Mathematical Programming for Nominating Exchange Students for International Universities: The Impact of Stakeholders' Objectives and Fairness Constraints on Allocations

Daniel Gartner[1] and Rainer Kolisch[2]

[1] School of Mathematics, Cardiff University, Senghennydd road, Cardiff CF24 4AG, United Kingdom, gartnerd@cardiff.ac.uk

[2] TUM School of Management, Technische Universität München, Arcisstr. 21, 80333 München, Germany, rainer.kolisch@tum.de

Initial submission: November 22, 2019

Decision: Major revision (March 15, 2020)

1st revision submitted: July 31, 2020

Decision: Accepted for publication (November 11, 2020)

Mathematical Programming for Nominating Exchange Students for International Universities: The Impact of Stakeholders' Objectives and Fairness Constraints on Allocations

Abstract

We consider the problem of nominating exchange students to attend international universities where places are limited. We take into account three objectives: The sending university aims to maximize the number of nominations, the students seek nomination for a highly preferred university and, finally, the receiving universities strive for excellent incoming students. Pairwise comparison of students should guarantee the following fairness: A student with higher academic achievements should be preferred over a student with lower academic achievements. We provide mathematical programming models of the nomination problem which maximize the overall objectives and guarantee different types of pairwise fairness. Several years of real data from a major school are employed to evaluate the models' performance including a benchmark against the heuristic that is used by the school. We show analytically and experimentally that the heuristic approach fails to guarantee some pairwise fairness. Our results reveal the following four insights: First, compared to the current approach, up to 6.6% more students can be nominated with our optimization model while ensuring all pairwise fairness perspectives. Second, on average, students are nominated with better academic achievements. Third, the problem instances can be solved to optimality within a fraction of a second even for large-size instances comprising more than 500 students and about 150 schools offering nearly 450 exchange places. This is important for its use in practice. Last, up to 17.9% more students can be nominated when considering the overall objective to maximize nominations.

Keywords: Education, Student placement, Multiple knapsack problem, Greedy algorithm

1 Introduction

Spending a semester abroad during undergraduate or graduate studies has become highly attractive to students world-wide (see Daly (2011); European Commission (2019)). As a central part of fostering higher education, outbound mobility programs typically cover traveling expenses and tuition fee waiving for student exchange between two universities in different countries. In order to receive financial support, it is required that the student's home university, henceforth denoted as sending university, has a contract with a foreign university, henceforth denoted as receiving university. In this bilateral agreement, it is specified the maximum number of students that can be sent and received between the two universities within a period of time.

When assigning students to exchange places, henceforth denoted as nomination, we distinguish three stakeholders: The sending university, the students and the receiving universities. Each stakeholder has its own interests. First, the sending university strives to nominate as many students as possible. Second, each student wants to be sent to the receiving university he or she mostly prefers and when competing for scarce exchange places, he expects the nomination to be fair. Third, the receiving universities want to receive excellent students where excellence is e.g. measured in terms of grade point average (GPA), language proficiency and motivation to stay at the receiving school. These different objectives as well as the competition between students for scarce exchange places lead to the problem of fairly nominating students.

We formalize interests of the stakeholders with three objectives and two fairness perspectives. Then, we formulate the problem by using mathematical programming. Our binary programs maximize the different objectives separately while the different types of fairness are employed as constraints. Afterwards, we provide a statistical analysis of the students' application behavior observed in several application years at a major business school in Europe. Finally, in an experimental study, the solutions of the mathematical programs are compared with a heuristic nomination procedure.

The results of our computational study demonstrate that with the proposed mathematical model, more students can be nominated by at the same time ensuring more fairness as compared to the current approach. Our main contributions are therefore firstly, that we provide an innovative approach to formulating different fairness perspectives when nominating students to exchange places employing mathematical programming. Secondly, we provide an order of fairness without neglecting the different objectives of the stakeholders of the nomination process. Finally, in our experimental study using real data, we give managerial insights which approach should be considered in which situation.

The remainder of the paper is structured as follows. The next section provides an overview of related publications in which we highlight similarities and differences with our work. Section 3 contains the problem description including the definition of fairness perspectives. Section 4 introduces the mathematical programming formulation of the nomination problems. Section 5 presents the nomination heuristic which is currently used by the school. A comprehensive computational study including a detailed fairness evaluation is provided in Section 6, followed by a discussion of the managerial relevance and concluding remarks in Section 7.

2 Literature Review

In the following we provide a literature review on matching and assignment problems in higher education and academia. We state how our problem differs from and extends this literature. Following the notion of Chu and Beasley (1997), the assignment problem is about jobs, which are assigned to agents such that the capacity of the agents is respected and that the assignment costs are minimized. In higher education and academia, examples for jobs and agents are students and universities, respectively. Students are assigned to universities which have a limited capacity. Other problems include students who have to be assigned to seminar theses or conference papers which have to be assigned to reviewers.

To structure our review of related work, we will take into account our aspects: Fairness, preferences, if each job has to be assigned to exactly one agent, and multiple objectives. Table 1 lists the relevant literature in terms of these four criteria.

Table 1: Related work

	Fairness	Preferences of		Job to agent assignment		Multi-objective
		Jobs (e.g. students)	Agents (e.g. universities)	1	$\neq 1$	
Abdulkadiroğlu and Sönmez (2003)	✓	✓	✓		✓	
Alvarez-Valdés et al. (2000)		✓	✓		✓	
Al-Yakoob and Sherali (2006)	✓	✓	✓		✓	
Badri (1996)		✓	✓		✓	✓
Bafail and Moreb (1993)		✓	✓	✓		✓
Balinski and Sönmez (1999)	✓	✓	✓		✓	
Bailey and Michaels (2019)		✓	✓	✓		✓
Behestian-Ardekani and Mahmood (1986)	✓	✓	✓	✓		✓
Breslaw (1976)		✓			✓	
Cechlárová et al. (2018)		✓			✓	
Chiarandini et al. (2019)	✓	✓	✓	✓		
Diebold and Bichler (2017)		✓	✓	✓		
Geiger and Wenger (2010)	✓	✓	✓		✓	✓
Graves et al. (1993)		✓			✓	
Jin et al. (2020)		✓	✓		✓	
Kenekayoro et al. (2020)		✓	✓	✓		
Kominers et al. (2010)		✓			✓	
Lee and Clayton (1972)		✓	✓		✓	✓
Manlove and O'Malley (2008)		✓	✓		✓	
Miyaji et al. (1987)		✓	✓	✓		✓
Othman et al. (2010)		✓			✓	
Özdemir and Gasimov (2004)		✓	✓	✓		✓
Reeves and Hickman (1992)		✓	✓	✓		✓
Saber and Ghosh (2001)		✓	✓		✓	
Sanchez-Anguix et al. (2019)	✓	✓	✓	✓		✓
Santarisi and Salhieh (2005)		✓			✓	
Schniederjans and Kim (1987)		✓			✓	✓
Sönmez and Ünver (2010)		✓			✓	
Toroslu and Arslanoglu (2007)		✓			✓	✓
Wang et al. (2019)		✓			✓	
Weitz and Jelassi (1992)		✓	✓	✓		✓

The table reveals that matching and assignment in higher education and academia have been studied extensively. We now highlight similarities and differences of our work with the literature using the four introduced criteria.

Fairness Papers that address fairness can be divided into two streams: Firstly, fairness from the individual students' perspective and fairness in terms of balancing for example lecturers' workload or group sizes. One of the seminal papers on considering fairness from the individual students' perspective is Balinski and Sönmez (1999) who introduce the concept of Pareto-optimality for the student placement problem. They also introduce a fairness mechanism which always ensures that students with better scores are assigned their higher preferences. By contrast, our models compare individual applicants' pairwise fairness and their preferences among each other, as Section 3.3 will reveal: We introduce a refined concept of fairness and mathematical programming to prevent unfair nominations on different fairness perspectives which is a major difference of our work as compared to Balinski and Sönmez (1999).

Matching of students to schools under consideration of choices and fairness has also been studied by Abdulkadiroğlu and Sönmez (2003) who propose a mechanism design approach. The difference to our approach is, however, that each of the applicants is nominated for at most one university. A third paper that considers fairness of individual students is Chiarandini et al. (2019) who define a student-project allocation as being fair if the most disadvantaged students are allocated to their highest preferences. The rational behind this is that the solution promotes a certain degree of egalitarianism in the outcome. This is different to our setting where students compete against each other for highly demanded and scarce exchange places.

Related work in the second stream, i.e. to fairly balance workload and group sizes, is Geiger and Wenger (2010). They aim for a fair balance of workload across lecturers in their student assignment problem. Similar to our approach, they employ a MIP where agents have a maximum capacity, which allows to assign multiple jobs. In contrast to our approach, they do not consider pairwise fairness between students. Al-Yakoob and Sherali (2006) propose a model that generates a class-faculty assignment. They define fairness as equity measure among faculty teaching similar loads. This is again different to our model because we are looking at fairness at the individual student level and prevent unfair assignment using our MIP approach. Behestian-Ardekani and Mahmood (1986) develop a model to assign students to groups for class projects. They define an allocation as being fair if experienced and inexperienced students may learn equally well. Both experienced and inexperienced students are assigned to groups in order to create an acceptable group balance. This is again different to our work because there, we do not define groups of students or groups of universities. Finally, Sanchez-Anguix et al. (2019) define fairness as a workload balancing measure across supervisors of student projects rather than considering fairness between two students who compete for scarce places.

Preferences Several papers address preferences either by the jobs (e.g. students), agents (e.g. supervisors or universities) or two-sided. One-sided preferences have been modeled by, for example, Cechlárová et al. (2018) who, similarly to our work, model preferences of students by lexicographic ordering them. However, they assign course bundles to students and not students to at most one university. Papers that take into account two-sided preferences are, again, Abdulkadiroğlu and Sönmez (2003) in which students have preferences for universities but also universities seek for excellent students to be admitted to their programs. To the best of our knowledge, there is no literature that takes into account only one-sided preferences from an agent's perspective, i.e. without considering the students' preferences.

Job to agent assignment Job to agent assignment can be differentiated into i) whether one job must be assigned to exactly one agent or not and ii) whether an agent has a capacity for one job or several jobs. If the assignment is optional and multiple agents with individual capacities are considered, we have a multiple knapsack problem (see Section 3.4). Cechlárová et al. (2018) is an example paper in which applicants have preferences over bundles of courses. The applicants are then assigned to courses each having a limited capacity. Also, Jin et al. (2020) consider this feature when addressing the reviewer assignment problem, in which reviewers' capacities are scarce which is modeled using a multiple knapsack constraint. In contrast, a problem where each job has to be assigned to exactly one agent is given in Geiger and Wenger (2010) where each student must be assigned to exactly one topic. This is different to our problem because it is not guaranteed that students are nominated for a university. Another difference to our work is, that the interaction between jobs (students) assigned to the same agent (lecturer) is not relevant for our problem. Also the balancing of workloads or allocations is not our objective.

Multi-criteria optimization In terms of multi-criteria optimization, three approaches can be found in the assignment and matching literature for higher education and academia: Aggregation of objective functions by the weighted sum method, reference point methods such as goal programming, and separate treatment of objective functions. An example of objective function aggregation by the weighted sum method is Al-Yakoob and Sherali (2006) who seek to minimize the total dissatisfaction and inequity costs associated with assigning faculty members to classes. The total students' dissatisfaction cost is combined with the sum of differences in dissatisfaction levels between faculty members having identical teaching loads. A weight factor is introduced to reflect the relative priorities of the different measures. Examples of the goal programming approach are Badri (1996) and Bafail and Moreb (1993). Finally, an example for treating the objective functions separately is Chiarandini et al. (2019), who tackle the problem of student to project allocation. This is similar to our approach where we solve different combinations of one objective function and fairness constraints.

As a conclusion, the models and methods proposed in this paper can be differentiated from the literature as follows: With respect to fairness, we refine and extend the fairness concept of Balinski and Sönmez (1999). Our approach takes into account an ordering of preferences of both students and universities both in the objective function as well as in the constraints. The impact of the different objective functions and the constraints is then evaluated in an experimental study. With respect to job to agent assignment, our work has similarities with the Multiple Knapsack Problem as considered in Jin et al. (2020) who have limited capacities of reviewers in their manuscript to reviewer assignment problem. Finally, with respect to multi-criteria optimization, we provide a test design that solves the different objective functions and constraints separately instead of aggregating the different objectives by the weighted sum method as, e.g., undertaken in Jin et al. (2020).

3 Problem description and fairness perspectives

In this section, we first provide a problem description including all necessary notation. The relevant objectives for the nomination problem are discussed in Section 3.2. Afterwards, Pareto efficiency and the pairwise fairness perspectives are introduced in Section 3.3.

3.1 Problem description

Let $\mathcal{S}$ denote the set of students and $\mathcal{U}$ the set of universities. Each student $s \in \mathcal{S}$ applies for at least one university $u \in \mathcal{U}$ which has a capacity $c_u \in \mathbb{N}$ for receiving students. Let $\mathcal{A}$ be a set of applications in which tuple $(s, u) \in \mathcal{A}$ represents an application of student s for university u. Similar to the concept of course bidding, studied e.g. in Sönmez and Ünver (2010), students have preferences which are expressed by ranks. Let $r_{s,u}$ denote the ordinal scaled rank of application $(s, u) \in \mathcal{A}$. Without loss of generality, we assume that ranks for universities are totally ordered for each student $s \in \mathcal{S}$ who applies for more than one university, i.e. $\{r_{s,u} < r_{s,u'} | \forall (s, u), (s, u') \in \mathcal{A} : u \neq u'\}$ meaning that for each pair of universities $u, u' \in \mathcal{U} : u \neq u'$, student s prefers one university u strictly more than another university u'. Table 2(a) provides an example. It reveals that student s_1 ranks university u_1 higher than university u_2. Moreover, students s_2 and s_3 rank university u_1 equally.

In addition to ranks, we take into account the aptitudes of students. Aptitudes are measured on the basis of student's application documents. The application documents and as a result aptitudes can be either university-dependent or university-independent. University-dependent application documents are for example language certificates and motivation letters. We denote university-dependent aptitudes as $a_{s,u}$ for application $(s,u) \in \mathcal{A}$. The lower the student s's aptitude $a_{s,u}$ the better his qualification for university u. Real-valued scores are normalized between 0 and 1 and Appendix 2 provides details of the calculation. In contrast, the university-independent aptitude is measured by the student's grade point average (GPA), denoted by g_s.

We also take into account the number of semesters student s has already spent at a university abroad. The rationale behind this integer parameter, denoted by e_s, is because the aim is to prioritize students who have spent less semesters abroad as compared to their competitors.

In the following, we represent a feasible solution for the nomination process by binary vector $\boldsymbol{x} = (x_{s,u})_{(s,u)\in\mathcal{A}}$; $x_{s,u} = 1$, if application $(s,u) \in \mathcal{A}$ turns into a nomination, i.e. student s is nominated to university u and 0 otherwise. Set $\mathcal{X}$ contains all feasible solutions.

3.2 Objectives

As stated above, we have three objectives reflecting the interests of the different stakeholders. The objective of the sending university is to maximize the number of nominated students in order to provide for as many students as possible an international experience (nomination-oriented objective). The receiving universities want to receive highly motivated students (rank-oriented objective) as well as highly qualified students (aptitude-oriented objective). In general, the objectives are conflicting. In order to demonstrate the different objectives and to show that they are conflicting let us introduce the following example.

Example. Assume we have students $\mathcal{S} := \{1,2,3\}$ and universities $\mathcal{U} := \{1,2,3\}$ of which each has a capacity of $c_u = 1$. Table 2(a) gives the ranks $r_{s,u}$ of the applications $(s,u) \in \mathcal{A}$ where empty cells mean that student s does not apply for university u.

Table 2: Ranks and nominations for an example problem

(a) Ranks of students

$r_{s,u}$	$u_j \in \mathcal{U}$		
$s_i \in \mathcal{S}$	u_1	u_2	u_3
s_1	2	3	1
s_2	1		2
s_3	1	2	

(b) Solution which does not meet the nomination-oriented objective

$\boldsymbol{x}$	$u_j \in \mathcal{U}$		
$s_i \in \mathcal{S}$	u_1	u_2	u_3
s_1	0	0	1
s_2	0		0
s_3	1	0	

(c) Solution which does not meet the rank-oriented objective

$\boldsymbol{x}$	$u_j \in \mathcal{U}$		
$s_i \in \mathcal{S}$	u_1	u_2	u_3
s_1	1	0	0
s_2	0		1
s_3	0	1	

Tables 2(b) and (c) provide two different solutions. The solution provided in Table 2(b) is not optimal with respect to the nomination-oriented objective because one more student could have been nominated. The solution provided in Table 2(c) is not optimal with respect to the rank-oriented objective because students are nominated for universities which are not ranked highest. More precisely, by swapping assigments (s_1, u_1) and (s_2, u_3) to (s_1, u_3) and (s_2, u_1) for two students the rank for students s_1 and s_2 can be improved from 2 to 1 which gives an average rank of $\frac{4}{3}$ compared to 2.

3.3 Pareto-efficiency and pairwise fairness

3.3.1 Pareto-efficiency

We will define Pareto-efficiency and Pareto dominance following Balinski and Sönmez (1999). However, the difference is that we consider multiple objectives. We denote solution $\boldsymbol{x} \in \mathcal{X}$ Pareto-efficient if there is no solution $\boldsymbol{x}' \in \mathcal{X} \setminus \{\boldsymbol{x}\}$ which is preferred by at least one student $s' \in \mathcal{S}$ because he receives a nomination with lower rank while all other students $s \in \mathcal{S} \setminus \{s'\}$ are indifferent between solutions $\boldsymbol{x}$ and $\boldsymbol{x}'$ because for them, the rank of the assigned university does not change. A solution $\boldsymbol{x}$ is Pareto-dominated by another solution $\boldsymbol{x}'$ if it is not Pareto-efficient.

Example. Consider a solution $\boldsymbol{x} \in \mathcal{X}$ in which at least one student $s' \in \mathcal{S}$ is not nominated to any university, i.e. $\sum_{(s',u) \in \mathcal{A}} x_{s',u} = 0$. If there is at least one application $(s', u') \in \mathcal{A}$ and $\sum_{(s,u') \in \mathcal{A}} x_{s,u'} < c_{u'}$, i.e. at least one exchange place of university u' is still available, solution $\boldsymbol{x}$ is Pareto-dominated by solution $\boldsymbol{x}' \in \mathcal{X}$ with $x'_{s',u'} = 1$ and $x'_{s,u} = x_{s,u} \forall s \in \mathcal{S}, u \in \mathcal{U} \setminus \{u'\}$.

3.3.2 Trivial pairwise fairness

Exchange semester-based pairwise fairness ensures that a student with a higher number of already undertaken exchange semesters is only nominated for a university if all students applying for the same university with a lower number of exchange semesters are nominated.

Rank- and aptitude-based pairwise fairness. Consider applications $(s, u), (k, u) \in \mathcal{A}$ of student $s, k : s \neq k$ with the same number of exchange semesters, i.e. $e_s = e_k$. To compare student s and student k with respect to their ranks of and aptitudes for university u, we distinguish the 5 cases **a**, ..., **e** shown in Table 3.

Table 3: Cases when two students apply for the same university

Let applications $(s,u),(k,u) \in \mathcal{A}$: $s \neq k, e_s = e_k$	$a_{s,u} < a_{k,u}$	$a_{s,u} = a_{k,u}$	$a_{s,u} > a_{k,u}$
$r_{s,u} < r_{k,u}$	**a**	**c**	
$r_{s,u} = r_{k,u}$	**b**	**d**	
$r_{s,u} > r_{k,u}$	**e**		

In terms of a fair nomination, cases **a**, **b** and **c** are trivial because student s dominates student k and hence should receive the exchange place. The case **d** when the students s and k have equal aptitudes and ranks is resolved by employing the GPA score g_s and g_k as a tie-breaker. In the following, we treat the non-trivial case **e**.

3.3.3 Non-trivial pairwise fairness

Obviously, in case **e** we have no strict dominance of one student s compared to another student k because s is better in terms of aptitude but k is superior in terms of rank.

Strong pairwise fairness. A nomination $x \in \mathcal{X}$ fulfills the strong pairwise fairness property for applications (s,u) and (k,u) iff application (k,u) turns into a nomination as long as application (s,u) turns into a nomination or as long as student s is nominated for any university that is ranked by s with a lower rank than university u and vice versa. The subset of feasible solution vectors fulfilling the strong pairwise fairness perspective is denoted by $\mathcal{X}_s \subset \mathcal{X}$.

Table 4 gives an example of this property. Balinski and Sönmez (1999) have shown that there is no nomination model or algorithm which yields a nomination $\boldsymbol{x} \in \mathcal{X}$ being strong pairwise fair and Pareto-efficient.

The strong pairwise fairness is in favor of the best students. As a consequence, the chance to study abroad for students with low grades is reduced. To overcome this drawback, consider the following weakened definition of fairness for the nomination process.

Weak pairwise fairness. A nomination $x \in \mathcal{X}$ fulfills the weak pairwise fairness property for applications (s,u) and (k,u) iff application (k,u) turns into a nomination as long as application (s,u) turns into a nomination or student s is nominated for any other university within the student's set of applications. The subset of feasible solution vectors fulfilling the weak pairwise fairness perspective is denoted by $\mathcal{X}_w \subset \mathcal{X}$.

A nomination (process) which respects the weak pairwise fairness perspectives always receives a Pareto-efficient solution.

Proposition 1. *A solution $\boldsymbol{x} \in \mathcal{X}_w$ which is weak pairwise fair is always Pareto-efficient.*

A solution $\boldsymbol{x} \in \mathcal{X}_s$ which is strong pairwise fair is also weak pairwise fair, but not vice versa. This leads to the following result.

Proposition 2. *Let n_{SPF} and n_{WPF} be the maximum numbers of nominations under strong and weak pairwise fairness conditions, respectively. The solution space taking into account strong pairwise fairness is a non-proper subset of the solution space taking into account weak pairwise fairness, i.e. $\mathcal{X}_s \subseteq \mathcal{X}_w$. As a result, we observe $n_{SPF} \leq n_{WPF}$.*

The following example demonstrates the two pairwise fairness principles.

Example. Consider the set of students $\mathcal{S} := \{1, 2, 3\}$, the set of universities $\mathcal{U} := \{1, 2, 3\}$, each with a capacity of $c_u = 1$, as well as the students' ranks and aptitudes as given in Tables 4(a) and (b), respectively.

Table 4: Example for the nomination of students according to the weak and strong pairwise fairness

(a) Ranks

$r_{s,u}$	$u \in \mathcal{U}$		
$s \in \mathcal{S}$	u_1	u_2	u_3
s_1	2	1	
s_2	1	2	3
s_3	2	1	

(b) Aptitudes

$a_{s,u}$	$u \in \mathcal{U}$		
$s \in \mathcal{S}$	u_1	u_2	u_3
s_1	0.6	0.9	
s_2	0.8	0.7	0.7
s_3	0.7	0.8	

(c) Strong pairwise fair solution

$\boldsymbol{x}$	$u \in \mathcal{U}$		
$s \in \mathcal{S}$	u_1	u_2	u_3
s_1	1	0	0
s_2	0	1	0
s_3	0	0	0

(d) Pareto-efficient and weak pairwise fair solution

$\boldsymbol{x}$	$u \in \mathcal{U}$		
$s \in \mathcal{S}$	u_1	u_2	u_3
s_1	0	1	0
s_2	0	0	1
s_3	1	0	0

In Table 4(c), the nomination based on the strong pairwise fairness is not Pareto-efficient as one more student can be nominated while ensuring the weak pairwise fairness as shown in the Pareto-efficient and weak pairwise fair solution of Table 4(d). However, the solution provided in Table 4(c) is strong pairwise fair because the following can be observed: Students s_1 and s_2 pairwise compete for university u_1. Student s_1 is nominated for this university because his aptitude is better than or equal to the one of student s_2. Moreover, student s_1's rank for this university is lower than or equal to the rank of student s_2. Similar conditions hold true for university u_2. Since student s_3's aptitudes are worse than the ones of students s_1 and s_2 and capacity is scarce, he cannot be nominated.

3.4 Complexity of special and general problems

In some universities and application settings, it may occur that students are allowed to give only one preference in their set of applications. Then, an algorithm can be described as follows: Consider universities in the order $u_1, u_2, \ldots, u_{|\mathcal{U}|}$. For each university u_i, sort applicants for this university depending on the objective function criterion. Then, nominate the applicants until capacity is used up. Running through the $|\mathcal{U}|$ universities and sorting can be done in polynomial time.

In the general case, however, we have a General Assignment Problem which is a variant of the NP-hard Multiple Knapsack Problem: Each student and university correspond to a job and agent, respectively. To prove NP-hardness, construct a directed graph with one source and one sink node. Students and universities represent nodes. The source connects all student nodes with a directed edge. All university nodes are connected with the sink node with a directed edge. Now, connect student's applications with universities with a directed edge. Since each student must not be connected to multiple universities, we strive for (the NP-hard problem of) maximizing unsplittable flows (Martens and Skutella (2006)) between the source and sink node because we maximize nominations, preferences or aptitudes.

4 Model formulations

In this section, we present a model which covers the objectives as well as the fairness definitions presented in the previous section.

4.1 Objectives and base model

Employing the decision variables

$$x_{s,u} = \begin{cases} 1, & \text{if application } (s,u) \in \mathcal{A} \text{ becomes a nomination} \\ 0, & \text{otherwise} \end{cases}$$

we can formulate the nomination-oriented, rank-oriented and aptitude-oriented objectives z_{N}, z_{R}

and z_{A} introduced in Section 3.2 in Equations (1), (2) and (3), respectively.

$$\text{Maximize } z_{\mathrm{N}}(x) = \sum_{(s,u)\in\mathcal{A}} \bar{n}\cdot x_{s,u} \tag{1}$$

$$\text{Maximize } z_{\mathrm{R}}(x) = \sum_{(s,u)\in\mathcal{A}} \bar{r}_{s,u}\cdot x_{s,u} \tag{2}$$

$$\text{Maximize } z_{\mathrm{A}}(x) = \sum_{(s,u)\in\mathcal{A}} \bar{a}_{s,u} \cdot x_{s,u} \tag{3}$$

Since ranks and aptitudes are the smaller the better, we recalculate and normalize them. We introduce the objective function coefficients $\bar{n}$, $\bar{r}_{s,u}$ and $\bar{a}_{s,u}$ for one normalized nomination, rank and aptitude objective, respectively. This ensures that z's dimensions are between 0 and 1. An explanation of the calculation is given in Section 2 of the Supplementary Materials. By adding constraints (4)–(5) and decision variables (6) to the objectives, the problem turns out to be a special zero-one Multiple Knapsack Problem (MKP) which is NP-hard (Kellerer et al. (2004) and Martello and Toth (1990)).

$$\sum_{(s,u)\in\mathcal{A}} x_{s,u} \leq 1 \qquad \forall s \in \mathcal{S} \tag{4}$$

$$\sum_{(s,u)\in\mathcal{A}} x_{s,u} \leq c_u \qquad \forall u \in \mathcal{U} \tag{5}$$

$$x_{s,u} \in \{0,1\} \qquad \forall (s,u) \in \mathcal{A} \tag{6}$$

Constraints (4) ensure that each student is nominated to at most one university while constraints (5) ensure that the capacity of each university is not exceeded. Binary decision variables are defined by (6). Alternatively, our problem can be formulated as an Assignment Problem (see e.g. Burkard et al. (2009)) by introducing dummy universities with unlimited capacity. However, this would require additional variables.

4.2 Pairwise fairness constraints

The semester-based pairwise fairness constraints read as follows:

$$x_{k,u} \leq \sum_{(s,l)\in\mathcal{A}} x_{s,l} \qquad \forall (s,u),(k,u) \in \mathcal{A} : s \neq k \quad \text{and} \quad e_s < e_k \tag{7}$$

Constraints (7) ensure that if two students s and k compete for the same university u, the student k with the larger number of exchange semesters is not nominated unless the student s with the lower number of exchange semesters is nominated for this or any other university in his set of applications.

$$
x_{k,u} \leq \sum_{(s,l)\in\mathcal{A}:r_{s,l}\leq r_{s,u}} x_{s,l} \qquad
\begin{array}{ll}
\forall (s,u),(k,u) \in \mathcal{A} : s \neq k,\ e_s = e_k \quad \text{and} & \\
((r_{s,u} < r_{k,u},\ a_{s,u} < a_{k,u}) \ \text{ or} & \text{case (a)} \\
(r_{s,u} = r_{k,u},\ a_{s,u} < a_{k,u}) \ \text{ or} & \text{case (b)} \\
(r_{s,u} < r_{k,u},\ a_{s,u} = a_{k,u}) \ \text{ or} & \text{case (c)} \\
(r_{s,u} = r_{k,u},\ a_{s,u} = a_{k,u},\ g_s < g_k)) & \text{case (d)}
\end{array}
\tag{8}
$$

Constraints (8) become relevant when the numbers of exchange semesters of two applicants are identical. Each line below the universal quantifier represents the cases **a**, **b**, **c**, and **d** as given in Table 3. For example, let the pairwise comparison of applications (s,u) and (k,u) satisfy condition **a**, i.e., application (s,u) represents a superior rank and aptitude score as compared to application (k,u), then application (k,u) is only allowed to turn into a nomination, if application (s,u) turns to a nomination for this or any higher ranked university. Set $\mathcal{X}$ contains all solutions described by constraints (4)–(8).

4.2.1 The weak pairwise fairness constraint

The weak pairwise fairness constraints ensure that, considering one university, the application of a student k with weaker aptitude turns into a nomination only if the application of a student s with better aptitude turns into a nomination to any university. This is expressed by the following additional constraints:

$$
x_{k,u} \leq \sum_{(s,l)\in\mathcal{A}} x_{s,l} \qquad
\begin{array}{l}
\forall (s,u),(k,u) \in \mathcal{A} : s \neq k, \\
r_{s,u} > r_{k,u}, a_{s,u} < a_{k,u}, e_s = e_k
\end{array}
\tag{9}
$$

Depending on the input parameters, the solution when incorporating the weak pairwise fairness constraint leads to a substantial reduction in nominations, compared to the nomination problem without constraints (9). We will empirically investigate on this in the experimental study in Section 6. Set $\mathcal{X}_w$ contains all solutions described by constraints (4)–(9).

4.2.2 The strong pairwise fairness constraint

The additional constraints for the strong pairwise fairness are

$$x_{k,u} \leq \sum_{(s,l)\in\mathcal{A}:r_{s,l}\leq r_{s,u}} x_{s,l} \qquad \begin{array}{l}\forall (s,u),(k,u)\in\mathcal{A}: s\neq k,\\ r_{s,u}>r_{k,u}, a_{s,u}<a_{k,u}, e_s=e_k\end{array} \tag{10}$$

Incorporating weak or strong pairwise fairness constraints into the model increases the problem size by $\frac{|\mathcal{A}|\cdot(|\mathcal{A}|-1)}{2}$ constraints. Set $\mathcal{X}_s$ contains all solutions described by constraints (4)–(8) and (10). The strong pairwise fairness constraint (10) leads to additional reductions in nominations, compared to the nomination problem without constraints (10) as our experimental study will reveal.

The number of decision variables of problem (4)–(8), regardless of the fairness constraints, is $|\mathcal{A}|$. The number of constraints is at most $|\mathcal{S}|+|\mathcal{U}|+\frac{|\mathcal{A}|\cdot(|\mathcal{A}|-1)}{2}$. Table 7 in Section 6 provides an overview of the problem sizes of real-world instances. The actual number of generated constraints, however, depends on the input data and cannot be stated generically.

5 Nomination heuristic

Currently, the school employs a nomination heuristic, for which the pseudo code is given in Algorithm 1. It is similar to Martello and Toth (1981)'s greedy heuristic for the Multiple Knapsack Problem with the following three major differences: i) all items' weights are equal to 1 ii) the knapsacks (in our case universities) are not ordered by increasing capacity and iii) no local improvement exchanges are performed.

Algorithm 1 Nomination heuristic

1: Initialize $x_{s,u} := 0 \quad \forall (s,u) \in \mathcal{A}$.
2: **for all** $i = 0, \ldots, \max\limits_{s \in \mathcal{S}} e_s$ **do**
3:     **for all** $r = 1, \ldots, \max_{(s,u) \in \mathcal{A}} r_{s,u}$ **do**
4:         **for all** $u \in \mathcal{U}$ (in arbitrary order) **do**
5:             $\mathcal{T}_u := \emptyset$;
6:             **for all** $(s,u') \in \mathcal{A} : u' = u, \sum\limits_{(s,u) \in \mathcal{A}} x_{s,u} = 0, r_{s,u'} = r, e_s = i$ **do**
7:                 Insert $(s,u') \cup \mathcal{T}_u$ by ascending $a_{s,u'}$ and g_s scores;
8:             **end for**
9:             **for** $(s,u) \in \mathcal{T}_u$ **do**
10:                 **if** $c_u > 0 \land x_{s,u} = 0$ **then**
11:                     $x_{s,u} := 1$;
12:                     $c_u := c_u - 1$;
13:                 **end if**
14:             **end for**
15:         **end for**
16:     **end for**
17: **end for**

The first outer loop (Lines 2–17) runs from zero undertaken exchange semesters to the maximum number of undertaken exchange semesters over all students. The second outer loop (Lines 3–16) runs from the highest to the lowest ranks. In the inner loop (Lines 4–15), all universities $u \in \mathcal{U}$ are processed in an arbitrary order. The first step in the inner loop is an ordering step (Lines 6–8): Applications with rank p for university u are inserted into a list (ordered set) $\mathcal{T}_u$ for each university $u \in \mathcal{U}$. The order is determined by ascending aptitude scores and the GPA sub-score is used as a tie breaker if two students' aptitudes are equal. The second step in the inner loop (Lines 9–14) is the nomination step: The loop runs through the sorted list $\mathcal{T}_u$. Students are nominated for university u, if they have not been nominated to any other university and if, u's capacity has not been used up. In what follows, we highlight two properties of the heuristic nomination.

Proposition 3. *The nomination heuristic ensures pairwise fairness* **a**,. . . , **d**.

Proof. First, we observe that the algorithm runs sequentially through all ranks. Assume two applications $(s,u), (k,u) \in \mathcal{A}$ with $s \neq k$ for university u and $r_{s,u} < r_{k,u}$, i.e. student s has a higher rank than student k for university u. In this case, the application of student s for university u would be considered in an earlier loop than the application of student k, regardless of both students' aptitudes. If student k is nominated to university u, then it follows that the capacity of u was not used up in the earlier loops. Consequently, student s was definitely nominated either for u or any other higher ranked university in one of those loops. Hence, pairwise fairness **a** and

c are ensured. PF condition **c** or **d** applies if student s and student k both submitted the same rank for university u. Thus, both applications would be in the same list $\mathcal{T}_u$. Since the list is sorted according to descending aptitudes and descending GPA scores as a tie-breaker, student k can only be nominated to university u if student s is nominated to u or any other higher ranked university, because student s's application is further up in the list.

Proposition 4. *The nomination heuristic fails to ensure weak pairwise fairness and strong pairwise fairness.*

Proof. We show this by contradiction. Consider two students s_1, s_2 and one university u_1 with capacity $c_{u_1} = 1$. Let $a_{s_1,u_1} < a_{s_2,u_1}$, i.e. student s_1 has a better aptitude for that university compared to student s_2. Assume u_1 is s_1's rank 3 university while it is s_2's rank 1 university, such that $r_{s_1,u_1} > r_{s_2,u_1}$. Then student s_2's application is considered in an earlier loop than student s_1's. If the capacity of university u_1 is already used up when student s_2's application is considered (because of the higher rank), and s_1 was not nominated to this or any other higher ranked university (note, we assume that there is only one university), student s_1 ends up not being nominated at all. This contradiction results in a violation of the weak pairwise fairness condition and, in consequence (see Proposition 2), in a violation of the strong pairwise fairness condition.

6 Experimental study

In the experimental study, we test the models with real-world data provided by a European Business School. The experimental study is organized as follows: Section 6.1 introduces the data and provides summary statistics. Section 6.2 presents an overview of our evaluation measures as well as the structured experimental design. A computation time analysis of real-word test instances is given in Section 6.3, followed by a presentation and an evaluation of the results in Sections 6.4–6.8.

6.1 Data analysis

We use seven years worth of data containing applications that the school received in the years 2008–2013 (years 1 to 6 in Table 5) and 2020 (year 7 in Table 5). Each year, students apply for exchange places. After the application deadline has passed, all application documents are evaluated in order to determine the parameters for the nomination approach presented in Section 5. Table 5 provides summary statistics of all application years. For each year, each student could apply for at most 3 universities.

Table 5: Summary statistics for each application year

Application year	1	2	3	4	5	6	7
Number of students $\|\mathcal{S}\|$	60	79	90	117	151	226	521
Number of universities $\|\mathcal{U}\|$	26	27	33	44	50	65	142
Number of applications $\|\mathcal{A}\|$	155	210	243	305	387	600	1,475
Number of exchange places $\sum_{u \in \mathcal{U}} c_u$	70	80	97	138	153	210	448
Number of students per exchange place	0.857	0.988	0.928	0.848	0.987	1.076	1.163

The table reveals an increasing demand and a shortage of exchange places for year 6 and 7. The average number of applications per student is quite constant between 2.56 and 2.83 within the seven years.

6.2 Evaluation measures and experimental design

All mathematical programs and the nomination heuristic used by the school are assessed using the same performance indicators: The number of nominations, the relative number of nominations to 1st, 2nd and 3rd rank universities, and the average aptitude of nominated students. In addition, we evaluate violations of pairwise fairness introduced in Section 3.3. Furthermore, we provide an analysis of the upper bound of students that can be nominated by relaxing all pairwise fairness constraints.

The 9 test setups that we use in order to evaluate the different objectives combined with the different fairness perspectives are shown in Table 6.

Table 6: Test setups are indicated by "•". "○" means that we do not evaluate the corresponding objective or fairness constraint

	Objective function			**Pairwise fairness**	
Setup	Nominations (z_{N})	Ranks (z_{R})	Aptitudes (z_{A})	Weak	Strong
1	•	○	○	○	○
2	•	○	○	•	○
3	•	○	○	○	•
4	○	•	○	○	○
5	○	•	○	•	○
6	○	•	○	○	•
7	○	○	•	○	○
8	○	○	•	•	○
9	○	○	•	○	•

The table shows that the three different objectives are paired with the different pairwise fairness constraints (off, weak and strong). As a consequence, we obtain 9 different setups given in Table 6.

For all setups, we incorporate the rank- and aptitude- as well as the exchange semester-based pairwise fairness as given by constraints (7) and (8), respectively (see also Section 3.3.3). The reason to incorporate semester-based pairwise fairness is that we want a fair comparison with the nomination heuristic used by the school which accounts for the exchange semester-based pairwise fairness. The reason to generate nine setups is that we perform a full-factorial test design by varying the three objectives given in equations (1)–(3) paired with in- and excluding the weak pairwise fairness as given by constraints (9) and the strong pairwise fairness, see constraints (10).

6.3 Computational results

All computations were performed on a 3.1 GHz personal computer (Intel Core i7-4940MX) with 32 GB RAM running a Windows 10 operating system. The models were coded in Java 1.8 using the 64 bit version of the application programming interface of IBM ILOG CPLEX 12.10. For the sake of computational comparison, we coded the nomination heuristic in Java, too. Table 7 shows the sizes of the test instances and the computation time in Milliseconds required to solve each test instance.

Table 7: Number of decision variables, constraints and computation times

Application year	1	2	3	4	5	6	7
# Decision variables	155	210	243	305	387	600	1,475
# Constraints							
Basic nomination problem (1), (4)–(6)	86	106	123	161	201	291	663
Incl. exchange semester-, rank-/aptitude-based and strong pairwise fairness	806	1,348	1,422	1,953	2,358	5,155	17,044
Computation time [ms]							
Nomination heuristic	2	1	3	2	2	2	10
MIP (Setup 1)	14	22	20	22	40	89	93
MIP (Setup 2)	13	32	34	34	60	161	194
MIP (Setup 3)	12	25	27	32	55	53	136
MIP (Setup 4)	16	21	21	23	63	42	97
MIP (Setup 5)	13	33	31	28	156	63	130
MIP (Setup 6)	19	21	28	40	60	58	127
MIP (Setup 7)	18	52	26	31	418	246	146
MIP (Setup 8)	17	19	37	36	58	62	135
MIP (Setup 9)	14	20	34	33	42	85	121

We observe that the computation times for solving each model for each application year is below one second time and that the solution time of the nomination heuristic is even faster. We can also observe that the number of additional pairwise fairness constraints increases more than the constraints of the basic nomination problem. One explanation for this phenomenon is that the number of applications increases faster than the number of available places per application.

6.4 Comparison with the nomination heuristic

The results of the nomination heuristic are provided in Table 8 and show that the ratio of nominations to all applicants is between 61% and 85%.

Table 8: Results of the nomination heuristic

Application year	1	2	3	4	5	6	7
# Nominations	51	59	69	90	120	146	318
Percentage of applying students	85.0	74.7	76.7	76.9	79.5	64.6	61.0
Percentage of nominations compared to setup 1	96.2	96.7	97.2	95.7	98.4	96.7	93.0
Percentage of nominations compared to setup 4	96.2	98.3	98.6	97.8	99.2	98.6	96.1
Percentage of nominations compared to setup 7	96.2	96.7	97.2	95.7	98.4	96.7	106.3
Percentage of rank 1 universities	74.5	74.6	73.9	76.7	77.5	70.5	74.2
Percentage of rank 2 universities	11.8	13.6	20.3	15.6	16.7	19.9	15.7
Percentage of rank 3 universities	13.7	11.9	5.8	7.8	5.8	9.6	10.1
Average aptitude of nominated students	0.43	0.56	0.53	0.50	0.76	0.57	0.42
Violations of exchange semester- and rank-/aptitude-based pairwise fairness	0	0	0	0	0	0	0
Violations of weak pairwise fairness	6	16	9	20	16	69	159
Violations of strong pairwise fairness	7	23	9	28	19	73	174

Due to Proposition 3, the exchange semester- as well as the rank- and aptitude-based pairwise fairness are never violated. However, a considerable number of violations of weak and strong pairwise fairness occur and as a consequence, some students are not nominated at all although their aptitude is better than the lowest aptitude of the students, which have been nominated to the same university. This demonstrates that the nomination heuristic is ill-suited for the fair nomination of exchange students.

6.5 Results of the upper bound on nominations

Maximizing objective function (1) in combination with constraints (4)–(5), and decision variables (6) gives an upper bound on the number of nominations. Its results are presented in Table 9 and show that in none of the application years a 100% nomination rate is obtained.

Table 9: Results of the upper bound problem (1), (4)–(6)

Application year	1	2	3	4	5	6	7
# Nominations (upper bound)	57	66	81	101	134	164	375
Percentage of applying students	95.0	83.5	90.0	86.3	88.7	72.6	72.0
Percentage of rank 1 universities	42.1	40.9	28.4	31.7	41.8	29.3	29.1
Percentage of rank 2 universities	38.6	24.2	42.0	37.6	32.1	33.5	30.1
Percentage of rank 3 universities	19.3	34.8	29.6	30.7	26.1	37.2	40.8
Violations of semester-based pairwise fairness	0	0	0	7	0	44	178
Violations of pairwise fairness **a**	31	31	75	85	91	171	337
Violations of pairwise fairness **b**	46	57	74	75	95	196	321
Violations of pairwise fairness **c**	0	0	0	0	0	1	16
Violations of pairwise fairness **d**	0	1	0	0	0	0	11
Violations of weak pairwise fairness	3	17	14	29	9	63	143
Violations of strong pairwise fairness	9	32	23	37	33	90	216

On average, more than 4 out of 5 students can be nominated using this approach. A major drawback however is the substantially large number of fairness violations. For example, there are 75 violations of the rank- and aptitude-based pairwise fairness (case **b**) in the year 4.

6.6 Evaluation of the nomination-, rank- and aptitude-based objective

We compare the results of the nomination heuristic with the MIP solution of setups 3, 6 and 9 which enforce strong pairwise fairness when optimizing the nomination-, rank- and aptitude-oriented objective, respectively. The results are given in Tables 10–12. We also show the results of setup 1 in Table 13 in the Appendix in which the nomination-oriented objective is pursued without weak or strong pairwise fairness.

Table 10: Results of the MIP applied to setup 3 (nomination-oriented objective, strong pairwise fairness constraints)

Application year	1	2	3	4	5	6	7
# Nominations	53	61	71	93	121	149	339
Percentage of applying students	88.3	77.2	78.9	79.5	80.1	65.9	65.1
Percentage of nominations compared to setup 1	100.0	100.0	100.0	98.9	99.2	98.7	99.7
Percentage of nominations compared to the upper bound	93.0	92.4	87.7	92.1	90.3	90.9	90.4
Percentage of rank 1 universities	58.5	54.1	63.4	61.3	65.3	50.3	52.8
Percentage of rank 2 universities	22.6	27.9	25.4	24.7	23.1	32.9	23.6
Percentage of rank 3 universities	18.9	18.0	11.3	14.0	11.6	16.8	23.6
Average aptitude of nominated students	0.44	0.56	0.53	0.51	0.77	0.58	0.49
Violations of strong pairwise fairness	0	0	0	0	0	0	0

The tables show a substantial drop-off in nominations as compared to the upper bound on nominations that can be achieved by relaxing all pairwise fairness constraints, see Table 9 in Section 6.5.

Table 11: Results for setup 6 (rank-oriented objective, strong pairwise fairness constraints)

Application round	1	2	3	4	5	6	7
# Nominations	53	61	71	93	121	149	337
Percentage of applying students	88.3	77.2	78.9	79.5	80.1	65.9	64.7
Percentage of nominations compared to setup 4	100.0	98.3	98.5	99.9	100.0	99.3	97.6
Percentage of nominations compared to the upper bound	93.0	92.4	87.7	92.1	90.3	90.9	89.9
Percentage of rank 1 universities	62.3	59.0	66.2	63.4	68.6	53.7	54.0
Percentage of rank 2 universities	20.8	24.6	26.8	25.8	22.3	32.9	26.7
Percentage of rank 3 universities	17.0	16.4	7.0	10.8	9.1	13.4	19.3
Average aptitude of nominated students	0.44	0.56	0.53	0.51	0.77	0.58	0.49
Violations of pairwise fairness	0	0	0	0	0	0	0

When maximizing the rank-based objective z_{R} which is considered in setup 6 and Table 11, we observe a higher nomination rate to rank 1 universities as compared to the rate observed in setup 3. For example, when using setup 6 in application year 2, 59.0% of the nominated students are sent to the highest ranked university as compared to 50.8% using setup 3 for the same application year. Another observation is that in four of six application years, more nominations occur as compared to setup 4 with rank-oriented objective but no pairwise fairness in place. One explanation of this phenomenon is that the strong pairwise fairness ensures that excellent students are nominated rather than nominating students who rank universities higher but have not as good aptitudes.

Table 12: Results for setup 9 (aptitude-oriented objective, strong pairwise fairness)

Application year	1	2	3	4	5	6	7
# Nominations	53	61	71	93	121	149	297
Percentage of applying students	88.3	77.2	78.9	79.5	80.1	65.9	57.0
Percentage of nominations compared to setup 7	100.0	100.0	100.0	98.9	99.2	98.7	99.3
Percentage of nominations compared to the upper bound	93.0	92.4	87.7	92.1	90.3	90.9	79.2
Percentage of rank 1 universities	58.5	52.5	62.0	59.1	64.5	51.7	52.2
Percentage of rank 2 universities	22.6	29.5	26.8	25.8	24.0	31.5	25.3
Percentage of rank 3 universities	18.9	18.0	11.3	15.1	11.6	16.8	22.5
Average aptitude of nominated students	0.45	0.57	0.54	0.52	0.77	0.58	0.67
Violations of pairwise fairness	0	0	0	0	0	0	0

Table 12 gives the results when maximizing aptitude-weighted nominations in setup 9 under strong pairwise fairness constraints. We observe only a slight improvement of the average of aptitudes of nominated students. However, the scores of nominated students are, on average, strictly equal or lower than the ones of setups 3 and 6.

6.7 On the impact of relaxing weak or strong pairwise fairness

Table 13 gives the results of Setup 1 which pursues the nomination-based objective without enforcing weak or strong pairwise fairness. Remarkably, for all years, more students can be nominated as compared to the nomination heuristic. We also observe that a consistent pattern exists in a sense that more students are nominated for rank 1 universities as compared to rank 2 universities. The same holds true when comparing the percentage of rank 2 university nominations as compared to rank 3 university nomination. In the heuristic approach, this pattern cannot be observed for year 1 where more students are nominated for rank 3 universities as compared to rank 2 universities. Another observation is that in 4 of the 6 application years, less rank 3 nominations occurred in Setup 4 as compared to the nomination heuristic. However, the rank 1 nomination percentage is in Setup 4 lower as compared to the nomination heuristic. Finally, the number of weak and strong pairwise fairness violations are in 10 of 12 cases lower when comparing Setup 4 with the violations of the nomination heuristic.

Table 13: Results for setup 1 (nomination-oriented objective without weak or strong pairwise fairness)

Application year	1	2	3	4	5	6	7
# Nominations	53	61	71	94	122	151	340
Percentage of applying students	88.3	77.2	78.9	80.3	80.8	66.8	65.3
Nominations for rank 1 universities	33	34	45	58	81	86	192
Percentage of nominations	62.3	55.7	63.4	61.7	66.4	57.0	56.5
Nominations for rank 2 universities	11	16	17	23	27	40	75
Percentage of nominations	20.8	26.2	23.9	24.5	22.1	26.5	22.1
Nominations for rank 3 universities	9	11	9	13	14	25	73
Percentage of nominations	17.0	18.0	12.7	13.8	11.5	16.6	21.5
Average aptitude of nominated students	0.43	0.56	0.52	0.50	0.77	0.57	0.46
Violations of weak pairwise fairness	2	0	2	6	1	22	49
Violations of strong pairwise fairness	2	3	2	11	3	29	53

6.8 Graphical evaluation of the results

Figure 1 graphically compares setups 1, 2 and 3 with the nomination heuristic for application year 6. We chose setups 1–3 because each of the corresponding models maximizes the number of nominations.

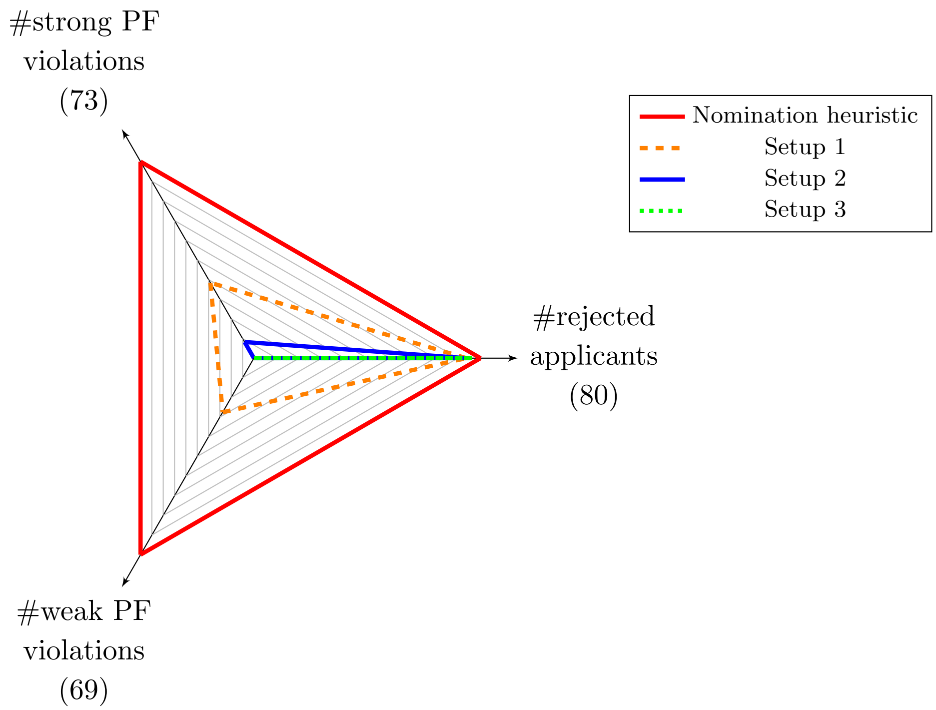

Figure 1: Kiviat diagram representing the solution characteristics for setup 1, 2 and 3 as compared to the approach currently employed in practice with maximum values in brackets

The diagram shows that the nomination heuristic considerably violates the weak and strong pairwise fairness. Moreover, the number of pairwise fairness violations employing setup 1 is still quite substantial. However, as the computational results show, we observe less than half of the fairness violations as compared to the nomination heuristic. A more detailed analysis of setup 2 reveals that, naturally, no weak pairwise fairness violations occur anymore and that the strong pairwise fairness violation drop to only 6 (see Table 15 in the Appendix). Another interesting observation is that the number of rejected applicants is one less when comparing setup 3 with the nomination heuristic. This means that our mathematical program has two major benefits and policy implications: More nominations and more fairness.

7 Managerial relevance and conclusions

7.1 Managerial relevance

This work has been originated by a major European business school. When starting the student exchange, a team member of the international office of the school undertook the nomination manually spending a considerable time for assigning the applications of up to 100 students to about the same number of partner schools (see the first three years in Table 5). However, as the number of applications and the number of partner schools grew, the time it took for processing physical applications and manually assigning applications to exchange places of partner schools was not bearable anymore. The school then decided to implement a web-based digital application system including a nomination heuristic mimicking the manual assignment undertaken so far. The logic of the heuristic is depicted in Section 5.3. Due to the digitization of the process, assignments were now derived within hours instead of days and the school was able to handle an increasing number of students and partner universities (see years 4–7 in Table 5). However, with the increase of nominations a growing number of students complained about violation of the strong pairwise fairness constraints. Complaints were in particular undertaken by very good students who competed for the limited places at highly ranked international business schools. The assignments to these schools were scrutinized in particular by the very good students. While in the early years of the student exchange, the main goal of the school was to increase the number of nominations in order to become an internationally visible business school, the school now had to manage the complaints of students about the violations of the strong pairwise fairness constraints. Since the complaints are from the best students, it is important for the school to minimize complaints and thus strong pairwise fairness violations. Hence, the demand for an advanced assignment system became apparent. Currently, the MIPs proposed in this paper have been embedded into a decision support system (DSS) and introduced to the school for deriving information about the optimum number of nominations, the total rank of assigned students, and the total aptitude of assigned students as expressed in the objective functions (1)–(3) when considering different fairness constraints. The main criteria for the school are the maximization of the number of nominations and the minimization of the strong pairwise fairness violations. At a strategic level, the DSS is judged as helpful in order to assess the value of partner schools, in particular potential new partner schools, as well as assessing the impact of changes of the number of exchange places of existing partner schools.

7.2 Conclusions

In this paper, we addressed the problem of fairly nominating students to exchange places at foreign universities. We first provided formal definitions of three objectives which are related to the three stakeholders of the nomination process. Afterwards, we introduced the concepts of weak and strong pairwise fairness and proved that weak pairwise fairness ensures Pareto efficient solutions. Next, we developed mathematical models with fairness constraints which we embedded in a full factorial experiment design varying different nomination objectives and fairness perspectives. We employed a nomination algorithm currently in use as a baseline and measured the computational performance of each approach based on real-world data. Finally, we analyzed the results based on overall performance measures and broke down the results by different evaluation metrics such as violation of fairness.

We have shown that real-world problems can be solved within less than a second time by using our mathematical programs. Our experiments confirm that maximizing the number of nominations while ensuring fairness are conflicting goals. Another result of this study is that the nomination heuristic used by the school does not guarantee fairness.

Future work could consider a strategic model that helps evaluating in which regions and for which universities a school should consider developing new partnerships. Also, the effect of bounding the maximum allowed number of applications per student on the number of nominations could be addressed.

References

Abdulkadiroğlu, A. and Sönmez, T. (2003), 'School choice: A mechanism design approach', *American economic review* **93**(3), 729–747.

Al-Yakoob, S. M. and Sherali, H. D. (2006), 'Mathematical programming models and algorithms for a class–faculty assignment problem', *Eur. J. Oper. Res.* **173**, 488 – 507.

Alvarez-Valdés, R., Crespo, E. and Tamarit, J. M. (2000), 'Assigning students to course sections using tabu search', *Ann. Oper. Res.* **96**(1–4), 1–16.

Badri, M. A. (1996), 'A two-stage multiobjective scheduling model for [faculty-course-time] assignments', *Eur. J. Oper. Res.* **94**, 16 – 28.

Bafail, A. O. and Moreb, A. A. (1993), 'Optimal allocation of students to different departments in an engineering college', *Comput. Ind. Eng.* **25**, 295–298.

Bailey, M. D. and Michaels, D. (2019), 'An optimization-based dss for student-to-teacher assignment: Classroom heterogeneity and teacher performance measures', *Decision Support Systems* **119**, 60–71.

Balinski, M. and Sönmez, T. (1999), 'A tale of two mechanisms: Student placement', *Journal of Economic Theory* **84**(1), 73–94.

Behestian-Ardekani, M. and Mahmood, M. A. (1986), 'Development and validation of a tool for assigning students to groups for class projects', *Decis. Sci.* **17**, 92–113.

Breslaw, J. A. (1976), 'A linear programming solution to the faculty assignment problem', *Socio-Econ. Plan. Sci.* **10**, 227 – 230.

Burkard, R., Dell'Amico, M. and Martello, S. (2009), *Assignment Problems*, Society for Industrial and Applied Mathematics, Philadelphia.

Cechlárová, K., Klaus, B. and Manlove, D. F. (2018), 'Pareto optimal matchings of students to courses in the presence of prerequisites', *Discrete Optimization* **29**, 174–195.

Chiarandini, M., Fagerberg, R. and Gualandi, S. (2019), 'Handling preferences in student-project allocation', *Annals of Operations Research* **275**(1), 39–78.

Chu, P. C. and Beasley, J. E. (1997), 'A genetic algorithm for the generalised assignment problem', *Computers & Operations Research* **24**(1), 17–23.

Daly, A. (2011), 'Determinants of participating in australian university student exchange programs', *Journal of Research in International Education* **10**(1), 58–70.

Diebold, F. and Bichler, M. (2017), 'Matching with indifferences: A comparison of algorithms in the context of course allocation', *European Journal of Operational Research* **260**(1), 268–282.

European Commission (2019), 'Lifelong Learning Programme'. https://ec.europa.eu/education, last accessed: November 22nd, 2019.

Geiger, M. J. and Wenger, W. (2010), 'On the assignment of students to topics: A variable neighborhood search approach', *Socio-Economic Planning Sciences* **44**(1), 25–34.

Graves, R. L., Schrage, L. and Sankaran, J. (1993), 'An auction method for course registration', *Interfaces* **23**, 81–92.

Jin, J., Niu, B., Ji, P. and Geng, Q. (2020), 'An integer linear programming model of reviewer assignment with research interest considerations', *Annals of Operations Research* **291**, 409–423.

Kellerer, H., Pferschy, U. and Pisinger, D. (2004), *Knapsack Problems*, Springer Berlin / Heidelberg.

Kenekayoro, P., Mebine, P. and Zipamone, B. G. (2020), 'Population based techniques for solving the student project allocation problem', *International Journal of Applied Metaheuristic Computing (IJAMC)* **11**(2), 192–207.

Kominers, S., Ruberry, M. and Ullman, J. (2010), Course allocation by proxy auction, *in* A. Saberi, ed., 'Internet and Network Economics', Vol. 6484 of *Lecture Notes in Computer Science*, Springer Berlin / Heidelberg, pp. 551–558.

Lee, S. M. and Clayton, E. R. (1972), 'A goal programming model for academic resource allocation', *Manag. Sci.* **18**, B395–B408.

Manlove, D. F. and O'Malley, G. (2008), 'Student-project allocation with preferences over projects', *Journal of Discrete Algorithms* **6**(4), 553 – 560.

Martello, S. and Toth, P. (1981), 'Heuristic algorithms for the multiple knapsack problem', *Computing* **27**(2), 93–112.

Martello, S. and Toth, P. (1990), *Knapsack problems: Algorithms and computer implementations*, John Wiley and Sons, Chichester.

Martens, M. and Skutella, M. (2006), 'Flows on few paths: Algorithms and lower bounds', *Networks* **48**(2), 68–76.

Miyaji, I., Ohno, K. and Mine, H. (1987), ‘Solution method for partitioning students into groups’, *Eur. J. Oper. Res.* **33**, 82–90.

Othman, A., Sandholm, T. and Budish, E. (2010), Finding approximate competitive equilibria: efficient and fair course allocation, *in* ‘Proceedings of the 9th International Conference on Autonomous Agents and Multiagent Systems: volume 1 - Volume 1’, AAMAS ’10, International Foundation for Autonomous Agents and Multiagent Systems, Richland, SC, pp. 873–880.

Reeves, G. R. and Hickman, E. P. (1992), ‘Assigning mba students to field study project teams: A multicriteria approach’, *Interfaces* **22**, 52–58.

Saber, H. M. and Ghosh, J. B. (2001), ‘Assigning students to academic majors’, *Omega* **29**, 513–523.

Sanchez-Anguix, V., Chalumuri, R., Aydoğan, R. and Julian, V. (2019), ‘A near pareto optimal approach to student–supervisor allocation with two sided preferences and workload balance’, *Applied Soft Computing* **76**, 1–15.

Santarisi, N. S. and Salhieh, S. M. (2005), ‘An experimental method for solving the problem of academic course scheduling’, *Engineering Sciences* **16**, 35–50.

Schniederjans, M. J. and Kim, G. C. (1987), ‘A goal programming model to optimize departmental preference in course assignments’, *Comput. Oper. Res.* **14**(2), 87 – 96.

Sönmez, T. and Ünver, M. U. (2010), ‘Course bidding at business schools’, *International Economic Review* **51**(51), 99–123.

Toroslu, I. H. and Arslanoglu, Y. (2007), ‘Genetic algorithm for the personnel assignment problem with multiple objectives’, *Information Sciences* **177**, 787 – 803.

Wang, F., Jagadeesan, R. and Kominers, S. D. (2019), ‘Optimizing reserves in school choice: A dynamic programming approach’, *Operations Research Letters* **47**(5), 438–446.

Weitz, R. R. and Jelassi, T. M. (1992), ‘Assigning students to groups: A multi-criteria decision support system approach’, *Decis. Sci.* **23**, 746–757.

Özdemir, M. S. and Gasimov, R. N. (2004), ‘The analytic hierarchy process and multiobjective 0-1 faculty course assignment’, *Eur. J. Oper. Res.* **157**, 398–408.

Supplementary Materials to "Fair Nomination of Exchange Students"

1 Notation and abbreviations

Notation	
$\mathcal{A}$	Set of applications with application $(s, u) \in \mathcal{A}$
$\mathcal{X}$	Set of solutions with solution $\boldsymbol{x} \in \mathcal{X}$
$\mathcal{X}_f$	Subset of solutions for fairness $f \in \{w, s\}$
$\mathcal{S}$	Set of students with student $s \in \mathcal{S}$
$\mathcal{T}_u$	Ordered list of students nominated for university $u \in \mathcal{U}$
$\mathcal{U}$	Set of universities with university $u \in \mathcal{U}$
$a_{s,u} \in [1; 5]$	Aptitude strength of application $(s, u) \in \mathcal{A}$
$\bar{a}_{s,u}$	Normalized aptitude of application (s, u)
c_u	Capacity of university $u \in \mathcal{U}$
e_s	Number of exchange semesters student $s \in \mathcal{S}$ has spent abroad prior to the application process
f	Fairness, e.g. $f = W$ refers to the weak pairwise fairness
g_s	Grade point average of student $s \in \mathcal{S}$
m	Objective, e.g. $m = N$ refers to the nomination-based objective
n_{SPF}	Maximum numbers of nominations under SPF
n_{WPF}	Maximum numbers of nominations under WPF
$\bar{n}$	Normalization constant for the number of applicants
$r_{s,u}$	Ranks corresponding to application $(s, u) \in \mathcal{A}$
$\bar{r}_{s,u}$	Normalized rank of application (s, u)
$x_{s,u}$	1 if application (s,u) becomes a nomination, 0 otherwise
$z_m(x)$	Objective function value for objective m given solution x
z_A	Aptitude-oriented objective
z_N	Nomination-oriented objective
z_R	Rank-oriented objective
Abbreviations	
GPA	Grade point average
MIP	Mixed-integer program
MKP	Multiple Knapsack Problem
PF	Pairwise fairness
SPF	Strong pairwise fairness
WPF	Weak pairwise fairness

2 Recalculation and normalization of the objective function weights

Let $\bar{n} = \frac{1}{|\mathcal{S}|}$ and $r^{\max} = \max\limits_{(s,u)\in\mathcal{A}} r_{s,u}$. Each application's normalized rank $\bar{r}_{s,u} \quad \forall (s, u) \in \mathcal{A}$ is determined by:

$$\bar{r}_{s,u} = \frac{1 + r^{\max} - r_{s,u}}{r^{\max} \cdot |\mathcal{S}|}. \tag{11}$$

Similarly, let $a^{\max} = \max_{(s,u)\in\mathcal{A}} a_{s,u}$. Then, each application's score $\bar{a}_{s,u} \quad \forall (s,u) \in \mathcal{A}$ is determined by:

$$\bar{a}_{s,u} = \frac{1 + a^{\max} - a_{s,u}}{a^{\max} \cdot |\mathcal{S}|}. \tag{12}$$

This ensures that each objective function z_N, z_R or z_A's values are within interval $[0, 1]$.

3 Results for setups 2, 4, 5, 7 and 8

The computational results for setups 2, 4, 5, 7 and 8 are provided by Tables 15, 16, 17, 18, and 19, respectively.

Table 15: Results of the MIP applied to setup 2 (nomination-oriented objective with weak and without strong pairwise fairness constraints)

Application year	1	2	3	4	5	6	7
# Nominations	53	61	71	93	122	151	340
Percentage of applying students	88.3	77.2	78.9	79.5	80.8	66.8	65.3
Nominations for rank 1 universities	31	32	44	56	78	79	182
Percentage of nominations	58.5	52.5	62.0	60.2	63.9	52.3	53.5
Nominations for rank 2 universities	12	18	18	22	28	44	79
Percentage of nominations	22.6	29.5	25.4	23.7	23.0	29.1	23.2
Nominations for rank 3 universities	10	11	9	15	16	28	79
Percentage of nominations	18.9	18.0	12.7	16.1	13.1	18.5	23.2
Average aptitude of nominated students	0.44	0.56	0.53	0.51	0.77	0.58	0.48
Violations of exchange semester- and rank-/aptitude-based pairwise fairness	0	0	0	0	0	0	0
Violations of weak pairwise fairness	0	0	0	0	0	0	0
Violations of strong pairwise fairness	0	1	0	6	1	6	7

Table 16: Results of the MIP applied to setup 4 (rank-oriented objective without weak or strong pairwise fairness constraints)

Application year	1	2	3	4	5	6	7
# Nominations	53	60	70	92	121	148	329
Percentage of applying students	88.3	75.9	77.8	78.6	80.1	65.5	63.1
Percentage of nominations compared to the heuristic	103.9	101.7	101.4	102.2	100.8	101.4	103.5
Percentage of rank 1 universities	67.9	70.0	71.4	72.8	76.9	66.9	69.0
Percentage of rank 2 universities	17.0	20.0	22.9	20.7	17.4	24.3	19.1
Percentage of rank 3 universities	15.1	10.0	5.7	6.5	5.8	8.8	11.9
Average aptitude of nominated students	2.87	2.30	2.28	2.19	2.30	2.19	2.31
Violations of exchange semester- and rank-/aptitude-based pairwise fairness	0	0	0	0	0	0	0
Violations of weak pairwise fairness	4	13	8	15	16	64	128
Violations of strong pairwise fairness	5	16	8	21	19	64	144

Table 17: Results of the MIP applied to setup 5 (rank-oriented objective with weak and without strong pairwise fairness constraints)

Application year	1	2	3	4	5	6	7
# Nominations	53	61	71	93	122	150	335
Percentage of applying students	88.3	77.2	78.9	79.5	80.8	66.4	64.3
Nominations for rank 1 universities	33	37	47	60	85	84	191
Percentage of nominations	62.3	60.7	66.2	64.5	69.7	56.0	57.0
Nominations for rank 2 universities	11	14	19	22	25	44	84
Percentage of nominations	20.8	23.0	26.8	23.7	20.5	29.3	25.1
Nominations for rank 3 universities	9	10	5	11	12	22	60
Percentage of nominations	17.0	16.4	7.0	11.8	9.8	14.7	17.9
Violations of exchange semester- and rank-/aptitude-based pairwise fairness	0	0	0	0	0	0	0
Average aptitude of nominated students	0.44	0.56	0.53	0.51	0.77	0.58	0.48
Violations of weak pairwise fairness	0	0	0	0	0	0	0
Violations of strong pairwise fairness	0	2	0	2	3	3	6

Table 18: Results of the MIP applied to setup 7 (aptitude-oriented objective without weak or strong pairwise fairness constraints)

Application year	1	2	3	4	5	6	7
# Nominations	52	60	70	92	122	151	299
Percentage of applying students	86.7	75.9	77.8	78.6	80.8	66.8	57.4
Nominations for rank 1 universities	29	34	45	55	79	80	161
Percentage of nominations	55.8	56.7	64.3	59.8	64.8	53.0	53.8
Nominations for rank 2 universities	12	15	15	23	28	42	70
Percentage of nominations	23.1	25.0	21.4	25.0	23.0	27.8	23.4
Nominations for rank 3 universities	11	11	10	14	15	29	68
Percentage of nominations	21.2	18.3	14.3	15.2	12.3	19.2	22.7
Average aptitude of nominated students	0.45	0.57	0.54	0.52	0.77	0.58	0.67
Violations of exchange semester- and rank-/aptitude-based pairwise fairness	0	0	0	0	0	0	0
Violations of weak pairwise fairness	0	0	0	0	0	0	2
Violations of strong pairwise fairness	2	1	2	2	2	6	10

Table 19: Results of the MIP applied to setup 8 (aptitude-oriented objective with weak and without strong pairwise fairness constraints)

Application year	1	2	3	4	5	6	7
# Nominations	52	60	70	92	122	151	299
Percentage of applying students	86.7	75.9	77.8	78.6	80.8	66.8	57.4
Nominations for rank 1 universities	29	33	45	56	78	80	163
Percentage of nominations	55.8	55.0	64.3	60.9	63.9	53.0	54.5
Nominations for rank 2 universities	12	16	15	21	29	42	71
Percentage of nominations	23.1	26.7	21.4	22.8	23.8	27.8	23.7
Nominations for rank 3 universities	11	11	10	15	15	29	65
Percentage of nominations	21.2	18.3	14.3	16.3	12.3	19.2	21.7
Average aptitude of nominated students	0.45	0.57	0.54	0.52	0.77	0.58	0.67
Violations of weak pairwise fairness	0	0	0	0	0	0	0
Violations of strong pairwise fairness	2	1	1	6	2	6	6