

Broadening the market design approach to school choice¹

Estelle Cantillon, June 2017

Forthcoming Oxford Review of Economic Policy

Abstract: School choice refers to policies that allow parents' preferences to be an input to the decision of which school a student will attend. A rich body of research has developed over the past 10-15 years to study mechanisms that implement school choice. This literature has mostly taken the inputs of school choice – preferences, priorities and capacities – as exogenous. More recently, researchers have sought to embed the school choice problem into its wider context, thereby broadening the scope of market design questions and enriching the analysis. This article discusses current school choice policy issues in light of this recent literature and outlines remaining open questions.

1. Introduction

School choice refers to policies that allow parents' preferences to be an input to the decision of which school a student will attend. The motivations for introducing school choice vary widely across countries but generally fall within one of the three following categories: to account for preference heterogeneity in the presence of a diverse educational offer (different curricula, different pedagogies, ...), to ensure equality of access to schools when schools' qualities differ, or to introduce competition between schools with the idea that it will increase quality.

Choice is rarely unconstrained, however. Even when preferences are the main input to the decision, limited capacity puts constraints on the number of admitted students in a school. The existence of peer effects or other externalities can also be reasons why, from a policy perspective, one might want to put constraints on choice and influence schools' student intake. Such concerns motivate current policies, in many countries, that seek to foster greater ethnic, academic or socio-economic diversity in schools.

In market design, the school choice problem refers to the question of how to allocate students to schools, accounting for preferences, school capacities and policy objectives.² The result will be an enrolment procedure that determines who and how parents can participate and how students will get assigned to schools. Inherent to the market design approach to school choice is the explicit account for parents' and schools' incentives.

Several excellent surveys of the current state of the literature and practice on school choice already exist (see e.g. Abdulkadiroğlu and Sönmez, 2013, Abdulkadiroğlu, 2013, and Pathak, 2016). My point of departure in this article will be the school choice problem as *one* element of the educational system, and its related markets. Indeed, all the inputs to the school choice problem (preferences, school capacities, priorities) are influenced by other educational policies and other markets among which the market for teachers and the residential housing market. Likewise, enrolment procedures, through the incentives that they create and the allocation of students that they generate, can also impact these related markets and the performance of the educational system.

¹ Acknowledgements: This piece has been inspired by my interactions over the years with policymakers and other stakeholders of school choice. They have forced me to take on their perspectives and concerns. I thank Julien Grenet, Scott Kominers, Thayer Morrill, Juan Pereyra, Alex Teytelboym and the editors of OXREP for their helpful comments and suggestions on a prior version.

² Implicit in this framing of the school choice problem is the understanding that money should not and does not play a role in determining the allocation.

Looking at these interactions is a recent but active area of research in market design. It is also an exciting area of research because it bridges the longstanding literature on school choice in education economics with the market design approach to school choice, to the benefit of both, and because it broadens the set of questions that market design research can address, beyond the original mostly redistributive question of “who goes to which school.”

2. The market design approach to school choice

The market design approach to school choice finds its origins in the seminal papers by Balinski and Sönmez (1999) and Abdulkadiroğlu and Sönmez (2003). These papers showed how the school choice problem can be cast in terms of the general two-sided matching problem introduced by Gale and Shapley (1962).³ In the canonical school choice model, there are schools on one side and students on the other side. Students have preferences over schools. Schools do not have preferences over students, but must assign places according to some exogenously imposed priorities in case of excess demand. The question is how to assign each student to a school, subject to capacity. Balinski and Sönmez (1999)’s insight was that these priorities play a similar role to preferences in the two-sided matching model. The main differences are that, in the school choice problem, we only care about the welfare of one side (the students) and we do not need to worry about the incentives for schools to reveal their “preferences” since they are given. Abdulkadiroğlu and Sönmez (2003) further described what made the school choice problem special in the class of two-sided matching problems: the fact that priorities are typically coarse, i.e. they define classes of students with priorities over others (e.g. siblings have priorities over other students) but do not provide a complete ordering of students, and the fact that students will typically benefit from different priorities at different schools. I would add to this list the common understanding, in most countries, that the State or some other public authority must guarantee a place to all school-age students.

An important insight from this literature is that school choice is not a 0-1 variable. The actual organization of how choice can be exerted, and how places are allocated to students, matters for its consequences. In this section, I briefly illustrate the market design approach to analysing the properties of these procedures and summarize the main findings of the literature to date. I then describe common performance measures for school enrolment procedures. These measures include the metrics commonly used in the literature but also other metrics motivated by practices and concerns from the field. The last subsection describes a few examples where insights from the market design approach have informed school choice policy reforms.

2.1. School enrolment procedures

A school enrolment procedure has typically three ingredients: (1) a set of rules for participating and expressing preferences over schools, (2) priorities that determine which student has priority over another one, at each school, and (3) an algorithm that determines which student goes where, given his parents’ reported preferences, his priority status at each school, and school capacities.

Participation rules describe who is eligible to apply for a school and how they can do so. They have first order effects on the ability of the procedure to satisfy preferences. A first issue is participation. Enrolment procedures cannot guarantee equal access to quality schools unless all eligible parents participate. Is participation automatic or does it require an action by the parents? Are eligible

³ This is not the only way to approach the school choice problem. One could have framed it as one of maximizing an objective function based on students’ reported preferences and school capacities. The current approach artificially separates the normative analysis of the algorithm from the normative discussion about priorities. It explains why the design of priorities has received little attention so far.

parents solicited to apply or should they take the initiative? How easy is it to participate? The timing of applications as well as its ease of use may also affect participation.

A second issue concerns parents' ability to express their preferences. The more they can express about their preferences over schools, the better the procedure will be able to take them into account. Typically, parents are asked for an ordinal ranking over schools. The number of schools they can rank is often limited however, which Haeringer and Klijn (2009) have shown can introduce a need for parents to strategize and drop schools that they feel unlikely to get any ways. In principle, information about parents' preference intensities could be relevant too but I am not aware of any school choice procedure that elicits such information directly, though some procedures do reward behavior that expresses preference intensities. Interestingly, Estonia only asks parents to declare a school as acceptable or not, rather than their preferences over schools (Lauri et al., 2014)

Information is important. Understanding the choice options is no easy task and figuring out how to rank them is even harder. There is great heterogeneity across countries and school districts in the information provided to parents, going from a simple list of schools with minimal administrative information, to detailed information about school offering, school performance, student body, and prior admission outcomes. The existing literature shows that parents are responsive to the information provided and, for example, tend to rank higher better performing schools (Hastings, Kane and Staiger, 2009, Koning and Van der Wiel, 2013, Burgess et al., 2015). However, the evidence also suggests that parents do not make equal use of available information and that low income families disproportionately benefit from clear and targeted information (Hastings and Weinstein, 2008).

Priorities describe which student has priority over which other student at each school, and under what conditions. This is an inherently redistributive question because a student wins and a student loses. Therefore, priorities ideally translate the political objectives of the school district. Some objectives give rise to an *absolute priority*, i.e. a priority independently of the number of students benefiting from this priority. Priorities for siblings are a common example of an absolute priority. They are meant to facilitate families' school logistics. Other political objectives give rise to a *conditional priority*: students who meet some criteria benefit from a priority up to a certain level, typically defined by a quota. Diversity objectives fall within this category. The Boston Public Schools' personalized choice sets which guarantee, for each student, access to a limited number of top quartile performing schools is another example (see Shi, 2015 for a description).

When the district pursues several objectives, the structure of priorities will reflect the trade-off among these different objectives. Some priority structures imply a strict hierarchy of priorities (e.g. siblings have priority over non siblings, then, within each category, students are ranked on the basis of distance). Other priority structures use points or other ways to balance the different objectives.

Political objectives often result in what is called priority classes, i.e. groups of students with a priority over another group, rather than an ordering of individual students. A tie-breaking rule is then needed. Lotteries are a common tie-breaking rule. It is important to note that even though priorities and tie-breaking rules both contribute to ranking students in order of priority, they do not have the same status: tie-breaking rules do not reflect any political objective. They are only there because of the need to rank two students when there is only one place left. For this reason, they are typically part of the design discussion.

Once preferences and priorities are given, there are still many ways to match students to schools. To illustrate some of these possibilities, consider the following example with three schools and three students (each school has one place). Adam and Bob prefer school A to school B to school C. Chloe prefers school B to school A to school C. Moreover, Adam benefits from a sibling priority at school A. These preferences are represented in Figure 1.

Figure 1: Adam, Bob and Chloe's preferences over schools (priority from which they benefit in parenthesis)

<u>Adam</u>	<u>Bob</u>	<u>Chloe</u>
-------------	------------	--------------

School A (sibling)	School A	School B
School B	School B	School A
School C	School C	School C

This example shares two common features with real school markets: there is a school (school A) that is the most popular and for which demand exceeds supply (school A is the first choice of Adam and Bob but has only one place). There is also a “bad school” (school C) that is least preferred by all students. Note also that the priorities only dictate that Adam has priority over other students for the place at school A. A tie-breaker will therefore be needed to rank other students.

Let us now examine how different algorithms and tie-breaking rules contribute to the final outcome. Under the “**first-come first-served**” rule still used in many smaller cities in the world, time is used as the tie-breaker (where time corresponds to the position in a queue, the time stamp of an online application or of a phone call). Priorities are implemented with an early registration period and so Adam will be able to apply early to school A, and get registered. Applications from other students are accepted (space permitting) as they come. After Adam’s registration, school A is full and Bob and Chloe both prefer school B over school C. Who will get the coveted place at school B will depend on Bob and Chloe’s parents’ ability to act fast or to queue, sometimes for several days, in front of school B. Bob’s parents also need to be well informed since if they first apply to school A (their true first choice), they are likely to lose precious time to have a chance to get a place in school B.

Another common algorithm is the **immediate acceptance algorithm** (also known as the Boston mechanism or the first-preference-first algorithm). The idea is simple and a priori appealing: the algorithm attempts to maximize the number of students who get their first choices (ignoring other choices of students), then the number of students getting their second choices and so on. Multi-stage implementations of the algorithm exist but for simplicity let us consider a single stage implementation where parents submit from the very beginning a list of schools in decreasing order of preference. The algorithm requires a tie-breaking rule when the existing priorities (here: the sibling priority) are silent about how to rank students. Suppose the outcome of this tie-breaking rule is the following:

Figure 2: Priorities at school A, B and C after the application of a tie-breaking rule

<u>School A</u>	<u>School B</u>	<u>School C</u>
Adam (sibling)	Bob	Adam
Bob	Chloe	Bob
Chloe	Adam	Chloe

If parents submit their true preferences as described in Figure 1, the algorithm will assign Adam to school A (since he has priority over Bob) and Chloe to school B (her first choice) in the first stage. By doing so the algorithm has managed to give their first choices to two students. Once this is done, only school C has a free place and this is what Bob gets. He therefore gets his worst outcome.

But note that if Bob’s parents understand this, they should indicate school B as their first choice, instead of their true first choice. In that case, when the algorithm tries to allocate the maximum number of first choices, Adam gets school A as before but now, Bob gets school B, a better outcome than if he were truthful.

This simple example illustrates a general feature of the immediate acceptance algorithm and its variants: because students’ priorities depend on how they rank schools, they have an incentive to deviate from submitting their true preferences. Essentially, what parents should do in the immediate acceptance algorithm is trade off their true preferences against the probability of getting in. If the chance of getting in their preferred school is low and it is higher in their second choice school, they should rank that second school first if they are risk averse and that second choice is good enough.

Research based on the analysis of submitted preferences in cities using the immediate acceptance algorithm has shown that not all parents have this ability to strategize. Abdulkadiroğlu et al. (2006) documented that black children and students from low socio-economic status in Boston were overrepresented among those ranking second a school that reaches capacity in the first round (a dominated strategy) and ending up unassigned. Calsamiglia and Güell (2014) have shown that, in Barcelona, most parents indicate a safe school (the school where they have priority) as their first choice. Those who do not (and thus supposedly indicate their true first choice) often have an outside option, such as the option to attend a private school. They can thus take the risk involved in reporting their true first choice. More recently, researchers have applied flexible structural approaches that allow for strategic and naïve behavior to preference submission data. While the exact findings vary across settings and methods, all papers find a significant fraction of strategic parents (He, 2016, Agarwal and Somaini, 2016, Calsamiglia et al., 2016).

The final example is based on the following intuitive process: parents apply wherever they want. If there is excess demand at one school, that school uses the priorities (and a tie-breaker if needed) to decide whom to accept. If a student is accepted in several schools, she keeps the offer from the school she prefers and rejects the others. The freed places can then be offered to the next students on the waiting list and the process continues until all places have been offered. A decision by a student to reject an offer is final but she can wait until the end of the process (in the hope of getting a preferred offer) to definitely accept an offer.

The centralized version of this algorithm (which takes preferences and priorities as inputs, carries out all the offers/rejections using these and generating a final allocation) is called the **school-proposing deferred acceptance algorithm** (DA) (Gale and Shapley, 1962). It is called “school-proposing” because we start from schools’ priorities when deciding whom to accept. Acceptance is said to be deferred because a student’s acceptance of an offer is not final until the very end. Variants of the school-proposing deferred acceptance algorithm are currently used in England (Department of Education, 2014), in Finland (Salonen, 2014) and for high schools in France (Hiller and Tercieux, 2013).

In our example, with the priorities of Figure 2, the school-proposing DA will proceed as follows if Adam, Bob and Chloe submit their true preferences. In the first stage of the algorithm (these stages take a few seconds on a computer), school A and C offer a place to Adam and school B offers a place to Bob. Adam who gets two offers, rejects school C (again, in a centralized and computerized version of this, it is not Adam who actually rejects, but the computer uses the preferences submitted by Adam to reject school C’s offer *on Adam’s behalf*). Bob on the other hand holds on for now to his unique offer. In stage 2, only school C has a free place and it offers it to the second student on its list, i.e. Bob. Bob now has two offers in his hands and since he prefers school B he rejects the offer from school C. In stage 3, school C offers its free place to Chloe who holds on to it for lack of a better alternative. Since all places have been offered, the algorithm ends and the students definitely accept the offer they have in hand. Note also that none of the students can do better than submitting their true preferences in this example so this is a (Nash) equilibrium outcome.⁴

As these three examples illustrate, even when students’ preferences and priorities are fixed, the final outcome still depends on which algorithm is used: while Adam always gets school A irrespective of the algorithm,

- (1) Bob and Chloe get B or C, depending on their parents’ *ability to act fast* under the first-come-first-served algorithm

⁴ This is not a general property of the school-proposing DA.

- (2) Bob and Chloe get B or C, depending of Bob's parents *ability to strategize* under the first-preferences-first algorithm, and
- (3) Bob gets B and Chloe gets school C for sure under the school-proposing DA algorithm.

2.2. What is a successful procedure?

Since the list of possible algorithms is very long and practice shows that school districts rarely lack imagination, a natural question that arises is how to choose among them. The market design literature has approached this question by specifying desirable criteria that the algorithm or the final assignment should satisfy and studying the algorithms that meet these criteria. In their seminal paper, Abdulkadiroğlu and Sönmez (2003) proposed three such criteria: ex-post efficiency, absence of justified envy and strategyproofness. **Ex-post efficiency** is a measure of the extent to which a procedure respects preferences. An assignment is said to be ex-post efficient if there does not exist any other feasible assignment which every student weakly prefers and at least one student strictly prefers. So, in the example of Figure 1, an algorithm giving Chloe a place in school A and Bob a place in school B would be inefficient: both Chloe and Bob would prefer swapping places. Note that ex-post efficiency is a very weak requirement because it abstains from comparing the gains and the losses of individual students. So in particular, it does not allow us to compare a situation where all three students get their second choices with a situation when two get their first choices and one gets his third choice. **Absence of justified envy** rules out accepting a student in a school when another one, who actually prefers this school and has priority over that student, is not accepted.⁵ It is a measure of the extent to which the procedure respects priorities. The third desideratum that Abdulkadiroğlu and Sönmez (2003) proposed is **strategyproofness**, a property that says that parents cannot do better than submitting their true preferences to the algorithm. Strategyproofness is desirable for four reasons. First, strategyproofness puts all parents on an equal level playing field: as the experience in Boston shows, not all parents are equally able to “game” the system effectively when the rules require parents to strategize. Pathak and Sönmez (2008) have later formally shown that parents submitting their preferences truthfully in the immediate acceptance algorithm are effectively penalized relative to parents who strategize (see also Dur et al., 2016a, for empirical evidence). Second, by eliciting true preferences, strategyproof procedures are also better able to respect parents' preferences. Third, strategyproof algorithms make it easier to advise parents since the best strategy for them is to report their true preferences. Last but not least, strategyproof algorithms generate preferences data that can be used to monitor the adequacy between supply and demand in the district.

The bad news is that these desiderata are not compatible. The tension arises from the conflict between preferences and priorities. To see this, consider again the preferences of Figure 1 but now suppose that priorities are given by Figure 3 (notice how the example is constructed to create a tension between priorities and preferences: even though school A is the preferred school for Bob and Adam, Chloe – whose first choice is school B - is the one who has highest priority at school A.) In this example, there is only one assignment that does not lead to justified envy: Chloe should get school A, Adam school B and Bob school C.⁶ However, this assignment is inefficient because Chloe prefers school B and Adam prefers school A.

⁵ The terms fairness and stability (even if it is a bit stronger than absence of justified envy) are equally used in the literature.

⁶ To see this, first note that Chloe cannot be assigned to school C since she would have justified envy for whoever gets assigned to school A. Likewise Adam cannot get assigned to school C because he would have justified envy for whoever gets school B. This means school C must be assigned to Bob. To make sure Bob (for

Figure 3: Priorities at school A, B and C (alternative)

<u>School A</u>	<u>School B</u>	<u>School C</u>
Chloe	Adam	Chloe
Bob	Bob	Adam
Adam	Chloe	Bob

Much of the literature since then has sought to identify ways to approach these desiderata as closely as possible. Two candidate algorithms have stood out: the student-proposing deferred acceptance algorithm and the top trading cycle algorithm. The **student-proposing DA** is the mirror version of the school-proposing DA, starting with students' preferences instead of school priorities. Thus in the first round, students apply to their first choice schools. Applications are temporarily accepted up to the capacity of each school (using priorities in case of excess demand). In the second round, rejected students in the first round apply to their second choice school. Each school considers the new applications together with the applications received in the first round and accept them up to capacity following priorities. The process goes on until all students are assigned or preference lists are exhausted. The student-proposing DA is strategyproof and leads to outcomes without justified envy. Indeed, the fact that both new and temporarily accepted applications are reconsidered at all stages ensures that a student will never be accepted at a school if another student for which the school is the best available choice has higher priority. Deferred acceptance also ensures that there is no benefit from misreporting preferences. With an appropriately chosen tie-breaking rule, the student-proposing DA can also be shown to be the most efficient in the class of strategyproof algorithms that generate allocations without justified envy (Gale and Shapley, 1962, Abdulkadiroğlu et al., 2009).⁷ The school-proposing DA has formed the basis of recent school choice reforms in several US cities including Boston and New York.

Where the student-proposing DA always respects priorities, the **top trading cycle algorithm**, inspired by Shapley and Scarf (1974), gives precedence to preferences when preferences and priorities conflict. It works as follows. In the first round, students point to their first choice and schools point to their highest priority student. A cycle is formed whenever a school points to a student who points to another school who points to a student, and so on, until we reach a student who points back to the first school. Students in a cycle are assigned to the school to which they were pointing and removed from the system. Capacities of schools in a cycle are adjusted accordingly. In the second round, students point to their first choice school among those with residual capacity and schools point to their highest priority student among remaining students. Cycles are again removed. The process continues until all students are assigned or there is no more acceptable schools for the unassigned students. It is easily shown that the top trading cycle is strategyproof and ex-post efficient, but because students "trade" priorities in the cycles, exchanging a place in a school in which they have a high priority for a school where they have a lower priority, it can lead to assignments with justified envy.

Ex-post efficiency, absence of justified envy and strategyproofness are not the only natural desiderata to think about when designing a school choice procedure. The academic literature has explored other desiderata and their consequences for the comparison between different algorithms, including ex-ante efficiency (Abdulkadiroğlu et al., 2011), manipulability (Pathak and Sönmez, 2013), and the maximization of first choices (Dur et al., 2016c).

whom school C is his last choice) does not have justified envy for the places at school A and B, these places must be assigned to a student who has higher priority than Bob: Chloe for school A and Adam for school B.

⁷ Recall that tie-breakers are only necessary to implement the algorithm but do not reflect any political objective. Therefore, tweaking the tie breaking rule is a natural place to try and resolve the tension between priorities and efficiency. For a slightly different approach, see also Kesten (2010).

Table 1: Performance metrics for enrolment procedures

	Indicator	Why it is a relevant measure of performance?	Limitations and other remarks
Quantitative measures of performance	Participation rate	Low participation rates signal that the procedure was not able to reach out to all parents or that loopholes exist to secure a place outside of the procedure	
	% successful applications (received an assignment)	A high level of successful applications signals the system's ability to meet parents' preferences	If the system is saturated, even a good enrolment procedure will not be able to deliver a high success rate
	% successful applications accepted (take up rate)	A high acceptance rate for successful applications signals that outcomes are indeed desirable for students and parents	A low acceptance rate for successful applications can also signal a high level of interdependency with other school markets (multiple registrations) and loopholes
	% first choice, % second choice, ...	Measure the procedure's ability to meet parents' desires.	In procedures that do not elicit truthful preferences, these statistics are meaningless and should not be used to measure performance. When preferences are polarized or the system is saturated, even a good procedure will not be able to deliver a high percentage of first choices
	Fraction of student who could benefit from mutually beneficial reallocation of places (efficiency)	This fraction measures the level of inefficiency of the procedure	This measure is only meaningful when the procedure elicits truthful preferences
	% of students with justified envy	This fraction measures the degree to which priorities are violated	Used by empirical researchers mostly
	% of applications leading to an appeal	Appeals signal unclear rules and/or incorrect applications of these rules	Social norms and the effectiveness and fairness in the treatment of appeals also influence appeal rates (if appeals are always rejected then there is no reason to appeal, even if this is justified)
	Transfer rate (% of applications in years other than natural entry points)	Transfers signal parents' unhappiness with their current school assignment	Some transfers are driven by changing preferences (realization that the school does not meet their needs) or by households' events (e.g. divorce, move). A low transfer rate may also reflect very restrictive conditions on school transfers.
	Elapsed time between application and decision	A short time between the application and the publication of the results decreases the uncertainty and stress borne by parents	
	Achievement of other quantitative objectives	For example, measures of social diversity in intake.	
Qualitative measures of performance	Ease of use for parents (incl. strategic simplicity of the algorithm, availability of adequate information)	Complex procedures are a hurdle for effective participation by all parents	What is an easy procedure will often depend on the context
	Ease of use for school administrators	An effective procedure should not overburden school administrations	
	Objectivity and verifiability of priority criteria	Non-objective and unverifiable priority criteria create opportunities for circumventing the spirit of the law	The objectivity and verifiability of criteria will often depend on the existence of reliable student databases.
	Achievement of other qualitative objectives		

On the empirical side, researchers have developed metrics to evaluate and compare the actual performance of school enrolment procedures. Most of these metrics critically rely on the preference submissions being truthful, so that the assignments can be evaluated on the basis of the submitted preferences, or on the ability of the researcher to recover preferences from preference submissions. Some of these metrics directly measure the desirable property (efficiency, absence of justified envy) but others – such as the distribution of preference ranks of the assignment – capture these

properties indirectly. A common empirical finding is that the distribution of preference ranks of the assignments does not differ much across algorithms (Pathak, 2016). While this may suggest that the algorithm is of second order importance, recent work by Pycia (2017) suggests another explanation. The design of these performance measures – their invariance to permutations of students’ identities and outcomes – is bound to neutralize any difference between the most common algorithms when the market is large enough.

Interestingly, school districts tend to look at broader measures of performance. They often care about participation rates and the ease of use by parents because getting all eligible students to effectively participate is essential for the legitimacy of the procedure. Likewise, they see a low take-up rate, i.e. a low fraction of assigned students who actually confirm their registration, as a red flag. Table 1 summarizes the most common metrics used by school districts, their relationship, if any, with the criteria used in the literature, and their limitations.⁸

Incorporating these broader dimensions of performance requires that we further open up the box of the school choice problem. Before doing so, however, it is useful to recall that we do not need to identify algorithms (and procedures) that best perform in all circumstances. The existing literature tells us that this is impossible. What we need is to be able to identify – *given the specificities of the school district at hand* – the algorithm and school choice procedure that best meet their needs. In fact, Abdulkadiroğlu and Sönmez (2003, p. 738) already cast the choice between the student-proposing and top trading cycle algorithm as a choice that depends “*on the structure and interpretation of the priorities. In some applications, policy makers may rank complete elimination of justified envy before efficiency (...) In other applications, the top trading mechanism may be more appealing.*” Some of the recent market design literature has gone further in that direction by analysing how specificities of the market may impact the relative performance of different algorithms (see e.g. Arnosti, 2016, Ashlagi and Nikzad, 2016, Ashlagi et al., 2017)

2.3. Market design in practice

The market design approach to school choice is not only an academic subject. It was an integral part of school enrolment reforms in New York in 2004 and in Boston in 2005. Until 2003, the New York City Department of Education (NYCDOE) operated a decentralized enrolment procedure for public high schools in a very diverse educational landscape. Students could apply to at most 5 schools. The scheme operated over three rounds. In the first round, schools ranked the applicants according to their admission criteria and sent admission letters. Applicants were then asked to keep at most one offer. Freed places were offered to students on the wait list during a second round, and then a third round. That process left around a third of students unassigned by the end of the third round. Moreover, the limit on the number of schools that could be listed (relative to the number of schools available – around 500) and the fact that schools could take students’ submitted preferences into account when making decisions raised concerns about incentives for parents. These issues motivated the NYCDOE to solicit advice from Alvin Roth and collaborators. The new scheme is based on the student-proposing deferred algorithm. It is centralized and allows students to list up to 12 schools on their applications. These two features alone were critical for the 40% increase in the number of assigned students as part of the main round of the procedure (Abdulkadiroğlu et al., 2005). Abdulkadiroğlu et al. (2015) provide a thorough evaluation of the impact of the reform. Take-up rates of successful applications went up by 7.2% and the authors find that all categories of students benefitted on average from the reform.

⁸ It is based on a sample of reports and news announcements from local authorities in the UK, US and Belgium as well as Coldron et al. (2008) and country profiles on www.matching-in-practice.eu.

In Boston, the central issue was the incentives for strategizing that the existing mechanism – the immediate acceptance algorithm – induced. Abdulkadiroğlu et al. (2006) showed that this was particularly detrimental to students of low socio-economic background. This convinced the Boston Public School leadership to switch to the strategyproof deferred acceptance algorithm. Post-reform evidence shows that the number of schools listed by parents on their applications has increased (Pathak, 2006). This suggests that parents have adapted to the new algorithm which does not penalize truthful reporting of preferences and that the new scheme can take this richer preference information into account.

Market design ideas were also influential in other school enrolment reforms even when market designers were not on the front seat. A central motivation for the 2007 reform of the school admission code in England was to ensure the fairness of admission criteria. This included ruling out the immediate acceptance algorithm and its variants, that “*made the system unnecessarily complex for parents*” (Department of Education and Skills, 2007, p. 7). In 2010, I made a presentation in front of heads of school admissions consultation platforms in Flanders that explained that their single quota for students of low socio-economic status was not implementing their stated policy objective of promoting diversity in school intakes. The idea percolated the political decision system and led to the introduction of a double quota for school enrolment in Flanders starting in 2013. Other examples of recent school choice reforms are described in Pathak (2016) and on the website Matching-in-Practice.eu.

3. On preferences, priorities and capacities

Early work and design efforts have taken preferences, priorities and capacities as exogenous and have sought to improve the ways students were assigned to schools based on these inputs. However, none of these inputs are fully exogenous in practice. Priorities and capacities result from policy choices. Preferences are formed on the basis of the information available to parents. In this section, I review the evidence about the endogeneity of these inputs and describe how accounting for their endogeneity both changes the way we should think about the impact of enrolment procedures, and enlarges the set of design dimensions to think about. These new design dimensions are summarized in Table 2.

3.1. The preference formation process

There are at least two ways in which preferences over schools are endogenous to the school enrolment procedure. First, parents will typically need to discover their preferences over schools through a time-consuming and costly process. They will need to gather information and possibly visit the premises. They will need to compare the options. Lovenheim and Walsh (2014) document how changes in the school choice environment in the US increase information gathering by parents. Frictions in that process are likely to result in ex-post unsatisfactory assignments (explaining some of the later transfers and unconfirmed registrations that school districts record) or even non assignments if some parents did not list enough schools.⁹

The design of the enrolment procedure can help reduce these frictions. One way to do this is to ensure that information about schools is readily available and accurate. Existing practices vary widely in this respect. In some school districts, the only information available about schools is their locations and the programs they offer. The consequence is that parents use other means – often imperfect – to infer information about the attributes of the schools that they care about. A second

⁹ In procedures that require parents to be strategic, the informational burden of participating effectively is further increased.

way to reduce the burden of the preference discovery process is to provide decision-support tools or appropriately reduce parents' choice set so that they can focus attention on the relevant schools. There are some innovative practices in this respect but, to my knowledge, no systematic analysis of how information and smart user interfaces can help increase the performance of school choice procedures.¹⁰ Finally, an interesting paper by Harless and Manjunah (2015) shows that different algorithms can create different incentives for information acquisition (see also Bade, 2015).

The second reason why preferences are partly endogenous is that parents care about the quality of peers and teachers, on top of exogenous quality attributes of schools (Rothstein, 2008, Burgess et al., 2015). To the extent that the school choice procedure changes the mix of peers or the teaching body, preferences over schools are likely to change. Calsamiglia et al. (2017) show that the immediate acceptance algorithm can exacerbate social segregation when students have preferences over peers, an effect that deferred acceptance does not have. The intuition is simple. High quality peers value schools with other high quality peers more than others. Because equilibrium behaviour in the immediate acceptance algorithm depends on intensities of preferences, these students will take the risk of applying to the (ex-post) high quality schools whereas their lower quality peers will prefer a safer option. This leads to a self-confirming equilibrium with segregation. In work in progress, I show that an enrolment policy designed to increase social diversity will also decrease the level of polarization of preferences over schools if parents value the presence of high socio-economic status peers. The reason is that these desirable peers will be spread out across schools reducing the endogenous vertical differentiation of schools. Policies that promote social diversity might then kill two birds with one stone: increase diversity and increase the ability of the system to meet the parents' preferences since parents are more likely to prefer different schools.

Changes in the student body can also impact perceived school quality – and thus preferences - indirectly through teachers. Terrier (2014) explains that teachers in France flee schools with a concentration of low socio-economic and immigrant background students and Pop-Eleches and Urquiola (2013) provide systematic evidence that teachers in Romanian secondary schools sort in a manner consistent with a preference for higher achieving students (see also Karbownik, 2016, for evidence in Sweden). This, in turn, is likely to feed back on preferences over schools.

3.2. The design of priorities

The funding and organisation of the education system is an important determinant of priorities. For example, the local funding and organisation of public schools in the US encourage residence-based priorities and a large autonomy for deciding on priorities.¹¹ In Romania and Belgium where education is funded at the state and language-community levels respectively, priorities are decided at that level and are thus common across cities (Pop-Eleches and Urquiola, 2013, Cantillon, 2015).

When local authorities have some freedom in deciding priorities, diversity of practices prevails. For example, the English admissions code defines guidelines for setting priorities but the final choice is left to organizing bodies, including, for public schools, to the local authority. Looked after children and special need students often have the first priority in local public schools but similarities stop there. Some local authorities give priority to siblings. Others give priority to distance or students living in the catchment area. Understanding the political economy of priorities is an interesting open question.

¹⁰ A remarkable example is the Discover website of the Boston Public Schools: <http://www.discoverbps.org/>

¹¹ Here and in the rest of the text, we use the term “public school” to refer to schools that are both publicly financed and publicly run (typically by the local authority).

Even if some objectives are exogenously set, there remain many bureaucratic and technical decisions involved before priorities and priority structure can be fully specified. This is where market design can play a role.

The first question is how to best translate political objectives into priorities. In some cases, this is obvious. If the political objective is to favour students living in the catchment area of the school, giving them priority over other students will do the job. Translating diversity goals into priorities is not so obvious. School districts have sought to foster diversity by reserving slots for minority students or giving them extra priority points. As we will see in the next section, such policies can sometimes have perverse effects. Another example is the use of blocks of places reserved for different sets of students for which Dur et al. (2016b) have showed that the order in which those places are allocated has first order effects on the effectiveness of the policy.

A second question is the choice of criteria that qualify students for priorities. Ideally, they should be easy to verify and capture a personal and intrinsic characteristic of the student that the priority seeks to favour. Verifiability ensures that students cannot claim priorities for which they are not eligible. An interesting variant of the problem arises when proof that a student meets a criterion is self-reported. Under-reporting can happen for lack of documentation and was a concern for policymakers in Flanders when they tried to provide favourable treatment to disadvantaged students. They addressed this problem by enlarging the list of criteria to be eligible for the “low socio-economic status” priority. Under-reporting can also happen for strategic reasons when it provides a student with a more favourable treatment, a situation described by Aygün and Bó (2016) in the context of college admissions in Brazil.

The use of proxies instead of personal and intrinsic characteristic to determine priorities is common but has two major shortcomings. First, proxies reduce the effectiveness of the policy since the policy favours students that may not be exactly those that are targeted. Ellison and Pathak (2016) illustrate this distortion in the context of race-neutral alternatives to race quotas in the Chicago Public Schools. Another common proxy is census-track data to determine the socio-economic status of a student. Absent legal restrictions (as in the case of race-based policies in some US states), access to high quality administrative data for those in charge of implementing the procedure can alleviate these problems. A second shortcoming of proxies is that they can create room for strategic behaviour that undermine the objective of the policy. An example is Texas top ten percent plan that guaranteed admission to any in-state public university to students who graduated in the top 10% of their high school classes. The policy was originally intended improve access to college for disadvantaged and minority students but Cullen et al. (2013) document that this led some students to strategically select high schools so that they could end up among the top 10 percent.

A third design question related to priorities is the desirability (or not) of coordination on priorities. Several countries, and especially England where school admissions cover both publicly and privately operated schools, leave significant freedom to schools to decide on their priorities. Are school-specific priorities a good thing? The existing literature has indirectly tackled that question when looking at the pros and cons of common, versus independent, tie-breaking rules across schools (see e.g. Abdulkadiroğlu et al., 2009). The common understanding at this stage is that a single tie-breaking rule allows the system to satisfy more top choices but sometimes at the cost of leaving more students unassigned. The size of the effect is empirically small. How these results translate to school-specific priorities is an open question. A priori, school-specific priorities may reduce the correlation of priorities, akin to the effect of multiple tie-breaking rules in the existing literature. However, because school-specific priorities do not operate at the margin, unlike tie-breaking rules, they may have a greater impact on the allocation. Additionally, priorities are strategic choices of schools. In a world where parents are uncertain about their preferences over schools, they can act as signals. Moreover, priorities generate an externality on the student intake of other schools. These elements suggest that schools’ self-interest may not coincide with the interest of the school district or students as a whole.

Table 2: New design questions from endogenizing preferences, priorities and capacities

Preferences	Priorities	Capacities
Information to parents	Translation of political objectives into priorities (or choice rules)	Capacity expansion as part of the mechanism
User interface and decision tools	Eligibility criteria	Strategic curriculum design
Incentives to acquire information built in the procedure	Coordinated versus independent priorities	
Impact of the allocation on preferences and, in turn, on outcomes		
Priorities as signal		

3.3. Flexible capacities

School capacities are largely fixed in the short run but it is often easy to add one or two students in class if needed. Currently, capacity is fixed in most school enrolment procedures, even if capacity can be adjusted ex-post (after the assignment has been done) if needed to accommodate unassigned students when education is compulsory and the procedure covers public schools (Shi, 2015 and Basteck et al., 2015, describe the process for doing so in Boston and in Frankfurt and Berlin respectively). Obviously, such ex-post capacity adjustments are inefficient because the added capacity might have been more valued by another student who ended up assigned.

Alternatively, some school districts use slightly inflated capacities during their assignment procedures to account for students failing their year or leaving during the summer. Clearly, inflating school capacities can only smooth the allocation problem. An open question is to what extent more significant capacity expansions (such as the construction of modular classrooms) could be systematically built into the algorithm before deciding on the final assignment.

Another way in which capacities can be seen as flexible is the design of school curricula. When schools see themselves competing for students, curriculum design is a choice variable that they can easily change. Anecdotal evidence from Belgium and France suggests that schools use curriculum design to attract or deter specific student types. Offering Latin and Greek attracts students from advantaged background. So do bilingual programs. The “industrial organisation” of school competition is an interesting open question.

In the long run of course, new schools can be built. One benefit of school enrolment procedures that elicit truthful preferences is that these data can be used to identify where and what kind of new capacity is needed.

4. Current challenges

Armed with this richer view of school enrolment policies and their inputs, we are now ready to revisit some of the big outstanding policy (and research) debates regarding school choice.

4.1. The scope of the market

In many cities, different school admission schemes run in parallel, with typically a centralized scheme for public schools and separate processes for private schools. Parallel schemes mean that multiple registrations cannot be prevented. This leads to no-shows in September and inefficient assignments. In New York City for example, no shows represent about 10% of the assigned places (Feigenbaum et al., 2017). Parallel schemes also lead to uncoordinated enrolment procedures and

increase the cost of applying for parents. It is obvious therefore that students would benefit from a single integrated school enrolment procedure. Such integrated schemes exist in several countries including Belgium, England, Hungary and Romania. In other countries, the debate is open. For example, the integration of the so-called “charter schools” – publicly financed but independently managed schools – into the centralized scheme for public schools is on the agenda in several US cities but progress is slow.

Why is integration not taking place? An obvious answer is political. These schools are unwilling to join a centralized coordinated scheme because they already attract the public they are interested in and entering the scheme could only lead to a loss of autonomy in the long run.

Understanding where priorities come from suggests other reasons. First, these schools may not want to bind themselves to the existing priorities for public schools because they do not reflect their preferences and the procedure in place cannot accommodate school-specific priorities. A related reason is that the procedure in place cannot guarantee that priorities are fully respected. Pathak (2016) explains that the New Orleans Recovery School District switched from the top trading cycle algorithm – which is efficient but may lead to violations of priorities – to a variant of the deferred acceptance algorithm in order to attract participation from neighboring schools that were keen on maintaining their existing admission criteria.

Ekmekci and Yenmez (2016) propose a more strategic reason for the lack of interest of schools outside of the main scheme to join. By staying outside, they lower the quality of the final assignment in the centralized match (relative to the situation where they would be included). This, in turn, increases applications to schools outside of the scheme, improving the quality of their intake. Absence of application costs is critical for this explanation. When applying to schools is costly, Espinoza et al. (2017) show that schools that have similar quality to those of schools already in the centralized scheme have an incentive to join as this increases the quality of their applicant pool. Using data from the Chilean university admissions, they find that decentralized admissions reduce both the participation by low income students and the efficiency of the allocation.

The governance of centralized enrolment procedures can be an answer. In Flanders for example, the Law provides guidelines for admissible priorities and procedures but leaves it to local school admissions “consultation platforms” to coordinate the details of the implementation. These platforms cover all publicly financed schools, whether publicly or privately run. All schools and other stakeholders of the local education system (unions, parents’ associations, minority associations, ...) are represented and decisions are taken at a majority. Experience shows that such platforms have helped schools understand their interdependence, i.e. the fact that admission practices in one school impacts others, and have fostered creative solutions while maintaining buy-in by all parties involved.

4.2. Diversity policies

At the heart of the diversity debate is the fundamental question of equality of opportunities. International surveys such as PISA emphasize both the average and the standard deviation of student performance and regularly point to the correlation between individual performance and socio-economic characteristics as an indicator of the inequity of education systems. School segregation is seen as one of the drivers to this source of inequity and it is therefore no surprise that diversity is high on the agenda in many countries. In France, for example, the recent Hollande government had made increasing social diversity in middle schools a priority. Options on the table include the redesign of catchment areas, changes in school curricula, coordination between public

and private schools and ... the introduction of choice.¹² In the UK, equality of opportunities is also central to the current debate on whether to expand the set of schools allowed to select their students (the co-called “grammar schools”) and what condition to impose on the socio-economic characteristics of their student intake.

Understanding how school enrolment procedures can contribute to diversity is an active area of research. It requires that we allow priorities to part of the design discussion. Early research showed that an inflexible interpretation of diversity is incompatible with stability. Intuitively, if we insist on an exact balance between different student types (for example: advantaged and disadvantaged), we may end up refusing a disadvantaged student even when there is space left, simply because no advantaged student wants to join the school to balance his presence (Kojima, 2012). Echenique and Yenmez (2015) have recently shown how diversity objectives can be defined to be compatible with stability. One special case is the ideal point rule where diversity objectives are defined as targets and admissions are decided based on preferences and priorities such that the final assignment brings the composition of the student intake as close as possible to that target. A practical application of this ideal point rule is the current double quota system introduced in Flemish schools in 2013. Two groups of students - advantaged and disadvantaged - are defined based on predetermined criteria. Places in each school are divided into two groups, in a proportion that corresponds to the target proportion of each type of students. Advantaged students have priority over disadvantaged students for the “advantaged places” and disadvantaged students have priority over advantaged students for the “disadvantaged places.” Students first fill the places that correspond to their types and then fill the other places, if space permits. This system brings the composition of the student intake as close as possible to the announced target, based on reported demand.

A challenge that all enrolment procedures that take parents’ preferences as an input face is to what extent they can significantly improve diversity. After all, if students from different types submit applications to different schools and capacity is sufficient, then priority structures designed to foster diversity will not alter the composition of the student body very much. Verhaeghe and Goetmaeckers (2013) evaluate the first year of application of the double quota system in the city of Ghent in Belgium. They find that diversity did increase in a majority of schools but overflow of excess demand from these schools led to a decrease in diversity in other schools due to the structure of preferences (see Wouters, 2016, for an extension to the whole of Flanders). This shows that well designed enrolment procedures can improve diversity. It is nevertheless wise to consider that going beyond these improvements will require acting on preferences.

5. Concluding comments

There is no such thing as “school choice” or “no school choice.” The details of how school choice is organized matter for how school choice meets the goal of taking parents’ preferences into account. The market design approach to school choice has proved useful in highlighting the trade-offs involved in each option and coming up with solutions that best meet the needs of school districts.

To a large extent however, the existing market design literature has focused on the redistributive dimension of the problem. Embedding the school choice problem into its wider context significantly enlarges the scope for market design. In this article, I have shown how endogenizing preferences, priorities and capacities can change the performance of school enrolment procedures and how it opens up an entirely new range of design questions, including how to translate policy goals into priorities, how to design enrolment procedures and user interfaces that support parents’ decisions

¹² Admissions to middle schools in France combines residence-based assignment for public schools and decentralized admissions for private schools. These parallel systems are thought to contribute to school segregation, especially in large cities where privately-operated, publicly-financed schools represent a sizable fraction of the market (34% in the Paris area).

and elicit preferences, and what level of coordination on priorities and admissions to choose. We are only at the beginning of this research agenda but the fact that many market design researchers are actively involved with school districts worldwide is promising.

References

- Abdulkadiroglu, A., Agarwal N. and Pathak P. (2015), 'The Welfare effects of coordinated assignment: Evidence from the NYC HS Match', *NBER Working Paper*, 21046
- Abdulkadiroğlu, A., Pathak, P., and Roth, A. E. (2009), 'Strategy-proofness versus efficiency in matching with indifferences: Redesigning the NYC high school match', *American Economic Review*, **99**(5), 1954-1978.
- Abdulkadiroğlu A., Pathak P., and Roth A.(2005), 'The New York City High School Match', *The American Economic Review*, **95**(2), Papers and Proceedings, 364-367.
- Abdulkadiroğlu A., Pathak P., Roth A., and Sönmez, T. (2006), 'Changing the Boston School Mechanism: Strategyproofness as equal access', *NBER Working Paper*, 11965.
- Abdulkadiroğlu, A., Che Y.K., and Yasuda, Y. (2011), 'Resolving the conflicts in School Choice: the Boston mechanism reconsidered', *American Economic Review*, **101**(1), 399–410.
- Abdulkadiroğlu, A., and Sönmez, T. (2003), 'School choice: A mechanism design approach', *American Economic Review*, **93**(3), 729-747.
- Abdulkadiroğlu, A., and Sönmez, T. (2013), 'Matching Markets: Theory and Practice', in D. Acemoglu, M. Arellano and E. Dekel (eds), *Advances in Economics and Econometrics*, Econometrics Society, vol 1, 3-47.
- Abdulkadiroğlu, A. (2013), 'School Choice', in: N. Vulkan, Alvin E. Roth and Zvika Neeman (eds), *The Handbook of Market Design*.
- Agarwal, N., and Somaini, P. (2016), 'Demand analysis using strategic reports: An application to a school choice mechanism', mimeo.
- Arnosti, N. (2016), 'Centralized clearinghouse design: A quantity-quality tradeoff', mimeo.
- Ashlagi, I., and Nikzad, A. (2015), 'What matters in tie-breaking rules? How competition guides design', mimeo.
- Ashlagi, I., Kanoria Y. and Leshno J. (2017), 'Unbalanced random matching markets: the stark effect of competition', forthcoming *Journal of Political Economy*
- Aygün, O., and Bó, I., (2016), 'College admission with multidimensional privileges: The Brazilian affirmative action case', mimeo.
- Bade, S. (2015), 'Serial dictatorship: the unique optimal allocation rule when information is endogenous', *Theoretical Economics*, **10**, 385–410.
- Balinski, M., and Sönmez, T. (1999), 'A Tale of two Mechanisms: Student Placement', *Journal of Economic Theory*, **84**(1), 73-94.
- Basteck, C., Huesmann, K., and Nax, H. (2015), 'Matching practices for secondary schools – Germany', *MiP Country Profile 21*.
- Burgess, S., Greaves, E., Vignoles, A. and Wilson, D. (2015), 'What Parents Want: School Preferences and School Choice', *The Economic Journal*, **125**(587), 1262-1289.
- Calsamiglia, C., and Güell, M. (2014), 'The illusion of school choice: Empirical evidence from Barcelona', mimeo.
- Calsamiglia, C., Fu, C., and Güell, M. (2016), 'Structural estimation of a model of school choices: The Boston mechanism versus its alternatives', mimeo.
- Calsamiglia, C., Martinez Mora, F. and Miralles, A. (2017), 'Sorting in public school districts under the Boston Mechanism', mimeo.
- Cantillon, E. (2015), *Matching Practices for Secondary Schools – Belgium (French-speaking region)*, MiP Country Profile 22.

- Coldron, J., Tanner E., Finch S., Shipton L., Wolstenholme C., Willis B., Demack S. and Stiell B. (2008), 'Secondary school admissions', report to the department of Children, Schools and Families (DCSF-RR020), 246 p.
- Cullen, J. B., Long, M.C., and Reback R. (2013), 'Jockeying for position: Strategic high school choice under the Texas' top ten percent plan', *Journal of Public Economics*, **97**, 32-48.
- Department for Education and Skills (2007), 'School Admissions Code', 132 pp., available at: <http://webarchive.nationalarchives.gov.uk/20071204130111/dfes.gov.uk/sacode/> (Accessed 03.06.2017)
- Department of Education (2014), 'School Admissions Code', 46 pp., available at: <https://www.gov.uk/government/publications/school-admissions-code--2> (Accessed 20.02.2017)
- Dur, U., Hammond, R. and Morrill, T. (2016a), 'Identifying the Harm of Manipulable School-Choice mechanisms', forthcoming *American Economic Journal: Economic Policy*.
- Dur, U., Kominers, S.D., Pathak, P. and Sönmez, T. (2016b), 'Reserve design: Unintended consequences and the demise of Boston's walk zones', mimeo.
- Dur, U., Mennle, T. and Seuken, S. (2016c), 'First-choice maximizing school choice mechanisms', mimeo.
- Echenique, F. and Yenmez, B. (2015), 'How to control controlled school choice', *American Economic Review*, 105(8), 2679-2694.
- Ekmekci, M. and Yenmez, B. (2016), 'Integrating schools for centralized admissions', mimeo.
- Ellison, G. and Pathak P. (2016), 'The efficiency of race-neutral alternatives to race-based affirmative action: evidence from Chicago's exam schools', mimeo
- Espinoza, R., Lee Soohyung, Lopez, H. (2017), 'Endogenous market formation: Theory and evidence from Chilean college admissions', mimeo.
- Feigenbaum, I., Kanoria Y., Lo I., Sethuraman, J. (2017), 'Dynamic matching in school choice: efficient seat allocation after late cancellation', mimeo.
- Gale, D. and Shapley, L. (1962), 'College admissions and the stability of marriage', *American Mathematical Monthly*, **69**(1), 9-15.
- Haeringer, G. and Klijn, F. (2009), 'Constrained school choice', *Journal of Economic Theory*, **144**(5), 1921-1947.
- Harless, P., and Manjunath, V. (2017), 'Learning matters: Reappraising object allocation rules when agents strategically investigate', *International Economic Review*, forthcoming.
- Hastings, J., Kane, T. and Staiger, D. (2009), 'Heterogeneous Preferences and the Efficacy of Public School Choice', *NBER Working Paper*, 2145.
- Hastings, J. and Weinstein, J. M. (2008), 'Information, school choice, and academic achievement: Evidence from two experiments', *The Quarterly Journal of Economics*, **123**(4), 1373-1414.
- He, Y. (2016), 'Gaming the Boston School Choice Mechanism in Beijing', mimeo.
- Hiller, V. and Tercieux, O. (2013), 'Matching practices in secondary schools – France', MiP country Profile 16.
- Karbownik, K. (2016), 'The effects of student composition on teacher turnover: Evidence from an admission reform', mimeo.
- Kesten, O. (2010), 'School choice with consent', *Quarterly Journal of Economics*, **125**(3), 1297-1348.
- Kojima, F. (2012), 'School choice: impossibilities for Affirmative Action', *Games and Economic Behavior*, 75(2), 685-93.
- Koning, P. and Van der Wiel, K. (2013), 'Ranking the Schools: How quality information affects school choice in the Netherlands', *Journal of the European Economic Association*, **11**(2), 466-493

- Lauri, T., Pöder, K. and Veski, A. (2014), 'Matching practices for elementary schools – Estonia', *MiP Country Profile 18*.
- Lovenheim, M. and Walsh P. (2014), Does choice increase information: Evidence from online school search behavior, mimeo
- Pathak, P. (2016), 'What really matters in designing school choice mechanisms', invited lecture to *11th World Congress of the Econometric Society*.
- Pathak, P. and Sönmez, T. (2008), 'Leveling the playing field in school choice: sincere and sophisticated players in the Boston mechanism', *American Economic Review*, **98**(4), 1636-52.
- Pathak, P. and Sönmez, T. (2013), 'School admissions reform in Chicago and England: Comparing mechanisms by their vulnerability to manipulation', *American Economic Review*, **103**(1), 80-106.
- Pop-Eleches, C. and Urquiola, M. (2013), 'Going to a better school: effects and behavioral responses', *American Economic Review*, **103**(4), 1289-1324.
- Pycia, Marek (2017), 'Invariance and matching market outcomes', mimeo
- Rothstein, J. (2006), 'Good Principals or Good Peers: Parental Valuation of School Characteristics, Tiebout Equilibrium, and the Incentive Effects of Competition Among Jurisdictions', *American Economic Review*, **96**(4), 1333-1350.
- Salonen, Mikko (2014), 'Matching Practices for Secondary Schools – Finland', *MiP Country Profile 19*.
- Shapley, L., and Scarf, H. (1974), 'On cores and indivisibility', *Journal of Mathematical Economics*, **1**(1), 23-37.
- Shi, P. (2015), 'Guiding school-choice reform through novel applications of operations research', *Interfaces* **45**(2), 117-132.
- Terrier, C. (2014), 'Matching practices for secondary public school teachers – France', *MiP Country Profile 20*.
- Verhaeghe, J-P and Goetmaeckers J. (2013), Evaluation of the enrolment procedure in Ghent elementary schools (in Dutch), evaluation report submitted to the Flemish Community.
- Wouters, Thomas (2016), Segregation and school enrolment policy, mimeo