Inefficiency-Manipulability Tradeoff in the Parallel Mechanism

Jerry Anunrojwong [*][†]

October 23, 2019

Abstract

Most school choice and other matching mechanisms are based on deferred acceptance (DA) for its incentive properties. However, non-strategyproof mechanisms can dominate DA in welfare because manipulation in preference rankings can reflect the intensities of underlying cardinal preferences. In this work, we use the parallel mechanism of Chen and Kesten, which interpolates between Boston mechanism and DA, to quantify this tradeoff. While it is previously known that mechanisms that are closer to Boston mechanism are more manipulable, we show that they are also more efficient in student welfare if school priorities are weak. Our theoretical results show the inefficiency-manipulability tradeoff in the worst case, while our simulation results show the same tradeoff in the typical case.

1 Introduction

Most school choice and other matching mechanisms are based on the student-proposing deferred acceptance (DA) algorithm (Gale and Shapley, 1962; Abdulkadiroğlu and Sönmez, 2003). These mechanisms are adopted because they have two attractive properties on incentives and welfare. First, on incentives, it is strategyproof for students, which makes it simple for students to rank schools. This good incentive property of DA motivated the Boston Public School System to adopt DA in place of Boston (Immediate Acceptance) Mechanism. Second, on welfare, the resulting matching is stable with respect to to student preferences and school priorities; this means that if a student prefers some school to her assigned school, she must have a lower priority than every student assigned to that school. It follows that the student-proposing DA is Pareto efficient, but this efficiency guarantee only holds if schools have strict priorities over students, and this assumption almost never holds. Rather, schools often have a few broad criteria to rank students, such as sibling attendance and walk zones, and break ties randomly within the same priority class to produce strict priorities required by DA. When schools have weak priorities, DA still has guarantees on incentives, but not on efficiency.

In this paper, we show that between Boston mechanism and DA, there is a *quantitative* tradeoff between *inefficiency* (which is a welfare property) and *manipulability* (which is an incentives

[*]Massachusetts Institute of Technology and Chulalongkorn University, `jerryanunroj@gmail.com`

[†]Supported in part by NSF Award 1841550, "EAGER: Algorithmic Approaches for Developing Markets."

property). *More manipulable mechanisms are more efficient in student welfare.* We assume that each student has a fixed cardinal preference over each school, but she is only allowed to report her ordinal preferences to the mechanism. To show this tradeoff, we analyze a family of matching mechanisms called the *parallel mechanism* formalized by Chen and Kesten (Chen and Kesten, 2017).[1] This family of mechanisms has one parameter d and interpolates between Boston mechanism ($d = 1$) and DA (d = number of schools).[2] Moreover, Chen and Kesten (2017) showed that if $d < d'$, then the d-parallel mechanism was "more manipulable" than the d'-parallel mechanism in the sense of Pathak and Sönmez (2013). Earlier works qualitatively compared the non-strategyproof Boston mechanism with the strategyproof DA (Abdulkadiroğlu et al., 2011, 2015), which were two extreme mechanisms of this entire family. Our analysis of the parallel mechanism gives a more fine-grained handle on efficiency. We quantify the inefficiency of the mechanism with respect to d and show that the mechanism with lower d (hence more manipulable) is also more efficient.

We focus on the parallel mechanism for three reasons. First, every mechanism in the parallel mechanism family is comparable in terms of manipulability, and the ranking is neatly captured in terms of the parameter d. Lower d means more manipulability. Efficiency guarantees for the parallel mechanism that are stronger for lower d can then be interpreted that higher manipulability and higher efficiency go hand in hand (with lower d). Since the comparison of manipulability is qualitative,[3] we cannot illustrate the inefficiency-manipulability tradeoff for any pair of mechanisms because they are generally not comparable. The choice of parallel mechanism allows us to sidestep this issue. Second, the parallel mechanism elegantly interpolates between two well-known and widely used mechanisms, Deferred Acceptance (DA) mechanism and Boston (Immediate Acceptance) Mechanism by combining the deferred acceptance and the immediate acceptance features. The analysis of the parallel mechanism therefore sheds light on the effects of deferred versus immediate acceptance features, which are central to market and mechanism design problems in other contexts. Third, the parallel mechanism itself is used in practical settings, most notably in university admissions in China. Despite its common use in practice, only Chen and Kesten (2017) has analyzed this mechanism. Our work can be viewed as a theoretical and empirical welfare analysis of this understudied mechanism in its own right.

Since all mechanisms in the parallel mechanism family except DA are not strategyproof, students can misreport their preferences to the mechanism. We are interested in evaluating student welfare assuming that students play a *Nash equilibrium* of the induced preference revelation game. This is in contrast with some previous works in economics that analyze the resulting matching assuming truthful play even in non-strategyproof mechanisms.[4] Empirical evidence suggests that

[1] Note that it is not true in general that more manipulability corresponds to more efficiency. For example, our results do not imply that the student-proposing DA mechanism is more efficient than the school-proposing DA, and we could construct an arbitrarily bad but manipulable mechanism that have poor efficiency properties. However, we believe that the tradeoff holds for reasonable "averages" of Boston and DA that are used in practice. For example, an interested reader can check that the proof of the theoretical guarantees in Section 3 still holds if the school assignment after the first round (of d schools) is arbitrary. So the exact form of the parallel mechanism is not that important. We commit to a specific form of the parallel mechanism described in Subsection 2.7.3 for expositional simplicity.

[2] If the number of schools is not known, DA corresponds to $d = \infty$. Chen and Kesten (2017) called this parameter "permanency-execution period" and denoted it by e, but we changed it to d to avoid confusion with the constant $e \approx 2.718$.

[3] We formally discuss how we compare manipulability across mechanisms in subsection 2.8.

[4] For example, Chen and Kesten (2017) compares stability of matching in the parallel mechanism relative to reported preferences, not true preferences.

students do behave strategically and misreport their preferences in non-strategyproof mechanisms (Abdulkadiroğlu and Sönmez, 2003; He, 2017; Calsamigliay et al., 2017), so analyzing welfare at Nash equilibrium is more justified than analyzing welfare assuming truthful reports.

The following argument shows the intuition behind our inefficiency-manipulability tradeoff result. DA does not have good welfare properties when school priorities are weak because the best each student can do is to truthfully report ordinal preferences and the mechanism has no way to respond to intensities of underlying cardinal preferences. Strategyproofness of DA therefore directly contributes to its inefficiency. When money cannot be used to calibrate students' valuations, observing how students make tradeoffs between choices is an alternative. For example, if a student is willing to give up a high probability opportunity of getting school A for a low probability opportunity of getting school B, we can infer that this student values school B a lot more than school A, and vice versa. When students misreport their preferences, they compare various possible choices of rankings under the random draws of school priorities. The result of this manipulation transmits some cardinal information to the mechanism in equilibrium even when students only report ordinal rankings. The more manipulable a mechanism, the more opportunities there are of transmitting cardinal preference information in this way. In contrast, under a strategyproof mechanism such as DA, the mechanism has no way of knowing and responding to such cardinal information. In this example, a student will only report that she prefers B over A, but the mechanism does not know how much.

Our preferred notion of efficiency is approximate Pareto efficiency, proposed by Immorlica, Lucier, Weyl and Mollner (Immorlica et al., 2017). A mechanism is γ-approximately efficient ($\gamma \geq 1$) if for any allocation of the mechanism, there is no alternative allocation such that every student is at least γ times better off. This notion of efficiency is appropriate in our setting because (1) in the spirit of nontransferable utility setting, it does not compare utilities across students, and (2) it is "fair" in a Rawlsian veil-of-ignorance sense in that an improvement only counts if it helps everyone including the least fortunate members of society. Note that even though student preferences are assumed to be fixed and known, so the students play a complete information game, weak priorities mean that student assignment is random because it depends on the random draws of school priorities. The relevant welfare metric is not the ex post welfare with respect to strict priorities, as implicitly assumed in earlier work, but ex ante welfare, averaging over random draws. Henceforth, whenever we talk about student welfare, we always mean ex ante welfare in this sense. For more discussion about this notion of efficiency, see subsection 2.5.

We assume that schools have weak priorities. This crucial assumption drives our results, so we discuss it here. We assume that students have explicit cardinal preferences over schools and are strategic, while schools have weak priorities and are non-strategic. This reflects the reality of school assignment systems in most places, where school priorities are determined by regulations and are known to all. It also implicitly prioritizes student welfare over school welfare, because weak priorities mean schools "do not care as much" about which students they accept. Our approximate Pareto efficiency welfare metric reflects this: it only concerns about each student's welfare and school welfare is not even defined in our model. In our theoretical results, we assume that schools have no priorities. This might seem like an overly restrictive assumption, but it is theoretically reasonable considering our approximate Pareto efficiency welfare metric only cares about student welfare, so schools should not have intrinsic priorities or else school welfare needs to be accounted for as well. This reasoning underlies a technical reason for assuming no priorities: because we have to. Approximate Pareto efficiency *does not hold* in general with weak priorities.

More detailed discussions can be found in subsection 2.6.

In the theoretical results section, we prove that when students are a continuum and schools have no priorities, the d-parallel mechanism is $(d + 1)$-approximately Pareto efficient. We also show an example that the factor has to be at least $d/2$, so linear dependence on d is necessary. When students are discrete, an analogous proof gives a slightly weaker $2d + 1$ bound. The main idea of the proof is to formalize how students trade off a probabilistic claim on one school versus another as outlined in the introduction.

In the simulations section, we give an algorithm to compute the Nash equilibrium of the parallel mechanism given student preferences and school priorities. We also show how to compute the Pareto welfare loss of a given equilibrium. We find that efficiency loss increases with d.

The rest of the paper proceeds as follows. The next subsection discusses related work. Section 2 describes our model. Section 3 presents theoretical results. Section 4 presents simulation results. Section 5 concludes.

1.1 Related Work

After Abdulkadiroğlu and Sönmez (2003) formulated school choice as a mechanism design problem, most works analyzed variants of Boston mechanism and DA under complete information and strict priorities. Abdulkadiroğlu et al. (2009); Kesten and Ünver (2010); Kesten and Kurino (2017) investigated relationships between efficiency and strategyproofness. Ergin and Sönmez (2006); Kojima (2008) analyzed structural properties of the Nash equilibria of Boston mechanism and argued that DA has better welfare than Boston mechanism. Miralles (2009); Erdil and Ergin (2008); Abdulkadiroğlu et al. (2011, 2015); Kesten and Ünver (2015) showed that the previous conclusion depends critically on strict priorities assumption which is often violated; under weak priorities, DA incurs substantial cost in student welfare. Harless (2014) proposed a slight variant of the Boston mechanism to improve its incentives property, while Abdulkadiroğlu et al. (2015) proposed a slight variant of DA to improve its efficiency. Both mechanisms do not generate a family of interpolating mechanisms, so they are not suitable for our exercise.

We focus primarily on welfare only on one side of the market (student welfare), and in the theoretical results assume that schools have no priorities. These are standard implicit assumptions in the substantial literature on one-sided matching or object allocation without money (Bhalgat et al., 2011; Filos-Ratsikas et al., 2014; Adamczyk et al., 2014; Christodoulou et al., 2016; Aziz et al., 2016). Similar to present work, Christodoulou et al. (2016) studied the equilibrium behavior of non-truthful mechanisms, but with respect to the utilitarian welfare objective. Other standard mechanisms for object allocation are the top trading cycles (TTC) algorithm (proposed by Shapley and Scarf (1974) and extended to school choice settings by Abdulkadiroğlu and Sönmez (2003); Pycia and Ünver (2017)), random serial dictatorship (RSD), and probabilistic serial algorithm (proposed by Bogomolnaia and Moulin (2001) and shown to be equivalent to DA and RSD in the continuum model with no priorities by Che and Kojima (2010)).

We compare manipulability across mechanisms using the notion from Pathak and Sönmez (2013). Mennle and Seuken (2018) proposed another relaxation of strategyproofness, where a mechanism is partially strategyproof if truthful reporting is a dominant strategy for agents whose preference intensities differ sufficiently between any two objects. This definition is less appropriate in our setting because we want our welfare guarantee to hold for every student preference profile.

Even though students have cardinal preferences, we focus on ordinal mechanisms that elicit only ordinal preferences from students. In contrast, cardinal mechanisms elicit cardinal preferences. The most well-known such mechanism is the pseudo-market mechanism of Hylland and Zeckhauser (1979), which was extended to settings with priorities by He et al. (2018). However, in settings without money, eliciting cardinal preferences is often difficult, impractical, and not robust (Carroll, 2018; Huesmann and Wambach, 2016; Ehlers et al., 2016), so this work focuses on ordinal mechanisms.

Mennle and Seuken (2017) also studied the tradeoff between strategyproofness and efficiency in assignment problems. Unlike our work, they used ordinal dominance under truthful reports to capture efficiency. They interpolated between two mechanisms by a convex combination, viewing each mechanism as a lottery; they called this a hybrid mechanism. Analyzing hybrid mechanisms in our framework (approximate Pareto efficiency under Nash equilibrium play) is an interesting research question, but is beyond the scope of this paper.

2 Model

2.1 Students and Schools

There is a (finite or infinite) set of students, denoted by $\mathcal{S}$, and a finite set of schools, denoted by $C = \{1, 2, \ldots, n\}$. There are n schools. For each school $j \in C$, let $c_j \in \mathbb{N}$ be the capacity of j (the number of seats at school j).

Each student $i \in \mathcal{S}$ has a type θ that specifies her value v^i_j for a single seat at school j. Students are unit-demand, that is, student i's value for a subset $C \subseteq C$ is $\max_{j \in C} v^i_j$. We assume that all values are nonnegative: $v^i_j \in [0, \infty)$. Write $\Theta = [0, \infty)^n$ for the space of types. Students are strategic and can misreport their preferences. Schools are non-strategic and the way school priorities are formed is common knowledge. Throughout this work, we assume that priorities are weak, and if the matching mechanism demands strict priorities, ties are broken randomly within tiers.

When all students are different, every student has her own type. The notion of types is most useful in the continuum model when a mass of students is identical, and only the assignment probability of that type matters.

While we use the language of students and schools, our setting applies more broadly to the problem of allocating indivisible heterogeneous goods without money. Example domains include allocating affordable apartments to tenants and allocating kidneys to patients. Apartments have "priorities" over tenants as determined by law, and kidneys have "priorities" over patients as determined by medical compatibility. In these domains, social good is the main objective, and the use of money is repugnant.

2.2 Lottery

A randomized assignment, or *lottery*, is a randomized mapping $\sigma : \mathcal{S} \to \Delta(C)$ from students to schools, where $\Delta(C)$ denotes probability distributions over the elements of C. Given a lottery σ, we denote by σ^i_j the probability that student i is matched to school j. A lottery σ is *feasible* if it respects capacities, i.e. for each school j, $\sum_{i \in \mathcal{S}} \sigma^i_j \leq c_j$. The value enjoyed by student i in a feasible lottery σ, written $v^i(\sigma)$, is her expected value from the assigned school: $v^i(\sigma) = \sum_{j \in C} v^i_j \sigma^i_j$.

2.3 Continuum Model

We now describe a continuum of students, which can be viewed as the large market limit in the sense of Azevedo and Leshno (2016).

In the continuum model, the capacity of school j, c_j, can take any non-negative real value (not just integers). The set of students $\mathcal{S}$ is described by a measure ρ over the set of types Θ. We assume that ρ is atomless and Lebesgue integrable.

A lottery can now be described as a mapping from types to distributions over goods $\sigma : \Theta \to \Delta(C)$. We write σ_j^θ for the probability that a student of type θ is matched to school j, and we sometime abuse notation and write σ_j^i to mean σ_j^θ where θ is the type of student i. A lottery is feasible if it respects capacities with respect to the measure ρ over types: $\int_{\theta\in\Theta} \sigma_j^\theta d\rho \le c_j$ for all $j \in C$.

2.4 Nash Equilibrium

We assume that agents play a complete-information Nash Equilibrium. The game is a complete information game because we assume that every student's cardinal preferences for schools are fixed and known to everyone. A Nash equilibrium is a set of reported preferences from all students such that each student's report is optimal given all the other students' reports. More formally, let x_i be student i's report, $\mathbf{x}$ be all students' reports, and $\mathbf{x}_{-i}$ be the reports of all students except i. Let $v^i(\cdot)$ be student i's expected payoff given the reported preferences. Then, $\mathbf{x}$ is a Nash equilibrium if and only if for all students i, $v^i(x_i, \mathbf{x}_{-i}) \ge v^i(\hat{x}_i, \mathbf{x}_{-i})$ for any report $\hat{x}_i$.

2.5 Approximate Pareto Efficiency

We use following notion of approximate Pareto efficiency to quantify the inefficiency of different mechanisms. This definition is central to both our theoretical analysis (section 3) and our simulations (section 4).

Definition 1 (Immorlica et al. (2017))**.** *For $\gamma \ge 1$, a feasible lottery σ is γ-approximately Pareto efficient if there is no other feasible lottery σ' such that $v^i(\sigma') \ge \gamma v^i(\sigma)$ for all students $i \in C$, with strict inequality for a positive measure of students.*

There is a long tradition in mechanism design with money to use utilitarian welfare to represent social welfare. However, utilitarian welfare entails comparing utilities across people, which can be inappropriate in settings where there is no money as numeraire such as ours. Welfare comparisons should be made on each student separately, and this is necessarily a multi-objective problem. Immorlica et al. (2017) proposed the above definition to capture comparisons within all students in a single number in a maximin sense. Given that many matching markets are motivated by social good applications, this maximin notion is fair in that an improvement must make everyone better off; an alternative lottery that makes only some students much better off does not count as an improvement.[5] Moreover, we can check that γ-approximate efficiency in utilitarian welfare implies γ-Pareto efficiency, but no worst-case utilitarian efficiency guarantee is possible because

[5]Maximin is not the only notion of fairness in money-free settings. Another notion that has received attention in the algorithmic game theory community is Nash Social Welfare (the product of utilities).

we can always scale any studentfis utility up arbitrarily and make that student the only important student. Our simulation similar to that of Section 4 shows that utilitarian efficiency loss also generally decreases with d.

The notion of approximate Pareto efficiency extends the notion of Pareto efficiency, and reduces to it when $\gamma = 1$. Our approximation result holds for any equilibrium of the parallel mechanism. In this sense it is a *price of anarchy* result, in which our approximation factor holds in the worst case over student preference profiles and equilibrium selection.

2.6 Weak Priorities and No Priorities

As discussed in the introduction, we assume that schools have weak priorities, and this crucial assumption drives the inefficiency-manipulability tradeoff result. If school priorities are strict, then student-proposing DA is strategyproof for students and produces a student-optimal stable matching, so it performs well both on efficiency and incentives.

School priorities are typically very weak. For example, Boston Public Schools (BPS) prioritize students based only on a small number of coarse features such as sibling attendance and "walk zone" (favoring students whose residences are close to the school), and many students are in the same priority class (Abdulkadiroğlu et al., 2011). Most other public school systems, including New York City (Abdulkadiroğlu et al., 2009), Amsterdam, Chile (Ashlagi and Nikzad, 2017), Chicago and England (Pathak and Sönmez, 2013), also have weak priorities, and ties are broken randomly by lottery. Much of the prior theory in matching assume that both parties have strict preferences, mainly because ties are viewed as a knife-edge phenomenon in economic applications like labor markets. In contrast, weak priorities in school choice are an institutional reality, with important incentives and welfare consequences.

In our theoretical results, we assume that schools have no priorities; that is, schools break ties randomly among students and do not prefer any student in particular. The no priorities assumption is more than just for theoretical tractability; it links our work to the substantial literature on one-sided matching (object assignment) and it is the most theoretically coherent choice. If schools have priorities, even weak ones, then the welfare metric should capture both student welfare and school welfare and how to trade off one against another. But the approximate Pareto efficiency welfare metric only cares about student welfare, so it makes sense to assume that schools do not prefer one student over another so we do not have to worry about school welfare. If schools have priorities, the tradeoff between student welfare and school welfare necessarily involves interpersonal welfare comparison and is context-specific, so we do not propose such a metric in our general model.

The discussion about student versus school welfare is not merely theoretical. Approximate Pareto efficiency fails to hold under weak priorities precisely for that reason, so *the no priorities assumption cannot be relaxed in the worst case.* Consider the following example. Let ϵ be a small number. There is a continuum of students indexed by $[0, 1]$ and there are 2 schools, school 1 and school 2, each with capacity $1/2$. Students in $[0, 1/2]$ have value 1 for school 1, and ϵ for school 2. Students in $[1/2, 1]$ have value ϵ for school 1 and 1 for school 2. School 1 has two priority classes, $[1/2, 1]$ before $[0, 1/2]$ and considers the second class unacceptable. School 1 has two priority classes, $[1/2, 1]$ before $[0, 1/2]$ and considers the second class unacceptable. School 2 has two priority classes, $[0, 1/2]$ before $[1/2, 1]$ and considers the second class unacceptable. Then under any parallel mechanism (or really any matching that respects school priorities), Students in

$[0, 1/2]$ get school 2 with value ϵ and students in $[1/2, 1]$ get school 1 with value ϵ, but a $1/\epsilon$-Pareto improvement is possible by assigning students in $[0, 1/2]$ to school 1 and students in $[1/2, 1]$ to school 2. (The Pareto improvement in our definition is not required to respect school priorities, only capacities.) Since $1/\epsilon$ can be arbitrarily large, there does not exists γ such that the mechanism is γ-approximately Pareto efficient even when the priorities are intuitively as weak as it can be – two equally-sized big priority classes.[6] Intuitively, in our example all the welfare gains go to schools, but school welfare does not count in our welfare metric, so we can improve student welfare by a very large factor transferring welfare from schools to students.

Even though it is too much to expect that the inefficiency-manipulability tradeoff holds always, as argued in the previous paragraph, it is still possible that the tradeoff holds for "typical" student preference profiles and "typical" weak school priorities. Our simulations in Section 4 suggest that this is indeed the case.

2.7 Mechanisms

2.7.1 Boston (Immediate Acceptance) Mechanism

In the Boston mechanism, each student submits a strict preference ranking over schools, and schools have strict priorities over students. Initially all students are unassigned.

Round $t \geq 1$: Consider the remaining unassigned students. For each school j with q_j^t remaining available seats, consider only those students who have listed it as their t-th choice. Those q_j^t students among them with the highest j-priority are assigned to school j.

2.7.2 Deferred Acceptance (DA) Mechanism

In student-proposing Deferred Acceptance mechanism, each student submits a strict preference ranking over schools, and schools have strict priorities over students. Initially all students are unassigned.

Round 1: Each student applies to the top school on her preference list. For each school j, up to c_j applicants who have the highest j-priority are tentatively assigned to school j. The remaining applicants are rejected.

Round $t \geq 2$: Each student rejected from a school at step $t - 1$ applies to the top school on her preference list that has yet to reject her. For each school j, up to c_j students who have the highest j-priority among the new applicants and those tentatively on hold from an earlier step, are tentatively assigned to school j. The remaining applicants are rejected.

[6] While Abdulkadiroğlu et al. (2011) showed that Boston mechanism dominates DA in ex ante welfare with no priorities, Troyan (2012) showed that the domination result failed when schools have weak priorities. Both papers are not technically related to this work, but they suggest a similar conceptual idea that weak priorities can destroy worst-case welfare comparisons that exist with no priorities.

2.7.3 Parallel Mechanism

Now we discuss in detail the parallel mechanism in Chen and Kesten (2017), the central mechanism of our paper. Intuitively, the parallel mechanism does immediate acceptance in blocks of d schools, and deferred acceptance within each block. The block in round t consists of school lists of $(td + 1)$ to $(td + d)$ choice of each unassigned student. $d = 1$ is Boston mechanism, and $d = n$ is DA. Formally, the d-parallel mechanism, denoted $\mathbb{P}_d$, proceeds as follows.

Initially, all students are unassigned. In round $t \geq 0$:

(1) Each unassigned student from the previous round applies to her $(td+1)$-st choice school. Each school j considers its applicants. Those students with the highest j-priority are tentatively assigned to school j up to its capacity. The rest of the applicants are rejected.

(2) Each rejected student, who is yet to apply to her $(td + d)$-choice school, applies to her next choice. If a student has been rejected from all her first $(td + d)$ choices, then she remains unassigned in this round and does not make any applications until the next round. Each school j considers its applicants. Those students with the highest j-priority are tentatively assigned to school j up to its capacity. The rest of the applicants are rejected.

(3) The round terminates whenever each student is either assigned to a school or is unassigned in this round, i.e., she has been rejected by all her first $(td + d)$ choice schools. At this point, all tentative assignments become final and the remaining capacity of each school is reduced by the number of students permanently assigned to it.

Apart from its theoretical significance, the parallel mechanism is also used in the centralized Chinese college admissions. For example, Shanghai uses $d = 2$; Jiangsu uses $d = 3$; Hainan uses $d = 6$; Tibet uses $d = 10$ (Chen and Kesten, 2017).

2.8 Comparing Manipulability Across Mechanisms

Pathak and Sönmez (2013) gave a definition that allows us to compare manipulability across mechanisms. Intuitively, a mechanism ψ is at least as manipulable as another mechanism φ if for any preference profile, if students are not truthful under φ, then students are also not truthful under ψ.

Definition 2 (Pathak and Sönmez (2013))**.** *A profile (set of reported preferences) t is* vulnerable *under mechanism φ if there exists a student i that strictly prefers to report $t_i' \neq t_i$ over t_i if all other players' reports follow the profile t_{-i}.*

A mechanism ψ is at least as manipulable *as mechanism φ if any profile that is vulnerable under φ is also vulnerable under ψ.*

A mechanism ψ is more manipulable *than mechanism φ if ψ is at least as manipulable as φ, and there is a set of students, allocations, and a profile t where t is vulnerable under ψ but not under φ.*

In a sense, this is the only definition of comparing manipulability that does not require us to specify which preference profile "counts more" in the definition than another. If we want to compare two mechanisms ψ and φ that are not comparable under Definition 2, then there is a profile p that is vulnerable under ψ but not φ, and there is a profile p' that is vulnerable under

φ but not ψ. So an alternative definition needs to judge whether p counts more than p' when it comes to comparing sets of profiles for manipulability comparison. Since we take preferences as given and fixed in our model, we do not think that type of judgement is justified, and Definition 2 is sufficient for our purposes.

This definition is qualitative; some pairs of mechanisms are not comparable, and we cannot say that a mechanism is "a lot more" manipulable than another. Nevertheless, Chen and Kesten (2017) shows that within the parallel mechanism family, every pair of mechanisms can be compared: those with lower d (that are closer to Boston mechanism and further away from DA) are more manipulable.

Proposition 1 (Chen and Kesten (2017))**.** *If $d' > d$, then $\mathbb{P}_d$ is more manipulable than $\mathbb{P}_{d'}$.*

In sections 3 and 4, we will show that the mechanisms with lower d are also more efficient in ex ante welfare, establishing the key claim that more manipulable mechanisms are more efficient within this family.

3 Theoretical Results

Throughout this section, we assume that schools have no priorities. This assumption is necessary; approximate efficiency does not hold for all preference profiles when priorities are merely weak. For a detailed discussion of this assumption, see subsection 2.6.

Theorem 1 is our main theoretical result, establishing approximate efficiency guarantee for the parallel mechanism that degrades linearly with d. Theorem 2 shows that this linear dependence is optimal.

Theorem 1. *Assume that schools have no priorities. Under any student preference profile and Nash equilibrium play, if students are a continuum, the d-parallel mechanism is $(d+1)$-approximately Pareto efficient; if there are a finite number of students, the d-parallel mechanism is $(2d+1)$-approximately Pareto efficient.*

Proof. We first consider the continuum case and prove the $(d + 1)$ bound. Later in the proof, we will see that we can adapt the proof to get a $(2d + 1)$ bound for the discrete case.

Consider an equilibrium of the mechanism, and write $x^i_{j_1 \dots j_d}$ for the indicator variable that student i lists schools $j_1, \dots, j_d$ in that order as her first d schools. We focus on this set of indicator variables because, as it turns out, we only need to focus on the first round. The efficiency guarantee holds if the allocation after the first round is arbitrary. Let σ be the lottery generated by the mechanism. We use the notation σ^i_j to be the probability that student i gets school j. Let $A \subseteq C$ be the set of schools that are fully allocated by the mechanism.

We will first show that if $A = C$ then σ is d-approximately Pareto efficient. Note that in the first round, the algorithm proceeds like deferred acceptance, and students do not lose priorities on schools as long as they list those schools within the first d schools. The continuum of students allows us to write q_j for the probability that if a student proposes to school j, she will get in. We can view q_j as a "unit" of school j that any student can get by applying.

Note that if a student lists schools $j_1, \dots, j_d$ in that order, then the probability of this student getting into school j_r, $1 \le r \le d$, is $(1 - q_{j_1}) \cdots (1 - q_{j_{r-1}})q_{j_r}$ because she must propose to schools $j_1, \dots, j_{r-1}$, is rejected from all of them, proposes to school j_r and is accepted.

First, we prove that $\sum_j \sigma_j^i / q_j \le d$. It follows from the following calculation.

$$\begin{aligned}
\sigma_j^i &= \sum_{r=1}^{d} \sum_{a_1,\dots,a_{d-1}} x^i_{a_1\dots a_{r-1} j a_r \dots a_{d-1}} q_j \prod_{k=1}^{r-1} (1 - q_{a_k}) \\
&\le \sum_{r=1}^{d} \sum_{a_1,\dots,a_{d-1}} x^i_{a_1\dots a_{r-1} j a_r \dots a_{d-1}} q_j.
\end{aligned}$$

Therefore,

$$\sum_j \frac{\sigma_j^i}{q_j} \le \sum_{r=1}^{d} \sum_{a_1,\dots,a_{d-1},j} x^i_{a_1\dots a_{r-1} j a_r \dots a_{d-1}} \le \sum_{r=1}^{d} 1 = d$$

We claim that σ is d-approximately Pareto efficient. Suppose for the contradiction that there exists a lottery μ such that $v^i(\mu) \ge d v^i(\sigma)$ and the inequality is strict for a positive measure of students. Define $y^i = \sum_j \mu_j^i / q_j$ be the total number of units obtained by student i in μ. Note that if $y^i < d$, then μ^i is dominated by a convex combination of outcomes that student i could have obtained in different strategies. The equilibrium condition therefore implies that $v^i(\mu) < d v^i(\sigma)$, but this contradicts d-approximate Pareto dominance. We conclude that $y^i \ge d$ for all i, and $y^i > d$ for some positive measure of students. However,

$$\int_i y^i \mathrm{d}\rho = \int_i \left(\sum_j \frac{\mu_j^i}{q_j} \right) \mathrm{d}\rho = \sum_j \frac{1}{q_j} \int_i \mu_j^i \mathrm{d}\rho \le \sum_j \frac{1}{q_j} \int_i \sigma_j^i \mathrm{d}\rho = \int_i \left(\sum_j \frac{\sigma_j^i}{q_j} \right) \mathrm{d}\rho \le \int_i d \mathrm{d}\rho$$

where the first inequality comes from the fact that each school is fully allocated under σ and the second inequality is a previously proved claim. But this contradicts the fact that $y^i \ge d$ for all i, and $y^i > d$ for some positive measure of students. We therefore conclude that σ is d-approximately Pareto efficient.

We now consider the case $A \ne C$. Let μ be any lottery and assume $v^i(\mu) \ge (d+1) v^i(\sigma)$ for all i and that this inequality is strict for a positive measure of students. Note that

$$v^i(\mu) = \sum_{j \in A} \mu_j^i v_j^i + \sum_{j \notin A} \mu_j^i v_j^i$$

For the second summation term, note that

$$\sum_{j \notin A} \mu_j^i v_j^i \le \max_{j \notin A} v_j^i \le v^i(\sigma)$$

where the first inequality comes from $v_j^i \le \max_{j \notin A} v_j^i$ and $\sum_j \mu_j^i \le 1$ and the second inequality comes from the equilibrium condition since student i can select school $\arg\max_{j \notin A} v_j^i$ and get accepted with certainty. Therefore,

$$d v^i(\sigma) \le \sum_{j \in A} \mu_j^i v_j^i$$

Write v_*^i for the value obtained by student i conditional on being rejected from her top d choices and thus unassigned in the first round. If student i lists $j_1, \ldots, j_d$, then

$$v^i(\sigma) = \sum_{r=1}^{d} \sigma_{j_r}^i v_{j_r}^i + \left(1 - \sum_{r=1}^{d} \sigma_{j_r}^i\right) v_*^i = \sum_{r=1}^{d} \sigma_{j_r}^i (v_{j_r}^i - v_*^i) + v_*^i$$

The payoffs of a student who listed schools in A (among top d choices) are therefore a linear shift of her payoff in a modified economy where the set of schools is A and the value of student i for school j is $v_j^i - v_*^i$, which is necessarily non-negative. The allocation continues to be an equilibrium in this modified game, and by definition of A, all schools are fully allocated. The earlier analysis therefore implies that σ is d-approximately Pareto efficient in this modified market. On the other hand,

$$\sum_{j\in A} \mu_j^i (v_j^i - v_*^i) \geq \left(\sum_{j\in A} \mu_j^i v_j^i\right) - v_*^i \geq d v^i(\sigma) - v_*^i \geq d \sum_{j\in A} \sigma_j^i \left(v_j^i - v_*^i\right)$$

for all i that lists schools in A, and the inequality must be strict for a positive measure of students, which is a contradiction.

Consider now the case of a finite number of students. The proof proceeds similarly to the continuum case. We first claim that if $A = C$, that is, all schools are fully allocated, then σ is $2d$-approximately Pareto efficient. If a student lists schools $j_1, \ldots, j_d$ in that order, then the probability of this student getting into school j_r, $1 \leq r \leq d$ is still at most q_{j_r}, which is sufficient to prove that $\sum_j \sigma_j^i / q_j \leq d$ as in the continuum case. Now we define $y^i = \sum_j \mu_j^i / q_j$ as in the continuum case. We use continuity in the above proof to argue that any student i receives the full q_j unit of school j. In the discrete setting, student i can get at least half a unit of school j, $q_j/2$, by applying school j in the first round. This is because we assume that school j is allocated, so at least one student must have applied to school j in the first round. When student i deviates to apply to school j in the first round as well, the number of students who apply to school j in the first round at most doubles. We therefore derive the contradiction analogously if we assume that there is a lottery μ such that $v^i(\mu) \geq 2dv^i(\sigma)$ and the inequality is strict for some student i. The argument that extends the $2d$-approximate Pareto efficiency in the case $A = C$ to $(2d+1)$-approximate Pareto efficiency in the case $A \neq C$ is completely analogous. □

Theorem 2. *There exists an instance of the d-parallel mechanism that is not $d/2$-approximately Pareto efficient.*

Proof. We exhibit an explicit example, adapted from Immorlica et al. (2017). There are d schools, and school j has 2^{j-1} seats, $1 \leq j \leq d$. There are d types of students, and type i has value $v_j^i = 1 + (n - j + 1)\epsilon$ for schools $1 \leq j \leq i$ and $v_j^i = (n - j + 1)\epsilon$ for $j > i$, where ϵ is negligibly small. Every student has value zero for being unassigned. In other words, student preferences are completely aligned (lower-numbered schools are preferred) and type i has value around 1 for schools number i and below, and 0 otherwise. Schools have no priorities. The parallel mechanism reduces to DA in this case because the number of schools is equal to d. Since schools have no priorities and student preferences are aligned, DA reduces to Random Serial Dictatorship. Each student of type i gets value 1 if she gets some schools $j \leq i$ which has $\sum_{j=1}^{i} 2^{j-1} = 2^i - 1$ total number of seats, and there are $d2^d$ students, each of whom is equally likely to get it. so each student of type i gets an ex ante value $(2^i - 1)/d2^d$ from the mechanism. An alternative allocation that randomly assigns school i to students of type i gives type i an ex ante value of $2^{i-1}/2^d$, which is a factor of bigger than $d/2$ improvement.

□

Recall the intuition that immediate acceptance increases efficiency and deferred acceptance decreases it. It is no surprise that the bad example in Theorem 2 maximizes the deferred acceptance feature by setting $n = d$ and reducing it to DA.

We note here that the proofs of both theorems depend only on the first round of the parallel mechanism. Therefore, we still get the inefficiency bound between $d/2$ and $d+1$ for any mechanism that permanently assigns students to schools using deferred acceptance on the first d schools in the preference list of each student and assigns remaining students arbitrarily thereafter.

Let γ_d be the largest number such that the d-parallel mechanism is γ_d-approximately Pareto efficient with a continuum of students and no priorities. Theorem 1 and 2 together show that $d/2 \le \gamma_d \le d + 1$, so the inefficiency factor γ_d grows linearly in d. The following questions are open. First, can we calculate γ_d exactly, or give tighter bounds for γ_d? Second, is γ_d increasing in d?

4 Simulations

In this section, we use numerical simulations to illustrate the performance of the parallel mechanism with different parameters across a selection of environments. The main takeaway is that efficiency degrades with d in a wide range of environments.

4.1 Settings

Reported simulation results in this section are from the setting used by Abdulkadiroğlu et al. (2015). There are 5 schools, each with capacity of 20 seats, and 100 students. The welfare metric of interest is approximate Pareto efficiency. We consider two cases, no priorities and weak priorities, and the results are reported in Table 1 and 2, respectively.

We construct students' preferences for schools as follows. First, we draw the unnormalized preference values. The unnormalized value of student i for school j is given by $\tilde{v}^i_j = \alpha u_j + (1-\alpha)u^i_j$, where $\alpha \in [0, 1]$ is a fixed parameter of the given setting. u_j is a common value component, and u^i_j is a private value component. For each i and j, u_j and u^i_j are drawn independently from the uniform distribution on $[0, 1]$. Second, we normalize the preference values: the normalized value of student i for school j is given by

$$v^i_j = \frac{\tilde{v}^i_j - \min_{j'} \tilde{v}^i_{j'}}{\max_{j'} \tilde{v}^i_{j'} - \min_{j'} \tilde{v}^i_{j'}}.$$

Under this normalization, each student has value zero and one for the least preferred and most preferred schools, respectively. This normalization is sometimes referred to as "unit-range" in the literature (Filos-Ratsikas et al., 2014; Adamczyk et al., 2014). Subsequent results are invariant to affine transformations of the unnormalized preference values.

The parameter α determines the strength of the common value relative to the private value component in a given setting. The case $\alpha = 1$ is the pure common value case, and every student shares the same cardinal preferences. The case $\alpha = 0$ is the pure private value case, and students have independent preferences.

For each $\alpha \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$, we construct 100 preference draws as above. For each such preference draw, we compute an equilibrium of the parallel mechanism with parameter $d \in \{1, 2, 3, 4, 5\}$ using modified best-response dynamics outlined in the next subsection. $d = 1$ corresponds to Boston mechanism, and $d = 5$ corresponds to DA.

Lastly, we construct school priorities. In Table 1 *no priorities*, every school has one priority class and ties are broken randomly within each class. In Table 2 *weak priorities*, each school has 2 priority classes. Schools $1, 2, 3$ rank students $\{1, 2, \dots, 50\}$ before $\{51, 52, \dots, 100\}$, and schools $4, 5$ rank students $\{51, 52, \dots, 100\}$ before $\{1, 2, \dots, 50\}$.[7]

4.2 Equilibrium Computation

Each set of preference draw and mechanism (i.e. each d) induces a complete information game, and the action space of each student consists of all $5! = 120$ permutations of the 5 schools. We compute a Nash equilibrium using the process described below.

(1) We initialize each student's play to her truthful ordinal report. At any time step, keep track of the current play $\mathbf{x}$.

(2) Pick a student i uniformly at random. Each student has 120 possible reports; we estimate student i's ex ante utility under all such reports. For each possible report $\tilde{x}_i$ of i, estimate the ex ante utility of student i fixing everyone else's report at the current play. We estimate the ex ante utility by drawing 2000 random priorities for the schools. For each set of priority draw, we feed the school priorities and students' reported preferences $(\tilde{x}_i, \mathbf{x}_{-i})$ into the mechanism. The mechanism gives an allocation, and we can compute each student's utility from that allocation. We then take the average over all priority draws to get an estimate and the standard error of each student's ex ante utility. We change the report of student i to the report that has the highest estimated ex ante utility, that improves the ex ante utility of the original report by at least 0.01 with non-overlapping confidence intervals.

(3) If the last 500 previous iterations of step (2) have an average utility improvement of less than 0.01, return $\mathbf{x}$ and terminate. Else, repeat step (2).

For DA, we can stop at step (1) and estimate the ex ante utilities directly.

4.3 Pareto Efficiency Welfare Loss

For each set of preference draw and d, subsection 4.2 gives us an equilibrium $\mathbf{x}$. We ask the following question. What is the largest $\gamma \geq 1$ such that there exists a randomized allocation $\mathbf{y}$ such that every student is at least γ times better off under $\mathbf{y}$ than under $\mathbf{x}$.

[7] There are many ways to specify weak priorities and this specification is not special; it is meant to be an illustration that the welfare loss generally increases with d even with weak priorities.

	$d=1$	$d=2$	$d=3$	$d=4$	$d=5$
$\alpha = 0.1$	0.0	0.0	0.0	0.0	0.0
$\alpha = 0.3$	0.0	0.0	0.0	0.0	0.0
$\alpha = 0.5$	0.1	0.2	0.4	2.0	2.2
$\alpha = 0.7$	2.5	4.7	7.7	8.7	8.8
$\alpha = 0.9$	3.2	4.8	5.5	6.2	6.3

Table 1: Pareto Welfare Loss (%), No Priorities

	$d=1$	$d=2$	$d=3$	$d=4$	$d=5$
$\alpha = 0.1$	0.0	0.0	0.0	0.0	0.0
$\alpha = 0.3$	0.0	0.0	0.0	0.0	0.2
$\alpha = 0.5$	0.1	0.2	0.4	1.0	1.3
$\alpha = 0.7$	2.3	4.7	6.9	8.1	10.0
$\alpha = 0.9$	3.3	4.1	4.5	5.7	5.8

Table 2: Pareto Welfare Loss (%), Weak Priorities

We first describes how to implement a subroutine that, given γ, checks whether such a $\mathbf{y}$ exists. This subroutine can be implemented as a feasibility test of a linear program, which can be efficiently solved. The variables of the linear program are y_j^i, the probability that student i is allocated school j under $\mathbf{y}$. The constraints are

$$0 \le y_j^i \le 1 \quad \forall i, j; \qquad \sum_j y_j^i \le 1 \quad \forall i; \qquad \sum_i y_j^i \le 20 \quad \forall j; \qquad \sum_j y_j^i v_j^i \ge \gamma \sum_j x_j^i v_j^i \quad \forall i$$

We can then use this subroutine to approximate the optimal γ with a modified binary search. We report $(1 - 1/\gamma) \times 100$ as the percent welfare loss. For each fixed (α, d), we generate 100 preference draws based on α, and we can compute the optimal γ for each such draw. Averaging over all preference draws gives an average Pareto welfare loss in percent for each (α, d).

4.4 Discussion

Table 1 and 2 show Pareto welfare loss in percent with no and weak priorities respectively, for each setting parameter α, across different mechanisms $d \in \{1, 2, 3, 4, 5\}$. We can see that in each setting, Pareto welfare loss goes up as d goes up from 1 to 5. That is, as the mechanism shifts closer to DA than Boston mechanism, it becomes more inefficient even as it becomes less manipulable.

While not the focus of this paper, we also note that α has a major impact on the efficiency of mechanisms across the range of d. When α is small (0.1 and 0.3 in our example), the allocation is nearly Pareto efficient. Intuitively, when α is small, student preferences are mostly uncorrelated. A Pareto improvement must improve on all these mostly uncorrelated frontiers, which is hard. In contrast, when α is large and preferences are mostly correlated, improvement in only one dimension can make every student better off.

We interpret the simulation result as the inefficiency-manipulability tradeoff in *typical* settings, that is, when student preferences arise in a natural and interpretable way. This result, therefore, complements the theoretical results that show the tradeoff in the *worst-case* setting, when preferences can be adversarially chosen to maximize the inefficiency of a given mechanism. Note that the behaviors of mechanisms in the typical case and in the worst case are different phenomena and one does not follow from the other. While we show that efficiency loss increase with d for both cases, the magnitudes are not comparable. For example, with $d = 5$, the lower bound $d/2$ and the upper bound $(d + 1)$ worst case guarantees correspond to Pareto loss of 60% and 83.3% when our typical case shows the largest loss at 8.8%. Importantly, the gap between theoretical results and simulations is not because the bound is not tight, but rather that worst-case instances are not “typical.”[8]

Our system – 100 students and 5 schools – is relatively small compared to most school choice systems. We use this system for three reasons. First, this system is a benchmark used before by previous papers that compute equilibria of non-strategyproof school choice mechanisms (Abdulkadiroğlu et al., 2015; Immorlica et al., 2017). Our simulation results can therefore be directly compared with theirs. Second, equilibrium computation is very computationally intensive. For n schools, the strategy space of each student has size $n!$, one for each possible ordering of the schools in her reported preference list. The modified best response dynamics also takes time that grow with the dimension of the strategy space $n!$, which is super-exponential in n. These issues make computing Nash equilbrium for even moderate values of n computationally difficult. Third, we find in the theoretical section that bounds for a continuum of students are tighter than that for discrete students. Even though those bounds are for worst-case instances, they suggest that larger systems are “better-behaved,” so if the tradeoff holds for small systems, we might expect it to hold in larger systems as well.

5 Conclusion

We show that the worst case inefficiency guarantee of the d-parallel mechanism degrades linearly with d: between $d/2$ and $d+1$. This result suggests that mechanisms that are closer to Boston mechanism than DA (lower d, hence more manipulable) are also more efficient in ex ante welfare. We corroborate this insight in typical settings by simulations.

We can interpret our result as a recommendation that policy makers should consider whether the parallel mechanism with intermediate d might be better than existing mechanisms at achieving desired goals. Historically, most places that abandoned the Boston mechanism adopted DA in place; such moves swapped one extreme end of the spectrum (most efficient, most manipulable) with the other extreme (least efficient, least manipulable). Depending on context, a parallel mechanism with intermediate d might improve efficiency while still being “good enough” in terms of incentives. While this paper shows general properties of the parallel mechanism, a decision whether to implement one requires more detailed modeling such as Shi (2016) that takes into account local institutional features and a better understanding of preferences of affected families and schools. Furthermore, the inefficiency-manipulability tradeoff identified in this paper should be considered as one factor among many in school choice design. Practical policy decisions depend also on strategic simplicity, “fairness” broadly defined, effects on non-strategic students, accurate data collection and policy analysis, and even how easy it is to describe the mechanism to different stakeholders and gain their support. Apart from possible policy relevance, our takeaway message is twofold.

First, our work presents the first theoretical and empirical analysis of the efficiency of the parallel mechanism, despite its common use in practice. Even for the Boston mechanism, the most well-known non-strategyproof mechanism in this family, existing works focus on structural properties of Nash equilibria rather than quantifying inefficiencies of equilibria. Our work can thus be viewed as a contribution to this understudied area in its own right.

[8]This issue is rather common in algorithm design. For exmaple, the simplex algorithm is very fast in practice but has worst-case exponential time. Beyond worst-case analysis is interesting but much more challenging.

Second, our work is the first to show that the amount of manipulability, and not just the possibility of it, affects the efficiency of the mechanism. The parallel mechanism can then be viewed more as a technical tool to interpolate between Boston mechanism and DA, an illustration to show how efficiency smoothly degrades from the Boston mechanism end of the family to the DA end. The parallel mechanism is only one possible choice. Nevertheless, there is no canonical way to interpolate between mechanisms, and the parallel mechanism is a natural choice for two reasons: (1) it does immediate acceptance in each round like Boston mechanism, but within each round it does deferred acceptance, (2) every two parallel mechanisms can be compared in terms of manipulability.

Our work suggests several directions for future research. The most immediate questions are to close the gap between the lower and the upper bound of the worst case bound, and to analyze the efficiency loss with real school choice data. Further afield, our work shows a connection between manipulability and efficiency via the parameter d and existing manipulability comparisons. A natural connection that directly translates manipulability to efficiency would be ideal. Moreover, if we hope to use manipulability to improve efficiency, students must actually know how to correctly manipulate and reach some sort of equilibrium. Learning that might be possible in a repeated game, but most students only have one chance to play. The preferred mechanism should therefore be "strategically simple" (perhaps in the sense of Börgers and Li (2018)) even if it is not strategyproof. Further theory, empirics, and experiments are needed to make the manipulability idea practical.

References

Atila Abdulkadiroğlu, Yeon-Koo Che, and Yosuke Yasuda. 2011. Resolving Conflicting Preferences in School Choice: The "Boston Mechanism" Reconsidered. *American Economic Review* (2011).

Atila Abdulkadiroğlu, Yeon-Koo Che, and Yosuke Yasuda. 2015. Expanding "Choice" in School Choice. *American Economic Journal: Microeconomics* (2015).

Atila Abdulkadiroğlu, Parag A. Pathak, and Alvin E. Roth. 2009. Strategy-proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match. *American Economic Review* (2009).

Atila Abdulkadiroğlu and Tayfun Sönmez. 2003. School Choice: A Mechanism Design Approach. *American Economic Review* (2003).

Marek Adamczyk, Piotr Sankowski, and Qiang Zhang. 2014. Efficiency of truthful and symmetric mechanisms in one-sided matching. *International Symposium on Algorithmic Game Theory (SAGT)* (2014).

Itai Ashlagi and Afshin Nikzad. 2017. *What Matters in Tie-breaking Rules? How Competition Guides Design.* Technical Report.

Eduardo M. Azevedo and Jacob D. Leshno. 2016. A Supply and Demand Framework for Two-Sided Matching Markets. *Journal of Political Economy* (2016).

Haris Aziz, Jiashu Chen, Aris Filos-Ratsikas, Simon Mackenzie, and Nicholas Mattei. 2016. Egalitarianism of random assignment mechanisms. *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (2016).

Anand Bhalgat, Deeparnab Chakrabarty, and Sanjeev Khanna. 2011. Social Welfare in One-sided Matching Markets without Money. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX-RANDOM)* (2011).

Anna Bogomolnaia and Hervé Moulin. 2001. A new solution to the random assignment problem. *Journal of Economic Theory* (2001).

Tilman Börgers and Jiangtao Li. 2018. *Strategically Simple Mechanisms.* Technical Report.

Caterina Calsamigliay, Chao Fuz, and Maia Güell. 2017. *Structural Estimation of a Model of School Choices: the Boston Mechanism vs. Its Alternatives.* Technical Report.

Gabriel Carroll. 2018. On mechanisms eliciting ordinal preferences. *Theoretical Economics* (2018).

Yeon-Koo Che and Fuhito Kojima. 2010. Asymptotic equivalence of probabilistic serial and random priority mechanisms. *Econometrica* (2010).

Yan Chen and Onur Kesten. 2017. Chinese College Admissions and School Choice Reforms: A Theoretical Analysis. *Journal of Political Economy* (2017).

George Christodoulou, Aris Filos-Ratsikas, Soren Kristoffer Stiil Frederiksen, Paul W. Goldberg, Jie Zhang, and Jinshan Zhang. 2016. Social welfare in one-sided matching mechanisms. *International Conference on Autonomous Agents and Multiagent Systems (AAMAS)* (2016).

Lars Ehlers, Dipiyoti Majumdar, Debasis Mishra, and Arunava Sen. 2016. *Continuity and Incentive Compatibility in Cardinal Voting Mechanisms.* Technical Report.

Aytek Erdil and Haluk Ergin. 2008. What's the Matter with Tie-breaking? Improving Efficiency with School Choice. *American Economic Review* (2008).

Haluk Ergin and Tayfun Sönmez. 2006. Games of School Choice under the Boston Mechanism. *Journal of Public Economics* (2006).

Aris Filos-Ratsikas, Søren Kristoffer Stiil Frederiksen, and Jie Zhang. 2014. Social welfare in one-sided matchings: Random priority and beyond. *International Symposium on Algorithmic Game Theory (SAGT)* (2014).

David Gale and Lloyd S. Shapley. 1962. College Admissions and the Stability of Marriage. *The American Mathematical Monthly* (1962).

Patrick Harless. 2014. *A School Choice Compromise: Between Immediate and Deferred Acceptance.* Technical Report.

Yinghua He. 2017. *Gaming the Boston School Choice Mechanism in Beijing.* Technical Report.

Yinghua He, Antonio Miralles, Marek Pycia, and Jianye Yan. 2018. A Pseudo-Market Approach to Allocation with Priorities. *American Economic Journal: Microeconomics* (2018).

Katharina Huesmann and Achim Wambach. 2016. *Constraints on Matching Markets Based on Moral Concerns.* Technical Report.

Aanund Hylland and Richard Zeckhauser. 1979. The Efficient Allocation of Individuals to Positions. *Journal of Political Economy* (1979).

Nicole Immorlica, Brendan Lucier, Glen Weyl, and Joshua Mollner. 2017. Approximate Efficiency in Matching Markets. *Proceedings of the 13th International Conference on Web and Internet Economics (WINE)* (2017).

Onur Kesten and Morimitsu Kurino. 2017. *Strategy-proof Improvements upon Deferred Acceptance: A Maximal Domain for Possibility.* Technical Report.

Onur Kesten and M. Utku Ünver. 2010. School Choice with Consent. *Quarterly Journal of Economics* (2010).

Onur Kesten and M. Utku Ünver. 2015. A theory of school choice lotteries. *Theoretical Economics* (2015).

Fuhito Kojima. 2008. Games of School Choice under the Boston Mechanism with general priority structures. *Social Choice and Welfare* (2008).

Timo Mennle and Sven Seuken. 2017. *Hybrid Mechanisms: Trading Off Strategyproofness and Efficiency of Random Assignment Mechanisms.* Technical Report.

Timo Mennle and Sven Seuken. 2018. *Partial Strategyproofness: Relaxing Strategyproofness for the Random Assignment Problem.* Technical Report.

Antonio Miralles. 2009. School Choice: The Case for the Boston Mechanism. *International Conference on Auctions, Market Mechanisms and Their Applications (AMMA)* (2009).

Parag A. Pathak and Tayfun Sönmez. 2013. School Admissions Reform in Chicago and England: Comparing Mechanisms by Their Vulnerability to Manipulation. *American Economic Review* (2013).

Marek Pycia and M. Utku Ünver. 2017. Incentive compatible allocation and exchange of discrete resources. *Theoretical Economics* (2017).

Lloyd Shapley and Herbert Scarf. 1974. On cores and indivisibility. *Journal of Mathematical Economics* (1974).

Peng Shi. 2016. *Assortment Planning in School Choice.* Technical Report.

Peter Troyan. 2012. Comparing school choice mechanisms by interim and ex-ante welfare. *Games and Economic Behavior* (2012).