact on the solutions for more complex problems (three and more players) remains to be investigated. What would be an efficient algorithm in that case?

Moreover the setup of our test environment is quite limited. Currently we have only tested the case of two participants, who had four items to share. Making the task more complex can reveal new insights on the algorithms. Next steps on the test settings should be:

1. Expanding the task for 3 and more participants
2. Expanding the item space and therefore the valuation lattice will be more complex
3. Extending the initial task beyond the territory division to allow more freedom, i.e. item set values between $[0;1]$

References

- [1] Hugo Steinhaus, The problem of fair division, *Econometrica*, 16, 101-104 (1948)
- [2] Steven J. Brams, Alan D. Taylor. *Fair Division: From cake-cutting to dispute resolution*, Melbourne (1996)
- [3] William L. Oberkampf, Jon C. Helton. *Evidence Theory for Engineering Applications Appeared in Engineering Design Reliability Handbook*, Danvers (2004)
- [4] Hugo Steinhaus. *Mathematical Snapshots*, New York (1969)
- [5] Hung T. Nguyen, Vladik Kreinovich, How to divide a Territory: A New Simple Formalism for Optimization of Set Functions, *International Journal of Intelligent Systems*, 223-251 (1999)
- [6] Duncan Luce, Howard Raiffa. *Games and Decisions: Introduction and critical survey*, Dover (1989)
- [7] Ken Binmore. *Natural Justice*, Oxford (2005)

B Publication 2

Veski, A. (2014). Price of Invisibility: Statistics of centralised and decentralised matching markets. In MacKerrow, E., Terano, T., Squazzon, F., and Sichman, J. S., editors, *Proceedings of the 5th. World Congress on Social Simulation*, pages 18–29, Sao Paulo

Price of Invisibility: Statistics of centralised and decentralised matching markets

Andre Veski*

Tallinn University of Technology, Tallinn, Estonia
andre.veski@ttu.ee

Abstract. We simulate a model of a decentralized and a centralised two-sided matching market in order to compare the efficiency of the two mechanisms. Also we parametrise the preference structure with preference list correlation and length. We use well known better response dynamic matching for a decentralised marketplace. We compare the decentralised mechanism to a centralised clearing house based on the Deferred-Acceptance mechanism. We use median rank and the number of unassigned agents to measure the efficiency of a matching. We found that median rank is statistically at most five times worse on a decentralised market, which occurs when correlation in preferences is small and list are long. Also with longer preference lists the decentralised mechanism has many unassigned agents and the matching is unstable, whereas a centralised mechanism computes a stable matching and usually with almost no unassigned agents.

Introduction

Two-sided matching markets have been studied quite extensively in the past half a century, starting with the National Residency Matching Program in US and the seminal result by David Gale and Lloyd Shapley [7] on a stable matching mechanism. This mechanism has proved useful in many entry-level job-markets (see e.g. [13]) and school choice markets (e.g. [1,12]). The general model is a two-sided market, where both sides have preferences or priorities over the agents on the other side.

The main benefit of a matching mechanism is that it is centrally applied. All the market participants report their preferences to a central clearing house that then can compute an optimal matching using for example Gale-Shapley Deferred-Acceptance algorithm [7]. Optimal usually means that the result is the best possible stable matching for one side of the market as the optimality can not be guaranteed for both sides. Stability is defined as a situation where participants do not have an incentive to deviate - there is no participant on the other side that the agent would prefer to his current match and that would also prefer him.

* This research was supported by European Social Fund's Doctoral Studies and Internationalisation Programme DoRa, and grant ETF8997

Roth [13] observed that in some situations where there was not a central clearing house in place, market participants still executed a very similar algorithm as proposed by Gale and Shapley. A major drawback of the execution was that usually it was time-capped, i.e. at some point the market had to be closed. This meant that the algorithm execution might not have finished and resulting matching may not be stable.

We model a situation where agents randomly interact and always select a better, compared to their current, match if available. This is also called a better response dynamic in [2]. Agents do not know all participants from the other side of the market, but rather become aware and form preference of them as they meet. Then they decide to change their match or not. In a real market setting there are always some cost related to a change, but in our current model we do not consider it.

Our main aim and contribution is to understand the benefit of having a centralised mechanism instead of a decentralised. Our secondary goal is to study the effect of preference structure on the outcome for both market mechanisms.

1 Market models

In a two-sided matching market there is a set $\mathcal{A} = \{a_1, \dots, a_n\}$ of agents on one side and a set $\mathcal{B} = \{b_1, \dots, b_n\}$ of agents on the other side. Although the number of agents on both sides can be different, we consider here only cases where the two sets are of equal size, i.e. the market is balanced. Each agent a_i from \mathcal{A} has a strict preference relation \succ_{a_i} over agents in \mathcal{B} , and similarly for $b_j \in \mathcal{B}$ there is a preference relation \succ_{b_j} over agents in \mathcal{A} . A matching μ is a mapping from $\mathcal{A} \cup \mathcal{B}$ to itself, such that for every $a \in \mathcal{A}$, is matched to $\mu(a) \in \mathcal{B} \cup \{a\}$, and similarly for $b \in \mathcal{B}$, $\mu(b) \in \mathcal{A} \cup \{b\}$. When an agent is matched to itself, $\mu(a) = a$ or $\mu(b) = b$ respectively indicates that they are in fact unmatched. Also for every $a, b \in \mathcal{A} \cup \mathcal{B}$, $\mu(a) = b$ implies $\mu(b) = a$. A matching is unstable if there is at least two agents, a blocking pair, a and b from opposite sides of the market such that $b \succ_a \mu(a)$ and $a \succ_b \mu(b)$. A matching is stable if it is not unstable.

1.1 Decentralised random matching model

For modelling a decentralised market we use NetLogo ([15]). As usual in NetLogo agents are positioned on a grid, we use the default size of 33×33 , with each position occupied by at most one agent. After the preference profiles have been set, at each time step an agent selects a random free position on the grid in the distance of 10 positions, if the position is occupied the agent remains in his current position for this round. The grid has a periodic boundary condition, meaning when an agent selects a position out of the grid, he is just moved to opposite side of the grid as the grid is toroidal.

After all agents have found a new position, we find the closest neighbour to all agents, from opposite side of the market, on the neighbouring 8 positions. If

there is more than one at the same distance, one is chosen randomly. Note that some agents might not have a neighbour at each step.

After all agents have been assigned a neighbour they start a transaction with the selected neighbour. First they both check if they are on each other's preference lists. If they are on each other's lists, the agents are myopic, they compare their current assignment with the neighbouring agent and decide to change their match if the neighbour is higher in their preference list (Algorithm 1). This is described as random better response dynamics in [2]. In one timestep t some agents might have multiple transactions if they were selected as closest neighbour by multiple agents, but only the best match will remain.

Algorithm 1 Better response dynamic

Require: $a, b, \succ_a, \succ_b, \mu$
Ensure: μ is a matching
 $m_a \leftarrow \mu(a), m_b \leftarrow \mu(b)$
if $b \succ_a m_a$ **and** $a \succ_b m_b$ **then**
 $\mu(m_a) \leftarrow m_a, \mu(m_b) \leftarrow m_b$
 $\mu(a) \leftarrow b, \mu(b) \leftarrow a$
end if
return μ

1.2 Centralized matching model

We compare our results from the decentralised matching model to one that would be centrally managed. In this case both parties submit their preferences to a central clearing-house that then outputs a matching. A classical algorithm used is the Deferred-Acceptance algorithm (DA), discovered by Gale and Shapley in [7].

The DA algorithm always produces a stable matching as showed by [7]. Furthermore depending the side that initiates the proposing sequence also obtains an optimal stable matching where agents from that side are matched to the best possible partner. The A-proposing algorithm is presented in Algorithm 2, similarly we can construct B-proposing algorithm that outputs stable matching optimal for B. These matchings may be, but necessarily are not the same, as there can be multiple stable matchings. It is also known that in general that A-side optimal stable matching is the worst stable matching for agents in \mathcal{B} and vice-versa, see for example [13].

2 Selected parameters

The preference structure is parametrised by the length of the preference list (k) and correlation between the preference lists (c). The preference list limit k is the

Algorithm 2 A-Proposing deferred-acceptance

Require: $\mathcal{A}, \mathcal{B}, \succ_a, \succ_b$

Ensure: μ is a matching

```
while There are unmatched agents in  $\mathcal{A}$  with proposals do
  for all  $a \in \mathcal{A}$  and  $\mu(a) = a$  and  $\succ_a$  is not empty do
     $b \leftarrow \text{FIRST}(\succ_a, \mathcal{B})$  // Most preferred match for  $a$  in  $\mathcal{B}$ 
     $\mu(a) \leftarrow b, \mu(b) \leftarrow \{\mu(b), a\}$ 
  end for
  for all  $b \in \mathcal{B}$  do
     $a \leftarrow \text{FIRST}(\succ_b, \mu(b))$ 
     $\mathcal{A}' \leftarrow \mu(b) \setminus \{a\}, \mu(b) \leftarrow a$ 
    for all  $a \in \mathcal{A}'$  do
       $\mu(a) \leftarrow a, \succ_a \leftarrow \succ_a$  without  $b$ 
    end for
  end for
end while
return  $\mu$ 
```

same for all agents. Although the preferences themselves are not necessarily the same, they might be correlated to some degree or might be totally random as defined by parameter c .

In reality the limit on the preference list might be artificial, due to limited processing capacity of the clearing mechanism. Or it might also be driven by agents themselves, due to the cost of additional information processing. It is hard to evaluate all the alternatives in a market, so they settle on listing or evaluating just a few. Or when considering the labour market, the list might be limited because some agents lack the skill to be matched with some jobs. Similar limitations to length of preference have been studied in [16,11].

High correlations between preference lists are usually driven by similar information people receive over alternatives and also similar value systems. It is observed in [14] that high correlations limit the size of the core of stable matchings. In some aspects correlation has been investigated in [4], where they look at fully correlated preference (uniform) lists and the effect on convergence to stability.

There have been additional studies on the effect of correlation. Mostly correlation is defined as a utility function of the agents in the form $u_{a_i}(b_j) = \beta \cdot \xi(b_j) + \xi_{a_i}(b_j)$ ([3,6,5]) and then sorted to obtain preference ordering. The parameter β is the correlation parameter and in case of $\beta = 0$ we recover the uncorrelated preferences. The $\xi(n_j)$ is global popularity of the agent b_j and $\xi_{a_i}(b_j)$ is the agent a_i specific utility for agent b_j . Note that β can be arbitrarily large, thus it is hard to have fully correlated preference lists. In ([5]) they define a similarity measure for preference lists after generation, but usually the results are far from total correlation.

Algorithm 3 Random permutation

Require: $n, k \in [0, 1], c \in [0, 1]$

Ensure: p is a permutation of unique numbers

$p \leftarrow 1, 2, 3, \dots, n, j \leftarrow n, l \leftarrow k \cdot n$

while $j > 0$ **do**

$r \leftarrow$ random number between 0 and 1

$q \leftarrow \lfloor j \cdot r^{1-c} \rfloor + [c \neq 1]$

$t \leftarrow p_q, p_q \leftarrow p_j, p_j \leftarrow t$

$j \leftarrow j - 1$

end while

return $\{p_1, p_2, \dots, p_l\}$

Table 1: Parameter spaces

Parameter	Values
Knowledge limit (k)	0.02, 0.1, 0.2, ..., 1
Correlation (c)	0, 0.1, 0.2, 0.3, 0.4..., 1

2.1 Generating the preferences

The preferences are generated using algorithm 3 with parameters k , c and n . This algorithm is taken from [9] and modified to generate correlated lists with parameter c . The algorithm starts with master list of n numbers. Then it iterates over the list from end to beginning, each time at position j randomly selecting a position $q \in [1, j]$ to exchange values with. The correlation parameter c just states how biased the randomly selected position is, higher values indicated that the exchange position is selected closer to the current position j . With $c = 0$ the selection is uniformly probable over all positions, until finally at $c = 1$ the exchange position is always the active position and all the generated lists are exactly the same.

2.2 Parameter space

We run our simulation for multiple combinations of parameters k and c with the pattern as in Table 1, altogether $11 \cdot 11 = 121$ combinations. Each set of parameters is executed 100 times, to account for some variability. Since we are interested in effects of k and c we fix the market size to $n = 100$ for both sides of the market. So altogether there are 200 participants in the market. There is some research on unbalanced markets [3], where there are more participants on one side of the market. Also there is most likely significant interaction with our selected parameters and the size and the balancedness of the market, but this currently remains future work.

3 Results

We define two measures for the outcome of a matching mechanism: median rank and the number of unassigned agents. And Price of Invisibility measure for comparing the two matching mechanisms.

Definition 1 *The rank of b_j in preference relation \succ_{a_i} of an agent a_i is defined as $\rho_{a_i}(b_j) = |\{b' : b' \succ_{a_i} b_j\}|$. Similarly we define the rank of a_i in b_j preference list by $\rho_{b_j}(a_i)$.*

Definition 2 *Given a matching μ then side A 's median rank r_a is defined as*

$$P(\rho_a(\mu(a)) \leq r_a) \geq \frac{1}{2}, P(\rho_a(\mu(a)) \geq r_a) \leq \frac{1}{2}, \forall a \in \{a' : \mu(a') \neq a', a' \in \mathcal{A}\} . \quad (1)$$

Similarly we define a median rank r_b for agents on side B .

Definition 3 *Given a matching μ the number of unassigned agents is defined as*

$$u = |\{a : a \in \mathcal{A}, \mu(a) = a\}| + |\{b : b \in \mathcal{B}, \mu(b) = b\}| . \quad (2)$$

Definition 4 *Given a better response dynamic matching μ_{brp} and a deferred-acceptance matching μ_{da} , the Price of Invisibility, with respect to metric $f(\cdot)$, is defined as*

$$PoI_f = \frac{f(\mu_{brd})}{f(\mu_{da})} . \quad (3)$$

We currently cut-off the execution of the decentralised mechanism after 5000 steps. Ideally we would run the simulation until we reach a stable state, but the random model does not find a stable solution in polynomial time [2]. Also when we observe the details we see that already after 2500 steps the rate of change is really slow - Figure 1.

Fig. 1: Convergence speed





Fig. 2: Probability of a stable solution



3.1 Stability

As mentioned the decentralised market takes exponential time to find a stable match so we cut-off the execution at $t = 5000$ steps. Our selected random better response decentralised matching model operates by satisfying blocking pairs with each transaction. As these are not guaranteed to be the best blocking pairs that are satisfied some blocking pairs containing one of the agents might still remain.

The results of probability of stability with different parameters is presented on Figure 2. Similar results are reported in [4], where they look at $k \cdot n \leq 8$.

3.2 Median Rank


We select the median rank of a matching as a descriptive statistic for a matching. Main reason is that the distribution of matched ranks is exponential, most receive their first some their second and then the number of agents decays by rank, and median is much better statistic for an exponential distribution than the mean rank. Secondly median has a much better interpretation to it as half of the agents received a better and half a worse rank than the median, but there is hard to find an agent who received the average rank. Another alternative would be the the rate parameter of the exponential distribution, but the parameter describes more the skewness of the distribution than the outcome.

On Figure 3¹ we have plotted the median rank as a function of correlation and length of preference lists, respectively c and k . Interesting observations are that in a decentralised market both sides, A and B , have very similar median ranks, which can be expected as neither side has a definite advantage over the other. Also median rank from decentralised mechanism is really close to the median rank for the accepting side of a centralised DA matching.

These observations are also confirmed by other papers - [8] and [10] show that when the preference lists are short, even on one side, the set of stable matchings is likely to be small, and the difference in ranks is also small, which we observe

¹ Surfaces are smoothed with local regression

Fig. 3: Median rank dependence on k and c in centralised and decentralised markets



when $k \leq 0.4$. In [14] it is also observed empirically that the size of the core is small when preference lists are short.

Next we estimate the relationship to get the expected median rank in a matching as a function of c and k . Looking at various multiplicative functional forms, we arrive at the following form that has a reasonable trade-off between accuracy and complexity. In a later section we also have an overview of residuals. Fitting the parameters for decentralised matching we obtain (4).

$$r^d \approx e^{3.9-2 \cdot k+1.9 \cdot c \cdot k} \cdot k \quad . \quad (4)$$

For the centralised (DA) matching expected median rank is different and depends on the proposing side. We denote the proposing side with a and the side b is the accepting/rejecting side as before. Similarly as before fitting the parameters we obtain (5).

$$r_a^c \approx e^{4.1-3.9 \cdot k+3.8 \cdot c \cdot k} \cdot k, r_b^c \approx e^{3.9-2.1 \cdot k+2.1 \cdot c \cdot k} \cdot k \quad . \quad (5)$$

We can now approximately estimate the proportional difference in median ranks in decentralised and centralised markets. Price of Invisibility on median rank is defined as equation (6) where r^d and r^c are the median ranks from decentralised and centralised models respectively. With this we can estimate how much worse is expected median rank from a decentralised market compared to a centralised one.


$$PoI_r = \frac{r^d}{r^c} \quad . \quad (6)$$

For the proposing side a we obtain (7).

$$PoI_{r_a} \approx e^{-0.2+1.9 \cdot k-1.9 \cdot c \cdot k} \quad . \quad (7)$$

First we observe that when $k = 1$ and $c = 1$ we actually obtain ratio < 1 , which indicates that the decentralised matching might be better. Initially this

Fig. 4: Unassigned probability dependence on k and c in centralised and decentralised markets



seems counter-intuitive, but as we will observe in the next section there is a hidden cost, the number of agents who are unmatched in this case tends to be very large compared to a centralised match.

Also we observe that the greatest PoI_{r_a} occurs when k is big and c is small. For example when $k = 1$ and $c = 0$ then $PoI_{r_a} \approx e^{1.7} \approx 5$.

For the receiving side we obtain (8).

$$PoI_{r_b} \approx e^{-0.1 \cdot k - 0.2 \cdot k \cdot c} \leq 1 . \quad (8)$$

The median rank does not differ much for side B in decentralised and centralised markets, the PoI_{r_b} is very close to 1. But again we shall observe that the difference in the number of unassigned agents is significant. The PoI_{r_a} of averages over 100 executions is on Figure 5a.

3.3 Unassigned agents

A critical metric of a matching market is the number of unassigned agents. A centralised matching scheme guarantees that we have a minimal number of unassigned agent. In a decentralised market this is not always the case, as agents make choices using the better response dynamics, we are not usually guaranteed to have a minimal number of unassigned agents because of the dynamics.

We start by fitting the relationship of unassigned agents as a function of k and c in a centralised market. We observe on Figure 4 that the general form of the relationship is sigmoidal for each k , with minimum unassigned value (lower asymptote A) as $c = 0$ and maximum value (upper asymptote K) at $c = 1$. Thus we fit the parameters (A, K, B, M) of a generalised logistic equation (9). The asymptote parameters depend on k and logistic equation depends on c .

$$u \approx A + \frac{K - A}{1 + e^{-B(c-M)}} . \quad (9)$$

We also obtain different fitting parameter values for stable and unstable matchings on decentralised mechanism. For stable and centralised market we obtain the following lower and upper asymptotes as in (10).

$$A^c = e^{5.5-10 \cdot k}, K^c = 200(1-k) . \quad (10)$$

As the length of the preference lists k grows the lower asymptote ($c = 0$) decays exponentially and becomes effectively zero after $k > 0.6$. In the case when agents have fully correlated preferences ($c = 1$) the decay in unassigned agents is merely linear as does not reach zero until the preference lists contain all the agents $k = 1$.

After the unassigned values are normalised using lower and upper asymptotes to be between 0 and 1 we fit a logistic function. We obtain that $B^c = 4 + 25 \cdot k$ and $M^c = \frac{1+27 \cdot k}{5+24 \cdot k}$. M is the position of maximum growth and B is the rate of growth. When k is small say $k = 0.1$ we observe that $M = 0.5$, so when $c \leq M$ the number of unassigned agents is closer to the lower asymptote and when $c \geq M$ the number of unassigned agents is closer to upper asymptote. If the growth rate B is large, it means that most of the growth happens near M . As k increases the growth rate B also becomes more rapid as correlation c passes the critical point M .

Next we look at the number of unassigned agents in a decentralised matching market. The number of unassigned is rather different due to the dynamics of the assignment process. We observe that the results are quite different when we obtain a stable match and when not, so we fit two separate logistic models. When we obtain a stable state we observe that the number of unassigned agents very closely resembles that of a centralised matching. On the other hand in the case of unstable matchings, approximately where $k > 0.3$, in decentralised matching we obtain the asymptotes in (11).

$$A^d = 30, K^d = 160(1-k) + 30 . \quad (11)$$

We observe that we can expect to have at least about 30 (15%) of unassigned agents as the lower asymptote indicates - even when we have full preference lists $k = 1$. Interestingly the lower asymptote is constant, which is an evidence that the number of unassigned agents is rather caused by the market mechanism than the inherent preference structure.

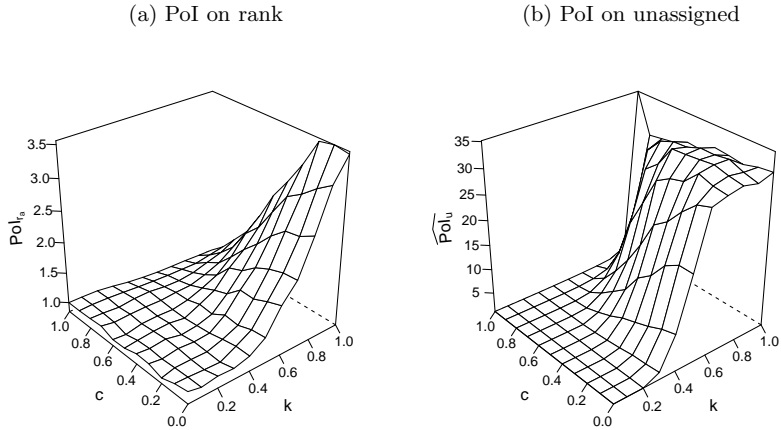
Similarly for unstable states in decentralised matching we obtain the position of maximum growth as $M^d = \frac{-4+55 \cdot k}{54 \cdot k}$ and growth rate $B^d = 54 \cdot k$.

Finally we estimate the average Price of Invisibility on number of unassigned agents as $PoI_u = \frac{u^d}{u^c}$. First we obtain lower and upper bounds based on previously obtained asymptotes where $c = 0$ and $c = 1$ respectively in (12).

$$PoI_u \approx \begin{cases} 1 & \text{if stable} \\ \in \left(\frac{30}{e^{5.5-10 \cdot k}}, \frac{160(1-k)+30}{200(1-k)} \right) & \text{if unstable} \end{cases} . \quad (12)$$

The PoI_u of averages over 100 executions is on Figure 5b. Since the number of unassigned agents in a centralised market can be zero in many circumstances the figure is plotting $\widehat{PoI}_u = \frac{\widehat{u^d}+1}{\widehat{u^c}+1}$.

Fig. 5: Price of Invisibility



3.4 Error analysis

The fitted functional forms of the relationships between matching metrics should be considered as approximations of the actual relationships. The error in the predictions tends to vary with the parameters. We define two error metrics:

Definition 5 *Root mean square error is defined as*

$$RMSE = \sqrt{\frac{\sum_1^n (\hat{y} - y)^2}{n}}. \quad (13)$$

Definition 6 *Normalised root-mean-square error is defined as*

$$NRMSE = \frac{RMSE}{y_{max} - y_{min}}. \quad (14)$$

In Table 2 we observe that the overall errors are small for r_a and u .

Table 2: Error levels

Mechanism	Function	RMSE	Normalised RMSE
Centralised	r_a	1.9	4.0 %
	r_b	3.1	6.3 %
	u	4.7	2.3 %
Decentralised	r_a	2.6	5.6 %
	r_b	2.6	5.8 %
	u	6.3	3.3 %

4 Conclusion and further research

Main conclusions

- the critical value to have a stable match with better response dynamic is to list less than 20 alternatives
- in a decentralised market there is always a significant number of unassigned agents - about 15%
- the median rank in the decentralised matching process produces very similar results as a centralised matching for the receiving side in DA

This research could be further extended by adding

- additional parameters for the size and balancedness of the market
- spatial properties like population density on a grid

References

1. Abdulkadiroglu, A., Sönmez, T.: School Choice: A Mechanism Design Approach. *American Economic Review* 93(3), 729–747 (2003)
2. Ackermann, H., Goldberg, P.W., Mirrokni, V.S., Röglin, H., Vöcking, B.: Uncoordinated two-sided matching markets. In: *Proceedings of the 9th ACM conference on Electronic commerce*. pp. 256–263. ACM Press, New York, USA (2008)
3. Ashlagi, I., Kanoria, Y., Leshno, J.D.: Unbalanced Random Matching Markets : The Stark Effect of Competition (2013)
4. Biró, P., Norman, G.: Analysis of stochastic matching markets. *International Journal of Game Theory* 42(4), 1021–1040 (2012)
5. Boudreau, J.W., Knoblauch, V.: Marriage Matching and Intercorrelation of Preferences. *Journal of Public Economic Theory* 12(3), 587–602 (2010)
6. Caldarelli, G., Capocci, A.: Beauty and distance in the stable marriage problem. *Physica A: Statistical Mechanics and its Applications* 300(1-2), 325–331 (2001)
7. Gale, D., Shapley, L.S.: College Admissions and the Stability of Marriage. *The American Mathematical Monthly* 69(1), 9–15 (1962)
8. Immorlica, N., Mahdian, M.: Marriage, Honesty, and Stability. In: *Proceedings of the 16th annual ACM-SIAM symposium on Discrete algorithms*. pp. 53–62 (2005)
9. Knuth, D.E.: *Seminumerical Algorithms*. Addison-Wesley, Reading, MA, 3rd edn. (1997)
10. Kojima, F., Pathak, P.A.: Incentives and Stability in Large Two-Sided Matching Markets. *American Economic Review* 99(3), 608–627 (2009)
11. Laureti, P., Zhang, Y.C.: Matching games with partial information. *Physica A: Statistical Mechanics and its Applications* 324(1-2), 49–65 (2003)
12. Pathak, P., Sönmez, T.: School Admissions Reform in Chicago and England: Comparing Mechanisms by Their Vulnerability to Manipulation. Tech. Rep. 16783, National Bureau of Economic Research, Cambridge, MA (2011)
13. Roth, A.E.: Deferred acceptance algorithms: history, theory, practice, and open questions. *International Journal of Game Theory* 36(3-4), 537–569 (2008)
14. Roth, A.E., Peranson, E.: The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design. *American Economic Review* 89(4), 748–780 (1999)
15. Wilensky, U.: NetLogo (1999), <http://ccl.northwestern.edu/netlogo/>
16. Zhang, Y.C.: Happier world with more information. *Physica A: Statistical Mechanics and its Applications* 299(1-2), 104–120 (2001)

C Publication 3

Veski, A. and Pöder, K. (2017). Zero-intelligence agents looking for a job.
Journal of Economic Interaction and Coordination

Zero-intelligence agents looking for a job

André Veski¹  · Kaire Pöder²

Received: 25 May 2016 / Accepted: 1 June 2017
© Springer-Verlag GmbH Germany 2017

Abstract We study a simple agent-based model of a decentralized matching market game in which agents (workers or job seekers) make proposals to other agents (firms) in order to be matched to a position within the firm. The aggregate result of agents interactions can be summarised in the form of a Beveridge curve, which determines the relationship between unmatched agents, unemployed job seekers and vacancies in firms. We open the black box of matching technology, by modelling how agents behave (make proposals) according to their information perception. We observe more efficient results—in the form of a downward shift of the Beveridge curve in the case of simple zero-intelligent agents. Our comparative statics indicate that market conditions, such as the heterogeneity of agents' preferences, will also shift the Beveridge curve downwards. Moreover, market thickness affects movement along the Beveridge curve. Movement right-down along the curve if there is an increasing number of agents compared to positions within firms. Furthermore, we show that frictions in re-matching, such as commitment to a match, could be another factor shifting the Beveridge curve toward the origin.

Keywords Matching market · Computational experiment · Decentralised matching · Job search · Beveridge curve

✉ André Veski
andre.veski@ttu.ee
Kaire Pöder
kaire.poder@ebs.ee

¹ Tallinn University of Technology, Tallinn, Estonia

² Estonian Business School, Tallinn, Estonia

1 Introduction

Market economies in general experience large employment fluctuations and average unemployment rates differ between countries. The underlying job search and matching theory (the Diamond–Mortensen–Pissarides canonical framework) provides a conceptual explanation for some aspects of the relationship between vacancies and unemployment known as the Beveridge curve. The core of job-search and matching models is built on the assumption that the external rate of job creation and destruction, but also worker reallocation, determine the steady-state equilibrium of number of unmatched workers and jobs (unassigned agents in our model) (Mortensen and Pissarides 1999). Because of search and recruiting costs, hiring and firing costs and other forms of matching-specific costs, decentralised markets create inefficiencies. This matching technology is implicitly characterised by its matching function, which summarises the trading technology between agents, their actions and choices that eventually bring them together into productive matches (Petrongolo and Pissarides 2001). In the relevant “matching function literature”, it is stressed that such a theoretical tool is useful because it allows to reduce the complexity of information imperfection, heterogeneous agents and congestion into a tool-kit similar to the production function or money demand function. However, the interactions or matching technology are still rather treated as a black box.

We open this black box by providing a simple agent-based model of a decentralized market game in which agents (workers or job seekers) make proposals to the agents on the other side of the market (firms) in order to be matched to available positions. In our computational experiments, a market game is identified by three components: the preference structure of agents, market conditions, i.e. the relative number of positions and workers, and the behaviour of agents (workers or firms) based on the information they have about their own preferences and options in the market.

Search models in labour economics rely on three pillars: the decision of workers, the decision of firms and wage setting mechanism. We concentrate on the first two pillars. Thus, our model belongs to the literature on agent-based partial labour market models (e.g. Neugart and Richiardi 2012) which use microsimulation to explain stylised facts about the labour market. These agent based “micro-to-macro” models give insights into labour markets in the form of partial or general models. In the latter labour markets are embedded into larger economic models. We are mainly interested in literature aiming to reproduce the Beveridge curve with search and coordination in a partial agent-based model. Thus, models (Richiardi 2004, 2006; Riccetti et al. 2015) aimed at explaining the job search on the basis of wages or more general market models that interact with the embedded labour market (Dawid et al. 2014; Deissenberg et al. 2008), that lie outside of our scope.

The partial agent-based models have been developed for replicating stylised facts from real labour markets, such as the negative-sloped Beveridge curve in the unemployment-vacancy (u, v)-space. Fagiolo et al. (2004), followed later by Silva et al. (2012), reproduce a Beveridge curve in a partial agent-based model and come up with a standard explanation that frictions and the institutional setting affect the position of the curve.

Moreover, giving up the assumption of rational expectations about the behaviour of agents has produced fruitful insights. [Tassier and Menczer \(2008\)](#) investigate job-hunting via social networks. They find that random social networks spread vacancy information better and thus achieve lower unemployment. [Neugart \(2004\)](#) uses an urn-ball matching model on a small scale (30–50 agents) and endogenous matching function workers send applications randomly, however the Beveridge curve is closer to origin than one would expect in large markets. Similarly [Richiardi \(2004\)](#) employs a similar model with wages and produces a Beveridge curve further from the origin. While [Richiardi \(2006\)](#) models labour supply in a general setting by a non-equilibrium, adaptive agent-based model of heterogeneous workers and firms, with on-the-job searching, endogenous entrepreneurial decisions and endogenous wage and income determination. The latter is able to reproduce a number of stylized facts generally accepted in labour economics and industrial organization, including the negative-sloped Beveridge curve. Also, in this set-up, the matching process is based on random applications from job seekers for vacancies in the single labour market. Furthermore, this model allows for on-the-job searching, meaning that assigned agents can get job offers as well.

There are some other search-and-match models (e.g. [Gabriele 2002](#); [Deissenberg et al. 2008](#); [Boudreau 2010](#)). Some of them also produce a Beveridge curve, but usually investigate different aspect of matching. Often their aim is to study some other aspects, like stratification and use different underlying assumptions (e.g. centralised matching or perfect information).

Our modelled agents can be considered myopic—they make proposals and only accept better proposals without additional strategic thinking. In our base model, agents apply for a random position and a matching occurs only when both sides find their new partner preferable to their current match. This is similar to the Zero-Intelligence (ZI) model from financial markets ([Gode and Sunder 1993](#); [Chen and Tai 2010](#)). Our behavioural models are an extended version of the better response dynamic proposed by [Knuth \(1997b\)](#) and further analysed by [Roth and Vate \(1990\)](#) and [Ackermann et al. \(2008\)](#).

Nevertheless, the agents will not be able to find an equilibrium (stable) matching in a micro sense. It is rather a steady-state equilibrium in a macro sense, as the size (number of matched agent-pairs) of the matching converges. In general, we show how the assumptions about the information available to agents, the structure of their preferences and market thickness determine the shape and placement of the Beveridge curve. Thus, our approach not only differs from the framework of Diamond-Mortensen-Pissarides, but also from recent discussions on job search and matching efficiency ([Veracierto 2011](#); [Shimer 2013](#); [Sahin et al. 2014](#)) where the persistent empirically observed adverse shifts (outward shifts of the Beveridge curve) have been explained mostly by the deterioration in the efficiency of the matching technology. We use a partial agent-based model as classified by [Neugart and Richiardi \(2012\)](#) to develop aggregate regularities from the micro-behaviour of individual agents in order to illustrate the position of the Beveridge curve in the (u, v) -space. Our model is able to reproduce some well-known stylised facts from labour market literature like the negative sloped Beveridge curve. Moreover, we explain the shift of the Beveridge curve not only as related to information available to players, but also including a level of heterogeneity

of the agents' preferences. We model the latter on the basis of the correlation between the preference lists of agents to indicate their similarity (or common understanding of a good job position). In addition, we allow for different lengths of the preference lists. We show that short lists and high correlation shift the curve to the upper-right, further from the origin. In addition, we include a parameter for market thickness—the balance between market sides—indicating the ratio of positions to workers. Market thickness models the effect of interactions with other markets through job creation and destruction rates, i.e. the out-of-equilibrium state of a general market model. We see that this determines the position on the Beveridge curve.

The comparative statics, unexpectedly, show that the noise behaviour produces Beveridge curves closest to the origin, compared to the more informed behaviours. This appears to be the effect of lower number of acceptable re-matches from random proposals. With more informed behaviours agents make proposals to acceptable matches, thus they have a higher re-matching rate. As each re-match creates one and potentially breaks two existing matches, this leads the Beveridge curve away from the origin. To elaborate this, we conduct additional experiments with frictions and find that greater obstacles to re-matching produce Beveridge curves closer to the origin.

In addition, we view our decentralised market game as an abstract model which in addition to studying agent-based interactions in job search and matching markets can be used in alternative settings, e.g. decentralised school or university choice. In all these cases, an institution or a central authority—a clearing house or similar—is missing and agents on both sides constantly have to react to new proposals and responses.

We continue as follows. First, we introduce the set-up of our model, concentrating on matching behaviour which includes search and commitment costs as well as the assumptions behind preference formation. Second, we analyse the convergence of the matching procedures and check if the resulting matching distributions correspond with the expectation. Third, we show our results by indicating how a Beveridge curve depends on the search behaviour and the structure of preferences. Additionally we show how introducing frictions could improve the size of the matching, albeit at cost, in terms of less preferred matched, for the agents. Finally, we conclude by discussing our contribution to the discourse in search and matching literature.

2 Models

2.1 Behaviours

In order to translate the neoclassical matching function into an agent-based version, we employ the framework offered by [Guerrero and Axtell \(2011\)](#). This relies on three orders of assumptions, which include rationality, agent homogeneity and “non-interactiviness” in an agent-based model. The categorisation of the features is presented in [Table 1](#).

Beginning with the first order assumptions about the nature of the interaction mechanism, we can see that the literature about the matching function often treats this procedure as a black box. In an agent-based model, the interaction is a central ques-

Table 1 Matching model assumptions

	Matching function	Our agent-based model
<i>First order assumptions</i>		
Interaction mechanism	Implicit (black box)	Explicit (decentralised or centralised matching)
<i>Second order assumptions</i>		
Rationality	No explicit reference. Functional form captures information imperfections	Explicit, exact notion of how proposals are made and accepted
Equilibrium	Equilibrium is inefficient due to negative externalities (e.g. congestion). Imbalance indicates skill mismatch and locational differences	Equilibrium is inefficient because of assumptions about behaviour and information
Agent types	Representative (homogeneous) agent(s), heterogeneous sectors	Heterogeneous (different preferences)
<i>Third order assumptions</i>		
Technological advances through time	Technological advances in matching shift “Beveridge curve” (search is less costly)	No technological advances. Preferences are static, do not adapt
Supply shocks or business cycles	Affect job creation and destruction rates, but are generally treated as an empirical question	No explicit job creation and destruction, but the market thickness, having more jobs than workers or vice versa, is of central interest
Contracts	A rate of matching and unmatching	A relationship can be broken whenever a better match is found (<i>and the contract has expired</i>)
Transaction costs	A limitation on matching rate	Once a better match is found for both parties, there is no additional cost for changing the match

tion for the investigation. In financial double-auction markets, Zero-Intelligence (ZI) interaction models have been fruitful in investigating aggregate market phenomena (e.g. [Gode and Sunder 1993](#); [Farmer et al. 2005](#); [Ladley 2012](#)). We employ a similar approach for modelling search in the labour market.

ZI model ([Ladley 2012](#); [Chen and Tai 2010](#)) is useful because it allows us to decouple the behaviour of an agent from the market structure. Moreover, we are interested in whether similarly to non-strategic agents (e.g. [Farmer et al. 2005](#); [Gode and Sunder 1993](#)), interesting market phenomena can be produced in the current context. To our knowledge, these types of models have not been studied for job search. There are macro-level studies that concentrate on modelling unemployment and vacancies (e.g. [Petrongolo and Pissarides 2001](#); [Mortensen and Pissarides 1999](#)). There also exist agent-based models of wage equilibrium (e.g. [Guerrero and Axtell 2011](#)) and job search on social networks [Guerrero and Axtell \(2013\)](#); [Zhou et al. \(2014a\)](#); [Hoefler and Wagner \(2012\)](#), but there is no simple model for job search.

Table 2 Explored behaviour models

Rationality (2nd order)	Interaction, proposer (1st order)	
	Random side	A-proposing
Noise proposal	Noise proposal	Noise proposal A
Better proposal	Better proposal	Better proposal A
Blocking proposal	Blocking proposal	Blocking proposal A

The labour market consists of two sets of agents workers and firms (or positions within a firm). The main behavioural aspect is how the match is initiated, i.e. the worker-position pair selection. We study models where the proposing power is either only on one side of the market (A-proposing models) or where it is proportionally shared (random agent proposing). In other words, either workers always make proposals, proposals are made interchangeably, or firms always make proposals, depending on who is considered side A. We call the non-proposing agent the responding agent. In the centralised Deferred-Acceptance markets (Gale and Shapley 1962), the matching is always optimal and stable for the proposing side, while it is the worst possible level for the responding side (e.g. Knuth 1976; Roth and Sotomayor 1990). However, in many practical (Roth and Peranson 1999) and large markets (Immorlica and Mahdian 2015), the difference seems to be small and the effect of market thickness is much greater (Ashlagi et al. 2013b). The matched rank structure may also affect the size of the matching as proposing probabilities are different under random agent proposing.

In fact, we investigate several models (Table 2), where a second-order assumption of ZI is characteristic of the base model. The ZI model is called the Noise proposal model wherein two random agents, one from each side, are selected, but a matching transaction occurs only when the new match is an improvement over their current matches. Similarly, in financial markets a deal is only accepted when the offered price is above the reserve price for both sides, i.e. the buyer and the seller, otherwise the price is offered at random. Thus, in our mechanism, there are only pairwise interactions and a transaction occurs when the reserve offer is met on both sides.

For comparative purposes, we include behaviours, where agents know more than in the ZI model, but less than in the better response dynamic. In our Better proposal model, we assume that agents, e.g. workers, know of a better match or a position that would also be suitable for them and thus do not make proposals in a wholly random manner. This can be considered similar to the Zero-Intelligence Plus model (Chen and Tai 2010). Agents only make proposals to a better match, i.e. a position higher on their preference list than their current (reserve) match. Thus, the proposing side does not even consider non-acceptable matches. In contrast, in the Noise proposal model an offer is made to a random agent on the preference list, and might actually not be acceptable to the proposer. This is learned in the transaction. In other words, in the Better proposal model, even agents with a current high ranking match have a high probability of a new match, which only depends on the responding side finding it acceptable. On the other hand, in the Noise proposal model a transaction probability would also be lower for proposing agents, if their current match is high on their preference list.

In order to further extend the information pool available to the agent, we use the Blocking proposal model, where agents make proposals to their random blocking partner, which is equivalent to the better response dynamic proposed by Knuth (1997b). A blocking pair is formed by two agents from the opposite sides of the market who, if they met, would prefer to be matched to each other. Here agents only make a proposal to other agents when they know it would be accepted by the other party, i.e. the blocking pair, in the current state of the market. The match can still be broken in the future if either of the agents finds a more preferred partner.

Although we investigate multiple behaviour models of the proposing agents, we always assume that the responder is a ZI agent. He only and always accepts proposals made by agents higher than the current match on his preference list. In addition, the existence of an information aggregating institution is implicitly assumed in Better and Blocking proposal behaviour. For example finding a potential blocking pair in the Blocking proposal model can be thought of as being supported by an institution.

We are only interested in studying the aggregate results of the search behaviour, therefore we simplify most of the third-order assumptions. The preferences of the agents are fixed, so they do not adapt to market conditions during the search. There is also no creation or destruction rate of new agents or positions, nor any external shocks that might trigger such destructions or creations. We do, however, explicitly model market thickness. The market is considered thick when there is exactly the same number of agents on both sides, so all the agents can potentially be matched. If there are more agents on one side, the market will not be thick and there will always be some agents unassigned. In addition, thickness has an impact on the search outcome, as agents from the smaller side have more options to choose from. Thickness is also an indicator of disequilibrium, as the number of jobs is not equal to the number of workers, characterising exogenous dynamics of job creation, destruction, etc.

Finally, in our main experiments there are no limitations on matching with a more preferred partner, i.e. no commitments to contracts or any transaction costs for changing a match. However, in Sect. 3.5 we show the effect of frictions of enforcing different types of obstacles, including contractual ones, on re-matching.

More formally, we employ a model similar to that used in modelling centralised two-sided matching markets (e.g. Roth 2008). There is a set $\mathcal{A} = \{a_1, \dots, a_{n_A}\}$ of agents on one side and a set $\mathcal{B} = \{b_1, \dots, b_{n_B}\}$ of agents on the other side. The number of agents on both sides can differ ($n_A \neq n_B$) depending on market thickness. Each agent a_i from \mathcal{A} has a strict preference relation \succ_{a_i} over agents in \mathcal{B} , and similarly for $b_j \in \mathcal{B}$ there is a preference relation \succ_{b_j} over agents in \mathcal{A} . A matching μ is a mapping from $\mathcal{A} \cup \mathcal{B}$ to itself, so that every $a_i \in \mathcal{A}$, is matched to $\mu(a_i) \in \mathcal{B} \cup \{a_i\}$, and similarly for $b_j \in \mathcal{B}$, $\mu(b_j) \in \mathcal{A} \cup \{b_j\}$. When an agent is matched to itself, $\mu(a_i) = a_i$ or $\mu(b_j) = b_j$ respectively indicates that they are in fact unmatched. Being matched to oneself is the least preferred option for all the agents. Agents would still find only the acceptable positions in their preference list, which might not contain all positions. Similarly for the position, only some agents might be acceptable. In addition, for every $a_i, b_j \in \mathcal{A} \cup \mathcal{B}$, $\mu(a_i) = b_j$ implies $\mu(b_j) = a_i$.

A matching is unstable if there are at least two agents a_i and b_j from opposite sides of the market so that $b_j \succ_{a_i} \mu(a_i)$ and $a_i \succ_{b_j} \mu(b_j)$ —a blocking pair. A matching is stable, if it is not unstable. In Table 3 we have listed the notation for quick reference.

Table 3 Notation

Symbol	Description
\mathcal{A}	Preferences of agents on side A
\mathcal{B}	Preferences of agents on side B
a_i	Preference profile for agent i , $a_i \in \mathcal{A}$
b_j	Preference profile for agent j , $b_j \in \mathcal{B}$
n_A	Number of agents on side A
n_B	Number of agents on side B
θ	Market thickness $\theta = \frac{n_B}{n_A}$
k	Length of preference lists
c	Correlation of preferences
τ	Re-matching friction length
μ	Matching
s	Size of matching, counted in pairs of agents
u	Unassigned percentage on side A , $u = 1 - \frac{s}{n_A}$
v	Unassigned percentage on side B , $v = 1 - \frac{s}{n_B}$
$r(\mu(i))$	Matched rank of agent i in matching μ
\tilde{r}_a	Median matched rank of agents in \mathcal{A}
\tilde{r}_b	Median matched rank of agents in \mathcal{B}
ρ_i	Number of blocking pairs for agent i
$\bar{\rho}_i$	Number of blocking pairs with unmatched agents for agent i
$\tilde{\rho}_i$	Number of blocking pairs with matched agents for agent i

Algorithm 1 General Proposal Dynamic**Require:** $\mathcal{A}, \mathcal{B}, \mu$ **Ensure:** μ is a matching $p \leftarrow \text{SelectProposer}(\mathcal{A}, \mathcal{B})$ $m_p \leftarrow \mu(p)$ $r \leftarrow \text{SelectResponder}(p)$ $m_r \leftarrow \mu(r)$ **if** $p \succ_r m_r$ **and** $r \succ_p m_p$ **then** $\mu(m_p) \leftarrow m_p, \mu(m_r) \leftarrow m_r$ $\mu(r) \leftarrow p, \mu(p) \leftarrow r$ **end if****return** μ

With the notation in Table 3 we can present a General Proposal Dynamic in Algorithm 1. The SelectProposer() and SelectResponder() procedures are distinct for each of the described models in Table 2. The SelectProposer() selects a random agent from set $\mathcal{A} \cup \mathcal{B}$ or \mathcal{A} depending on whether the behaviour model is Random side or A-proposing. The SelectResponder() selects an agent from the preferences of the proposer, and the actual selection depends on whether the behavioural model is the Random, Better or Blocking proposal.

We study the macro-level convergence properties of the search behaviour. On an individual agent level, the market needs not to be in equilibrium. There have been studies on the equilibrium of decentralised matching processes. [Niederle and Yariv \(2009\)](#) study such applications where firms and workers have aligned preferences and show the conditions for having a stable matching in equilibrium. [Haeringer and Wooders \(2011\)](#) examine equilibrium behaviour with slightly different models, where agents cannot be re-matched to a previously rejected partner, but their model is otherwise similar to ours as agents have to respond immediately. [Diamantoudi et al. \(2015\)](#) look at stability when agents make a commitment to a partner, which can either be only for a certain period, or an infinite commitment so that participants exit the market, or only a one-sided commitment. They show that having a requirement for firms to commit to an employee can result in unstable matchings in equilibrium. In our models we mostly concentrate at the no-commitment scenario, except in the frictions scenario. [Pais \(2008\)](#) analyses the equilibrium with limited information about preferences. [Eriksson and Häggström \(2007\)](#) also study decentralised matching, but they do not make any underlying assumptions about how the matching is reached. Instead, they measure the degree of instability in some random matchings. However, if some decentralised matching model is assumed, the resulting matching would not be a uniform selection of all the possible matchings.

In addition there have been experimental studies ([Echenique and Wilson 2009](#); [Echenique and Yariv 2013](#)) about decentralised matching markets with human subjects which show that stability tends to be a prevalent outcome, but is not always guaranteed. The interesting aspect in those cases is human behaviour, which usually also restricts the size of the experiments, which tend to be small—10–20 participants. [Zhou et al. \(2014b\)](#) use real-world data from small and large on-line matching markets and study the statistical regularities of those matchings, mainly how the size of the markets relates to the size of the matching. This is also what we are interested in. Unfortunately they do not count the size of the two sides of the markets, but only the overall size.

2.2 Preferences

Nevertheless, we look at heterogeneous agents with various degrees of correlation in their preferences and the availability of matches. We model a situation where the preferences of the agents are all idiosyncratic (effectively random) and agents find all partners acceptable. Yet, we also look at some structural constraints. Firstly, we introduce limited preference lists indicating that an agent finds only a fraction of the partners acceptable. Second, preference lists are somewhat correlated, or in extreme cases, preferences are exactly the same, indicating common tastes.

In the real world, correlated preferences show “popular tastes”, e.g. all agents have similar preferences for high paying jobs or are interested in simple assignments, etc. The length of the preference list, however, indicates the probability of an agent being found unacceptable, even if a certain agent would be the only candidate. So a person without a pilot’s licence would never be employed as a pilot. Thus, shorter preference lists imply that not all agents are acceptable to a particular position.

Algorithm 2 Correlated permutation

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

We assume that agents have strict preferences for agents (workers or positions) from the opposite side of the market. In the simplest case preferences are random, i.e. each agent has a totally idiosyncratic preference ordering. In general we can think of more structured preferences in a society, parametrised by the length of the preference list (k) and the correlation between the preference lists (c). In our experiments, the preference list limit k is set to be the same for all agents. Correlated preferences are from a global preference ordering. The degree of correlation is also the same for all agents, but the preference ordering is not necessarily the same when comparing two agents.

We generate the preferences using Algorithm 2 with parameters k, c and n . This algorithm is a modified version of a random permutation algorithm from Knuth (1997a) to generate correlated preferences with parameter c . The algorithm starts with a master list of n numbers (agents). Then it iterates the list from beginning to end, each time at position j randomly selecting a position $q \in [j + 1, n]$ to exchange values with. The correlation parameter c states how biased the randomly selected position is, higher values indicate that the exchange position is selected closer to the current position j . With $c = 0.0$ the selection is uniformly probable over all positions, until finally at $c = 1$ the exchange position is always the active position and all the generated lists are exactly the same. There is one global ordering of agents for each side of the market that is used for generating correlated preferences.

The power of uniform distribution \mathcal{U}^{1-c} used to randomly select the exchange positions while generating the correlated preference list is proportional to the *Beta* distribution with parameters $Beta(\frac{1}{1-c}, 1) \sim \mathcal{U}^{1-c}$. In Fig. 1, we see the probabilities of having a particular value at some position in a list of 10 values between 0 and 9. Each box is a position in a list and displays the probability of having a certain value in that position. We see that when $c = 1.0$ then all positions have a 100% probability of having the same value and when $c = 0.0$ then all values in any position are uniformly probable.

In Fig. 2, we compute for comparison the mean Spearman ρ and Kendall τ correlation coefficients over all the preference lists. We compute two types of means over the correlation coefficients, first compared to the initial global ordering and then a mean over pairwise correlations among a random sample of preference lists. We see that the pairwise means are always below when compared to the correlation with the global ordering. This is because although all the preferences are a similar distance




Fig. 1 Preference probabilities with degrees of correlation




Fig. 2 Spearman ρ and Kendall τ of generated preferences

from the global list, the generated lists might still be far from each other, i.e. have a lower correlation. That is the case with small degrees of correlation. Still, when the correlation $c = 1.0$ all the lists are exactly the same and the ρ and τ values are also 1.0.

In reality, the limit of the preference list might be due to skill mismatch in position requirements and for agents based on utility. Similar limitations on the length of preferences have been studied in Zhang (2001) and Laureti and Zhang (2003). We consider the preferences to be “known” to the agents only in terms of the behaviour model employed. So, for example, in the Better proposal model, agents would select a random proposal that is an improvement over their current match, but it might not be their most preferred match.

We do not study societies where, in general, positions and workers might have aligned preferences as in [Niederle and Yariv \(2009\)](#). High correlations between preference lists are usually driven by people receiving similar information about alternatives and also due to similar value systems. It is observed by [Roth and Peranson \(1999\)](#) that high correlations limit the size of the core of stable matchings. Certain aspects of correlation have been investigated by [Biró and Norman \(2012\)](#) that looks at fully correlated preference lists by varying the length of preference lists and its effect on convergence to stability.

There have been additional studies on the effect of correlation. Generally, correlation is defined as the agent's utility function in the form $u_{a_i}(b_j) = \beta \cdot \xi(b_j) + \xi_{a_i}(b_j)$ ([Ashlagi et al. 2013a](#); [Caldarelli and Capocci 2001](#); [Boudreau and Knoblauch 2010](#)) and then sorted to obtain a preference ordering. The parameter β is the correlation parameter and in case of $\beta = 0$ we would recover the uncorrelated preferences. The $\xi(n_j)$ is the global popularity of the agent b_j and $\xi_{a_i}(b_j)$ is the specific utility of agent a_i for agent b_j . It should be noted that β can be arbitrarily large, thus it is hard to have fully correlated preference lists. ([Boudreau and Knoblauch \(2010\)](#)) define a similarity measure for preference lists after generation, but usually the results are still far from fully correlated ($c = 1.0$) preferences.

3 Results

3.1 Computational experiments and convergence

All our experiments are carried out with $n_A = |\mathcal{A}| = 1000$ agents. We vary the market thickness $\theta = \frac{n_B}{n_A} \in [0.5, 2.0]$, which varies in the number of agents on side \mathcal{B} $n_B = |\mathcal{B}| \in [500, 2000]$. We do 300,000 matches. Figure 3 shows that with all the behaviour models and various values of market thickness, the matching size converges to a steady-state equilibrium. This does not mean that there are no changes in the matching. Individual agents still change their matches whenever their behavioural mechanism conditions are met. In the experiments, the result was never a stable matching without




Fig. 3 Convergence over time

blocking pairs. Therefore, small fluctuations always occur in the matching, but Fig. 3 demonstrates that this does not have a significant effect on the macro-level of matching.

To begin with, in Sects. 3.2 and 3.3, only the effect of market thickness is explored, with full ($k = 1.0$) and uncorrelated ($c = 0.0$) preference lists. In Sect. 3.4 we study more structured ($k \in [0.02, 1.0]$ and $c \in [0.0, 1.0]$) preferences and their impact on the Beveridge curve. Further in Sect. 3.5 we study the effect of some frictions on the matching.

3.2 Analysis of convergence conditions

Our probabilistic analysis considers the simpler situation where the market is thick ($\theta = 1.0$), all agents have full ($k = 1.0$) and uncorrelated preference lists ($c = 0.0$). With limited preference lists, the analysis would not hold and would need to be augmented with probabilities of having certain agents on a preference list. Similarly, with correlation, we would need to assume some probability of having certain agents higher on the preference lists. Calculations are much more simplified, when we can assume this to be of uniform probability for all the relevant agents.

There are four types of events that can occur in all of the decentralised matching behaviour models:

- e_1 Two previously unmatched agents are matched. The size of the matched population increases by one on both sides and nobody becomes unmatched. The net change in the size of the matching will be one.
- e_2, e_3 One unmatched agent (either from \mathcal{A} or \mathcal{B}) is matched to another matched agent. The matched population increases by one, but one previously matched agent now becomes unmatched due to the divorce of the already matched agent. The net change in the matching size will be zero.
- e_4 Two already matched agents are matched to each other and consequently two divorces occur. The net change in the size of the matching will be minus one.

We are interested in understanding the convergence of the size of the matching. Since for events e_2 and e_3 the net change in the matching is zero, we are not interested in those events. The size of the matching changes only with the events e_1 and e_4 and has converged when $\Delta s = P(e_1) - P(e_4) \rightarrow 0$. We analyse the probabilities of the events e_1 and e_4 for all of the six decentralised behaviour models. This is similar to the model in (Mortensen and Pissarides 1999, p. 1185). However, Mortensen and Pissarides (1999) analyse the model on a macroscopic level with transition probabilities on a Markov chain. Yet, we analyse the model on an agent level, where the transition probabilities depend on the states of the agents.

Noise proposal and noise proposal A Whenever two unmatched agents meet, they always prefer to be matched rather than unmatched, given our assumptions about preferences. Hence, the probability for event e_1 is the probability for two unmatched agents to meet as in (1).

$$\widehat{P}(e_1) = \left(1 - \frac{s}{n_A}\right) \left(1 - \frac{s}{n_B}\right) \quad (1)$$

The probability of event e_4 is selecting two matched agents that prefer to be matched. This firstly depends on selecting two matched agents, one from \mathcal{A} and the other from \mathcal{B} . Secondly, the selected agents would both have to be higher on each other's preference lists than their current match. The latter is an average over all the matched agents. This is summarised in (2).

$$\widehat{P}(e_4) = \left(\frac{s}{n_A}\right) \left(\frac{s}{n_B}\right) \left(\frac{1}{s} \sum_{i \in s, i \in B} P(r(\mu(i)) > x) P(X = x)\right) \cdot \left(\frac{1}{s} \sum_{i \in s, i \in A} P(r(\mu(i)) > x) P(X = x)\right) \tag{2}$$

Since the probability of selecting an agent in a particular position is uniform, $P(X = x) = \frac{1}{n}$, and the number of agents n is either n_A or n_B , depending on which side we are looking at, we can simplify (2), which results in (3).

$$\widehat{P}(e_4) = \left(\frac{s}{n_A^2}\right) \left(\frac{s}{n_B^2}\right) \left(\frac{1}{s} \sum_{i \in s, i \in B} P(r(\mu(i)) > x)\right) \cdot \left(\frac{1}{s} \sum_{i \in s, i \in A} P(r(\mu(i)) > x)\right) \tag{3}$$

Better proposal A The probability of event e_1 is the same as with Noise behaviour. The probability that an unmatched agent a_i is selected from \mathcal{A} and then the probability that the agent will select an unmatched agent is $b_i \in \mathcal{B}$. Since agent a_i has a full preference list, the selection is made from the entire set \mathcal{B} . This results in the probability of two unmatched agents being selected, as expressed in Eq. (4).

$$\widehat{P}(e_1) = \left(1 - \frac{s}{n_A}\right) \left(1 - \frac{s}{n_B}\right) \tag{4}$$

To find the probability of event e_4 of Better proposal A, we first have to take the probability of selecting a matched agent from \mathcal{A} . Then the selected agent a_i will randomly select an agent from the set of better matches on its preference list. The matching is successful only if the selected agent from \mathcal{B} side finds a_i acceptable as well. This means, by definition, that the two agents have to form a blocking pair. With $\tilde{\rho}_i$ we count the number of blocking pairs with another matched agent from \mathcal{B} for agent a_i . This results in probability for event e_4 as in Eq. 5.

$$\widehat{P}(e_4) = \left(\frac{s}{n_A}\right) \left(\frac{1}{s} \sum_{i \in \mu, i \in A} \frac{\tilde{\rho}_i}{r(\mu(i))}\right) \tag{5}$$

Better proposal When an agent from either side can act as a proposer, we only need to weigh the proposers selection probabilities by the size of the respective agent-sets. For the probability of event e_1 , this would result in Eq. (6), which simplifies to the same result as (4).

$$\begin{aligned} \widehat{P}(e_1) &= \left(\frac{n_A}{n_A + n_B} + \frac{n_B}{n_A + n_B} \right) \left(1 - \frac{s}{n_A} \right) \left(1 - \frac{s}{n_B} \right) \\ &= \left(1 - \frac{s}{n_A} \right) \left(1 - \frac{s}{n_B} \right) \end{aligned} \tag{6}$$

The probability of event e_4 for Better proposal behaviour is again similar to the Better proposal A. We first take the probability of selecting a matched agent from either from \mathcal{A} or \mathcal{B} . If the proposing agent is selected, then selecting an accepting responder has the same probability as (5), but over all of the agents in $\mathcal{A} \cup \mathcal{B}$, which results in total probability as in Eq. (7).

$$\begin{aligned} \widehat{P}(e_4) &= \frac{n_A}{n_A + n_B} \left(\frac{s}{n_A} \right) \left(\frac{1}{s} \sum_{i \in \mu, i \in \mathcal{A}} \frac{\tilde{\rho}_i}{r(\mu(i))} \right) \\ &\quad + \frac{n_B}{n_A + n_B} \left(\frac{s}{n_B} \right) \left(\frac{1}{s} \sum_{i \in \mu, i \in \mathcal{B}} \frac{\tilde{\rho}_i}{r(\mu(i))} \right) \\ &= \frac{1}{n_A + n_B} \sum_{i \in \mu} \frac{\tilde{\rho}_i}{r(\mu(i))} \end{aligned} \tag{7}$$

Blocking proposal A The probability of event e_1 depends on selecting an unmatched agent a_i from \mathcal{A} and then agent a_i selecting an unmatched agent from among its blocking pairs. On average, this results in probability as in Eq. (8). When all agents have full preference lists, we could simplify even further with $\bar{\rho}_i = n_B - s$ as all unmatched B agents would be blocking pairs for any unmatched A agent.

$$\widehat{P}(e_1) = \left(1 - \frac{s}{n_A} \right) \left(\frac{1}{n_A - s} \sum_{i \notin \mu, i \in \mathcal{A}} \frac{\bar{\rho}_i}{\rho_i} \right) = \frac{1}{n_A} \sum_{i \notin \mu, i \in \mathcal{A}} \frac{\bar{\rho}_i}{\rho_i} \tag{8}$$

Similarly the probability of event e_4 depends on selecting a matched agent $a_i \in \mathcal{A}$ and this agent a_i selecting a blocking pair with a matched agent from among all blocking pairs, including the ones with and unmatched agent. This results in probability as in Eq. (9).

$$\widehat{P}(e_4) = \left(\frac{s}{n_A} \right) \left(\frac{1}{s} \sum_{i \in \mu, i \in \mathcal{A}} \frac{\tilde{\rho}_i}{\rho_i} \right) = \frac{1}{n_A} \sum_{i \in \mu, i \in \mathcal{A}} \frac{\tilde{\rho}_i}{\rho_i} \tag{9}$$

Blocking proposal Similarly to the Better proposal behaviour, we need to weigh the probabilities in Eqs. (8) and (9) against the probabilities of selecting an agent either from \mathcal{A} or \mathcal{B} . This results in probabilities as in Eqs. (10) and (11) for e_1 and e_4 respectively.

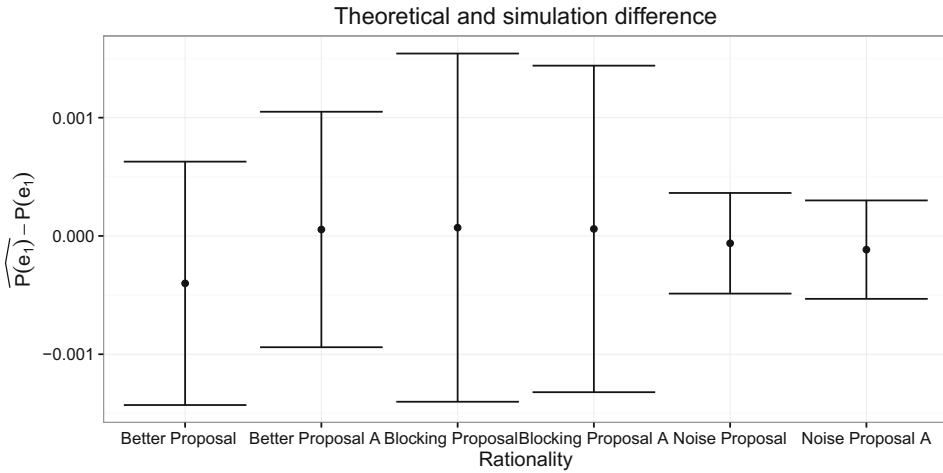


Fig. 4 Comparison of expected and experimental probabilities of event e_1

$$\begin{aligned}
 \widehat{P}(e_1) &= \frac{n_A}{n_A + n_B} \left(1 - \frac{s}{n_A}\right) \frac{1}{n_A - s} \sum_{i \notin \mu, i \in A} \frac{\bar{\rho}_i}{\rho_i} \\
 &\quad + \frac{n_B}{n_A + n_B} \left(1 - \frac{s}{n_B}\right) \frac{1}{n_B - s} \sum_{i \notin \mu, i \in B} \frac{\bar{\rho}_i}{\rho_i} \\
 &= \frac{1}{n_A + n_B} \sum_{i \notin \mu} \frac{\bar{\rho}_i}{\rho_i} \tag{10}
 \end{aligned}$$

$$\begin{aligned}
 \widehat{P}(e_4) &= \frac{n_A}{n_A + n_B} \frac{s}{n_A} \frac{1}{s} \sum_{i \in \mu, i \in A} \frac{\tilde{\rho}_i}{\rho_i} + \frac{n_B}{n_A + n_B} \frac{s}{n_B} \frac{1}{s} \sum_{i \in \mu, i \in B} \frac{\tilde{\rho}_i}{\rho_i} \\
 &= \frac{1}{n_A + n_B} \sum_{i \in \mu} \frac{\tilde{\rho}_i}{\rho_i} \tag{11}
 \end{aligned}$$

In Figs. 4 and 5, we compare the results of the probabilistic matching estimations from the structural properties of the matchings $P(\cdot)$ from the specified equations and actual observed probabilities $P(\cdot)$ from computational experiments. The figures show the average difference in these probabilities with 99% confidence bounds on normal distributions. We see that with all the behaviours, the statistical difference between the estimated and observed probabilities is close to zero and is always within the 99% bound (Figs. 4, 5). This indicates that the structural properties of the matchings are as expected.

In Fig. 6 we investigate the convergence of the matching process. The process converges when $\Delta s = P(e_1) - P(e_4) \rightarrow 0$. This figure demonstrates that the difference $P(e_1) - P(e_4) \approx 0$ is statistically within the 99% confidence bound. This is not to say that the matching freezes. There are still new matches made as well as broken. Rather the statistical properties of the matching, in terms of size, distribution of obtained rank, and the distributions of blocking pairs, converge and become stationary.




Fig. 5 Comparison of expected and experimental probabilities of event e_4




Fig. 6 Statistical difference in $P(e_1)$ and $P(e_4)$

3.3 Beveridge curve and the movement along the curve

We look at the Beveridge curve without any structure in preferences, that is $c = 0.0$ and $k = 1.0$. Figure 7 contains data points for all six of the behaviours.

The thickness (θ) of the market sides determines where on the Beveridge curve the steady-state of the matching is situated. In Fig. 7 the lines represent some examples of market thickness. When the market is thick, i.e. we have equal number of agents on both sides ($\theta = 1.00$), the result will lie on the 45 degree line. When the market is biased toward one or the other side, we move along the Beveridge curve to the upper left or lower right. Different values for thickness can be considered the effects of outside influences, e.g. economic state that influence the job destruction and creation rates. So the curve is the result of out-of-equilibrium state of the wider marker.




Fig. 7 Beveridge curves

Contrary to intuition, in Fig. 7, we see that the noise behaviour produces Beveridge curves closest to the origin, compared to the more informed behaviours. The largest, closest to the origin, matching outcome is obtained when agents behave randomly in the market, as in the two Noise proposal behaviours. Moreover, it is not relevant how the proposing power is distributed, shared by both sides or concentrated on one side, however, the results are the same on average. It has to be stressed that the size of the matching is much smaller when agents exhibit more informed behaviours. By proposing to more preferred agents (Better proposing behaviours) or to agents proposers know will accept (Blocking proposing behaviours) lead to more re-matches and breaking existing ties. This is the result of a matching transaction being much more likely in the latter two cases than with noisy proposals. Furthermore, each time a re-matching occurs potentially two existing matches are broken, which results in a lower overall size of the matching. In all behaviours, when two unmatched agents meet they would be matched thus increasing the overall matching, however, better informed agents make many more swaps of their partner and reach a steady-state with a smaller size.

It is also clear that in any of the mechanisms, which side proposes selection does not affect the size of the match. This might be an indication of the fact that being a proposer is not that relevant in large markets, as has been also discovered in large centralised matching markets (Roth and Peranson 1999; Immorlica and Mahdian 2015).

Depending on market thickness, the ratio of free agents is higher on the larger side on the market. In case of Better and Blocking response behaviours, the change in free agents depends linearly on market thickness. In the case of the zero-intelligence Noise proposing behaviour, the relation of the number of free agents with thickness is not



Fig. 8 Beveridge curve and correlation in preferences c



Fig. 9 Beveridge curve and length of preferences k

linear, but is akin to a square root function. So the size of the matching increases faster with Noise proposing behaviour when the market is becoming thicker ($\theta \rightarrow 1.00$) compared to other behaviours.

3.4 Shifts in the Beveridge curve

The Beveridge curve is concerned with the size of the matching. Most empirical curves show a relationship between unemployment and vacancies, and it is never the case that neither of them is zero. This is usually attributed to the structural properties of preferences—some workers are not suitable for some jobs. We observe a similar effect of having structure in preferences. In addition, we show that the shift can also be the effect of search behaviour.

It has been long assumed that shifts of the Beveridge curve are due to the structure of preferences in the labour market (Abraham and Katz 1986; Blanchard et al. 1989; Mortensen and Pissarides 1999; Sahin et al. 2014): namely, where workers can and would like to work, and similarly who the employers would like to hire. If agents are low on the preference lists for every position or not on the list at all, it is very hard for them to find a match.

We model the preferences of the agents in terms of correlation c in a society and the length k of the preference lists. Both these factors play a significant role in how good, large, the match is. In the experiments in Figs. 8 and 9, we vary c and k simultaneously to understand their interaction effects. If preference lists are very correlated (Fig. 8), the matching tends to be small, which is also the case when the lists are short (Fig. 9). However, it also appears that either of these features can determine the location of the Beveridge curve on its own. Conversely to the trend even when the lists are short, but with low correlation, large matches can result. This can also occur when the lists



Fig. 10 Beveridge curve and maximum potential matching

are long and correlation is high. Naturally, with long lists and low correlation, the matching is the largest. So the relationship in preference list parameters (c and k) and the size of the matching is not straightforward.

To simplify thinking about the structural properties of the preferences of agents, we use the maximum potential matching to determine the effect of correlations and limited lists on a matching. The maximum potential matching is computed using the Hopcroft-Karp algorithm (e.g. [Cormen et al. 2004](#), p. 696) in networkX library ([Hagberg et al. 2008](#)). This matching is then compared to the maximum matching with no correlations or limits on preferences and this comes down to the number of agents in the smaller side of the market. Thus, when $n_A = 1000$ and $n_B = 500$, the maximum matching can be $s_m = 500$. However, this can never be obtained because the preferences are somewhat structured. The maximum matching returns the size that can be obtained given the preference structure—the potential size p_m . Figure 10 looks at the effect of $\varphi = \frac{p_m}{s_m}$ on the Beveridge curve. We see that φ is close to 1 when we find a large matching with ZI and close to 0 when the matching with ZI is small. The latter may reflect certain skill mismatches in the market similarly to the established stylised facts in macroeconomic literature.

3.5 Effect of a re-matching friction

The main effect of the behaviours on the size of the matching originates from the differences in probabilities that a transaction ends with a successful re-matching. In the Noise proposal behaviour, a pair of agents is randomly selected, whereas in Better and Blocking proposal behaviours, only the agent on one side is randomly selected. In the latter two, the randomly selected agent only makes proposals that they already find acceptable. ZI-agents make proposals to random agents from the other side of the market and learn its ranking during the transaction. Thus, they might eventually reject the match.

Therefore, to disentangle the effect of the Noise proposal behaviour in a comparative context, we tweak our model slightly by adding a friction timer τ to an individual match. This would lower the probability of a transaction to result in a successful rematch. We investigate the friction effect with preferences, where $c = 0.0$ and $k = 100\%$. A new matching is only accepted, when the timer condition is satisfied for both agents forming a match. We implement three types of friction timers:




Fig. 11 Beverage curve of lazy agents




Fig. 12 Beverage curve of patient agents




Fig. 13 Beverage curve of greedy agents

1. After making a new match *lazy* agents wait for τ iterations before accepting a new offer;
2. After making a new match *patient* agents wait for τ proposals before accepting a new offer;
3. After making a new match *greedy* agents accept only matches that are τ positions higher than their current match.

In Figs. 11, 12 and 13 we show the results of lazy, patient and greedy agents respectively. In our experiments the friction is fixed for all agents. We see that by introducing frictions to agents before allowing them to be re-matched, the resulting match becomes larger and the Beveridge curve shifts closer to the origin. This is true for all our modelled behaviours. However, the effect is significantly greater for Better and Blocking proposal behaviours with lazy and patient agents (Figs. 11, 12). This is caused by the initial lower re-matching probability already present in frictionless Noise proposal behaviour. Moreover, the types of frictions cause some overlap. In frictionless Better and Blocking proposal behaviour, once a pair of agents was selected, the re-matching probability was higher compared to the Noise proposal, so the effect of the friction is also greater.




Fig. 14 Median rank for A-side lazy agents

For lazy and patient agents, the effects of the friction are similar. For a patient agent to re-match, they would have to be selected on τ occasions, whereas a lazy agent would have to wait for τ iterations. If the selection probabilities of an agent are the same in both cases, it should be straightforward to scale the results of lazy agents to the matching size of patient agents.

Furthermore, regarding patient (Fig. 12) and greedy (Fig. 13) agents with large $\tau \approx 100$, the resulting Beveridge curve is close to the origin for all behaviours. With a slightly unbalanced market, the number of free agents on the smaller side becomes effectively zero.

The greedy agent type of friction with smaller values for τ does not significantly improve the size of the matching (Fig. 13). Greedy agents accept a re-match only when it improves their position by at least τ ranks. They would still accept any match if they were unmatched. Similarly to behaviours without the frictions, the Better and Blocking proposal behaviours still result in more unmatched agents than the Noise proposal, as agents would tend to accept proposals more often. Also, the probability of selecting a match for a greedy agent that is a τ -improvement over their current match appears high for the selected τ as the size of the matching does not increase significantly (Fig. 13).

However, the effect of these frictions is not all positive. The re-matching friction increases the size of the matching, but decreases the rank of the matched agents. In Fig. 14 we show the results of the median rank (\tilde{r}_a) over matched agents from the lazy agent experiments, weighted by the length of the preference list (n_b). So we see the relative matched median rank as a percentage over the agents' entire preference list, where a lower value indicates that the median agent has a more preferred match. First, we observe that with the Noise Proposal behaviour and minimal friction the median rank is in about top 20% position. And surprisingly slightly higher with Blocking Proposal. However, when only A-side has proposing power this side achieves better median rank. When we increase the waiting time the median rank always worsens and the effect is more pronounced in the Blocking Proposal A behaviour.

The re-matching friction also has significant interaction effect with market thickness (θ). With noisy behaviour the agents on the smaller side have the power to get matched to more preferred agents, regardless of the friction. However, with Blocking Proposal behaviour the effect depends on who is on which side. For agents on the larger side, the friction has an adverse effect, i.e. the median agent has a less preferred match. Conversely, for agents on the smaller side ($\theta > 1.0$), the longer waiting time will result in more preferred matches.

As friction waiting times decrease the expected matched rank of an agent, it might not be rational for agents to participate in such a market. So market participants would advocate for lowering friction, for the reward of being better matched. However, they would also taking an additional risk of being left unmatched.

4 Conclusion and discussion

Recent contributions to the economists' understanding of the micro-foundations of the Beveridge curve have enriched the early work of [Blanchard et al. \(1989\)](#). However, substantial gaps remain in our understanding of both the impact of the matching technology as well as the process of including mechanisms which affect the Beveridge curve. Thus, we contribute to this research gap by studying the micro-foundations underlying the Beveridge curve.

We translated the framework of Diamond-Mortensen-Pissarides to an agent-based model, with the intention of explaining both the movements along the Beveridge curve and the shifts (location) of the curve itself. Our simple model shows a two-sided decentralized market game with three key determinants—preferences, information and market conditions. Thus, it may be argued that instead of explicitly modelling labour market institutions, we implicitly include features of institutions by modelling the various behaviours of agents. Our agents can have degrees of heterogeneous or completely homogeneous preferences. The structure of the preferences indicates a notion associated with the possible mismatch of the skills of workers across jobs. There might be high demand for the same jobs and same workers, which form the source of the mismatch. We have multiple approaches to model preferences. Firstly, agents can be heterogeneous with random preferences and full-length preference lists. Secondly, preferences can be correlated to some degree which is common to all agents. Thirdly, the length of the preference list of the agents can vary, which indicates that not all positions are acceptable or not all agents are suitable for certain positions. This allows us to model the limitations of structural unemployment.

The cornerstone of our analysis is our assumption about information. Information determines how the market game is played. Generally our agents are myopic—at each stage of the game, they make random decisions and accept better proposals without any alternative strategic thinking. Agents do not learn. However, we studied different behavioural models. In our initial Noise proposal (zero-intelligence) model, agents make random proposals. For comparative purposes, we constructed two alternative decision models—the Better proposal and the Blocking proposal model. In the Better proposal model, agents randomly make proposals only to a more preferred agent than their current match. In the Blocking proposal model, agents only make proposals to a random blocking pair, indicating that the proposal is always accepted.

Through the computational experiments, we found the aggregate number of vacancies and unmatched agents which constitute the Beveridge curve. We have three relevant agent related dimensions that explain the position of the curve and/or the current position along the curve—the correlation of preferences, the length of the preference lists, and the assumptions about the decision-making mechanisms of the agents. For comparative statics, we first showed that low correlation (heterogeneous agents)

will shift the Beveridge curve downward and long lists of preferences have a similar effect. We also observed that the assumptions about the decision-making behaviour affect the location of Beveridge curve considerably. Noise proposing models shift the Beveridge curve toward the origin compared to the Better or Blocking proposal models. This insight can be interpreted in light of the search and wait unemployment concept—zero-intelligence agents make random proposals that are not always accepted, while more advanced players make better proposals, thus resulting in a better match for the agent, but smaller matching overall.

In addition, we were interested in the effect of market thickness. This is the indicator for measuring the balance between market sides, i.e. equal number of jobs (agents) and worker agents indicates a thicker market. We demonstrated that thickness affects movement along the Beveridge curve. For instance, in the case of random preferences, we move right-down along the curve if there is an decreasing number of positions (job offerors) compared to agents (job seekers). This shows that the Beveridge curve is mostly the result of out-of-equilibrium dynamics in interrelated markets, affecting job creation and destruction rates. It appeared that regarding the Better and Blocking proposal mechanisms, changing market thickness simply means shifting the number of free agents or positions from one side of the market to the other. On the other hand, when agents make proposals randomly and market thickness becomes closer to one, the decrease in the rate of free agents is not linear, but a square root of free agents from the other side. Therefore, each additional position has a larger effect than one additional match, meaning that it creates opportunities for more agents to be matched. As the Better and Blocking proposals implicitly model search institutions, e.g. job hunters, it seems that these have a decreasing effect on employment.

The investigation of Noise behaviour revealed that the decreasing effect on unemployment and vacancies is related to limiting the probability of re-matches. Additional experiments showed that by enforcing some obstacle, friction, on the termination of the contract brings the Beveridge curve closer to the origin. These frictions might also have a basis from human psychology, as a sense of duty might limit an agent's willingness to terminate an existing contract. The longer the obstacle lasts, the closer to the origin the Beveridge curve locates. However, we also saw that frictions affect the matched rankings, i.e. with stronger friction the expected matched rankings decrease. Balancing the number of matched agents and their matched rank remains for further research.

Our approach had several simplifying assumptions: no transaction costs, no search and matching costs, no agency, homogeneous behaviour, and no dynamics (behaviour learning, new agents or change in preferences). Despite this, we open a path of research in agent-based modelling in order to contribute to the search and matching literature. Modelling matching technology by including some kind of a job board or alternative agency to the agent-based model remains a challenge for the future research.

Acknowledgements We thank Péter Biró for helpful comments during the early stages of this work and also the participants of the Santa Fe Institute Complex Systems Summer School 2015 for fruitful discussions. We are also grateful of constructive feedback from two anonymous referees, based on which the paper was significantly improved.

References

- Abraham K, Katz LF (1986) Cyclical unemployment: sectoral shifts or aggregate disturbances? *J Polit Econ* 94(3):507–522
- Ackermann H, Goldberg PW, Mirrokni VS, Röglin H, Vöcking B (2008) Uncoordinated two-sided matching markets. In: Proceedings of the 9th ACM conference on Electronic commerce. ACM Press, New York, pp 256–263
- Ashlagi I, Kanoria Y, Leshno JD (2013a) Unbalanced random matching markets. In: Proceedings of the fourteenth ACM conference on Electronic commerce. ACM Press, New York, pp 27–28
- Ashlagi I, Kanoria Y, Leshno JD (2013b) Unbalanced random matching markets: the stark effect of competition. <http://web.mit.edu/iashlagi/www/papers/UnbalancedMatchingAKL.pdf>. Accessed 14 Sept 2016
- Biró P, Norman G (2012) Analysis of stochastic matching markets. *Int J Game Theory* 42(4):1021–1040
- Blanchard OJ, Diamond P, Hall RE, Yellen J (1989) The Beveridge curve. *Brook Pap Econ Act* 1(1989):1–76
- Boudreau JW (2010) Stratification and growth in agent-based matching markets. *J Econ Behav Organ* 75(2):168–179
- Boudreau JW, Knoblauch V (2010) Marriage matching and intercorrelation of preferences. *J Public Econ Theory* 12(3):587–602
- Caldarelli G, Capocci A (2001) Beauty and distance in the stable marriage problem. *Phys A* 300(1–2):325–331
- Chen SH, Tai CC (2010) The agent-based double auction markets: 15 years on. In: Takadama K, Cioffi-Revilla C, Deffuant G (eds) *Simulating interacting agents and social phenomena*, chap 9. Springer, Tokyo, pp 119–136
- Cormen TH, Leiserson CE, Rivest RL, Stein C (2004) *Introduction to algorithms*, 2nd edn. MIT Press, Boston
- Dawid H, Gemkow S, Harting P, Hoog SVD, Neugart M (2014) An agent-based macroeconomic model for economic policy analysis: the Eurace@ unibi model. Working papers in economics and management 01-2014
- Deissenberg C, van der Hoog S, Dawid H (2008) EURACE: a massively parallel agent-based model of the European economy. *Appl Math Comput* 204(2):541–552
- Diamantoudi E, Miyagawa E, Xue L (2015) Decentralized matching: the role of commitment. *Games Econ Behav* 92:1–17
- Echenique F, Wilson AJ (2009) Clearinghouses for two-sided matching: an experimental study. Social science working paper 1315. California Institute of Technology
- Echenique F, Yariv L (2013) An experimental study of decentralized matching. <http://people.hss.caltech.edu/%7EYariv/papers/ExpDecentralizedMatching.pdf>. Accessed 25 May 2016
- Eriksson K, Häggström O (2007) Instability of matchings in decentralized markets with various preference structures. *Int J Game Theory* 36(3–4):409–420
- Fagiolo G, Dosi G, Gabriele R (2004) Matching, bargaining, and wage setting in an evolutionary model of labor market and output dynamics. *Adv Complex Syst* 07(02):157–186
- Farmer JD, Patelli P, Zovko II (2005) The predictive power of zero intelligence in financial markets. *Proc Natl Acad Sci* 102(6):2254–2259
- Gabriele R (2002) Labor market dynamics and institutions: an evolutionary approach. LEM working paper series 2002/07
- Gale D, Shapley LS (1962) College admissions and the stability of marriage. *Am Math Mon* 69(1):9–15
- Gode DK, Sunder S (1993) Allocative efficiency of markets with zero-intelligence traders: market as a partial substitute for individual rationality. *J Polit Econ* 101(1):119–137
- Guerrero OA, Axtell RL (2011) Using agentization for exploring firm and labor dynamics. In: *Emergent results of artificial economics*. Springer, Berlin, chap 12, pp 139–150
- Guerrero OA, Axtell RL (2013) Employment growth through labor flow networks. *PLoS ONE* 8(5):e60808
- Haeringer G, Wooders M (2011) Decentralized job matching. *Int J Game Theory* 40(1):1–28
- Hagberg AA, Schult DA, Swart PJ (2008) Exploring network structure, dynamics, and function using networkX. In: Varoquaux G, Vaught T, Millman J (eds) *Proceedings of the 7th python in science conference*, Pasadena, pp 11–15
- Hoefler M, Wagner L (2012) Locally stable matching with general preferences. <http://arxiv.org/abs/1207.1265>. Accessed 22 Aug 2016

- Immorlica N, Mahdian M (2015) Incentives in large random two-sided markets. *ACM Trans Econ Comput* 3(3):1–25
- Knuth DE (1976) *Mariages stables*. Les Presses de l'Université de Montréal, Montréal
- Knuth DE (1997a) *Seminumerical algorithms*, 3rd edn. Addison-Wesley, Reading
- Knuth DE (1997b) *Stable marriage and its relation to other combinatorial problems*. American Mathematical Society, Providence
- Ladley D (2012) Zero intelligence in economics and finance. *Knowl Eng Rev* 27(02):273–286
- Laureti P, Zhang YC (2003) Matching games with partial information. *Phys A* 324(1–2):49–65
- Mortensen DT, Pissarides CA (1999) Job reallocation, employment fluctuations and unemployment. In: Taylor JB, Woodford M (eds) *Handbook of macroeconomics*, vol 1. North Holland, Amsterdam, chap 18, pp 1171–1228
- Neugart M (2004) Endogeneous matching functions: and agent-based computational approach. *Adv Complex Syst* 07(02):187–201
- Neugart M, Richiardi M (2012) Agent-based models of the labor market. Laboratorio Riccardo Revelli working paper no. 125
- Niederle M, Yariv L (2009) Decentralized matching with aligned preferences. NBER working paper series 14840
- Pais J (2008) Incentives in decentralized random matching markets. *Games Econ Behav* 64(2):632–649
- Petrongolo B, Pissarides CA (2001) Looking into the black box: a survey of the matching function. *J Econ Lit* 39(2):390–431
- Riccetti L, Russo A, Gallegati M (2015) An agent based decentralized matching macroeconomic model. *J Econ Interact Coord* 10(2):305–332
- Richiardi M (2004) A search model of unemployment and firm dynamics. *Adv Complex Syst* 07(02):203–221
- Richiardi M (2006) Toward a non-equilibrium unemployment theory. *Comput Econ* 27(1):135–160
- Roth AE (2008) Deferred acceptance algorithms: history, theory, practice, and open questions. *Int J Game Theory* 36(3–4):537–569
- Roth AE, Peranson E (1999) The redesign of the matching market for American physicians: some engineering aspects of economic design. *Am Econ Rev* 89(4):748–780
- Roth AE, Sotomayor MAO (1990) *Two-sided matching: a study in game-theoretic modeling and analysis*. Cambridge University Press, Cambridge
- Roth AE, Vate JHV (1990) Random paths to stability in two-sided matching. *Econometrica* 58(6):1475–1480
- Sahin A, Song J, Topa G, Violante GL (2014) Mismatch unemployment. *Am Econ Rev* 104(11):3529–3564
- Shimer R (2013) Job search, labour force participation, and wage rigidities. In: Acemoglu D, Arellano M, Dekel E (eds) *Advances in economics and econometrics: theory and applications: tenth world congress*, chap 5. Cambridge University Press, New York, pp 197–234
- Silva ST, Valente JMS, Teixeira AAC (2012) An evolutionary model of industry dynamics and firms' institutional behavior with job search, bargaining and matching. *J Econ Interact Coord* 7(1):23–61
- Tassier T, Menczer F (2008) Social network structure, segregation, and equality in a labor market with referral hiring. *J Econ Behav Organ* 66(3–4):514–528
- Veracierto M (2011) Worker flows and matching efficiency. *Econ Perspect* 35(4):147–169
- Zhang YC (2001) Happier world with more information. *Phys A* 299(1–2):104–120
- Zhou B, He Z, Jiang LL, Wang NX, Wang BH (2014a) Bidirectional selection between two classes in complex social networks. *Sci Rep* 4:7577
- Zhou B, Qin S, Han XP, He Z, Xie JR, Wang BH (2014b) A model of two-way selection system for human behavior. *PLoS ONE* 9(1):e81424

D Publication 4

Veski, A. and Pöder, K. (2016). Strategies in Tallinn school choice mechanism. *Research in Economics and Business: Central and Eastern Europe*, 8(1):5–24

Strategies in the Tallinn School Choice Mechanism

Andre Veski

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

Kaire Põder

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

Abstract

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

JEL codes: D02, D04, D47, D82

Keywords: market design, school choice, preference revelation, stability

1. Introduction

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

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

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

As with the shortcomings of the Boston mechanism, which has created a rule of thumb for submitting the preferences strategically (Ergin and Sönmez, 2006; Pathak and Sönmez,

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

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

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

2. Background: Tallinn School Market

Over the years, some schools in Tallinn have become over-subscribed. These selective schools have inter-district admissions to primary school and have all introduced aptitude entrance tests (hereinafter exam schools). For intra-district comprehensive schools (hereinafter regular schools), the tradition has been a central or semi-central catchment-based allocation based on an application (single preference or multiple preferences) from the

family. Rejected offers were not treated centrally – each school and student should find the match independently.

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

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

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

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

2.1. Tallinn Mechanism

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

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

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

Table 1: Number of Reported Preferences under the Tallinn Mechanism

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

Source: Authors' calculation

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


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

10

2.2. Example of Deciding What to Report

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

Figure 1: Extensive Form Reporting Game



Source: Authors' illustration

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

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

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

Table 2: Utilities

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

Source: Authors' calculation

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

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

- for reporting one school $E_1[u_1(k)] = \frac{1}{2}(0.358 + 0) = .179$
- for reporting the first two schools $E_2[u_1(k)] = \frac{1}{3}(0.358 + 0.333) = .230$
- for reporting the first three schools $E_3[u_1(k)] = \frac{1}{4}(0.358 + 0.333 + 0.309) = .250$

We see that reporting all three preferences maximised utility. With utility function $u_2(k)$ the expected utilities are:

- for reporting one school $E_1[u_2(k)] = \frac{1}{2}(0.658 + 0) = .329$
- for reporting the first two schools $E_2[u_2(k)] = \frac{1}{3}(0.658 + 0.333) = .330$
- for reporting the first three schools $E_3[u_2(k)] = \frac{1}{4}(0.658 + 0.333 + 0.009) = .250$

As Figure 1 illustrates the game, where under the expected utility maximisation assumptions, parents obtain higher utility by reporting only one or two schools with $u_2(k)$.

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

2.3. Deferred-Acceptance Mechanism

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

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

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

In general the Deferred-Acceptance works as follows:

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

...

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

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

3. Model

3.1. Environment

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

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

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

Table 3: Fixed Societal Parameters

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

Table 4: Variable Societal Parameters

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

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

3.2. Preferences

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

We generate the preferences using the Algorithm 1 with parameters k , c and m . This algorithm is a modified version of a random permutation algorithm (Knuth, 1997, p. 145) to generate correlated preferences with parameter c . The algorithm starts with a master list of n numbers (agents). Then it iterates the list from beginning to end, each time at position j randomly selecting a position $q \in [j + 1, n]$ to exchange values with. The correlation parameter c illustrates how biased the randomly selected position is; higher values indicate that the

exchange position is selected closer to the current position j . With $c = 0.0$ the selection is uniformly probable over all positions, until finally at $c = 1$ the exchange position is always the active position and all the generated lists are exactly the same. There is one global ordering of agents for each side of the market that is used for generating correlated preferences.

3.3. Utility Function

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

Algorithm 1 Correlated permutation

Require: $m, k = 15, c \in [0, 1]$

Ensure: p is a permutation of unique numbers

$p \leftarrow 1, 2, 3, \dots, n, j \leftarrow 0$

while $j < k$ **do**

$r \leftarrow [0.0, 1.0]$ uniform random number between 0 and 1

$q \leftarrow [m - (m - j) \cdot r^{1-c}]$

$t \leftarrow p_q, p_q \leftarrow p_r, p_r \leftarrow t$

$j \leftarrow j + 1$


end while

return $\{p_1, p_2, \dots, p_k\}$

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

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

Figure 2: Exponential Utility Function



Source: Authors' illustration

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

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

3.4. Genetic Algorithm Optimisation

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

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

Here we do not assume that the manner of genetic algorithms is in reality how humans learn. We only employ it for computational tractability, as exploring the entire strategy-space

for 3,000 agents is resource consuming. However, there are studies that use a form of genetic algorithm as a model for learning (see e.g. Unver, 2001; Roth, 2002; Unver, 2005) and is also observed as exhibiting features with human subjects (e.g. Arifovic, 1994, 1996; Duffy, 2006).

Algorithm 2 Simple Genetic Algorithm – single iteration

Require: A set of agents, u agents utilities

Ensure: A is a set of agents

```

 $n \leftarrow |A|$ 
 $s \leftarrow \sum_{\alpha \in A} U_{\alpha}$ 
 $p \leftarrow \{ \frac{U_{\alpha}}{s}, \forall \alpha \in A \}$  {selection probabilities}
 $i \leftarrow 0$ 
for all  $r_1, r_2 \in \text{Select}(A, p, n)$  do
    {select with probability  $p$ , with replacement  $n$  pairs of strategies}
     $\alpha_i \text{ Crossover}(r_1, r_2)$  {assign new strategy to agent  $\alpha_i$ }
    if  $\text{RandomNumber}() < 0.05$  then
         $\text{Mutate}(\alpha_i)$ 
    end if
     $i \leftarrow i + 1$ 
end for
return  $A$ 
    
```

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

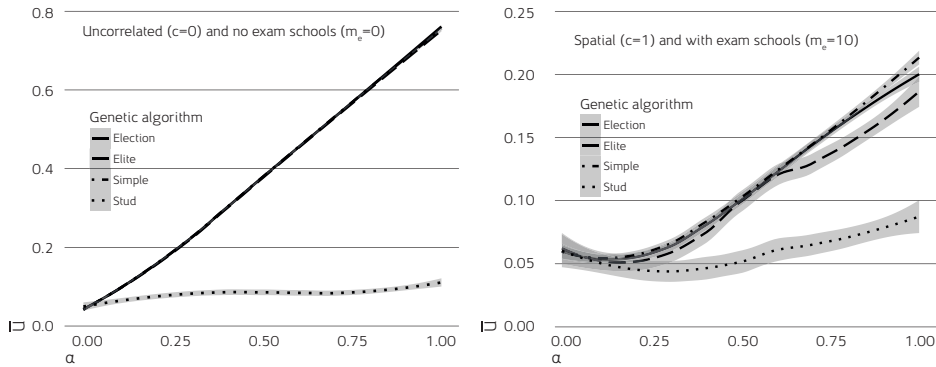
A strategy in the case of the Tallinn mechanism is simply a bit-string. A bit-string is a series of 1-s and 0-s, which respectively stand for reported and not reported preference. Since we limit our agent’s preferences to $k = 15$, the length of the bit-string is 15 bits. Since the Tallinn mechanism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

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

We evaluate four versions of genetic algorithms: simple genetic algorithm; genetic algorithm with election; genetic algorithm with stud selection; and genetic algorithm with elitism. The last three are slight modifications of the simple genetic algorithm. In the election modification,

the agents remember their previous strategy and the corresponding utility. Before the selection operation in the next allocation, each agent picks the strategy with a higher utility from the previously remembered and the newly evaluated strategies (e.g. Riechmann, 2001). In the stud selection, we pick the top 20% of strategies with higher utility and always set one of the strategies in the cross-over operator to be in the top 20% (Simon, 2013). In addition, we ignore the bottom 10% of strategies. In elitism, we keep the top 20% of strategies fixed and only use the remaining strategies in the crossover (Simon, 2013).


Figure 3: Mean Utility



Source: Authors' calculation

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

Figure 4: Ratio of Variance and Mean Utility



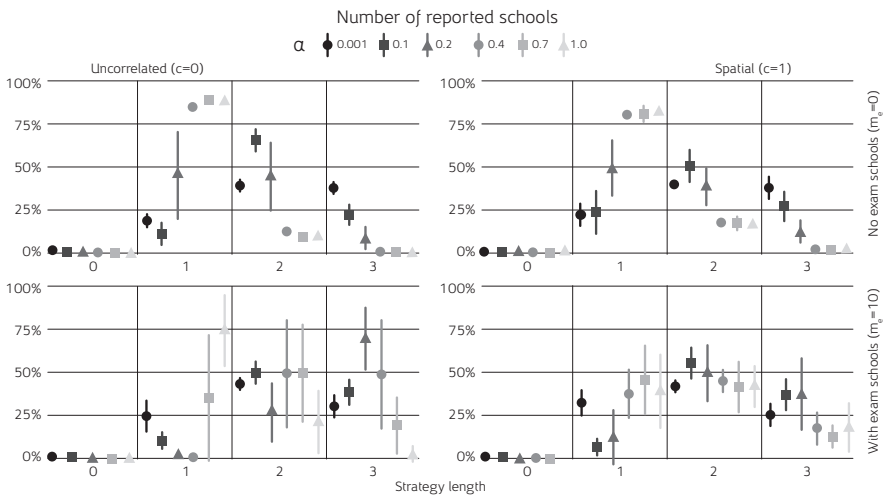
Source: Authors' calculation

4. Results

4.1. Expected Utility Maximising Strategies

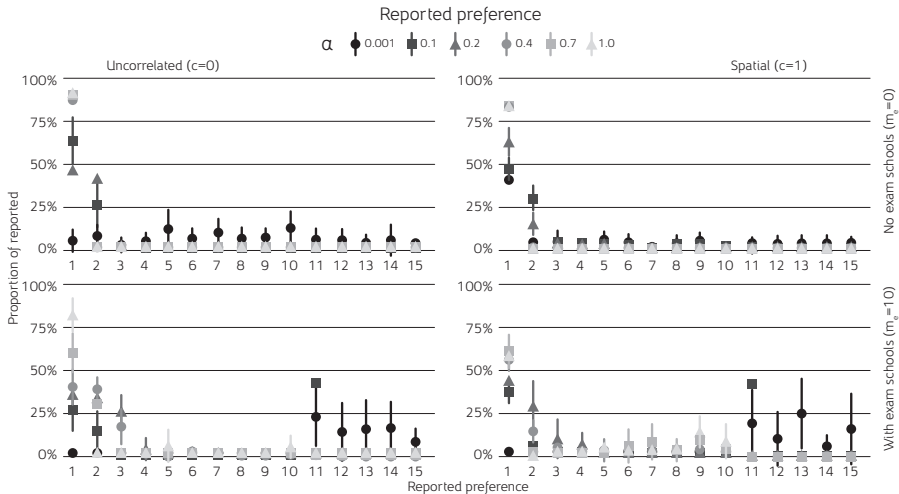
The reported results are divided into four cases. In all of the figures illustrating the results, in the upper left corner the results with no correlation (random) preferences and no exam schools are indicated; in the upper right corner, the results with spatial (2D Euclidean) preferences and no exam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

Figure 5: Reported Strategy Length



Source: Authors' calculation

Figure 6: Reported Preference by Utility Coefficient α



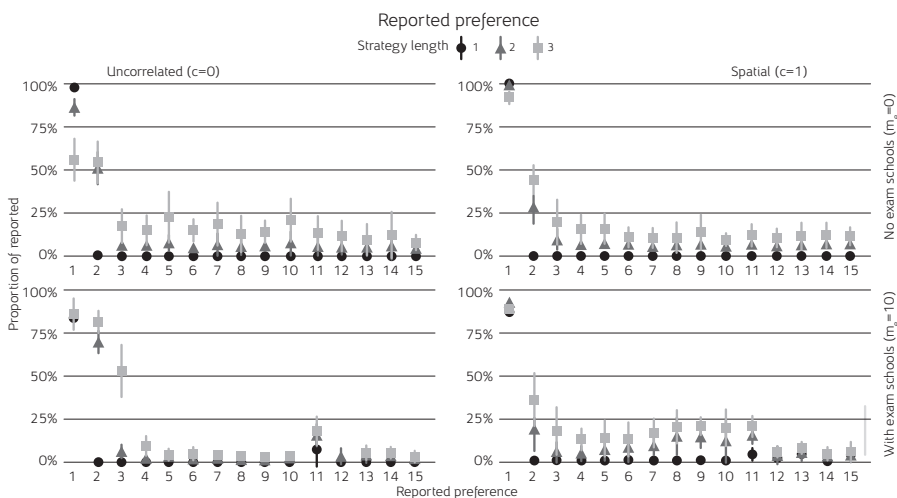
Source: Authors' calculation

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

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

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

Figure 7: Reported Preference by Strategy Length



Source: Authors' calculation

We are also concerned with what to report. Figures 6 and 7 show the preferences reported by the agents' α and the strategy length. We see that without exam schools it is almost always ($\approx 90\%$) optimal for $\alpha \geq 0.4$ to report from the top of the preference list, namely just their first preference. In regard to exam schools and random preferences, optimal reporting depends more on α , but generally the top three schools are reported. Figure 6 illustrates that in the case


of exam schools and spatial preferences with high α , it would be better to report something from the higher and lower ends of exam schools, skipping the middle. Reporting schools lower on the preferences lists probably indicates that those agents would be otherwise unassigned, due to high demand, so they gain at least some utility. For medium α , the first three preferences are almost equally good. For indifferent agents, $\alpha \approx 0.0$, it would be best to randomly pick some schools from the list of regular schools. Also for agents with $\alpha = 0.1$, it would be beneficial to specify their most preferred exam school and most preferred regular school.

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

4.2. Social Welfare

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

Figure 8: Unassigned Agents




Source: Authors' calculation

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

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

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



Source: Authors' calculation

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

In Figure 9 we show the expected utility under the two mechanisms. Expected utility is higher in the Tallinn mechanism compared to the DA results. This is because the strategies of the agents maximise their expected utility, which is greater under the Tallinn mechanism than under DA. A similar result was discovered in the manipulable Boston mechanism (Abdulkadiroğlu et al. 2011). This leads to the conjecture that manipulable mechanisms

provide the option to maximise an agent's expected utility at the risk of being unassigned or assigned to a low ranked preference. Yet, as a result, a large number of agents are unassigned in the Tallinn mechanism.

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

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

5. Conclusion and Discussion

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

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

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

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

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

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

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

References

- Abdulkadiroğlu, A., Agarwal, N., and Pathak, P.A. 2015. The Welfare Effects of Coordinated Assignment: Evidence from the NYC HS Match. Technical Report May, National Bureau of Economic Research, Cambridge, MA.
- Abdulkadiroğlu, A., Che, Y.-K., and Yasuda, Y. 2011. Resolving Conflicting Preferences in School Choice: The Boston Mechanism Reconsidered. *American Economic Review*, Vol. 101, No. 1, pp. 399–410.
- Abdulkadiroğlu, A., Pathak, P. A., and Roth, A. E. 2005a. The New York City High School Match. *American Economic Review*, Vol. 95, No. 2, pp. 364–367.
- Abdulkadiroğlu, A., Pathak, P. A., and Roth, A. E. 2009. Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match. *American Economic Review*, Vol. 99, No. 5, pp. 1954–1978.
- Abdulkadiroğlu, A., Pathak, P. A., Roth, A. E., and Sönmez, T. 2005b. The Boston Public School Match. *American Economic Review*, Vol. 95, No. 2, pp. 368–371.
- Abdulkadiroğlu, A., Pathak, P. A., Roth, A. E., and Sönmez, T. 2006. Changing the Boston School Choice Mechanism. *National Bureau of Economic Research, Working Paper Series*, 11965.

- Abdulkadiroğlu, A. and Sönmez, T. 2003. School Choice: A Mechanism Design Approach. *American Economic Review*, Vol. 93, No. 3, pp. 729–747.
- Arifovic, J. 1994. Genetic Algorithm Learning and the Cobweb Model. *Journal of Economic Dynamics and Control*, Vol. 18, No. 1, pp. 3–28.
- Arifovic, J. 1996. The Behavior of the Exchange Rate in the Genetic Algorithm and Experimental Economies. *Journal of Political Economy*, Vol. 104, No. 3, pp. 510–541.
- Aygün, O. and Bo, I. 2013. College Admission with Multidimensional Reserves: The Brazilian Affirmative Action Case. Available at: <https://www2.bc.edu/inacio-bo/AygunBo2013.pdf> (Accessed 02.08.2016).
- Balinski, M. and Sönmez, T. 1999. A Tale of Two Mechanisms: Student Placement. *Journal of Economic Theory*, Vol. 84, No. 1, pp. 73–94.
- Bogomolnaia, A. and Laslier, J.-F. 2007. Euclidean Preferences. *Journal of Mathematical Economics*, Vol. 43, No. 2, pp. 87–98.
- Chen, S.-H. (Ed.). 2002. *Genetic Algorithms and Genetic Programming in Computational Finance*. Springer US, Boston, MA.
- Chen, S.-H., Kampouridis, M. and Tsang, E. 2011. Microstructure Dynamics and Agent-Based Financial Markets. In: Bosse, T., Geller, A., and Jonker, C. M. (Eds.). *Multi-Agent-Based Simulation XI*, pp. 121–135. Springer: Berlin Heidelberg.
- Chen, S.-H. and Tai, C.-C. 2010. The Agent-Based Double Auction Markets: 15 Years On. In: Takadama, K., Cioffi-Revilla, C., and Deffuant, G. (Eds.). *Simulating Interacting Agents and Social Phenomena*, pp. 119–136. Springer: Japan, Tokyo.
- Duffy, J. 2006. Agent-Based Models and Human Subject Experiments. In: Tesfatsion, L. and Judd, K. L. (Eds.). *Handbook of Computational Economics Volume 2 - Agent-Based Computational Economics*, ch.19, pp. 948–1012. North-Holland.
- Dur, U.M., Kominers, S.D., Pathak, P.A., and Sönmez, T. 2013. The Demise of Walk Zones in Boston: Priorities vs. Precedence in School Choice. *NBER Working Paper Series* 18981.
- Erdil, A. and Ergin, H. 2008. What's the Matter with Tie-Breaking? Improving Efficiency in School Choice. *American Economic Review*, Vol. 98, No. 3, pp. 669–689.
- Erdil, A. and Kumano, T. 2013. Prioritizing Diversity in School Choice. Available at: http://www.matching-in-practice.eu/wp-content/uploads/2013/09/Erdil-Prioritizing_Diversity.pdf (Accessed 02.08.2016).
- Ergin, H. and Sönmez, T. 2006. Games of School Choice under Boston Mechanism. *Journal of Public Economics*, Vol. 90, pp. 215–237.
- Fragiadakis, D. and Troyan, P. 2013. Market Design under Distributional Constraints: Diversity in School Choice and Other Applications. Available at: <http://tippie.uiowa.edu/economics/tow/papers/troyan-spring2014.pdf> (Accessed 02.08.2016).
- Franek, J. and Kresta, A. 2014. Judgment Scales and Consistency Measure in AHP. *Procedia Economics and Finance*, Vol. 12, pp. 164–173.
- Gale, D. and Shapley, L.S. 1962. College Admissions and the Stability of Marriage. *The American Mathematical Monthly*, Vol. 69, No. 1, pp. 9–15.
- Knuth, D.E. 1997. *Seminumerical Algorithms*. Addison-Wesley, Reading, MA, 3rd edition.
- Kominers, S. D. and Sönmez, T. 2013. Designing for Diversity in Matching. Boston College Working Papers in Economics 806.
- McAfee, Preston, R. and Mcmillan, J. 1996. Analyzing Airwaves Auction. *The Journal of Economic Perspectives*, Vol. 10, No. 1, pp. 159–175.

- Milgrom, P.R. 2000. Putting Auction Theory to Work: The Simultaneous Ascending Auction. *Journal of Political Economy*, Vol. 108, No. 2, pp. 245–272.
- Mullin, F.J. and Stalnaker, J.M. 1952. The Matching Plan for Internship Placement: A Report of the First Year's Experience. *Journal of Medical Education*, Vol. 27, No. 3, pp. 193–200.
- Pathak, P.A. and Shi, P. 2013. Simulating Alternative School Choice Options in Boston. Technical report, MIT School Effectiveness and Inequality Initiative.
- Pathak, P.A. and Sönmez, T. 2008. Leveling the Playing Field: Sincere and Sophisticated Players in the Boston Mechanism. *American Economic Review*, Vol. 98, No. 4, pp. 1636–1652.
- Pathak, P.A. and Sönmez, T. 2013. School Admissions Reform in Chicago and England: Comparing Mechanisms by their Vulnerability to Manipulation. *American Economic Review*, Vol. 103, No. 1, pp. 80–106.
- Pöder, K. and Lauri, T. 2014. When Public Acts like Private: The Failure of Estonia's School Choice Mechanism. *European Educational Research Journal*, Vol. 13, No. 2, pp. 220–234.
- Riechmann, T. 2001. Learning in Economics. Springer-Verlag: New York.
- Romero-Medina, A. 1998. Implementation of Stable Solutions in a Restricted Matching Market. *Review of Economic Design*, Vol. 3, No. 2, pp. 137–147.
- Roth, A.E. 1982. The Economics of Matching: Stability and Incentives. *Mathematics of Operations Research*, Vol. 7, No. 4, pp. 617–628.
- Roth, A.E. 1984. The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory. *The Journal of Political Economy*, Vol. 92, No. 6, pp. 991–1016.
- Roth, A.E. 2002. The Economist as Engineer: Game Theory, Experimentation, and Computation as Tools for Design Economics. *Econometrica*, Vol. 70, No. 4, pp. 1341–1378.
- Roth, A.E. 2008. Deferred Acceptance Algorithms: History, Theory, Practice, and Open Questions. *International Journal of Game Theory*, Vol. 36, No. 3-4, pp. 537–569.
- Simon, D. 2013. *Evolutionary Optimization Algorithms: Biologically-inspired and Population-based Approaches to Computer Intelligence*. Wiley.
- Ünver, M.U. 2001. Backward Unraveling Over Time: The Evolution of Strategic Behavior in the Entry Level British Medical Labor Markets. *Journal of Economic Dynamics and Control*, Vol. 25, No. 6-7, pp.1039–1080.
- Ünver, M.U. 2005. On the Survival of some Unstable Two-Sided Matching Mechanisms. *International Journal of Game Theory*, Vol. 33, No. 2, pp. 239–254.
- Varian, H. 2006. *Intermediate Microeconomics*. W. W. Norton & Company, New York, seventh edition.
- Veskioja, T. 2005. *Stable Marriage Problem and College Admission*. PhD thesis, Tallinn University of Technology.
- Zhu, M. 2014. College Admissions in China: A Mechanism Design Perspective. *China Economic Review*, Vol. 30, pp. 618–631.

E Publication 5

Veski, A., Biro, P., Pöder, K., and Lauri, T. (2017). (forthcoming) Efficiency and fair access in kindergarten allocation policy design. *Journal of Mechanism and Institution Design*

Efficiency and fair access in kindergarten allocation policy design

André Veski

Tallinn University of Technology, andre.veski@ttu.ee

Péter Biró

Hungarian Academy of Sciences, biro.peter@krtk.mta.hu

Kaire Põder

Estonian Business School, kaire.poder@ebs.ee

Triin Lauri

Tallinn University, triin.lauri@tlu.ee

Acknowledgements

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 comment of two anonymous reviewers, especially for the recommendation to include results with TTC algorithm.

Efficiency and fair access in kindergarten allocation policy design

Abstract

We study kindergarten allocation practices in an Estonian municipality, Harku. Based on our recommendations, the allocation practices in Harku were redesigned in 2016. The new mechanism provides a child-optimal stable matching, with priorities primarily based on siblings and distance. We evaluate seven policy designs based on the 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 the 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 allocation mechanism in Estonian kindergartens, and potentially elsewhere as well.

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 and Sönmez, 2013; Pathak and Shi, 2013; Ergin and 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 allocation mechanism. However, our theoretical founding relies on the mechanism design literature motivated by related applications, such as school choice (Abdulkadiroğlu and Sönmez, 2003; Abdulkadiroğlu et al., 2005a), college admissions (Biró et al., 2010b; 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 case is 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 not considered or these are 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 the 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 1.2. This mechanism is known to be strategy-proof, and the parents were encourage 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 the distance is used in the priorities.

The classical 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 and 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 schools (Calsamiglia and Güell, 2014, Section 5.1)². In Boston, another relative distance policy was proposed recently by Shi (2015), mainly in order to reach the goal of the city 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 that were typically for nearby kindergartens.

¹In 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.

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).”

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 2 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 and 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 secondary schools). Therefore, we found it important to investigate the effects of the 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 in the probabilities of getting placed in the first choice schools.

In general our results indicate that preference structures, more precisely their endogeneity on proximity, influence the policy design. However, we advocate for 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 dominates others by our fairness criteria, especially when preferences of the families are aligned with policy priorities.

We structure our paper as follows. In Section 1 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 2 we define seven alternative policies and descriptive statistics of our data, including our results from computational experiments. Finally, we discuss additional mechanism design challenges with some policy recommendations in Section 3 and give conclusions in Section 4.

1 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 allocated to 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, which is different from some applications of two-sided markets. In college admissions for example (Gale and 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 and 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 and Sönmez, 2013), and are sometimes even prohibited by law. Thus, stability is a crucial property for most applications.

While there is 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 and Klijn (2009), backed by evidence from laboratory experiments (Calsamiglia et al., 2010) and by practical applications across the world (Pathak and Sönmez, 2013). In addition to advocating for DA, the main policy implications of these studies indicate that for 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 child-birth. 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 the 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.

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 2.2 for policy design alternatives
Matching mechanism	
Decentralised mechanism which has some properties of Serial dictatorship and Boston (Immediate-Acceptance)	Deferred-Acceptance

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 on the top of their list and remained unassigned in the first round, 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 and Sönmez, 2003) and the procedure that had been used in Harku until 2015 was very similar to this.

1.1 Building a mechanism for kindergarten seat allocation

Our redesign of Harku kindergarten allocation mechanism inspired by 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 the child. However, during the subsequent three years, the preferences of the families could have changed. That 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 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 an imminent 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 introduced clearly defined priority metrics and a centralised allocation system that ensures that the criteria are always followed. Moreover, instead of unstable and manipulable Immediate-Acceptance mechanism we proposed the child-proposing DA. This is a standard method for school choice (Abdulkadiroğlu and Sönmez, 2003), which eliminates justified envy, and gives incentive for the families to state their true preferences.

1.2 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 higher. If the number of available places is not exactly divisible by four, then some type of priorities might have more positions available than others.

³<https://www.haldo.ee/>

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 anymore. 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.

2 Policy design

2.1 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 (Matching in Practice, 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 a designer, 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 everyone to get into 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 proximity and siblings are common practice in the case of school and kindergarten choice design, often favoured as the means to sustain community cohesively 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 in their most preferred kindergartens.

2.2 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 multiple 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 (Chen, 2012), is the single lottery. We will investigate both in our computational experiments. This question is discussed further by Ashlagi and Nikzad (2015) and Pathak and 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.

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

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.
- 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.

Table 4: Distance priorities

	absolute (DA1)	walk-zone (DA2, DA3)	3-closest (DA4)	relative (DA7)
k_1	$c_2 \prec c_1 \prec c_4 \prec c_3$	$c_2 \prec \{c_1, c_3, c_4\}$	$\{c_1, c_2, c_4\} \prec c_3$	$c_1 \prec \{c_2, c_4\} \prec c_3$
k_2	$c_4 \prec c_3 \prec c_2 \prec c_1$	$\{c_2, c_3, c_4\} \prec c_1$	$\{c_1, c_2, c_3, c_4\}$	$c_4 \prec \{c_1, c_2, c_3\}$
k_3	$c_2 \prec c_3 \prec c_4 \prec c_1$	$\{c_2, c_3\} \prec \{c_1, c_4\}$	$\{c_1, c_2, c_3, c_4\}$	$c_2 \prec \{c_1, c_3\} \prec c_4$
k_4	$c_3 \prec c_2 \prec c_4 \prec c_1$	$c_3 \prec \{c_1, c_2, c_3\}$	$c_3 \prec \{c_1, c_2, c_4\}$	$c_3 \prec \{c_1, c_2, c_4\}$

2.3 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 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.

Table 5: Harku allocation

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

By using a simpler policy such as the DA1, we saw that there are fewer families receiving a place in 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 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 will not 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.

2.4 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

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

Table 6: Year 2016 comparison of policies using reported preferences

Policy	Mean preference	First	Unassigned	Mean distance (km)	With siblings
Harku	1.68	115	0	4.24	95 %
DA 1	1.76	110	0	4.26	100 %
DA 2 (M) ^a	1.85 (0.01)	98.75 (0.61)	0	4.59 (0.02)	100 % (0.0 %)
DA 2 (S)	1.72 (0.01)	108.05 (0.61)	0	4.44 (0.01)	100 % (0.0 %)
DA 3 (M)	1.83 (0.01)	98.30 (0.79)	0	4.51 (0.02)	95 % (0.25 %)
DA 3 (S)	1.72 (0.01)	107.75 (0.38)	0	4.45 (0.02)	96 % (0.3 %)
DA 4 (M)	1.91 (0.01)	89.25 (1.06)	0	4.53 (0.02)	100 % (0.0 %)
DA 4 (S)	1.75 (0.01)	104.85 (0.7)	0	4.49 (0.01)	100 % (0.0 %)
DA 5	1.76	110	0	4.26	100 %
DA 6	1.76	110	0	4.26	100 %
DA 7 (M)	1.78 (0.01)	107.60 (0.47)	0	4.30 (0.01)	100 % (0.0 %)
DA 7 (S)	1.76 (0.01)	107.75 (0.47)	0	4.31 (0.01)	100 % (0.0 %)

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

preferences, we use the locations and the information on the siblings from the 2016 preference data. The detailed description of how we generated 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 run 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 to 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 $\Pr(r_i \succ r_j \mid d_i < d_j) \rightarrow 1.0$, it produces one of the highest average rank score, one similar to other policies such as DA1, DA5 and DA6. Surprisingly, when the preferences of families are close to random, with conditional probability of $\Pr(r_i \succ r_j \mid d_i < d_j) \rightarrow 0.5$, 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

⁵smoothed with local polynomial regression

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

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 $\Pr(r_i > r_j \mid d_i < d_j) \rightarrow 0.5$. 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.

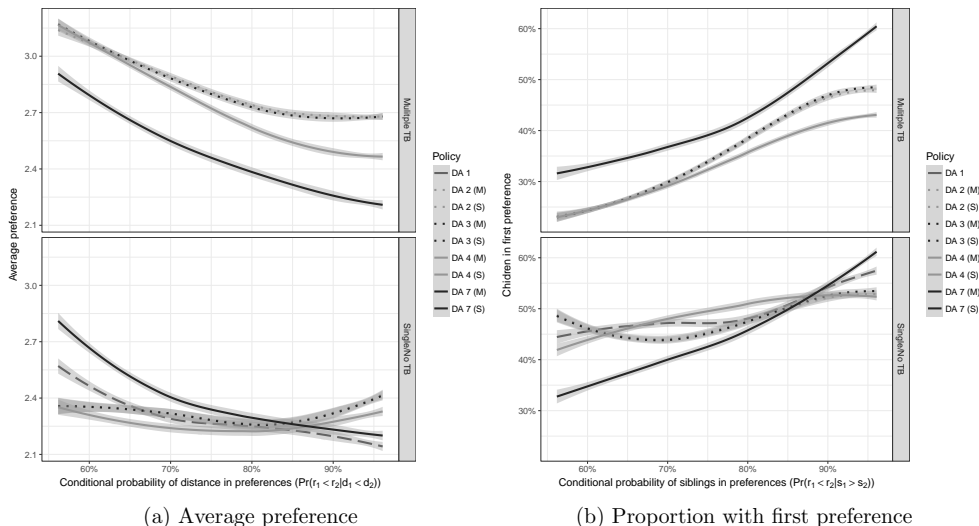


Figure 1: Conditional probability of distance

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 $\Pr(r_i > r_j \mid d_i < d_j) > 0.9$, 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 2: Average distance and conditional probability of distance in preferences

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. Expectedly, 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 often, if preferences are random, the multiple tie-breaking lotteries have a lower average distance than single lotteries. A small improvement of 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 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, then DA7 is better by 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 siblings 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 the 45 degree line, indicating that policies that are below this have some children,

who would prefer a kindergarten with sibling, assigned to a different kindergarten. Multiple lottery policies seem to be better at placing children in the same kindergarten with siblings.




Figure 3: Fairness of access

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 for fairness, or fair (equal) access to kindergartens, which is similar to the 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 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. 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 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 this case with DA7, the fairness is almost as good as with DA4, average distance was a significant improvement over DA4, and average allocated rank very close to DA1.

3 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% (Ministry of Social Affairs, 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 and Kesten, 2014) and also to the design of appeal processes (Dur and 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, and this disturbs 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 for free or with a tuition fee. Furthermore, US cadets (Sönmez and 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 over 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 and 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 smoothly 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, it is still debated 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.

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: a) it is almost costless (as in Hungary); b) there are some significant costs such as renting or having a flat in the area; or c) the family truly has to live there, as 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 they already have attending siblings, 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.

4 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 appeared 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 still 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, the result of DA7 is one of the best policy designs in all aspects. DA7 gives many families their first preference, it has the shortest average distance and even 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 infeasible.

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.

A few interesting aspects of designing a more flexible mechanism might improve the allocation 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 in 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 have an utility function of the form,

$$u_{ij} = \alpha_j + \sum_k \beta_k \cdot x_{kij} + \epsilon_{ij} \quad (1)$$

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,

$$\begin{cases} Pr_{ij}(\text{ranking} = 1) = \frac{e^{u_{ij}}}{\sum_{r=1}^7 e^{u_{ir}}} \\ Pr_{ij}(\text{ranking} = 2) = \frac{e^{u_{ij}}}{\sum_{r=2}^7 e^{u_{ir}}} \\ \dots \\ Pr_{ij}(\text{ranking} = 6) = \frac{e^{u_{ij}}}{\sum_{r=6}^7 e^{u_{ir}}} \end{cases} \quad (2)$$

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.

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 Google (2015)
walking_distance_m	walking distance between family’s home and kindergarten, based on Google (2015)
driving_distance_sec	driving time between family’s home and kindergarten, based on Google (2015)
driving_distance_m	driving distance between family’s home and kindergarten, based on Google (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	$\log(\textit{walking_distance_sec})$
sqrt_walking_distance_sec	$\sqrt{\textit{walking_distance_sec}}$
log_walking_distance_m	$\log(\textit{walking_distance_m})$
sqrt_walking_distance_m	$\sqrt{\textit{walking_distance_m}}$
log_driving_distance_sec	$\log(\textit{driving_distance_sec})$
sqrt_driving_distance_sec	$\sqrt{\textit{driving_distance_sec}}$
log_driving_distance_m	$\log(\textit{driving_distance_m})$
sqrt_driving_distance_m	$\sqrt{\textit{driving_distance_m}}$
log_haversine_distance_m	$\log(\textit{haversine_distance_m})$
sqrt_haversine_distance_m	$\sqrt{\textit{haversine_distance_m}}$

To obtain a full ranking of kindergarten we use the probabilities from (2). For counterfactual 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.

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

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).

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

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





Figure 4: Predictive features

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

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 $\Pr(r_i > r_j \mid d_i < d_j) = 0.81$, $i \neq j$. In 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

¹⁰1.96 standard deviations, 95% probability

B Allocated preferences

Table 9: Year 2016 allocated preference comparison

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)
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)

^aFor policies with lotteries in parenthesis are the standard errors

C Map of the municipality




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

D Results with Top Trading Cycles (TTC)

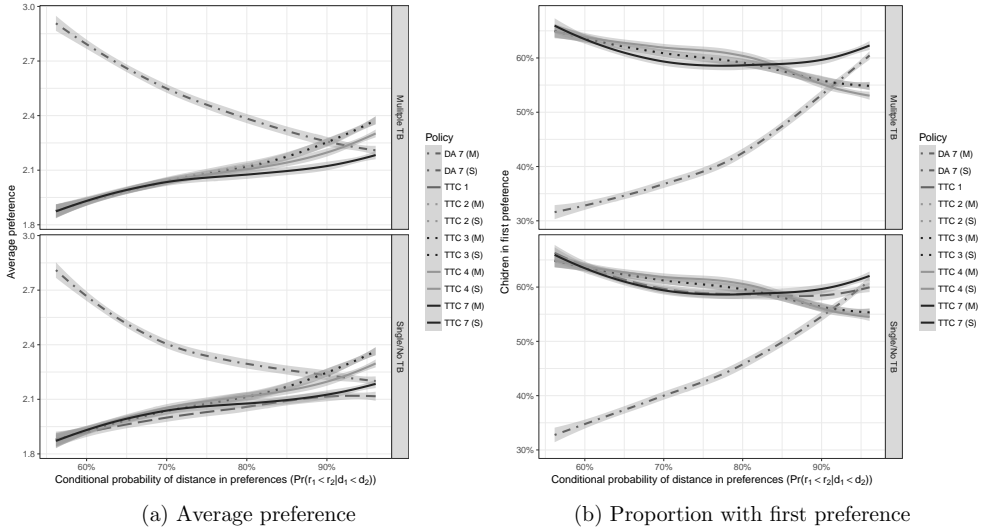


Figure 7: Conditional probability of distance (TTC)

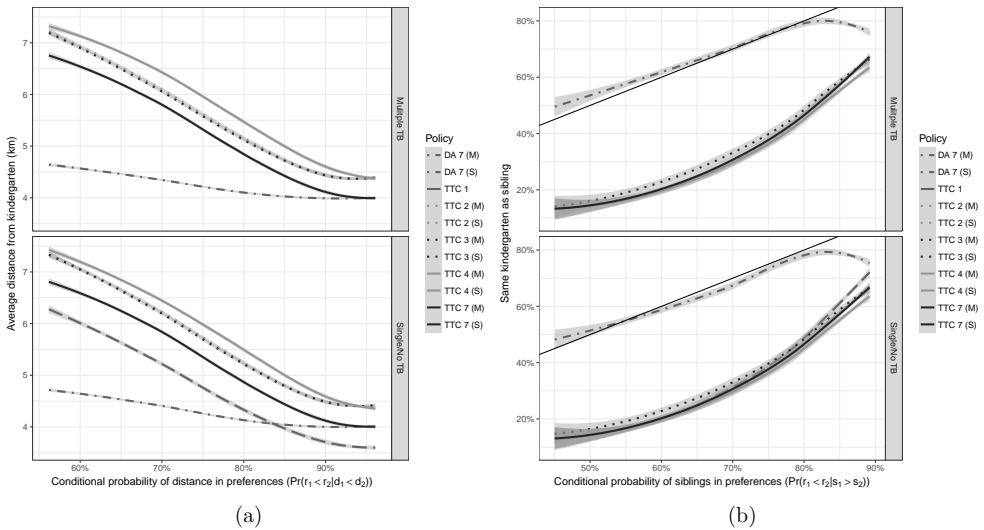


Figure 8: Average distance and conditional probability of distance in preferences (TTC)




Figure 9: Fairness of access

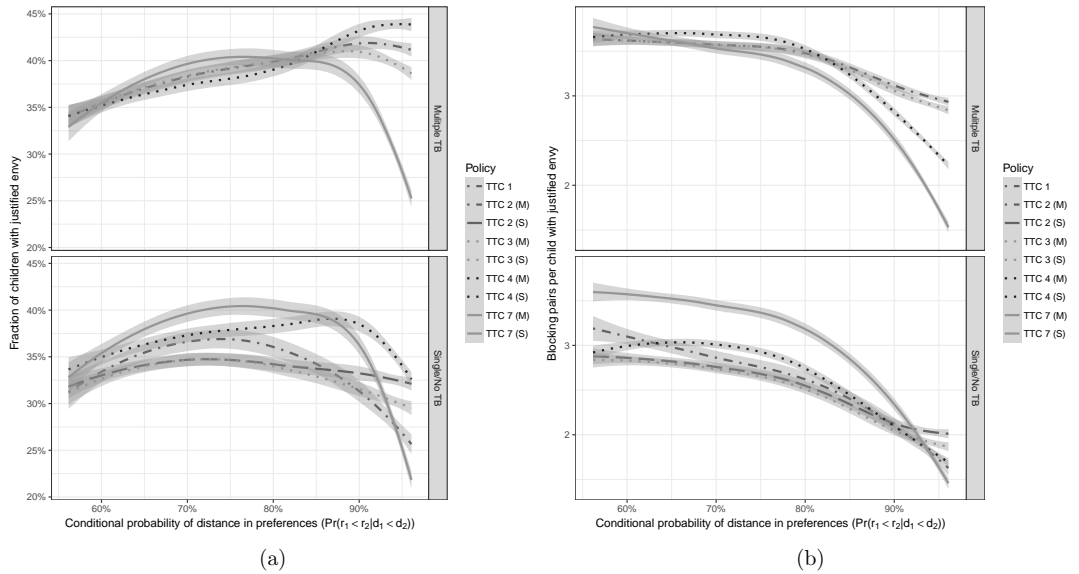


Figure 10: Justified envy and blocking pairs (TTC)

Table 10: Year 2016 comparison of policies using reported preferences (DA and TTC)

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 %)		
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)

^aJE - Justified Envy, BP - Blocking Pairs

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

References

- Abdulkadiroğlu, A., Pathak, P. A., and 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., and Sönmez, T. (2005b). The Boston Public School Match. *American Economic Review*, 95(2):368–371.
- Abdulkadiroğlu, A. and Sönmez, T. (2003). School Choice: A Mechanism Design Approach. *American Economic Review*, 93(3):729–747.
- Andersson, T. (2017). Matching Practices for Elementary Schools - Sweden. MiP Country Profile 24 http://www.matching-in-practice.eu/wp-content/uploads/2017/01/MiP_Profile_No.24.pdf.
- Ashlagi, I., Kanoria, Y., and Leshno, J. D. (2013a). Unbalanced random matching markets. In *Proceedings of the fourteenth ACM conference on Electronic commerce*, pages 27–28, New York. ACM Press.
- Ashlagi, I., Kanoria, Y., and Leshno, J. D. (2013b). Unbalanced random matching markets: the stark effect of competition. <http://web.mit.edu/iashlagi/www/papers/UnbalancedMatchingAKL.pdf> (Accessed 14.09.2016).
- Ashlagi, I. and Nikzad, A. (2015). What matters in tie-breaking rules? How competition guides design. Unpublished working paper.
- Biró, P., Fleiner, T., Irving, R. W., and 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., and McBride, I. (2014). The Hospitals / Residents Problem with Couples: Complexity and Integer Programming Models. In Gudmundsson, J. and Katajainen, J., editors, *Experimental Algorithms*, pages 10–21. Springer International Publishing.
- Biró, P., Manlove, D. F., and Mittal, S. (2010b). Size versus stability in the marriage problem. *Theoretical Computer Science*, 411(16-18):1828–1841.
- Blum, Y., Roth, A. E., and Rothblum, U. G. (1997). Vacancy Chains and Equilibration in Senior-Level Labor Markets. *Journal of Economic Theory*, 76(2):362–411.
- Calsamiglia, C. and 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., and 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 2016-01-18.

- Chen, S.-H., Chang, C.-L., and Du, Y.-R. (2012). Agent-based economic models and econometrics. *The Knowledge Engineering Review*, 27(2):187–219.
- Combe, J., Tercieux, O., and 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. <https://cran.r-project.org/web/packages/mlogit/vignettes/mlogit.pdf> (Accessed 12.12.2006).
- de Haan, M., Gautier, P. A., Oosterbeek, H., and van der Klaauw, B. (2015). The performance of school assignment mechanisms in practice. IZA Discussion Papers 9118.
- Dur, U. M. and 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. and Kesten, O. (2015). The Appeals Process in NYC High School Match. Unpublished working paper.
- Dur, U. M., Kominers, S. D., Pathak, P. A., and Sönmez, T. (2013). The Demise of Walk Zones in Boston: Priorities vs. Precedence in School Choice. NBER Working Paper Series 18981.
- Ergin, H. and Sönmez, T. (2006). Games of school choice under Boston Mechanism. *Journal of Public Economics*, 90:215–237.
- Gale, D. and Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15.
- Google (2015). The Google Maps Distance Matrix API. <https://maps.googleapis.com/maps/api/distancematrix/>. accessed 2015-12-31.
- Haeringer, G. and Klijn, F. (2009). Constrained school choice. *Journal of Economic Theory*, 144(5):1921–1947.
- Hafalir, I. E., Yenmez, M. B., and Yildirim, M. A. (2013). Effective affirmative action in school choice. *Theoretical Economics*, 8(2):325–363.
- Kennes, J., Monte, D., and 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*. American Mathematical Society, Providence.
- Matching in Practice (2016). Matching Practices in Europe. <http://www.matching-in-practice.eu/> (Accessed 27.12.2016).

- McDermid, E. J. and Manlove, D. F. (2010). Keeping partners together: algorithmic results for the hospitals/residents problem with couples. *Journal of Combinatorial Optimization*, 19(3):279–303.
- Ministry of Social Affairs (2015). Peretoetuste, teenuste ja vanemapuhkuste roheline raamat. Technical report, Ministry of Social Affairs.
- Paat, G., Kaarna, R., and Aaviksoo, A. (2011). Meditsiiniliste erivajadustega laste üldhariduse korralduse analüüs. Technical report, Praxis Centre for Policy Studies, Tallinn.
- Pathak, P. A. and Sethuraman, J. (2011). Lotteries in student assignment: An equivalence result. *Theoretical Economics*, 6(1):1–17.
- Pathak, P. A. and Shi, P. (2013). Simulating Alternative School Choice Options in Boston. Technical report, MIT School Effectiveness and Inequality Initiative.
- Pathak, P. A. and Sönmez, T. (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).
- Rawls, J. (1971). *A Theory of Justice*. Harvard University Press, Cambridge.
- 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).
- Shi, P. (2015). Guiding School-Choice Reform through Novel Applications of Operations Research. *Interfaces*, 45(2):117–132.
- Sönmez, T. and 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., and Pennell, H. (2004). School admissions and 'selection' in comprehensive schools: policy and practice. *Oxford Review of Education*, 30(3):347–369.

Acknowledgements

I am grateful to my supervisor, Professor Emeritus Leo Võhandu, for introducing me to the topic of fair division in my bachelor studies, from which grew out my wider interests for mechanism design and matching.

I also wish to thank my second supervisor, Professor Kaire Põder, for including me into her projects on practical matching applications. Many ideas covered in this thesis were inspired by challenges encountered in there projects. Without these the thesis would not be what it is today.

I am especially thankful to Péter Biró for inspiring discussions and encouragement throughout the years of I have been working on matching markets. Without his advice it would had probably taken even longer for me to finish. I am also thankful to the Head of Informatics Institute, Rein Kuusik, for encouragement during the extended years of my studies.

I am also thankful for the participants of 2011 Lipari School on Computational Complex Systems and Santa Fe Institute 2015 Complex Systems Summer School for inspiring ideas and critical discussions that helped be better understand the interactions in markets and arising complexities.

I grateful to the language editor of this thesis and some of my papers, Aylin Gayibli, for the countless hours it must have took to improve my text and work with the \LaTeX source code. Obviously, all the remaining errors are my own.

I am thankful for Estonian IT Academy scholarship for PhD students and Estonian National Doctoral School in Information and Communication Technologies for funding. Parts of this work were supported by Norway Grants V653 and European Cooperation in Science and Technology (COST) action IC1205.

Above all, I wish to thank Eve, for constant support and encouragement during the seemingly never-ending writing. Most of all, Mirtel and Oskar, for showing time and again me that all this is worth the effort.

Abstract

Two-sided markets have been extensively studied in many centralised situations, e.g. hospital-residence matching, school choice, kidney exchange, etc. Matching problems in general are becoming increasingly relevant and even more so in partially decentralised settings, e.g. allocating spare rooms to guest and taxis to passenger. Motivated primarily by the matching market cases in Estonia, i.e. school choice in Tallinn and kindergarten allocation in Harku, we study decentralised matching markets, strategic behaviour and the design of a centralised matching clearing-house.

In order to study decentralised matching, we propose three behavioural agent-based models ranging from random to more informed actions. The two more informed behaviours are previously known from the literature as best and better response dynamics, while the randomised behaviour is novel. We find that to coordinate a good match in a decentralised setting can prove to be difficult. Usually, a significant amount of agents remain unmatched and others are matched to a lower preference than they would be in a deferred-acceptance based centralised matching. In some settings with a small number of agents, short and correlated preferences, there is a high likelihood of finding a stable matching, but it diminishes quickly as preference lists become longer. Additionally, we find that noisy randomised behaviour often has better aggregate outcomes, a lower number of unassigned agents and median agent matched to a higher preference.

The school choice mechanism in Tallinn is partly centrally managed, but incentivises families to misrepresent their preferences. We use genetic algorithms to find a near-equilibrium distribution of strategies in this market. For learning strategies we propose a model for preferences and a cardinal utility function. We find that by using a multiplicative form for the utility function, we manage to recreate some empirical observations, namely how many preferred schools are reported to the central clearing-house. Additionally, we observe that the expected utility may be greater for some families under the manipulable Tallinn school choice mechanism compared to a central allocation by the deferred-acceptance based mechanism. Similar observations have been made in the literature for the manipulable Boston immediate-acceptance mechanism. Therefore, we can conjecture that this might be a property of a certain class of manipulable mechanisms.

Finally, we design an allocation mechanism for kindergartens in Harku. There are three main goals for allocation: considering the allocation preference rank of families, proximity to home and having siblings in the same kindergarten. We propose a design based on the deferred-acceptance algorithm. However, we find that the goals can be considered using different approaches. We propose seven policy designs that were evaluated on the basis of the average allocated rank, distance, siblings in the same kindergarten and fairness, i.e. the chance of being allocated to kindergartens ranked higher on the preference list. The evaluations are made based on collected preferences and parametrised counter-factual preferences. We find that with current preferences, the deterministic as well as some lottery-based policies produce very similar results to each other in terms of the allocated rank. Average distance is the lowest when policies are deterministic, while lotteries entail more fairness. However, when the preference structure changes, some policies can have an increasing effect the allocated preference rank, while still maintaining short distances between homes and kindergartens.

In general, we observed that centralised allocations are superior to decentralised mechanisms. If possible, it would be desirable to design a clearing-house to be simple and strategy-proof for the participants. However, we also observed that some participant might benefit from a manipulable clearing-house. Furthermore, we saw that certain behaviours in decentralised matching produce larger matching and match higher-ranked agents. As the sharing economy motivates more markets where goods are substitutes and need to be individually evaluated, these decentralised matching markets become increasingly important. Moreover, these markets are privately managed, as opposed to public markets. Therefore, considering how to manage and balance those markets in terms of efficiency or revenue is a significant challenge for further research.

Kokkuvõte

Kahepoolseid sobitusturgusid on põhjalikult uuritud mitmetes tsentraliseeritud rakendustes nagu haiglate-residentide sobitus, koolivalik, organivahetus jms. Enamasti leiab neis olukordades rakendust optimaalne edasilükatud vastuvõtu mehhanism. Sobitusülesanded üldiselt muutuvad järjest olulisemaks ja isegi rohkem osaliselt detsentraliseeritud jagamis- ja platvormimajanduses. Peamiselt motiveeritud sobitusturgudest Eestis, nagu koolivalik Tallinnas ja lasteaiakohtade jagamine Harkus, uurime detsentraliseeritud sobitusturge, strateegilist käitumist ja lõpuks turudisaini võimalusi tsentraliseeritud sobitamiseks.

Detsentraliseeritud sobituse uurimiseks pakume välja kolm agentide käitumismudelit skaalal juhuslikust kuni informeeritud agendi käitumiseni. Informeeritud mudelid on kirjandusest varem tuntud kui parema ja parima vastuse dünaamikad, kuid juhusliku käitumise mudel on uudne sobitusturgude kontekstis. Tulemused näitavad, et suurel hulgal agentidel on keeruline detsentraliseeritud turul leida hea sobitus. Enamus olukordades jääb suur hulk agente ilma sobituseta ja ülejäänud on sobitatud madalama eelistusega kui seda oleks tsentraliseeritud turul. Samas, mõnes olukorras, kus on vähe agente väheste eelistustega, on suur võimalus neil ka omavahel detsentraliseeritult leida stabiilne sobitus, kuid see võimalus kahaneb kiirelt agentide arvu kasvades. Üllatusena selgub, et agentide juhuslik käitumine on turu kui terviku mõttes isegi parem kui informeeritumad käitumised – suurem arv agente on sobitatud ja seda ka eelistatuma agendiga. Lisaks leiame, et see on põhjustatud tänu väiksemale uuesti sobituse võimalusele juhusliku käitumisega mudelis. Lisades takistuse uuesti sobituseks ka informeeritumates mudelites, saame tulemuseks sarnased tulemused nagu juhusliku käitumisega mudelis.

Tallinnas on koolivaliku mehhanism osaliselt keskselt juhitud, kuid see ei motiveeri siiski vanemaid esitama tõeseid koolieelistusi. Me kasutame geneetilisi algoritme, et leida tasakaalupunkti lähedasi strateegiaid eelistuste esitamiseks. Strateegiate õppimiseks pakume eelistuste ja kardinaalse heaolufunktsiooni mudeli. Kasutades multiplikatiivset heaolufunktsiooni kuju, suudame taastada mõned empiirilised tulemused, nagu mitu kooli on kasulik esitada. Lisaks selgub, et paljudel peredel (agentidel) on kasulik Tallinna koolivaliku mehhanism, kui optimaalne edasilükatud vas-

tuvõtu mehhanism. Sarnane tulemus avaldub ka kirjandusest tuntud koheste vastuvõtu ehk Bostoni mehhanismis, kus mõnel agendil strateegiliselt manipuleerides on võimalik saavutada parem tulemus kui optimaalses jagamises. Siit tekib hüpotees, et ehk on see omadus mingil klassil manipuleeritavatel mehhanismidel. Viimases osas pakume võimaliku disaini Harkus vallas lasteaiakohtade jagamiseks. Harkul on jagamisel kolm eesmärki: arvestada perede eelistustega; pakkuda kodulähedane lasteaiakoht; ning õed-vennad samas lasteaias. Hindame võimalikke prioriteetide seadmise viise, kasutades jagamiseks edasilükatud vastuvõtu mehhanismi. Lisaks eeltoodud kriteeriumitele, hindame ka sobituse õiglust läbi pere (agendi) tõenäosusliku võimaluse saada koht eelistatuimas lasteaias. Kasutame hindamiseks nii kogutud eelistusi kui ka mudelipõhiseid eelistusi, et arvestada ka tundlikkust eelistuste muutumisele. Kogutud eelistustega selgub, et suur osa prioriteetide seadmise viise annavad sarnase määratud eelistuse positsiooni tulemuse. Keskmise kaugus lasteaiast on madalaim deterministlike prioriteetidega samas kui, oodatult, on õiglus suurem loteriipõhiste prioriteetidega. Seevastu, kui eelistuste struktuur peaks muutuma, siis mõned prioriteetide disainid määravad rohkem lapsi eelistatud lasteaeda, samas säilitades väikest kaugust lasteaiast ja kodu vahel.

Üldiselt kinnitasime kirjanduses varem esitatud väidet, et tsentraalsed jagamised on mitmes mõttes paremad kui detsentraalsed turud. Tsentraalne turg peaks olema osalistele lihtne, et strateegiliselt mõtlemata oleks võimalik siiski saavutada kõigile parim tulemus. Vastasel juhul saavad ainult mõned, osavamad, agendid endale parema tulemuse ning teistel tuleb teine kord leppida kehvema tulemusega, kui optimaalsemalt käitumisega õnnestuks saavutada. Siiski näeme ka, et detsentraalseid turgusid saab mõningal määral juhtida, kas läbi parema info, tehnilike tõkete või eelistuste struktuuri haldamisega. Samas on nende juhtimismehhanismide mõju väga sarnane ja reaalse maailma olukorras ei ole parima vahendi valimine üheselt määratud, mis ka kinnitab põhilist kompleksüsteemide raskesti haldamise omadust. Jagamismajanduse olukorras tekib rohkem turge, kus jagatavad kaubad on väga erinevad ja vajavad eraldi hindamist. Seega muutuvad detsentraalsed sobitusturud veelgi olulisemaks. Lisaks on need turud enamasti kasumit maksimeerivad eraettevõtted, mitte sotsiaalset heaolu maksimeerivad avalikud turud, mistõttu nende juhtimine ja tasakaalustamine on keeruline väljakutse edasiseks uurimiseks.

Curriculum Vitae

Contact data

Name: Andre Veski
Email: andre.veski@ttu.ee

Education

2008 – 2017 Tallinn University of Technology (TUT),
PhD studies in computer science
2005 – 2008 TUT, MSc studies in computer science
2000 – 2005 TUT, BSc studies in computer science

Special courses

2015 Santa Fe Institute Complex Systems Summer School
2013 Summer School on Matching Problems, Markets, and
Mechanisms
2011 Lipari School on Computational Complex Systems,
Game Theoretic approach to Computational Complex
Systems

Professional employment

2014 – 2016 Tallinn University of Technology; junior researcher
2006 – 2014 Skype Technologies; data warehouse developer, data
analyst, researcher
2003 – 2005 Elisa Mobiilsideteenused AS, systems engineer
2000 – 2002 OÜ Jaotuskeskus, IT specialist

Research activity

Veski, A., Biro, P., Pöder, K., and Lauri, T. (2017). (forthcoming) Efficiency and fair access in kindergarten allocation policy design. *Journal of Mechanism and Institution Design*

Veski, A. and Pöder, K. (2017). Zero-intelligence agents looking for a job. *Journal of Economic Interaction and Coordination*

Veski, A. and Pöder, K. (2016). Strategies in Tallinn school choice mechanism. *Research in Economics and Business: Central and Eastern Europe*, 8(1):5–24

Pöder, K., Lauri, T., and Veski, A. (2016). Does School Admission by Zoning Affect Educational Inequality? A Study of Family Background Effect in Estonia, Finland, and Sweden. *Scandinavian Journal of Educational Research*, pages 1–21

Pöder, K., Lauri, T., Karmo, K., Veski, A., Roosalu, T., and Simm, K. (2015). *Lasteaiakohtade jagamine Soovitud kohalikele omavalitsustele*. Gutenbergi Pojad, Tallinn

Veski, A. and Pöder, K. (2015). Primary School Choice in Tallinn: Data and Simulations. TUTECON Working Paper No. WP-2015/1

Pöder, K., Veski, A., and Lauri, T. (2014). Eesti põhikooli- ja gümnaasiumivõrgu analüüs aastaks 2020. Technical report, PRAXIS Poliitikauuringute keskus

Veski, A. (2014). Price of Invisibility: Statistics of centralised and decentralised matching markets. In MacKerrow, E., Terano, T., Squazzon, F., and Sichman, J. S., editors, *Proceedings of the 5th. World Congress on Social Simulation*, pages 18–29, Sao Paulo

Veski, A. (2012). Some issues in Multi-Agent Resource Allocation. In *Proceedings of the 6th Annual Conference of the Estonian National Doctoral School in Information and Communication Technologies*, pages 101–104, Tallinn. TUT Press

Veski, A. and Vöhandu, L. (2011). Two Player Fair Division Problem with Uncertainty. In Barzdins, J. and Kirikova, M., editors, *Frontiers in Artificial Intelligence and Applications*, pages 394–407. IOS Press, Amsterdam

Veski, A. and Võhandu, L. (2010). Another View on Territory Fair Division. In Barzdins, J. and Kirikova, M., editors, *Databases and information systems : proceedings of the Ninth International Baltic Conference*, pages 261–276, Riga. University of Latvia Press

Elulookirjeldus

Kontaktandmed

Name: Andre Veski
Email: andre.veski@ttu.ee

Hariduskäik

2008 – 2017 Tallinn Tehnikaülikool (TTÜ),
doktoriõpingud informaatikas
2005 – 2008 TTÜ, magistriõpingud informaatikas
2000 – 2005 TTÜ, bakalaureuseõpingud informaatikas

Täiendkoolitus

2015 Santa Fe Institute Complex Systems Summer School
2013 Summer School on Matching Problems, Markets, and
Mechanisms
2011 Lipari School on Computational Complex Systems,
Game Theoretic approach to Computational Complex
Systems

Teenistuskäik

2014 – 2016 Tallinn Tehnikaülikool; nooremteadur
2006 – 2014 Skype Technologies; andmelao arendaja,
andmeanalüütik, teadur
2003 – 2005 Elisa Mobiilsideteenused AS, süsteemiinsener
2000 – 2002 OÜ Jaotuskeskus, IT spetsialist

Teadustegevus

Veski, A., Biro, P., Põder, K., and Lauri, T. (2017). (forthcoming) Efficiency and fair access in kindergarten allocation policy design. *Journal of Mechanism and Institution Design*

Veski, A. and Põder, K. (2017). Zero-intelligence agents looking for a job. *Journal of Economic Interaction and Coordination*

Veski, A. and Põder, K. (2016). Strategies in Tallinn school choice mechanism. *Research in Economics and Business: Central and Eastern Europe*, 8(1):5–24

Põder, K., Lauri, T., and Veski, A. (2016). Does School Admission by Zoning Affect Educational Inequality? A Study of Family Background Effect in Estonia, Finland, and Sweden. *Scandinavian Journal of Educational Research*, pages 1–21

Põder, K., Lauri, T., Karmo, K., Veski, A., Roosalu, T., and Simm, K. (2015). *Lasteaiakohtade jagamine Soovitud kohalikele omavalitsustele*. Gutenbergi Pojad, Tallinn

Veski, A. and Põder, K. (2015). Primary School Choice in Tallinn: Data and Simulations. TUTECON Working Paper No. WP-2015/1

Põder, K., Veski, A., and Lauri, T. (2014). Eesti põhikooli- ja gümnaasiumivõrgu analüüs aastaks 2020. Technical report, PRAXIS Poliitikauuringute keskus

Veski, A. (2014). Price of Invisibility: Statistics of centralised and decentralised matching markets. In MacKerrow, E., Terano, T., Squazzon, F., and Sichman, J. S., editors, *Proceedings of the 5th. World Congress on Social Simulation*, pages 18–29, Sao Paulo

Veski, A. (2012). Some issues in Multi-Agent Resource Allocation. In *Proceedings of the 6th Annual Conference of the Estonian National Doctoral School in Information and Communication Technologies*, pages 101–104, Tallinn. TUT Press

Veski, A. and Vöhandu, L. (2011). Two Player Fair Division Problem with Uncertainty. In Barzdins, J. and Kirikova, M., editors, *Frontiers in Artificial Intelligence and Applications*, pages 394–407. IOS Press, Amsterdam

Veski, A. and Võhandu, L. (2010). Another View on Territory Fair Division. In Barzdins, J. and Kirikova, M., editors, *Databases and information systems : proceedings of the Ninth International Baltic Conference*, pages 261–276, Riga. University of Latvia Press

**DISSERTATIONS DEFENDED AT
TALLINN UNIVERSITY OF TECHNOLOGY ON
INFORMATICS AND SYSTEM ENGINEERING**

1. **Lea Elmik**. Informational Modelling of a Communication Office. 1992.
2. **Kalle Tammemäe**. Control Intensive Digital System Synthesis. 1997.
3. **Eerik Lossmann**. Complex Signal Classification Algorithms, Based on the Third-Order Statistical Models. 1999.
4. **Kaido Kikkas**. Using the Internet in Rehabilitation of People with Mobility Impairments – Case Studies and Views from Estonia. 1999.
5. **Nazmun Nahar**. Global Electronic Commerce Process: Business-to-Business. 1999.
6. **Jevgeni Riipulk**. Microwave Radiometry for Medical Applications. 2000.
7. **Alar Kuusik**. Compact Smart Home Systems: Design and Verification of Cost Effective Hardware Solutions. 2001.
8. **Jaan Raik**. Hierarchical Test Generation for Digital Circuits Represented by Decision Diagrams. 2001.
9. **Andri Riid**. Transparent Fuzzy Systems: Model and Control. 2002.
10. **Marina Brik**. Investigation and Development of Test Generation Methods for Control Part of Digital Systems. 2002.
11. **Raul Land**. Synchronous Approximation and Processing of Sampled Data Signals. 2002.
12. **Ants Ronk**. An Extended Block-Adaptive Fourier Analyser for Analysis and Reproduction of Periodic Components of Band-Limited Discrete-Time Signals. 2002.
13. **Toivo Paavle**. System Level Modeling of the Phase Locked Loops: Behavioral Analysis and Parameterization. 2003.
14. **Irina Astrova**. On Integration of Object-Oriented Applications with Relational Databases. 2003.
15. **Kuldar Taveter**. A Multi-Perspective Methodology for Agent-Oriented Business Modelling and Simulation. 2004.
16. **Taivo Kangilaski**. Eesti Energia käiduhaldussüsteem. 2004.
17. **Artur Jutman**. Selected Issues of Modeling, Verification and Testing of Digital Systems. 2004.
18. **Ander Tenno**. Simulation and Estimation of Electro-Chemical Processes in Maintenance-Free Batteries with Fixed Electrolyte. 2004.

19. **Oleg Korolkov.** Formation of Diffusion Welded Al Contacts to Semiconductor Silicon. 2004.
20. **Risto Vaarandi.** Tools and Techniques for Event Log Analysis. 2005.
21. **Marko Koort.** Transmitter Power Control in Wireless Communication Systems. 2005.
22. **Raul Savimaa.** Modelling Emergent Behaviour of Organizations. Time-Aware, UML and Agent Based Approach. 2005.
23. **Raido Kurel.** Investigation of Electrical Characteristics of SiC Based Complementary JBS Structures. 2005.
24. **Rainer Taniloo.** Ökonoomsete negatiivse diferentsiaaltakistusega astmete ja elementide disainimine ja optimeerimine. 2005.
25. **Pauli Lallo.** Adaptive Secure Data Transmission Method for OSI Level I. 2005.
26. **Deniss Kumlander.** Some Practical Algorithms to Solve the Maximum Clique Problem. 2005.
27. **Tarmo Veskiõja.** Stable Marriage Problem and College Admission. 2005.
28. **Elena Fomina.** Low Power Finite State Machine Synthesis. 2005.
29. **Eero Ivask.** Digital Test in WEB-Based Environment 2006.
30. **Виктор Войтович.** Разработка технологий выращивания из жидкой фазы эпитаксиальных структур арсенида галлия с высоковольтным p-n переходом и изготовления диодов на их основе. 2006.
31. **Tanel Alumäe.** Methods for Estonian Large Vocabulary Speech Recognition. 2006.
32. **Erki Eessaar.** Relational and Object-Relational Database Management Systems as Platforms for Managing Softwareengineering Artefacts. 2006.
33. **Rauno Gordon.** Modelling of Cardiac Dynamics and Intracardiac Bio-impedance. 2007.
34. **Madis Listak.** A Task-Oriented Design of a Biologically Inspired Underwater Robot. 2007.
35. **Elmet Orasson.** Hybrid Built-in Self-Test. Methods and Tools for Analysis and Optimization of BIST. 2007.
36. **Eduard Petlenkov.** Neural Networks Based Identification and Control of Nonlinear Systems: ANARX Model Based Approach. 2007.
37. **Toomas Kirt.** Concept Formation in Exploratory Data Analysis: Case Studies of Linguistic and Banking Data. 2007.
38. **Juhan-Peep Ernits.** Two State Space Reduction Techniques for Explicit State Model Checking. 2007.

39. **Innar Liiv**. Pattern Discovery Using Seriation and Matrix Reordering: A Unified View, Extensions and an Application to Inventory Management. 2008.
40. **Andrei Pokatilov**. Development of National Standard for Voltage Unit Based on Solid-State References. 2008.
41. **Karin Lindroos**. Mapping Social Structures by Formal Non-Linear Information Processing Methods: Case Studies of Estonian Islands Environments. 2008.
42. **Maksim Jenihhin**. Simulation-Based Hardware Verification with High-Level Decision Diagrams. 2008.
43. **Ando Saabas**. Logics for Low-Level Code and Proof-Preserving Program Transformations. 2008.
44. **Ilja Tšahhиров**. Security Protocols Analysis in the Computational Model – Dependency Flow Graphs-Based Approach. 2008.
45. **Toomas Ruuben**. Wideband Digital Beamforming in Sonar Systems. 2009.
46. **Sergei Devadze**. Fault Simulation of Digital Systems. 2009.
47. **Andrei Krivošei**. Model Based Method for Adaptive Decomposition of the Thoracic Bio-Impedance Variations into Cardiac and Respiratory Components. 2009.
48. **Vineeth Govind**. DfT-Based External Test and Diagnosis of Mesh-like Networks on Chips. 2009.
49. **Andres Kull**. Model-Based Testing of Reactive Systems. 2009.
50. **Ants Torim**. Formal Concepts in the Theory of Monotone Systems. 2009.
51. **Erika Matsak**. Discovering Logical Constructs from Estonian Children Language. 2009.
52. **Paul Annus**. Multichannel Bioimpedance Spectroscopy: Instrumentation Methods and Design Principles. 2009.
53. **Maris Tõnso**. Computer Algebra Tools for Modelling, Analysis and Synthesis for Nonlinear Control Systems. 2010.
54. **Aivo Jürgenson**. Efficient Semantics of Parallel and Serial Models of Attack Trees. 2010.
55. **Erkki Joasoon**. The Tactile Feedback Device for Multi-Touch User Interfaces. 2010.
56. **Jürgo-Sören Preden**. Enhancing Situation – Awareness Cognition and Reasoning of Ad-Hoc Network Agents. 2010.
57. **Pavel Grigorenko**. Higher-Order Attribute Semantics of Flat Languages. 2010.
58. **Anna Rannaste**. Hierarcical Test Pattern Generation and Untestability Identification Techniques for Synchronous Sequential Circuits. 2010.

59. **Sergei Strik**. Battery Charging and Full-Featured Battery Charger Integrated Circuit for Portable Applications. 2011.
60. **Rain Ottis**. A Systematic Approach to Offensive Volunteer Cyber Militia. 2011.
61. **Natalja Sleptšuk**. Investigation of the Intermediate Layer in the Metal-Silicon Carbide Contact Obtained by Diffusion Welding. 2011.
62. **Martin Jaanus**. The Interactive Learning Environment for Mobile Laboratories. 2011.
63. **Argo Kasemaa**. Analog Front End Components for Bio-Impedance Measurement: Current Source Design and Implementation. 2011.
64. **Kenneth Geers**. Strategic Cyber Security: Evaluating Nation-State Cyber Attack Mitigation Strategies. 2011.
65. **Riina Maigre**. Composition of Web Services on Large Service Models. 2011.
66. **Helena Kruus**. Optimization of Built-in Self-Test in Digital Systems. 2011.
67. **Gunnar Pihoo**. Archetypes Based Techniques for Development of Domains, Requirements and Software. 2011.
68. **Juri Gavšin**. Intrinsic Robot Safety Through Reversibility of Actions. 2011.
69. **Dmitri Mihhailov**. Hardware Implementation of Recursive Sorting Algorithms Using Tree-like Structures and HFSM Models. 2012.
70. **Anton Tšertov**. System Modeling for Processor-Centric Test Automation. 2012.
71. **Sergei Kostin**. Self-Diagnosis in Digital Systems. 2012.
72. **Mihkel Tagel**. System-Level Design of Timing-Sensitive Network-on-Chip Based Dependable Systems. 2012.
73. **Juri Belikov**. Polynomial Methods for Nonlinear Control Systems. 2012.
74. **Kristina Vassiljeva**. Restricted Connectivity Neural Networks based Identification for Control. 2012.
75. **Tarmo Robal**. Towards Adaptive Web – Analysing and Recommending Web Users` Behaviour. 2012.
76. **Anton Karputkin**. Formal Verification and Error Correction on High-Level Decision Diagrams. 2012.
77. **Vadim Kimlaychuk**. Simulations in Multi-Agent Communication System. 2012.
78. **Taavi Viilukas**. Constraints Solving Based Hierarchical Test Generation for Synchronous Sequential Circuits. 2012.

79. **Marko Kääramees.** A Symbolic Approach to Model-based Online Testing. 2012.
80. **Enar Reilent.** Whiteboard Architecture for the Multi-agent Sensor Systems. 2012.
81. **Jaan Ojarand.** Wideband Excitation Signals for Fast Impedance Spectroscopy of Biological Objects. 2012.
82. **Igor Aleksejev.** FPGA-based Embedded Virtual Instrumentation. 2013.
83. **Juri Mihhailov.** Accurate Flexible Current Measurement Method and its Realization in Power and Battery Management Integrated Circuits for Portable Applications. 2013.
84. **Tõnis Saar.** The Piezo-Electric Impedance Spectroscopy: Solutions and Applications. 2013.
85. **Ermo Täks.** An Automated Legal Content Capture and Visualisation Method. 2013.
86. **Uljana Reinsalu.** Fault Simulation and Code Coverage Analysis of RTL Designs Using High-Level Decision Diagrams. 2013.
87. **Anton Tšepurov.** Hardware Modeling for Design Verification and Debug. 2013.
88. **Ivo Mürsepp.** Robust Detectors for Cognitive Radio. 2013.
89. **Jaas Ježov.** Pressure sensitive lateral line for underwater robot. 2013.
90. **Vadim Kaparin.** Transformation of Nonlinear State Equations into Observer Form. 2013.
92. **Reeno Reeder.** Development and Optimisation of Modelling Methods and Algorithms for Terahertz Range Radiation Sources Based on Quantum Well Heterostructures. 2014.
93. **Ants Koel.** GaAs and SiC Semiconductor Materials Based Power Structures: Static and Dynamic Behavior Analysis. 2014.
94. **Jaan Übi.** Methods for Coopetition and Retention Analysis: An Application to University Management. 2014.
95. **Innokenti Sobolev.** Hyperspectral Data Processing and Interpretation in Remote Sensing Based on Laser-Induced Fluorescence Method. 2014.
96. **Jana Toompuu.** Investigation of the Specific Deep Levels in p -, i - and n -Regions of GaAs p^+pin-n^+ Structures. 2014.
97. **Taavi Salumäe.** Flow-Sensitive Robotic Fish: From Concept to Experiments. 2015.
98. **Yar Muhammad.** A Parametric Framework for Modelling of Bioelectrical Signals. 2015.
99. **Ago Mõlder.** Image Processing Solutions for Precise Road Profile Measurement Systems. 2015.

100. **Kairit Sirts.** Non-Parametric Bayesian Models for Computational Morphology. 2015.
101. **Alina Gavrijaševa.** Coin Validation by Electromagnetic, Acoustic and Visual Features. 2015.
102. **Emiliano Pastorelli.** Analysis and 3D Visualisation of Microstructured Materials on Custom-Built Virtual Reality Environment. 2015.
103. **Asko Ristolainen.** Phantom Organs and their Applications in Robotic Surgery and Radiology Training. 2015.
104. **Aleksei Tepljakov.** Fractional-order Modeling and Control of Dynamic Systems. 2015.
105. **Ahti Lohk.** A System of Test Patterns to Check and Validate the Semantic Hierarchies of Wordnet-type Dictionaries. 2015.
106. **Hanno Hantson.** Mutation-Based Verification and Error Correction in High-Level Designs. 2015.
107. **Lin Li.** Statistical Methods for Ultrasound Image Segmentation. 2015.
108. **Aleksandr Lenin.** Reliable and Efficient Determination of the Likelihood of Rational Attacks. 2015.
109. **Maksim Gorev.** At-Speed Testing and Test Quality Evaluation for High-Performance Pipelined Systems. 2016.
110. **Mari-Anne Meister.** Electromagnetic Environment and Propagation Factors of Short-Wave Range in Estonia. 2016.
111. **Syed Saif Abrar.** Comprehensive Abstraction of VHDL RTL Cores to ESL SystemC. 2016.
112. **Arvo Kaldmäe.** Advanced Design of Nonlinear Discrete-time and Delayed Systems. 2016.
113. **Mairo Leier.** Scalable Open Platform for Reliable Medical Sensorics. 2016.
114. **Georgios Giannoukos.** Mathematical and Physical Modelling of Dynamic Electrical Impedance. 2016.
115. **Aivo Anier.** Model Based Framework for Distributed Control and Testing of Cyber-Physical Systems. 2016.
116. **Denis Firsov.** Certification of Context-Free Grammar Algorithms. 2016.
117. **Sergei Astatpov.** Distributed Signal Processing for Situation Assessment in Cyber-Physical Systems. 2016.
118. **Erkki Moorits.** Embedded Software Solutions for Development of Marine Navigation Light Systems. 2016.
119. **Andres Ojamaa.** Software Technology for Cyber Security Simulations. 2016.
120. **Gert Toming.** Fluid Body Interaction of Biomimetic Underwater Robots. 2016.

121. **Kadri Umbleja.** Competence Based Learning – Framework, Implementation, Analysis and Management of Learning Process. 2017.
122. **Andres Hunt.** Application-Oriented Performance Characterization of the Ionic Polymer Transducers (IPTs). 2017.
123. **Niccolò Veltri.** A Type-Theoretical Study of Nontermination. 2017.
124. **Tauseef Ahmed.** Radio Spectrum and Power Optimization Cognitive Techniques for Wireless Body Area Networks. 2017.