学位請求論文要約

慶應義塾大学大学院経済学研究科

畠山 哲太郎

About the Thesis

This thesis ("Essays in Market Design") consists of four essays in *market design,* which focus on both theoretical and practical issues in real-world *matching markets* such as high school or college admissions. In the situations we study, a central planner (policymaker) allocates scarce resources (school seats) to students via a direct mechanism. Resources are indivisible and monetary transfers are prohibited to exclude the effects of wealth inequality. The planner has several potential objectives such as allocative efficiency or fairness and student incentives in reporting their preferences. Which objective to emphasize depends on the situation.

The projects in this thesis are conducted under several different motivations: For example, Chapter 2 provides an answer to an open question in the literature whereas Chapter 4 designs desirable market institutions using mechanism design approach. The following table summarizes the contents.

Table 1: Summary of contents

	Central objective			Type of	Type of
Chapter	Efficiency	Fairness	Incentives	question	analysis
1	O		O	Practical	Mechanism design
2	O		O	Practical	Equilibrium analysis
3	O			Theoretical	Equilibrium analysis
4		O		Practical	Mechanism design

Chapter 1: “When is a sequential school choice system (non-)deficient?”

In recent decades, market design has been successful in analyzing and improving real-life resource allocations where monetary transfers are not allowed, for example, school choice (Abdulkadiroglu and Sonmez 2003) and college admissions (Balinski and Sonmez 1999). While much of the existing literature concentrates on (and contributes by) designing *static* allocation systems, several systems in reality allocate resources *sequentially* in the sense that some are allocated first (this month) and the remaining are next (the next month).

In many regions of the United States, for example, Boston, Chicago, and New York, student enrollment systems for exam, charter, magnet, private, and specialized schools are operated separately from the enrollment system for regular (mainstream) schools with different timelines. Every student can participate in multiple systems where different systems adopt different mechanisms (e.g., student-proposing deferred acceptance or serial dictatorship) and different priority structures (e.g., lottery numbers or exam scores). Therefore, one student may participate in the exam school system and proceed to the regular school system only when she is not assigned to any exam school, while another student may only participate in the regular school system.

In China, the college admissions system comprises several stages where admissions for top-tier (the most prestigious) colleges are followed by admissions for lower-tier (less prestigious) colleges. In Japan, a similar sequential system is used in employment exams for public officers, that is, exams for the most prestigious national employee positions are followed by those for less prestigious prefecture/city-level employee positions. Although extensively practiced, as shown in Dur and Kesten (2019), a sequential enrollment/admission has deficiencies, including inefficiency of allocations (wastefulness) and student incentives in reporting true preferences (non-straightforwardness).

These examples pose a puzzle: a sequential assignment is not desirable in theory, but it is so prevalent in practice. In fact, static resource allocation is not that easy in some practices. According to Manjunath and Turhan (2016), a journalist says that it is difficult to unify public and private school enrollment systems in Milwaukee because they are in competition for the same students and have no incentives for coordination. Toward addressing this puzzling observation, we ask if there are circumstances when these deficiencies can be avoided. We show that despite the general result, under a restrictive but realistic condition (the *tiered* preference domain), a sequential assignment overcomes the above deficiencies. More interestingly, outside the tiered domain, the general impossibility result always holds (Proposition 2). Thus, our results also highlight the difficulty of sequential assignments in general diverse domains.

By slightly modifying the setting of Dur and Kesten (2019) and Andersson et al. (2018), this paper introduces a problem of designing an *allocation schedule*: Given a fixed number of stages (available dates) in which objects can be allocated, a mechanism designer (a local government) chooses in which stage to allocate each object to agents. The designer has two standard objectives: non-wastefulness (no object should be disposed if there is someone who wishes it) and straightforwardness (no lie should be profitable). We assume that objects (schools) are not strategic and they can coordinate enough to follow a fixed schedule. To focus on the design of the allocation schedule, we assume that each mechanism used at each stage satisfies certain desirable properties. If we pick one stage and allocate all the objects in that stage, because it is a static system, we can readily achieve both non-wastefulness and straightforwardness. While we wish to allocate objects in two or more stages, Proposition 1 shows a negative result in general domains. This means that there are cases in which any schedule comprising multiple stages leads to wastefulness and non-straightforwardness. Afterward, we search for a condition for a positive result, that is, a condition under which non-wasteful and straightforward allocation schedules exist. The main result is that a special type of domain, the so-called "tiered domain," is necessary and sufficient for a possibility (Proposition 2). Roughly speaking, this domain contains two types of profiles: the one with all agents partially agreeing on which objects are good (bad) and the other with agents preferring only specific objects to being unassigned.

Proposition 2 gives us an implication on when sequential assignment costs (does not cost) wastefulness or non-straightforwardness. The tiered domain is highly restricted but it naturally arises under specific circumstances. In Chinese college admissions (Japanese employment exams), it is commonly presumed that everyone prefers a more prestigious college (position). As long as these presumptions are realistic, prestige-based preference tiers naturally arise, and thus, non-wastefulness and straightforwardness are both achievable via a sequential assignment. Interestingly, the allocation schedule in reality is actually "better tiers in earlier stages," which achieves non-wastefulness and straightforwardness. This observation indicates that the schedule in reality may not be problematic in theory. In contrast, under more general, diverse preference profiles where no tiers exist, such a positive result does not hold.

Chapter 2: "School choice with limited student commitment" (coauthored with Kenzo Imamura)

Within the last two decades, school choice has become a widespread practice in educational policy all over the world. Both theory and evidence support school choice by allocative efficiency and fairness of school seats as scarce resources. In such a practice, a district allocates school seats to students based on student submitted preference ranking over schools ("preference list") without monetary transfers. Both in theory and practice, the central allocation mechanism is the so-called *deferred acceptance* mechanism (Gale and Shapley 1962), which is strategy-proof for students and fair in eliminating justified envies among students.

In many districts, students need to follow the constraint on the maximal length of preference list. This constraint makes real-life allocation mechanism ("constrained" deferred acceptance) different from standard ("unconstrained") deferred acceptance under which each student can report any number of schools. Moreover, a constrained deferred acceptance fails to satisfy desirable properties of the deferred acceptance. Since students face a strategic choice among including popular and non-popular schools in their list, this mechanism is clearly not strategy-proof. The allocation the mechanism outputs is not assured to be fair (at least under true preferences).

Faced to the gap lying between theory and practice, this paper asks the following question: Is there any justification for the use of a *constrained* deferred acceptance, rather than the *unconstrained* deferred acceptance? The answer of this question is crucial for us, market designers, when we advise policymakers. If there is a convincing justification, we can delegate the choice among constrained and unconstrained mechanisms to each policymaker; Otherwise, we should recommend unconstrained deferred acceptance.

Stylized models of school choice employ implicit assumptions on student choice behavior. A major assumption is about student commitment on consuming what she is assigned in the mechanism. In other words, we assume that students participating in the school choice system never *withdraw* their assigned school seats. While this assumption seems a natural proxy of reality, the following nature of the problem indicates the potential impact of withdrawal behavior on the (fair) allocation of school seats.

1. Students have *outside options* such as seats for private or specialized schools, which are allocated outset of the school choice system, or home schooling options.
2. Students are not totally confident about their own preferences. Instead, they learn their own preferences through costly *information acquisition* activities.
3. Given 1 and 2 above, it can be a natural best response for a student to secure a school seat first and then consider whether to attend that school or to consume the outside option.

In a practice similar to school choice, withdrawals are a major concern of the policymaker. In AY2022-2024, the authors were consulted by two Universities in Japan providing foreign exchange programs to undergraduate students. Their practice of assigning students to programs is essentially the same as school choice. Especially, they both use a constrained deferred acceptance mechanism. One of the Universities experienced 5-10% of annual withdrawal rate these years. Administrative managers of these Universities tell us the following: They wish and try to fill all the capacities for their partner institutions they have. However, since some students withdraw their assigned seats, several seats remain vacant. To make matters worse, some partner institutions impose a penalty on withdrawals. Once its student withdraws after matched to that institution, its capacity for that institution would be reduced at least for several years.

We conjecture that constrained deferred acceptance may be more immune to student withdrawal behavior than the unconstrained deferred acceptance. Our main result (Theorem 1) provides a support for this conjecture, that is, Theorem 1 gives a rationale to the use a constrained deferred acceptance. The main logic behind the result is quite simple: Since students can report a small number of schools, it is not optimal for students to report a school which may become unacceptable. Instead, they should fill the list by schools that are more likely to be acceptable.

This paper provides a new theoretical model in which students strategically choose their withdrawal behavior as well as reported preference list on schools. To capture the strategic situation each student faces, we employ a Bayesian framework of incomplete information about others' types. Each student type is a pair of valuations over schools (and outside option) and priority ranks for schools. Type distribution has a full-support on a rich domain and it is not necessarily i.i.d. Key features of the model are the following.

1. The uncertainty about each student's own preference is represented by the uncertainty about the value of the outside option. The true value realizes after school seats are assigned to students following a known distribution. Some schools can be acceptable or unacceptable depending on the true value of the outside option.
2. Before the preference submission, each student can learn the true value of the outside option (thus her true preferences) by a costly information acquisition ("search") activity.
3. After a school seat is assigned, the value of the outside option realizes and each student chooses which one to consume. If he chooses the outside option, assigned school seat is withdrawn/disposed.

We examine student equilibrium behavior on withdrawals under constrained or unconstrained deferred acceptance. Specifically, we define a new property of a mechanism: A mechanism is *withdrawal-proof* if, as long as the search cost is small enough, in any Bayes Nash equilibrium, no student disposes her assigned seat with positive probability. Proposition 1 characterizes when withdrawal-proofness holds under two extreme mechanisms: when there are m schools, the one is the unconstrained deferred acceptance (DA^m) and the other is the constrained deferred acceptance where each student can submit at most *one* school (DA^1). Proposition 1 shows that students are much more likely to dispose what they are assigned under DA^m than they are under DA^1.

We then compare the unconstrained deferred acceptance DA^m with practical class of constrained deferred acceptance DA^k (k=1, 2, ..., m-1) where each student can submit at most k schools. Using the measure on mechanism's vulnerability to withdrawals we propose, Theorem 1 shows that any constrained deferred acceptance DA^k is *less vulnerable to withdrawals than* the unconstrained deferred acceptance DA^m. In line with our conjecture, Theorem 1 gives rationale for constrained deferred acceptance mechanisms. To our best knowledge, this result provides a new advantage of a constrained mechanism.

Combined with well-known results in the literature, our finding implies a new trade-off in mechanism selection for school choice problems. Under complete information, the unconstrained deferred acceptance outperforms constrained deferred acceptances in many aspects. Especially, the former is fair (eliminates justified envy) and non-manipulable (strategy-proof) while the latter is not. However, on the vulnerability to withdrawals under incomplete information, any constrained deferred acceptances are better than the unconstrained deferred acceptance. We emphasize that our results do not subtract the nice properties of the unconstrained deferred acceptance. Instead, the results provide a richer options to policymakers. When withdrawals are a critical issue for the policymaker, we can recommend the use of a constrained deferred acceptance.

Chapter 3: "Cardinal utilities and welfare comparison of matching mechanisms" (coauthored with Kenzo Imamura)

Matching theory has played a crucial role in institutional reforms in real-life markets. Particularly, school choice systems are one of the most successful applications of matching theory. Over the past twenty years, many school districts worldwide have reformed their systems based on this theory. The deferred acceptance (DA) mechanism and the immediate acceptance (IA) mechanism are the two standard mechanisms in practice.

In the literature, the advantages of the DA over the IA has been widely studied (Abdulkadirouglu and Sonmez 2003). The DA is a strategy-proof mechanism, whereas the IA is not. Strategy-proofness is essential from an equity perspective, as it ensures that less informed or less sophisticated students are not disadvantaged compared to their more sophisticated peers (Pathak and Sonmez 2008). In addition, the DA is stable (or fair), whereas the IA is not. These discussions have contributed to shifts from the IA to the DA in school choice system reforms.

The IA has advantages over the DA. One of the main advantages of the IA is its ability to take into account students' *cardinal utilities*. Under the IA mechanism, school seats are allocated according to each student's ranking of the school: students who rank a school first are admitted first, followed by those who rank it second if seats are still available, and so forth. This feature implies that a student faces a risk when top-ranking a popular school: while it reduces a student's chances at their less preferred schools, a popular school rejects them with a high probability. Thus, the IA is not strategy-proof. However, this feature allows the IA to screen students who strongly prefer popular schools from those who do not strongly prefer popular schools over others. Abdulkadirouglu et al. (2011) showed that the IA has a welfare gain over the DA: in a stylized environment, every student is weakly better off in any symmetric equilibrium of the IA compared to the dominant strategy equilibrium of the DA.

We study this welfare gain. In practice, many mechanisms other than the DA and the IA are widely used. For instance, the Chinese parallel mechanisms are implemented in college admissions in China (Chen and Kesten 2017), while the Taiwan mechanisms are used for high school admissions in Taiwan (Dur et al. 2022). These mechanisms, like IA, are not strategy-proof (except for extremes). In addition, similar to IA, they appear to account for cardinal utility. This raises the following question: do these mechanisms offer welfare gains over the DA?

We show that two conditions, *non-wastefulness* and *anonymity*, are key to achieving this welfare gain. Let us explain this in more detail. We follow the setting of Abdulkadirouglu et al. (2011). First, every school has no priority. Therefore, it is necessary to tie-break using methods such as lotteries. Next, we assume students have identical ordinal preference rankings while allowing their cardinal preferences (vNM utilities) to differ. Lastly, we consider an incomplete information setting. Specifically, students know neither the priorities (after tie-breaking) nor preferences of other students. In this situation, we show that if a mechanism satisfies non-wastefulness and anonymity, then every student is weakly better off in any symmetric equilibrium of the mechanism compared to the dominant strategy equilibrium of the DA.

An implication is that many mechanisms can achieve the same welfare gains as that of the IA. Since we assume no priorites, anonymity is a weak condition. As long as random tie-breaking is used, the mechanism satisfies this property. In addition, all of the above mechanisms satisfy non-wastefulness.

We discuss mechanisms that violate either of these properties. First, we consider the constrained mechanisms that limit the length of the list students can submit. These mechanisms are widely used in practice. For instance, when the DA and the IA are used in practice, these constraints are almost always imposed. A constrained mechanism does not satisfy non-wastefulness. We illustrate that the welfare gain may not be achieved under a constrained mechanism. Next, we consider mechanisms that do not satisfy anonymity. For instance, such mechanisms may arise if tie-breaking is based on a deterministic way. Similarly, welfare gains may not be achieved under these mechanisms.

Chapter 4: "Who makes the cut? Endogenous priority design for heterogeneous groups of agents" (coauthored with Onur Kesten and Morimitsu Kurino)

In numerous real-world allocation scenarios, priorities serve as the cornerstone for distributing sought-after goods and services. This is evident across various sectors: healthcare professionals must arrange patient surgeries, government authorities must determine immigration priorities, educational institutions must rank applicants, and during crises like pandemics, critical resources such as vaccines and ventilators must be allocated according to predetermined criteria.

The vast literature on assignment problems often take priorities as a primitive given exogenously at the outset and focus on the design of institutions for equitable and efficient allocations in light of given priorities. Nevertheless, it is inevitable that carefully designed institutions on the premise of equity can only be fair inasmuch the underlying priorities are the result of a transparent, objective and consistent protocol.

In this paper, we are interested in the problem of *endogenously* designing priorities to align with the primary objectives of a planner. Unlike situations where agents share a common attribute, such as test scores for university admissions or arrival times for service queues, prioritization becomes notably challenging when agents possess diverse attributes or belong to different identity groups. For instance, a major dilemma faced by policymakers during the COVID-19 pandemic pertained to determining the vaccination order among various groups of individuals in society including healthcare personnel, essential workers, the elderly individuals etc. The absence of consensus led to the emergence of varied rationing protocols in various countries and regions, highlighting the complexity of prioritization in heterogeneous populations (Pathak et al. 2024).

Similar examples are abundant. When scheduling surgeries at a hospital, prioritizing patients becomes imperative. Within each surgical category, scoring-based systems are often utilized to rank patients based on the urgency and severity of their conditions, facilitating the allocation of limited surgical resources effectively. In allocating research funds across disciplines, government bodies determine which applications to award grants. Each application undergoes review and ranking within its respective discipline, with a fixed total number of grants available. The challenge lies in deciding which researchers and disciplines should receive funding, taking into account the unique needs and priorities of each field. In allocating visas or green cards, applications from various job categories or nationalities are considered. Within each job category, candidates are evaluated based on merit, and a fixed number of visas or green cards are allocated. Allocating sabbatical positions within a university faculty presents similar complexities, as applicants come from diverse departments seeking leaves at the faculty level. Decisions regarding sabbatical awards are made centrally, with applicants within each department ranked against their peers. Here, the task is to prioritize individuals for sabbaticals, considering the academic contributions and needs of each department while maintaining fairness. Furthermore, affirmative action policies in school admissions highlight the importance of prioritizing students from various backgrounds, such as ethnicity and socioeconomic status. When faced with binary choices, schools aim to admit students from multiple applicant types to promote diversity and equity. However, in scenarios involving more than two groups, schools must also determine the composition of each class, balancing factors like academic merit and inclusivity to create a diverse student body that reflects the institution's values. These scenarios illustrate the multifaceted nature of prioritization in allocation problems, ranging from binary choices to complex, multi-group considerations. Throughout these allocation processes, the challenge lies in devising prioritization mechanisms that account for the heterogeneity of agents while ensuring fairness and efficiency.

In this paper, we aim to elucidate the principles behind endogenously designing priorities in allocation problems, offering insights into assignment problems at the intersection of economics, management policy, and social justice. Our formal model builds on a problem faced by universities in Japan to assign students from various faculties to exchange programs. Keio University in Tokyo specifically consulted the third author to develop a protocol to solve this problem. This application constitutes the basis for our canonical model and the terminology we adopt in the rest of the paper.

Universities all over the world have exchange programs in which millions of students with various majors study abroad at partner universities. The common allocation method for assigning students to exchange programs is a choice based on priority, also commonly referred as a *serial dictatorship* mechanism in the assignment literature. Before applying this mechanism, a complete priority order must be constructed. While GPA is a natural basis for determining priorities, the process is further complicated by the fact that students actually come from different faculties with different class sizes and different grading metrics. Universities often address this complication by resorting to ad-hoc internal decision processes. For example, universities such as Keio usually hold face-to-face interviews with each student to prioritize all exchange program applicants. Nonetheless, the sheer administrative cost of conducting such interviews led to the reform in which the centralized prioritization process proposed in this paper was adopted.

We formulate a *prioritization problem* as the problem of assigning a set of agents to an exogenously given K priority groups of arbitrary sizes. Each agent prefers to be in a higher (lower-indexed) group. Each agent inherently belongs to a category, e.g., the faculty a student is affiliated with. Existing priorities are only partial in that only agents from the same category can be ranked, i.e., any two agents across different categories are incomparable. A systematic solution to this problem is a prioritization *protocol* that returns a complete and transitive (not necessarily linear) priority over agents, or equivalently, an assignment of agents to K priority groups. When agents are otherwise indistinguishable, to achieve fairness in an ex ante sense we also allow for prioritization protocols that allow for lotteries. Our goal is to identify the types of prioritization protocols that emerge if one insists on minimal forms of efficiency and fairness. Importantly, when prioritization becomes a central design issue, it expands the concept of *choice rule* that have attracted much attention in recent matching theory and market design.

As a minimal form of efficiency, we consider *non-wasteful* assignments where no student is assigned to a group while a higher-priority group has a vacancy. An assignment *respects priorities* if no student is assigned to a higher priority group than a higher priority student from the same category. Our first result is a complete characterization of the set of non-wasteful deterministic assignments that respect priorities (Proposition 1).

We consider two plausible metrics for gauging the fairness of assignments depending on how one evaluates two students from different categories. A natural starting point is an *absolute* priority metric, which considers the rank of the student within her category, regardless of the category's size. In many applications, population size is a major determinant for the resource distribution. For example, in public health funding, resources like vaccines, hospital beds, medical staff, and funding are distributed proportionally to the size of different areas or groups. Government funding for public schools is often allocated based on the number of students. Public services such as police, fire departments, and public transportation allocate resources based on the population size of different districts or neighborhoods. Seats in legislative bodies (e.g., House of Representatives in the US) are often allocated based on population size. Motivated by the abundance of applications where population size matters, we define a *relative* priority metric. It refines the absolute priority metric by considering the percentiles agents in which agents lie with respect to their categories, i.e., rank is defined relative to the size of the category.

Our central fairness notion for prioritization protocols is *equal treatment of students with equal priority metrics*: any pair of students who have the same priority metric must receive the same probability of getting assigned to a given priority group.

Our main result shows that two types of prioritization protocols stand out when the planner requires non-wastefulness, respecting priorities, consistency and equal treatment of equals while using randomization only to break ties. These protocols assign probability distributions over deterministic assignments. Each deterministic assignment in the support is obtained by allowing students to successively compete to fill the positions in the priority groups starting with the highest priority group moving down to the lowest. Each student is equipped with a "filling speed." Students fill positions based first on priority and then speed. When two students from two different categories are tied to fill a position in all relevant characteristics, a tie-breaker determines who gets the position.

When fairness is pursued in an absolute sense, the protocol has to be a *uniform prioritization protocol*: each student's speed is identical and a uniformly random tie-breaker is used (Theorem 1). On the other hand, when fairness is pursued in a relative sense, e.g., students in similar percentiles in their respective categories should be treated the same, the protocol has to be a *proportional prioritization protocol*: each student's speed is proportional to the size of her category and a uniformly random tie-breaker is used (Theorem 2).

It is important to emphasize that the protocols we characterize "minimally" rely on randomization. More specifically, the filling algorithms that constitute the backbone of the two protocols do not treat the positions in the priority groups as a continuum mass of students. Instead, protocols allow students to fill positions one by one to generate deterministic assignments and randomization is applied only "locally" to students that are otherwise indistinguishable.

A major challenge in our approach is that the fairness notion requires making comparisons across deterministic assignments, e.g., do similar agents appear in similar groups with same frequency? This makes working with random assignments (i.e., assignment matrices showing the probabilities with which category members end up in which group) rather intractable. To overcome this challenge, one needs to directly construct lotteries over deterministic assignments rather than resorting to standard machinery of random assignment decomposition, e.g., Birkhoff von Neuman type of methods (Birkhoff 1946; von Neumann 1953) are inapplicable. In this sense, despite the use of similar terminology, our protocols are markedly different than well-known stochastic rules such as the probabilistic serial (Bogomolnaia and Moulin 2001) or competitive equilibrium (Hylland and Zeckhauser 1979) that have attracted much attention in the assignment literature.