

Costly information acquisition in centralized matching markets*

Rustamdjan Hakimov[†] Dorothea Kübler[‡] Siqu Pan[§]

March 22, 2023

Abstract

When applying to university, students and their parents devote considerable time acquiring information about university programs in order to form preferences. We explore ways to reduce wasteful information acquisition, that is, to help students avoid acquiring information about out-of-reach schools or universities, using a market design approach. Focusing on markets where students are ranked by universities based on exam scores, we find that, both theoretically and experimentally, a sequential serial dictatorship mechanism leads to higher student welfare than a direct serial dictatorship mechanism. This is because the sequential mechanism informs students about which universities are willing to admit them, thereby directing their search. Our experiments also show that the sequential mechanism has behavioral advantages because subjects deviate from the optimal search strategy less frequently than under the direct mechanism. Furthermore, providing historical cutoff scores under the direct mechanism can increase student welfare, especially when the information costs are high, although the observed effect is weaker than that of a sequential mechanism.

*Our special thanks go to Nina Bonge, who helped us program the experiments, and to Jennifer Rontganger for copy editing. We are grateful to the editor and three referees for their insightful and constructive comments and clear guidance. We would like to thank Inácio Bó, Jacob Leshno, Ignacio Rios, Renke Schmacker, Ran Shorrer, Xin Zhao, Georgy Artemov, and the participants of the market design seminar at the University of Lausanne, the 17th Matching-in-Practice workshop, the ESA 2021 Global Around-the-Clock Virtual Conference, the 20th Annual SAET Conference, the Social, Behavioral and Experimental Economics (SBEE) seminar at the University of Michigan, and the seminars at Wirtschaftsuniversität Wien, Queensland University of Technology, Monash University, and the University of Alabama for their valuable comments. Financial support from the DFG through CRC TRR 190 (Kübler), the Leibniz SAW project MADEP (Kübler), the University of Melbourne (Pan), and the Swiss National Science Foundation project #100018_189152 (Hakimov) is gratefully acknowledged.

[†]University of Lausanne & WZB Berlin Social Science Center, Internef 536, Quartier de Chamberonne, CH-1015, Lausanne, Switzerland; email: rustamdjan.hakimov@unil.ch

[‡]WZB Berlin Social Science Center, Technische Universität Berlin, and CESifo, Reichpietschufer 50, 10785 Berlin, Germany; email: dorothea.kuebler@wzb.eu

[§]Department of Economics, The University of Melbourne, VIC 3010, Australia; email: siqi.pan@unimelb.edu.au

Keywords: Matching market, deferred acceptance, serial dictatorship, information acquisition, game theory, lab experiment

JEL classification: C92, D47

1 Introduction

Every year, many students and parents devote considerable time and effort to screening universities and study programs. These activities include searching university websites and brochures, consulting university rankings and independent guidebooks, as well as talking to current students, alumni, and counsellors. Parents and students gather information regarding academic quality and the various programs offered by universities, about costs and scholarships, the facilities and location, housing availability, etc. This information helps them form preferences in the choice of the universities and programs. A similar behavior can be observed in school choice when parents and children have to decide which schools to put on their wish list.

If parents and prospective students search too little or in the wrong places, this can lead to an educational mismatch. It is often difficult to consider the right universities and to rank them properly. While search behavior has been studied in empirical work that takes the organization of the market as given, the question arises as to how admission procedures can be designed to facilitate the search process for students and direct them toward appropriate and realistic options.

Despite its importance, the need to collect information in order to form preferences in matching markets has not received much attention so far. The study of centralized and decentralized markets has mainly focused on the stability and efficiency of outcomes, and the incentive properties of the mechanisms. Regarding preferences, it is typically assumed that parents and students can rank universities at no cost. Such models are at odds with the observed activities of parents and students and the wealth of information on websites and in books and brochures. The assumption of costless preference formation has led to a strong emphasis on mechanisms that elicit the complete rank-order list from applicants. However, such mechanisms may turn out to be suboptimal when students have to first collect costly information on the universities to be able to rank them.¹

Our paper aims to study which matching mechanisms generate high student welfare when information acquisition to form preferences is costly. As a first step, we use a simple model to derive optimal search strategies and demonstrate welfare differences between mechanisms given optimal search strategies. As a second step, we empirically evaluate whether the mechanisms that perform better according to the theory are also superior in a laboratory experiment. We conduct the experiments because in our setting as well as in real life, search strategies can be complicated and their complexity may vary across different mechanisms. The experimental method allows us to investigate whether people search optimally under the different mechanisms, and thus which mechanism results in the highest welfare empirically.

To avoid wasteful information acquisition, it is important for parents and prospective students to search only among universities that are within their reach—that is, universities that would accept

¹It could be argued that the cost of information acquisition is low compared to the benefit of choosing a suitable study program. However, when the number of programs is high, gaining full knowledge of all programs is impossible or very costly in terms of time and effort. In this case, students have to decide which programs to acquire information on, which is captured by the search costs in the model.

the students in the course of the procedure. We say that such universities belong to the students' budget set. The budget set depends on the preferences or priorities of universities, determined, for example, by the rank in admissions exams or the school GPA, as well as the preferences and choices of other students. We focus on situations where the priorities of universities are aligned and are common knowledge, as is the case in single-exam or GPA-based university admissions where all universities rank students in the same way (e.g., in China, Turkey, Denmark, Sweden, Tunisia, and Germany). Additional applications include school choice based on a single average grade, for instance, in Berlin (Basteck et al., 2021); centralized labor markets, for example, for doctors at public hospitals in France;² and other allocations based on priority and queues, such as office allocations in new buildings (Baccara et al., 2012).

We model the formation of preferences as a costly process of information acquisition. Our search technology is motivated by centralized university admission systems that rely on students' ordinal university rankings, and allows students to learn their ordinal preferences of the universities.³ In our model, students' prior beliefs before a search follow a tiered structure. Specifically, all students prefer a university in a better tier to a university in a worse tier but their within-tier preference follows a uniform distribution, which means that all ordinal rankings of universities in a tier are equally likely. Universities in the same tier have the same capacity while those in different tiers may have different capacities.⁴

The exact search technology is not crucial for our main theoretical results regarding student welfare. However, we chose a technology that allows for a closed-form solution of the optimal search strategy. This is useful for the experimental implementation. In our model, each student can acquire information about her own within-tier preferences using the following search technology. In the first step of the search, the student can pay a cost to compare any two universities in that tier. Then, for an additional cost, a student can choose another university in that tier and learn how it compares to the two universities previously investigated. Thus, to learn the relative ranking of m universities in a tier, a student has to pay the information cost $(m - 1)$ times. Students can stop at any step in the search process and can choose to search in zero, one, or multiple tiers. We assume that students can also apply to universities for which they have not acquired information, since there is typically no such requirement in university admissions.

In our environment where all universities rank students in the same way, we consider two mechanisms for implementing the serial dictatorship rule, namely one where students simultaneously

²See <https://www.anemf.org/wp-content/uploads/2019/07/procedure-des-choix.pdf>, last accessed 4.7.2021.

³Ordinal preferences are enough to determine the optimal submission strategy, as we consider strategy-proof mechanisms. This would not be the case, for instance, in the Immediate Acceptance mechanism, where cardinal utilities might influence the equilibrium strategy.

⁴A tiered prior structure can be found in China where universities belong to exogenously determined tiers, see Chen and Kesten (2017), and where the ranking of universities within tiers differs between students. Even when there is no official categorization of universities into tiers, tiers can exist due to a general understanding regarding university quality. The assumption of an equal size of universities can be interpreted as a situation of uncertainty where students are not informed about the number of seats before acquiring information and therefore assume that capacities are equal.

submit their preference lists, called the Direct Serial Dictatorship (DirSD) mechanism, and one in which students sequentially select universities in their order of priority, called the Sequential Serial Dictatorship (SeqSD) mechanism.⁵

The DirSD mechanism has many desirable properties (Svensson, 1999; Abdulkadiroğlu and Sönmez, 1998), and variations of it are employed in Australia, China, Turkey, Greece, Denmark and Sweden for university admissions. In DirSD, all students simultaneously submit their rank-order lists to the clearinghouse. Then, the serial dictatorship algorithm is run where the student ranked highest is assigned to her most-preferred program, the student ranked second to her most-preferred program with open seats, and so forth. There is no opportunity to learn about the preferences and choices of other students. Students can only guess what their budget set is, and their optimal search strategies are based on their expectations regarding others' choices. We consider DirSD as the baseline mechanism, and study two approaches to improve the welfare of students relative to this baseline.

The first approach is a sequential mechanism, SeqSD where students take decisions sequentially in order of priority. The student with the highest priority chooses a university first, then the student with the second highest priority chooses a university among the universities that still have vacant seats, etc. Under this mechanism, students do not face any uncertainty about their budget set: when it is their turn, students observe their budget set, and can pick the most-preferred option. Importantly, students can acquire information about their preferences after they have learned their budget set in SeqSD. To our knowledge, this exact mechanism is not used for school or college admissions in practice, but some countries employ similar dynamic mechanisms in which students can observe their set of offers. For instance, France switched to a sequential university-proposing deferred-acceptance mechanism in 2018. In this mechanism, students receive offers from the universities over several weeks. Tunisia uses a three-step SD, in which the cohort is divided into three groups based on priority orders. Starting from the highest-ranked group, the three groups sequentially submit preference lists and are then assigned using the SD mechanism. After the assignment of each group, the remaining vacancies are published before the next group submits their preference lists. Germany and the Chinese province of Inner Mongolia also use dynamic mechanisms that reveal partial information on the budget set to the participants. SeqSD is also adopted to allocate faculty offices in buildings in US professional schools, with lower ranked participants observing the choices of their peers with higher priority (Baccara et al., 2012). To our knowledge, we are the first to investigate the effects of a sequential mechanism on student welfare when taking information acquisition into consideration.

The second approach is to provide historical cutoff scores under DirSD (Cutoff), which is motivated by a common practice in real-life markets. Whenerent cohorts have similar preferences over universities and similar distributions of exam scores, historical cutoffs contain useful information

⁵When all universities rank students in the same way, the outcome of the Serial Dictatorship rule is the same as that of the Top Trading Cycles rule and the Gale-Shapley Deferred Acceptance rule.

on the selectivity of the universities. Thus, by observing historical cutoffs, students receive information on their admission chances, that is, on their budget set. Cutoff is less time-consuming than SeqSD, especially in larger markets, because students search and choose offers one-by-one under SeqSD, but submit preferences simultaneously under Cutoff. However, given that SeqSD provides accurate information on the students' budget set, cutoffs can be less effective than SeqSD in reducing wasteful information acquisition if the distributions of preferences and exam scores fluctuate between years. For instance, Ajayi and Sidibe (2020) show that the correlation between the school cutoffs in 2007 and 2008 in Ghana was 0.84 for all schools, and as low as 0.37 for less selective schools. Information on the cutoffs is widely used in practice (see Immorlica et al. (2020) for an overview of college admission systems where the cutoffs are published). However, to our knowledge there is no empirical evidence yet showing the causal effects of historical cutoff provision on student welfare, and we aim to close this gap.

Based on a model of information acquisition and in line with the intuition described above, we show that student welfare under SeqSD is always higher than or equal to welfare under DirSD. This prediction does not depend on a particular search technology. It also holds true for every student in the market regardless of her priority. The model allows us to derive exact predictions concerning optimal search and submission strategies with potentially incomplete preference information under DirSD and SeqSD.

The theoretical results show that the sequential mechanism can improve student welfare compared to the direct mechanism. However, this benefit may not be fully captured by the theory. To derive the optimal search strategy in the direct mechanism, a student needs to form correct beliefs about the probability of each possible composition of her budget set and, based on these beliefs, to weigh the benefit of an additional search against its cost. This is a complicated problem to solve, and inexperienced participants in real markets may deviate substantially from searching optimally. The sequential mechanism significantly mitigates the search complexity by giving students precise information on their budget set. Thus, it might have additional behavioral benefits over the direct mechanism, which the theory does not capture.⁶ Cutoffs can also provide useful information on the budget set to students and thereby simplify the search. However, a student under Cutoff not only needs to form beliefs about the search strategies of previous generations, but also needs to consider how the reactions of students in her generation impact her budget set. This may be challenging for students, and an empirical test is warranted. Because search costs and student preferences are usually not observable, it is difficult to identify the optimality of search strategies from field data. Therefore, a laboratory experiment where we control both the costs and the preferences seems appropriate to address our research questions.

We designed experiments to compare DirSD, SeqSD, and Cutoff in a centralized university

⁶Note that SeqSD is behaviorally simpler than DirSD for the students, even when we assume students have complete knowledge of their own preferences, because it is obviously strategy-proof (Li, 2017). We study whether the behavioral benefits of SeqSD extend to the search strategies.

admissions experiment where we vary the market environment in terms of the monetary cost of information acquisition (high cost or low cost) and the priors of students regarding the quality of the universities (one tier or two tier). First, we find that student welfare is highest under SeqSD, second-highest under Cutoff, and lowest under DirSD, with all differences being significant. The improvement of SeqSD relative to DirSD is significant in all environments. SeqSD leads to higher payoffs from the resulting matching and lower information acquisition costs than DirSD. The improvement of Cutoff relative to DirSD is driven by high-cost environments. Cutoff and DirSD lead to similar average payoffs from the matching, but the information acquisition costs are significantly lower under Cutoff. We observe significant deviations from the optimal search strategies in both directions (over- and under-search) in DirSD and SeqSD. However, the deviations are significantly less frequent in SeqSD. Thus, SeqSD leads to higher welfare not only due to the information on the budget set but also due to fewer strategic mistakes in the search strategies. As for Cutoff, we do not have point predictions, but we observe that participants avoid information acquisition for universities with cutoff scores higher than their score, especially when the search cost is high. Moreover, compared to DirSD and Cutoff, subjects under SeqSD make fewer strategic mistakes in submission decisions given the search. Thus, in addition to the theoretically predicted benefits of SeqSD, it is easier for participants to follow the optimal search and submission strategies under SeqSD than under DirSD. If students have more complex priors than the tiered structure considered in this paper, say, students are not exactly sure to which tier a certain school belongs, the behavioral benefits of SeqSD can be expected to be even more important. This is because a more complex prior structure makes it more challenging for a student in DirSD to form beliefs about her own budget set when making search decisions, but this effect is absent in SeqSD as belief formation is not necessary in the sequential mechanism where the budget set is known and therefore independent of the priors.

Finally, we run treatments to test the robustness of our findings to the search technology and the level of search costs. Instead of the technology described above where participants learn about their ordinal ranking of the programs, they pay to learn the cardinal value of each university. We also choose two intermediate search costs between the low and the high cost in our original treatments. The results concerning the superior performance of SeqSD relative to DirSD are largely replicated. However, we find that the provision of cutoffs does not lead to significantly higher student welfare. While participants spent less money on information acquisition in Cutoff relative to DirSD, the payoffs from the matching to universities were somewhat lower under Cutoff than DirSD. This echoes the finding in our original treatments that the improvement of Cutoff relative to DirSD is driven by the high-cost environments and confirms that cutoff provision is effective when the cost of information acquisition is sufficiently high.

Related Literature

Our paper contributes to the recent literature on information acquisition in matching markets, which includes theoretical studies such as Bade (2015), Immorlica et al. (2020), Chen and He

(2022), and Artemov (2021), and experimental studies like Chen and He (2021) and Koh and Lim (2022).

Bade (2015) studies the house allocation problem and shows that when information acquisition for one’s own preferences is costly, serial dictatorship is the unique ex-ante Pareto optimal mechanism among all strategy-proof and nonbossy mechanisms. The paper allows for multiple levels of information acquisition through partitions of the state space but does not explicitly model the search process.

Immorlica et al. (2020) show that it is beneficial for students to know their budget set or set of feasible options before acquiring costly information about universities. This creates incentives to wait until the market has resolved before searching for information and accepting an offer. As a result, information deadlocks arise when there is a cycle of students in which each student’s information acquisition decisions depend on the demand of others. The analysis of Immorlica et al. (2020) suggests that facilitating efficient price discovery by publishing cutoffs can improve student welfare. We find empirical support for this theoretical result. Additionally, we provide causal empirical evidence of the difference between direct and sequential mechanisms with regard to information acquisition.

Chen and He (2021, 2022) compare students’ incentives to acquire information under the immediate acceptance mechanism and the student-proposing deferred acceptance mechanism. In both mechanisms, students have to submit rank-order lists upfront and do not receive information on their budget set. In their theoretical contribution, Chen and He (2022) show that only the immediate acceptance mechanism incentivizes students to learn their own cardinal and the other participants’ preferences. In experiments, Chen and He (2021) find that, overall, the willingness-to-pay for information is too high across treatments, lowering aggregate welfare. In contrast, we compare direct and sequential serial dictatorship mechanisms and study the effects of cutoff provision with respect to information acquisition and student welfare. Also, Chen and He (2021, 2022) model information acquisition as a binary choice: acquiring zero or full information. We develop a sequential search model in which agents can choose various stopping points. This captures search processes in many real-life scenarios, and it also provides us with rich data on search patterns.

Artemov (2021) finds that informational incentives provided by a random serial dictatorship mechanism fall short of the social optimum in most cases due to externalities in information acquisition, and he proposes policies to improve social welfare. Noda (2020) investigates the optimal disclosure policy regarding the choice set—that is, the set of objects available, under a random serial dictatorship mechanism. The paper concludes that the full-disclosure policy is typically Pareto inefficient due to the presence of a positive externality in information acquisition. Similar to Chen and He (2021, 2022), Artemov (2021) and Noda (2020) also simplify the search process to a binary choice of acquiring zero or full information. Harless and Manjunath (2018) consider a setup where information acquisition is costless but each agent can choose to learn his value for only one object. They show that the top trading cycles rule outperforms serial

dictatorship in terms of fairness, though the allocation might not be Pareto efficient. Bó and Ko (2020) consider colleges' incentives to acquire information on the quality of applicants. They show that when screening costs are low, all schools acquire more information about applicants, but this does not improve the quality of the admitted pool for the lower-ranked colleges, as the best students are more likely to be assigned to better colleges.

Recent empirical work on school choice (Narita, 2018) and university admissions (Grenet et al., 2022) provides indirect evidence of students searching and learning about their preferences during the application process. Narita (2018) studies re-application behavior for high schools in New York City, and documents that a considerable proportion of students who receive their first choice decide to re-apply. Such changes in demand create a welfare loss if they cannot be accommodated by the market. Similarly, Grenet et al. (2022) document that university programs whose offers are received earlier are more likely to be ranked higher than programs whose offers arrive later. This can be explained by the students' costly search regarding the programs.

Cutoff scores have been studied by Azevedo and Leshno (2016) in a two-sided matching framework with demand and supply. They show that cutoff scores can be interpreted as prices, such that at any vector of cutoffs that equates supply and demand, the demand function yields a stable matching. The cutoffs are also used in studies employing a regression-discontinuity design to estimate the causal impact of university or school attendance (see, for instance, Abdulkadiroğlu et al., 2014; Hastings et al., 2013; Zimmerman, 2019; Lufade, 2019). Ajayi et al. (2020) run a field experiment in Ghana providing participants in school choice with extensive information, including information on admission chances. The information provision changes the application behavior, but it is hard to conclude which part of the information intervention drives the effect. To our knowledge, our paper is the first to empirically investigate the effect of providing historical cutoffs on search and market outcomes.

Our paper relates to the recent literature pointing out the advantages of dynamic mechanisms. Li (2017) introduces the concept of obvious strategy-proofness and shows that truth-telling is an obviously dominant strategy under SeqSD, while this is not the case under DirSD. Pycia and Troyan (2019) show that SeqSD is the only mechanism which is efficient, fair, and obviously strategy-proof. Li (2017) compares SeqSD and DirSD in the lab and finds that a significantly higher proportion of participants use the truth-telling strategy under SeqSD than under DirSD. Klijn et al. (2019) and Bó and Hakimov (2020) present similar results for the comparison of sequential and direct versions of the deferred acceptance mechanism, and Bó and Hakimov (2023) for the top trading cycles mechanism.⁷ An additional argument in favor of dynamic mechanisms is that they can be more transparent for the participants than direct mechanisms (Hakimov and Raghavan, 2022), and, unlike direct mechanisms, they can reach a stable allocation when applicants have preferences

⁷Other related experiments on matching markets are surveyed in Hakimov and Kübler (2021) and Pan (2020). In an auction context, most empirical studies confirm that bidding behavior is more consistent with theoretical predictions in dynamic formats than in static formats. See, for instance, Kagel et al. (1987), Kagel and Levin (2001; 2009), and Engelmann and Grimm (2009).

not only over colleges but also over their peers (Cox et al., 2021).

Several papers analyze dynamic mechanisms used in practice where offers, acceptances, and information can be spread out over time. Bo and Hakimov (2022) and Gong and Liang (2016) analyze university admissions mechanisms used in Brazil and Inner Mongolia, respectively. Both mechanisms are sequential and include the provision of intermediate cutoff scores to students. Dur et al. (2017) analyze the submission mechanism used in the Wake County Public School System where parents can wait and observe how many others have applied to certain schools in order to gauge their chances of getting a seat. Luflade (2019) studies university admissions in Tunisia, where the SD mechanism is implemented in three sequential stages, and documents that the sequential implementation has a positive effect on the students' welfare.

A set of papers (Das and Li, 2014; Kadam, 2015; Lee and Schwarz, 2017) analyzes decentralized matching markets where agents acquire information about their preferences through interviews. The search and matching literature also explores the role of costly search in decentralized matching markets. For example, Shimer and Smith (2000), Atakan (2006), and Eeckhout and Kircher (2010) focus on sequential and directed search while Chade and Smith (2006) and Shorrer (2019) investigate simultaneous search. From a different perspective, Rastegari et al. (2013) and Drummond and Boutilier (2014) consider eliciting information about agents' preferences with a minimal number of interviews which are just enough to ensure that a stable match is found.

Most theoretical work on matching markets assumes that applicants know their priority at schools or universities, based, for example, on grades or entrance exams. This assumption is relaxed when studying the consequences of not publicizing the results of entrance exams before students have to submit their rank-order lists (Lien et al., 2016, Pan, 2019). In our study, we provide full information on the priorities of students at universities but vary information on the preferences and behavior of others.

2 Theoretical analysis

In this section, we present a model to analyze the strategies and welfare of students under the Direct Serial Dictatorship mechanism (DirSD), the Sequential Serial Dictatorship mechanism (SeqSD), and the Direct Serial Dictatorship mechanism with cutoff provision (Cutoff). First, we modify the standard school choice problem described by Abdulkadiroğlu and Sönmez (2003) to allow students to acquire information on their own preferences before and during the matching process.⁸ Next, we provide a detailed description of the procedures of DirSD, SeqSD, and Cutoff. We then discuss and compare these three procedures in terms of information acquisition, preference submission, and student welfare.

⁸The school and college admissions markets investigated in this paper are more similar to the school choice problem than to the college admissions problem (Gale and Shapley, 1962) in the matching literature because universities are not strategic and we focus on students when conducting the welfare analysis.

2.1 A university admissions problem

Students want to be assigned a seat at one of the universities. Each student has strict preferences over all universities and each university has strict priorities over all students. There is a maximum capacity at each university, but the total number of seats exceeds the total number of students. We consider an environment in which every university’s priority ordering over students is determined by exam rankings. Formally, the university admissions problem consists of:

1. A set of students $I = \{i_1, \dots, i_n\}$, $n \geq 2$.
2. A set of universities $C = \{c_1, \dots, c_m\}$, $m \geq 2$.
3. A capacity vector $q = (q_1, \dots, q_m)$.
4. A vector of students’ ranks $r = (r_1, \dots, r_n)$ in an exam, where r_i denotes the rank of student i among all students (with 1 being the highest rank). The ranking determines their priority ordering at every university.
5. A list of strict student preferences $\succ_I = (\succ_{i_1}, \dots, \succ_{i_n})$. The preference relation \succ_i of student i is a linear order over C , where $c \succ_i c'$ means that student i strictly prefers university c to university c' . Students prefer any university to remaining unmatched.⁹
6. For each student $i \in I$, a set of cardinal utilities $u_i = \{u_i^1, \dots, u_i^m\}$ associated with her ordinal preferences: student i receives u_i^j when assigned to a university ranked j th in her preference relation \succ_i . For any $1 \leq j < j' \leq m$, $u_i^j > u_i^{j'}$.¹⁰

It is common knowledge that every student knows her own exam rank. We assume that all market participants are risk neutral. Let Ω be the set of all linear orders over C . The preference relation \succ_i of student i is randomly and independently drawn from Ω following a prior probability distribution.

Specifically, we assume the following prior structure. Universities belong to different “tiers,” ranked from better to worse. All students have the same between-tier preference: they all prefer any university in a better tier to any university in a worse tier. However, students may have different within-tier preferences: each student’s preference over universities in the same tier follows a uniform distribution, that is, it is equally likely to be any linear order over these universities.

⁹The assumptions that the total number of seats exceeds the total number of students, that students prefer any university to remaining unmatched, and that universities prefer any student to leaving a seat empty are made to simplify the exposition. They are not necessary for our discussion and can be relaxed easily.

¹⁰Our main theorem regarding student welfare allows for alternative ways of modeling cardinal utilities (see Section 5.1), but the approach presented here corresponds to our main experimental setting. In the model, the cardinal utilities $u_i = \{u_i^1, \dots, u_i^m\}$ are fixed and known in advance. Given that these utilities are associated with student i ’s ordinal preference \succ_i , the student only needs to acquire information about how she ranks the universities. For example, student i knows that the university of her first preference gives her a utility of u_i^1 . She can discover which university is her first preference by acquiring information about her ranking of the universities \succ_i . A similar setup is used by Coles and Shorrer (2014).

Formally, let $\{T_t\}_{t=1,2,\dots,\tau}$ be a partition of the set of universities C . For any $c \in T_t$, $c' \in T_{t'}$, and $i \in I$, we have $c \succ_i c'$ if $t < t'$. That is, all students prefer any university in T_1 to any university in T_2 , prefer any university in T_2 to any university in T_3 , and so on. This between-tier preference is common knowledge to the entire market. A student can learn about the realization of her own within-tier preferences at a cost, but she cannot learn about other students' within-tier preferences. Moreover, the information acquired by each student is their private information. This prior structure allows for both a common and a private factor in students' preferences. For example, in many real-life university admission contexts, there is a consensus as to which universities belong to the top tier, to the second tier, and so on. But students' tastes over universities within a tier may vary depending on the location, family culture, personal tastes, etc.

We define the *budget set* B_i as the set of all universities available to student i . That is, student i can be assigned to university $c \in B_i$ if she so desires, and cannot be assigned to any university in the complement set $C \setminus B_i$.¹¹ In Section 2.4, we will discuss how the budget set of a student is determined under each procedure.

2.2 Matching procedures

Direct serial dictatorship mechanism (DirSD)

Every student simultaneously submits her rank-order list of universities. DirSD considers students in the order of their exam ranking.

Step 1: The student who is ranked first in the exam (with the highest score) is assigned a seat at the first choice on her submitted list.

In general, Step κ ($\kappa \geq 2$) can be described as follows.

Step κ : The student ranked κ th is assigned a seat at her best choice that still has vacant seats.

The procedure terminates when every student has been considered. Students can acquire information about their own preferences before submitting their rank-order lists.

Sequential serial dictatorship mechanism (SeqSD)

Under SeqSD, students sequentially select universities in the order of their exam ranking.

Step 1: The student who is ranked first in the exam (with the highest score) selects one university out of all universities. She is assigned a seat at this university.

In general, Step κ ($\kappa \geq 2$) can be described as follows.

¹¹We borrowed this terminology from Immorlica et al. (2020).

Step κ : The student ranked κ th selects one university out of all the universities that still have vacant seats. She is assigned a seat at this university.

The procedure terminates when every student has been considered. Students can acquire information about their own preferences before and after it is their turn to select a university.¹²

Direct serial dictatorship mechanism with a cutoff provision (Cutoff)

The matching procedure of Cutoff is the same as DirSD except that before the procedure starts, all students observe the scores that were necessary to be accepted by the universities in previous years.

We use the terms “preference submission” or simply “submission” to refer to the students’ interaction with a mechanism—that is, submitting a rank-order list under DirSD and Cutoff and picking a university under SeqSD.

2.3 Search technology

Information acquisition can take place before the submission of preferences under DirSD and SeqSD, and after observing the cutoff information but before the submission of preferences under Cutoff. Each student i , with information cost $k_i > 0$, can acquire information on her own within-tier preferences in each tier T_t using the following search technology:

Step 1: For a cost of k_i , the student can choose any two universities from tier T_t and learn which of the two universities is ranked higher in her own preference relation. Thus, for a cost of k_i she can learn the relative ordering of two universities in T_t .

Step 2: For an additional cost of k_i , the student can choose a third university from tier T_t and learn how it compares to the two universities previously investigated. Thus, for a total cost of $2k_i$ she can learn the relative ordering of three universities in T_t .

...

Step $(|T_t| - 1)$: For an additional cost of k_i , the student can learn how the $|T_t|$ -th university is compared to the $(|T_t| - 1)$ universities previously investigated. Thus, for a total cost of $(|T_t| - 1)k_i$ she can fully discover her preferences over universities in T_t .

Students can choose to stop at any step in the above process and can choose to search in zero, one, or multiple tiers. With a total cost of $\sum_{t=1}^T (|T_t| - 1)k_i$, the student can obtain full knowledge of her own preferences. We use this search technology to model a student who is investigating universities

¹²We adopted a slightly different description of SeqSD in our experiments to facilitate understanding. We framed it as students sequentially receiving offers from universities that still have vacant seats and being asked to accept one offer.

one after the other. After investigating each additional university in a tier, she knows how to compare it to all the universities she had previously investigated in the same tier. Before searching all universities in a tier, she can only learn the relative ranking of the universities investigated, but not their exact ranks among all the universities in that tier. In our setting, investigating only one university in a tier does not carry any information because the cardinal utility a student receives from a university is determined by the rank of that university in her preferences. Therefore, we start the process by having each student choose two universities in a tier to compare. One of these two universities could be a university with which the student is already familiar, for instance, a local university, and every investment in information informs her of a new university.

This search technology captures important features of search in many real-life scenarios. In centralized college admissions, students typically choose from hundreds of universities and programs. The characteristics of universities are multidimensional, which makes it natural to consider relative comparisons between universities. Studying the details of a certain program should allow a student to compare this university to those she has already investigated. But it might be unrealistic to assume she can be sure that this is her most-preferred university among all options, including those still unexplored. It may also be challenging for her to pin down her exact valuation of the university. Recent evidence suggests that people have a hard time reporting cardinal preferences and are much better at reporting ordinal preferences (Budish and Kessler, 2022), likely because they do not consider cardinal values when making decisions. Therefore, we believe that our search technology is realistic, and it is also relatively simple and easy to understand for experiment participants. However, our main theorem regarding student welfare does not rely on a particular search technology. We also check the robustness of our experimental results against an alternative search technology (Section 5.1).

2.4 Preference submission

We start our analysis by characterizing students' optimal preference submission strategy under DirSD, SeqSD, and Cutoff given partial or full knowledge of their own preferences. Then, we show that if every student adopts the optimal submission strategy, a student never has an incentive to search in more than one tier, and each tier can be considered as an independent market.

Proposition 1. (1) *Truth-telling is an optimal submission strategy under DirSD and Cutoff. That is, under DirSD and Cutoff, it is optimal for a student to rank universities according to the expected utilities (from high to low) in her rank-order list.*

(2) *Truth-telling is an optimal submission strategy under SeqSD. That is, when it is her turn to choose, it is optimal for a student to select the university with the highest expected utility out of all the universities that still have vacant seats.*

The proof of Proposition 1 is straightforward. In each step of DirSD, the student whose turn it is is assigned to the highest-ranked university in her submitted list from those that still have

vacant seats. Therefore, a student is never better off by ranking a university with a lower expected utility above one with a higher expected utility. The same intuition applies for Cutoff except that students may have different expectations due to the provision of information on historical cutoff scores. In SeqSD, a student whose turn it is chooses from all the universities that are available to her, so it is optimal to simply choose the university with the highest expected utility. Proposition 1 relies on the strategy-proofness of DirSD and SeqSD in environments with complete information.

Due to the assumption of aligned between-tier preferences, we can further characterize the truth-telling strategies under the three procedures.

Proposition 2. *A student who adopts the truth-telling strategy, regardless of her knowledge about her within-tier preferences, always*

(1) *ranks any university in a better tier above any university in a worse tier in her submitted rank-order list under DirSD and Cutoff, and*

(2) *chooses a university in the best tier among those available to her under SeqSD.*

We use q_{T_t} to denote the total capacity of all universities in T_t , that is, $q_{T_t} = \sum_j q_j$ for all j such that $c_j \in T_t$. From the strategies characterized above, we know that in equilibria with truth-telling strategies under all three procedures, students with the exam rank $r \leq q_{T_1}$ are admitted to universities in T_1 , students with $q_{T_1} \leq r \leq q_{T_1} + q_{T_2}$ are admitted to universities in T_2 , and so on. In general, students with $\sum_{i=1}^{t-1} q_{T_i} \leq r \leq \sum_{i=1}^t q_{T_i}$ are admitted to universities in T_t .

Proposition 2 implies that we can also categorize students by tiers: we say that those with $\sum_{i=1}^{t-1} q_{T_i} \leq r \leq \sum_{i=1}^t q_{T_i}$ are “tier- t students” since they would be admitted to a tier- t university as long as all students adopt the truth-telling strategy. For a tier- t student, universities in a better tier $T_{t'} (t' < t)$ are not available to her. On the other hand, although universities in a worse tier $T_{t'} (t' > t)$ are available to the student, she can always secure a seat at a tier- t university by adopting the truth-telling strategy. This means a tier- t student only needs to consider universities in T_t when choosing strategies; universities in other tiers are essentially irrelevant for her. Therefore, for any given $t = 1, 2, \dots, \tau$, we can consider all the tier- t universities and tier- t students as an independent market, and this market is identical to a one-tier market with uniform priors. We summarize the predictions regarding information acquisition

Proposition 3. *In a market with multiple tiers,*

(i) *a tier- t student only searches among tier- t universities if she chooses to acquire information; and*

(ii) *her search strategy among tier- t universities is the same as that in a one-tier market with only tier- t universities.*

Because each student only searches in one tier and her search strategy in that tier is the same as that in a one-tier market, we can focus our subsequent analysis of information acquisition on the simple case of a one-tier market.

2.5 Belief updating

In this section, we characterize how students update their beliefs when searching in a one-tier market. With uniform within-tier priors, when all universities belong to the same tier, each student i 's realized preference relation \succ_i is equally likely to be any linear order in Ω . In other words, all m universities are equally likely to be of any rank in \succ_i . Students can use the above-mentioned search technology to obtain full knowledge of their own preferences with $(m - 1)$ search steps.

Suppose student i stops searching at step α ($\alpha = 1, 2, \dots, m - 1$), and the set of universities she has chosen to search is given by C^S ($C^S \subseteq C$). This implies that for a cost of αk_i , the student learns the relative ranking of the $(\alpha + 1)$ universities in C^S , denoted as \succ_i^S . She can then eliminate the possibility of all linear orders in Ω that are not consistent with this ranking, and redistribute the probability uniformly among the remaining rankings.¹³

With uniform within-tier priors, student i has the same expected utility for each university before searching, which is given by

$$V_i(0) = \frac{1}{m} \sum_{j=1}^m u_i^j.$$

After conducting α steps of search, her updated expected utility for the university that is relatively ranked γ th in C^S according to \succ_i^S ($\gamma = 1, \dots, \alpha + 1$) is given by

$$V_i^\gamma(\alpha) = \sum_{j=1}^m f^\gamma(j, \alpha) u_i^j,$$

in which

$$\begin{aligned} f^\gamma(j, \alpha) &= \frac{\binom{j-1}{\gamma-1} \binom{m-\alpha-1}{j-\gamma} (j-\gamma)! \times \binom{m-j}{\alpha+1-\gamma} (m-j-\alpha-1+\gamma)!}{\binom{m}{\alpha+1} (m-\alpha-1)!} \\ &= \frac{\binom{j-1}{\gamma-1} \binom{m-j}{\alpha-\gamma+1}}{\binom{m}{\alpha+1}} \end{aligned}$$

calculates the probability that the γ th-ranked university in \succ_i^S is ranked j th in \succ_i . Intuitively, because this university is ranked γ th among the $(\alpha + 1)$ searched universities, if it is ranked j th in the student's complete preference ordering \succ_i , we can identify $(\gamma - 1)$ out of the j universities ranked above it and $(\alpha + 1 - \gamma)$ out of the $(m - j)$ universities ranked below it in \succ_i . The first term of the

¹³Formally, we say a linear order $\omega \in \Omega$ is *consistent* with \succ_i^S if $c \succ_i^S c'$ implies $c\omega c'$, $\forall c, c' \in C^S$.

numerator $\binom{j-1}{\gamma-1} \binom{m-\alpha-1}{j-\gamma} (j-\gamma)!$ is the number of possible orderings of the universities ranked above it and the second term $\binom{m-j}{\alpha+1-\gamma} (m-j-\alpha-1+\gamma)!$ is the number of possible orderings of the universities ranked below it. The denominator $\binom{m}{\alpha+1} (m-\alpha-1)!$ is the permutation of all universities after knowing the relative ranking of $(\alpha+1)$ of them.¹⁴ When $\alpha = m-1$, the student has full knowledge of her own preferences, and thus $V_i^\gamma(m-1) = u_i^\gamma$. When $\alpha < m-1$, the student's expected utility for those unsearched universities in $C \setminus C^S$ remains the same as the prior $V_i(0)$. We illustrate the belief updating process using the following example.

Example 1. Consider a one-tier market with three universities $C = \{c_1, c_2, c_3\}$. The first row of Table 1 lists all six possible preference orders over C (where we omit the subscript i to refer to student i). Without acquiring any information ($\alpha = 0$), a student holds the uniform prior belief. That is, she assigns the same probability $1/6$ to each of the six linear orders, which means that the expected utility is the same for all three universities:

$$E[c_1|\alpha = 0] = E[c_2|\alpha = 0] = E[c_3|\alpha = 0] = V(0) = \frac{1}{3}(u^1 + u^2 + u^3).$$

1st preference (u^1)	c_1	c_1	c_2	c_2	c_3	c_3
2nd preference (u^2)	c_2	c_3	c_1	c_3	c_1	c_2
3rd preference (u^3)	c_3	c_2	c_3	c_1	c_2	c_1
$\Pr[\succ \alpha = 0]$	$1/6$	$1/6$	$1/6$	$1/6$	$1/6$	$1/6$
$\Pr[\succ \alpha = 1, c_3 \succ c_2]$	0	$1/3$	0	0	$1/3$	$1/3$
$\Pr[\succ \alpha = 2, c_1 \succ c_3 \succ c_2]$	0	1	0	0	0	0

Table 1: Belief Updating Example

Suppose the student chooses to conduct the first step of searching ($\alpha = 1$). Suppose she picks universities c_2 and c_3 and learns that $c_3 \succ c_2$. Now she is able to eliminate all orders over C that are inconsistent with $c_3 \succ c_2$, and she redistributes the probability uniformly among the remaining

¹⁴Note that $f^\gamma(j, \alpha) = 0$ if $j < \gamma$ or $j > m - \alpha + \gamma - 1$ because there have to be at least $(\gamma - 1)$ universities ranked above the γ th-ranked university in \succ_i^S and at least $(\alpha + 1 - \gamma)$ universities ranked below it.

orders (see the third row of Table 1). Her updated expectation is

$$\begin{aligned} E[c_1|\alpha = 1, c_3 \succ c_2] &= V(0) = \frac{1}{3}(u^1 + u^2 + u^3), \\ E[c_2|\alpha = 1, c_3 \succ c_2] &= V^2(1) = \sum_{j=1}^3 f^2(j, 1) u^j = \frac{1}{3}u^2 + \frac{2}{3}u^3, \\ E[c_3|\alpha = 1, c_3 \succ c_2] &= V^1(1) = \sum_{j=1}^3 f^1(j, 2) u^j = \frac{2}{3}u^1 + \frac{1}{3}u^2, \end{aligned}$$

with $E[c_3|\alpha = 1, c_3 \succ c_2] > E[c_1|\alpha = 1, c_3 \succ c_2] > E[c_2|\alpha = 1, c_3 \succ c_2]$. Suppose the student continues searching ($\alpha = 2$) and learns that $c_1 \succ c_3 \succ c_2$. She can then further eliminate the possibility of any order inconsistent with $c_1 \succ c_3 \succ c_2$ (see the fourth row of Table 1) and obtain full knowledge of her preferences:

$$\begin{aligned} E[c_1|\alpha = 2, c_1 \succ c_3 \succ c_2] &= V^1(2) = u^1, \\ E[c_2|\alpha = 2, c_1 \succ c_3 \succ c_2] &= V^3(2) = u^3, \\ E[c_3|\alpha = 2, c_1 \succ c_3 \succ c_2] &= V^2(2) = u^2. \end{aligned}$$

The example demonstrates an important feature of the updating process. In particular, after discovering $c_3 \succ c_2$ at $\alpha = 1$, the student realizes that c_2 cannot be her favorite university while c_3 cannot be her least favorite. As she still holds the prior belief about the rank of c_1 , she now prefers c_3 to c_1 and c_1 to c_2 in expectation. The proposition below states that this is a general property: a student always prefers the higher-ranked searched universities to the unsearched universities and prefers the unsearched universities to the lower-ranked searched universities. This feature of the students' posterior beliefs allows us to characterize the students' strategies regarding preference submission and their expected utility functions under the three procedures.

Proposition 4. *In a one-tier market, for any $i \in I$, there exists a threshold rank $\hat{\gamma}_i(\alpha)$ at which $V_i^\gamma(\alpha) > V_i(0)$ for all $\gamma \leq \hat{\gamma}_i(\alpha)$ and $V_i^\gamma(\alpha) \leq V_i(0)$ otherwise. Thus, there is a step-specific threshold rank that splits all searched universities into two groups. The searched universities that are ranked higher than or equal to the threshold are preferred to all unsearched universities. All unsearched universities are weakly preferred to the searched universities that are ranked lower than the threshold.*

The complete proof of Proposition 1 can be found in Appendix A.1. In Example 1, the threshold rank after the first search step is $\hat{\gamma}_i(1) = 1$, meaning that the university ranked first among those searched has a higher expected utility than the unsearched university, while the university with a rank below one has a lower expected utility than the unsearched university. Thus, the higher-ranked searched university c_3 has a higher expected utility than the unsearched university c_1 , and

the lower-ranked searched university c_2 has a lower expected utility than the unsearched university c_1 .

2.6 Budget set and information acquisition

Recall that student i 's budget set B_i is defined as the set of all universities available to her. Under DirSD and Cutoff, if $c \in B_i$, i can be assigned to c unless she is assigned to another university that is ranked higher than c in her submitted rank-order list; if $c \notin B_i$, i cannot be assigned to c no matter how she ranks c in her submitted list. Under SeqSD, when i selects universities, c would be available to her if $c \in B_i$, and unavailable if $c \notin B_i$.

Student i 's budget set is determined by her exam rank and the submission strategies of those who are ranked above her. For instance, consider a market with three universities $C = \{c_1, c_2, c_3\}$, each of which has two seats. Under DirSD, the budget set of the student ranked third in the exam depends on the submitted rank-order lists of the two students ranked above her. If, for example, they both place university c_3 at the top of their lists, the budget set of the student ranked third contains only c_1 and c_2 . Since all rank-order lists are submitted simultaneously under DirSD, a student decides what to search for based on the ex-ante probability distribution of her budget set $\left\{P_i(\tilde{B})\right\}_{\tilde{B} \subseteq C}$, that is, $P_i(\tilde{B}) = \Pr[B_i = \tilde{B}]$, $\tilde{B} \subseteq C$. This requires students to form beliefs about the submission strategies of higher-ranked students. In contrast, under SeqSD a student selects the preferred university after the higher-ranked students have made their choices. She therefore observes the realization of her budget set before she does a search and makes her preference submission decisions.

We consider Cutoff as a case that comes in between DirSD and SeqSD in terms of students' knowledge about their budget set before the search. Under Cutoff, students are provided with the scores needed for university admission in previous years. Such information can be helpful to determine the budget set. However, this information is often not precise enough to pin down the exact budget set for every student in the market because the distribution of student preferences, the distribution of exam scores, and the capacities of universities may change from year to year. Therefore, after observing historical cutoffs student i decides on her search based on the updated beliefs, which may still be non-deterministic.

To simplify the subsequent analysis, we make the following assumptions about the market structure and the students' strategies.

Assumption 1. (1) *All universities in the same tier have the same capacity; universities in different tiers can have different capacities.*

(2) *At each step of the search process, a tier- t student is equally likely to investigate any of the unsearched tier- t universities under DirSD, and is equally likely to investigate any of the unsearched tier- t universities that are available to her under SeqSD.*

(3) *If a tier- t student did not search all tier- t universities under DirSD, she is equally likely to choose any relative order over the unsearched tier- t universities in her submitted rank-order list. If a tier- t student did not search any tier- t universities that are available to her under SeqSD, she is equally likely to select any one of these tier- t universities.*

Assumption 1 can be considered as an anonymity assumption in that universities in the same tier are not labeled and thus the search is random among unsearched universities. To illustrate, consider a one-tier market. Assumption 1, together with the uniform within-tier prior structure and Proposition 1, implies that a student i , who by assumption cannot observe what another student $i' \neq i$ has learned about her preferences $\succ_{i'}$, believes that i' is equally likely to submit any ranking in Ω under DirSD and is equally likely to select any university in C under SeqSD. In other words, under DirSD and SeqSD, for student i the submission strategy of i' always follows a uniform distribution, regardless of the search strategy of student i' . Thus, a student does not need to consider how much information other students acquire in equilibrium when forming beliefs about the submission strategies of others under DirSD and SeqSD.

Therefore, the anonymity assumption and uniform within-tier priors are mainly used to simplify the derivation of the ex-ante probability distribution of a student's budget set, which is needed to pin down the optimal search strategy under DirSD. Without these assumptions, students may need to form beliefs about the search strategies of others under DirSD. But even with these simplifying assumptions, decision-making under DirSD is challenging especially for the lower-ranked students, because they have to consider the submission strategies of other students.

In appendices A.2 and A.3, we present the strategies of information acquisition given the search technology. Under SeqSD, the marginal benefit of an additional search step among the available universities decreases, and we characterize the optimal stopping point of the search process. In contrast, under DirSD the marginal benefit of an additional search step may be non-monotonic. The optimal information acquisition strategy is not necessarily unique, but it is unique for the parameters that we chose in the experiment.

These arguments regarding optimal search under DirSD shed light on the complexity of deriving the optimal search strategy under Cutoff. After observing the cutoffs, students update their beliefs about their budget set. In a one-tier market, suppose student i observes that universities c and c' have low cutoff scores and thus believes that they are available to her. She may then be more likely to search c and c' compared to other unsearched universities, which violates part (2) of Assumption 1. As a result, she may also be more likely to rank c or c' as her top choice rather than other universities under Cutoff. Because the cutoff scores are public information, another student $i' \neq i$ can infer that the submission strategy of student i does not necessarily follow a uniform distribution. Therefore, i' may need to consider how much and what information i chooses to acquire in equilibrium when forming beliefs about her own budget set. This significantly complicates the theoretical analysis. In this paper, we focus on the empirical exploration of subjects' search

behavior when observing historical cutoffs. In the experimental section, we test whether subjects adopt the simple strategy of limiting their search to universities that are within their budget set based on the historical cutoffs.

2.7 Student welfare

Our main theorem compares student welfare between SeqSD and DirSD. We focus on the welfare comparison at the ex-ante stage—that is, before students acquire any information about their preferences—and assume that all students acquire information optimally and adopt the truth-telling submission strategy under both procedures.

Theorem 1. *Every student is weakly better off under SeqSD than under DirSD ex ante if all students acquire information optimally and adopt the truth-telling submission strategy.*

While the complete proof of Theorem 1 can be found in Appendix A.4, the key intuition comes from the difference in the amount of information students have when they make their decisions. Under DirSD, all students simultaneously submit their rank-order lists. Thus, they do not have any opportunity to identify their budget sets by learning about the choices of others. In contrast, the dynamic nature of SeqSD can provide additional information about which universities other students have chosen and which ones are still available in the market, thus helping students to identify their budget sets. By focusing on a search within the budget set, a student can reduce wasteful information acquisition and can thus be weakly better off under SeqSD. A student has the same welfare under SeqSD and DirSD if she can perfectly identify her budget set under DirSD (for example, if she is the highest-ranked student) or if the information cost is so large that she does not search at all under either procedure. This explains why the relationship is weak.

Importantly, the proof of Theorem 1 does not rely on a particular search technology. Also, it holds true for any ex-ante distribution over a student’s budget set, which means that it applies to every student regardless of the exam rank.

The welfare comparison between Cutoff and the other two procedures depends on the specific setting and implementation. On the one hand, historical cutoffs can be helpful to determine a student’s budget set, reduce wasteful information acquisition, and thus improve student welfare. On the other hand, public information about cutoffs can change the search behaviors of higher-ranked students, which may have a spillover effect on a student’s budget set and affect her welfare. For example, a student may have a smaller budget set under Cutoff if higher-ranked students focus their search on similar universities after observing the cutoff information. It is challenging to derive theoretical predictions on the overall effect using the discrete model in this paper. Therefore, we focus on empirically investigating the effects of historical cutoffs on student welfare using our experimental data.¹⁵

¹⁵Using a continuous model, Artemov (2021) analyzes the spillover effects of information acquisition and Immorlica et al. (2020) study how market-clearing cutoffs can facilitate efficient information acquisition.

3 Experimental Design

We conducted an experiment to test the predictions of our information acquisition model, and to compare the welfare of students under the three centralized matching procedures. The three procedures (treatments) are studied in four different environments that are characterized by the one and two-tier preferences of students and two different costs of information acquisition.

3.1 Setup

In the experiments, 12 students competed for 12 seats at six universities. Each university had two seats. All universities ranked students based on the exam scores. The score of each student was randomly and independently drawn from a uniform distribution between 1 and 100. Students were played by experimental subjects while the universities were not strategic and their actions were simulated by the computer. Students knew their score and the rank of their score among the other students in their group.¹⁶

Participants received 40 AUD for being assigned to their most-preferred university, 34 AUD for the second most-preferred university, 28 AUD for the third most-preferred university, etc., and 10 AUD for the least-preferred university. At the beginning of each round, participants did not know their preferences over universities but were told that each ranking was equally likely. They had the opportunity to acquire costly information about their own ordinal preferences. The exact timing, technology, and costs of this search process varied between environments and treatments.

Each session consisted of 24 participants who were split into two groups of 12 for the entire session. We used fixed matching groups to increase the number of independent observations. Each round represented a new university admissions process for the students. In total, there were eight rounds in the experiment. At the end of the experiment, one round was drawn randomly to determine the subjects' payoff.

3.2 Treatments

We conducted three treatments between subjects by varying the centralized allocation procedure and the information provided:

In **Treatment DirSD**, the direct serial dictatorship mechanism is adopted. Participants could learn their preferences at a cost before the start of the procedure, i.e., before submitting the rank-order lists to the system.

In **Treatment SeqSD**, the sequential serial dictatorship mechanism is adopted. Students could search before or after observing the set of available universities.¹⁷

¹⁶Although students under DirSD and SeqSD only need to know their ranks to make decisions, we also provided them with scores because this enabled us to conduct the treatment with cutoffs scores.

¹⁷Participants could also start searching before observing the set of available universities, and continue with the search after observing it.

In **Treatment Cutoff**, we provide historical cutoff scores under the direct serial dictatorship mechanism. All participants observed the cutoffs of all universities of the previous cohorts before the procedure started. To generate the cutoffs, we used the results of the DirSD sessions. More precisely, we provided the average cutoff scores from all previous DirSD markets in the same environment. Students with the same rank in the previous sessions had the same preferences over universities as in the current session, which was explained to the participants in the current session. Note that this design choice ensures the informativeness of the cutoff scores and their relevance for the optimal search strategies (see section 3.5 for details). After learning the cutoffs, subjects could acquire information about their own preferences and then submit their rank-order list, just as in DirSD.

3.3 Environments

There were four environments in the experiments. The environments varied in two dimensions, namely the prior about the quality of the universities and the cost of information acquisition.

Dimension 1: One-tier versus two-tier preferences. The first dimension of our environments was the number of tiers in the preferences of students. That is, we varied the prior belief of students about the position of universities in their preferences, and thus the degree of correlation of preferences between students. We considered two preference structures:

Two tiers. The six universities A1, A2, A3, B1, B2, and B3 belonged to two different tiers: universities A1, A2, and A3 belonged to tier A, and universities B1, B2, and B3 belonged to tier B. Every student preferred a university in tier A to a university in tier B, which was common knowledge. Students could have different preferences within each tier. For each tier, the within-tier ordinal preferences of each student were independently and randomly drawn from the set of all possible orderings of the three universities in that tier. Each ordering was equally likely.

For each university tier, the search process was as follows:

1. For a cost of $\$X$, a student could pick any two universities belonging to the same tier and learn which of these two universities was ranked higher in her preference ordering. Thus, for a cost of $\$X$ a student could learn the relative ordering of two universities within a tier.
2. For an additional cost of $\$X$, a student could learn how the third university from the same tier compared to the two universities that she had chosen previously. Thus, for a cost of $\$2X$ a student could learn her full ranking of universities within a tier.

The same process applied to both tiers of universities. Thus, for a total cost of $\$4X$ a student was able to obtain full knowledge of her preferences.

One tier. The six universities, namely A, B, C, D, E, and F, belonged to one tier. The ordinal preferences of each student were independently and randomly drawn from the set of all possible orderings of the six universities. Each ordering was equally likely to be drawn.

The search process was as follows:

1. For a cost of $\$X$, a student could pick any two universities and learn which of these two universities was ranked higher in her preference ordering. Thus, for a cost of $\$X$ a student could learn the relative ordering of two universities.
2. Next, for an additional cost of $\$X$, a student could learn how a third university compared to the two that she had chosen previously. Thus, for a cost of $\$2X$ a student could learn the relative ordering of three universities.
3. ...
4. Finally, for an additional cost of $\$X$, a student could learn the preference ordering of all six universities. Thus, for a total cost of $\$5X$ a student was able to obtain full knowledge of her preferences.

We used the two-tier environment for two reasons. First, the strategies are more straightforward than in the case of one tier, as each tier essentially represents a smaller separate market, and our theoretical results apply to each tier separately (Proposition 3). In equilibrium, the six higher-ranked students are assigned to tier-A universities while the six lower-ranked students are assigned to tier-B universities. Thus, each student in equilibrium only has incentives to search within one tier. Moreover, markets with multiple tiers are more realistic in university admissions than one-tier markets, since in practice students often agree on how to group universities in terms of quality, but may have different preferences within each group. We used one-tier environments because they generate higher variation in the optimal search strategies, both between students of different ranks and between treatments.

Dimension 2: Cost of information acquisition. The second dimension of our environments was the cost of information acquisition (the value of X). We considered two cost levels:

Low cost ($X=\$0.5$) and High cost ($X=\2.3).

When varying the cost, the predictions regarding the centralized admission procedures varied greatly. The exact parameters were chosen such that the optimal search strategies ranged from full search to no search, depending on the rank of a student. Thus, we allowed for deviations in both directions, namely under-search and over-search. Furthermore, the low-cost environment that induces full search by all students (except those with the lowest score in each tier) guarantees the informativeness of the cutoff scores under optimal search.

Table 2 presents the summary of the design by sessions, including the treatments and the order of the environments. The first four rounds were always the rounds with two tiers, while rounds

five to eight were the rounds with one tier. We chose to fix this order, since the rounds with two tiers are simpler in terms of finding the optimal search strategies for subjects, and since we do not intend to directly compare the two environments. Within the two-tier and one-tier environments, there were two rounds with high costs and two rounds with low costs. To control for order effects regarding the costs, we used two different orders: either two consecutive high-cost rounds preceded the two low-cost rounds or vice versa.¹⁸

Sessions	Round	1	2	3	4	5	6	7	8
1, 2, 9	Treatment	DirSD							
	Tier	2				1			
	Cost	High	High	Low	Low	High	High	Low	Low
4, 7, 10	Treatment	DirSD							
	Tier	2				1			
	Cost	Low	Low	High	High	Low	Low	High	High
3, 5, 12	Treatment	SeqSD							
	Tier	2				1			
	Cost	High	High	Low	Low	High	High	Low	Low
6, 8, 11	Treatment	SeqSD							
	Tier	2				1			
	Cost	Low	Low	High	High	Low	Low	High	High
16, 17, 18	Treatment	Cutoff							
	Tier	2				1			
	Cost	High	High	Low	Low	High	High	Low	Low
13, 14, 15	Treatment	Cutoff							
	Tier	2				1			
	Cost	Low	Low	High	High	Low	Low	High	High

Table 2: Summary of sessions by treatments and environments

3.4 Experimental procedures

Across all treatments, we assigned the same randomly generated preferences to students with the same rank in the corresponding rounds and environments. For instance, the preferences of a student with rank 1 were the same in round 1 of sessions 1, 2, 11, 3, 5, 12, 16, 17, 18, and in round 3 of all other sessions. This is because we implemented two different orderings of the costs (see Table 2). For the cutoffs from the previous sessions of the DirSD treatment to be informative for the subjects in the Cutoff treatment, a correlation between the preferences of the two cohorts was necessary. This correlation was created by students with the same rank who had identical preferences in all treatments.

¹⁸Note that we do not randomize the order of tiers. Instead, we always start with the simpler environment with two tiers. This is because we are not interested in the comparative statics between environments and causal evidence on the effect of one or two tiers on market behavior.

We used the same randomly generated scores in DirSD and SeqSD but regenerated the scores for the Cutoff treatment. This was explained to the participants in the Cutoff treatment. This design ensures that the cutoffs were informative about the competitiveness of universities, but did not provide perfect information due to the fluctuation in the distribution of scores. It also prevented the cutoffs from directly informing the students about the preferences of the previous cohort.¹⁹

The experiment was conducted in the Experimental Economics Laboratory of the University of Melbourne (E²MU) and was programmed using z-Tree. Upon entering the lab, subjects were provided with experimental instructions for the treatment in the two-tier environment. Before the start of the one-tier environment, an additional set of instructions was distributed. In total, we conducted 18 sessions with 24 subjects each. Thus, we had 432 participants with average earnings of 28 AUD. The sessions lasted around 80 minutes.

3.5 Hypotheses

The main goal of the experiment is to compare the three treatments across the different environments. Our design is not aimed at comparisons between environments under the same matching procedure. Instead, we compare the different matching procedures in a variety of market environments.

We start with the search behavior of students.

Hypothesis 1 (Search): *Participants in DirSD and SeqSD acquire information about their own preferences following the predictions of the model. In Cutoff, students are less likely to search the universities with cutoffs higher than their score compared to the universities with cutoffs lower than their score.*

The optimal search strategies under DirSD and SeqSD are provided in figures 4 and 5 in Appendix B.1. The predictions regarding search depend not only on the treatment, but also on each student’s rank. Because the rank essentially determines the student’s budget set, the benefit of searching varies greatly between ranks. For instance, it is an optimal strategy to never search for a rank-12 student, as she always receives the last available seat. Similarly, a rank-6 student never searches in the two-tier environment, since she prefers to be matched to the last available seat in tier A. Note that the search incentives depend on the size and probability distribution of the budget set. This explains why the incentive to search does not necessarily decrease for lower-ranked students and why students do not always search weakly less under SeqSD than under

¹⁹Without regenerating scores for the Cutoff treatment, a subject in a Cutoff session could have directly observed the allocation of her “copy” and thus infer the realization of her own preferences. For example, if a subject with score 85 learns that the cutoff score of University A was 85, she would know that her copy was allocated to University A. This would have affected her decisions regarding information acquisition and preference submission. This is not the type of information cutoffs carry in real markets, but is a result of a small and discrete market implemented in the experiment.

DirSD, given their rank.

In contrast to DirSD and SeqSD, it is challenging to derive point predictions for the optimal search strategies in the Cutoff treatment. The main reason is that given the public information on cutoffs, students update their beliefs about their budget sets, and their prior beliefs before information acquisition can be complex.²⁰ Recall from Section 2.4 that uniform within-tier priors are important for keeping the derivations of the optimal search strategies theoretically tractable. Without them, every student needs to consider what information other students acquire in equilibrium when forming beliefs about her budget set. We therefore choose to focus on the empirical exploration of the effects of cutoffs in this paper. Specifically, we investigate a simple strategy in the Cutoff treatment: whether subjects are less likely to search the universities with cutoffs higher than their scores. This helps us understand whether subjects use the cutoff information to narrow down their budget sets and base their search strategies on it.

Hypothesis 2 (Submission): *In DirSD and Cutoff, students submit their preferences in order of the decreasing expected utility of universities, given the updated beliefs after the search. In SeqSD, a student who has searched selects the highest-ranked university among those searched.*

The prediction follows directly from Proposition 1.

Next, we turn to the our main interest—the welfare of the students.

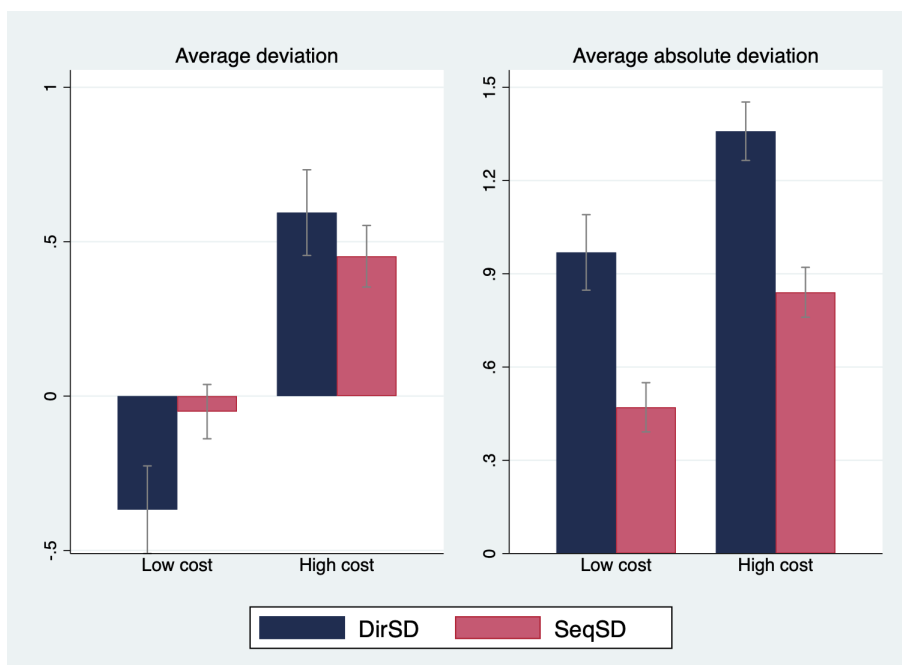
Hypothesis 3 (Welfare): *In terms of student welfare, the following relationships hold for all students in all environments: $DirSD \leq Cutoff \leq SeqSD$.*

The comparison between DirSD and SeqSD follows from Theorem 1. As explained in Section 2.7, it is challenging to theoretically compare Cutoff with the other two treatments due to the potential externality in information acquisition. We therefore focus on the empirical exploration of the welfare effects of cutoff provision. As we consider Cutoff to be a case that comes in between DirSD and SeqSD in terms of students’ knowledge about their budget set prior to searching, we hypothesize that Cutoff is also between DirSD and SeqSD in terms of student welfare.

4 Experimental Results

We start with the analysis of individual strategies and then present the market outcomes. We can pool the sessions that differ with respect to the order of information costs, since this order does not significantly affect the main variables of interest (see Table B.3 in Appendix B.3). All results reported are significant at the 5% level if not stated otherwise. For all tests, we use the p-values

²⁰For example, consider a one-tier market. Suppose that based on the cutoff information, other students can infer that the fourth ranked student knows that her budget set is more likely to include universities A and B than the other universities. Then they would know this student is more likely to search universities A and B and, as a result, is more likely to rank A or B as her top choice than other universities. Thus, the submission strategy of this student no longer follows a uniform distribution, which affects the beliefs of other students about their own budget sets.



Notes: The left panel presents average deviations with under-search as a negative value and over-search as a positive value. The right panel presents average absolute deviations, that is, the average absolute differences between the optimal and the observed number of searches. Vertical gray bars represent the 95% confidence intervals.

Figure 1: Deviations from the optimal number of searches by treatments

of the coefficient of the treatment dummy in regressions on the variable of interest. Standard errors are clustered at the level of matching groups, and the sample is restricted to the treatments that are of interest for the test. We use the sign “>” between treatments to express “significantly higher,” and the sign “=” to express “no significant difference.”²¹

4.1 Individual behavior

4.1.1 Search strategies

In this section we study the subjects’ search strategies across the different treatments and environments.

First, we present the main results on the optimal search in DirSD and SeqSD. (The detailed analysis of search strategies by ranks and environments is presented in Appendix B.1.) We find that in all treatments, the average number of searches is significantly higher in environments with low costs than in environments with high costs. In the low-cost environments in DirSD, all subjects

²¹Due to a mistake in our code of the z-tree program, in one market of treatment SeqSD (the second round with two tiers and high cost), rank 12 students were asked to choose a university before rank 11 students. In order to keep the comparison of welfare balanced between procedures, we deleted the observations for the students ranked 11th and 12th in this round of all treatments. The inclusion of these data or the exclusion of these subjects from analyses does not change the qualitative results.

except those ranked last in each tier are predicted to obtain full knowledge of their preferences, but they under-search on average. Unlike in DirSD, in the low-cost environments in SeqSD, the deviations from the predicted number of searches are small. In the high-cost environments, none of the subjects are predicted to obtain full knowledge about their preferences, and we find that, on average, they search too much in both DirSD and SeqSD.

The left panel of Figure 1 presents the average deviation from the predicted number of searches for low and high-cost environments in DirSD and SeqSD. In low-cost environments, students under-search on average, with significantly more under-search in DirSD than in SeqSD. In high-cost environments, we observe over-search on average. Note, however, that under-search is the only possible deviation for all but the lowest-ranked students in each tier in the low-cost environments, and thus it has to be interpreted with caution.²² Our results concerning over-search in high-cost environments are in line with previous experimental findings on information acquisition (see Chen and He, 2021, for school choice, Bhattacharya et al., 2017, for voting, and Gretschno and Rajko, 2015, for auctions).

Averaging over positive and negative deviations can substantially mask actual deviations from the theory. Therefore, we also consider absolute deviations. The right panel of Figure 1 presents the average absolute deviation from the predicted number of searches for low- and high-cost environments in DirSD and SeqSD. The average absolute deviation is significantly lower in SeqSD than in DirSD independent of the costs. The difference is significant for the pooled sample ($p < 0.01$) and for each environment separately ($p < 0.01$). Thus, SeqSD not only leads to lower search costs in theory, but it also induces behavior in the lab which is more in line with an optimal search than under DirSD. One possible explanation for this result is that the optimal search strategy is more straightforward under SeqSD than under DirSD.

Next, we turn to the search behavior in the Cutoff treatment. On average, when the cost is low, the search under Cutoff is not significantly different from DirSD in the two-tier environment ($p = 0.32$), but is significantly lower than under DirSD in the one-tier environment ($p < 0.01$). When the cost is high, the search under Cutoff is significantly lower than under DirSD ($p < 0.01$ for both one- and two-tier environments), and under SeqSD for the one-tier environment ($p < 0.01$), but not for the two-tier environment ($p = 0.16$). Thus, the participants rely on cutoffs more in high-cost environments than in low-cost environments. As we do not form point predictions for the optimal search under Cutoff, we use regressions to analyze the empirical patterns of the subjects' reaction to cutoffs. Specifically, we investigate the simple strategy identified in Prediction 2, that subjects are less likely to search the universities with cutoffs higher than their own score compared to universities with cutoffs below their own score. Table 3 presents the marginal effects of the probit model for information acquisition about a university, depending on the cutoff of this

²²The finding of over-search in low-cost and under-search in high-cost environments is in line with recent findings of Descamps et al. (2021). However, in their experiment under-search is not due to the equilibrium being a corner solution.

university.

	Cutoff all	Two tiers & low cost	Two tiers & high cost	One tier & low cost	One tier & high cost	Cutoff all	Two tiers	One tier
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Cost of search	-0.11 (0.01)					-0.12 (0.01)	-0.12 (0.01)	-0.13 (0.01)
Dummy for two tiers	-0.16 (0.01)					-0.15 (0.01)		
Cutoff higher, dummy	-0.08 (0.02)	-0.02 (0.03)	-0.11 (0.03)	-0.05 (0.03)	-0.13 (0.03)			
Cutoff higher, difference						-0.009 (0.001)	-0.013 (0.001)	-0.004 (0.001)
Cutoff lower, difference						-0.004 (0.0003)	-0.009 (0.0004)	-0.0005 (0.0007)
Observations	6768	1728	1584	1728	1728	6768	3312	3456

Note: Marginal effects of probit regressions regarding information acquisition about a university in Cutoff. “Cutoff higher, dummy” is a dummy that is equal to one if the cutoff score of the university minus the score of the student is greater than zero. “Cutoff higher, difference” is equal to the cutoff score of the university minus the score of the student if the difference is positive and zero otherwise. “Cutoff lower, difference” is equal to the student’s score minus the cutoff score of the university if the difference is positive and zero otherwise. Standard errors are clustered at the level of matching groups.

Table 3: Probability of information acquisition about a university depending on the cutoff

Model (1) of Table 3 presents the results for all environments of the Cutoff treatment. The coefficient of “Higher cutoff, dummy” is negative and statistically significant. Thus, on average, participants are less likely to search among the universities with cutoff scores higher than their score. This suggests that participants believe that these universities are less likely to be within their budget set. We consider each environment separately and find that the effect is strongest in the environments with high information costs, see models (3) and (5). In contrast, in low-cost environments, the effect is either not significant or only marginally significant, see models (2) and (4). Model (6) considers the absolute difference between a cutoff and a student’s score for all environments. The higher the cutoff is relative to the student’s score, the less likely it is that the student will acquire information about this university. However, students are also less likely to acquire information about universities the further the cutoff is below their own score. The magnitude of the effect is smaller, but still significant. Models (7) and (8) study two- and one-tier environments separately. The effect of lower cutoffs is significant in the two-tier environments. This can be explained by students in ranks 1 to 6 not searching tier-two universities. In one-tier environments, the lower cutoff scores do not decrease the probability of search significantly, which can be rationalized by the independence of the preferences.

We summarize our findings regarding individual search strategies in the following result.

Result 1 (Search strategies):

- *In low-cost environments, the average number of searches in SeqSD is not statistically different from the predicted optimal strategy, while there is significant under-search in DirSD. In high-cost environments, there is over-search in DirSD and SeqSD, with larger deviations from the optimal strategy in DirSD.*
- *The average absolute deviation from the optimal search strategy is lower in SeqSD than in DirSD.*
- *Students are less likely to search universities with cutoffs higher than their score compared to universities with cutoffs below their score, especially in the high-cost environments.*

4.1.2 Submission strategies

In this section, we analyze the subjects' strategies for ranking and choosing universities. When a participant has learned her preferences completely, the optimal submission strategy is to list all universities in order of her true preferences in DirSD and Cutoff, and to select the most-preferred university from the available ones according to her true preferences in SeqSD. Note that a student's submitted list in DirSD and Cutoff is relevant only up to her guaranteed university. For example, a rank 4 participant is guaranteed a seat at the university of her second preference, since each university has two seats. Similarly, a rank 7 participant is guaranteed a seat at the university of her fourth preference. If a student does not have full knowledge of her preferences, propositions 1 and 2 present the optimal submission strategies. Depending on the treatment, the optimal submission strategy entails the following behavior:

- In DirSD, under the optimal submission strategy, universities are ranked in decreasing order of expected values. In keeping with Proposition 4, this implies listing the higher-ranked searched universities above all unsearched ones, followed by the lower-ranked searched universities.²³ Students are indifferent between unsearched universities such that the order among them does not affect optimality. If a student does not search any university, any list is optimal (respecting tiers). When identifying optimal submission strategies, we only consider a student's submitted list up to her guaranteed university. For instance, for a student with rank 1, only the top choice is considered when evaluating the optimality of her strategy.
- In SeqSD, the optimal submission strategy implies choosing the highest-ranked university among those searched from the set of available universities. If a student does not search any available university, any choice is optimal (respecting tiers).

²³In our experimental setting, if the number of searched universities is even, the higher-ranked half of the searched universities should be listed above the unsearched universities, followed by the lower-ranked half of the searched universities. If the number of searched universities is odd, the optimal submission strategy is the same, but the middle-ranked searched university is treated like an unsearched university.

- In Cutoff, the optimal submission strategy is the same as in DirSD. Note, however, that the cutoffs might lead to multiple optimal strategies. For instance, if a student believes, based on the cutoff, that a university is not in her budget set, she is indifferent with respect to how to rank this university. Thus, unsearched universities that are not in the budget set can be placed anywhere in the submitted rank-order list. As a benchmark, we compare the strategies in Cutoff to the optimal strategies in DirSD, thereby potentially underestimating the proportion of optimal strategies in Cutoff, since we neglect the fact that universities not in the budget set can be placed anywhere in the rank-order list.

Table 4 presents the proportion of optimal submission strategies conditional on the subjects' search behavior. The highest rate of optimal submission strategies is observed in SeqSD with differences being significant in all environments relative to DirSD and Cutoff. Note that finding the optimal submission strategy under SeqSD does not require the ability to compare the expected utilities of all searched and unsearched universities. In contrast, the overall proportions of optimal submission strategies are 80.4% and 75.7% in DirSD and Cutoff, respectively, with the difference between them being significant for the pooled sample but not in any of the environments separately. In DirSD and Cutoff, the deviation from the optimal submission strategies might be driven either by the participants' attempt to manipulate the rank-order lists submitted to the mechanism or by the complexity of comparing the expected utilities of the searched to the unsearched universities. We find that when subjects have full knowledge of their preferences, the rate of manipulations of the submitted lists is only 8.2% in DirSD and 14.7% in Cutoff. With partial knowledge, the rates of reports where the relative positions of the searched universities are misreported are 5.7% in DirSD and 7.6% in Cutoff, which are just 29% and 31% of all suboptimal strategies in DirSD and Cutoff, respectively.²⁴ The remaining deviations from optimal strategies (71% in DirSD and 69% in Cutoff) are driven by the incorrect rankings of unsearched universities relative to searched ones.

Thus, the complexity of comparing the expected utilities of searched to unsearched universities contributes to the higher-than-predicted difference in welfare between SeqSD and DirSD. In line with this observation, the rate of optimal submission strategies is significantly higher in the environments with low costs than with high costs under DirSD and Cutoff, because more subjects acquire full information when the search cost is low. However, there is no significant difference in the optimality of submission strategies between high and low-cost environments under SeqSD.

Result 2 (Submission strategies): *In all environments, the proportion of optimal submission strategies is significantly higher under SeqSD than under DirSD and Cutoff.*

²⁴The rate of truthful reporting given that students have full knowledge of their preferences is higher than in Li (2017) and Bó and Hakimov (2023), where subjects misreport their rank-order lists in 33% and 28% of cases, respectively, under DirSD. This difference may be driven by the fact that those who acquire full information about their preferences in our experiment are a selected sample. Alternatively, our participants have to decide on their search strategies and may therefore give less thought to submission strategies, which can improve truthfulness. As for SeqSD, the rate of truthful behavior is 98% in Bó and Hakimov (2023), which is similar to ours, and higher than the 83% in the first round of Li (2017), who does not report the average over all rounds.

		Treatment			p -value for test of equality		
		DirSD (1)	SeqSD (2)	Cutoff (3)	DirSD=SeqSD (4)	DirSD=Cutoff (5)	SeqSD=Cutoff (6)
Optimal strategies given search	Two tiers & low cost	78.8%	97.9%	<i>74.0%</i>	0.00	0.22	0.00
	Two tiers & high cost	79.2%	98.5%	<i>72.3%</i>	0.00	0.08	0.00
	One tier & low cost	85.4%	99.3%	<i>78.1%</i>	0.00	0.07	0.00
	One tier & high cost	78.1%	98.6%	<i>78.1%</i>	0.00	1.00	0.00
	All	80.4%	98.6%	<i>75.7%</i>	0.00	0.03	0.00

Notes: For the tests in columns 4-6, we use the p -values for the coefficient of the treatment dummy in the probit regression of the optimal submission strategy. Standard errors are clustered at the level of matching groups, and the sample is restricted to the treatments that are of interest for the test. The proportions of optimal submission strategies in Cutoff are italicized as the prediction ignores that multiple optimal strategies can exist.

Table 4: Proportions of optimal submission strategies by treatments and environments

In the next section, we investigate whether the observed differences in individual behavior across procedures translate into differences in market outcomes.

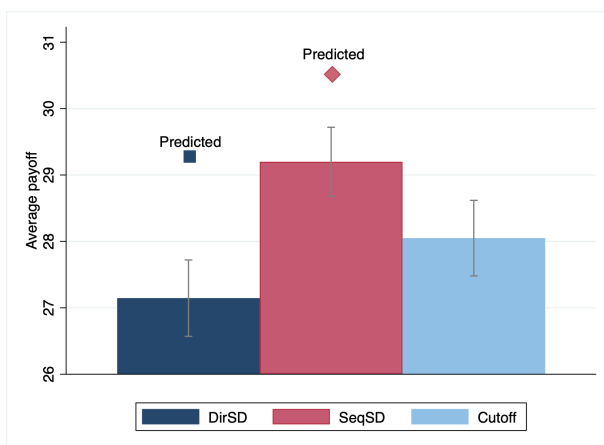
4.2 Market outcomes

4.1.1. Welfare

Figure 2 shows the average payoffs of participants by treatments, aggregated over all environments. The average payoff is highest in SeqSD, with the difference to DirSD_< and to Cutoff being significant. The markers in Figure 2 indicate the theoretical predictions for the average payoffs in DirSD and SeqSD, showing that the higher welfare of participants under SeqSD compared to DirSD is in line with the theoretical predictions. However, in both SeqSD and DirSD welfare is lower than predicted: in DirSD, the average payoffs are 2.2 AUD lower than predicted, while in SeqSD the difference is 1.2 AUD. This can be due to either the suboptimal search strategies or suboptimal submission strategies discussed in the previous section. In the Cutoff treatment, as hypothesized, we observe that the average payoffs of students are higher than in DirSD ($p = 0.05$) and lower than in SeqSD ($p < 0.01$).²⁵

To understand in which environments the Cutoff and SeqSD procedures have the greatest advantage over DirSD, Table 5 presents the average payoffs of participants by treatments for each tier and cost combination. First, SeqSD has significantly higher average payoffs than DirSD in all environments, see column (4) for the p -values. This confirms our theoretical prediction. Regarding the policy of providing historical cutoffs, we observe that cutoffs significantly improve the welfare of students relative to DirSD in the environments with high search costs. In these environments,

²⁵Since we do not calculate a point prediction for the Cutoff treatment, there is no marker for the predicted average payoff under Cutoff in Figure 2.



Notes: Vertical gray bars represent the 95% confidence intervals. The square marker indicates the theoretical prediction for DirSD. The diamond marker indicates the theoretical prediction for SeqSD. The y-scale is in AUD.

Figure 2: Average payoffs by treatments

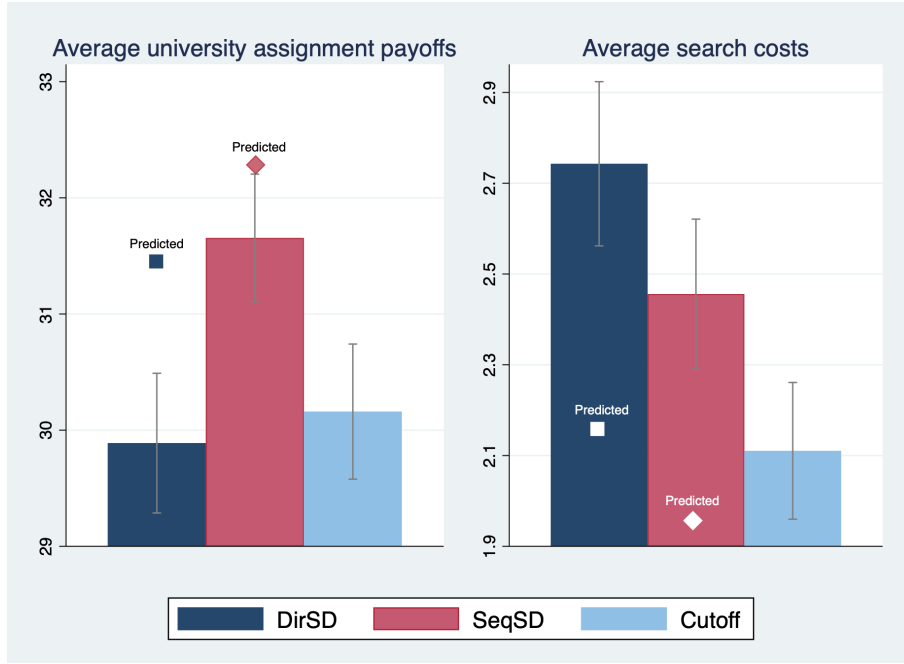
the increase in welfare under Cutoff is similar to the increase under SeqSD relative to DirSD. SeqSD still generates higher welfare, but the difference is not significant. In the environments with low costs, the average payoffs of participants are significantly higher in SeqSD than in Cutoff. In all treatments, the average payoffs are significantly higher in environments with low costs than in environments with high costs.

	Treatment			<i>p</i> -value for test of equality		
	DirSD (1)	SeqSD (2)	Cutoff (3)	DirSD=SeqSD (4)	DirSD=Cutoff (5)	SeqSD=Cutoff (6)
Two tiers & low cost	26.7	27.6	26.9	0.01	0.53	0.02
Two tiers & high cost	25.0	26.2	26.1	0.00	0.01	0.61
One tier & low cost	32.6	34.9	32.4	0.01	0.79	0.00
One tier & high cost	24.1	27.9	26.7	0.00	0.00	0.07
All	27.1	29.2	28.0	0.00	0.05	0.00

Notes: For the tests in columns 4-6, we use the *p*-values for the coefficient of the treatment dummy in the OLS regression of payoffs on this dummy with standard errors clustered at the level of matching groups and with a sample restricted to the treatments that are of interest for the test.

Table 5: Average payoffs of subjects by treatments and environments

Next, we consider the two components of welfare separately. Figure 3 presents the average payoffs of participants from the university assignments and the average search costs by treatments. It emerges that the welfare benefits of SeqSD relative to DirSD can be attributed to both sources, namely more efficient matching outcomes and lower costs of information acquisition. Both treatment differences are predicted by the theory and turn out to be significant in the experiment ($p < 0.01$). At the same time, in both DirSD and SeqSD, participants receive lower than predicted



Notes: Vertical gray bars represent the 95% confidence intervals. Square markers indicate the theoretical predictions for DirSD. Diamond markers indicate the theoretical predictions for SeqSD. The y-scales are in AUD.

Figure 3: Average payoffs from the assignment and search costs by treatments

payoffs from the university assignment, despite higher than predicted search costs on average. The left panel shows that the difference between the predicted and realized payoffs from the assignment is higher in DirSD than in SeqSD in the sense that, in SeqSD, the predicted payoff lies in the 95% confidence interval of the realized payoffs, which is not the case in DirSD.

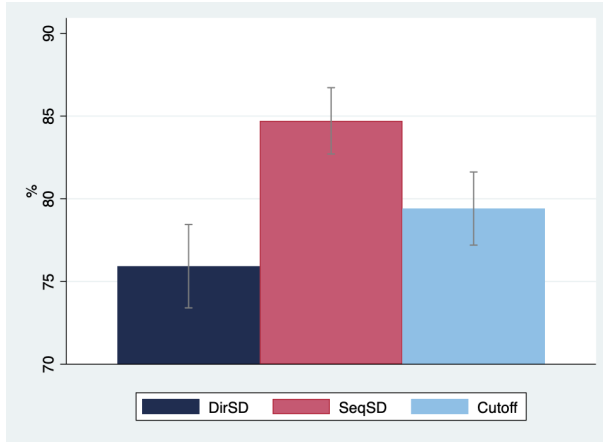
There is no significant difference in the assignment payoffs between DirSD and Cutoff ($p = 0.58$) while search costs are significantly lower under Cutoff than under DirSD ($p < 0.01$). Note that search costs are even lower under Cutoff than under SeqSD ($p = 0.01$).

The average payoff of participants is sensitive to market parameters, such as the preference draw and the costs of information acquisition. To make sure these factors do not bias the comparison between treatments, we check the robustness of our results with the help of a normalized efficiency measure as follows:

$$\text{Normalized Efficiency} = \frac{\text{Actual sum of payoffs} - \text{Minimum sum of payoffs}}{\text{Maximum sum of payoffs} - \text{Minimum sum of payoffs}} \in [0, 1].$$

We calculate the maximum sum of payoffs assuming every student has full information on her preferences (zero search cost) and receives the university allocation as in the stable match.²⁶ For the minimum sum of payoffs, we used the worst payoff based on 1,000 simulations of random

²⁶Note that this allocation is also Pareto efficient, and we calculate it by running SD under truthful preference submission.



Notes: Vertical gray bars represent the 95% confidence intervals. The scale is in percentage points.

Figure 4: Average normalized efficiency by treatments

strategies of participants assuming maximum search costs.²⁷

Figure 4 presents the average normalized efficiency by treatments. It is highest in SeqSD, with the difference to DirSD and to Cutoff being significant. Table 6 in Appendix B.2 shows the analogue of Table 5 for the normalized efficiency measure. The comparison of treatments in each of the environments is qualitatively the same in terms of statistical significance, showing the robustness of our main result to a different measure of welfare.

We summarize these findings in the following result.

Result 3 (Welfare)

(i) For the average payoff of participants and the normalized efficiency, the following relationships hold: $SeqSD > Cutoff > DirSD$.

(ii) For the average payoff of participants from the university assignments, the following relationships hold: $SeqSD > DirSD$, $SeqSD > Cutoff$, and $DirSD = Cutoff$ s. For search costs, the following relationships hold: $DirSD > SeqSD > Cutoff$.

5 Alternative search technology

The search technology studied so far relies on the acquisition of ordinal information. It captures important features of search in real-life scenarios. For example, centralized college admissions markets have a large number of programs with multidimensional characteristics, which makes it natural to focus on relative comparisons between programs when searching. Budish and Kessler (2022) provide evidence that people find it particularly difficult to report cardinal preferences and

²⁷We determine the minimum payoff using simulations, since we are unaware of an algorithm that finds the matching with the minimum sum of payoffs. Note that the submission of the opposite of truthful preferences at the individual level leads to the lowest expected individual payoff but can still result in a non-minimal sum of payoffs. This is due to positive externalities of the worst placements of higher-ranked students on lower-ranked students.

are much better at reporting ordinal preferences in an assignment problem, likely because they do not consider cardinal values when comparing different options in these types of environments.

However, in some important applications, search might uncover the cardinal value of an option. For instance, in a school choice setting, visiting a school might reveal the cardinal value of the school to students and parents. In this section, we introduce additional experiment sessions designed to check the robustness of our results against an alternative search technology, under which each step of the search leads to the discovery of a university's cardinal value. We focus our discussion on one-tier markets as the generalization to markets with multiple tiers does not depend on a particular search technology.

5.1 Theoretical analysis

Adopting a different search technology does not change our theoretical results regarding preference submission, tiered priors, and student welfare because the proofs of propositions 1 to 3 and Theorem 1 do not rely on a particular search procedure. In this section, we provide a theoretical analysis of the search strategies in a one-tier market under the alternative search framework.

To model the search technology that operates on cardinal utilities, we replace item 6 of the university admissions problem defined in Section 2.1 with the following:

- 6'. For each student $i \in I$, a set of cardinal utilities $u_i = \{u_i^{c_1}, \dots, u_i^{c_m}\}$: student i receives u_i^c when assigned to university c . For any $c, c' \in C$, $u_i^c > u_i^{c'}$ if and only if $c \succ_i c'$.

Corresponding to the assumption of uniform within-tier priors, we further specify that for any student $i \in I$ and any university $c \in C$ in the one-tier market, u_i^c is independent and identically distributed according to a uniform distribution $U[a, b]$ with $b > a \geq 0$. In each step of search, with a cost of k_i , the student chooses a school c to investigate and discovers the realization of u_i^c . With m steps of search and a total cost of mk_i , she obtains full knowledge of her own preferences. The student can choose to stop at any step in the search process or to not search at all.

Under SeqSD, each student i observes the realization of her budget set B_i before conducting a search and therefore only has incentives to search and select universities within B_i . This significantly simplifies information acquisition under SeqSD. Suppose that at a given step of the search, the set of universities that student i has chosen to search is given by C_i^S . Among the searched universities, let \bar{c}_i^S be the one with the highest utility and let its utility be \bar{u}_i^S , that is, $\bar{u}_i^S = \max_{c \in C_i^S} u_i^c$. In Appendix A.5, we characterize the optimal search strategy under SeqSD as a simple stopping rule that depends on k_i and \bar{u}_i^S in Proposition 7 and we provide the proof.

Under DirSD, a student i chooses her search strategy based on k_i and the ex-ante probability distribution of her budget set $\left\{P_i(\tilde{B})\right\}_{\tilde{B} \subseteq C}$. Unlike the simple stopping rule under SeqSD, the optimal search strategy is history-dependent. That is, whether a student should stop searching under DirSD may depend on the utilities of *all* the universities she has searched (instead of only the

highest discovered utility \bar{u}_i^S under SeqSD) because the student can be assigned to any university in her rank-order list. As a result, the number of contingencies grows rapidly as the student conducts each additional step of the search, and it is more challenging to derive the optimal search strategy under DirSD than under SeqSD. Therefore, instead of generally characterizing the optimal search strategy under DirSD, we numerically calculate it for every student in our experimental setup and describe the calculations in Appendix A.6. For the same reason explained in Section 2.6, we do not provide point predictions for search strategies under Cutoff but instead test whether subjects adopt a simple strategy of narrowing down their search using the cutoff information.

5.2 Experimental design

For every student in the experiment, the payoff of each university was drawn independently and randomly from the set $\{\$0, \$1, \$2, \dots, \$30\}$. There is one tier, and the cost of investigating an additional university was \$1 in the low-cost environment and \$1.5 in the high-cost environment. We chose these two search costs so we could have a wide range of predictions of the optimal search strategy depending on a student’s rank. The experiment consisted of four rounds with two consecutive rounds of each environment. The order of the two environments varied between sessions. The remaining design and procedure followed our main experiment described in Section 3. We ran 12 independent groups per treatment, with 96 subjects per treatment and 288 subjects in total. The sessions were conducted in the same lab as our main experiment, with average earnings of 21.72 AUD.

5.3 Experimental results

5.3.1 Search behavior

Each participant’s optimal number of searches is calculated depending on her entire search history, that is, all the cardinal utilities of the universities she has previously investigated. Figure 5 presents the frequency distribution of search strategies for DirSD and SeqSD. We skip the distribution for Cutoff for which we do not have point predictions regarding optimal search.

The left side of Figure 5 presents the distribution of search strategies in DirSD. Over-search is most common (54%), while 43% searched optimally. We observe only 3% of subjects who under-search in DirSD. The distribution is very different for SeqSD (right panel of Figure 5). Most subjects searched optimally (66%), while 19.5% over-searched and 14.5% under-searched in SeqSD. Thus, subjects are significantly more likely to search optimally in SeqSD than in DirSD ($p < 0.01$) and significantly more likely to under-search ($p < 0.01$). As for over-search, participants are significantly more likely to over-search in DirSD than SeqSD ($p < 0.01$). Moreover, conditional on over-search, participants search an average of 3.5 universities too many in DirSD and 2.2

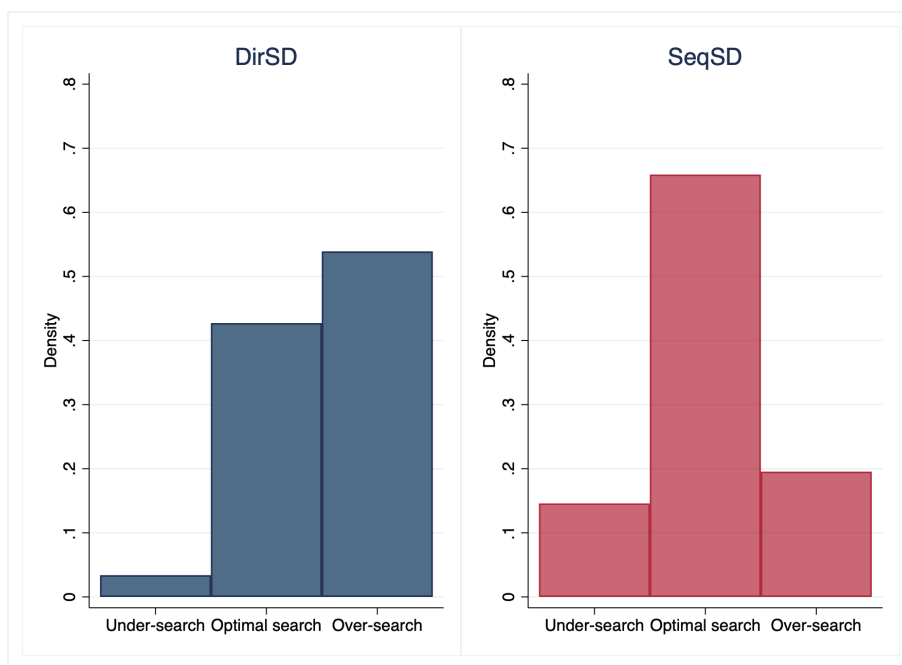


Figure 5: Optimality of search by treatments

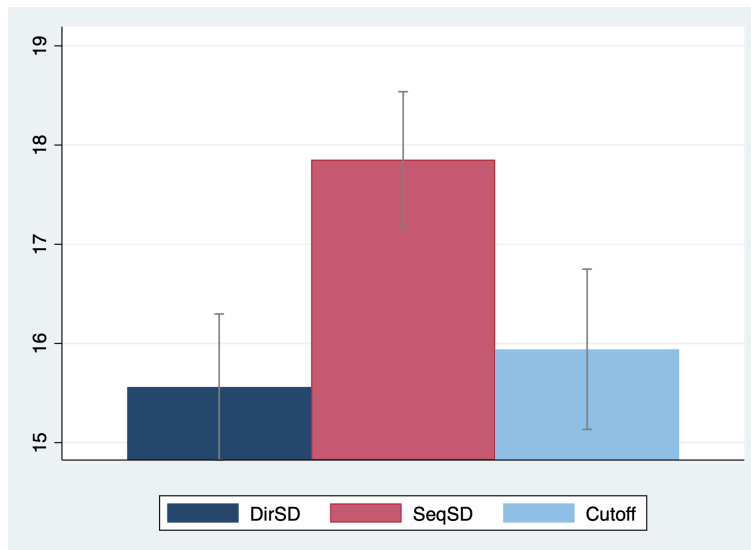
universities in SeqSD ($p < 0.01$).²⁸ Thus, the results from the main experiment are replicated: the search is more in line with the prediction in SeqSD than in DirSD. Under Cutoff, we do not have point predictions and thus use regressions to analyze the subjects' responses to the cutoffs, see Appendix B.4. Again, we replicate the pattern from the main experiment that students are less likely to search universities with cutoffs higher than their score compared to universities with cutoffs lower than their score.

In terms of submission strategies, we find that the proportion of optimal strategies is significantly higher under SeqSD than under DirSD and Cutoff, which is consistent with our results from the main experiment. Conditional on subjects' search behavior, the proportion of optimal submission strategies is 97.9% under SeqSD, 84.6% under DirSD, and 80.2% under Cutoff. The proportion under SeqSD is significantly higher than under DirSD ($p < 0.01$) and Cutoff ($p < 0.01$), while the difference between DirSD and Cutoff is not significant ($p = 0.21$).

Lastly, we analyze student welfare. Figure 6 presents the average payoffs under the three treatments.²⁹ SeqSD leads to the highest welfare among all treatments, the difference to both other treatments being significant ($p < 0.01$). This is in line with the findings from our main experiment. However, there is no significant welfare difference between DirSD and Cutoff. One potential explanation is that the difference between DirSD and Cutoff is mainly driven by the high-cost environment in our main experiment, while the two cost levels in the current experiment

²⁸In the case of under-search, we do not know the realization of a student's optimal stopping point because we do not observe the realization of the discovered cardinal utilities if the student would have continued searching. Therefore, here we do not compare the extent of under-search in DirSD and SeqSD.

²⁹Due to the history dependence of optimal search, we cannot derive theoretical point predictions.



Notes: Vertical gray bars represent the 95% confidence intervals. The y-scale is in AUD.

Figure 6: Average payoffs by treatments

are relatively low.

We can consider the two components of welfare separately, namely the payoffs from university assignments and the search costs. SeqSD leads to the highest average payoff from university assignments ($p=0.04$ relative to DirSD, $p<0.01$ relative to Cutoff), while the difference between Cutoff and DirSD is not significant ($p=0.07$). The average search cost is highest under DirSD ($p<0.01$ relative to both SeqSD and Cutoff), while the difference between SeqSD and Cutoff is not significant ($p=0.83$). Again, we replicate the result from the main experiment that the welfare benefits of SeqSD relative to DirSD can be attributed to both sources, namely more efficient matching outcomes and lower costs of information acquisition. Cutoff achieves a similar matching outcome to DirSD with lower search costs.

Summing up, most of the findings from the main experiment are robust to the alternative search technology and the variation in search costs. SeqSD leads to the highest welfare and students make significantly fewer mistakes both in their search and submission strategies. However, the benefits of Cutoff relative to DirSD are less robust. We still find that Cutoff leads to significant savings on search costs but the overall welfare improvement from cutoffs is small and seems to depend on the level of search costs.

6 Discussion: Practicalities of sequential mechanisms

Enabled by the digitization of assignment procedures and the possibility to communicate and coordinate repeatedly in real time through online platforms, sequential mechanisms have become more common in practice. Our results show that such mechanisms can improve student welfare in

markets with costly information acquisition about preferences.

However, a sequential procedure may last substantially longer than a direct mechanism. There is a tradeoff between the additional time required to run the sequential procedure and the benefit it provides in terms of information acquisition. This tradeoff is hard to model and to address in a lab experiment, as it is difficult to compare the cost of a longer wait to the gain in the quality of the assignment *ex ante*.

In many countries, students applying to colleges conduct research online to learn about the differences between programs with respect to academic quality, costs and scholarships, facilities and location, housing availability, etc. These search activities may not be very time-consuming, such that giving students enough time to search through a sequential mechanism is feasible. For example, there is evidence from Germany that students search for information during the admissions procedure (Grenet et al., 2022). In university admissions in France, students have between two to five days after receiving an offer to search for information and compare offers. The procedure has been run successfully for some years, and the main admissions are completed within 50 days (Hakimov et al., 2022).

Note that the tradeoff for students depends on their ranking and preferences. In a sequential mechanism, lower-ranked students wait longer for their turn and thus incur higher waiting costs. But they may also benefit more from the sequential mechanism in terms of lower search costs, because their budget sets shrink more during the procedure. In general, the design of a sequential mechanism should take the specific market conditions into consideration, and policymakers need to decide on some key features. First, it is important to determine the amount of time each student is given to make a decision when it is her turn in the sequential procedure. A longer decision time allows students to acquire more information about their preferences, but may result in a longer matching process. For example, in the first year that France switched to a sequential mechanism, the market took too long to clear due to too long decision times. An appropriate time constraint can increase welfare because it can also prevent over-search, a tendency of students that we observe in the data. Second, when the market is large, it is unlikely that students can move one after the other. However, sequential decisions by groups of students, depending on their scores, are realistic. Tunisia uses a sequential version of SD, with students with top, middle, and lower grades deciding sequentially, after the choices of the previous group are finalized. Similarly, the Chinese province of Inner Mongolia uses a dynamic mechanism where the choices of students are finalized sequentially by groups, depending on students' scores. Students with higher scores leave the procedure before students with lower scores (Gong and Liang, 2016). Intuitively, sequential decisions based on groups preserve the benefits of sequential mechanisms from the perspective of information acquisition, as lower-ranked groups of students do not have to acquire information about universities that have been filled by higher-ranked students. When implementing a sequential mechanism by groups, policymakers need to determine the size of the groups. A larger group size can shorten the matching process, but may decrease the benefit of a sequential mechanism with

respect to information acquisition. The more students move simultaneously, the less certain they are of being accepted by a university, even if it is still available when the members of the group have made their choices.

To sum up, it is possible to implement a sequential procedure in various environments as long as the design is tailored to the specific context. Nevertheless, a direct mechanism may be a better choice for markets where information acquisition is not a big concern. For example, school choice markets in relatively small cities can have few options such that students are already well informed about their preferences. Apart from the shorter time it takes, another benefit of a direct mechanism is that it does not require online coordination. Thus, in places where students and parents have limited internet access, a direct mechanism with historical cutoff provision might be a good option, especially when information acquisition costs are high.

7 Conclusions

We explore how students search university programs, how wasteful information acquisition can be reduced, and how student welfare can be improved in a market where students are ranked by universities based on exam scores. In theory, a sequential serial dictatorship mechanism leads to less wasteful information acquisition and higher student welfare than a direct serial dictatorship mechanism. We find that the theory underestimates these benefits, as the participants in the experiment make superior decisions both when searching and when choosing among universities under sequential serial dictatorship compared to direct serial dictatorship. We also find that the provision of cutoffs can increase student welfare, especially when information costs are high, although the effects of cutoffs are not as strong as the effects of using a sequential mechanism. With cutoffs, we observe that participants follow a simple strategy and avoid searching universities with cutoffs higher than their scores, especially when the costs are high. This simple strategy results in higher welfare with cutoff provision than under direct serial dictatorship without cutoff provision in high-cost environments.

The symmetry of universities and the tiered prior structures are simplifying assumptions that do not describe all school and college admissions markets in practice.³⁰ We have chosen the relatively simple environment for reasons of tractability and to derive predictions for the optimal search strategies in the direct mechanism. With a more complex market structure, a student in the direct mechanism needs to consider the search strategies of higher-ranked students when forming beliefs about her budget set. Such reasoning is not necessary in the sequential mechanism where the budget set is known. Therefore, we conjecture that a more complex market structure may increase the behavioral benefits of SeqSD and strengthen our results regarding the advantages of

³⁰For example, we do not study a market in which students are not exactly sure which tier a certain school belongs to. That is, students in our model do not face a tradeoff between exploring “risky” versus “safe” options, which can be considered as an important feature of costly information acquisition in some real-life markets.

the sequential mechanism. In line with this, we have suggestive evidence that the benefit of SeqSD relative to DirSD increases when the market grows and when the quality of programs is more uncertain.

Finally, if it is not possible to implement a sequential mechanism, providing historical cutoffs may improve welfare. This policy is used in some countries, and our study provides empirical support for it, especially when the costs of information acquisition are high.

References

- Abdulkadiroğlu, Atila, Joshua Angrist, and Parag Pathak (2014) “The elite illusion: Achievement effects at Boston and New York exam schools,” *Econometrica*, 82 (1), 137–196.
- Abdulkadiroğlu, Atila and Tayfun Sönmez (1998) “Random serial dictatorship and the core from random endowments in house allocation problems,” *Econometrica*, 66 (3), 689–701.
- Abdulkadiroğlu, Atila and Tayfun Sönmez (2003) “School Choice: A Mechanism Design Approach,” *The American Economic Review*, 93 (3), 729–747, <http://dx.doi.org/10.2307/3132114> 10.2307/3132114.
- Ajayi, Kehinde F, Willa H Friedman, and Adrienne M Lucas (2020) “When Information is Not Enough: Evidence from a Centralized School Choice System,” Technical report, National Bureau of Economic Research.
- Ajayi, Kehinde and Modibo Sidibe (2020) “School choice under imperfect information,” Technical report, mimeo.
- Artemov, Georgy (2021) “Assignment mechanisms: common preferences and information acquisition,” *Journal of Economic Theory*, 198, 105370.
- Atakan, Alp E (2006) “Assortative matching with explicit search costs,” *Econometrica*, 74 (3), 667–680.
- Azevedo, Eduardo M and Jacob D Leshno (2016) “A supply and demand framework for two-sided matching markets,” *Journal of Political Economy*, 124 (5), 1235–1268.
- Baccara, Mariagiovanna, Ayşe İmrohoroğlu, Alistair J Wilson, and Leeat Yariv (2012) “A field study on matching with network externalities,” *American Economic Review*, 102 (5), 1773–1804.
- Bade, Sophie (2015) “Serial dictatorship: The unique optimal allocation rule when information is endogenous,” *Theoretical Economics*, 10(2), 385–410.
- Basteck, Christian, Bettina Klaus, and Dorothea Kübler (2021) “How lotteries in school choice help to level the playing field,” *Games and Economic Behavior*.

- Bhattacharya, Sourav, John Duffy, and SunTak Kim (2017) “Voting with endogenous information acquisition: Experimental evidence,” *Games and Economic Behavior*, 102, 316–338.
- Bó, Inácio and Rustamdjan Hakimov (2020) “Iterative versus standard deferred acceptance: Experimental evidence,” *The Economic Journal*, 130 (626), 356–392.
- Bo, Inacio and Rustamdjan Hakimov (2022) “The iterative deferred acceptance mechanism,” *Games and Economic Behavior*, 135, 411–433.
- Bó, Inacio and Rustamdjan Hakimov (2023) “Pick-an-object mechanisms,” *Management Science*.
- Bó, Inácio and Chiu Yu Ko (2020) “Competitive screening and information transmission,” *Journal of Public Economic Theory*.
- Budish, Eric and Judd B Kessler (2022) “Can market participants report their preferences accurately (enough)?” *Management Science*, 68 (2), 1107–1130.
- Chade, Hector and Lones Smith (2006) “Simultaneous search,” *Econometrica*, 74 (5), 1293–1307.
- Chen, Yan and Yinghua He (2021) “Information acquisition and provision in school choice: an experimental study,” *Journal of Economic Theory*, 197, 105345.
- (2022) “Information acquisition and provision in school choice: a theoretical investigation,” *Economic Theory*, 74 (1), 293–327.
- Chen, Yan and Onur Kesten (2017) “Chinese college admissions and school choice reforms: A theoretical analysis,” *Journal of Political Economy*, 125 (1), 99–139.
- Coles, Peter and Ran Shorrer (2014) “Optimal truncation in matching markets,” *Games and Economic Behavior*, 87, 591–615.
- Cox, Natalie, Ricardo Fonseca, and Bobak Pakzad-Hurson (2021) “Do Peer Preferences Matter in School Choice Market Design? Theory and Evidence.”
- Das, Sanmay and Zhuoshu Li (2014) “The role of common and private signals in two-sided matching with interviews,” in *International Conference on Web and Internet Economics*, 492–497, Springer.
- Descamps, Ambroise, Sébastien Massoni, and Lionel Page (2021) “Learning to hesitate,” *Experimental Economics*, 1–25.
- Drummond, Joanna and Craig Boutilier (2014) “Preference elicitation and interview minimization in stable matchings,” in *Twenty-Eighth AAAI Conference on Artificial Intelligence*.

- Dur, Umut, Robert G Hammond, and Onur Kesten (2017) “Sequential School Choice: Theory and Evidence from the Field and Lab,” Technical report, Working Paper, North Carolina State University.
- Eeckhout, Jan and Philipp Kircher (2010) “Sorting and decentralized price competition,” *Econometrica*, 78 (2), 539–574.
- Engelmann, Dirk and Veronika Grimm (2009) “Bidding behaviour in multi-unit auctions—an experimental investigation,” *The Economic Journal*, 119 (537), 855–882.
- Gale, D. and L. S. Shapley (1962) “College Admissions and the Stability of Marriage,” *The American Mathematical Monthly*, 69 (1), 9–15.
- Gong, Binglin and Yingzhi Liang (2016) “A dynamic college admission mechanism in Inner Mongolia: Theory and experiment,” Technical report, Working paper, 2016. 2, 10.
- Grenet, Julien, YingHua He, and Dorothea Kübler (2022) “Preference Discovery in University Admissions: The Case for Dynamic Multioffer Mechanisms,” *Journal of Political Economy*, 130 (6), 000–000.
- Gretschko, Vitali and Alexander Rajko (2015) “Excess information acquisition in auctions,” *Experimental Economics*, 18 (3), 335–355.
- Hakimov, Rustamdjan and Dorothea Kübler (2021) “Experiments on centralized school choice and college admissions: a survey,” *Experimental Economics*, 24, 434–488.
- Hakimov, Rustamdjan and Madhav Raghavan (2022) “Improving Transparency and Verifiability in School Admissions: Theory and Experiment,” Working Paper.
- Hakimov, Rustamdjan, Renke Schmacker, and Camille Terrier (2022) “Self-censoring in college applications.”
- Harless, Patrick and Vikram Manjunath (2018) “Learning matters: Reappraising object allocation rules when agents strategically investigate,” *International Economic Review*, 59 (2), 557–592.
- Hastings, Justine S, Christopher A Neilson, and Seth D Zimmerman (2013) “Are some degrees worth more than others? Evidence from college admission cutoffs in Chile,” Technical report, National Bureau of Economic Research.
- Immorlica, Nicole S., Jacob D. Leshno, Irene Y. Lo, and Brendan J. Lucier (2020) “Information Acquisition in Matching Markets: The Role of Price Discovery,” mimeo, CRC TRR 190.
- Kadam, Sangram V (2015) “Interviewing in matching markets,” Technical report, Working Paper.

- Kagel, John H, Ronald M Harstad, and Dan Levin (1987) “Information Impact and Allocation Rules in Auctions with Affiliated Private Values: A Laboratory Study,” *Econometrica*, 55 (6), 1275–1304.
- Kagel, John H and Dan Levin (2001) “Behavior in Multi-Unit Demand Auctions: Experiments with Uniform Price and Dynamic Vickrey Auctions,” *Econometrica*, 69 (2), 413–454.
- (2009) “Implementing efficient multi-object auction institutions: An experimental study of the performance of boundedly rational agents,” *Games and Economic Behavior*, 66 (1), 221–237.
- Klijn, Flip, Joana Pais, and Marc Vorsatz (2019) “Static versus dynamic deferred acceptance in school choice: Theory and experiment,” *Games and Economic Behavior*, 113, 147–163.
- Koh, Youngwoo and Wooyoung Lim (2022) “When to Visit: Information Acquisition in College Admissions,” Technical report, Working Paper.
- Lee, Robin S and Michael Schwarz (2017) “Interviewing in two-sided matching markets,” *The RAND Journal of Economics*, 48 (3), 835–855.
- Li, Shengwu (2017) “Obviously strategy-proof mechanisms,” *American Economic Review*, 107 (11), 3257–87.
- Lien, Jaimie W, Jie Zheng, and Xiaohan Zhong (2016) “Preference submission timing in school choice matching: testing fairness and efficiency in the laboratory,” *Experimental Economics*, 19 (1), 116–150.
- Lufade, Margaux (2019) “The value of information in centralized school choice systems (job market paper).”
- Narita, Yusuke (2018) “Match or Mismatch? Learning and Inertia in School Choice,” Unpublished.
- Noda, Shunya (2020) “Strategic Experimentation with Random Serial Dictatorship,” *Available at SSRN 2956548*.
- Pan, Siqu (2019) “The instability of matching with overconfident agents,” *Games and Economic Behavior*, 113, 396–415.
- (2020) “Experiments in market design,” in *Handbook of Experimental Game Theory*: Edward Elgar Publishing.
- Pycia, Marek and Peter Troyan (2019) “A Theory of Simplicity in Games and Mechanism Design,” Working Paper.

- Rastegari, Baharak, Anne Condon, Nicole Immorlica, and Kevin Leyton-Brown (2013) “Two-sided matching with partial information,” in *Proceedings of the fourteenth ACM conference on Electronic commerce*, 733–750.
- Shimer, Robert and Lones Smith (2000) “Assortative matching and search,” *Econometrica*, 68 (2), 343–369.
- Shorrer, Ran I (2019) “Simultaneous search: Beyond independent successes,” in *Proceedings of the 2019 ACM Conference on Economics and Computation*, 347–348.
- Svensson, Lars-Gunnar (1999) “Strategy-proof allocation of indivisible goods,” *Social Choice and Welfare*, 16 (4), 557–567.
- Zimmerman, Seth D (2019) “Elite colleges and upward mobility to top jobs and top incomes,” *American Economic Review*, 109 (1), 1–47.