Operations Research

Publication details, including instructions for authors and subscription information:
http://pubsonline.informs.org

Improving the Chilean College Admissions System

Ignacio Rios , Tomás Larroucau , Giorgiogiulio Parra , Roberto Cominetti

informs
http://pubsonline.informs.org/journal/opre

OPERATIONS RESEARCH
Vol. 69, No. 4, July–August 2021, pp. 1186–1205
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Crosscutting Areas

Improving the Chilean College Admissions System

Ignacio Rios,[a] Tomás Larroucau,[b] Giorgiogiulio Parra,[c] Roberto Cominetti[d]

[a] Naveen Jindal School of Management, University of Texas at Dallas, Richardson, Texas 75080; [b] Department of Economics, University of Pennsylvania, Philadelphia, Pennsylvania 19104; [c] Departamento de Evaluación, Medición y Registro Educacional (DEMRE), Universidad de Chile, 8320000 Santiago, Chile; [d] Facultad de Ingeniería y Ciencias, Universidad Adolfo Ibáñez, 9250000 Santiago, Chile

Contact: ignacio.riosuribe@utdallas.edu, https://orcid.org/0000-0002-8526-9353 (IR); tomasl@sas.upenn.edu (TL); roberto.cominetti@uai.cl, https://orcid.org/0000-0001-9442-344X (RC)

Received: February 8, 2018
Revised: January 31, 2019; August 19, 2020
Accepted: November 3, 2020
Published Online in Articles in Advance: March 24, 2021

OR/MS Subject Classification: networks/graphs: matchings; education systems: operations
Area of Review: OR Practice

https://doi.org/10.1287/opre.2021.2116

Abstract. In this paper we present the design and implementation of a new system to solve the Chilean college admissions problem. We develop an algorithm that obtains all applicant/program pairs that can be part of a stable allocation when preferences are not strict and when all students tied in the last seat of a program (if any) must be allocated. We use this algorithm to identify which mechanism was used in the past to perform the allocation, and we propose a new method to incorporate the affirmative action that is part of the system to correct the inefficiencies that arise from having double-assigned students. By unifying the regular admission with the affirmative action, we have improved the allocation of approximately 2.5% of students assigned every year since 2016. From a theoretical standpoint, we show that some desired properties, such as strategy-proofness and monotonicity, cannot be guaranteed under flexible quotas.

Funding: This research was supported by the Núcleo Milenio Información y Coordinación en Redes [ICM/FIC P10-024F].
Supplemental Material: The electronic companion is available at https://doi.org/10.1287/opre.2021.2116.

Keywords: college admissions • stable assignment • flexible quotas • nonstrict preferences

1. Introduction

A centralized mechanism to match students to programs[1] has been used in Chile since the late 1960s by the Departamento de Evaluación, Medición y Registro Educacional (DEMRE), the analogue of the American College Board. Every year, more than 250,000 students participate in the system, which includes more than 1,400 programs in 41 universities. This system has two main[2] components: a regular admission track, where all students who graduated from high school can participate, and an affirmative action policy that aims to benefit underrepresented groups by offering them reserved seats and economic support. More specifically, to be considered for the reserved seats and the scholarship—called the Beca de Excelencia Académica, or simply BEA—a student must belong to the top 10% of his class, he must graduate from a public/voucher school, and his family income must be among the lowest four quintiles of the national income distribution.

When the affirmative action was introduced in 2007, the procedure to match students to programs relied on a black-box software that could not be updated to incorporate this new feature. Hence, the authorities decided that the admission of BEA students would be run after the admission of regular students. Because BEA students can apply to both regular and reserved seats,[3] running the process sequentially introduces inefficiencies. For instance, a BEA student can be assigned to two different programs, and the seat that this student decides not to take cannot be reallocated to another student. Because of this problem, more than 1,000 vacancies were not filled every year, mainly affecting students from underrepresented groups.

In this paper we provide a "reverse-engineering" approach to correct these inefficiencies. The reason why we start from the current system and do not simply propose a complete redesign is that DEMRE wanted to keep the current rules and incorporate the reserved seats keeping the system as close as possible to its current state. Hence, there were two practical challenges to address: (1) to identify the mechanism that was currently being used and (2) to modify this mechanism to unify the admission tracks. To address these challenges and eliminate the aforementioned inefficiencies, our first goal was to identify the mechanism "inside the black-box." Going by the rules of the system, we had enough evidence to think that the desired outcome was a stable matching. In addition, we realized that, unlike other systems, all students tied in the last seat had to be admitted, so quotas had to be *flexible* in order to allocate them. With these features in mind, we implemented an algorithm based on an algorithm introduced by Baïou and Balinski (2004) that obtains all applicant/program pairs that can be part of a stable assignment, along with two extreme allocations: the student-optimal and the university-optimal stable assignments. By comparing the results of our algorithm with the actual assignment of past years, we found that the mechanism used was the university-proposing deferred acceptance, with the special feature of flexible quotas to allocate all tied students in the last seat. Furthermore, we show that, unlike the case with strict preferences, the Chilean mechanism is not strategy-proof nor monotone. Nevertheless, we argue that flexible quotas do not introduce a major strategic concern, given the large size of the market.

After identifying the algorithm being used, our next goal was to integrate both systems in order to maximize the utilization of vacancies. To solve this problem, we introduce a new approach where each type of seat (regular or reserved) is assumed to belong to a different program with its own capacity and requirements, and students benefited by the affirmative action can apply to both.

This research is the outcome of an ongoing multiyear collaboration with DEMRE (2012–2020), aiming to improve the Chilean college admissions system. All the solutions described in this paper were adopted and implemented starting in 2014 with a pilot phase. In 2015, the system switched to a student-optimal mechanism with flexible quotas, and in 2016, the unified allocation was finally adopted. On the basis of simulations in 2014–2015 and the results of the final implementation in 2016, we find that our implementation has improved the allocation of approximately 2.5% of the students assigned each year. Furthermore, our "white-box" implementation made the admission process fully transparent and reduced the execution time from over five hours to a couple of minutes. This improvement in transparency and performance has allowed the evaluation and introduction of different policies (e.g., the inclusion of the high school class rank as an admission factor; see Larroucau et al. (2015), among others) that otherwise could not have been studied.

The rest of this paper is organized as follows. Section 2 provides a background on the Chilean tertiary education system and the college admissions process. In Section 3 we discuss the closest related literature. In Section 4 we develop a model that formalizes the problem, describe the mechanisms, and present their properties. We discuss the implementation in Section 5. Finally, we provide concluding remarks in Section 6.

2. The Chilean College Admissions System

Tertiary education in Chile is offered by 149 institutions,[4] which can be classified as three types: (i) universities (60), which have the exclusive right to award academic degrees—bachelor's, master's, and doctorate—and offer academic programs that require a previous degree, such as medicine and law; (ii) professional institutes (IPs) (43), which offer professional/technical programs that lead to a professional/technician qualification; and (iii) technical schooling centers (CFTs) (46), which exclusively offer vocational programs leading to a technician qualification. These institutions not only differ in the type of programs they offer but also differ in their programs' duration[5] and their application requirements. In particular, IPs and CFTs only require a secondary education license for admission, and some professional institutes may select their students based on their grades in high school. By contrast, most universities require students to take a series of standardized tests (Prueba de Selección Universitaria, or PSU). These tests include math, language, and a choice between science and history, providing a score for each of them. The performance of students in high school gives two additional scores, one obtained from the average grade in high school (*notas de enseñanza media*, or NEM) and a second that depends on the relative position of the student among his or her cohorts (*ranking de notas*, or rank).[6]

The admissions process to these institutions is semicentralized, with the most selective universities having their own centralized system and the remaining institutions carrying their admission processes independently. In the centralized system, which is organized by the Consejo de Rectores de las Universidades Chilenas (CRUCH),[7] students submit a single application list, and a centralized algorithm simultaneously performs the allocation to all participating programs. To participate in this system, universities must (i) certify their quality, (ii) guarantee that their controllers are nonprofit organizations, and (iii) agree with the terms and conditions, such as publishing their requirements for admission and the number of seats offered for each program, among other things. On the other hand, IPs, CFTs, and the universities that are not part of the centralized system run their admissions independently.[8]

In this paper we focus on the centralized part of the system, whose timeline is summarized in Figure 1. The process starts in May, when each program defines the specific requirements that must be met by applicants to be acceptable, such as the minimum application score or minimum tests scores. In addition,

Figure 1. Timeline of the Centralized Process

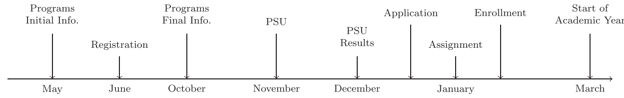

each program freely[9] defines the weights assigned to each score and also the number of seats offered for (i) the regular process, where all students compete, and for (ii) the special admission track related to the affirmative action policy (BEA process/track). Programs have until October to update this information. In June, students must register to take the PSU, which takes place at the end of November. Scores are published by the end of December, and right after this, the application process starts. Students have five days to submit their list of preferences, which can contain at most 10 different programs. These programs must be listed in strict order of preference.

Each program's preference list is defined by first filtering out all applicants that do not meet the specific requirements. Then, students are ordered in terms of their application scores, which are computed as the weighted sum of the applicants' scores and the weights predefined by each program. Note that two candidates can obtain the same application score, and therefore programs' preferences are not necessarily strict.

Considering the preference lists of applicants and programs, as well as the number of seats offered in both admission tracks, DEMRE runs an assignment algorithm to match students and programs. Specifically, the regular process is solved first considering all applications and the regular seats. Once the regular process is done, the BEA process is solved considering the reserved seats, the students shortlisted for the scholarship (BEA students), and those students' applications to programs that (i) rank higher in their preference list than the program they were assigned to in the regular process (if any) and (ii) wait-listed them in the regular process. As a result, BEA students can be *double-assigned*; that is, they can get assigned to a program in the regular process and to another program—a strictly preferred program—in the BEA process. DEMRE reports both allocations, and double-assigned students are allowed to enroll in any of their two assigned programs.

DEMRE performs the matching for both processes using a black-box software for which no information is available regarding the specific algorithm used. Instead, the following description is provided for the regular process:[10]

SORTING OF APPLICANTS PER PROGRAM AND ELIMINATION OF MULTIPLE ALLOCATIONS:

(a) Once the final application score is computed, candidates will be ordered in strict decreasing order based on their scores in each program.

(b) Programs complete their vacancies starting with the applicant that is first in the list of candidates, and continue in order of precedence until seats are full.

(c) If an applicant is selected in his first choice, then he is erased from the lists of his 2nd, 3rd, 4th, until his last preference. If he is not selected in his first choice, he is wait-listed and moves on to compete for his 2nd preference. If he is selected in this preference, he is dropped from the list of his 3rd to his 10th choice, and so on. In this way, it is possible that a student is selected in his 6th preference and wait-listed in his top five preferences; however, he will be dropped from the lists of his preferences 7th to 10th.

(d) This procedure to select candidates is the result of an agreement between the universities to have a unified and integrated process, so that no student is admitted by more than one program. Nevertheless, a student can be wait-listed in more than one program if his score is not enough to be admitted.

(e) All candidates that apply and satisfy the requirements of the corresponding program and institution will be wait-listed.

THEREFORE, IT IS FUNDAMENTAL THAT APPLICANTS SELECT THEIR PROGRAMS IN THE SAME ORDER AS THEIR PREFERENCES.

This description suggests that the final allocation must be stable, in the sense that there is no student and program that simultaneously prefer to be matched together rather than to their matches in the proposed assignment. Indeed, as the results are public, and

Figure 2. (Color online) Nonstrict Preferences

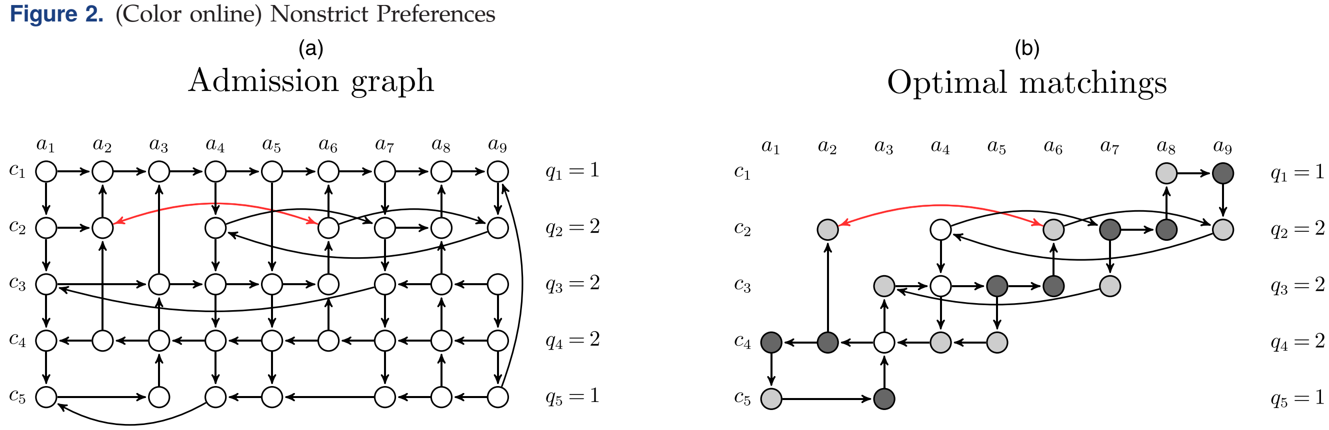

students can easily check whether their application scores are higher than that of the last student admitted in the program they prefer, legal problems may arise if the resulting matching was unstable. However, it is unclear from the description which specific stable assignment is implemented. Moreover, by analyzing the results of previous assignments, we realized that, in the case of a tie in the last seat of a program, the number of seats were increased in order to include all tied applicants. This feature of the system was confirmed by DEMRE, and it applies to both admission tracks (regular and BEA). From now on, we refer to this feature of the system as *flexible quotas*.[11]

The results of the assignment process are released by mid-January, and at this point, the enrollment process starts. In its first stage, which lasts for three days, students can enroll in the programs they were assigned (either in the regular process or in the BEA process). In its second stage, which lasts for one week, programs with seats left after the first stage can call students on their wait-lists and offer them the chance to enroll. This must be done in strict order of preference given by the application scores, and students must decline their enrollments in the first stage to enroll in a new program.[12] However, programs can decide not to call students if they have already filled a minimum number of regular seats,[13] and they are not forced to reallocate unassigned BEA seats. By the rules of the system, all seats left after the second stage of enrollment are lost, including those seats not taken by students with double-assignments.[14] This is the main source of inefficiency that we address in this paper.

3. Literature Review

Centralized admission systems have been increasingly used in recent years to carry out the assignment of students to schools and colleges. A variety of mechanisms have been studied, including the celebrated deferred acceptance (DA) algorithm (Gale and Shapley 1962), the immediate acceptance (Boston) algorithm (Abdulkadiroglu et al. 2005, Ergin and Sömnez 2006), and the top-trading cycles algorithm (Shapley and Scarf 1974). An important part of the literature in market design has been devoted to characterizing these mechanisms, mostly focusing on canonical examples that illustrate their properties. Another important body of literature studies real-life applications by combining the aforementioned mechanisms with specific rules, such as restrictions in the length of preferences, tie-breaking rules, and affirmative actions, among many others. In this paper we try to contribute to both by studying the Chilean college admissions problem.

The most closely related paper to ours is that by Biró and Kiselgof (2015), which analyzes the college admissions system in Hungary, where all students tied in the lowest-ranked group of a program are rejected if their admission would exceed the quota. This mechanism is opposed to the Chilean case, where the quota is increased just enough so that all tied students are admitted. Biró and Kiselgof (2015) formalize these ideas by introducing the concepts of H-stability and L-stability, which correspond to the rules in Hungary and Chile, respectively. They also provide a natural adaptation of DA to compute H-stability and L-stability based on ascending score limits, and they provide an alternative proof of the manipulability of H-stable and L-stable mechanisms. In a recent paper, Kamiyama (2017) presents a polynomial time algorithm to check whether a student can manipulate his preferences to obtain a better allocation. Our paper contributes to this strand of the literature by independently introducing the notion of L-stability, providing an algorithm based on Baïou and Balinski (2004) to find all pairs that may belong to an L-stable matching and implementing it to solve a real, large, and relevant problem.

Our paper is also related to the literature on affirmative action policies. Most of the research in this strand has focused on proposing mechanisms to solve the college admissions problem with diversity constraints and deriving properties such as stability, strategy-proofness, and Pareto optimality. From a theoretical perspective, Echenique and Yenmez (2015) point out that the main tension between diversity concerns and stability is the existence of complementarities, although the theory requires substitutability for colleges' choices. Abdulkadiroğlu (2007) explores the deferred acceptance algorithm under type-specific quotas and finds that the student-proposing DA is strategy-proof for students if colleges' preferences satisfy responsiveness. Kojima (2012) shows that majority quotas may actually hurt minority students. Consequently, Hafalir et al. (2013) propose the use of minority reserves to overcome this problem, showing that the deferred acceptance algorithm with minority reserves Pareto dominates the one with majority quotas. Ehlers et al. (2014) extend the previous model to account for multiple disjoint types, and they propose extensions of DA to incorporate soft and hard bounds. Other types of constraints are considered by Kamada and Kojima (2015), who study problems with distributional constraints motivated by a Japanese medical residency. The authors propose a mechanism that respects these constraints while satisfying other desirable properties such as stability, efficiency, and incentives.

Some authors have recently analyzed the impact of the order in which reserves are processed. Dur et al. (2018) analyze the Boston school system and show that the precedence order in which seats are filled has important quantitative effects on distributional objectives.

Dur et al. (2018) formalize our idea that processing reserved seats in a lower precedence order benefits BEA students. In a follow-up paper, Dur et al. (2020) characterize optimal policies when there are multiple reserve groups, and they analyze their impact using Chicago's system data.

Finally, our paper also contributes to the literature on designing large-scale clearinghouses. Institutional details and special requirements oftentimes forbid the use of tools directly taken from the theory, and other engineering aspects become relevant in the design process (Roth 2002). Roth and Peranson (2002) report the design of a new clearinghouse to organize the labor market for new physicians in the United States. Since 1997, when the new algorithm was finally adopted by the National Resident Matching Program (NRMP), more than 20,000 doctors have been matched to entry-level positions every year, and other labor markets have adopted the Roth–Peranson design, including dental, pharmacy, and medical residencies (see Roth (2002) for other examples). In the school choice context, Abdulkadiroğlu et al. (2005) describe the design of a new mechanism to match entering students to public high schools in New York. The new algorithm helped to dramatically reduce the number of students assigned to schools for which they had expressed no preference, and it has motivated the implementation of centralized clearinghouses in other school districts (e.g., Boston, Amsterdam, New Orleans, and Chicago, among others) and also nationwide (see Correa et al. (2019)). Closer to our setting, Baswana et al. (2019) describe the design and implementation of a clearinghouse to perform the allocation of students to technical universities in India. Their heuristic approach, which also allocates all tied students in the last seat and extends DA to accommodate the multiple types of seat reservations for affirmative action, has been successfully running since 2015.

4. Model

The following framework is assumed hereafter. Consider two finite sets of agents: programs, $C = \{c_1, \dots, c_m\}$, and applicants, $A = \{a_1, \dots, a_n\}$. Let $V \subseteq C \times A$ be the set of *feasible pairs*, with $(c, a) \in V$ meaning that student a has submitted an application to program c and meets the specific requirements to be admissible in that program. A *feasible assignment* is any subset $\mu \subseteq V$. We denote by $\mu(a) = \{c \in C : (c, a) \in \mu\}$ the set of programs assigned to a, and we denote by $\mu(c) = \{a \in A : (c, a) \in \mu\}$ the set of students assigned to program c. Each program c has a quota $q_c \in \mathbb{N}$ that limits the number of students that the program can admit. Moreover, program $c \in C$ ranks applicants according to a *total preorder* $\leq_c$ (i.e., a transitive relation in which all pairs of students are comparable). The indifference $a \sim_c a'$ denotes, as usual, the fact that we simultaneously have $a \leq_c a'$ and $a' \leq_c a$, and we write $a <_c a'$ when $a \leq_c a'$ but not $a \sim_c a'$. On the other side of the market, each applicant $a \in A$ ranks programs according to a *strict total order* $<_a$; that is, for any programs c, c' such that $c >_a \emptyset$ and $c' >_a \emptyset$ (i.e., c, c' are acceptable to student a), we have either $c <_a c'$ or $c' <_a c$.

A *matching* is a feasible assignment $\mu \subseteq V$ such that for each applicant a, the set of assigned programs $\mu(a)$ has at most one element, whereas for each program c, the set of assigned students $\mu(c)$ has at most q_c elements. A matching μ is *stable* if for all pairs $(c, a) \in V \setminus \mu$ we have that either the set $\mu(a)$ has an element preferred over c in the strict order $<_a$ or the set $\mu(c)$ contains q_c elements preferred over a in the strict order $<_c$. In the first case, applicant a likes the match proposed by μ better than c; in the second case, the program has all its vacancies filled with students strictly preferred than a. If both conditions fail simultaneously, a and c would be better off by being matched together rather than accepting the assignment μ, in which case (c, a) forms a *blocking pair*. In other words, a matching is stable if it has no blocking pairs.

In their seminal paper, Gale and Shapley (1962) introduce the DA algorithm, which returns the stable matching that is most preferred by agents on the proposing side. Hence, by changing the proposing side, DA allows to find two extreme stable matchings: the student-optimal and the university-optimal. However, there are many reasons why the clearinghouse may want a stable outcome that is different from the extreme ones. For instance, the clearinghouse may be concerned about fairness (e.g., Teo and Sethuraman 1998, Schwarz and Yenmez 2011), or it may prefer to benefit some specific agents in the market. In a recent paper, Dworczak (2016) introduces the concept of deferred acceptance with compensation chains (DACC), which generalizes DA by allowing both sides of the market to propose. The author shows that a matching is stable if and only if it can be obtained through a DACC algorithm and provides an algorithm that finds the stable matching given a sequence of proposers.

The aforementioned approach could be used to obtain all stable matchings by considering different sequences of proposers. However, this would require running the algorithm for each potential sequence, which is inefficient, especially when the core of stable outcomes is relatively small. Therefore, we adopt an alternative approach and extend the algorithm introduced by Baïou and Balinski (2004), which uses a graph representation of the admissions problem. Following their approach, an *instance* of the college admissions problem can be fully described in terms of a pair $\Gamma = (G, q)$, where $G = (V, E)$ is an admission graph consisting of a set of feasible nodes V on a grid $C \times A$

and a set of directed arcs $E \subseteq V \times V$ that represent programs and applicants preferences; q is a vector of quotas. Each row in the grid represents a program $c \in C$, and each column represents an applicant $a \in A$. The preferences of program c are encoded by horizontal arcs from (c, a) to (c, a') whenever $a \leq_c a'$ and those of student a by vertical arcs from (c, a) to (c', a) representing $c <_a c'$. For simplicity, the arcs that can be inferred by transitivity are omitted.

By exploiting this graph representation, the algorithm proposed by Baïou and Balinski (2004) recursively eliminates pairs $(c, a) \in V$ that are strictly dominated and thus cannot belong to any stable matching. More precisely, (c, a) is *a-dominated* if there are q_c or more applicants that have c as their top choice and dominate a in the strict preference order $<_c$. In this case, program c is guaranteed to fill its quota with applicants strictly above a, so that a has no chance to be assigned to c, and the pair (c, a) can be eliminated from further consideration.

Similarly, (c, a) is *c-dominated* if there is a program c' that places a among the top $q_{c'}$ applicants (i.e., less than $q_{c'}$ applicants are ranked strictly above a) and that is preferred by a over c (i.e., $c' >_a c$). In this case, applicant a is guaranteed to be assigned to a program ranked at least as high as program c' in his preference list, so the pair (c, a) cannot belong to any stable assignment.

As a result, the algorithm returns a *domination-free subgraph* $G^* = (V^*, E^*)$, with a node set $V^* \subseteq V$ that includes all nodes that are not strictly dominated. The domination-free-equivalent subgraph G^* contains all possible stable allocations, including the two most interesting (and extreme) cases: the *student-optimal* matching, μ^*_A, that assigns each applicant $a \in A$ to its best remaining choice in G^* (if any); and the *university-optimal* matching, μ^*_C, that assigns to each program $c \in C$ its q_c top choices in G^*. In this way, the algorithm by Baïou and Balinski (2004) returns not only the allocations that could be obtained using DA but also the nodes that could potentially be in other stable outcomes.

In order to apply this mechanism to the Chilean case, we extend it to incorporate two special features: (1) the existence of ties and flexible quotas and (2) the affirmative action policy. The next two sections describe how we incorporate these elements into the mechanism.

4.1. Ties and Flexible Quotas: FQ-Matchings

Suppose now that programs' preferences may not be strict and that programs are required to adjust their quotas to include all applicants tied in the last seat. More precisely, a program c may exceed its quota q_c only if the last group of students admitted are in a tie, and upon rejecting all these students, c ends up with unassigned seats. We also impose a nondiscrimination condition: an applicant a' who is tied with a student a admitted to a program c must himself be granted admission to c or to a more preferred program. The following definitions state these conditions formally.

Definition 1. An assignment μ satisfies *quotas-up-to-ties* if, for each program c and each $a \in \mu(c)$, the set of strictly preferred students assigned to c satisfies $|\{a' \in \mu(c) : a' >_c a\}| < q_c$.

Definition 2. An assignment μ satisfies *nondiscrimination* if whenever $a \in \mu(c)$ and $a' \sim_c a$ with $(c, a') \in V$, then $a' \in \mu(c')$ for some program $c' \geq_{a'} c$.

With these preliminary definitions, we introduce our notion of matching, which additionally requires that each applicant be assigned to at most one program.

Definition 3. A *matching with flexible quotas* (FQ-matching) is an assignment $\mu \subseteq V$ that satisfies quotas-up-to-ties and nondiscrimination and ensures that each applicant $a \in A$ is assigned to at most one program (i.e., $\mu(a)$ has at most one element).

Finally, the presence of ties and flexible quotas requires updating the notion of stability.

Definition 4. An FQ-matching μ is stable if it has no blocking pairs; that is, there are no pairs $(c, a) \notin \mu$ such that

- $\mu(c)$ has less than q_c applicants who are *strictly* preferred to a, and
- a prefers c over his current assignment, $\mu(a)$.

To compute a stable FQ-matching, we propose the following procedure. As in the algorithm by Baïou and Balinski (2004), we start by recursively removing all strictly dominated nodes, ensuring that students tied in the last place of a program are kept. Notice that the definitions of a-dominance and c-dominance directly extend to the case with ties and flexible quotas, as they involve strict dominance. Then, starting from the corresponding domination-free subgraph G^*, we can obtain a stable FQ-matching by assigning each student to a program. In particular, the two extreme allocations can be obtained by greedily assigning each student to his top preference (student-optimal) in G^* or each program to its most desired q_c students, including those tied in the last place (university-optimal).[15] Applying this procedure to the example in Figure 3(a), we obtain the reduced graph G^* presented in Figure 3(b) and the extreme assignments: μ^*_A (light gray) and μ^*_C (gray). This example shows that the inclusion of a single tie may considerably change the outcome (see Appendix 5 of the electronic companion for a detailed discussion of this example and how it compares to the case with no ties).

In Appendix A, we formally describe the algorithm that was finally implemented (in 2016), which is a faster version, as it only computes the student-optimal FQ-matching. To accomplish this, the algorithm recursively

Figure 3. Unified Process

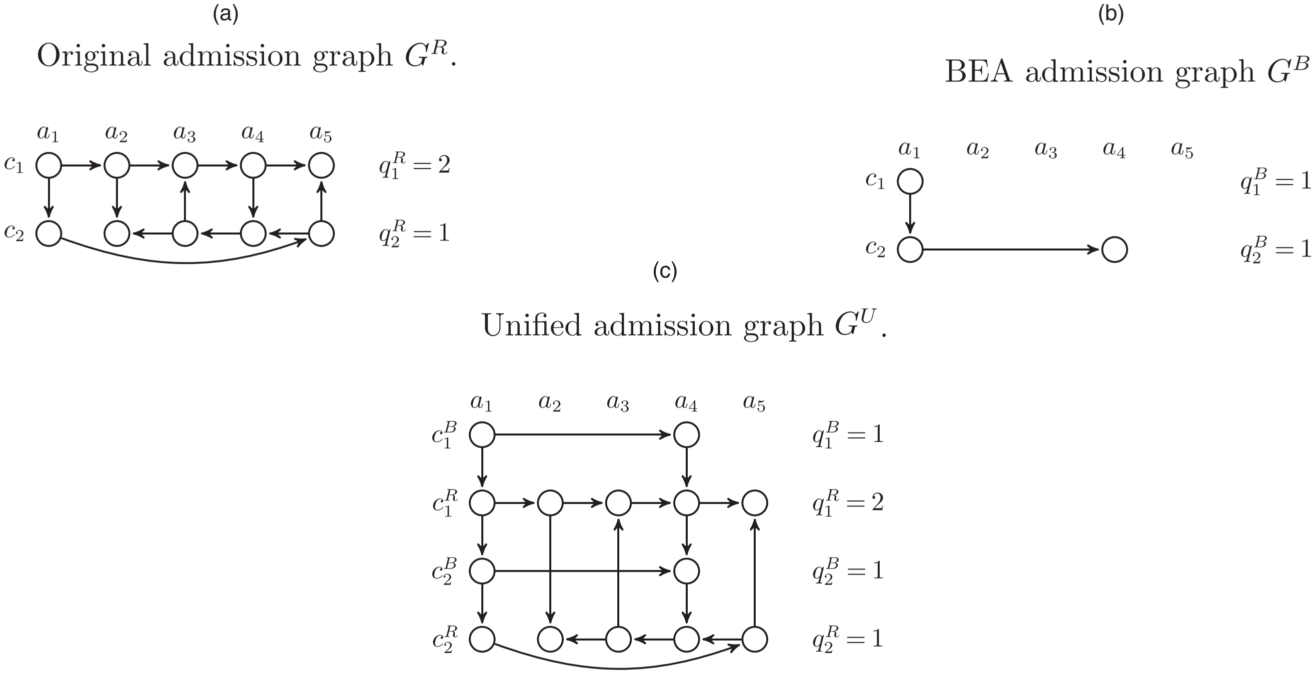

eliminates all nodes that are *a*-dominated, and it later assigns each student to his top choice in the resulting subgraph. This is a good alternative to other algorithms (e.g., DA or ascending score limits; see Biró and Kiselgof (2015)) to compute the student-optimal assignment when quotas-up-to-ties and nondiscrimination are required.

In Appendix 1 of the electronic companion, we show that the aforementioned procedure results in a stable FQ-matching, and in Appendix 2 of the electronic companion, we show that the two extreme stable FQ-matchings are optimal, but they lack two important properties: monotonicity and strategy-proofness (SP).[16] The lack of strategy-proofness can be troublesome because it may induce agents to misreport their preferences strategically, giving an unfair advantage to more sophisticated students. However, we conjecture and argue that the mechanism is *strategy-proof in the large* (SP-L), which means that students find it approximately optimal to submit their true preferences in a large market for any full-support independent and identical distribution of students' reports (see Azevedo and Budish (2018)). In fact, as Azevedo and Budish (2018) argue, the relevant distinction for practice in a large market is whether a mechanism is "SP-L versus not SP-L" and not "SP versus not SP" because students in a large market do not know what the realized reports of every other student are, so imposing an optimality of truthful reporting against every report realization (as in SP) is too strong. Thus, the lack of strategy-proofness is not a problem in our setting. Finally, in Appendix 3 of the electronic companion, we describe how our definition of FQ-matching relates to other notions of stability, such as weak, strong, super, and L-stable. Notice that if there are no ties in the preferences of programs, nondiscrimination holds trivially, whereas quotas-up-to-ties reduces to $|\mu(c)| \le q_c$; hence, FQ-matching coincides with the standard notion of stable matching.

Overall, the implementation of a stable FQ-matching involves a trade-off between potentially exceeding capacities and obtaining a better allocation for students in the Pareto sense. If universities do not want to arbitrarily discriminate students, and their marginal cost of increasing their capacity is low enough, allowing for ties and flexible quotas can be a sensible policy because it translates to a Pareto improvement for students,[17] and it eliminates any fairness concerns that can arise as a result of tie-breaking rules. We discuss this in more detail in Section 5.3.1.

4.2. Unifying Admission Tracks

Using the model described in the previous section, we can directly include the affirmative action and solve both admission tracks (regular and BEA) simultaneously.

To accomplish this, we consider a unified admission instance $\Gamma^U = (G^U, q^U)$, where each program $c \in C$ is split into two virtual programs, c^R and c^B, with q_c^R and q_c^B vacancies, that represent the regular and the BEA processes, respectively. The preferences of students that are not shortlisted for the scholarship remain unchanged. By contrast, each program in the preference list of a BEA student is also divided into the two virtual programs, giving a higher position in the preference list to the regular process; that is, for any two programs c_1, c_2 such that $c_1 >_a c_2$, the new preference order is $c_1^R >_a c_1^B >_a c_2^R >_a c_2^B$. We decided to use this order because DEMRE wanted to prioritize BEA students. Then, by applying to the regular seats first, BEA students with good scores can be admitted in regular seats, reducing the competition for reserved seats and therefore weakly increasing the total number of BEA students admitted in the system. This idea is formalized in Dur et al. (2013) and was recently extended to more reserve groups in Dur et al. (2020). In Appendix 1 of the electronic companion, we show that every student is weakly better off compared with the sequential solution. We illustrate this in Example 1.

Example 1. Consider the admission graph in Figure 3(a), and suppose that students a_1 and a_4 are shortlisted for the scholarship. In the sequential case, the regular process is run first considering the admission graph G^R in Figure 3(a) and quotas $q_c = q_c^R$, resulting in the allocation $\mu(\Gamma^R) = \{(c_1, a_4), (c_1, a_5), (c_2, a_2)\}$.[18] Then, the BEA instance $\Gamma^B = (G^B, q^B)$ is built considering only the shortlisted students and their preferences, where they were wait-listed in the regular process. Figure 3(b) illustrates the corresponding graph G^B. The resulting allocation for the BEA process is $\mu(\Gamma^B) = \{(c_1, a_1), (c_2, a_4)\}$, and therefore student a_4 is assigned to c_1 in the regular process and to c_2 in the BEA process, whereas a_3 remains unassigned. Independent of which option is taken by a_4, a seat that could have been otherwise assigned to a_3 will be lost.

The unified graph of this problem is shown in Figure 3(c). In this case we observe that there is a unique FQ-matching given by $\mu(\Gamma^U) = \{(c_1^R, a_3), (c_1^R, a_5), (c_1^B, a_1), (c_2^R, a_2), (c_2^B, a_4)\}$—that is, all applicants are assigned and no seats are lost. More important, every student is indifferent or better off compared with the sequential assignment.

Table 1. General Description: Programs

	2014	2015	2016
Universities	33	33	33
Programs	1,419	1,423	1,436
Regular seats	110,380	105,516	105,513
Reserved seats	4,394	4,422	4,295

5. Implementation

In this section we report the results of the implementation of this project. We start by providing a general description of the Chilean college admissions problem. Then, we describe the results of our first goal, which was to find the algorithm that has been used in Chile to perform the allocation. Finally, we close this section with the results of unifying the admission tracks and other additional side effects of this project.

5.1. General Description

In Table 1 we present general descriptives on the programs that are part of the centralized admission system. We observe that between 2014 and 2016, the number of universities did not change, although the number of programs slightly increased. However, we see that the number of seats available decreased over the years.

To describe the other side of the market, in Figure 4 we present the number of regular and BEA students at each stage of the admissions process (more details are available in Table 5 in Appendix 8 of the electronic companion). A *participant* is a student who registered to participate in the standardized national exam and who has at least one valid score.[19] Once the results of the national exam are published, participants have five days to submit their applications to the centralized clearinghouse. We refer to the students who

Figure 4. Evolution of Students Through Admission Process

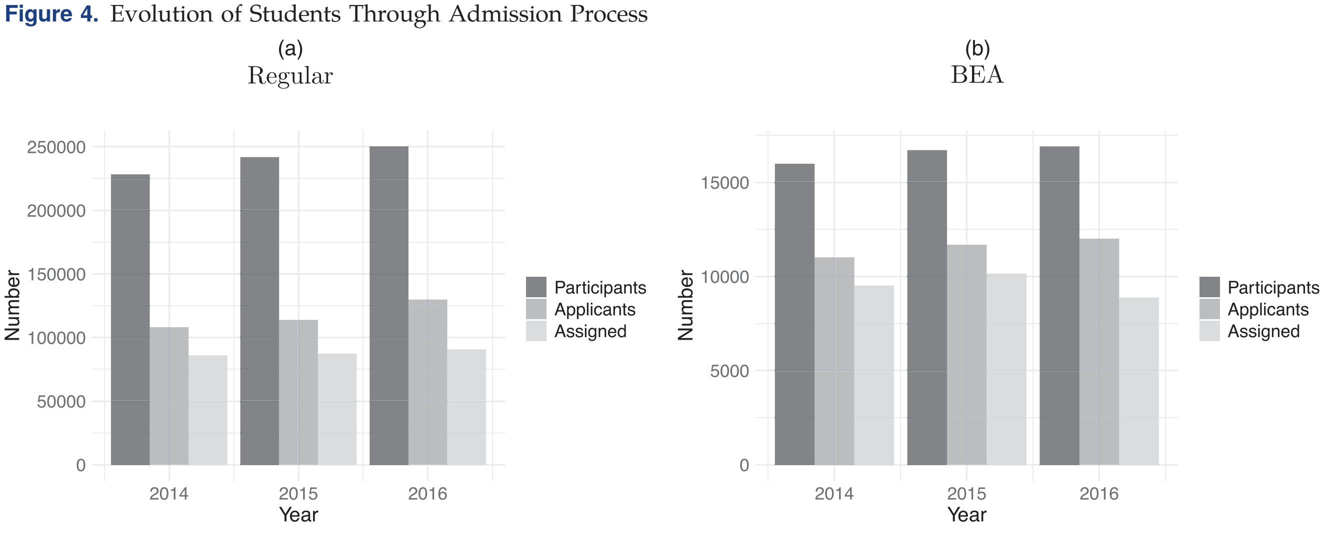

apply to at least one program that is part of the centralized system as *applicants*. Finally, we refer to students who were admitted by a program that is part of the centralized system as *assigned*.

First, we observe that the total number of participants has increased over the years, reaching a total of 267,231 participants in 2016. Second, comparing the number of participants with the number of applicants, we observe that close to half of the students who registered for the national exam applied to programs that are part of the system. The main reason for this is that CRUCH sets a minimum threshold of 450 points[20] for students to be eligible by any program that is part of the system, and because tests are standardized to have a mean of 500, roughly half of the students will not satisfy this admissibility condition.

Also related to the application process, in Figure 5 we show the distribution of applications per student for each year. The median number of applications is 4, and the share of each number of applications stays roughly constant across years. As students are restricted to submit a list with no more than 10 programs, we observe that between 5% and 10% of applicants submit a full list of 10 applications. Notice that some universities further restrict the number of programs to which a student can apply, as well as the position that an application can take in the applicant's list.[21] Theoretically, any restriction on the length of the application list will break strategy-proofness. Nevertheless, whether these constraints are binding or not in practice, and what the strategic implications are for students, are questions for future research.

Regarding the assignment, we observe that although the number of admitted students has increased, the system has become more competitive over the years as more students participate and slightly less vacancies are offered each year. In fact, the overall fraction of applicants who are assigned to some program is close to 80% in 2014, and it goes down to 78% in 2015 and 70% in 2016 (see Table B.1 in Appendix B for details). On the other hand, Figure 6, (a) and (b), shows the distribution of the preference of assignment for regular and BEA students, respectively. We see that close to 50% of students get assigned to their first reported preference and close to 90% get assigned to one of their first three preferences. Although both regular and BEA students exhibit the same pattern of assignment, notice that the latter get assigned consistently more to their first preference compared with regular students.

Finally, in Table B.2 in Appendix B we present detailed sociodemographic characteristics of the students who are assigned. The fact that about 23% of admitted students graduated from a private school is striking, considering that they represent only 12% of the total number of participants in the admission process. Similarly, students from the highest income group only represent 9% of the total number of participants, but they account for 18% of admitted students. These numbers shed some light on the huge inequalities in opportunities that characterize the Chilean college admissions process.

The point of having reserved seats is to alleviate these inequalities and favor underrepresented groups. In Figure 7 we compare sociodemographic characteristics of students who were assigned in the admissions process of 2016, separating regular applicants from BEA applicants. We observe that the fraction of female students is higher in the BEA group. In addition, we observe that the fraction of students with low income levels and from public schools is also higher in the BEA group. These results suggest that affirmative action has a positive effect at providing more opportunities to these underrepresented groups.

Figure 5. Distribution of Applications per Student

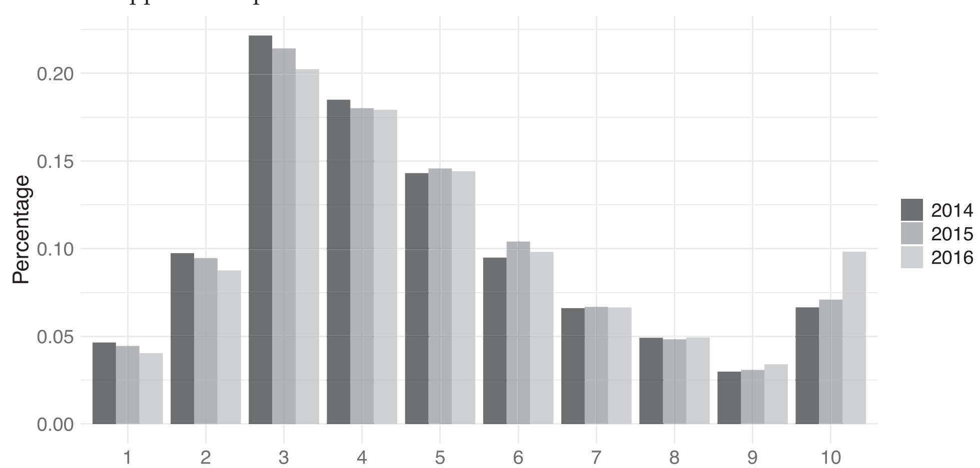

Figure 6. Preference of Assignment

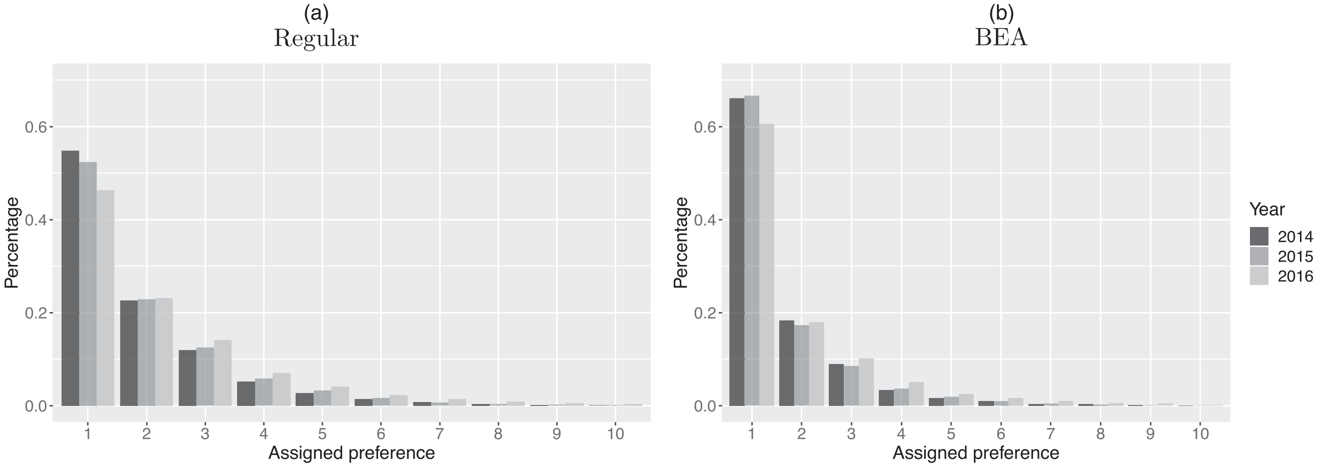

5.2. Identifying the Current Mechanism

Our first goal was to identify which mechanism has been used to solve the Chilean college admissions problem. After implementing the algorithm described in Section 4.1 and including all the constraints that are part of the system, we solved the admission instances from 2012 to 2014, comparing the FQ-student-optimal and FQ-university-optimal allocations with the official results obtained using DEMRE's black box. From these comparisons, the rules of the system, and evidence provided by DEMRE, we conclude that the allocation used is equivalent to the university-optimal FQ-matching, as the results are exactly the same for all the years considered.

Given that our algorithm returns the student-optimal and the university-optimal FQ-matchings for each instance, we can easily compare these two extreme allocations. Indeed, we find that the number of differences between these two allocations has been at most 10 (pairs student/program) since 2012. This suggests that the size of the core of stable assignments in the Chilean case is rather small, supporting the theoretical results in Roth and Peranson (2002) and Ashlagi et al. (2017).

Even though the number of differences is small, we proposed that DEMRE adopt a student-optimal FQ-matching because it benefits some students and, more important, because it is a message for students that the mechanism aims to give them the best possible allocation. DEMRE agreed with this view, and after a pilot version in 2014, it adopted the student-optimal version of our algorithm to perform the allocation in 2015.

5.3. Integrating Admission Tracks

Having identified the algorithm that is used to perform the allocation, our second goal was to integrate the admission tracks in order to alleviate the aforementioned inefficiencies.[22] To accomplish this, we implemented the framework described in Section 4.2 and ran it as a trial version during the admission processes of 2014 and 2015. From the results, we convinced DEMRE to adopt our unified student-optimal FQ-matching in 2016, and this allocation has been the official mechanism used since then. In this section we report the results from our simulations (2014 and 2015) and the actual impact of our implementation in 2016.

Figure 7. Characteristics of Assigned Applicants: Regular vs. BEA

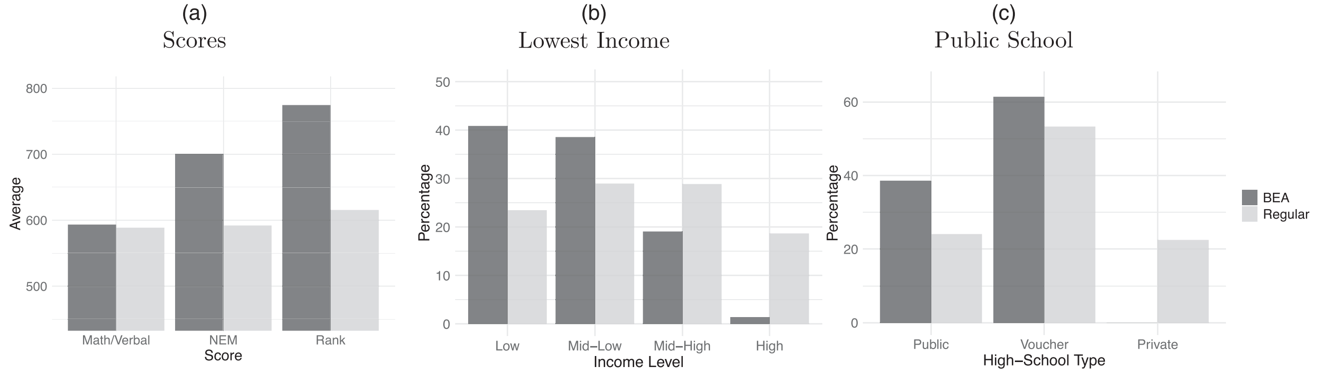

Table 2. Impact of Unified Assignment (2014–2016)

	2014	2015	2016
Double-assigned	1,100	1,180	1,127
Improvements	1,737	1,915	1,749
New Assigned	568	672	777

In Table 2 we present a summary of the results. The first row presents the number of students who would have been double-assigned under the old system. The second row presents the number of students who improved their assignment compared with the old system. Finally, the third row shows the number of students who are assigned to some program under the new system and were not assigned in the old system. We observe that the number of students who benefit from unifying the admission tracks is larger than the number of seats lost as a result of double-assignments. The reason is that a student who directly benefits releases a seat that can be used by another student, who, in turn, allows another student to take his old seat, and so on. This chain of improvements eventually ends either because there are no students wait-listed in that program or because they reach a student who was unassigned and therefore does not release another seat. Overall, we observe that about 2.5% of students who were admitted to a program in the system benefited from our implementation, and this number is relatively constant across years.

Improving the assignment of students is relevant because the probability of enrollment is increasing in the preference of assignment,[23] and most programs in the system have positive and high expected returns, which are measured in terms of the net present value of future earnings over the life cycle after graduation (see Lara et al. (2017)). Moreover, there is evidence that students who were assigned in low listed preferences have a higher probability of switching and dropping out of their programs, leading to lower on-time graduation rates (see Larroucau and Rios (2020)).

In Figure 8 we plot the benefits of unifying the admission tracks in 2016, measured in terms of preferences of assignment. For students who improve their assignment, we plot the distribution of the magnitude of their improvement based on their preference list in Figure 8(a). We observe that most students improve their allocation by getting assigned to the program listed immediately above the program to which they were previously assigned (improvement equal to 1). On the other hand, in Figure 8(b) we show the preference of assignment for those students who would not have been assigned under the old system but who end up being assigned by the new system. Most of these students benefit from the unified allocation by getting assigned to their top choice. A potential reason for these results is that an important fraction of these students apply to less than three programs.

We provide more details on the students who benefit from unifying the admission tracks in Table B.3 (see Appendix B). We first observe that most of the students who benefit from unifying the admission tracks are regular students. For instance, a total of 2,526 students benefited from the change in the algorithm in 2016 (see Table 2), among whom 2,405 are regular students and 121 are BEA students. The reason is that seats that were dropped by a BEA student with a double assignment are now used by other students, and this generates improvement chains that reach other (mostly regular) students.

In addition, in Figure 9 we compare some characteristics of those who improve (Improvements) with those who get assigned and would not under the old system (New Assigned) among regular students in the admission process of 2016. We observe that the latter group has lower scores and a larger fraction of students who come from lower income families.

Figure 8. Benefits of Unifying the Admission Tracks (2016)

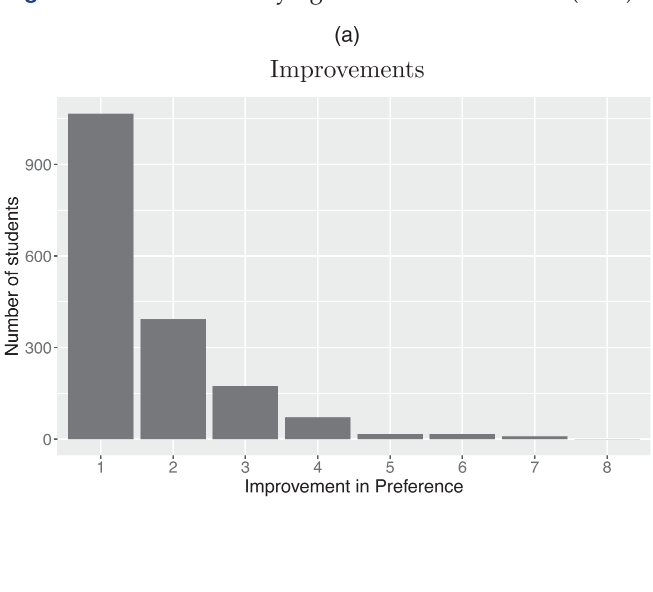

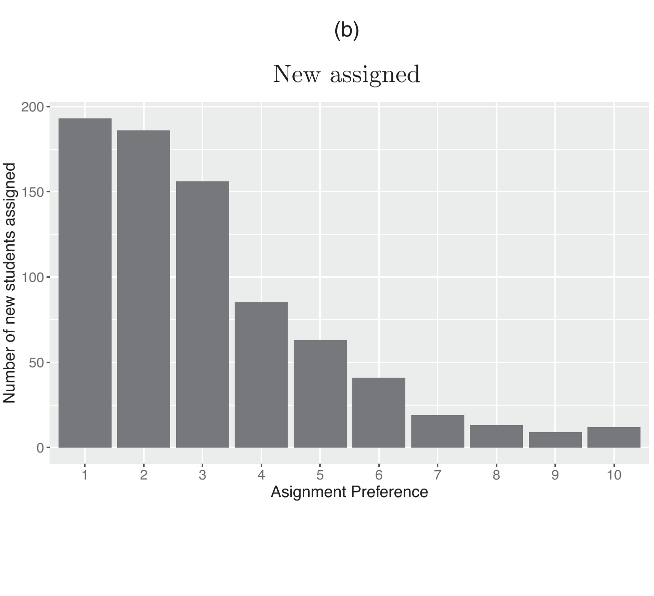

Figure 9. Characteristics of Students That Benefit: Improvements vs. New Assigned, Regular Students (2016)

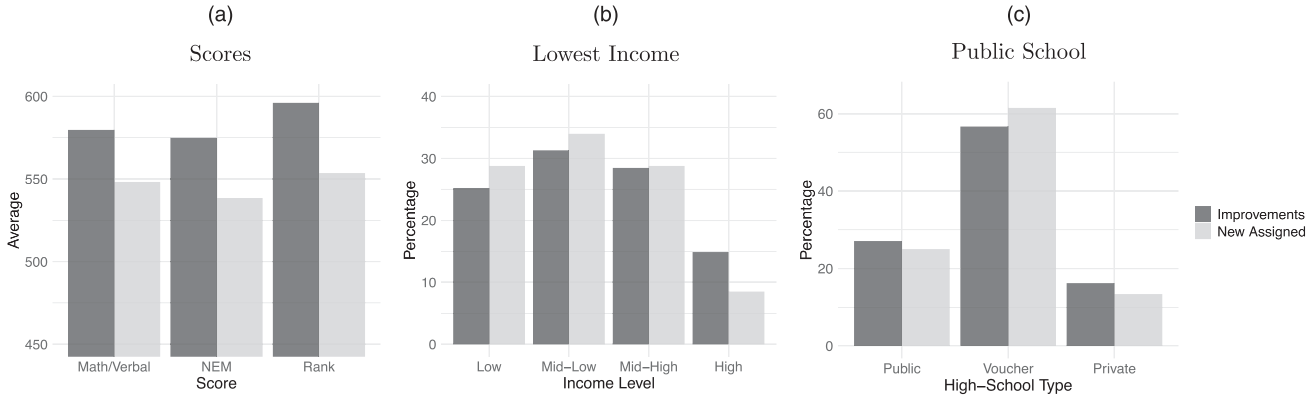

Moreover, from Table B.3 we observe that the fraction of female students is higher among New Assigned students compared with the Improvements group. The reason for this is that students who improved were also assigned under the old system, whereas those from the New Assigned group were not. Therefore, students from the Improvements group have, on average, higher scores, and these are positively correlated with family income.

The differences in terms of scores and demographics are also present if we compare these groups with the overall group of assigned students described in Table B.2. Indeed, previously assigned students have, on average, higher scores and higher family income than students who were benefited by the unified assignment. For instance, the share of assigned students coming from private high schools was about 24% for regular students, whereas it was close to 18% and 14% for students in the Improvements and New Assigned groups, respectively.

Another interesting result is that most of BEA students are assigned to regular seats, and more than half of the reserved seats remains unfilled (before the enrollment process begins); this pattern continues even after the implementation of the unified system (see Tables 1 and B.1). Indeed, we proposed DEMRE to transfer the unfilled reserved seats to students from the regular process, but they declined because some universities "were not open to this option."[24]

5.3.1. Effect of Flexible Quotas. One of the most distinctive features of the Chilean case is the use of flexible quotas. This approach belongs to the general category of *equal treatment policies* (Biró and Kiselgof 2015), where all tied students whose admission would exceed a program's capacity are either accepted or rejected, as is the case in Chile and Hungary, respectively.

An alternative approach is to break ties using a combination of randomness and administrative rules, as is the case in most school districts and in many college admissions settings, such as in Spain, Turkey, Germany, and France. A special case of these tie-breaking rules are those that exclusively depend on randomization, such as *single tie-breaking* (STB) and *multiple tie-breaking* (MTB).[25] Which approach to use is a relevant design decision, and it heavily depends on the characteristics of the problem.

We identify three dimensions that could help guide the decision-making process. First, the nature of agents' preferences plays a critical role. Indeed, if preferences are fine enough so that the number of ties is relatively small, the benefits of having flexible quotas—nondiscrimination and better allocation for students in the Pareto sense—may outweigh its costs, exceeding capacities. This would be the case in most college admissions settings, where preferences are built based on scores from exams and/or grades. By contrast, when preferences are rather coarse (e.g., in school choice settings with a limited number of priority groups), having flexible quotas may lead to large violations of initial capacities, which could make implementation unfeasible. Second, the level of heterogeneity in students' preferences is also relevant, as it prevents most of the ties from being concentrated in a small number of programs. Finally, a last element to consider is whether the allocation depends on factors that are perceived as relevant to the process. For example, waiting times for public housing, the condition of a patient for an organ transplant, or exam scores in college admissions are generally considered as fair factors, so using random tiebreakers may be considered arbitrary and discriminatory. In fact, the use of random orders to break ties in the admission process to high schools in Chile (see Correa et al. (2019)) has generated a strong debate, with opponents arguing that some measure of academic achievement should be considered instead.[26]

To illustrate the effect of having flexible quotas, we compare the official assignments to the results that would be obtained if ties were handled with other approaches. To our knowledge, this is the first paper to compare equal treatment policies with random tiebreakers, and it thus contributes to the literature that compares STB and MTB empirically (see Abdulkadiroğlu et al. (2009) and de Haan et al. (2015)) and theoretically (see Arnosti (2015) and Ashlagi et al. (2019)).

In Table 3 we report the number of extra seats required as a result of flexible quotas and how these are distributed across programs. We observe that the number of additional seats created is small (a maximum of 89 seats in 2015), which represents less than 0.1% of the total number of seats for each year. In addition, we observe that these seats are evenly spread across programs, as the maximum number of seats created by a given program is 3. Hence, we conclude that having flexible quotas does not involve a large cost for programs.

Table 3 also reports the number of students who benefit from having flexible quotas compared with an H-stable mechanism (Biró and Kiselgof 2015)—that is, one that rejects all tied students whose admission would exceed the program's capacity. The first group, Improvements, includes students who improve their assignment; the second, New Assignments, considers students who are assigned to some program in the official assignment (with flexible quotas) and would not be assigned under the alternative mechanism. Similarly, in Table 4 we summarize the benefits from having flexible quotas compared with breaking ties randomly using STB and MTB. These results are obtained from 100 simulations for each tie-breaking rule, and as before, we separate the students who benefit into two groups: Improvements and New Assignments.

As expected, all students weakly prefer their allocation under flexible quotas, and a significant number of students strictly prefer it compared with H-stability, STB, and MTB. In addition, we observe that the average number of students who benefit from flexible quotas largely exceeds the number of extra seats created, with between two and three students benefiting from each extra seat. Among these, we find that roughly 2/3 are students who improve their assignment, whereas 1/3 are students who would not be assigned if a random tiebreaker were used. All these results suggest that having flexible quotas benefits an important number of students without generating a substantial cost for programs, and we expect to find similar patterns in other college admissions systems that are similar to the Chilean case, such as those in Hungary, Turkey, and Spain. In Appendix 9 of the electronic companion, we further compare flexible quotas with other approaches that combine tie-breaking rules (STB and MTB) with increasing vacancies to match those resulting from flexible quotas, and we find that flexible quotas still outperform these benchmarks.

5.4. Additional Side Effects

In terms of running times, our implementation considerably outperforms the algorithm previously used by DEMRE. In fact, their black-box software takes up to five hours to return the final assignment, whereas our implementation solves the problem in less than two minutes on a standard laptop. This time reduction has had a significant impact because it allows for the evaluation of different policy changes in the system, such as the inclusion of new admission criteria, the impact of new instruments, and the redesign of affirmative action policies. In particular, the new algorithm was used to evaluate the effect of including the high school class rank as an admission factor through simulations, changing the conditions in which this new instrument is included (Larroucau et al. 2015). Furthermore, the efficiency gains have opened other directions for future research involving the evaluations of policies that could stress the system in the future—for instance, the impact of modifying the constraints on the size of the list or the weights on the admission factors, the inclusion of score bonuses for first-time applicants or penalties for reapplicants, and the implementation of admission quotas for underrepresented groups. The evaluation of these policies was not possible in the past because of the computational time involved.

Table 3. Impact of Flexible Quotas vs. H-Stability

	2014	2015	2016
Effects of flexible quotas			
Extra seats	61	89	67
Programs with flexible quotas	53	79	59
Maximum number of flexible quotas	2	3	3
Benefits compared with H-stability			
Improvements	124	317	212
New Assigned	68	134	99
Total	192	451	311

Table 4. Impact of Flexible Quotas vs. Random Tie-Breaking

	Benefits from flexible quotas					
	STB			MTB		
	2014	2015	2016	2014	2015	2016
Improvements	89.4	161.8	110.9	88.5	162.4	111.3
	(5.4)	(7.1)	(5.8)	(5.6)	(8.1)	(5.6)
New Assigned	32.1	63.9	52.4	31.9	63.8	52.2
	(2.0)	(2.3)	(1.6)	(2.0)	(2.3)	(1.7)
Total	121.6	225.6	163.3	120.5	226.2	163.6
	(5.0)	(6.4)	(5.9)	(4.8)	(7.7)	(5.8)

6. Conclusions

We investigate how the Chilean college admissions system works. There are two main features that make the Chilean system different from the classic college admissions problem: (i) preferences of colleges are not strict, and all students tied for the last seat of a program must be assigned, and (ii) the system considers a policy of affirmative action that is solved sequentially after the regular process. Students who benefit from affirmative action can be double-assigned, introducing a series of inefficiencies in the assignment and enrollment processes. Even though the authorities were aware of this problem, they could not solve it because they relied on a black-box software that could not be updated to incorporate the affirmative action.

To identify which mechanism was used, we developed an algorithm that finds all applicant/program pairs that can be part of a stable allocation satisfying the rules of the system (i.e., flexible quotas and nondiscrimination of tied students). We also introduced the notion of FQ-matching to account for these features, and we characterized its main properties. We showed that this mechanism leads to the optimal stable allocations satisfying flexible quotas and nondiscrimination, but it lacks monotonicity and strategyproofness. Nevertheless, we argue that this should not be a concern considering the large size of the market, which makes it practically impossible for students to benefit from manipulating their preferences to leverage the presence of ties and flexible quotas.

By comparing the results of our algorithm with historical data, we found that the algorithm that had been used in the past is a variant of the university-optimal stable assignment that satisfies flexible quotas and nondiscrimination. Even though the number of differences is small, we convinced DEMRE to switch to the student-optimal FQ-matching, which was finally adopted in 2015 after a pilot version in 2014.

Having identified the algorithm, we proposed a new method to incorporate the affirmative action by treating regular and reserved seats as different programs. The unified approach to solve the problem was adopted and implemented by DEMRE in 2016, after two years of analyzing its potential impact. The results of the implementation in 2016, as well as the pilot results in 2014 and 2015, show that about 2.5% of the total number of students that are admitted each year benefit from the unified assignment. Among the students who actually benefited in 2016, 30.8% would not have been assigned to any program under the old system, and 69.2% are students who improved compared with what they would get under the old system. The benefited students have, on average, lower scores and lower family income compared with the students that would have been assigned under the old system. Thus, the unified approach reduces both the inefficiencies of the mechanism and the inequality in the system.

In addition to its direct impact on students, the efficiency of our algorithm reduced the running time by two orders of magnitude relative to the old system, enabling to perform simulations to evaluate different policies oriented to make the admission process more fair and inclusive. Finally, our method helped to improve the transparency of the system and allowed other changes to be implemented on top of it.

Certainly, there are many directions for future work. While working on this project, we realized that many students do not apply to programs where their chances of admission are too low, even though the constraint on the length of their preference list is not binding. Hence, considering their reports as truthful would lead to serious biases in the estimation of preferences, leading to wrong evaluations of policies. We are currently working on a model of preferences that takes this fact into account (see Larroucau and Rios (2019)). Another direction that emerged from this project is focused on trying to understand why some students apply to programs where they have no chance of getting admitted, as they do not satisfy the requirements to be eligible. We are currently working on understanding why this is the case and designing changes to the application process aimed at reducing mistaken reports. Finally, another research question that arose from this project is how to modify the mechanism to better elicit the intensity of students' preferences and use this to improve students' retention and on-time graduation rates (see Larroucau and Rios (2020)).

Overall, we hope that the current results encourage the Chilean authorities to keep improving the system and that they motivate other college systems around the world to evaluate and adopt flexible quotas, as this would increase the overall efficiency of the processes and would also improve the welfare of students.

Acknowledgments

The authors thank the department editor, associate editor, and two anonymous referees for constructive comments that significantly improved this paper. The authors gratefully acknowledge the support of the DEMRE for providing essential information and data without which the present study could not have been completed. The authors also thank José Correa, Nicolás Figueroa, Eduardo Azevedo, Hanming Fang, Fuhito Kojima, Greg Macnamara, Javier Martínez de Albéniz, Rakesh Vohra, Xavier Warnes, and the committee cochairs and referees of the Doing Good with Good OR—Student Paper Competition for their constructive feedback.

Operations Research, 2021, vol. 69, no. 4, pp. 1186–1205, © 2021 INFORMS

Appendix A. Faster Algorithm for FQ-Matchings

An observation that might be exploited to improve the algorithm described in Section 4.1 is that not all dominated nodes need to be removed.

Indeed, a natural approach is to drop only those nodes that are causing violations of the properties of the assignment (e.g., stability, quotas-up-to-ties, and nondiscrimination). This is similar to the strategy used in the deferred-acceptance algorithm of Gale and Shapley (1962).

We describe the idea for the student-optimal FQ-matching μ_A. For each program c, we set up an ordered list L_c in which we will sequentially add and remove applicants. Initially, $L_c = \emptyset$. For each $a \in A$, we set a pointer t_a to its most preferred program, and if this pointer is not null, we push a into a stack S that contains the applicants with no program assigned yet.

We iterate as follows. We pop a student a from the stack S and insert it in the ordered list L_c of the most preferred program c. If $|L_c|$ exceeds q_c, we check quotas-up-to-ties and eventually remove the last group of students in L_c to ensure that this property holds by using the following procedure.

Procedure 1 (Check-Quotas-up-to-Ties)**.** Find the set T_c of applicants tied in the last position in L_c. If $|L_c| \le q_c + |T_c|$, we keep the list as it is; otherwise, each node (c, a') for $a' \in T_c$ is a-dominated, so we remove a' from L_c and update $t_{a'}$ to its next most preferred program, $c' <_{a'} c$. If such a c' exists, we push a' back into the stack S, and otherwise, we leave a' unassigned. If the tie T_c is removed, then $|L_c|$ reduces its size by $|T_c|$.

Algorithm 1 (Fast student-optimal FQ-matching)

Read instance Γ
Initialize $L_c = \emptyset$ for all c and stack $S = A$
while (S is nonempty) **do**
  $a \leftarrow \text{pop}(S)$
  $c \leftarrow t_a$
  Insert a into L_c
  if ($|L_c| > q_c$) **then**
    Check-Quotas-up-to-Ties (Procedure 1)
  end if
end while
return Assignment μ_A represented by the final lists L_c

Appendix B. Additional Results

Table B.1. General Description: Students

	Regular			BEA		
	2014	2015	2016	2014	2015	2016
Participants	228,318	241,873	250,320	15,990	16,710	16,911
Applicants	108,144	113,900	129,896	11,017	11,688	12,010
Assigned						
Regular seats	86,048	87,466	90,741	9,520	10,154	8,886
Reserved seats	—	—	—	1,325	1,404	1,345

Table B.2. General Description: Assigned

	Regular			BEA		
	2014	2015	2016	2014	2015	2016
Assigned						
Total	86,048	87,466	90,741	9,745	10,378	10,231
Gender						
Female (%)	49.5	49.5	50.2	57.6	58.4	59.5
Average scores						
Math/Verbal[a]	588	589.3	588.7	591.1	595.9	593.4
NEM[b]	586	588.4	592	696.8	696.7	700.5
Rank[c]	608.7	614.4	615.5	770.6	776.9	774.5
Income[d]						
[0, \$288] (%)	28.8	26.4	23.5	46	42.6	40.9
(\$288, \$576] (%)	26.7	27.5	29	34.6	37.1	38.6
(\$576, \$1,584] (%)	26.8	27.5	28.9	18.7	19.5	19.1
> \$1,584 (%)	17.8	18.6	18.7	0.7	0.8	1.4
High school						
Private (%)	23.4	23.4	22.5	0	0	0
Voucher[e] (%)	52.4	52.9	53.3	60.7	61	61.4
Public (%)	24.2	23.7	24.1	39.3	39	38.6

[a]Score constructed with the average math score and verbal score. For students using scores from the previous year, we considered the maximum of both averages.
[b]Score constructed with the average grade in high school.
[c]Score constructed with the relative position of the student among his orher classmates.
[d]Gross family monthly income in thousand Chilean pesos (nominal).
[e]Partially subsidized schools.

Table B.3. General Description: Impact of Unifying Admission Tracks

	Improvements						New Assigned					
	Regular			BEA			Regular			BEA		
	2014	2015	2016	2014	2015	2016	2014	2015	2016	2014	2015	2016
Assigned												
Total	1,592	1,791	1,640	145	124	109	548	647	765	20	25	12
Gender												
Female (%)	47.2	44.4	45.7	66.9	60.5	62.4	52.2	47.9	51.6	65	68	58.3
Average scores												
Math/Verbal[a]	584.2	582	579.7	579.8	573.7	584.1	560.8	551.2	548.2	564.8	561	558
NEM[b]	577.8	576.9	575	686	675.7	691.1	540	538.9	538.4	677.1	658	689.2
Rank[c]	600.2	600.6	596.1	755.8	752.9	770	554.8	556.4	553.5	753.2	734	764.1
Income[d]												
[$0, $288] (%)	28.1	27.4	25.2	48.3	41.9	48.6	30.1	32.3	28.8	60	36	75
($288, $576] (%)	30.2	29	31.3	29.7	39.5	37.6	34.5	32.9	34	25	32	8.3
($576, $1,584] (%)	27.4	29.3	28.5	20	18.5	13.8	26.1	24.3	28.8	15	32%	16.7
> $1,584 (%)	14.3	14.3	14.9	2.1	0	0	9.3	10.5	8.5	0	0	0
High school												
Private (%)	18.4	18.6	16.2	0	0	0	13.8	14.5	13.4	0	0	0
Voucher[e] (%)	56.7	57.3	56.7	63.4	61.3	65.1	59.5	61.4	61.5	60	72	58.3
Public (%)	24.9	24.2	27.1	36.6	38.7	34.9	26.7	24.1	25	40	28	41.7

[a]Score constructed with the average math score and verbal score. For students using scores from the previous year, we considered the maximum of both averages.
[b]Score constructed with the average grade in high school.
[c]Score constructed with the relative position of the student among his or her classmates.
[d]Gross family monthly income in thousand Chilean pesos (nominal).
[e]Partially subsidized schools.

Appendix C. Tie-Breaking

As discussed in Section 5.3.1, flexible quotas are an alternative to tie-breaking rules such as STB and MTB. However, the comparison between these alternatives is not completely fair, as flexible quotas may end up allocating a higher number of seats because of potential ties. One approach to make the comparison more fair is to increase the number of vacancies used as an input for STB or MTB, emulating the extra vacancies that result from flexible quotas. We consider three different approaches:

1. *By program*: We consider as vacancies the maximum between the original number of seats and the number of students assigned under flexible quotas (for each program).
2. *Uniform*: We consider as vacancies the original number of seats plus the extra seats resulting from flexible quotas uniformly distributed across programs. The redistribution of extra seats can be done by sampling with replacement a number of programs equal to the number of extra seats and then increasing the vacancies of each sampled program by one.
3. *Bootstrap*: Because the number of extra vacancies that result from flexible quotas is a random variable that depends on the distribution of applications and the distribution of scores, another approach is to estimate the distribution of extra seats that result from flexible quotas and then run STB or MTB sampling on both the random tiebreakers and the number of extra seats for each program. We implement this approach in two steps:

   a. We estimate the distribution of extra seats for each program and for each admission track using bootstrap,

Table C.1. Impact of Flexible Quotas: Uniform Redistribution

	STB			MTB		
	2014	2015	2016	2014	2015	2016
Improvements	77.0	140.2	91.3	77.3	140.9	91.1
	(6.5)	(9.3)	(6.3)	(6.8)	(8.9)	(6.2)
New Assigned	26.9	54.1	44.1	26.7	53.6	44.0
	(2.4)	(3.3)	(2.1)	(2.6)	(3.3)	(2.2)
Total (positive)	103.9	194.4	135.3	104.0	194.5	135.0
	(6.2)	(10.0)	(6.3)	(6.5)	(9.8)	(6.4)
Worsened	34.0	51.8	49.4	34.1	51.9	49.7
	(5.8)	(7.4)	(7.6)	(6.1)	(8.0)	(7.4)
Not Assigned	14.1	29.8	32.0	14.1	29.7	32.0
	(3.0)	(3.8)	(2.1)	(3.0)	(3.7)	(2.1)
Total (negative)	48.1	81.5	81.5	48.3	81.6	81.6
	(7.3)	(10.5)	(6.9)	(7.5)	(11.0)	(6.7)

Notes. Comparison of the official assignment (using flexible quotas) with the allocation obtained from 100 simulations using STB and 100 simulations using MTB after redistributing extra vacancies uniformly across programs. We report the average and standard deviation (in parentheses) of the number of students that benefit, separated into four groups: (1) improvements, for those students who improve their assignment; (2) new assignments, for those who are assigned under the official assignment but who result in being unassigned using STB/MTB; (3) worsened, for those students who worsen their assignment (but remain assigned), and (4) not assigned, for those who are not assigned under the official assignment but result in being assigned using STB/MTB. Finally, we also report the total number of students who benefit (Total (positive) = Improvements + New Assigned) and who are worse off (Total (negative) = Worsened + Not Assigned).

Table C.2. Impact of Flexible Quotas: Bootstrap

	STB			MTB		
	2014	2015	2016	2014	2015	2016
Extra seats	55.3	67.4	70.7	55.3	67.5	70.7
	(7.4)	(9.9)	(8.2)	(7.4)	(9.9)	(8.2)
Programs with extra seats	51.3	61.8	65.9	51.3	61.8	65.9
	(6.7)	(8.1)	(7.6)	(6.7)	(8.1)	(7.6)
Maximum number of extra seats	2.3	2.4	2.3	2.3	2.4	2.3
	(0.5)	(0.7)	(0.6)	(0.5)	(0.7)	(0.6)
Improvements	69.1	120.1	80.7	68.2	121.1	81.8
	(6.3)	(13.0)	(9.1)	(6.9)	(12.0)	(8.9)
New Assigned	24.5	47.3	37.2	24.9	46.9	37.8
	(2.7)	(3.8)	(3.0)	(2.6)	(4.1)	(2.7)
Total (positive)	93.6	167.4	117.9	93.1	168.0	119.5
	(7.0)	(15.0)	(10.7)	(6.9)	(13.8)	(10.4)
Worsened	68.0	64.6	85.7	67.7	65.0	85.4
	(13.4)	(13.4)	(14.4)	(12.5)	(13.1)	(15.3)
Not Assigned	20.9	26.6	37.4	21.0	26.4	38.0
	(5.0)	(6.1)	(6.0)	(4.8)	(5.8)	(5.8)
Total (negative)	88.9	91.3	123.1	88.6	91.5	123.4
	(16.0)	(17.1)	(17.0)	(15.1)	(16.7)	(17.7)

Notes. Comparison of the official assignment (using flexible quotas) with the allocation obtained from 100 simulations using STB and 100 simulations using MTB after adding extra vacancies obtained from the bootstrap procedure. We report average and standard deviation (in parentheses) of the number of students that benefit, separated into four groups: (1) improvements, for those students who improve their assignment; (2) new assignments, for those who are assigned under the official assignment but who result in being unassigned using STB/MTB; (3) worsened, for those students who worsen their assignment (but remain assigned); and (4) not assigned, for those who are not assigned under the official assignment but result in being assigned using STB/MTB. Finally, we also report the total number of students who benefit (Total (positive) = Improvements + New Assigned) and who are worse off (Total (negative) = Worsened + Not Assigned).

adapting the procedures described in Agarwal and Somaini (2018) and Larroucau and Rios (2019). In particular, for each bootstrap simulation, we sample with replacement a number of students equal to the number of students, along with their application and scores, and we add a random noise to each of their application scores distributed uniformly in the set $\{-0.1, -0.09, \ldots, 0, \ldots, 0.09, 0.1\}$. We use these random noises to prevent having an excessive number of ties; the bootstrap procedure is prone to this, as it uses students having the same application and scores if no noise is added. Moreover, we choose to sample noises from that subset of values because it leads to a similar expected number of extra vacancies to what we observe under flexible quotas. Then, for each bootstrap simulation, and considering the sampled students and their new application scores, we solve the allocation considering flexible quotas. As a result, for each program and each admission track, we obtain a number of extra seats that are allocated as a result of the flexible quotas, and we store these values.

b. We simulate the allocations obtained with STB or MTB. To accomplish this, at the beginning of each simulation we sample both the random tiebreakers and the number of extra seats for each program and each admission track (sampling a random simulation from the bootstrap), and then we run DA using either STB or MTB.

Our first result is that increasing the number of vacancies by program (option a) leads to exactly the same allocation as the one obtained with flexible quotas. This claim is formalized in Proposition C.1.

Proposition C.1. *For a fixed algorithm (student-optimal or university-optimal) and a given instance Γ with vector of vacancies $q = \{q_c\}_{c \in C}$, let $\tilde{q} = \{\tilde{q}_c\}$ be the number of seats allocated under flexible quotas, and let $\tilde{\mu}(q)$ be the flexible quotas' allocation considering as input the vector of vacancies q. In addition, let $\bar{\mu}(q)$ be the allocation obtained considering tie-breaking (STB or MTB) and a vector of vacancies q. Then,*

$$\tilde{\mu}(q) = \bar{\mu}(\tilde{q}).$$

Proof. We know that DA has a cutoff structure—that is, the allocation can be obtained from a vector of cutoffs (or score limits) and assigning each student to his or her highest preference among programs for which their application score is greater than or equal to the cutoff (see *Biró and Kiselgof (2015)*). Let $\{\tilde{s}_c\}$ be the vector of cutoffs that results from using flexible quotas, and let s_{ac} be the application score of student a in program c. Because scores are discrete and flexible quotas admit all students tied in last place, we know that there is a strictly positive difference between $\tilde{s}_c$ and the application score of the student on the wait-list with the highest score (if no students are on the wait-list, we simply assume that this difference is $+\infty$). Let δ be the smallest difference between a program's cutoff and the highest score on the wait-list:

$$\delta = \min_{c\in C}\{\tilde{s}_c - s_{ac} \;:\; a \in A,\; \tilde{\mu}_a <_a c\}.$$

Next, let $\bar{s}_c = \tilde{s}_c - \delta/2$ be a new set of cutoffs. Given the cutoff structure of the flexible quotas algorithm, we know that the allocations obtained given $\{\tilde{s}_c\}_{c\in C}$ and $\{\bar{s}_c\}_{c\in C}$ are exactly the same, because there are no wait-listed students whose score is in $[\bar{s}_c, \tilde{s}_c)$.

Next, suppose that we implement a tie-breaking rule where ties are broken by adding a small enough noise to the application scores. More formally, suppose that the tie-breaking rule considers as application scores $\bar{s}_{ac} = s_{ac} - \epsilon_{ac}$, where $\epsilon_{ac} \sim U[0, \delta/2]$ (this represents MTB; for STB, we would sample $\epsilon_a \sim U[0, \delta/2]$, but for the purpose of the proof, both are equivalent). By the cutoff structure of DA, we know that the allocation obtained with scores $\{\bar{s}_{ac}\}_{a\in A,c\in C}$ and cutoffs $\{\bar{s}_c\}_{c\in C}$ is equivalent to the allocation obtained under MTB considering as input vacancies $\tilde{q}$ and scores $\{\bar{s}_{ac}\}_{a\in A,c\in C}$. Moreover, this allocation is the same as the one obtained using the original set of scores $\{s_{ac}\}_{a\in A,c\in C}$ by construction. Hence, the allocation obtained from MTB is equivalent to the allocation obtained considering scores $\{s_{ac}\}_{a\in A,c\in C}$ and cutoffs $\{\bar{s}_c\}_{c\in C}$, which is equivalent to the allocation resulting from flexible quotas. □

On the other hand, if option b is implemented, it is possible that the allocation obtained from STB or MTB is different from the one obtained from flexible quotas. In Table C.1 we compare the results of 100 simulations that use as input the updated vacancies after distributing the extra seats uniformly across programs.

First, we observe that the number of students who benefit from having flexible quotas is greater than the number of students who are worse off compared with the allocations with tie-breaking and extra seats. In addition, we observe that the total number of students who benefit is smaller than reported in Table 4. This result was expected: in this case, the allocations with tie-breaking consider more vacancies as input, and thus it will obtain a better allocation compared with the case where vacancies are kept fixed.

The results comparing flexible quotas with STB or MTB, where the latter two include the simulated extra seats obtained from the bootstrap procedure, are reported in Table C.2. The first part of the table reports the average and standard deviation for the number of extra seats, the number of programs with extra seats, and the maximum number of extra seats that result from sampling the number of extra seats. The second part of the table reports the mean and standard deviation of the number of students that are benefitted and harmed by using flexible quotas instead of a tie-breaking rule that includes random extra vacancies.

We observe that using a random tiebreaker with a random number of extra seats leads to similar results than considering the original number of seats and flexible quotas, as the average total number of students who benefit is similar to the average total number of students who are worse off.

Endnotes

[1] In Chile, students apply directly to a major in a given university, such as medicine in the University of Chile. We refer to each program as a major-university pair.

[2] In addition to what we describe in this paper, each university has special admission programs for athletes and racial minorities, among others. In addition, there are other centralized admission tracks that were added to the system in 2017 that we do not address in this paper for simplicity.

[3] BEA students should be indifferent between regular and reserved seats because they obtain the scholarship regardless of how they were admitted, and there are no differences between these types of seats.

[4] We do not include military and police academies (7).

[5] IP and CFT programs tend to be shorter.

[6] Some programs such as music, arts and acting may require additional aptitude tests.

[7] CRUCH is the institution that gathers these universities and is responsible to drive the admission process, whereas DEMRE is the organization in charge of applying the admission tests and carrying out the assignment of students to programs.

[8] Many of these institutions run two admission processes: the first, and most significant in terms of vacancies, is simultaneous to the centralized process; the second takes place in late July/early August and grants admission for the second semester.

[9] This respects some basic criteria defined by CRUCH.

[10] This was directly translated from the document "Normas, Inscripción y Aspectos Importantes del Proceso de Admisión, 2013" (CRUCH 2013, p. 8).

[11] The legal origins of this requirement can be traced back to 1967, when a group of eight universities decided to change the exam used as part of the admission process (switching from the "Bachillerato" to the "Prueba de Aptitud Académica" (PAA) exam, which is the exam that preceded the PSU). As part of the transition, the University of Chile was in charge of implementing the system to process the applications, and as part of the basic requirements, they included that "each student should be evaluated under the same conditions and criteria." Since then, this rule has been interpreted as equal treatment of equals, which implies that the system cannot accept one student and reject another one having the same scores, and over the last 20 years, this requirement has been included explicitly as part of the rules of the regular and BEA admission processes.

[12] Students get a full refund of the enrollment fees if they decide to decline their enrollment in the first stage to enroll in a new program.

[13] To be more precise, the system differentiates between normal and overbooking seats. During the enrollment process, if an admitted student does not enroll, then wait-listed students are offered admission only up to the normal vacancies. Hence, only those students who were admitted before the enrollment process can use overbooking seats.

[14] In Appendix 6 of the electronic companion, we compare our results using enrollment data and show that the inefficiencies introduced by the double assignment are not addressed in the enrollment process.

[15] In Appendix 4 of the electronic companion, we show that the complexity of this procedure is $O(|V|^2)$.

[16] Appendices labeled with Arabic numerals can be found in the electronic companion; those labeled with capital letters can be found at the end of the paper.

[17] As long as we consider applications as fixed, allowing for ties and flexible quotas will weakly increase the number of seats per program, resulting in a Pareto improvement for students.

[18] In this example, the student-optimal and the university-optimal algorithms return the same allocation.

[19] This group also includes students who participated in the national exam in the previous year.

[20] This is the average of the math and language scores.

[21] For example, the University of Chile requires applicants to apply to at most four of its programs, and these applications must be listed within the top four positions in the applicant's list.

[22] In Appendix 6 of the electronic companion, we show that these inefficiencies are not eliminated by the enrollment process.

[23] In Figure 18 in Appendix 6 of the electronic companion, we show that the share of students who enroll after being assigned in one of their 10 listed preferences is decreasing in the number of assigned preference.

[24] As an anonymous referee pointed out, this is not necessarily a source of inefficiency. It could be the case that some universities are just willing to enroll more BEA students if they get more applicants than anticipated, but they are not willing to fill that capacity with extra regular students.

[25] Under STB, every program uses the same random ordering to break ties, whereas under MTB, each program uses its own random order.

[26] See Plaza Pública Cadem, Survey No. 262 (January 21, 2019; in Spanish), https://www.cadem.cl/wp-content/uploads/2019/01/Track-PP-262-Enero-S3-VF.pdf.

References

Abdulkadiroğlu A (2007) Controlled school choice. Working paper, University of Montreal, Montreal.

Abdulkadiroğlu A, Pathak PA, Roth AE (2005) The New York City high school match. *Amer. Econom. Rev.* 95(2):364–367.

Abdulkadiroğlu A, Pathak PA, Roth AE (2009) Strategy-proofness vs. efficiency in matching with indifferences: Redesigning the New York City high school match. *Amer. Econom. Rev.* 99(5): 1954–1978.

Abdulkadiroglu A, Pathak PA, Roth AE, Sömnez T (2005) The Boston Public School Match. *Amer. Econom. Rev.* 95(2):368–371.

Agarwal N, Somaini P (2018) Demand analysis using strategic reports: An application to a school choice mechanism. *Econometrica* 86(2):391–444.

Arnosti N (2015) Short lists in centralized clearinghouses. *Proc. 16th ACM Conf. Econom. Comput.* (ACM, New York), 751.

Ashlagi I, Kanoria Y, Leshno JD (2017) Unbalanced random matching markets: The stark effect of competition. *J. Political Econom.* 125(1):69–98.

Ashlagi I, Nikzad A, Romm A (2019) Assigning more students to their top choices: A comparison of tie-breaking rules. *Games Econom. Behav.* 115:167–187.

Azevedo EM, Budish E (2018) Strategy-proofness in the large. *Rev. Econom. Stud.* 86(1):81–116.

Baïou M, Balinski M (2004) Student admissions and faculty recruitment. *Theoret. Comput. Sci.* 322(2):245–265.

Baswana S, Chakrabarti PP, Chandran S, Kanoria Y, Patange U (2019) Centralized admissions for engineering colleges in centralized admissions for engineering colleges in India. *INFORMS J. Appl. Analytics* 49(5):338–354.

Biró P, Kiselgof S (2015) College admissions with stable score-limits. *Central Eur. J. Oper. Res.* 23(4):727–741.

Consejo de Rectores de las Universidades Chilenas (CRUCH) (2013) Proceso de admisión 2013: Normas, inscripción y aspectos importantes del proceso de admisión, 2013. Report 2, El Consejo de Rectores de las Universidades Chilenas, Santiago, Chile.

Correa J, Epstein R, Escobar J, Rios I, Bahamondes B, Bonet C, Epstein N, et al. (2019) School choice in Chile. *Proc. 2019 ACM Conf. Econom. Comput.* (ACM, New York), 325–343.

de Haan M, Gautier PA, Oosterbeek H, van der Klaauw B (2015) The performance of school assignment mechanisms in practice. Report IZA DP 9118, Institute of Labor Economics, Bonn, Germany.

Dur U, Pathak PA, Sönmez T (2020) Explicit vs. statistical targeting in affirmative action: Theory and evidence from Chicago's exam schools. *J. Econom. Theory* 187(May):Article 104996.

Dur U, Kominers SD, Pathak PA, Sönmez T (2013) Priorities vs. precedence in school choice: Theory and evidence from Boston. Working paper, Massachusetts Institute of Technology, Cambridge.

Dur U, Duke S, Parag K, Pathak PA, Sönmez T (2018) Reserve design: Unintended consequences and the demise of Boston's walk zones. *J. Political Econom.* 126(6):2457–2479.

Dworczak P (2016) Deferred acceptance with compensation chains. *Proc. 2016 ACM Conf. Econom. Comput.* (ACM, New York), 65–66.

Echenique F, Yenmez MB (2015) How to control controlled school choice. *Amer. Econom. Rev.* 105(8):2679–2694.

Ehlers L, Hafalir IE, Yenmez MB, Yildirim MA (2014) School choice with controlled choice constraints: Hard bounds vs. soft bounds. *J. Econom. Theory* 153(September):648–683.

Ergin H, Sömnez T (2006) Games of school choice under the Boston mechanism. *J. Public Econom.* 90(1–2):215–237.

Gale D, Shapley LS (1962) College admissions and the stability of marriage. *Amer. Math. Monthly* 69(1):9–15.

Hafalir IE, Yenmez MB, Yildirim MA (2013) Effective affirmative action in school choice. *Theoret. Econom.* 8(2):325–363.

Kamada Y, Kojima F (2015) Efficient matching under distributional constraints: Theory and applications. *Amer. Econom. Rev.* 105(1): 67–99.

Kamiyama N (2017) Strategic issues in college admissions with score-limits. *Oper. Res. Lett.* 45(2):105–108.

Kojima F (2012) School choice: Impossibilities for affirmative action. *Games Econom. Behav.* 75(2):685–693.

Lara B, Meller P, Valdés G (2017) Life-cycle valuation of different university majors: Case study of Chile. *Interciencia* 42(6): 380–387.

Larroucau T, Rios I (2019) Do short-list students report truthfully? Strategic behavior in the Chilean college admissions problem. Working paper, University of Pennsylvania, Philadelphia.

Larroucau T, Rios I (2020) Dynamic college admissions and the determinants of students' college retention. Working paper, University of Pennsylvania, Philadelphia.

Larroucau T, Mizala A, Rios I (2015) The effect of including high school grade rankings in the admission process for Chilean universities. *Pensamiento. Ed.* 52(1):95–118.

Roth AE (2002) The economist as engineer: Game theory, experimentation, and computation as tools for design economics. *Econometrica* 70(4):1341–1378.

Roth AE, Peranson E (2002) The redesign of the matching market for American physicians: Some engineering aspects of economic design. *Amer. Econom. Rev.* 89(4):748–780.

Schwarz M, Yenmez MB (2011) Median stable matching for markets with wages. *J. Econom. Theory* 146(2):619–637.

Shapley L, Scarf H (1974) On cores and indivisibility. *J. Math. Econom.* 1(1):23–37.

Teo CP, Sethuraman J (1998) The geometry of fractional stable matchings and its applications. *Math. Oper. Res.* 23(4):874–891.

Ignacio Rios is an assistant professor of operations management at Naveen Jindal School of Management, University of Texas at Dallas, working on the design and operations of centralized markets and platforms, incorporating behavioral aspects of their users. Previously, he obtained a BS in industrial engineering and an MS in operations management from University of Chile in 2014, and an MA in economics

and a PhD in operations, information and technology from Stanford University in 2020.

Tomás Larroucau is a PhD candidate in economics at University of Pennsylvania, working at the intersection of empirical market design, labor economics, and education. The core of his research is understanding the role of centralized allocation mechanisms for improving efficiency and equity in dynamic settings. He obtained a BS in industrial engineering and an MA in public policy from University of Chile in 2013, and an MA in economics from University of Pennsylvania in 2018.

Giorgiogiulio Parra is a senior engineer and researcher at Vice-Rectory of Information Technologies, Universidad de Chile. His working areas include web information support systems, monitoring, support of the assignment algorithm, and computer-based assessment initiatives. His main research interests include distributed and high-availability systems, learning analytics, and operations research.

Roberto Cominetti is a professor with the Faculty of Engineering and Sciences at Universidad Adolfo Ibáñez, Santiago, Chile. His research interests include convex analysis and game theory and their applications in transportation networks.