

Experiments on Matching Markets: A Survey^{*}

Rustamdjan Hakimov (University of Lausanne & WZB Berlin)

Dorothea Kübler (WZB Berlin & TU Berlin)

This version: April 2019

Abstract

The paper surveys the experimental literature on matching markets. It covers house allocation, school choice, and two-sided matching markets such as college admissions. The main focus of the survey is on truth-telling and strategic manipulations by the agents, on the stability and efficiency of the matching outcome, as well as on the distribution of utility.

JEL Classification: C92, D47, D83

Keywords: Experiments; matching markets; survey

^{*} The survey has benefited from frequent interactions with members of the Matching-in-Practice network. We would also like to thank Puja Bhattacharya, Inácio Bó, Marie-Pierre Dagnies, Bob Hammond, Bettina Klaus, Andy Schotter, and Roel van Veldhuizen for helpful comments on earlier drafts of the paper as well as Jennifer Rontganger for copy editing. Dorothea Kübler thanks Al Roth and Stanford University where part of this paper was written. Financial support from the DFG through CRC TRR 190 is gratefully acknowledged.

Table of contents

1. Introduction	3
2. Basic concepts of matching theory	5
3. One-sided matching	12
3.1 House allocation	12
3.2 School choice	15
3.2.1 Comparison of mechanisms	16
3.2.2 Learning and effect of market size	21
3.2.3 Nudging, chatting, and advice	25
3.2.4 Determinants of reporting strategies: biases, risk aversion, and cognitive ability	28
3.2.5 Dynamic mechanisms	32
3.2.6 Constrained rank-order lists	36
3.2.7 Affirmative action and quotas	38
3.2.8 Information acquisition	40
3.2.9 Preferences over mechanisms	40
3.2.10 Timing of the publication of exam scores	41
4. Two-sided matching	43
4.1 Static DA	44
4.2 Dynamic DA	46
4.3 Decentralized markets	49
4.4. Unraveling of markets	51
5. Conclusions	55
References	56

1. Introduction

For a long time, economists have focused on markets where prices coordinate demand and supply. However, there are many markets where prices do not determine who receives what. Examples include entry-level labor markets, school choice, university admission, social housing allocation, and kidney exchange. Over the past decades, the study of matching markets has become an active area of research. In all of these markets, agents have preferences over other agents or over objects they are matched to. For instance, workers have preferences over firms, and students have preferences over universities. The market designer collects these preferences from the agents and uses a mechanism to determine the matching which satisfies the designer's objectives like efficiency or fairness.

In the past decades, economists have been involved in re-designing centralized matching markets, canonical examples being the National Resident Matching Program for young doctors in the US (Roth and Peranson 1999) and school choice in Boston (Abdulkadiroğlu et al. 2005). A growing interest in the topic as well as novel questions arising when analyzing existing matching procedures have fueled rapid progress of research on the topic, and a considerable fraction of this research employs experiments.

This survey provides a comprehensive overview of the experimental literature on matching markets, complementing the survey article in the recently published handbook article (Roth 2015).² Almost all of the experiments are lab experiments. Field experiments on matching are faced with the difficulty that the preferences of participants are not known, but we report on two papers that find a way around this limitation. The goal of the survey is not only to summarize the main experimental findings, but also to identify what appear to be robust results across studies. To do so, we provide statistics across studies if possible, and also compare the results of related studies. Finally, by grouping the articles into a set of topics, we structure the current state of research.

Experiments play an important role in the study of matching markets. One important advantage of experiments is that preferences can be induced by the experimenter, and are therefore fully controlled for. For example, subjects' preferences are often induced by assigning different monetary payoffs for being matched to the different schools. Furthermore, experiments can create counterfactual situations and serve as testbeds for new mechanisms. A frequent concern regarding existing mechanisms is their incentive compatibility: if participants do not have a dominant strategy to

² In this survey we only touch briefly on the topics of unraveling and course allocation which are covered in depth by Roth (2015).

reveal their preferences truthfully, the game induced by the mechanism can become strategically complex. Thus, it can be unclear to what degree participants understand the incentive properties. The experimental revelation games induced by matching mechanisms shed light on the effectiveness of implementation in dominant strategies versus Nash equilibrium implementation. If the properties of the outcomes of a mechanism are analyzed under the assumption that participants state their true preferences or that they play the equilibrium strategy, this can lead to wrong conclusions regarding the desirability of the mechanism. For example, the efficiency of an allocation has to be calculated based on the true preferences of participants, which are hard to know from the data. Thus, experiments are a handy tool for the comparison of mechanisms, since they allow for testing whether subjects understand the incentive properties of the mechanisms and for comparing allocations based on the true preferences of participants.

Moreover, experiments enable the researchers to identify the factors that influence the agents' strategies, such as the information available to them regarding the preferences of other market participants, the size of the market, and so on. Also, experiments have permitted researchers to identify the causes of market failure, which is often impossible with observational data alone. Finally, student subjects are similar to the target population, especially in college admissions experiments. What in many studies simply serves as a convenience sample is the relevant sample for many matching experiments, which increases their external validity.

Market experiments are complex and have many degrees of freedom regarding the design. For example, matching markets are characterized by the preferences of participants, such as their degree of correlation, but also by the size of the market, by whether they are two- or one-sided, the amount of information provided about own and others' preferences, and the capacities of schools. For this reason and because the literature is still relatively young, there are fewer replications than in other areas of experimental economics. Nevertheless, we believe that there is a lot to be learned from relating the existing work to each other. Thus, we compare experiments that share a number of similarities even if they differ with respect to some features of the markets. Overall, we find a great level of consistency of the findings with clear patterns of behavior emerging. At the end of each section, we provide a short summary of the main findings.

Most of the results reported on in the papers that we review concern individual behavior – the input of subjects into the mechanisms. They consider whether participants report truthfully in the strategy-proof mechanisms and manipulate optimally in the ones where manipulations are part of an equilibrium. The rates of equilibrium reporting often have a direct effect on the properties of the resulting allocation; however, different subjects might have a different power to influence the allocation by submitting out-of-equilibrium preference reports to the designer. Moreover, subjects often have only a weakly dominant strategy of reporting truthfully in strategy-proof mechanisms, and thus not every deviation from truthful reporting influences the resulting allocation. For this reason,

some papers (typically papers that compare allocations reached by different mechanisms) look at stability and a measure of efficiency. Some papers emphasize efficiency, others stability, depending on the main interest and the mechanisms studied. For instance, if allocations reached under the student-proposing deferred acceptance (DA) mechanism are analyzed, the emphasis is on stability, as DA is predicted to produce stable allocations which do not have to be efficient. In the case of the top trading cycles mechanism (TTC), the emphasis is on efficiency, as TTC is predicted to produce Pareto-efficient allocations that are not necessarily fair.

The next section presents the basic model of a matching market as well as desirable properties of mechanisms and matching outcomes, namely stability, (Pareto) efficiency, and strategy-proofness. We introduce the five most common mechanisms, that is, the student-proposing deferred acceptance (DA), the school-proposing deferred acceptance (School-DA), the Boston or immediate acceptance (BOS), the top trading cycles (TTC), and the serial dictatorship (SD) mechanism. In section 3, we present the experiments on one-sided matching markets, i.e., matching problems where only one side of the market has preferences over the other side of the market, such as house allocation and school choice problems. We first discuss a number of experiments that compare mechanisms such as BOS, DA, and TTC with respect to the rate of truth-telling, efficiency, and the stability of the outcome. Furthermore, it is studied whether subjects learn to play the mechanisms and whether advice and coaching can be effective. This is followed by sections on the role of biases for reporting strategies, dynamic mechanisms, preferences over mechanisms, constraints regarding the length of the rank-order lists that can be submitted, and affirmative action. In section 4, we report on experiments about two-sided markets, i.e., matching markets where both sides of the markets have preferences over the matching partners, namely the marriage market and the college admissions model. In this section the main focus is on the role of the market rules for the stability of the outcome, for the distribution of payoffs between two sides of the market, and for market unraveling. In the last section, we briefly conclude and address what we perceive to be open questions and possible avenues for future research.

2. Basic concepts of matching theory

In this section, we provide a brief introduction to the central concepts and results of matching theory that are necessary to understand the experiments. For a thorough and detailed introduction of the theory, we refer the reader to the classic textbook by Roth and Sotomayor (1990).³

We use the set-up of students and schools to introduce the basic theoretical concepts and the matching mechanisms, since the majority of experiments have been conducted with this set-up. However, with a simple re-labeling of the agents, the model also applies to workers and firms if firms

³ The textbook by Guillaume Haeringer (2017) also offers a useful introduction, and it covers additional topics such as school choice and probabilistic assignments.

hire more than one worker, to students and universities, tenants and houses, and so on. Furthermore, the model can be modified in a number of ways. For example, if all schools have a capacity of one seat, the model captures a one-to-one matching problem and is therefore identical to the famous marriage problem.⁴ Here, we consider the set-up where one side of the market can accommodate more than one agent of the opposite side. This is called many-to-one matching.⁵ We will call the two sets of agents ‘students and schools’. The students are denoted by i , and the schools are denoted by s . Each student i wants to find a seat at a school s . Thus, each student has a strict ordinal preference over the set of schools (which might include the option of being unmatched). Each school has a strict ordinal preference over the set of students (which includes only the list of acceptable students⁶) and wants to accept at least one student. Each school has a maximum quota of students it can accept, q_s . A matching is a mapping that assigns each student i to a school s or leaves her unmatched, and it maps school s to student i if and only if student i is mapped to school s . The interpretation is that student i is only matched to school s if she chooses s and is chosen by s .⁷ The total number of students mapped to school s is no higher than q_s .

The matching game proceeds as follows: The designer asks all schools and students to report a rank-order list over the possible matching partners (i.e., to submit their ordinal preferences). The matching mechanism uses these rank-order lists to produce a matching and the agents are informed about the outcome.

Before turning to the matching mechanisms, we introduce some important properties of matching outcomes. A matching is *stable* (i) if every agent prefers the assigned matching partner to remaining unmatched, i.e., the student is matched to a school that she prefers to being unmatched, and the school is matched only to acceptable students, and (ii) if there is no school–student pair such that each prefers one another to their respective match. Stability is important because it precludes situations where students and schools would like to avoid being matched through the clearinghouse. A matching is *Pareto efficient* if there is no other matching which assigns every agent (student or school) a weakly better match and at least one agent a strictly better match. Furthermore, we call a matching *efficient* if the sum of the payoffs of the agents is highest among all matchings. Efficiency presupposes that the ordinal preferences of agents correspond to some cardinal payoff values.

⁴ In a marriage problem, there are two sets of agents, men and women. Each woman has strict ordinal preferences over the set of men (which might include the option of being unmatched). Each man has a strict ordinal preference over the set of women (which might include the option of being unmatched). Each man can be matched to at most one woman, and each woman can be matched to at most one man.

⁵ The set-up in which agents on each side can accommodate more than one agent of the opposite side is called many-to-many matching. We do not introduce it in this survey because we are not aware of any experiments using this set-up.

⁶ Acceptable students and acceptable schools refer only to the students and schools that lead to higher utility for schools/students than the outside option. Strictly speaking, it might be the case that a school prefers to have an empty seat to admitting a student. In this case the student is not acceptable.

⁷ This interpretation is true only for individually rational matchings, when no student is pushed to attend a school which is worse than her outside option. In what follows only individually rational mechanisms are considered.

Moreover, we can characterize some matching outcomes as assortative. If all students rank schools in the same way, i.e., they agree on the best school, the second-best school, and so on up to the worst school, and if at the same time all schools rank students in the same way, then there exists an assortative matching outcome. A matching outcome is *assortative* if the highest-ranked q_1 students are matched to the most-preferred school s_1 , the next highest-ranked q_2 students are matched to the second most-preferred school s_2 and so on. In a one-to-one matching problem, the assortative matching implies that for all agents that can be matched, the k -th student according to the ranking of schools is matched to the k -th best school according to the preferences of the students.

The mechanism designer is concerned not only with the properties of the allocation but also with the incentive properties of the mechanism. How complicated is it for agents to optimally submit their rank-order lists to the designer? One of the most desirable incentive properties is strategy-proofness. The mechanism is strategy-proof if truthful preference revelation is a (weakly) dominant strategy for agents. Thus, if the mechanism is strategy-proof, an optimal application strategy is straightforward for the agents. They should report their true preferences to the designer in the form of a rank-order list, which ensures them the best possible outcome (relative to alternative reports).

The distinction between one-sided and two-sided matching markets is essential. In a two-sided market (for which we just introduced the notation above), both students and schools strategically submit their preferences to the designer. This model is often appropriate for labor markets, marriage markets, and college admissions. When studying two-sided markets, the incentives and welfare of both sides are considered. The notion of stability is defined for such two-sided matching markets.

In one-sided markets, only one side of the market is strategic and has preferences over the agents or objects of the other side. In the context of students and schools, this means that only students have preferences over schools and act strategically, while schools are not strategic in their choice of students. This model is appropriate for the allocation of on-campus housing or for seats in public schools, for example. Houses and public schools do not have preferences over tenants or students, but instead have *priorities* which determine the rankings of agents. Unlike preferences, priorities are determined in advance by law or the mechanism designer and are not strategically reported to the mechanism. Examples are priorities for students who live in the neighborhood of a school or who have a sibling at the school. A one-sided matching problem is called housing market if it is a one-to-one matching problem and school choice if it is many-to-one. Thus, when studying one-sided markets, only the incentives and welfare of one side of the market, namely the strategic agents, are considered. The concept of stability of the two-sided matching model translates into the concept of *elimination of justified envy* in the one-sided matching model. The envy of student i toward student j regarding school k is *justified* if student j is assigned to school k , student i ranks school k higher than her assigned school, and student i has a higher priority than student j at school k .

In the following, we describe the five most important matching mechanisms in the literature. Only one of the five mechanisms presented possesses all three desirable properties: strategy-proofness, stability, and Pareto efficiency. It is the Serial Dictatorship mechanism. It is presented last, since it can be used only in markets where all agents on one side of the market have the same ordinal ranking of the agents on the other side. In the context of school choice, it implies that all schools rank all students in the same manner. The other four mechanisms can be used under any preferences, but do not possess all three desirable properties. In fact, it has been shown that such a mechanism does not exist (Alcalde and Barberà 1994). The first three mechanisms described, DA, School-DA, and Boston, are the most frequently used procedures in centralized labor markets, school choice, and college admissions. The fourth mechanism, TTC, is applied less frequently despite its efficiency.

For the description of the five mechanisms, we use the context of a school choice problem, that is, a many-to-one one-sided matching problem. The reason is that the bulk of experimental papers are based on this model by assuming that schools do not act strategically and that the mechanism uses strict priority rankings of schools that are exogenously given. Students report their preferences over schools in the form of rank-order lists. The mechanism also receives the rank-order lists of schools (priorities), and the capacities of schools (the maximum number of students each school can admit). The preferences and priorities are strict, and if not, the ties are broken arbitrarily. According to these reports and the priorities, the mechanism produces a matching.

Student-Proposing Deferred Acceptance Mechanism (DA)

Step 1: Each student applies to the school that is ranked first in her preference list. Each school admits acceptable students up to its capacity, following its priority order. The remaining students are rejected.

Step k , $k \geq 2$: Each student rejected in the previous step applies to the most-preferred acceptable school among those she has not yet applied to. Each school receiving applications considers the set of students it admitted in the previous step together with the set of new acceptable applicants. From this set, the school admits students up to its capacity, following its priority order. The remaining students are rejected.

End: The algorithm stops when no student is rejected, or all schools have filled their capacity. Any remaining students are unassigned.

Note that the allocation is temporary at each step until the last step.

Properties of DA

The student-proposing DA is strategy-proof in the one-sided matching model.⁸ Moreover, DA eliminates justified envy for the students, and the outcome Pareto dominates all other envy-free outcomes from the perspective of the students. However, DA is not Pareto efficient in the one-sided matching set-up.

School-Proposing Deferred Acceptance Mechanism (school-DA)

The mechanism receives the rank-order lists of students (preferences), the rank-order lists of schools (priorities), and the capacities of schools (the maximum number of students each school can admit).

Step 1: Each school offers seats to students with the highest priority up to its capacity. Each student accepts the best acceptable offer she has received, according to her preference list. The other schools are rejected.

Step k , $k \geq 2$: Each school rejected in the previous step makes offers to the students with the highest priority among those that have not rejected an offer from the school yet such that the number of accepted offers from previous steps and the number of new offers do not exceed capacity. Each student receiving offers considers the school she accepted in the previous step together with the set of new offers from schools. From this set, the student accepts the school that is highest on her preference list. All other schools are rejected.

End: The algorithm stops when no school is rejected, or all students have found a seat. Any remaining students are unassigned.

Note that the allocation is temporary at each step until the last step.

Properties of School-DA

The school-proposing DA, or short School-DA, is not strategy-proof. The School-DA eliminates justified envy by students. However, the School-DA is not Pareto efficient.

Boston mechanism (BOS)

⁸ Note, however, that by the Impossibility theorem (Roth 1982), there is no stable mechanism that is strategy proof for both sides of the market. Thus, DA is strategy-proof only in the one-sided matching set-up where the strategic agents are proposing.)

Step 1: Each student applies to the school that is ranked first in her preference list. Each school admits acceptable students up to its capacity, following its priority order. These assignments are final. The remaining students are rejected.

Step k , $k \geq 2$: Each student who was rejected in the previous step applies to the most-preferred acceptable school among the schools to which the student has not yet applied. Each school admits acceptable students up to its remaining capacity, following its priority order. These assignments are final. The remaining students are rejected.

End: The algorithm stops when no student is rejected, or all schools have filled the seats up to their capacity. All remaining students are unassigned.

Note that the allocation is final at each step of the mechanism.

Properties of BOS

BOS is not strategy-proof for the students. In Nash equilibrium, BOS eliminates justified envy by the students (Ergin and Sönmez 2006). However, the equilibrium requires strategic play by the students. If all students report truthfully, BOS produces a Pareto efficient, but possibly not a stable allocation (that is, it might not eliminate justified envy).

Top trading cycles (TTC)

Step 1: For each student, we point from this student to the school that is the most preferred by that student. If there is no such school, she points at herself, since she prefers to remain unmatched.

For each school, we point from the school to the student who has the highest priority for the school.

There must be at least one cycle of students and schools pointing at each other or a student pointing to herself.⁹ Every student in a cycle is assigned to the school she is pointing to or to herself if pointing to herself,¹⁰ and is removed. The remaining capacity of each school in the cycle is reduced by one and if it reaches zero, the school is removed.

⁹ The existence of at least one cycle is guaranteed since there are finitely many agents. A cycle can consist of one student pointing to herself, one student and one school pointing to each other, 4, 6, or another equal number of students and schools pointing at each other such that the first member of a cycle points to the second, the second to the third ..., and the last points to the first.

¹⁰ Being assigned to herself means that the student will remain unmatched.

Step k , $k \geq 2$: For each student, we point from the student to the acceptable school that is the most preferred by that student among the schools that are still present. If there is no such school, she points at herself.

For each school, we point from the school to the student who has the highest priority for the school among the acceptable students who are still present.

There must be at least one cycle. Every student in a cycle is assigned to the school she is pointing to or to herself and is removed. The remaining capacity of each school in the cycle is reduced by one and if it reaches zero, the school is removed.

End: The algorithm stops when all students or all schools have been assigned. Any remaining students are assigned to themselves.

Note that the allocation is final at each step of the mechanism.

Properties of TTC

TTC is strategy-proof for the students. TTC produces a Pareto efficient allocation, but it does not eliminate justified envy for the students (Abdulkadiroğlu and Sönmez, 2003).

The next mechanism relies on all schools ranking the students in the same way.

Serial Dictatorship mechanism (SD)

Step 1: The student at the top of the schools' priorities is assigned to the school at the top of her preference list. The student is deleted from the priority list and the capacity of the respective school is reduced by one. If capacity reaches zero, the school is removed from the all preference lists.

Step k , $k \geq 2$: The highest remaining acceptable student on the priority list of the schools is assigned to the acceptable school at the top of her preference list.

End: The procedure terminates when the list of priorities is exhausted, or all schools have capacity zero.

Note that the allocation is final at each step of the mechanism.

Properties of SD

SD is strategy-proof for the students, eliminates justified envy, and leads to the Pareto efficient allocation for students.

The experimental work that we review often studies the mechanisms above or modifications of these mechanisms. Some articles investigate other mechanisms that we then explain in the respective paragraphs.

3. One-sided matching

3.1 House allocation

The house allocation problem is a one-sided one-to-one matching problem. Houses do not have preferences over owners, and are thus not strategic players. Moreover, each house can only have one owner. Houses have only one copy, and agents have preferences over the houses. The most important concern in the house allocation literature is efficiency. Moreover, since one of the main motivations of the literature is the allocation of on-campus housing, an additional feature of the problem is that some agents can have prior claims over houses. Thus, the incentives of existing tenants to participate in the redistribution of houses are explored. In addition to DA, TTC, and SD, the so-called probabilistic serial mechanism (PS) was introduced by Bogomolnaia and Moulin (2001) for the house allocation problem.¹¹ PS is motivated by the fact that Random Serial Dictatorship (RSD)¹² is not ordinally efficient. This means that there exists an alternative way to randomly allocate the houses which would be preferred by all participants from an ex-ante perspective. Ex-ante means that participants prefer one set of probabilities to another when the probabilities of getting each house are determined by the mechanism, but the realization of these probabilities is still unknown. PS addresses this issue and produces an ordinally efficient probabilistic allocation if all agents report truthfully. The following theoretical predictions for the housing market are studied with the help of experiments:

- (1) TTC, DA, and RSD are strategy-proof, PS is not.
- (2) TTC and RSD are Pareto efficient for all participating agents while DA is not.

¹¹ PS is a procedure for the fair random assignment of houses. It collects rank-order lists of agents, and based on the submitted lists determines a probabilistic allocation. Thus, agents are assigned probabilities of receiving each house. This assignment is both envy-free (a random assignment is envy-free if every agent weakly prefers his own assignment to that of any other agent in terms of first-order stochastic dominance with respect to his reported ordinal preferences) and Pareto efficient. In the experiments, one of the realizations of the probabilistic matrix was drawn to determine the payoffs of participants. For details on PS see Bogomolnaia and Moulin (2001).

¹² RSD is a version of SD where the priorities are not known in advance, but are randomly determined by the designer in the allocation process. Thus, participants do not know ex ante who is at the top of the priority lists, who is second, and so on.

- (3) Participation in TTC is individually rational for all agents while this is not the case in RSD.
- (4) Although RSD is ex-post efficient, it is not ordinally efficient while PS is.

Note that due to the random priorities in RSD, existing tenants may be worse off when participating than when simply keeping their current endowment. This issue is resolved by TTC where existing tenants have the highest priority for their current house and are therefore guaranteed to get their current or a better house if they participate in the procedure.

The first experimental paper that studies the house allocation problem with existing tenants is Chen and Sönmez (2002). The authors compare two mechanisms in the lab: one is widely used in practice for the allocation of on-campus housing – random serial dictatorship with squatting rights (RSD) – and the other is TTC which was proposed by Abdulkadiroğlu and Sönmez (1999) and has never been used in practice. In RSD with squatting rights, each existing tenant can either keep the current house or participate in the procedure. The procedure implements RSD according to which each participant receives her top choice among the remaining houses in the random priority order.

The two mechanisms are compared between subjects in two environments, one with random preferences and one with correlated ('designed') preferences.¹³ Each environment has eight existing tenants and four newcomers. Subjects know only their preferences and not the preferences of the others. Subjects in the role of the existing tenants first had to decide whether they wanted to participate in the allocation mechanism or just keep their initial house. All newcomers and the existing tenants who decided to participate in the allocation procedure submitted rank-order lists to the mechanism, and the new allocation was determined. As predicted by theory, in both environments the participation of existing tenants is higher under TTC than under RSD, with 79% and 47% respectively. Truth-telling is measured as submitting full lists of the true preference ranking. The truth-telling rates given participation are higher under RSD, but the difference is not significant. Approximately 73.5% and 70.9% of subjects reported truthfully under RSD and TTC respectively.¹⁴ Finally, in all markets that were tested, the efficiency of TTC was higher than the efficiency of RSD, also in a large market environment which was included in the study as a robustness check. Thus, the study demonstrates that despite a matching mechanism being used in practice for a long time, the efficiency of the allocation can be improved by the introduction of a mechanism which incentivizes existing tenants to participate in the yearly re-allocation procedure.

¹³ The term "designed" environment was used in many subsequent studies to refer to preferences which are constructed in an attempt to create a correlation between the preferences of the experimental subjects.

¹⁴ The Serial Dictatorship (SD) mechanism was previously tested in the lab in a different context. SD works as RSD, with the difference that the priority order is predetermined, not random. Olson and Porter (1994) compare SD to different auction formats for the problem of allocating the objects to agents. Monetary transfers are allowed, and the main goal of the designer is to maximize the welfare. SD provided the lowest levels of welfare, as it is an ordinal mechanism; thus, the allocation does not account for potential differences in the intensity of cardinal preferences. As for individual behavior, 80% of subjects in 80% of the problems followed the dominant strategy of truthful reporting. In the other 20% of problems, the percent of truthful reporting was lower.

Due to the concern that the incomplete information environment where subjects do not know the preferences of others may be unrealistic, and to test the robustness of the findings, the authors also ran both mechanisms in a complete information environment and in a simpler environment with only three agents. The results of these experiments are presented in Chen and Sönmez (2004) and are qualitatively the same.

In a closely related paper, Guillen and Kesten (2012) test TTC versus a mechanism that was used for the allocation of on-campus housing at MIT. First, they show theoretically that the mechanism used at MIT is a natural adaptation of the DA mechanism for the house allocation problem. In this mechanism the priorities of agents are common for all houses and are determined randomly, except for the existing tenants who have the highest priority for the house they live in. Thus, both mechanisms that were tested are strategy-proof and incentivize existing tenants to take part in the allocation procedure. However, TTC has an edge over DA regarding efficiency. The authors used the environment employed by Chen and Sönmez (2002) with eight existing tenants and four newcomers. The results of the experiments differ from the results of Chen and Sönmez (2002) mostly with respect to the participation rate of existing tenants under TTC. Only 47.5% of existing tenants participated in the allocation compared to 79% in the study by Chen and Sönmez (2002), in contrast to the prediction of full participation. Under DA, this rate is significantly higher and equals 77.5%. This leads to an overall higher efficiency of DA, but the significance of the difference disappears once recombinant estimation is used (Mullin and Reiley 2006).¹⁵ As for truth-telling, the rates were higher for DA compared to TTC (80% versus 69%), but the difference is not statistically significant.

Hugh-Jones et al. (2014) test the Random Serial Dictatorship mechanism (RSD) against the Probabilistic Serial mechanism (PS). Their set-up differs from the other papers discussed in this section since there are no existing tenants who have priority: so-called squatting rights. The experimental results show that misreporting occurs in PS even in environments where it is not predicted. Some misreporting is also observed under the strategy-proof RSD. Depending on the environment, the outcome is significantly more efficient under PS than under RSD or no difference between the mechanisms is found.

Schmelzer (2018) tests TTC with random endowments¹⁶ against a simultaneous and a sequential version of RSD in the context of house allocation without existing tenants. In theory, both mechanisms are strategy-proof and ex-post Pareto efficient, and are thus equivalent (Abdulkadiroğlu and Sönmez 1999). In the sequential version of RSD, the agents do not submit rank-order lists to the mechanism

¹⁵ Recombinant estimation is a technique from statistics that allows for a robust estimation of group-level outcomes in one-shot games. Assuming that the strategies of subjects are independent of the identity of the partners (due to the one-shot nature of the interaction), one can recombine players from different sessions, and calculate an allocation for each recombination. The recombination technique leads to a distribution of possible outcomes and thus to a more robust estimation of treatment differences.

¹⁶ Random endowments imply that each agent has the highest priority at one of the houses, and this priority is randomly determined.

but choose a house among the remaining houses in the given priority order. In the experiment, subjects' choices were elicited using the strategy method: each subject had to choose one house from different menus of houses that could be realized. The results of the experiment by Schmelzer (2018) contradict the theoretical prediction regarding the equivalence of TTC and RSD: the Pareto efficient allocations are reached significantly more often under both versions of RSD than under TTC. The subjects of the experiment play additional games which allow the author to identify the subjects with a higher ability for contingent reasoning. The subjects with extremely low and extremely high levels of contingent reasoning are more likely to report truthfully.

To sum up, the literature on house allocation mechanisms is inconclusive when it comes to the rate of participation of existing tenants in TTC. As for the rates of truthful reporting, even in the simplest mechanism, RSD, some subjects do not follow the dominant strategy of truthful reporting, though the rates are typically higher than in TTC.

It is worth mentioning that the house allocation problem is closely related to the course allocation problem where bundles of courses are allocated to students. We refrain from surveying the work on course allocation and refer the reader to the comprehensive overview by Roth (2015).¹⁷

3.2 School choice

By far the largest part of the experimental literature on one-sided markets deals with the school choice model. Its central theoretical results have been studied experimentally. In school choice problems, the schools are assumed to be non-strategic, and welfare considerations only apply to the students. The following predictions hold:

- (1) The DA and the TTC mechanism are strategy-proof, while BOS and school-DA are not.¹⁸
- (2) The DA and School-DA mechanisms eliminate justified envy, while TTC and BOS with truthful preference revelation do not.
- (3) TTC is Pareto efficient as is BOS with truthful preference revelation, but DA and School-DA are not.

¹⁷ Another related literature studies object allocation without money in the context of booking systems, e.g., for appointments at public offices. Insights from matching can help in fighting undesirable properties of these systems. In a recent experiment by Hakimov et al. (2018), the authors show that a first-come first-served system for the booking of appointment slots can suffer from the intervention of scalpers. These firms often have a technological advantage relative to the seekers of slots and make profits by booking and then selling the slots to appointment seekers. The authors introduce an alternative booking system in which firms cannot make profits.

¹⁸ Note that in this section we consider only the students as agents. We therefore say that DA is strategy-proof. It is not strategy-proof for schools, but they are not considered as players in the one-sided matching set-up. School-DA is not strategy-proof for either for the students or the schools.

- (4) Under BOS the Nash equilibrium outcomes with complete information eliminates justified envy but is not Pareto efficient.

3.2.1 Comparison of mechanisms

The literature starts with the seminal paper of Chen and Sönmez (2006). Its experimental design is largely based on the design of their 2002 paper on house allocation. The experiment studies preference reporting under three alternative mechanisms, namely BOS, DA, and TTC, and compares the outcomes of these mechanisms from the perspective of efficiency and stability. BOS is used as a natural baseline, since it was actually used for school choice in Boston and New York. DA and TTC are the two leading mechanisms suggested by economists. The experiment was run in class and was paper-based. This allowed the authors to run fairly large markets, namely 36 participants competing for 36 seats in seven schools. Students came to class and received a set of instructions.

In order to determine the preference profiles, the authors used two alternative procedures. The ‘designed’ procedure was aimed at generating realistic preferences. In order to do so, each student’s ranking of the schools was generated by a utility function which depended on the school’s quality, proximity, and a random factor. The utility derived from the quality of the school was common for all students. To determine the utility from proximity, the authors first determined a district school for each student. Each student received utility from proximity for this school. In the second environment the preferences were randomly determined, and this environment was used as a robustness check. Based on the resulting rankings, fixed payoffs were assigned to each rank, such that there was no difference in the cardinality of preferences. As for priorities of schools, in both environments the highest priority was given to district students, and for all other students the priorities were determined by a lottery. In order to avoid trivial observations, the induced preferences of students were corrected in such a way as to ensure that the district school was never a student’s top choice. This design of the environment became common in many subsequent school choice experiments.

With the three matching mechanisms and two preference environments, the experiment by Chen and Sönmez (2006) follows a 3x2 design. The six treatments were run between subjects, meaning each subject participated in only one mechanism in one environment. The experiment involved incomplete information in that subjects in the experiment knew which school was their district school and could observe their own preferences, but they had no information about the preferences of the other students. The experiments were one-shot, meaning each subject played the game just once.

The main result regarding individual behavior is in line with the theoretical predictions: in both the designed and random environments, the proportion of truthful preference revelation under BOS

was significantly lower than the proportion of truthful preference revelation under either DA or TTC.¹⁹ This is one of the main insights of the paper, which was replicated in all subsequent studies. Additionally, it turned out that despite the strategy-proofness of both mechanisms, the proportion of truthful reporting was significantly higher in DA than in TTC, especially in the designed environment. The finding is surprising, since both mechanisms are strategy-proof. However, subsequent papers show that it is not robust to other environments and settings. For instance, in the baseline treatments of Klijn et al. (2010), which is a replication of Chen and Sönmez (2006), the proportion of truthful preference revelation under TTC is higher than under DA. The main result of Chen and Sönmez (2006), however, that the proportion of truthful preference revelation under DA is higher than under BOS, was replicated in the designed and random environments of the baseline treatments of Klijn et al. (2010) and in the designed environment of the baseline treatment of Koutout et al. (2018).

The authors also identified a common tendency in the manipulated reports, which they call “district school bias.” It refers to the finding that the district school (or safe school) is ranked higher in the reported list than in the true preferences. In BOS, 15.5% and 8% of subjects displayed the district school bias in the designed and random environments respectively. As for the analysis of allocations, the authors report a higher efficiency of allocations in DA mostly due to higher rates of truthful reporting. However, a subsequent analysis of their data by Calsamiglia et al. (2011) show that all mechanisms lead to similar efficiency levels when using recombinations from the robust recombination technique (Mullin and Reiley 2006), described in footnote 15. Note that the experimental design with an assignment of identical cardinal utilities to the first rank, the second rank, etc., preclude efficiency gains in BOS due to the possibility to express preference intensities.

One of the design features of the experiment by Chen and Sönmez (2006) that may be controversial is the decision to provide no information about the preferences of other students. While the participants of school choice procedures most likely do not know the exact preferences of their peers, it seems unlikely that they do not know anything about others’ preferences. Often parents know which schools are more popular than others. Pais and Pintér (2008) investigate the effect of providing richer information for participants in a set-up that is very similar to the one in Chen and Sönmez (2006). However, they implemented a smaller market, with five teachers competing for seats in three schools. In their 3x4 design, three mechanisms—namely BOS, DA, and TTC—were run under four different information conditions between subjects.²⁰

¹⁹ Truthful preference revelation means that a full list is submitted which corresponds to the true ranking for BOS. For TTC and DA truthful preference revelation in the study required only reported choices from the first to the district school to be truthful for DA and TTC. This is because the district school is a guaranteed school, and all choices below it are irrelevant for the allocation under DA and TTC.

²⁰ In zero information, participants only know their own preferences and the capacities of schools. In low information, subjects also know each school’s favorite student. In partial information, they additionally know the favorite students at each school up to capacity. In full information, they know the full priority lists of schools and the preferences of all students.

Pais and Pintér (2008) replicated the result of Chen and Sönmez (2006) that truthful preference revelation is higher in DA and TTC than in BOS in the same information condition as Chen and Sönmez (2006), namely their low information treatment, as well as in two other information environments with more information provided about the preferences of other participants. In the zero-information environment when students knew only their own preferences and did not know even the priorities of schools, there was no difference between the mechanisms with respect to truth-telling. The main take-away from the experiment is that subjects reacted to the additional information about the preferences of others and the schools' priorities by misrepresenting their preferences more frequently in all mechanisms. While the effect was strongest in BOS, it was also significant in DA and TTC. The findings for DA and TTC are not predicted by the theory and can be interpreted as suggesting that truthful revelation in the incomplete information environment represents an upper bound. The truth-telling rate in DA was just 66.7% (both with full and partial information), while it was 75% in TTC (partial information). Note that unlike in Chen and Sönmez (2006), TTC had, on average, a 12% higher rate of truthful reporting than DA in all treatments, with the difference being significant in three out of four environments.

The provision of information was also detrimental for the efficiency in BOS and DA, but not in TTC. Regarding comparisons across mechanisms, TTC led to a significantly higher efficiency of allocations than DA and BOS in the partial and full information conditions. Moreover, the provision of information did not have an effect on the stability of allocations. As predicted, DA led to the highest rates of stable allocations, but the difference is only significant in three out of four information conditions relative to TTC and two out of four information conditions relative to BOS. Despite its worse performance under complete information, DA still outperforms BOS in Pais and Pintér (2008) at least weakly from the perspective of truth-telling and stability.

Many Chinese provinces use a hybrid mechanism between DA and BOS – the so-called parallel mechanism (Chen and Kesten 2017) – which has been studied experimentally by Chen and Kesten (2016). The parallel mechanism uses choice-bands that determine the number of steps for which the allocations are tentative. Within a choice-band all assignments are tentative while they are final once a student is either assigned to a school or has been rejected by all his choices in this choice-band. Thus, both BOS and DA are nested in the parallel mechanism with choice-band sizes of 1 and infinity, respectively. The experiments are designed to test the theoretical predictions of Chen and Kesten (2017) with respect to the effects of the size of the choice-band on the amount of truth-telling, efficiency, and stability.

In a complete information environment with markets of either four or six schools, it is found that the parallel mechanism with two schools per choice-band induces truth-telling rates that are between those of BOS and DA, in line with the theory. With respect to efficiency, the results depend on the exact environments studied, and there is no clear ranking of the mechanisms, as predicted.

Finally, the observed stability of the matchings again supports the theory, with DA leading to more stable matchings than the parallel mechanism and BOS. The parallel mechanism induces (weakly) more stable matchings than BOS, depending on the markets considered. Thus, one can interpret the findings of the intermediate performance of the parallel mechanism as a successful robustness check for the superior performance of DA relative to BOS with respect to truthful reporting and stability. In a closely related paper, Chen, Jiang, and Kesten (2017) replicate the finding of Chen and Kesten (2016) for the set-up of six colleges with identical priorities over students, a characteristic of Chinese college admissions that are solely based on the centralized entrance exam. Moreover, the authors show that their theoretical and experimental findings are in line with data from the Sichuan province of China where BOS was changed to the parallel mechanism: students started to list more colleges, and the prestigious colleges were ranked as a top choice more often.

The comparison of DA and BOS holds up in larger markets where the number of schools is kept constant, but the number of students is 40 or 4,000 (Chen et al. 2017). In these experiments, some students are played by robots that use the strategies of real subjects who participated in previous sessions. More details about the experiments of Chen et al. are presented in the next subsection.

Dur et al. (2015) introduce another mechanism, the secure Boston mechanism (secure BOS), which can be understood as an intermediate mechanism between BOS and DA, just like the Chinese parallel mechanism. The authors note that BOS can be seen as a version of DA that is run on the modified priorities of schools according to the preference reports of students, such that the students who rank a school first move to the top of the priority list of that school. The secure BOS mechanism also modifies the original priorities of schools but keeps the most-preferred students of each school up to its capacity at the top of the priority list, independent of how students rank the school. The secure BOS runs DA on the modified priority lists. The secure BOS is not strategy-proof, but it is less manipulable than BOS. The intuition is that the students still have seats in their district schools guaranteed for them, even if they do not rank them first. The authors compare BOS and secure BOS in the lab, in a set-up where experimental subjects play against computerized players. Subjects knew the top choices submitted by the computer players. In line with the theoretical predictions, secure BOS led to fewer manipulations than BOS. However, the rates of truthful reporting were rather low with 28.4% and 17.1%, respectively. Secure BOS led to fewer instances of justified envy than BOS.

The experiments of this section demonstrate that the concerns regarding the manipulability and the inferior outcomes of BOS find support in the lab. TTC and DA appear to be superior mechanisms although the absolute levels of truth-telling and the decrease in truth-telling the more information is provided have raised concerns of whether the properties of DA and TTC are transparent enough for the participants. Evidence regarding this question comes from an experiment by Guillen and Hakimov (2017). The authors use TTC in a set-up where students play against computers in a matching market of four schools and four students. In every market three students are played by

computers, and one student is represented by an experimental subject. Each subject had to make two decisions. In the first decision subjects knew the preferences of the computer players and they were told that the computer players submitted their rankings truthfully. In the second decision, subjects were provided with different partial information on the possibly non-truthful strategies of computer players depending on the treatment.²¹ The two decision problems were presented on the same screen, so participants made both decisions simultaneously. While a great majority (85% of subjects) reported truthfully in the first situation, the rate of truth-telling was dramatically lower in the second decision. Only 31% of subjects were truthful in both decisions. Thus, the experiment demonstrates that subjects do not perceive truth-telling as a dominant strategy but are influenced by the behavior of others. This lends support to the concern that understanding the incentives of TTC and possibly other strategy-proof mechanisms is not straightforward for the participants.

While lower truth-telling rates under BOS relative to DA were found in all studies, it is possible that the strategies played in BOS are in line with the equilibrium prediction. This is crucial given the theoretical results of Abdulkadiroğlu et al. (2011), showing that the equilibrium of BOS can dominate DA from an ex-ante efficiency perspective, since it allows the students to express their cardinal utilities through the strategies played in BOS. Featherstone and Niederle (2016) run experiments comparing DA and BOS in two different environments. In the environment with correlated preferences where BOS has a unique equilibrium in non-truthful strategies, only 42.9% of the reports were consistent with the unique pure-strategy Bayes-Nash equilibrium in BOS. Moreover, 40% of reports were truthful under BOS, which is significantly lower than 80% of truthful reporting under DA. However, in the environment with uncorrelated preferences, BOS admits truth-telling by all students in the ordinal Bayes-Nash equilibrium. The results show that in this environment the truth-telling rate under BOS is 58%, which is not statistically different from 66% of truthful reports under DA, resulting in a higher efficiency of BOS. The authors interpret this result as a proof of concept that non-strategy-proof mechanisms with a truthful ordinal Bayes-Nash equilibrium might succeed in practice. As for DA, the truth-telling rates were 66% in the uncorrelated environments and 80% in the correlated environment.

Summing up this section, most of the studies find that the truth-telling rates under BOS are lower than under strategy-proof mechanisms. The comparison of truth-telling rates between DA and TTC is inconclusive. Moreover, most manipulations in BOS do not represent equilibrium play. DA mostly outperforms the other mechanisms in terms of stability while the comparison of mechanisms with respect to efficiency is inconclusive and depends on the environment. Despite the relative success of the strategy-proof mechanisms DA and TTC, the sensitivity of the rates of truthful reporting to the information provided and to the market environment raises concerns regarding the successful

²¹ Partial information on the strategies of the computer players included a statement that some players manipulate their reports in an unknown way in one treatment. In another treatment, the participants were given an exact report of one of the computer player's non-truthful choices.

implementation of these mechanisms in the field. One possible explanation of the relatively low truth-telling rates under the strategy-proof mechanisms DA and TTC is the absence of experience. The next section reports evidence on learning in BOS and DA.²²

3.2.2 Learning and effect of market size

This section focuses on the dynamics of the subjects' reports in experiments where they play DA or BOS repeatedly. Most of the papers mentioned in this section do not focus on learning, but they employ multiple rounds of matching markets such that the data can be used to study learning.

The baseline treatments of Ding and Schotter (2015) include repeated play of BOS and DA for 20 rounds. The market consisted of five students competing for three schools. Participants knew their own preferences and the priority schools of all students, but not the preferences of other students. The group was randomly re-matched and the ties in priorities were broken randomly in every round. The authors found a significant increase in truthful reporting in DA, from 64% to 77%, over 20 rounds. There was no significant change in the efficiency of the allocations, with a slight decrease of efficiency with experience. Note that the efficiency might be in conflict with stability, which is why an increase in truthful reporting does not necessarily translate into higher efficiency. There was no increase in truth-telling in BOS, as expected, since truth-telling is not a Nash equilibrium strategy. However, experience did not increase the efficiency of the allocation reached in BOS either.

In the baseline treatment of Chen and Kesten (2016), subjects played DA and BOS for 20 rounds, in either a four-school environment or a six-school environment. The experiments were run under complete information, where both the preferences of other players and the priorities of schools were known to participants. In the four-school environment, there was no significant learning under DA, but the truth-telling rates were relatively high at around 75% on average. As for the six-school environment, truth-telling under DA was lower at 47% on average, and there was a significant negative effect of experience on truth-telling. As for BOS, the truth-telling rates were lower than under DA in both environments (46% and 23% respectively in the four-school and the six-school environment), and experience had a slightly negative effect, that was significant only in the four-school environment.

In the baseline of Zhu (2015), subjects played DA for 15 rounds under complete information about preferences and priorities. The experiments were run in two environments. Each environment had three students and three schools with one seat each. In the first environment, there were no conflicts between top choices (uncorrelated preferences), and in the second environment preferences were correlated. Results show significant learning in both environments, with truthful reporting rates reaching around 75% in the final rounds of the experiment. Finally, in the baseline of Bó and Hakimov

²² To the best of our knowledge there are no studies in which subjects play TTC for multiple rounds.

(forthcoming) subjects played DA for 20 rounds. The preferences were generated anew every round, following a procedure inspired by the designed markets of Chen and Sönmez (2006). There were eight students and eight schools with one seat each in every round. The authors found a significant increase in truthful reporting when comparing the first 10 to the last 10 rounds of the experiments. Experience had a positive effect on truth-telling rates which increased from 38% in the first five rounds of the experiment to 56% in the last five rounds.

Figure 1 presents the average rates of truthful reporting in all studies using repeated DA, and the dynamics of truthful reporting by rounds. Summing up, the majority of studies find some evidence in favor of learning to report truthfully. There is an increase in truthful reporting in all but two studies, namely Chen and Kesten (2016) in the six-school environment and Chen et al. (2017) in environments with 40 human players. However, the levels of truth-telling vary between the studies.

One might conjecture that the longer the list to submit, that is, the more schools to choose from, the lower the truth-telling rates are. This conjecture is supported by Chen and Kesten (2016) when comparing their four-school environment to their six-school environment.²³ It is also in line with the levels of truthful behavior between the studies. For instance, all of the studies with high rates of truth-telling in Figure 1 have three or four schools to be ranked, while Bó and Hakimov (2017) with eight schools, Chen, Jiang, and Kesten (2017) with six schools, and Chen and Kesten (2016) with six schools display the lowest average truth-telling rates.

²³ This conjecture is also supported by Hakimov and Kesten (2018) who test TTC against another strategy-proof and Pareto efficient mechanism, equitable TTC. They use markets with three, four, and five schools. Truthful reporting is highest in the markets with three schools, and lowest in the market with five schools. There is no significant difference between rates of truthful reporting under the different mechanisms.




Figure 1. Dynamics of truthful reporting in DA experiments

Notes: Each line corresponds to a study. The legend first names the study, followed by the name of the treatment in case the DA was used in multiple treatments, followed by the number of schools that participants had to rank in that environment.

Figure 2 presents the regression of truth-telling on the number of schools that can be ranked for nine studies with at least 15 rounds of play. We chose 15 rounds, since this is the minimum length of matching experiments with repeated play, as can be taken from Figure 1. The coefficient for the length of the list is significant and negative. Nevertheless, due to many differences between the studies, this evidence is merely suggestive and might be worth testing systematically. Also, it is an open question as to why this relationship seems to hold, e.g., whether it is due to random choices by some subjects, implying that the longer the rank-order lists, the lower the probability of randomly picking the truthful strategy.

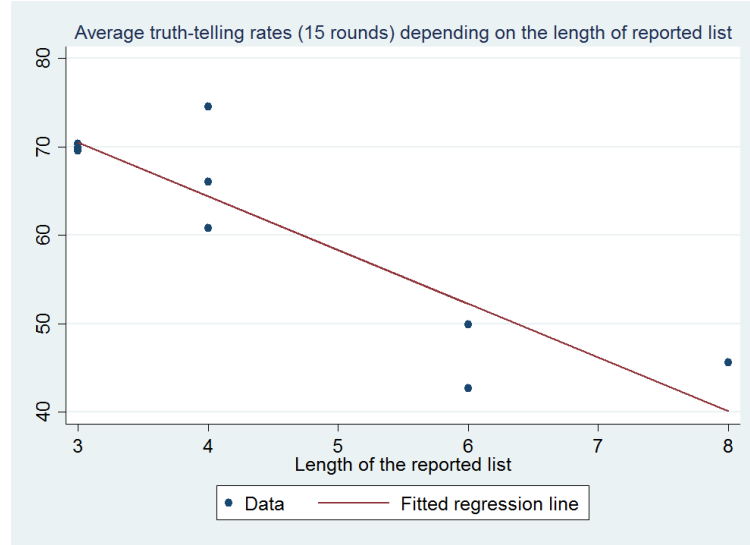


Figure 2: Truth-telling and number of schools that can be listed under DA

Notes: Each dot represents one study. The line displays the predicted rate of truth-telling from the linear regression of truth-telling rates on the length of the rank-order.

The effect of market size on behavior and on the properties of the allocation under DA and BOS are studied in the experiments of Chen et al. 2017. The authors keep the length of the rank-order list fixed (four schools) but increase the size of the match by increasing the number of students and the number of seats in each school. One environment replicates the four-school environment of Chen and Kesten (2016). The other two environments increase the number of students to 40 and 4,000, together with increasing the number of seats in each school to 10 or 1,000. Note that the number of students is increased by creating 10 or 1,000 students for each preference type of students in the four-school environment of Chen and Kesten (2016). To make the large-scale experiment possible, the authors run some treatments where students interact with robots. Robots either play the empirical strategies of other subjects or they report truthfully, depending on the treatment. The results show that in all environments the truth-telling rates under DA are higher than under BOS, while the proportion of students exhibiting justified envy is lower under DA than under BOS. No difference between mechanisms was found regarding efficiency. The theory predicts that the scale should not influence the subjects' strategies under BOS and DA. It is found that the increase in the scale from four to 40 students has a weakly significant and positive effect on truth-telling under DA and a significant negative effect on truth-telling under BOS. The increase from 40 to 4,000 students has a positive effect on truth-telling under DA and a negative effect on truth-telling under BOS, but these effects are not statistically significant. There is a small negative effect of the increase in market size from four to 40 on efficiency under both mechanisms, but no effect on stability. Finally, strategies of subjects are not significantly different if they play against human subjects or robots whose strategies are drawn from empirical human strategies, keeping the size of the market fixed at 40.

The results of Chen et al. 2017 together with the results summarized in Figure 2 imply that the overall effect of scale is unclear: while an increase in the market size, keeping the number of schools fixed, if at all has a positive effect on truth-telling under DA, an increase in the length of the rank-order list seems to have the opposite effect. This raises the question of how these two effects interact, since real-life markets often exhibit an increase in scale along both dimensions relative to typical experimental markets.

The commonly observed misreporting in strategy-proof mechanisms suggests exploring whether advice and communication between players can improve outcomes. The next section provides an overview of experiments investigating this question.

3.2.3 Nudging, chatting, and advice

An important aspect of market design is how the rules of the market are explained to the participants and what information they receive about the strategic properties of the mechanism. Experimental economists usually refrain from pointing out the optimal choice or the Nash equilibrium to participants, but this maxim does not necessarily hold for experiments in market design. The reason is that explanations and advice are part of the design of markets, and experiments can be useful for testing the effectiveness of providing such advice. For example, some studies explore systematically how participants can be taught to state their true preferences under a strategy-proof mechanism.

The first experimental paper on advice given to subjects in matching markets is the paper by Guillen and Hing (2014). The subjects played against three computer players under the TTC mechanism. In the baseline, they submitted their preferences in a one-shot game. In the other three treatments, the subjects received advice from a third party before submitting their preferences. This advice was either correct (report truthfully), wrong (think about realistic schools), or both pieces of advice were given at the same time. The advice was framed as advice from a third party in order to avoid experimenter demand effects and possible concerns regarding deception. Subjects were told that the information was found in a newspaper, or on parental forums, or was spread by word of mouth. The information given to the subjects was not deceptive, since the wrong advice was indeed taken from the Boston school board forum of parents. In all three treatments with advice, the effect on truthful reporting was detrimental. While the percentage of truthful reports was above 70% in the baseline without advice, it was only 50% in the case of correct advice, 28% in the case of wrong advice, and 42% in the case of both types of advice. The differences to the baseline treatment without advice are significant in the treatments with wrong advice and with both pieces of information, while only marginally significant in the case of correct advice. The most puzzling result of the paper is the negative effect of correct advice on truth-telling. Possibly, the subjects became suspicious due to the

source of the advice that was indicated to them. Moreover, the detrimental effect of two contradicting pieces of information on truth-telling points to the possibility that participants find advice in favor of manipulations more convincing than advice to report truthfully. The study highlights the importance of understanding how correct advice should be given to participants when they also receive wrong advice from their peers.

Another study on the effect of advice in TTC was conducted by Guillen and Hakimov (2018) in a field setting. The topics of semester projects were allocated among students in a microeconomics course. In order to identify the preferred topic of each student among three possible topics, the authors first asked the students to choose their most-preferred one. Later, the professor announced that the distribution of choices was not satisfactory and therefore an allocation procedure had to be used. In the baseline treatment, the students received the usual experimental instructions about TTC explaining the mechanism to them. In the second treatment, they were additionally given advice to report truthfully. In the third treatment they only saw the advice without learning the details of the mechanism. Contrary to Guillen and Hing (2014), the advice to report truthfully significantly increased the rate of truthful reporting from 81% to 94%. Interestingly, the disclosure of the mechanism reduced the rate of truth-telling among a subsample of subjects. Because the advice was given in a natural setting by the lecturer of the course, it may have come across as more natural and credible. However, a positive effect of advice has also been observed in a lab experiment. Braun et al. (2014) explained to their subjects the strategy-proofness of DA (and made available a verbal explanation of the proof), which led to more truthful reporting than in the treatment without advice. Thus, the source and the framing of the advice seem to matter.

Koutout et al. (2018) replicate the designed environment of Chen and Sönmez (2006) in the baseline under BOS and DA, and introduce strategic advice for both mechanisms in the main treatments. In DA with advice to report truthfully, the proportion of subjects reporting truthfully is 19 percentage points higher than in the baseline without the advice, and this difference is statistically significant. In BOS the advice included the statement that the truthful strategy is risky and one of the following two strategies were suggested: the risky strategy of ranking the true top choice first and ranking the district school second, or the safe strategy of ranking the district school first. The advice led to a significant increase in the proportion of subjects who played one of these two strategies, and a significant decrease in the proportion of subjects who submitted their preferences truthfully. Note that under DA advice decreased the number of blocking pairs, but slightly decreased the average payoff due to the conflict between stability and efficiency. Under BOS advice led to an increase in the number of blocking pairs and a decrease in efficiency.

In real markets, advice is often given by peers. Parents in school choice programs consult with the other parents participating in the mechanism or with parents whose child participated in the program in previous years. Ding and Schotter (2017) studied how the possibility to chat before

submitting one's preferences to the system influences the reports and the market allocation in DA and BOS. Each subject took two decisions in the experiment. The first decision was taken individually and the second decision after chatting with other subjects. Either the participants chatted with another participant with the same preferences, or with a participant who had different preferences. The main result is that with both chatting protocols, chatting increased the likelihood of subjects changing their reports, which in turn led to, on average, higher payoffs of subjects who chatted relative to those who did not chat both in DA and BOS. However, chatting had no significant effect of truth-telling under both mechanisms. Finally, there was no difference between truth-telling rates under BOS and DA in both phases.

In a companion paper, Ding and Schotter (2015) investigate the effect of intergenerational advice and attempt to mimic the communication between parents about their strategies with previous cohorts of parents. In the experiment subjects played either DA or BOS. The other dimension of treatment variation was the source of learning: subjects either played the same mechanism repeatedly for 20 rounds, or received advice from the previous generation of players but played the mechanism only once. In this intergenerational advice treatment, right after learning about their allocation the subjects were asked to give advice to the next group of participants. To incentivize them to give correct advice, subjects earned 20% of the payment of the subject to whom they gave the advice. Contrary to the increase in truth-telling rates when DA is played repeatedly, intergenerational advice led to a significant decrease in truthful reporting from 72% in the first five rounds to 44% in the last five rounds. In BOS, in contrast, the advice increased truthful reporting. In both DA and BOS, advice strongly increases the probability of the advised strategy being chosen. Based on a structural estimation, the authors demonstrate that any advice, even the advice to choose a dominated action, increases the probability of playing the advised strategy. Turning to the question of how convincing a certain piece of advice is, the advice to play the most frequent non-truthful strategy, namely exchanging the top and the second most-preferred choices in the reported lists, is followed most often. This advice increases the probability of playing the strategy from 32% to 74%, i.e., by 42 percentage points, while correct advice increases the probability of truthful reporting from 54% to 88%, i.e., by 34 percentage points. Note that in the experiment each subject received only one piece of advice, and these numbers are based on a between-subjects comparison.

Rees-Jones and Skowronek (2018) conducted a large experiment with medical students immediately after their participation in the medical residency match (NRMP) that relies on the DA mechanism. Unlike the other papers which implement advice in experiments, the authors investigate the effect of advice by surveying participants about the advice received in the NRMP. After participants submitted their rank-order lists in the experiments (see the next section), they were asked whether they had received advice from the medical school, NRMP, or other students, and if so, which kind of advice was given. The NRMP website turned out to be the most reliable source regarding the content

of advice, as 75% of students who reported receiving advice from NRMP report the correct advice. Other sources often provided mixed – correct and wrong – advice. In line with the field evidence of Guillen and Hakimov (2018), the authors find a positive effect of receiving correct advice from NRMP.

Summing up this section, in all but one study the correct advice increases the rates of truthful reporting. However, there is evidence that the subjects are more likely to follow wrong advice, namely to manipulate their preference reports, than the correct advice to state their preferences truthfully. Thus, two challenges emerge regarding the provision of advice in practice. First, it is essential to make sure that the advice coming from officials (the clearinghouse, schools, or hospitals) is correct, since it has a significant effect on choices. Second, it is necessary to improve the way in which such advice is delivered in order to make sure that it is more convincing than the advice of peers that tends to be wrong. The latter is a challenging empirical question that invites further research.

3.2.4 Determinants of reporting strategies – biases, risk-aversion, and cognitive ability

The papers presented in the previous sections demonstrate that a substantial share of subjects misreport their preferences under DA, despite experimental treatments aimed at limiting the submission of dominated strategies. In this section, we try to take a closer look at the types of strategies subjects used, and we survey possible determinants of truth-telling and manipulations that have been investigated in the literature.

Chen and Sönmez (2006) identified three types of biases that subjects display: a district school bias, a small school bias, and a similar preferences bias. The district school bias refers to a participant putting her district school higher up on the reported list than its position in the true preference order. Under BOS, the district school bias can be part of an equilibrium strategy. Participants with a small school bias move smaller schools down to a lower position than in the true preference order. Participants with a similar preferences bias put schools with the highest payoffs into lower positions. This manipulation is interpreted as subjects assuming that other subjects have the same or similar preferences. Participants did not know the preferences of others, nor the degree of the correlation of preferences.

In the experiments of Chen and Sönmez (2006), almost two-thirds of subjects misreported their preferences in line with the district school bias under BOS. Note that in the majority of cases the biases cannot be uniquely identified, which explains why the following proportions do not add up to 100%. As for the strategy-proof mechanisms, the district school bias was consistent with 34.8% and 31.5% of the misreported lists in the designed and the random environment under DA, and with 58.4% and 80.8% of the misreported lists in the designed and the random environment under TTC. The respective numbers for the small school bias are 84.9% and 59.5% for DA, and 91.8% and 46.4% for TTC. The numbers for the similar preferences bias are 84.9% and 62.6% in DA, and 80.6% and 75.6% in TTC.

Pais and Pintér (2008) study the district school and the small school bias. In their allocation problems, the small school bias and the similar preferences bias coincide, since the small schools are the most competitive. In the full information environment under DA, the district school bias was found in 17.8% of reported lists, and in addition, 8.9% of the lists were consistent with both the district school bias and the small school bias. Overall, the district school bias can explain 80.2% of misreported lists. In the full information environments under TTC, 8.9% of lists were consistent with the district school bias, which explains 67% of all misreported lists.

Despite the high percentage of reports under strategy-proof mechanisms that are explained by the small school bias and the similar preferences bias, a number of studies concentrate on the district school bias. One reason is that it is in line with the typical strategic advice given to participants for BOS in school choice procedures. Another reason is that many studies used schools with an equal number of seats, such that there are no small and big schools.

Unlike previous studies, Guillen and Hakimov (2017) found that only around 10% of manipulations in TTC are in line with the district school bias. One reason for the relatively small percentage of district-school bias manipulations might be that the district school was always at the bottom of the true preference list. The switch of the top two choices was the most common misrepresentation, which seems to be in line with the similar preferences bias of Chen and Sönmez (2006). However, the experiment by Guillen and Hakimov reveals that the cause of these misreports must be a different one. In Chen and Sönmez (2006), subjects did not know the preferences of other subjects, and thus the authors attributed the switch of the two top choices to the similar preferences bias, since they assumed that these switches were driven by the belief that participants might have similar preferences. In the case of Guillen and Hakimov (2017), the participants knew the other subjects' preferences and there was no conflict of top choices. Thus, the switch of the two top choices cannot be rationalized easily.

Further evidence of such irrational choices comes from the experiment by Ding and Schotter (2017) in which three out of five players have their district school as their second most-preferred school. First, they find that 56.9% of these players submitted preferences in line with the district school bias in the second phase of the experiment, which explains 96% of their misreports. Again, note that these submissions are also in line with the similar preferences bias. Second, the other two player types vary their reports in a manner that allows us to distinguish between the similarity of preferences bias and an irrational choice that cannot be rationalized by beliefs about other students' preferences. Both types had no priority at their second choice and reported it first on their list 60% and 52% of times, respectively. In the case of the type reporting the second choice first in 60% of the cases, the true second choice was not popular among other players, while in the case of the type misreporting in 52% of cases, it was the most-preferred choice of the other players. This provides another piece of evidence

that the switch of the two top choices cannot be rationalized by the similarity of preferences bias. Note that these manipulations are also not in line with the district school bias.

More evidence of switching the first and the second preference comes from Klijn, Pais, and Vorsatz (2013). They study the effect of preference intensities and risk aversion on application strategies under DA and BOS. Three participants competed for three seats in three schools. The payment for receiving the first choice and the last choice was fixed, while the value of the middle option changed between the treatments by being either closer to the top choice, in the middle between the top and the last choice or closer to the last choice. The safe school (the analogue of the district school) was always the least preferred by the subjects. In DA, 53% of reports were truthful, and this proportion did not vary significantly between treatments with different preference-intensities. Between 6% and 14% of reports under DA were in line with the district school bias, while the majority of misrepresentations were switches of the first and second choices. As all three participants had different most-preferred schools, once again these strategies are only consistent with irrational choices and not with the similarity of preferences bias. The frequency of this switching strategy was 34% but varied between conditions: it was 19% when the relative value of the second choice was the lowest and 43% when it was the highest. Thus, the higher the value of the second choice, the more frequent were irrational switching choices and the less truthful reporting was observed. A similar tendency of switching first and second preferences was observed in BOS, where this strategy can be in line with equilibrium. Moreover, the authors found a positive correlation between risk aversion (the switching point in the Holt and Laury task) and the propensity to submit the truthful strategy in DA. Note, however, that the effect was mostly driven by extremely risk-averse subjects, who switched to the less risky option in the Holt and Laury task with a 90% and higher probability of winning. The effect of risk aversion on the propensity to misreport in TTC was also studied by Guillen and Hakimov (2017), and no correlation of misreporting in TTC with the measure of risk-aversion was found. Basteck and Mantovani (2018) show a positive correlation of risk tolerance with payoffs in BOS.

To investigate the reasons for biased choices, several studies include some measure of the cognitive abilities of subjects. Guillen and Hakimov (2017) use the CRT and Wonderlic tests to measure cognitive ability. They find that subjects who performed well in these tests were more likely to report the preferences truthfully under TTC. Basteck and Mantovani (2018) study whether lower cognitive ability students are disadvantaged in BOS. The authors sorted subjects into two categories of cognitive ability, namely Low and High, based on the performance in a Raven matrices test in the first part of the experiment. Then, in order to make the preference profiles uncorrelated with ability, they assigned preference profiles to students such that half of the students with each preference profile are Low subjects and half are High subjects. They found that High subjects are more likely to truthfully report under DA and are less likely to do so under BOS. This led to higher earnings of High subjects relative to Low subjects in both mechanisms, but the difference is significantly higher in BOS, which confirms the

concern that BOS disadvantages students of low cognitive ability. These findings are complemented by the empirical work of Dur et al. (2018b) which quantifies the cost of sincere reports under BOS and shows that they are substantial.

In a follow-up paper, Basteck and Mantovani (2018b) investigate whether information about the popularity of the schools (in particular, the number of students who ranked the school first in their reported preferences) helps to level the playing field and close the gap between High and Low subjects under BOS. The authors use two different school choice problems. In the treatment with information, the proportion of Low subjects best-responding to the average play of others is higher than in the treatment without information in both problems. As for High subjects, there is a significant increase in the proportion of best responses in the treatment with information relative to no information in only one of the two problems. Despite a significant treatment difference in best responses between High and Low subjects, there is no significant difference in payoff gaps between treatments. The authors explain this finding by a higher propensity of High subjects to play the best response in high-stakes situations. Thus, while Low students improve, High students do so as well.

Hakimov and Bó (forthcoming) used an incentivized quiz for DA where subjects were paid two euros if they were able to correctly determine the allocation of a school choice problem. The authors found no correlation between truthful reporting and the ability to find the correct allocation when controlling for other factors such as the preference profiles and priorities. Instead, the main determinant of truthful reporting was the priority of the student: the higher the average priority, the more likely she was to report truthfully. This observation is in line with the district school bias identified in lab experiments, since high priority students can be sure to get into their most-preferred schools. It is also in line with the field observations of Hassidim et al. (2016) who show in the context of admissions to psychology programs in Israel, where applicants are ranked by the programs mostly based on their school grades, that applicants with bad grades are more likely to submit dominated rank-order lists to the DA mechanism than the applicants with good grades. However, this result from the field can be driven by differences in the priority and the cognitive ability of students. Moreover, Schmelzer (2018) found that subjects with very low and very high levels of contingent reasoning, as measured by choices in the beauty contest game, are more likely to report truthfully in RSD and TTC than subjects with intermediate levels of contingent reasoning.

Finally, Rees-Jones and Skowronek (2018) conducted a large experiment with 1,714 medical students immediately after their participation in the medical residency match. These students went through significant training and advice regarding the mechanism used (a modified version of DA that is strategy-proof for medical students). In the experiment, students were told that they would be allocated to hypothetical residency programs using the same mechanism as in NRMP, and they had access to a detailed explanation of the mechanism. The preferences of students were generated such that all students had the same preferences over five residency programs. The preferences of residency

programs were correlated with hypothetical test scores which were known to the students. However, the preferences were not uniquely defined by the test scores, and students were aware that every student could be assigned to every program with at least some positive probability. It turns out that 23% of students did not report their rank-order lists truthfully. This finding shows that preference misreports in DA can be observed for a highly relevant group of participants in a lab experiment. The authors also investigate some variables that influence misreporting. Similar to other studies, it is found that students with a lower performance in cognitive tests and students with lower perceived chances of being accepted to the best residency programs (students that were assigned low test scores in the experiment) are more likely to misreport their lists. The authors also asked participants whether they trusted NRMP to run the mechanism correctly, and 97% of participants indicated that they did trust the system. However, when asked whether the medical residency programs ranked students fairly, only 42% of participants agreed, which is negatively significantly correlated with truthful reporting.

From the papers in this section it can be taken that despite correlations of subjects' misreports with various measures of ability, clear evidence on the reasons for these misreports is still missing. In many studies, the modal manipulation is the switch of the two top choices, which is present even when it cannot be rationalized by the district school bias (safety motive) or by the similarity of preferences bias (motive to avoid competition). Misreporting could be due to a limited cognitive ability, to a demand effect resulting in the belief that truth-telling is too simple and therefore cannot be the optimal strategy in the experiment, or due to the perception of the deferred acceptance mechanism as immediate acceptance. Moreover, as pointed out above, most of the biases were documented in one-shot experiments. It would be interesting to look at the biases in repeated environments to see which of them are more persistent than others. Experiments with repeated play could allow for a cleaner identification of possible biases and their drivers.

3.2.5 Dynamic mechanisms

While most research has focused on direct mechanisms where students have to report their rank-order list before the algorithm is run, dynamic or iterative mechanisms can be an alternative. Motivated by the high rates of misreporting observed in lab experiments and in field studies with strategy-proof mechanisms (see Hassidim et al. 2016, Li and Pereyra 2017, Rees-Jones 2018), a number of papers test dynamic mechanisms. Unlike direct mechanisms, these mechanisms allow for multiple interactions between the participants and the designer. This allows for learning and for the correction of mistakes during the allocation process. Due to their dynamic nature, these mechanisms might provide feedback to the participants about intermediate allocations. However, the exact

implementation of the mechanisms differs largely between studies, and seemingly small details influence the theoretic properties of the mechanisms as well as their observed outcomes.

The first two studies reported on in this section by Klijn et al. (forthcoming) and Bó and Hakimov (forthcoming) consider the iterative DA mechanism. Instead of the submission of rank-order lists before the algorithm starts as in static DA, under iterative DA the proposing side makes one proposal at a time, and if it is tentatively accepted, the proposer cannot make any other proposals.²⁴ If it is rejected, the proposer is asked to make another proposal. Following the literature, we use the terms “iterative” and “dynamic” interchangeably to refer to this modification of DA. Klijn et al. (forthcoming) compare the student- and school-proposing DA to their iterative counterparts. The experiment implements a one-sided matching set-up, where the schools were played by the computer and always behaved truthfully (either by proposing in the order of priorities or by choosing among proposals according to their priorities). The authors used four different environments with complete information. Each environment had four students and four schools, with one seat each. Subjects played the same environment for six rounds in a row before switching to the next environment. The authors show that the strategy-proofness of the static student-proposing DA is lost in the dynamic version of the mechanism. In the school-proposing DA, the set of equilibrium outcomes for the static and dynamic versions coincide, but a wider range of behavior is supported in equilibrium in the dynamic version. The results of the experiment show that subjects switch the ranking of schools relative to the true preferences weakly more often in static mechanisms than in the dynamic versions. The overall truth-telling rate in dynamic DA was 55%, while in static DA it was only 47% and the difference was not statistically significant. Moreover, there is no significant increase in the number of stable allocations under the iterative student-proposing DA relative to its static version. In the school-proposing DA (where offers were made by the computer in the order of priorities), however, in two out of four environments the dynamic version leads to a significantly higher proportion of stable outcomes. Finally, the overall frequencies of stable matchings were 63.88% in the static student-proposing DA, 77.08% in the dynamic student-proposing DA, 68.75% in the static school-proposing DA, and 89.93% in the dynamic school-proposing DA.

In a closely related paper by Bó and Hakimov (forthcoming), the authors also compare static and iterative versions of DA, concentrating only on the student-proposing versions. They implemented incomplete information about the priorities. Subjects played each mechanism for 20 rounds, facing a new set of preferences and priorities in each round. Eight students competed for seats at eight colleges with one seat each. Unlike Klijn et al. (forthcoming), they find that under iterative DA stable allocations are reached significantly more often than under static DA. The difference is driven by a significantly higher proportion of subjects behaving truthfully in the iterative mechanism.

²⁴ Both studies use a set-up where schools only have one seat each. Therefore, the schools only make one proposal at a time.

The results of both Bó and Hakimov (forthcoming) and Klijn et al. (forthcoming) indicate a failure of the theoretical prediction of more truthful behavior under static mechanisms. The results of Bó and Hakimov (forthcoming) even show more truth-telling in the dynamic mechanism. The differences between the two studies may be due to the environments studied, namely markets with four schools in Klijn et al. versus eight schools in Bó and Hakimov. In environments with four schools, higher rates of truthful reporting are observed than in environments with eight schools under static DA, in line with our analysis of several studies in section 3.2.2. This makes it harder in the four-school environment to observe a significant difference between the mechanisms. Moreover, in environments with eight schools, one round of the iterative mechanism takes longer and thus provides more learning opportunities and more possibilities to correct mistakes within one round. Bó and Hakimov (forthcoming) investigate different possible explanations for the observed difference between static and dynamic DA and conclude that the advantage of the dynamic DA is the feedback it provides after every step (rejection from the previously applied university) and the possibility to re-strategize given this feedback. This allows subjects to realize that the strategy of skipping the most-preferred schools does not help and to make them abandon this strategy in favor of truthful behavior.²⁵

Additionally, Bó and Hakimov (forthcoming) conducted a treatment under the iterative DA where the tentative cut-off grades of students are posted after each step of the iterative DA. These grades reflect the minimum grade necessary to be accepted by a particular university at each step. The cut-off grades can only improve between the steps of the iterative DA, since only those applicants who were not tentatively accepted can reapply. The provision of these cutoffs leads to a significant increase in truthful behavior relative to the iterative version without the provision of cutoffs, but this increase does not translate into a significantly higher proportion of stable allocations.

The effect of intermediate information in the allocation process is also explored by other studies. Stephenson (2016) tests the effect of continuous feedback on allocations depending on the lists submitted by the participants. The subjects first submitted a report, were then able to revise it multiple times, and immediately received feedback about the allocation they would have reached given the current reports of all students. The treatments vary with respect to the mechanism used, namely BOS, DA, and TTC, and the frequency of the feedback, either after all participants submit their reports (discrete feedback), or already after tentative reports (continuous feedback). In all three tested mechanisms, the continuous feedback improved the rationality of the lists submitted and moved the allocations significantly closer to those predicted by the theory. A stable outcome was reached in 83% of markets under discrete feedback, and in almost 99% of markets under continuous feedback, which are the highest rates of stability observed in the literature.

²⁵ It is worth noting that these experiments are closely related to Echenique et al. (2016) who implemented dynamic DA in a two-sided matching market. However, they do not compare the results to the static version of DA. The paper is discussed in section 4.

Gong and Liang (2016) study the college admission mechanism of the Chinese province of Inner-Mongolia. It is a dynamic mechanism where students are given real-time feedback about the current allocation and are allowed to revise their choices. The mechanism is based on DA, but unlike in the iterative DA that was tested by Bó and Hakimov (forthcoming) and by Klijn et al. (forthcoming), subjects can revise their applications at any time. In this respect, the mechanism is similar to the continuous feedback explored by Stephenson (2016). The difference is that participants can only submit one choice at a time as opposed to a full rank-order list in Stephenson (2016). Additionally, the students are split into groups according to their grades. Each group has its own deadline for the final submission after which the allocations are finalized. The authors compare the dynamic mechanism to standard DA and BOS. In the environment with highly correlated preferences, the dynamic mechanism leads to significantly less stability and lower efficiency than in DA, while students misreport at a similar rate. In the low correlation environment, the dynamic mechanism is as stable and efficient as DA but has lower rates of misreporting.

Dur et al. (2018) investigate a modification of BOS where rank-order lists are submitted sequentially, and late movers can observe the submissions of previous students. This is motivated by the mechanism used in the Wake County Public School System. The equilibrium of BOS with sequential submission under complete information can lead to improvements in efficiency compared to standard BOS. In experiments with four students and four schools with one seat each, the authors compare the standard versions of BOS and DA to mechanisms where the preferences are submitted sequentially. The order of moves is predetermined in the experiment. The theoretical predictions hold: while there was no difference in efficiency between both versions of DA and standard BOS, BOS with sequential submissions reached the highest level of efficiency and the difference to the other mechanisms was significant. In DA, on average, 77% of students reported truthfully, and the rates were not different in a standard DA and a DA with sequential submissions.

A dynamic mechanism is also considered by Li (2017) who compares RSD in a standard and in a sequential version. In the standard version, subjects are asked to submit their full rank-order lists over all options. In the sequential version of RSD subjects have to pick their preferred choice from a set of options. The sequential version of RSD is obviously strategy-proof, a property of the incentive structure introduced by Li (2017). A truthful strategy is obviously dominant if, for any deviating strategy, starting from any earliest deviation from the truthful strategy, the best possible outcome from the deviating strategy is no better than the worst possible outcome from the truthful strategy. A mechanism is obviously strategy-proof if it has an equilibrium in obviously dominant strategies. Thus, in the sequential version of RSD it does not require contingent reasoning by the players to realize that truthful behavior is a dominant strategy, while it is required in the standard version of RSD (Li 2017). The experimental results show that higher rates of truthful behavior are observed in the dynamic version of RSD than in the static version, which can be explained by obvious strategy-proofness.

Summing up this section, the experimental evidence at least weakly favors dynamic mechanisms over their direct counterparts, the only exception being the study by Gong and Liang (2016) where the evidence is mixed. Note, however, that despite the number of papers testing mechanisms with dynamic features, the meaning of the term “dynamic” varies. Very broadly, dynamic mechanisms can be categorized into three groups:

- Subjects take decisions on their rank-order lists step-by-step, submitting one choice at a time, as in iterative DA studied by Klijn et al. (forthcoming) and Bó and Hakimov (forthcoming) or in sequential RSD studied by Li (2017).
- Subjects are allowed to revise their strategies (rank-order lists or just one choice) multiple times as in the experiments by Stephenson (2016) and Gong and Liang (2017).
- Subjects report rank-order lists one after another in a standard direct mechanism, learning about the strategies of players who chose before them, as in Dur et al. (2018).

While all these modifications seem to lead to improvements in the quality of the allocations relative to static mechanisms, the size of these improvements varies between studies and modifications. More research is needed to understand whether the benefits are robust and what their channels are.

3.2.6 Constrained rank-order lists

In many settings, the number of items that can be ranked on the preference lists is smaller than the number of options available. The effects of such constraints have been studied by Calsamiglia et al. (2010) for BOS, DA, and TTC. They used the designed and random markets of Chen and Sönmez (2006). In the designed treatment small schools and district schools are more preferred while preferences are uncorrelated in the random markets. There are 36 students that have to be assigned to seven schools. In the constrained treatment, only lists of up to three schools can be submitted while the lists can contain up to seven schools in the unconstrained environment. In case a student is not accepted by any of the three listed schools, she remains unassigned and receives a payoff of zero from the match. A 3x2 design was employed in a one-shot environment with incomplete information about the preferences of other applicants. As predicted, since DA and TTC are no longer strategy-proof for many students when the lists are constrained, it is found that subjects rank their safe district school higher and small schools lower in the constrained than in the unconstrained version of the two mechanisms. The district-school bias leads to fewer students not being assigned to their district school in all three mechanisms. Furthermore, efficiency is significantly lowered by the constraint in all three mechanisms. Regarding stability, DA performs better than BOS and TTC both under the constrained and the unconstrained version of the mechanism, but the constraint in DA significantly increases the

number of blocking pairs. The authors conclude that “removing the constraint will come at a small cost but will clearly improve the performance of the school choice mechanisms.”

The findings of Calsamiglia et al. (2010) were replicated for DA by Bó and Hakimov (forthcoming)²⁶ who compared standard DA to constrained DA as a robustness check, and find that constrained DA leads to stable matches significantly less often than unconstrained DA. Also, efficiency under constrained DA is significantly lower than under unconstrained DA. Interestingly, these relations also hold for the iterative versions of DA where the constraint means a limit on the maximum number of steps of the iterative procedure. The experimental results show that the constrained number of steps in iterative DA leads to lower rates of truthful behavior, stability, and efficiency relative to the unconstrained iterative DA, with all differences being statistically significant.

One rationale for constraining the length of the rank-order lists is the cost of dealing with many applications that can lead to congestion. He and Magnac (2017) use a field experiment to study how the costs of university programs incurred by inspecting student applications under DA can be reduced by restricting the number of choices on the rank-order list. They run a field experiment with 129 students applying for seven masters’ programs at Toulouse School of Economics. They compare constrained and unconstrained DA with a version of DA that includes a Pigouvian tax on each application that is supposed to internalize the externality imposed on the selection committees of the programs. The tax was implemented by requiring a motivation letter from the students for each application from the fourth one on. The students knew that either DA, DA with motivation letters, or constrained DA with only four programs would be implemented. Each student could submit a rank-order list for each of the mechanisms. The authors treat the submissions under unconstrained DA as the true preferences. This allows them to simulate the effect of the tax and the constrained lists on stability. While both DA with a Pigouvian tax and constrained DA significantly lower the number of applications to each of the programs, the constraint on the list leads to high distortions of stability and to some students being unassigned. Simulations and counterfactual analyses suggest that the small application cost is the best regime: while lowering the screening costs of the programs due to fewer applications, stability is unaffected.

Summing up, we conclude that there is robust evidence of a detrimental effect of a constraint on rank-order lists in DA on stability. However, once the constraint becomes endogenous and is implemented as a small tax, the detrimental effect disappears. This may be due to the effect that risk-averse participants overspend on the tax instead of dropping the good schools from their reported list. Thus, the participants possibly report more options than the student-optimal stable match requires, while in the case of constrained lists risk-aversion might drive them to not list the student-optimal stable match in favor of safer schools. This differential effect of a constrained list versus a Pigouvian tax on reporting behavior might be of interest for further studies.

²⁶ The results are in the appendix of the paper.

3.2.7 Affirmative action

In many school choice and university admissions procedures, affirmative action policies or quotas for certain groups of students play an important role. Also, lotteries are used to admit students. The goal can be to increase the enrolment of minority students, to foster diversity in schools, or to satisfy legal rules stipulating the use of such admission criteria. Experiments have been employed to understand the effects of the various ways to implement affirmative action policies in matching markets.

Two alternative approaches to affirmative action are majority quotas and minority reserves. Majority quotas specify the maximum share (or number) of seats in each school that can be allocated to majority students. Minority reserves specify the share (or number) of students for each school such that in case the number of minority students in the school is lower than this number, any minority student is preferred to any majority student. Before turning to the experiments, we summarize the main theoretical findings guiding the experiments. Kojima (2012) and Matsubae (2011) show that the introduction of majority quotas can result in undesirable effects for both majority and minority students under DA and TTC. The reason is that a majority student, who is rejected by her preferred school due to the quota and who therefore gets a seat at her second most-preferred school, may thereby take the seat of a minority student at this school, making the minority student worse off. Hafalir et al. (2013) analyze minority reserves and demonstrate that the reserves do not affect the strategy-proofness of DA and TCC, and are an improvement relative to majority quotas. In essence, the efficient implementation of affirmative action requires flexibility such that reserved seats that are not taken by minority students can be filled with majority students.

Two possibilities of implementing quotas in a flexible manner have been studied experimentally by Braun et al. (2014). The first relies on the existing procedure for university seats in medicine in Germany where quotas for certain groups of students are filled sequentially. First, the 20% of applicants with the best grades can submit a rank-order list of universities, thereby competing for seats reserved for the top-grade students. Then, the same students have a chance to participate in the allocation of general seats by submitting a potentially different preference list. The German procedure creates incentives for strategic behavior by the students who are eligible for the quota for applicants with the best grades, since the quotas are filled sequentially. Intuitively, the students with the best grades have two chances of being admitted: through the top-grade quota, or later when the remaining seats are distributed among all students. Sometimes a student is better off when she is not matched to a university right away through the quota for best-grade students, since she can be matched to a better university which rejected her in this quota but accepts her under the general quota later on. Thus, in equilibrium the students need to truncate their list for the top-grade quota that is

administered first. Thus, the strategic incentives are due to the sequential process of filling the quotas. The second procedure is a modified version of DA that is strategy-proof (Westkamp 2013) and that is similar in spirit to the minority reserves described by Hafalir et al. (2013). The experiment employed four markets that differ with respect to the correlation of preferences. Participants played each market three times in different roles, amounting to 12 rounds overall. The results show that many students fail to optimally truncate their preference list for the first quota, even when the truncation is a dominant strategy, and achieve worse outcomes than in the modified DA that is strategy-proof.

The theoretical results of Hafalir et al. (2013) were tested in the lab by Klijn et al. (2016a) who compare DA and TTC, each with and without minority quotas. The results show that the mechanism with reserves favors minority students, since they are less likely to form a blocking pair and have higher payoffs than in the mechanism without reserves. No effect on truth-telling rates was observed, except for an increase in truth-telling by some minority students in the case of DA with reserves when compared to DA. Overall, DA performed better than TTC regarding both stability and efficiency.

Kawagoe et al. (2018) employ experiments to compare majority quotas and minority reserves for DA. They used two environments. As predicted, in the first environment DA with minority reserves led to higher efficiency than DA with a majority quota, with no significant difference in the second environment. Both mechanisms are strategy-proof, but led to approximately 60% of truthful preference reporting with significant differences between mechanisms and environments.

Lotteries have been proposed to desegregate schools in the UK (School Admissions Code 2006) and they are also used in Berlin for this purpose. Basteck et al. (2018) study the existing school choice procedure that combines lotteries with the BOS mechanism and compare it to DA with an equivalent lottery quota. A certain proportion of seats at each oversubscribed school are allocated based on lottery draws. Thus, the policy reserves a certain proportion of seats for applicants with the highest lottery numbers, independent of their priority at the school (which is based on grades in Berlin). The BOS and DA mechanisms with a lottery quota are compared to BOS and DA without a lottery. In line with the theoretical predictions, truth-telling increases with the lottery in both mechanisms, and schools become less segregated. However, students of intermediate priority are less likely to receive seats at their preferred schools.

Summing up this section, the experimental evidence mostly confirms the theoretical predictions concerning the effects of quotas and reserves. If quotas are administered sequentially (as in the case of German university admissions), affirmative action can backfire and can destroy the incentives to report truthfully. On the other hand, lotteries can strengthen truth-telling both theoretically and empirically and lead to more mixed schools (as in the case of school choice in Berlin).

3.2.8 Information acquisition

The matching literature typically starts with the assumption that students know their own preferences. However, it is evident that forming preferences over a set of schools can be a time-consuming and complex task. Chen and He (2017) compare students' incentives to invest in learning their own preferences and the preferences of others under BOS and DA. They show that in theory students have incentives to find out both their own preferences and the preferences of other students under BOS. In contrast under DA, due to its strategy-proofness only one's own preferences matter for the optimal strategy. Furthermore, the willingness-to-pay for finding out about one's own cardinal preferences is predicted to be higher for BOS than for DA. The authors test these predictions in a lab experiment. The results show that the subjects' WTP for information on their own preferences is higher under BOS than under DA, as predicted. However, the WTP is significantly higher than what the theory predicts for DA. Regarding the WTP for information on the preferences of others, again subjects' WTP is higher under BOS than under DA, but it is significantly higher than predicted under both mechanisms. The welfare of the different information regimes is also studied. There are no significant differences regarding the efficiency of the allocations under the two mechanisms in the uninformed case. The free provision of information about the students' own preferences leads to significantly higher efficiency under both mechanisms, while providing information about the preferences of others for free has no effect on the efficiency in either of the mechanisms. Under regimes with information provision about own or others' preferences, however, the allocations reached under BOS are closer to the efficient one than the allocations under DA.

The results are related to the evidence in Pais and Pintér (2008). According to their results, information about the preferences of others decreases truthful reporting in DA. If this was the case in the experiment by Chen and He (2017), subjects would pay not only for information that is irrelevant, but also for information that makes them submit preferences less optimally. While Chen and He (2017) do not analyze their data with respect to this question, it can be expected that overinvestment in learning the preferences of others is even higher if this information leads to suboptimal reporting.

3.2.9 Preferences over mechanisms

It is possible that people have preferences over the matching procedures themselves and not just over outcomes. If this is the case, such preferences should be taken into consideration by the policy-makers when choosing among mechanisms. However, almost all existing studies on matching procedures implement a between-subjects design regarding the allocation mechanism. The main reason is probably that some of the mechanisms are not straightforward to understand, and the

instructions for only one mechanism are already quite involved. However, there are two experimental papers (Schmelzer 2016, 2018) that investigate subjects' preferences over mechanisms by using either two very similar mechanisms or one simple mechanism. Schmelzer (2016) studies DA with different tie-breaking rules for priorities. Motivated by recent policy debates, two common ways of dealing with ties due to coarse priorities are tested in the lab. The author elicits the preferences of subjects over DA with single and multiple tie-breaking by ensuring that the ex-ante outcomes (before the lottery) are the same by design. The subjects have to make two decisions, namely submit the preference lists under single and under multiple tie-breaking. Under single tie-breaking, all schools use the same lottery to rank students, while under multiple tie-breaking each school runs its own lottery. Without providing information about the allocation reached under each tie-breaking regime, the subjects can express their preferences (and pay 10 cents) for the procedure they prefer. One of the subjects is randomly chosen to determine the mechanism that will be applied. Though the majority of subjects are indifferent, among those who are not most express a preference for multiple tie-breaking.

In a second paper on choice between mechanisms (Schmelzer 2018), TTC is compared to RSD in a house allocation problem without existing endowments (see section 3 for a discussion of the results of the paper regarding truthful reporting). The hypothesis was that subjects might prefer RSD because of its simplicity relative to TTC. RSD is straightforward to explain, and thus it is possible to run it together with TTC in a within-subjects design without risking confusing the subjects. Around 40% of subjects are willing to pay a small amount to vote for one of the mechanisms, and the number of votes for each of the mechanisms is not significantly different.

To summarize this section, the elicitation of preferences over mechanisms has demonstrated that a large number of subjects are indifferent between the mechanisms. This might be driven by the fact that the mechanisms that were compared are similar from a theoretical perspective. It seems important to find ways to elicit preferences over more distinct mechanisms, like BOS, DA, and TTC. However, this is challenging since the preference over mechanisms has to be disentangled from the motive to reach the most-preferable assignment.

3.2.10 Timing of the publication of centralized exam scores

A seemingly small institutional detail such as the timing of information about the results from the university entrance exam can have implications for the matching outcome. Lien et al. (2016) run experiments inspired by a policy change in China regarding this issue. The universities accept students based on the exam scores. The experiment tests the hypothesis that not knowing the result of the exam when submitting one's preferences can lead to an ex-ante fair and efficient outcome. Ex-ante [ex-post] fairness means that there is no justified envy with respect to the expected [realized] exam

scores while efficiency takes into account the preference intensities (cardinal utilities). This hypothesis relies on the idea that people have an unbiased prior about their ability while the exam score is a noisy signal of it. Therefore, the noisy exam score can lead to a matching that is not stable with respect to ability under DA. The same holds under BOS if students know their realized exam score. However, if they do not know the score of the exam, they can only base their choice on their expected exam score which, by assumption, is a better measure of ability. In the experiments, students are informed of their true ability, that is, the average of their score distribution. It emerges that BOS where students do not know their exam score leads to the most efficient but least ex-post fair outcome, while there is no support for the prediction that it is ex-ante fair. Overall, it turns out that despite small markets (three students and three schools), the equilibrium strategies are often not played when students do not know their exam score.

Pan (forthcoming) questions the assumption by Lien et al. (2016) that people have an unbiased expectation of their ability before the exam score is published. She shows that biased self-assessments further weaken the ex-ante fairness of the matching under BOS. She runs experiments in a similar set-up, but instead of exogenously given priorities, the priorities in the mechanism were determined by a real-effort task. The theoretical prediction that BOS should lead to a higher percentage of ex-ante stable matchings under the regime of publishing the grades after the submission of preferences finds no support in the data, and DA outperforms BOS in all publishing regimes.²⁷ This is another piece of evidence for the strategic complexity of BOS. Despite the fact that in theory BOS can improve on DA with respect to efficiency, the prediction typically fails in experiments due to the low percentage of equilibrium outcomes reached under BOS (see, for instance, Featherstone and Niederle 2016).

Summing up this section, we can conclude that in situations where the benefits of a particular mechanism are based on the assumption of participants holding correct beliefs about their ability, the predictions often fail to find support in experiments. This is not surprising given the robust experimental evidence on biased self-assessments. This section emphasizes the importance of empirical evidence when recommending policies instead of solely basing the recommendations on the equilibrium solution. It shows that theoretical and empirical results lead to opposite recommendations, which is rare in matching experiments.

²⁷ Stability is defined here with respect to the abilities of students and not with respect to exam grades, which are a noisy measure of ability.

4 Two-sided matching

In two-sided matching markets, the agents on both sides of the market have preferences and can act strategically. Given that there is no stable and strategy-proof mechanism for both sides of the market (Roth 1982a), manipulations play an important role. Moreover, since the set of stable matchings has a lattice structure, the two sides of the market prefer different stable outcomes. The marriage market as a one-to-one matching market and the college-admissions model as a many-to-one model have been investigated experimentally. The section covers experiments on decentralized and dynamic two-sided markets with a variety of market rules, as well as centralized two-sided matching markets. With respect to matching mechanisms, the experimental literature focuses mainly on DA. This is because stability, unlike Pareto efficiency, can ensure the implementation of the resulting allocation, and avoids the renegotiation of matches afterwards. The following theoretical predictions hold:

(1) In DA, it is a weakly dominant strategy for the proposers to state their preferences truthfully unless colleges with multiple seats are proposing.²⁸ It is a weakly dominant strategy of receivers to report their first preference truthfully.

(2) The Nash equilibria of the game induced by DA lead to matchings that are stable with respect to the stated preferences.

(3) DA leads to the proposer-optimal stable matching.

Experiments are employed to test these theoretical properties, with a strong focus on the stability of market outcomes. A second focus is on the question of which of the stable matchings is selected if there is more than one and thus on the distribution of welfare. While the theory makes a clear prediction that the proposer-optimal matching should be reached, the experimental evidence does not univocally support this. The information available and – in dynamic two-sided markets – the exact market rules play an important role.

In the following, we organize the studies according to the market rules in place. Table 1 presents an overview of the studies summarized in sections 4.1 to 4.3 and shows the connection between the market rules and the experimental findings. We start out with experiments implementing the static DA mechanism where participants in the role of students and universities submit their rank-order lists, and the central clearinghouse played by the computer determines the matching. We then move on to dynamic implementations of DA where there are no submissions of rank-order lists but offers, acceptances, and rejections are made by the market participants one at a time, following the

²⁸ There is no stable and strategy-proof mechanism for colleges (Roth and Sotomayor 1990).

protocol of the static DA mechanism. Finally, we turn to unstructured interactions in two-sided matching markets, and to the possible unraveling of such unstructured markets.

4.1 Static DA

The earliest experiment on the marriage market studies the strategic misrepresentation of preferences by the participants under the student-proposing DA (Harrison and McCabe 1996). Markets with three or four students and three or four universities were played for 25 periods with complete information about the preferences of all market participants. Both the role of universities and of students was taken on by the experimental subjects, but the treatments varied the number of players that were computerized and always reported the true preferences. At least one market participant was played by the computer in every market, and markets with eight players were run with all students computerized. In the six-player markets, there were always two stable matchings, while in the four-player markets there were always four stable matches. Thus, the receiving side had incentives to misrepresent their preference reports in order to secure a more favorable stable matching.

In the environments with three students and universities where only one or two players were not computerized, the best response (either telling the truth for the proposing side or manipulating for the receiving side) was frequently chosen by the participants. However, with fewer computerized players and with markets of four students and universities, the subjects were not able to secure themselves favorable outcomes by manipulating their preferences. The efficiency of the matching was lower in markets with more human players, and experience did not help to reduce the efficiency losses. The number of blocking pairs was also higher in the larger markets and the markets with more human players.

The authors computed the payoff distance of the realized outcome to the student-optimal stable matching. They state that in markets with three players on each side, matchings were realized that were closer to the outcome preferred by the universities than to that of the students. However, in the markets with four players on each side, the matchings are closer to the student-optimal outcomes. To figure out what drives these results, the strategies of the universities given the two different market sizes can be compared in the treatments with computerized students who tell the truth. This analysis suggests that the strategic complexity for the universities increases from markets of six to markets of eight players and prevents them from optimal manipulations in the larger markets.

Market participants need to have sufficient information to effectively manipulate their preference lists. To investigate this systematically in a two-sided matching market with a centralized clearinghouse, a within-subjects design was employed by Pais et al. 2011 where each participant played DA under three different information conditions (no information except one's own preferences,

partial information, full information, always in this order). There were five students and three colleges, and two of the colleges offered two seats. The authors observe a considerable amount of preference manipulations by the colleges and the students with full information. In particular, under the student-proposing DA with full information only 56% of the students and 27% of the colleges report truthfully.²⁹ Almost all manipulations by the colleges (93%) were due to moving up a student that ranked the college highly in the true preference ordering. The manipulations by the students were predominantly due to moving up a college by which the student was ranked highly (a bias analogous to the district school bias) and moving down small colleges with fewer seats (small school bias), or a combination of both (73%). Truth-telling rates were higher in the treatments with less information available to the players. The stability rates ranged from 31% to 82% depending on treatments. The proportion of stable matches was 38% in the zero-information condition, 48% in the partial-information condition, and 2% in the full information condition.³⁰ In the zero-information condition, all stable outcomes were student-optimal, while with partial and full information the proportion of student and school optimal outcomes was approximately the same. Note that this is in line with the theoretical predictions, as colleges have the possibility to manipulate reports optimally only if they have enough information about the preferences of other players.

Strategic manipulations by the receiving side are also the focus of Castillo and Dianat (2016) who investigate the use of truncation strategies under DA. Truncation strategies are exhaustive in the sense that any matching that can be achieved with a misrepresentation that is not a truncation of one's preference list can also be achieved by a truncation (Roth and Rothblum 1999). When all agents can only truncate their list and no other misrepresentations are allowed, the optimal truncation strategy of an agent does not depend on the other agents' strategies. Thus, the complex market game is reduced to a single-person decision problem. A complete information environment is used, and the proposing side is computerized by stating the true preferences. Two behavioral predictions are tested, namely that the profitability and the riskiness of the truncation strategy affect its frequency of being chosen. However, it turns out that the observed truncation strategy of the receivers is not sensitive to the payoff differences between the different matching partners, contradicting the prediction. On the other hand, the hypothesized relationship between the riskiness and the likelihood of truncations finds support: Subjects whose best achievable matching partner is higher on their list are less likely to truncate their list above the best achievable partner, which would lead to no assignment. The lower the most-preferred achievable firm on the list is, the higher the probability that the subject will truncate the list optimally. A stable outcome is reached in 88% of the markets (but note that only

²⁹ Note that the authors also study BOS and TTC in a two-sided market and compare them to DA. The main result is that DA leads to the lowest proportion of stable matchings among the mechanisms. This finding is driven by low rates of truthful reporting by the proposing side, and suboptimal preference manipulations by the receiving side in DA.

³⁰ Partial information stands for subjects knowing their own preferences, the capacities of schools as well as the preference lists over students by each school up to capacity and the school that is top-listed by each student.

truncations above the best achievable matching partner can lead to unstable matchings in the markets considered). Finally, it is found that outcomes are closer to the receiver-optimal stable matching than to the proposer-optimal stable matching.

A closely related paper by Featherstone and Mayevsky (2015) also focuses on the strategies of the receiving side in DA. They run two environments: the first environment is characterized by multiple stable matchings, which implies that the receiving side has incentives to truncate the reported lists; in the second environment the stable matching is unique, which implies that the receiving side should report truthfully. In the experiments, the proposing side was played by computers, and truthful lists were always submitted. The results show that subjects do not differentiate between the two environments, and the rate of truthful reporting by the receiving side is not statistically different between the two environments. The difference increases in the last 10 rounds of the experiments but remains small (and only marginally significant). Surprisingly, a stable matching is never reached.

Summing up this section, a couple of findings are fairly robust in all experiments: the receiving side tends to manipulate the rank-order lists, but the success of the truncations and the rates are far from those predicted. Manipulations seem to be more successful when only a truncation of the rank-order list is allowed. Once the proposing side is played by human participants, the outcomes are even more distorted since both sides tend to deviate from equilibrium strategies. In the next section we will consider situations when the receiving side can react to the strategies of the proposing side.

4.2 Dynamic DA

Many matching markets are organized in a decentralized manner. However, even without a centralized clearinghouse, there can be rules that govern the process of offers and acceptances. If such rules are similar to the protocol of matching algorithms, the outcomes of decentralized markets can be described with the help of these algorithms. Thus, studying the dynamic version of DA is useful, since it can capture decentralized markets following the protocol of DA. Another feature of the dynamic version of DA is that it may be less demanding in terms of information collection and preference formation. Participants do not have to submit their full rank-order list but are only asked to name their preferred choice among a set of options. This may be one of the motivations for why policymakers have opted for dynamic mechanisms in real markets in recent years.³¹ Finally, for the sake of lab experiments the study of dynamic versions of DA has the advantage that when the preferences are induced by the experimenter, a dynamic procedure makes reporting the truth seem less artificial than in a static DA mechanism.

³¹ Examples are university admissions in Brazil and Inner Mongolia. See section 3.2.5.

The experiments of Haruvy and Ünver (2007) were designed to investigate the hypothesis that DA is a good predictor of behavior in repeated decentralized matching markets. Moreover, the study tests whether the amount of information of the side that receives offers affects the matching outcome. Workers and firms search for a matching partner by making and accepting offers bilaterally, without a clearinghouse as an intermediary. In the markets consisting of four firms and four workers, each firm can make an offer to one worker in every period. In the second part of a period, the workers decide whether they want to accept a contract offer or not. The market ends after a certain number of periods, where each of the periods is payoff-relevant. This feature of the game creates incentives for firms to skip offers that are rejected by the worker with a high probability. The final outcome, however, is still predicted to be stable, as contracts can be reneged upon in the next period and can also be repeated with the same worker. Three different markets requiring a different number of iterations to reach the stable matching were implemented in a within-subjects design. A 2x2 design was employed: all workers were either humans or computers programmed to state the truth, and low and high information conditions were implemented where market participants only had information about their own preferences, or about the preferences of all participants.

Haruvy and Ünver (2007) find that in all four conditions, the firm-optimal stable matching was reached in the majority of cases, confirming the hypothesis that DA can predict the outcome in certain decentralized settings. With non-strategic computerized workers, the firm-optimal stable matching was reached in 65% of the markets, while in 8% of the markets the worker-optimal stable allocation was reached. With participants playing the role of workers, the worker-optimal stable matching resulted in 16% of the markets, while the firm-optimal stable matching was still by far the modal outcome, with 72% of allocations.

A systematic analysis of strategic behavior under dynamic DA was conducted by Echenique, Wilson, and Yariv (2016). Their version of the mechanism follows the description of the DA algorithm, but offers, acceptances, and rejections are all executed by the subjects. In this respect the set-up is comparable to Haruvy and Ünver (2007), but in contrast to their study only the final allocation is payoff-relevant, and thus DA should predict individual behavior at each step of the game. Echenique et al. (2016) study different markets with eight subjects on each side and complete information, varying the number of stable matchings, and the cardinal representation of preferences. It is observed that proposers often skip entries of their preference list to avoid proposing to responders who do not rank them highly.³² Most of the responders in the experiment behave straightforwardly by accepting the best offer at any stage. But note that due to the dynamic implementation, the choices of the responders are restricted by the proposers' strategies. Across all markets, the average payoffs of proposers are closer to the average predicted payoff under the receiver-optimal stable matching than

³² This is consistent with the observations from school choice experiments in DA and dynamic DA where students move up on their list those schools that rank them highly.

under the proposer-optimal stable matching, which is a consequence of their skipping behavior. Finally, half of the markets lead to unstable outcomes, and only 29% of them are proposer-optimal.

The observed skipping behavior by proposers has interesting implications for patterns observed in the NRMP. In the NRMP, a high percentage of residents (proposers) receive their first choice (mostly above 50% in recent years) compared to their second choice (around 15%) and their third choice (below 10%). The authors argue that this is consistent with the decision heuristic that conflates the likelihood of matching with a certain partner and the preference for this partner.

An extensive form implementation of DA was also studied by Castillo and Dianat (2017). The main goal of their paper is to understand the impact of information of market participants about the preferences of others on their strategies and on the matching outcome. The paper is closely related to Pais et al. (2011), but runs DA in a dynamic setting, just like Echenique et al. (2016). Their main finding is that the distribution of information affects which outcome is selected, but it does not affect the stability of the outcome. As in Echenique et al. (2016), if proposers have full information about the preferences of the receivers, the proposers often skip preferred partners if they are not ranked highly by them. This explains the finding that the average distance to the responders' preferred stable outcome is smallest in the treatment with full information. Thus, the relationship between the stable matching selected and the amount of information is the same as in Pais et al (2011) for static DA but, in contrast to their findings, the stability of the outcome does not decrease in the amount of information available. This difference could be due to the dynamic implementation, giving participants flexibility to change their strategy in the process and thereby avoid instabilities. Additionally, the dynamic implementation allows receivers to observe that proposers manipulate in the receivers' interest, which leaves less room for profitable manipulations by the receivers. In static DA, receivers do not expect this and manipulate their lists, thereby failing to best respond to deviations from the predicted behavior of the proposers.

To sum up this section, the step-by-step implementation of DA limits the analysis of the strategies of the receiving side: the biggest difference to static mechanisms is found for the receivers who often best-respond to the deviating strategies of the proposers. However, we cannot distinguish whether receivers report truthfully because truth-telling has become the best response or because they do not understand the incentives for skipping and would have been truthful anyway. However, this possibility to react to the proposer's offer increases the proportion of stable allocations relative to the static mechanism in which both sides submit their strategies simultaneously. The effect of information is independent of DA being implemented as a static or dynamic mechanism: the more proposers know about the preferences of the other side of the market, the more they deviate from their optimal strategy, and thus the further away they move from the proposer-optimal toward the receiver-optimal stable matching. Note that similar findings exist for school choice problems where students react to priorities and tend to rank schools higher at which they have a higher priority.

4.3 Decentralized Matching

In this section, we consider markets without a clearinghouse where offers can be made at any time by both sides of the market. The proposed solution concept for such unstructured markets is the core, equaling the set of stable matchings (Roth and Sotomayor 1990). The predictions of cooperative game theory have been tested in different environments. An early contribution to the experimental literature has studied decentralized markets, motivated by the recruitment system of baseball players, the so-called free agent system. Nalbantian and Schotter (1995) investigate a market with three agents on each side who have to form matches and agree on prices. Agents only know their own preferences and the distribution of the preferences of all other agents. In their decentralized treatment, they seat participants in offices with a phone, providing them with their preference list over the three agents on the other market side and their phone numbers. All participants had five minutes to make calls and arrange contracts. Any contract that was formed was final and was announced to the other participants. This treatment is compared to a centralized (or simultaneous) treatment based on Demange and Gale (1985) where teams (and players) submit bids indicating the maximum (minimum) they would be willing to pay (must be paid) to be matched with each other. This information is used to solve for the vector of competitive or Walrasian prices implemented in the experiment.

While the outcomes of the two treatments do not differ significantly with respect to the efficiency of the outcomes, the sources of inefficiencies differ. The decentralized treatment where participants made phone calls never ended with unmatched agents but sometimes led to suboptimal matches such that, overall, 94.8% of the available surplus was captured. The centralized procedure (where all inefficiencies are due to misrepresentations by the agents) led to 89.4% of the surplus captured. (Note that the experiment was designed such that matching all agents would lead to at least 75% of the potential surplus.) Efficiency losses occurred in the centralized treatment because no matches were formed in a number of instances, but fewer suboptimal matches were formed than in the decentralized treatment.³³

The hypothesis that an unstructured two-sided matching market leads to a stable outcome was studied experimentally in a slightly more structured setting with eight agents on each side where both sides were allowed to make offers (Echenique and Yariv 2012). Each agent could accept and thereby hold up to one offer and make up to one offer even when holding an offer. This possibility of holding an offer “temporarily” makes the decentralized market more similar to DA, relative to the market when every acceptance is final, like in Nalbantian and Schotter (1995). A market ended after 30 seconds of inactivity. Complete information about the preferences of all participants was provided. A number of markets were designed that differed with respect to the number of stable matchings. On average, 76% of markets reached a stable allocation: 90% of markets with a unique stable matching,

³³ This paper is the only study on efficiency in the section on two-sided matching markets.

89% with two stable matchings, and 47% with three stable matchings. Note that these rates are higher than in the experiments described in the previous subsections (static and dynamic implementation of DA). This could be due to participants having the possibility to correct their mistakes, since there is no commitment to either an accepted offer or to an offer made before. Also, the possibility to observe the matches of others and thereby the blocking possibilities may be helpful in achieving stable outcomes. Another important factor is the absence of an exogenous time constraint, as the market ends only after 30 seconds of inactivity. This makes the initiation of a new round of offers less risky relative to a market where the deadline is exogenous.

Echenique and Yariv (2012) also found that if there are three stable matchings, the median stable matching was the most frequent outcome, i.e., the matching where each agent was matched to his median stable matching partner. While egalitarian preferences seem to play the main role for this outcome they do not trump stability, since unstable egalitarian outcomes were never realized. On the other hand, the cardinality of preferences played a role. If the median matching was not realized, the side that had more to lose tended to achieve its preferred matching. Regarding the dynamics within markets, it could be observed that blocking pairs disappeared quickly.

In a closely related paper, the information conditions and search frictions were varied to investigate their impact on the matching outcome in a decentralized setting (Pais et al. 2012). Unlike in Echenique and Yariv (2012), the markets had a pre-specified duration of four minutes. High information corresponded to full information about the preferences of all market participants while low information implied that agents only knew their own preferences. Search frictions were operationalized in two different ways: Making an offer involved a small cost, or an offer that was accepted led to a binding contract where the pair exited the market. Markets with five agents on each side were investigated. All agents could make offers (one at a time), and they could accept or reject offers whenever they wished. The markets employed had either two or three stable outcomes. A participant's screen showed the current matches in the market as well as the participant's own outstanding offer (if there was one).

The treatment of Pais et al. (2012) with low information and no frictions led to the highest number of offers made and the most stable outcomes. The market activity was reduced when offers were costly or when agents could not renege on accepted offers. Without frictions, the agents made offers in the order of their preferences whereas skipping the most-preferred partner(s) could be observed when offers were costly. Costly offers had a clear negative effect on the stability of the outcome. On the other hand, being committed to an accepted offer sped up convergence despite achieving less stability than without commitment. In markets with three stable outcomes, the median stable outcome was the one most frequently observed, as in Echenique and Yariv (2012).

The complete information no-frictions treatment of Pais et al. (2012) is comparable to the set-up of Echenique and Yariv (2012). Echenique and Yariv (2012) find that, on average, 76% of markets

with eight participants on each side reached a stable outcome, as did 67% of markets with 15 participants on each side. Pais et al. (2012) observe this for 69% of the markets with five participants on each side and only for 28% of markets with 10 participants on each side. Thus, in Pais et al. (2012) stability is lower and the negative effect of the increase in market size is stronger than in Echenique and Yariv (2012). Another difference is that in the baseline of Echenique and Yariv (2012) none of the stable matchings reached were extreme (proposer or receiver optimal), while in Pais et al. (2012) the extreme stable matches were more likely to be reached. We think these differences might be due to the different stopping rules for the markets: the inactivity rule versus the exogenous stopping time.

The experimental results on decentralized markets show that under ideal conditions (i.e., no search frictions) such unregulated markets can lead to more stable matchings than markets with a centralized clearinghouse employing a stable mechanism, especially when the number of stable matches is low. Moreover, decentralized markets seem to lead to the median stable outcome, unlike markets with dynamic and static DA where extreme matchings are selected more often. While these are interesting results, they need to be taken with caution, since the laboratory markets do not suffer from congestion and related problems. The introduction of frictions leads to a significant decrease in the stability of outcomes even in the lab. Moreover, the decentralized markets can fail for many other reasons as well. In the next section, the sources of market failure due to unraveling and their possible remedies are discussed.

4.4 Unraveling of markets

Markets function well when there are many potential transaction partners and there is sufficient time. Then, market participants can choose among trading partners and find the best match. But there are circumstances under which market participants have an incentive to make or accept offers before the market starts. These incentives can reinforce each other, inducing more people to form matches early on such that transactions become dispersed over time. When a market unravels, this can affect the efficiency and the stability of the matching outcome. There are two reasons for this: (a) agents can choose from fewer transactions at any given point in time, and (b) there is less information available on the quality of the matching partner. Such unraveling dynamics in markets can be countered by the establishment of a clearinghouse, such as in the case of the National Resident Matching Program in the US in the 1950s.

Various sources of unraveling have been studied, and some of the contributions have employed experiments. This group of papers is presented and discussed in the handbook chapter by Kagel and Roth (2016) in a compact manner, and we therefore only mention the contributions of these papers in passing. We then report on two papers not covered in the handbook chapter that focus on incomplete information as a source of unraveling.

The stability or instability of clearinghouses has been identified as one source of market unraveling with the help of data on British regional clearinghouses for young doctors and laboratory experiments (Kagel and Roth 2000). If a clearinghouse produces an unstable matching, agents who are likely to form blocking pairs have an incentive to match early. As a remedy for a market that has unraveled, a change in the rules governing early offers has been investigated experimentally. Encouraging the reneging of offers by young doctors discourages early offers (Niederle and Roth 2009) and can thereby roll back the unraveling process.

In the market for law clerks in the US, there are strict norms which make it almost impossible for law students to reject offers from judges they have applied to, unless they get another offer simultaneously. The importance of the possibility of rejecting offers to halt unraveling was identified with the help of an experiment (Haruvy, Roth, and Ünver 2006). It demonstrates that the social norm that coerces students to immediately accept an offer is causal for unraveling and, therefore, for lower quality matches. In the experiment, the norm was built into the design by not allowing students to reject offers of a particular period. In a control treatment where all offers could be rejected by students, significantly more matches are formed late through the clearinghouse when all information about the ability of the students had been revealed.

Uncertainty about who is on the long or on the short side of the market can cause the unraveling of markets as shown in a theoretical contribution by Li and Rosen (1998). By contracting early the agents can avoid the payoff risk that arises because of incomplete information regarding the productivity of each agent. Motivated by the match for gastroenterology fellowships, the market participants' uncertainty about the size of supply and demand was identified as a source of unraveling, again based on market data and data from controlled lab experiments (McKinney, Niederle, and Roth 2005). Niederle, Roth, and Ünver (2013) show the relevance of the uncertainty argument and the irrelevance of a shortage of workers per se for unraveling. They experimentally vary supply and demand in a market where each firm wants to hire one worker and where all agents agree on the quality of firms and workers (apart from a small idiosyncratic term in the utility function to break indifferences). They compare a situation with an 'excess supply' of high-quality workers to a situation where high-quality workers are on the short side (called 'comparable demand and supply' since there are many workers in total) and also to 'excess demand' where there are more firms than workers. In line with the subgame equilibrium predictions, they observe that the low-quality firms match early in the situation with comparable demand and supply when there exists a shortage of high-quality workers and quality is not known. In contrast, almost no early matches are formed in the excess supply and excess demand conditions.

Workers and firms in the model by Li and Rosen (1998) and the experiment by Niederle, Roth, and Ünver (2013) have symmetric information regarding the quality of workers: In the early stage, the workers' productivity is unknown with given probabilities, and in the later stage the productivity is

revealed. Thus, in the early stage, firms and workers base their decisions on the expected productivity of a worker. A situation where workers have superior information about their productivity relative to firms is studied by Dargnies, Hakimov, and Kübler (forthcoming). The aim of this study is to investigate the effect of the self-confidence of workers on market unraveling. The workers' productivity is determined by ranking workers according to their performance in a real-effort task. The exact productivity rank of the worker is revealed only in the later stage of the market, where the assortative matching among all remaining firms and workers is generated. Firms can make early offers and, if a worker accepts an offer, both leave the market. At this early stage, the workers' relative self-assessments regarding their performance in the real-effort task matter for their decision to accept or reject an offer. Two treatments are employed, one with a hard and one with an easy real-effort task that are predicted to generate relative under- and overconfidence, respectively. There is a significant treatment difference with respect to acceptances of early offers, implying more blocking pairs and lower efficiency in the treatment with the hard task and less confident workers. The beliefs of workers drive this difference in unraveling and stability, since, as predicted by psychological studies, participants display underconfidence regarding their relative performance after working on the hard task while they show overconfidence regarding their relative performance after the easy task.

Unraveling of markets has many different causes. This section shows that experiments have been useful for identifying some of them, in particular unstable clearinghouses, social norms governing the reneging of offers, and uncertainty about the quality of the workers, and miscalibrated beliefs of workers about their own quality.

Market rules	Study	Market size	Information	Special design features	Truth-telling	Efficiency/Stability	Distribution of payoffs
Static DA	Harrison, McCabe 1996	3+3, 4+4	Full	Some agents computerized (tell the truth) in 3+3, all students computerized in 4+4	[not available]	Efficiency lower, the more human agents; no blocking pairs in markets with computerized students and markets up to two human players; 16% blocking pairs in all other treatments	3+3 players: outcomes preferred by receivers; 4+4 players: outcomes preferred by proposers
	Pais, Pintér, Veszteg 2011	5+3 (2 schools with 2 seats)	No, partial, and full	Comparison with BOS and TTC	Negatively correlated with information (proposers and receivers)	Stable [Pareto efficient] outcomes: 38% [100%] in no, 48% [90%] in partial, 28% [96%] in full information	[not available]
	Castillo, Dianat 2016	4+4	Full	Proposers computerized; manipulations of responders restricted to truncations	Responders: 56% with truncations increasing in profitability and decreasing in riskiness	Stable outcomes (only over-truncation creates instability): 88%	Receiver [proposer]-optimal stable matching: 24% [21%]; payoffs closer to receiver-optimal matching
	Featherstone, Mayefsky 2015	5+5	Full	Proposers computerized; markets with unique or multiple stable matchings	Responders: 66% [57%] with unique[multiple] stable matching[s]	.47 [.49] blocking pairs per period with unique[multiple] stable matching[s]	[not available]
Dynamic DA	Haruvy, Ünver 2007	4+4	Low (own preferences) or high (all preferences)	Repeated game (every round payoff-relevant); receivers computerized [comp] or human	High for proposers ("going-down-the-list") and receivers	Stable outcome in at least 6 out of 10 last rounds (of 30): Low+human [comp]: 96% [67%] high+human [comp]: 79% [80%]	Receiver-optimal stable matching is rare
	Echenique, Wilson, Yariv 2016	8+8	Full	Dynamic implementation of DA protocol	Proposers skip entries of true ROL; receivers accept best available offer	Stable outcomes: around 50% of markets	Only 28% of stable outcomes are proposer-optimal
	Castillo, Dianat 2017	8+8	Full, no info, incomplete-prop./resp.	Dynamic implementation of DA protocol	First offer to first choice: 46% in full info, 77% in no info, 46% in incompl-prop., 58% in incompl-resp.	Stable outcomes: 55% in full info, 50% in no info, 25% in incompl.-prop., 54% in incomplete-resp.	[not available]
De-centralized Matching	Nalbantian, Schotter 1995	3+3	Own pref. and distrib. of others' preferences	All agents can call each other up and agree on wage (comparison with centralized Demange/Gale mech.)	[does not apply]	94.8% of available surplus captured; some suboptimal matches but all agents matched	[not available]
	Echenique, Yariv 2012	8+8	Full	Up to one offer can be made and one accepted by each agent; market ends when inactive for 30 seconds	[does not apply]	Stable outcomes: 85%	With three stable matchings, median matching most frequent
	Pais, Pintér, Veszteg 2012	5+5	Full and no info	Search frictions: cost of making offers or no reneging of accepted offers	Without search frictions offers follow preferences; skipping when offers are costly	Stability without frictions: 82% with no inf. and 69% with full inf.; stability lower when offers costly and no reneging; faster convergence without reneging	With three stable outcomes, median matching most frequent

Table 1. Two-sided experiments and the market rules

Notes: Full information means that all market participants know the preferences of all agents. Partial information means that agents get some information about the preferences of other agents. No information stands for agents knowing only their own preferences and the capacity of schools. Incomplete-proposer [-responder] means that proposers know their own preferences and the preferences of all proposers [responders] while responders have full information.

5 Conclusions

The purpose and style of matching experiments has changed over time. Many of the early experiments were tests of the theory. Horse races between different house allocation or school choice mechanisms were conducted. Another strand of the early literature deals with questions directly related to existing markets, such as the causes of market unraveling. Many recent studies deal with systematic biases in behavior that matter in matching markets, such as bounded rationality, biased self-assessments, etc. Moreover, recent work also focuses on the question of how the exact implementation of a mechanism, e.g., static versus dynamic, with or without advice, affects market outcomes. Thus, the matching literature has started to establish behavioral regularities that can be taken up by the theory, for example, suggesting the choice of solution concepts or the use of simple mechanisms.

Matching experiments are not only of interest for researchers working in the area of matching markets, but they also shed light on questions that are relevant for mechanism design in general. For example, the robust finding that people often do not play their dominant strategy under DA challenges the concept of implementation in dominant strategies. It is important to understand the causes of this behavioral regularity and the conditions under which dominant strategies are actually chosen. We believe that theories addressing the behavioral regularities will prove useful in systematizing these findings.³⁴ At the same time, market design is needed to come up with mechanisms that are simpler in the sense that more people understand that truth-telling is optimal.

This survey showcases research that has predominantly been conducted over the past 20 years. We believe that the dialogue between experimentalists, theorists, and practitioners in the domain of matching markets has been a fruitful one. In the past years, empirical work has entered the scene, often confirming the external validity of experimental findings. At the same time, we believe that there are many open questions that deserve closer scrutiny. A more rigorous consideration of certain established behavioral biases in the context of matching seems fruitful. For example, it could be investigated whether the district-school bias is due to loss aversion or an endowment effect. Also, self-deception and wishful thinking could play a role when making choices between programs. Moreover, the environment is changing constantly with new requirements and opportunities. For example, many clearinghouses have started to rely on more frequent interactions with the participants during the matching procedure through online systems. This opens up new possibilities, in particular with respect to the transparency of the matching procedure, but also creates new challenges. Finally, certain important applications, such as kidney exchange, have not yet been in the focus of experimental

³⁴ An example of such a theory is obvious strategy-proofness (Li 2017).

economists. The interaction between hospitals, doctors, and patients could be an interesting field for future research.

References

Abdulkadiroğlu, Atila, Parag A. Pathak, Alvin E. Roth, and Tayfun Sönmez. “The Boston public school match.” *American Economic Review* 95, no. 2 (2005): 368–371.

Abdulkadiroğlu, Atila; Che, Yeon-Koo and Yasuda, Yosuke. “Resolving conflicting preferences in school choice: The ‘Boston mechanism’ reconsidered.” *American Economic Review* 101.1 (2011): 399–410.

Abdulkadiroğlu, Atila and Sönmez, Tayfun. “House allocation with existing tenants.” *Journal of Economic Theory*, 1999, 88(2), pp.233–260.

Abdulkadiroğlu, Atila and Sönmez, Tayfun. “School choice: A mechanism design approach.” *American Economic Review*, 2003, 93(3), pp.729–747.

Alcalde, José and Barberà, Salvador. “Top dominance and the possibility of strategy-proof stable solutions to matching problems.” *Economic Theory*, 1994, 4(3), pp 417–435.

Basteck, Christian; Klaus, Bettina and Kübler, Dorothea. “How lotteries in school choice help to level the playing field,” 2018. Working paper.

Basteck, Christian and Mantovani, Marco. “Cognitive ability and games of school choice.” *Games and Economic Behavior*, 2018, 109, pp. 156–183.

Basteck, Christian and Mantovani, Marco. “Aiding applicants: Leveling the playing field within the immediate acceptance mechanism.” 2018b. Working paper.

Bó, Inacio and Hakimov, Rustamdjan. “Iterative versus standard deferred acceptance: Experimental evidence.” *The Economic Journal*, forthcoming

Bogomolnaia, Anna and Moulin, Hervé. "A new solution to the random assignment problem." *Journal of Economic Theory*, 2001, 100(2), pp. 295–328.

Braun, Sebastian; Dwenger, Nadja; Kübler, Dorothea and Westkamp, Alexander. "Implementing quotas in university admissions: An experimental analysis." *Games and Economic Behavior*, 2014, 85, pp.232–251.

Calsamiglia, Caterina; Haeringer, Guillaume and Klijn, Flip. "Constrained school choice: An experimental study." *American Economic Review*, 2010, 100(4), pp.1860–1874.

Calsamiglia, Caterina; Haeringer, Guillaume and Klijn, Flip. "A comment on "School choice: An experimental study." [J. Econ. Theory 127 (1) (2006) 202–231]. *Journal of Economic Theory*, 2011, 146(1), pp.392-396.

Castillo, Marco and Dianat, Ahrash. "Information effects in the dynamic Gale-Shapley mechanism: An experimental study." 2017. Working paper.

Castillo, Marco and Dianat, Ahrash. "Truncation strategies in two-sided matching markets: Theory and experiment." *Games and Economic Behavior*, 2016, 98, pp.180–196.

Chen, Li, and Juan Pereyra. "Self-selection in school choice." 2017. Working paper.

Chen, Yan; Jiang, Ming; Kesten, Onur; Robin, Stéphane and Zhu, Min. "Matching in the large: An experimental study." 2017. Working paper.

Chen, Yan and Sönmez, Tayfun. "Improving efficiency of on-campus housing: An experimental study." *American Economic Review*, 2002, 92(5), pp.1669–1686.

Chen, Yan and Sönmez, Tayfun. "An experimental study of house allocation mechanisms." *Economics Letters*, 2004, 83(1), pp.137–140.

Chen, Yan and Sönmez, Tayfun. "School choice: An experimental study." *Journal of Economic Theory*, 2006, 127(1), pp.202–231.

Chen, Yan; Jiang, Ming and Kesten, Onur. "Chinese college admissions reforms: Experimental and empirical evaluations." 2017. Working paper.

Chen, Yan and Kesten, Onur. “Chinese college admissions and school choice reforms: An experimental study.” 2016. Working paper.

Chen, Yan, and Kesten, Onur. “Chinese college admissions and school choice reforms: A theoretical analysis.” *Journal of Political Economy*, 2017, 125(1), pp. 99–139.

Chen, Yan and He, Yinghua. “Information Acquisition and Provision in School Choice: An Experimental Study.” 2017. Working paper.

Dargnies, Marie-Pierre; Hakimov, Rustamdjan and Kübler, Dorothea. “Self-confidence and unraveling in matching markets.” *Management Science*, forthcoming.

Demange, Gabrielle and Gale, David. “The strategy structure of two-sided matching markets.” *Econometrica*, 1985, 53(4), pp.873–888.

Ding, Tingting and Schotter, Andrew. “Learning and mechanism design: An experimental test of school matching mechanisms with intergenerational advice.” *New York University Manuscript*, 2015.

Ding, Tingting and Schotter, Andrew. “Matching and chatting: An experimental study of the impact of network communication on school-matching mechanisms.” *Games and Economic Behavior*, 2017, 103, pp.94–115.

Dur, Umut; Hammond, Robert G. and Kesten, Onur. “Sequential School Choice: Theory and Evidence from the Field and Lab.” 2018a. Working paper.

Dur, Umut; Hammond, Robert G. and Morrill, Thayer. “Identifying the Harm of Manipulable School-Choice Mechanisms.” *American Economic Journal: Economic Policy*, 2018b, 10(1), pp. 187–213.

Dur, Umut, Robert G. Hammond, and Thayer Morrill. “The secure Boston mechanism: Theory and experiments.” *Experimental Economics* (2015): 1–36.

Echenique, Federico; Wilson, Alistair. J. and Yariv, Leeat. “Clearinghouses for two-sided matching: An experimental study.” *Quantitative Economics*, 2016, 7, 449–482.

Echenique, Federico and Yariv, Leeat. "An experimental study of decentralized matching." 2012. Working paper.

Ergin, Haluk, and Tayfun Sönmez. "Games of school choice under the Boston mechanism." *Journal of Public Economics* 90.1-2 (2006): 215–237.

Featherstone, Clayton R. and Mayefsky, Eric. "Why do some clearinghouses yield stable outcomes? Experimental evidence on out-of-equilibrium truth-telling." 2015. Working paper.

Featherstone, Clayton R. and Niederle, Muriel. "Boston versus deferred acceptance in an interim setting: An experimental investigation." *Games and Economic Behavior*, 2016, 100, pp.353–375.

Gale, David, and Shapley, Lloyd S. "College admissions and the stability of marriage." *The American Mathematical Monthly*, 1962, 69(1), pp.9–15.

Gong, Binglin and Liang, Yingzhi. "A dynamic college admission mechanism in Inner Mongolia: Theory and experiment." 2016. Working paper.

Guillen, Pablo, and Hakimov, Rustamdjan. "The Effectiveness of top-down advice in strategy-proof mechanisms: A field experiment." *European Economic Review*, 2018, 101, pp. 505–511.

Guillen, Pablo, and Hakimov, Rustamdjan. "Not quite the best response: truth-telling, strategy-proof matching, and the manipulation of others." *Experimental Economics*, 2017, 20(3), pp.670–686.

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

Guillen, Pablo, and Kesten, Onur. "Matching markets with mixed ownership: The case for a real-life assignment mechanism." *International Economic Review*, 2012, 53(3), pp.1027–1046.

Haeringer, Guillaume. "Market Design. Auctions and Matching." MIT Press, 2017.

Hafalir, Isa E.; Yenmez, M. Bumin and Yildirim, Muhammed A. "Effective affirmative action in school choice." *Theoretical Economics* 8.2 (2013): 325–363.

Hakimov, Rustamdjan; Heller, C-Philipp; Kübler, Dorothea, and Kurino; Morimitsu. “How to avoid black markets for appointments at public offices.” 2018

Hakimov, Rustamdjan, and Onur Kesten. “The equitable top trading cycles mechanism for school choice.” *International Economic Review*, 2018, 59(4), pp.2219-2258.

Harrison, Glenn W. and McCabe, Kevin A. “Stability and preference distortion in resource matching: An experimental study of the marriage problem,” in R. Mark Isaac, ed., *Research in experimental economics*, Vol. 6. Greenwich, CT: JAI Press, 1996, pp. 53–129.

Haruvy, Ernan and Ünver, M. Utku. “Equilibrium selection and the role of information in repeated matching markets.” *Economics Letters*, 2007, 94(2), pp.284–289.

Haruvy, Ernan; Roth, Alvin E. and Ünver, M. Utku. “The dynamics of law clerk matching: An experimental and computational investigation of proposals for reform of the market.” *Journal of Economic Dynamics and Control*, 2006, 30(3), pp.457–486.

Hassidim, Avinatan; Romm, Assaf and Shorrer, Ran I. “Strategic behavior in a strategy-proof environment.” *Proceedings of the 2016 ACM Conference on Economics and Computation*. ACM, 2016.

He, Yinghua and Magnac, Thierry. “A Pigouvian Approach to Congestion in Matching Markets.” 2017. Working Paper.

Hugh-Jones, David; Kurino, Morimitsu and Vanberg, Christoph. “An experimental study on the incentives of the probabilistic serial mechanism.” *Games and Economic Behavior*, 2014, 87, pp.367–380.

Kagel, John H., and Roth, Alvin E. “The dynamics of reorganization in matching markets: A laboratory experiment motivated by a natural experiment.” *The Quarterly Journal of Economics*, 2000, 115(1), pp.201–235.

Kagel, John H. and Roth, Alvin E. (editors) “*Handbook of Experimental Economics*”, Volume 2, Princeton University Press, 2016

Kawagoe, Toshiji; Matsubae, Taisuke and Takizawa, Hirokazu. “The skipping-down strategy and stability in school choice problems with affirmative action: Theory and experiment.” *Games and Economic Behavior*, 2018, 109, pp. 212–239.

Klijn, Flip; Pais, Joana and Vorsatz, Marc. “Affirmative action through minority reserves: An experimental study on school choice.” *Economics Letters*, 2016a, 139, pp.72–75.

Klijn, Flip; Pais, Joana and Vorsatz, Marc. “Static versus dynamic deferred acceptance in school choice: A laboratory experiment.” *Games and Economic Behavior*, forthcoming.

Klijn, Flip; Pais, Joana and Vorsatz, Marc. “Preference intensities and risk aversion in school choice: A laboratory experiment.” *Experimental Economics*, 2013, 16(1), pp.1–22.

Kojima, Fuhito. “School choice: Impossibilities for affirmative action.” *Games and Economic Behavior* 75.2 (2012): 685–693.

Koutout, Kristine; Van der Linden, Martin; Dustan, Andrew and Wooders, Myrna. „Mechanism performance under strategy advice and sub-optimal play: A school-choice experiment.” 2018. Working paper

Lien, Jaimie W.; Zheng, Jie and Zhong, Xiaohan. “Preferences submission timing in school choice matching: Testing fairness and efficiency in the laboratory.” *Experimental Economics*, 2016, 19, pp.116–150.

Li, Shengwu. “Obviously strategy-proof mechanisms.” *American Economic Review*, 2017, 107(11), pp. 3257-87.

Li, Hao and Rosen, Sherwin. “Unraveling in matching markets.” *American Economic Review*, 1998, 88(3), pp. 371-387.

Matsubae, Taisuke. “Do the minority students benefit from affirmative action policy in school choice markets?: Reconsidered.” 2011. Working paper.

McKinney, C. Nicholas; Niederle, Muriel and Roth, Alvin E. “The collapse of a medical labor clearinghouse (and why such failures are rare).” *American Economic Review*, 2005, 95(3), pp. 878–889.

Mullin, Charles H. and Reiley, David H. “Recombinant estimation for normal-form games with applications to auctions and bargaining,” *Games and Economic Behavior*, 2006, 54, pp. 159–182.

Nalbantian, Haig R. and Schotter, Andrew. “Matching and Efficiency in the Baseball Free-Agent System: An Experimental Examination.” *Journal of Labor Economics*, 1995, 13(1), pp. 1–31.

Niederle, Muriel and Roth, Alvin E. “Market culture: how rules governing exploding offers affect market performance.” *American Economic Journal: Microeconomics*, 2009, 1(2), pp.199–219.

Niederle, Muriel; Roth, Alvin E. and Ünver, M. Utku. “Unraveling results from comparable demand and supply: An experimental investigation.” *Games*, 2013, 4(2), pp.243–282.

Olson, Mark and Porter, David. “An experimental examination into the design of decentralized methods to solve the assignment problem with and without money.” *Economic Theory*, 1994, 4(1), pp.11–40.

Pais, Joana; Pintér, Agnes and Veszteg, Robert F. “College admissions and the role of information: An experimental study.” *International Economic Review*, 2011, 52(3), pp.713–737.

Pais, Joana; Pintér, Agnes and Veszteg, Robert F. “Decentralized matching markets: A laboratory experiment,” Working Papers Department of Economics 2012/08, ISEG – School of Economics and Management, Department of Economics, University of Lisbon. 2012.

Pais, Joana, and Pintér, Agnes. “School choice and information: An experimental study on matching mechanisms.” *Games and Economic Behavior*, 2008, 64(1), pp.303–328.

Pan, Siqi. “The instability of matching with overconfident agents: Laboratory and field investigations.” *Games and Economic Behavior*, forthcoming.

Rees-Jones, Alex. “Suboptimal behavior in strategy-proof mechanisms: Evidence from the residency match.” *Games and Economic Behavior*, 2018, 108, pp. 317–330.

Rees-Jones, Alex and Samuel Skowronek. “An experimental investigation of preference misrepresentation in the residency match.” *Proceedings of the National Academy of Sciences* 115.45 (2018): 11471–11476.

Roth, Alvin E. “Experiments in Market Design,” Chapter 5, in: *Handbook of Experimental Economics* (volume 2), Kagel, John H. and Roth, Alvin E (eds.), 2015, pp.290–346.

Roth, Alvin E. "The economics of matching: Stability and incentives." *Mathematics of operations research*, 1982, 7(4), pp.617–628.

Roth, Alvin E. "The evolution of the labor market for medical interns and residents: a case study in game theory." *Journal of Political Economy*, 1984, vol. 92, pp. 991–1016.

Roth, Alvin E., and Elliott Peranson. "The redesign of the matching market for American physicians: Some engineering aspects of economic design." *American economic review* 89.4 (1999): 748–780.

Roth, Alvin E. and Rothblum, Uriel G. "Truncation strategies in matching markets—in search of advice for participants." *Econometrica*, 1999, 67(1), pp.21–43.

Roth, Alvin E. and Sotomayor, Marilda A. "Two-sided matching: A study in game-theoretic modeling and analysis." *Econometric Society Monograph Series*, Cambridge University Press, 1990.

Stephenson, D. G. "Continuous feedback in school choice mechanisms." 2016. Working paper.

Schmelzer, André. "Single versus multiple randomization in matching mechanisms." 2016. Working paper.

Schmelzer, André. "Strategy-proofness of stochastic assignment mechanisms." *Journal of Mechanism and Institution Design* 3(1), 2018.

Westkamp, Alexander. "An analysis of the German university admissions system." *Economic Theory*, 2013, 53, pp. 561–589.

Zhu, Min. "Experience transmission: Truth-telling adoption in matching." 2015. Working paper.