Ties matter: improving efficiency in course allocation by introducing ties

Chen, Ning and Li, Mengling

Nanyang Technological University, Singapore

25 April 2013

Ties Matter: Improving Efficiency in Course Allocation by Introducing Ties

By Ning Chen and Mengling Li*

We study the course allocation system at Nanyang Technological University, where students submit strict preferences for courses and courses have implicit preferences for students. This formulates a many-to-many matching problem. We show the inefficiencies of the current mechanism and propose new competing mechanisms called Pareto-improving draft and dictatorship mechanisms, which introduce ties into students' preferences. Our mechanisms generate (group) stable and Pareto-efficient allocations, and the dictatorship mechanism can be implemented truthfully. Simulations on real data show that introducing ties into students' preferences can significantly improve efficiency, and the draft mechanism outperforms the dictatorship mechanism despite that the former is non-strategyproof.

JEL: C78, D82, I23

Course allocation is a classic many-to-many matching problem in which a set of courses are allocated to a set of students who have multiunit course demands. Allocating courses equitably and efficiently has proven a challenging market design problem and there has been little in the literature to address it due to a variety of difficulties. First, students do not have proprietary rights over courses; therefore, buying or selling courses through money transfer is strictly prohibited. Second, the use of an unauthorized computer program to gain unfair advantage over other students in securing courses is unacceptable. Third, students may have preferences not only for individual courses but also for combinations of courses, and these preferences may include ties (i.e., indifferences); thus, they may exhibit

* Chen: Division of Mathematical Sciences, School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore (e-mail: ningc@ntu.edu.sg); Li: Division of Mathematical Sciences, School of Physical and Mathematical Sciences, Nanyang Technological University, Singapore (e-mail: mli6@ntu.edu.sg).

complicated strategic behavior. Finally, because most applications involve a large number of students and courses under a strict deadline to produce a feasible allocation, an efficient computation is critical. Consequently, the course allocation problem and the more general many-to-many matching problems are substantially more difficult to address than one-to-one and many-to-one allocation problems.[1]

In practice, there are two main types of course allocation mechanisms employed in educational institutions.

- Preference-ranking mechanisms, e.g., the draft mechanism at Harvard Business School (HBS), in which students submit ordinal preferences for courses. Budish and Cantillon (2012) showed that the draft mechanism is manipulable in theory and manipulated in practice; but, interestingly, the mechanism outperforms the strategyproof alternative, which implies that strategyproofness has both benefits and costs. The authors further proposed a proxy draft mechanism that is shown to generate better efficiency. Later, Kominers et al. (2010) proposed a new proxy mechanism that simplifies students' strategic decision by directly incorporating their manipulation strategy into the mechanism. The mechanism is Pareto-efficient and resistant to strategic manipulations observed in the extant data. However, other unobserved manipulations may still exist.
- Bidding mechanisms, e.g., the mechanism at the Ross School of Business at the University of Michigan (UMBS), in which students bid for courses. In the UMBS mechanism, bids submitted by students play a dual role—to infer the preferences of both students and courses. These two roles can easily conflict and result in unnecessary efficiency loss. Sönmez and Ünver (2010) proposed an alternative Gale-Shapley Pareto-dominant mechanism that asks students to submit their preferences for courses in addition to bids. The mechanism is confirmed to have superior efficiency in both field and laboratory studies by Krishna and Ünver (2008). However, the mechanism is not strategyproof, which can prompt additional concerns about efficiency loss.

[1]The generalization to multiunit demand on both sides is nontrivial as the properties and structures of many-to-many matchings behave rather differently from one-to-one and many-to-one matchings (Roth and Sotomayor 1990, Sotomayor 1999, Echenique and Oviedo 2006). Further, the presence of a few many-to-many demands may change the matching for all agents completely (see Example 2.2 in Echenique and Oviedo (2006)). There have been a growing number of instances of many-to-many matching markets in recent years, such as social lending markets (Chen and Ghosh 2011) and online labor markets. It is therefore important to study the general many-to-many matching models to understand these marketplaces.

Note that while the common objectives of course allocation mechanisms are efficiency and equity, the practice of using different mechanisms in different educational institutions indicates that the current mechanisms are neither well understood nor satisfactory.

In this paper, we examine the course allocation system for around 13,000 undergraduate students at Nanyang Technology University (NTU) in Singapore and explore its potential improvement. In NTU's current mechanism, students submit strict *preferences* for individual courses in different categories and courses have predetermined *preferences* (with ties) for students, which are essentially the *priorities* of each student. A centralized mechanism then determines allocations by considering student-course pairs in an order based on the priorities of students and their course preferences. This is a preference-ranking mechanism and is similar to the Boston Student Assignment Mechanism (Abdulkadiroğlu and Sönmez 2003). The details of the mechanism are deferred to the subsequent section.

The NTU's mechanism does exhibit some nice properties. An allocation generated through this mechanism is both *pairwise stable* (i.e., there is no student-course blocking pair) and students-sided *group stable* (i.e., there is no group of students such that each of them can strictly improve his/her assignment by swapping courses among themselves). Note that the latter is a simple implication of the strict preferences of students and a homogeneity property exhibited by the students' priorities. However, the mechanism is not strategyproof in theory and is manipulated in practice.[2] Furthermore, breaking ties arbitrarily at random (as courses' preferences for students have ties) may result in severe efficiency loss. A similar issue was recently addressed for many-to-one matchings (Erdil and Ergin 2006, Erdil and Ergin 2008, Abdulkadiroğlu et al. 2009).

A critical ingredient in NTU's mechanism is that students are enforced to submit *strict* preferences. Indeed, much of the literature on many-to-many matching markets assumes strict preferences when studying solution concepts.[3] In prac-

[2] As shown in a field survey of over 1,200 students at NTU (see Figure 1), 19 percent indicate that they would manipulate their preferences.

[3] Roth (1984) showed that the set of pairwise stable matchings is nonempty with substitutable preferences, and that one-sided optimal stable matchings exist. Blair (1988) proved that the set of pairwise stable matchings forms a lattice structure; its properties have been investigated by Alkan (2001) and Alkan (2002). Martínez et al. (2004) presented an algorithm that finds all pairwise stable matchings. A pairwise stable matching, however, need not be group stable and need not even be Pareto-optimal with responsive preferences (Roth and Sotomayor 1990). Further, Sotomayor (2011) presented examples

tice, however, there are various matching markets in which agents are not able to strictly rank their prospective partners for a variety of reasons (e.g., incomplete information). Our survey data (see Figure 1) show that in the course allocation at NTU, 76 percent of students prefer to have ties in their preference lists. This motivates us to design a course allocation mechanism for many-to-many matching markets that includes *ties* in both courses' and students' preferences, in the hope of satisfying more students' demand and increasing overall efficiency.

Before designing our mechanism, we examine which solutions are desired in course allocation. On the one hand, we still hope to keep pairwise and student-sided stability to capture fairness by taking students' preferences and priorities into account. On the other hand, we require that the allocation be *Pareto-efficient*, which to some extent qualifies the overall efficiency of an allocation. Pairwise stability and Pareto efficiency (together called *Pareto stability* as suggested by Sotomayor (2011)), provide a natural solution benchmark for matching markets in the presence of ties. Note that Pareto stability and one-sided group stability generally cannot coexist (Claim 1). However, in our application, with the homogeneity property of the students' priorities, we show that Pareto stability implies one-sided group stability (Proposition 2).

The question then becomes whether a Pareto-stable matching always exists, and how to find one efficiently. Introducing ties into preferences results in dramatic changes to the properties of stable matchings. In particular, *stability no longer guarantees Pareto efficiency.*[4] The question of finding a Pareto-stable matching has recently been addressed by Erdil and Ergin (2006, 2008) for many-to-one matchings. They presented a stability-preserving Pareto-improving algorithm by eliminating augmenting paths/cycles. However, their algorithm fails to work for many-to-many matchings, in which Pareto improvement does not necessarily preserve stability (see Example 5).

We construct an efficient algorithm that computes a Pareto-stable many-to-many matching. Our result immediately implies the existence of such a match-

in which the set of group stable matchings is empty. Echenique and Oviedo (2006) later gave conditions under which the group stable matchings set is nonempty and can be approached through an algorithm.

[4]For example, there are two men m_1, m_2 and two women w_1, w_2, where m_1 strictly prefers w_1 to w_2, but all others are indifferent amongst their possible partners. The matching $(m_1, w_2), (m_2, w_1)$ is stable, but not Pareto-efficient because m_1 can be reassigned to w_1 and m_2 to w_2 without negative effect.

ing. Our algorithm works for general arbitrary preferences for individuals in the presence of ties. The algorithm, from a high-level overview, builds on the idea of Roth and Vande Vate (1990), who provided an alternative to the deferred acceptance algorithm that computes a stable (one-to-one) matching. The details of the algorithm are presented in Section IV.

Our algorithm, like those of Erdil and Ergin (2006, 2008), is a Pareto-stable matching mechanism, which is generally not strategyproof (see Example 2 of Erdil and Ergin (2008)). In the course allocation applications, we construct two additional competing allocation mechanisms: the *Pareto-improving draft mechanism* and the *Pareto-improving serial dictatorship mechanism*, based on the homogeneity of students' priorities. Both mechanisms compute a Pareto-stable matching and the latter can be implemented in a truthful manner, which makes it a dominant strategy for students to report their true preferences. This significantly simplifies the strategic consideration of students. The truthful mechanism is based on the combination of the random serial dictatorship mechanism and augmenting paths/cycles elimination. While the dictatorship mechanism exhibits the property of strategyproofness, it brings out some potential fairness issues by giving too much priority to the lucky students who get a high random order while callously disregarding the preferences of those who belong to the same priority group but receive a low random order. Therefore, with some simple welfare measures, including the average rank of assigned courses and the total number of unassigned students, we compare the performance of the Pareto-improving draft and dictatorship mechanisms using real course registration data.

To quantify efficiency improvement that allows ties in students' preferences and compare the two competing allocation mechanisms, simulations are performed using real course registration data from three academic years at NTU: 2010, 2011 and 2012 at NTU. A survey has revealed that a majority of students prefer to have two or three levels of preference. Given the strict preferences submitted by students, we therefore divide the students' preferences into two or three levels at random and assume that students are indifferent between courses at the same level. We run simulations in different environments introducing ties and for different categories of courses using the draft and dictatorship mechanisms. The simulation results show that (1) the efficiency in terms of total number of allocations and total number of unassigned students can be significantly improved by introducing ties into students' preferences; (2) the draft mechanism outperforms the dictatorship mechanism in terms of the total number of allocations, the average rank and the total number of unassigned students, even though it is not strategyproof.

Organization. The remainder of this paper is organized as follows. Section I describes NTU's current course allocation system and mechanism. Section II formally defines the many-to-many matching model with ties. Section III discusses the solution concepts considered in this paper. The algorithm that computes a Pareto-stable matching in general frameworks and two specific mechanisms for the course allocation problem are described in Section IV. The simulation results and discussions are presented in Section V. We conclude our work in Section VI.

I. NTU Course Allocation: Matching with Preferences

In course allocation, while it is natural to assume that students have preferences for courses, educational institutes usually decide a priority ranking for individual students for each course. For instance, a course may take prior considerations for final year graduating students who need the course to fulfill graduation requirements, or those who failed the course in the preceding semesters and need to take the course again in the current semester. The priority ranking can be considered a course's implicit preference for students.

The current course allocation system at NTU solicits preferences from both students and courses to decide allocations. Specifically, in addition to the major core courses required by each department, there are two types of general education requirements that are open to students from all departments: Prescribed Electives (PE) and Unrestricted Electives (UE).[5] Students submit up to five courses under a *strict* order as preferences for both PE and UE, respectively, to the system. The implicit preferences of courses for individual students (i.e., their *priorities*) are formed according to the following hierarchy (from highest to lowest):

[5]The detailed description of the curriculum structure of NTU is deferred to Appendix A.

1)
   - Students with only PE courses remaining to fulfill, if the course is a PE type.
   - Students with only UE courses remaining to fulfill, if the course is a UE type.
   - Students with only PE or UE courses remaining to fulfill, if the course is both a PE and UE type.
2) Final year students.
3) Students of special programs (e.g., accelerated bachelor degree).
4) Year 3 students.
5) Year 2 students.

Note that some of the students are at the same level on the priority lists; that is, there are *ties* (i.e., indifferences) in the courses' preference lists. In addition, it can be seen that all of the courses have the same preferences for students (called *homogeneous*). In particular, all of the courses (PE or UE) can list all those students who have only PE or UE courses to fulfill in the first priority. (This is because a student with only PE/UE courses to fulfill will not submit any UE/PE courses, respectively. The individual rationality property of feasible assignments allows us to unify the preferences of courses in such a way.)

PROPOSITION 1: *The preferences of courses for students are homogeneous.*[6]

Each course has a prespecified capacity constraint due to resource limitations, e.g., the size of a classroom. In addition, every student has a capacity constraint as well, which sets an upper bound on the number of registered courses allowed. Specifically, except for final year graduating students, a student will be allocated, at most, one course for PE and one course for UE. In addition, for all students (including those in their final year), a maximum of 24 academic units (including major core courses) can be registered within one semester, which corresponds to about 7 courses.

Given the capacity constraints and preferences of both students and courses, the system allocates the courses to the students using a mechanism, which runs

[6]Precisely, some courses may have special predetermined preferences for students (e.g., the course Economics of Manufacturing takes prior consideration of students majoring in Material Sciences and Engineering). In such cases, some vacancies are reserved in advance by the corresponding departments, and the preferences of courses are still homogeneous in the system.

for PE and UE courses separately and sequentially, first for PE, followed by UE. (If a course is in both PE and UE, its capacity will be predivided for PE and UE, respectively, by the administration.) The mechanism is described as follows.

```
NTU mechanism
```

```
Consider students priority by priority.  For all students that are at the same
priority level, consider courses according to their preferences.  That is, for
k = 1,...,5,

  • Consider all student-course pairs in which the student is at the considered
    priority level and the course is as his/her k-th choice, and assign courses
    to students amongst these pairs subject to the capacity constraints (with ties
    broken randomly).
```

The mechanism is simple to implement and quite similar to the Boston Student Assignment Mechanism as described by Abdulkadiroğlu and Sönmez (2003). The main difference is that the Boston Student Assignment Mechanism runs simultaneously for all students instead of running separately for students in different priority groups. A serious shortcoming of the Boston Student Assignment Mechanism is that students with high priorities at specific schools lose their priorities unless they list these schools as their top choices. As a consequence, truth-telling is not a dominant strategy and students can easily manipulate their preferences. Indeed, in a recent field survey (see Figure 1) of NTU's current course allocation system, from a pool of over 1,200 participating students, 19 percent indicated that they did not place course preferences truthfully. Manipulation is a key challenge to the mechanism, and may result in severe efficiency loss.

A. *Inefficiency of the Mechanism*

While the NTU mechanism guarantees fairness to some extent by considering the preferences of students and courses, there are a large number of students who are not assigned any course due to competitiveness and the mechanism, which results in considerable efficiency loss (see Table 1 for the data summary of course allocation statistics for the three academic years: 2010, 2011 and 2012). Note that the allocation statistics were collected after certain manual adjustments, which have already corrected some of the inefficiency issues in the results directly from the mechanism. Inefficient course allocations result in appeals and complaints

Table 1—Course Registration Statistics for Semester 1 of Academic Year 2010-2012.

	Course	Vacancies	Students	Allocations	Percentage
2010	PE	8,118	9,660	6,844	84.31
	UE	14,178	10,672	8,366	59.01
	PE+UE	22,296	13,471	15,210	68.22
2011	PE	10,092	8,546	6,919	68.56
	UE	15,342	8,836	9,128	59.50
	PE+UE	25,434	11,321	16,047	63.09
2012	PE	9,305	8,157	6,729	72.32
	UE	16,260	8,731	7,170	44.10
	PE+UE	25,565	11,541	13,899	54.37

Percentage is measured as the total number of allocations over the total vacancies. For instance, for the 2010 data, the PEs and UEs have 8,118 and 14,178 vacancies while 9,660 and 10,672 students demand them, respectively. The data indicate that the current systems resulted in an allocation of 6,844 PEs and 8,366 UEs. About 29.15 percent and 21.61 percent of the students who requested PEs and UEs failed to get any allocation, respectively. Note that in the 2011 data, the number of UE allocations is larger than that of students, because a considerable amount of students get more than one UE registered in the process of manual adjustments after the implementation of the mechanism.

submitted by unsatisfied students, and can even cause the deferral of graduation for final year students. Due to the inefficiency of the NTU mechanism, many unsatisfied students and program coordinators spend a tremendous amount of time and effort manually seeking better allocations. Despite this effort, some students can still end up with an unfavorable outcome in which no course is registered.

Our survey statistics (Figure 1) show that 65 percent of students find the current system unsatisfactory. Another important phenomenon illustrated in the survey is that 76 percent of the students think that they should be allowed to indicate ties (i.e., indifferences) in their course preferences. In practice, while students may have strict preferences for some courses, such strictness is not very sensitive in the sense that most students are usually only concerned with whether one of his/her desired courses is registered, but not exactly which one. In other words, it is reasonable to assume that students' preferences should have ties. As a consequence, we consider introducing ties into students' preferences for courses, with the objective of reducing the flaws in the current system and improving overall efficiency for students.

1) Are you satisfied with the current course allocation system?
   - Yes --- 35 percent
   - No --- 65 percent

2) Are you *indifferent* between two or more courses that you are interested in (i.e., among these courses, you do not strictly prefer one to another)?
   - Yes --- 76 percent
   - No --- 24 percent

3) If your answer is "Yes" to the previous question, how many levels would like to have in your preference structure (assuming you are indifferent between the courses in the same level)?
   - One level (all indifferent) --- 10 percent
   - Two levels --- 35 percent
   - Three levels --- 23 percent
   - Four levels --- 8 percent

4) What is the number of courses you would like to put on your preference list?
   - Five --- 75 percent
   - Ten --- 13 percent
   - Other --- 12 percent

5) Do you submit courses to the system according to your true preferences?
   - Yes --- 81 percent
   - No --- 19 percent

Figure 1. Survey Questions and Results for 1,200 students at NTU.

II. Model

Given the fact that both students and courses have capacities and incomplete preferences with ties, we consider a two-sided many-to-many matching model with a set of men M and a set of women W.[7] Throughout this paper, we use $m \in M$ to denote a man, $w \in W$ to denote a woman and $x, y, z \in M \cup W$ to denote any individual agent (man or woman). For each agent $x \in M \cup W$, let $c_x \in \mathbb{N}$ be his/her capacity, which is the maximum number of agents on the opposite side that can be matched to x. The presence of capacities allows us to assume, without loss of generality, that $|M| = |W| = n$, as dummy agents with $c_x = 0$ can be added to the market.

Each man $m \in M$ has a *preference* list P_m ranking individual women, denoted

[7] As our model and technical results work for more general settings, we use the terms 'men' and 'women', following the seminal work of Gale and Shapley (1962), to describe the model. In our application, men can represent students and women can represent courses, or vice versa.

by $\succ$ and $=$, where $w_1 \succ w_2$ means that m strictly prefers w_1 to w_2, and $w_1 = w_2$ means that m is indifferent between them[8]. We say m weakly prefers w_1 to w_2 if either $w_1 \succ w_2$ or $w_1 = w_2$, denoted by $w_1 \succeq w_2$. Every two women in P_m are comparable and the preference is assumed to be transitive. The preference P_m gives a partial list of individual women that are acceptable to m (i.e., m does not want to be matched with any woman that is not on the list). For example, a possible preference list for m is $P_m : (w_1 = w_2 \succ w_3 = w_5)$. Here, m is indifferent between w_1 and w_2, prefers either of them to w_3 and w_5, amongst which m is indifferent, and finds all other women unacceptable. The preference list P_w for each woman $w \in W$ is defined similarly.[9] We use $E = \{(m, w) \mid m \in P_w, w \in P_m\}$ to denote the set of mutually acceptable pairs.

Note that the preference lists P_m and P_w as defined above are for individuals. Because agents can have capacities greater than one, we also need to define preferences for subsets. Considering our motivating application course allocation, from the viewpoint of courses, they only care about the interests of individual students and do not necessarily have preferences over subsets of students. (In particular, the preferences of courses for individual students are according to the students' need to take the courses, rather than their identities and scores.) Students in NTU's current system are only allowed to submit preferences for individual courses. Indeed, with a large number of students and courses, specifying preferences for subsets results in complicated comparisons of alternatives and lengthy preference lists, yielding an inapplicable system. In addition, while students may have preferences for subsets of courses, this mostly occurs among those with prerequisite relations, which cannot be registered in the same semester. For example, one must take the course 'Microeconomic Principles' before being allowed to take the course 'Industrial Organization'. Within one semester the main concern of every student is thus individual courses.

Therefore, in our study we assume that preferences are *responsive with a lattice structure*. That is, given a subset $S \subseteq W$ and two women $w, w' \notin S$, m prefers

[8]Precisely, there is a subscription m in the notations, i.e., $\succ_m$ and $=_m$. For simplicity we omit the subscription when it is clear from the context.

[9]In the application of course allocation at NTU, students submit two preference lists, one for PE and one for UE. Here, for simplicity, we assume that each individual has only one preference for the other side. Our technical results can be easily generalized to the settings in which one has multiple preferences. Details are provided in Section IV.C.

$S \cup \{w\}$ to $S \cup \{w'\}$ if and only if m prefers w to w'. (Note that it is allowed that w or $w' = \emptyset$; in particular, m prefers $\emptyset$ to any woman $w \notin P_m$, prefers any woman $w \in P_m$ to $\emptyset$, and is indifferent between any $w, w' \notin P_m$.) In other words, for any two sets that differ in only one woman, m prefers one that contains the more preferred woman and is indifferent between them if he is indifferent between the two women. The preferences of women are defined similarly. In addition, the responsive preference is transitive, i.e., if m prefers S_1 to S_2 and S_2 to S_3, he also prefers S_1 to S_3. Note that this preference for subsets only constitutes a partial order, precisely, a complete distributive lattice. That is, two alternatives are comparable only if they are an ancestor-descendant relation in the lattice.[10] In Appendix B, we give an example to show responsive preference with a lattice structure. This model of preferences with multiunit capacity has been used in, e.g., Erdil and Ergin (2006). It is simple, because agents only need to express preferences for individuals, and is arguably natural for settings in which the benefit from a partner to an agent does not depend on the agent's remaining partners.

Given the preferences of all of the agents, our objective is to establish a multiunit pairing between men and women, called an *assignment* (or a many-to-many matching). An assignment is denoted by $\mu = (\mu_{mw})_{m \in M, w \in W}$, where $\mu_{mw} = 1$ means that m and w are matched and $\mu_{mw} = 0$ otherwise. A *feasible* assignment is one that satisfies the following conditions: $\sum_w \mu_{mw} \leq c_m$ and $\sum_m \mu_{mw} \leq c_w$, and $\mu_{mw} = 1$ only if m and w are mutually acceptable. All of the assignments considered in this paper are feasible. Further, for any $x \in M \cup W$, we denote by $\mu(x)$ the set of individuals matched to x in the assignment μ.

III. Solution Concepts

Given the many-to-many matching model defined above, the next question is which assignment is desired. While we cannot satisfy the demands of all students, we hope to provide a solution that is in favor of most students. In this section, we examine a number of solution concepts from different perspectives in course allocation.

[10]Note that the responsive preference defined in Roth and Sotomayor (1990) is a complete ranking for all compatible subsets, unlike ours, which forms a lattice structure.

A. *Pairwise Stability*

In many two-sided matching models, e.g., student placement and school choice (Sönmez and Ünver 2011), one desired property of an allocation is the *elimination of justified envy.* That is, whenever a student m prefers the allocation of another student m', m should not rank higher than m' in the priority list of courses. Consider the following example.

EXAMPLE 1: *There are two students* m_1, m_2 *and two courses* w_1, w_2 *each with unit capacity. Their preferences are shown on the left-hand side below.*

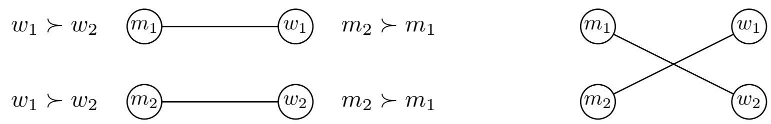

In the first allocation, m_2 *is not assigned to his first choice* w_1*, which lists him with a higher priority than* m_1*. In such a case,* m_2 *can simply ask:*

> *"I am more eligible for the course; why was I not assigned?"*

Hence, a fairer solution is to allocate w_1 *to* m_2 *(see the above figure on the right-hand side). Now* m_1 *is not assigned to his first choice* w_1*, but he can use the following reasoning:*

> *"I did not get the course, but all those who got it are (more) eligible."*

Equivalently, if a course (i.e., w_2*) has not been assigned to a student (i.e.,* m_2*) who has a higher priority, then the student should be assigned to another course that he prefers (more).*

The issues illustrated by the example above are captured by stability, a solution concept first proposed by Gale and Shapley (1962) in the application of marriage markets and college admission. The formal definition of stability in our many-to-many matching model is as follows.

DEFINITION 1 (Pairwise stability): *We say that a feasible assignment* $\mu = (\mu_{mw})$ *is (pairwise)* stable *if there is no mutually accepted pair* $(m, w) \in E$ *(called a blocking pair) with* $\mu_{mw} = 0$*, satisfying one of the following conditions:*

- *Both m and w have leftover capacity;*
- *m has leftover capacity and there is m', $\mu_{m'w} = 1$, such that w strictly prefers m to m'; or w has capacity remaining and there is w', $\mu_{mw'} = 1$, such that m prefers w to w';*
- *There are m' and w', $\mu_{mw'} = 1$ and $\mu_{m'w} = 1$, such that m strictly prefers w to w' and w strictly prefers m to m'.*

In course allocation, the stability concept ensures a certain level of fairness among students in the sense that no blocking pairs can upset the structure of a matching. It creates a balance in the competition among students and their priorities in each course.

Note that both members of a blocking pair are able to strictly improve their assignments respectively by matching with each other (and possibly breaking some of the current assignments). A stable assignment always exists, and can be found using a variant of Gale-Shapley's deferred acceptance algorithm (Gale and Shapley 1962) for computing one-to-one stable matchings (e.g., by making c_x copies for each individual $x \in M \cup W$ with the same preference list and breaking ties randomly).

B. Pareto Efficiency

Another important criterion in our application is social welfare, which measures the overall efficiency of an allocation. Consider the following two examples.

EXAMPLE 2: *There are two students m_1, m_2 and two courses w_1, w_2 each with unit capacity. Their preferences are shown on the left-hand side below.*

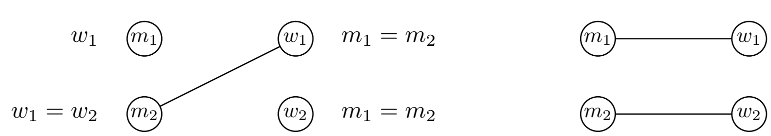

The first allocation in which only m_2 is assigned to a course w_1 is stable. However, a more efficient allocation is on the right-hand side, where both m_1 and m_2 are assigned to a course. Note that the second allocation is also stable.

EXAMPLE 3: *There are two students m_1, m_2 and two courses w_1, w_2 each with unit capacity. Their preferences are shown on the left-hand side below.*

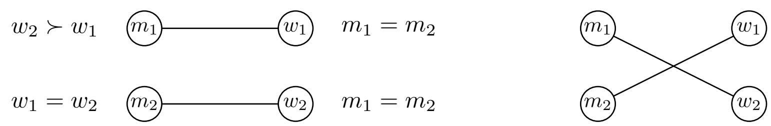

While both of the allocations are stable, the second is more desirable in the sense that student m_1 is assigned to a course he likes more and the assignments of other individuals remain at the same preference level.

In the two examples above, the second allocation dominates the first one in the sense that someone's allocation is strictly improved while no one is worse off. This property is captured by the notion of Pareto efficiency. An allocation that is not Pareto-efficient implies that certain changes in allocations may result in some individuals being better off with no individual being worse off, therefore leading to an efficiency improvement. The formal definition is given below.

DEFINITION 2 (Pareto efficiency): *Given a feasible assignment $\mu = (\mu_{mw})$, we say that $\mu' = (\mu'_{mw})$ is a* Pareto improvement *of μ if for all $x \in M \cup W$, x weakly prefers $\mu'(x)$ to $\mu(x)$, and the preference is strict for at least one agent. An assignment μ is called* Pareto-efficient *if it does not have any Pareto improvement.*

Note that in the original Gale-Shapley stable marriage model with strict preferences, stability implies Pareto efficiency (Roth and Sotomayor 1990). However, when indifferences (ties) are allowed, stability no longer guarantees Pareto efficiency. Further, even if preferences are strict, in many-to-many matching models with responsive preferences, a stable matching need not be Pareto-efficient (Roth and Sotomayor 1990).

This definition is *two-sided* Pareto efficiency, i.e., it considers the social welfare of both students and courses. Another related notion is *one-sided* Pareto efficiency, i.e., it only considers the social welfare of students. In course allocation and general many-to-many matchings, both definitions are reasonable and may find their applications.[11] In the following, we discuss the existence and computa-

[11]In the course allocation at NTU, the priorities of courses mainly reflect students' need to take the courses; thus, student-sided Pareto efficiency may better capture social efficiency. However, in the

tion of solutions with respect to both definitions.

C. Group Stability

The notion of pairwise stability described above captures fairness in terms of a pair of student and course. However, students may still end up with an unsatisfying assignment. Consider the following example.

EXAMPLE 4: *There are two students* m_1, m_2 *and two courses* w_1, w_2 *each with unit capacity. Their preferences are shown on the left-hand side below.*

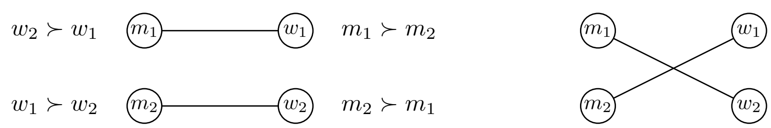

The first allocation is stable and Pareto-efficient. Note that since m_1 *prefers* m_2*'s assignment and* m_2 *prefers* m_1*'s assignment, they can swap courses with each other, resulting in the second allocation. Although the swapping is not a Pareto improvement, it improves both students' welfare. Therefore, given that students' welfare is our main consideration, the second allocation is more desirable.*

To capture the issue illustrated by the above example, we define a property called (one-sided) group stability: if an assignment is group stable, there is no group of students such that each of them can be strictly better off through re-assignment within the group. The formal definition is as follows (recall that M denotes the set of students).

DEFINITION 3 (Group stability): *Given a feasible assignment* $\mu = (\mu_{mw})$*, we say that* μ *is blocked by a coalition* $S \subseteq M$ *if all members in* S *are able to get a strictly better assignment by swapping matchings within* S*. Formally, consider the submarket formed by* S *and* W *with capacity* $c_w - |\{m \in M \setminus S \mid w \in \mu(m)\}|$ *for each* $w \in W$*, there is a feasible assignment* μ' *such that* $\mu'(m) \succ_m \mu(m)$ *for all* $m \in S$*. We say an assignment is (one-side)* group stable *if it is not blocked by any coalition.*

applications where the preferences of courses are set by individual departments or lecturers, two-sided Pareto efficiency might be a better solution concept.

In the definition above, the capacity of each $w \in W$ is the remaining capacity of w for S. If an allocation is group stable, then there is no way for any subset of students to swap courses to improve everyone's assignment. In other words, there always exists a student in the coalition whose assignment cannot be improved; that student then has no incentive to form the coalition and the original allocation is group stable.

In many-to-one matching, it is well known that group stability is equivalent to pairwise stability given responsive preferences (Lemma 5.5, Roth and Sotomayor (1990)). (Note that the notion of group stability defined in Roth and Sotomayor (1990) considers the improvement of both sides.) However, such equivalence does not hold for one-sided group stability even in one-to-one matching, as illustrated by Example 4. Indeed, we can show the following impossibility result.

CLAIM 1: *There is an instance in which pairwise stability and one-sided group stability cannot hold simultaneously even if all individuals have unit capacity.*

Note that if a matching is student-sided Pareto-efficient, it is also students-sided group stable. Thus, the above claim also implies that *pairwise stability and one-sided Pareto efficiency cannot coexist in general*; this fact is also illustrated by Example 2.31 of Roth and Sotomayor (1990).

However, if one side has homogeneous preferences (i.e., all preferences are the same), as is the case in our course allocation application, the two stability notions can coexist. (Further, given such a condition, one-sided Pareto efficiency implies two-sided Pareto efficiency; see more discussions in Section IV.E.)

PROPOSITION 2: *In a many-to-many matching market, if one side has homogeneous preferences (i.e., courses), then a pairwise stable and two-sided Pareto-efficient assignment is group stable for the other side (i.e., students).*

D. Incentive Compatibility

Another important issue in determining an allocation is strategic considerations, i.e., whether it is a dominant strategy for individual agents to report their private preference truthfully. While it is well known that there is one-sided truthfulness for one-to-one and many-to-one matching models (Roth and Sotomayor 1990), the truthfulness only applies to the side with unit demand, i.e., the "one" side. Hence, in the general many-to-many matching model, we cannot expect to have a truthful mechanism that always generates a stable matching.

However, many stable matching mechanisms work quite well in practice even though theoretically they are not incentive compatible. One theoretical support for such phenomena is that many markets of interests can be modeled as large markets. Incentive compatibility in large markets has been studied in, e.g., exchange economy (Roberts and Postlewaite 1976), double auctions (Cripps and Swinkels 2006, Fudenberg et al. 2007), and the probabilistic serial mechanism (Kojima and Manea 2010). Most of these studies show either that the gain from manipulations converges to zero or that an equilibrium behavior converges to truth-telling. In two-sided matching markets, several studies have analyzed the incentive compatibility of large markets. Roth and Peranson (1999), Immorlica and Mahdian (2005) and Kojima and Pathak (2009) showed that the Gale-Shapley deferred acceptance algorithm becomes increasingly hard to manipulate as the number of participants increases.

The application of course allocation usually involves a large amount of students and courses. Motivated by previous studies on strategic behavior in large markets, we expect similar results to hold. It would be an interesting future direction to explore the incentive properties of large many-to-many matching markets.

In the course allocation at NTU, as discussed earlier, the priorities of courses are homogeneous. Such a property allows us to design an incentive compatible mechanism that satisfies the above solution conditions. The mechanism is presented in Section IV.E.

In the following discussions, unless specified explicitly, 'stability' refers to pairwise stability, 'group stability' refers to one-sided group stability and 'Pareto efficiency' refers to two-sided Pareto efficiency. Further, we use the notion *Pareto stability* (Sotomayor 2011) to denote matchings that are both two-sided Pareto-efficient and pairwise stable.

IV. Pareto-Stable Matching Mechanisms

For many-to-one matching with ties, a Pareto-stable assignment always exists and can be computed using the algorithm of Erdil and Ergin (2006, 2008). The algorithm relies on two observations. First, an assignment has a Pareto improvement only if the assignment graph has an augmenting path or cycle (formally defined in Section IV.A). Second, and more critically, any Pareto improvement to a stable assignment preserves stability. These observations immediately imply an algorithm to find a Pareto-stable assignment: starting from any stable assignment, keep making Pareto improvements by eliminating augmenting paths and cycles until none remains, and the resulting matching is both stable and Pareto-efficient.

In a many-to-many matching market, if only one side has ties, the same stability preserving result still holds.

CLAIM 2: *In a many-to-many matching market, if only one side has ties, Pareto improvement preserves stability.*

When both sides of a market have ties, however, we observe that the second critical property fails. That is, a Pareto improvement to a stable assignment need not preserve stability even when one side has homogeneous preferences, as the following example shows.

EXAMPLE 5 (Pareto improvement does not preserve stability): *Consider the example in the following figure, where m_2 and w_2 have a capacity of two each and other agents all have unit capacity.*

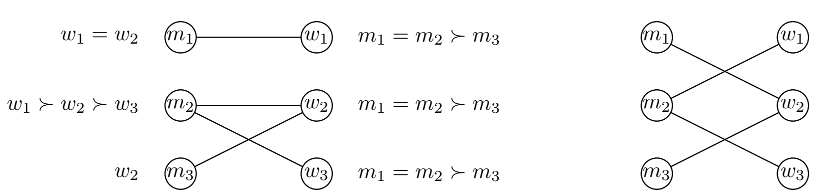

The assignment on the left-hand side is stable, and the assignment on the right-hand side is a Pareto improvement where m_2 strictly improves his assignment *and no one is worse off. However, the assignment on the right is unstable as* m_2 *and* w_2 *would like to match with each other rather than* w_3 *and* m_3*, respectively, i.e., it is a blocking pair.*

The example above shows that the approach of starting with an arbitrary stable assignment and making Pareto improvements does not work, because this need not preserve stability. Thus, all previous approaches (e.g., Erdil and Ergin (2006, 2008)) computing Pareto-stable assignments in variant models fail. Further, for the given stable assignment in the above example (left figure), there is only one Pareto improvement (right figure); thus, the problem cannot be solved by a careful selection of Pareto improvements.

It is thus unclear whether a Pareto-stable assignment exists in many-to-many matching markets with general arbitrary preferences for individuals. In this section we give a confirmative answer to this question by showing an algorithm that computes a Pareto-stable assignment. We first provide a characterization of Pareto efficiency, and then describe the algorithm. With respect to the course allocation at NTU with homogeneous preferences, we at the end provide two implementable Pareto-stable matching mechanisms.

A. *Characterization of Pareto Efficiency*

Given the connection between matching and network flow, it is not surprising that the existence of augmenting paths and cycles in an assignment is closely related to whether it can be improved, i.e., its Pareto efficiency. The main difference in the context of stable assignment is that nodes have preferences in addition to capacities. Thus, augmenting paths and cycles must improve not only the size of an assignment, but also its quality, as determined by node preferences. The formal definitions are as follows.

DEFINITION 4 (Augmenting Path): *Given an assignment* $\mu = (\mu_{mw})$*, we say that* $[m_0, w_1, m_1, \ldots, w_\ell, m_\ell, w_{\ell+1}]$ *is an* augmenting path *if (i)* $\sum_w \mu_{m_0 w} < c_{m_0}$ *and* $\sum_m \mu_{m w_{\ell+1}} < c_{w_{\ell+1}}$*, (ii)* $\mu_{m_k w_k} = 1$ *and* $\mu_{m_{k-1} w_k} = 0$ *for all* k*, and (iii)* m_k *weakly prefers* w_{k+1} *to* w_k *and* w_k *weakly prefers* m_{k-1} *to* m_k*.*

The first condition states that the capacities of m_0 and $w_{\ell+1}$ are not exhausted. The second condition states that pairs alternatively are not and are in the current assignment μ along the path. The last condition ensures that we are able to achieve a Pareto improvement by reassigning matches according to the augmenting path. That is, removing all pairs (m_k, w_k) and matching all pairs (m_k, w_{k+1}) produces a feasible assignment, which is a Pareto improvement over μ (where no one is worse off and m_0 and $w_{\ell+1}$ are better off).

DEFINITION 5 (Augmenting Cycle): *Given an assignment $\mu = (\mu_{ij})$, we say that $[m_1, w_2, m_2, \ldots, w_\ell, m_\ell, w_1, m_1]$ is an* augmenting cycle *if (i) $\mu_{m_k w_k} = 1$ and $\mu_{m_k w_{k+1}} = 0$ for all k (where $w_{\ell+1} = w_1$) (ii) m_k weakly prefers w_{k+1} to w_k and w_k weakly prefers m_{k-1} to m_k, and at least one of these preferences is strict.*

Again, we are able to match all pairs (m_k, w_{k+1}) and unmatch all pairs (m_k, w_k) in an augmenting cycle to get a Pareto improvement. For a given assignment, an augmenting path or cycle can be found easily using a network flow approach.

The following lemma characterizes the relation between stable assignment and augmenting path and cycle (its proof is the same as the one for a many-to-one matching market (Erdil and Ergin 2008)).

LEMMA 1: *A feasible assignment is Pareto-efficient if and only if it has no augmenting path or cycle.*

B. *Computing a Pareto-Stable Matching*

Our algorithm builds on the idea of Roth and Vande Vate (1990), who provide an alternative to the deferred acceptance algorithm to compute a stable (one-to-one) matching. Their algorithm can be interpreted as follows. Assume that all women are present at the beginning, and men 'arrive' one by one. We start with the empty matching. When a new man m arrives, match him to a most preferred woman w with whom he forms a blocking pair, if any; if this woman was already matched to a man m', set m' free and consider him as the next arriving man; the algorithm runs iteratively until all men have arrived. Because every woman who changes her partner in this process gets a strict improvement and no woman ever becomes worse off, the algorithm terminates, and the final matching is stable, because by construction the matching at every man's arrival is stable.

In our algorithm, all individuals are initially available; women are with *full* capacities and men are with *null* capacity. We consider all men one by one and increase their capacities unit by unit. When the capacity of a man is increased by one, we do a sequence of reassignments such that the resulting matching satisfies the following invariants (with respect to the current considered capacities):

- *Stability preserving:* it is always stable.
- *Women improving:* the assignment of any women does not become worse off.
- *No augmenting cycle:* it does not contain any augmenting cycle.

An important idea in our algorithm to derive Pareto efficiency is that in the process of reassignments, no augmenting cycles have ever been introduced in the matching. However, we allow the existence of augmenting paths. The key component of our algorithm is a subroutine for eliminating augmenting paths while preserving stability (and introducing no augmenting cycles). Having constructed a matching that is stable and contains no augmenting cycles, we apply the subroutine to eliminate augmenting paths in a stability preserving fashion, which finally yields a Pareto-stable matching as characterized by Lemma 1. The high-level structure of the algorithm is described below.

ALG-PARETO-STABLE algorithm

- Initialization:
  - There are no assigned matches (i.e., $\mu = 0$) between M and W.
  - All women have their full capacities available.
  - Let $\mathbf{d} = (d_m)_{m \in M}$ be a virtual capacity vector of men; initially $d_m = 0$ for $m \in M$.
- While there is $m \in M$ such that $d_m < c_m$, run INCREASE-CAP($\mathbf{d}$).
- While there is an augmenting path P, run ELIMINATE-PATH(P).
- Return the final assignment μ.

We have the following claim. (The details of the two subroutines and their analyses are rather technical and thus are deferred to Appendix D.)

THEOREM 1: *For any many-to-many matching market with arbitrary preferences, a Pareto-stable assignment always exists and can be efficiently computed by the algorithm* ALG-PARETO-STABLE.

C. Multiple Preferences

In the class registration application at NTU, students submit two preferences for PE and UE, respectively. Our algorithm continues to work for the settings with such multiple preferences. In general, for each individual x, the other side of the market is divided into (not necessarily disjoint) partitions $S_1(x), S_2(x), \ldots, S_\ell(x)$, where x has a preference (again, can be incomplete and have ties) and a capacity c_{xk} for each partition $S_k(x)$. Further, x has a universal capacity c_x that bounds the total number of partners that can be matched to x among all partitions. Observe that there are now two types of capacity constraints for each individual: a universal one and a local one for each partition. Without loss of generality, we can assume that $c_x \geq c_{xk}$ for any k. Note that there could be no relation between c_x and $\sum_k c_{xk}$. The preference model discussed in the previous sections corresponds to the special case when there is only one partition (i.e., $\ell = 1$), and the NTU class registration application corresponds to the case with two partitions (i.e., $\ell = 2$).

In this extension, the preference of every agent is restricted to every partition. While partitions are not necessarily disjoint, we assume that the preference lists of all partitions of an individual are disjoint. This is in accordance with our course allocation motivation in which the preferences that a student submits for PE and UE are required to be disjoint. Hence, for a given assignment $\mu = (\mu_{mu})$, m and w form a blocking pair if both of them strictly prefer each other to one of their assigned partners in the same partition (or if they have remaining capacities).

Our objective is again to find a Pareto-stable assignment, which can be computed by the same mechanisms described in the previous sections. Observe, however, that the preference of every agent is essentially with respect to each of its partitions. Hence, the definition of augmenting path and cycle will be changed accordingly, i.e., for any node in the path/cycle, its two neighbors must be from the same partition.

D. Pareto-Improving Draft Mechanism

In the course allocation problem at NTU, the homogeneity of the course preferences helps us to design implementable Pareto-stable matching mechanisms, in addition to the algorithm above that works for general many-to-many matching frameworks. In this subsection, we present a mechanism called the *Pareto-improving draft mechanism*, which incorporates the HBS draft mechanism (Budish and Cantillon 2012) with the augmenting paths/cycles elimination process. The mechanism is described as follows (assume that there are L priority levels of students).

```
Pareto-improving draft mechanism
--------------------------------------------------------------------------
For ℓ = 1,...,L:  In each round ℓ, consider students in the ℓ-th priority group.
Consider all these students one by one in a random order.
    • Each considered student receives his/her most preferred course among the
      remaining available courses (under the capacity constraint and breaking ties
      randomly).
    • After the assignment, consider all students who have been assigned courses
      (including those in higher priority groups) and all courses, eliminate
      student-sided augmenting paths/cycles until there is none left.
```

The Pareto-improving draft mechanism first considers students according to their priorities, and among students who are in the same priority level, assigns one course at a time over a series of rounds, according to the students' preferences. This is similar to the HBS draft mechanism and NTU's current mechanism. The main difference is that with ties in students' preferences, we also do a sequence of augmenting paths/cycles eliminations to derive Pareto efficiency. Specifically, in the mechanism, we use student-sided augmenting paths/cyles elimination[12] to ensure that each considered student, among all feasible assignments, is matched with his/her best possible assignment while not hurting all previously matched students. By the rule of the mechanism and the one-sided Pareto efficiency

[12]The student-sided augmenting paths/cycles here only consider the welfare of students, i.e., an augmenting path/cycle only requires that no students are worse off and at least one student is strictly better off. This is similar to the definition in Erdil and Ergin (2008), but a bit different from the ones defined earlier in Section IV.A in which welfare takes both sides of the market into account.

[12]For the first step when considering the first choice of all students (in any priority group), the augmenting paths/cycles eliminations are with respect to the unit capacity of all students; for all other steps, the capacities of students are two, as assumed in the model.

characterization of Erdil and Ergin (2008), the mechanism actually computes a students-sided Pareto-efficient matching. Further, the following theorem says that it satisfies two-sided Pareto efficiency and stability, i.e., Pareto stability.

THEOREM 2: *The Pareto-improving draft mechanism computes a Pareto-stable matching.*

E. *Pareto-Improving Serial Dictatorship Mechanism*

The Alg-Pareto-Stable algorithm and Pareto-improving draft mechanism presented above involve tie-breaking and may lead to different outcomes for different implementations. In general, a Pareto-stable matching mechanism is not strategyproof, even for one-to-one matching markets with ties (see Example 2 in Erdil and Ergin (2008)). In our setting, the same observation holds for the draft mechanism even when all courses have homogeneous preferences, as the following example shows.

EXAMPLE 6: *There are two students m_1, m_2 with a capacity of two each and three courses w_1, w_2, w_3 each with unit capacity. Their preferences are as follows:*

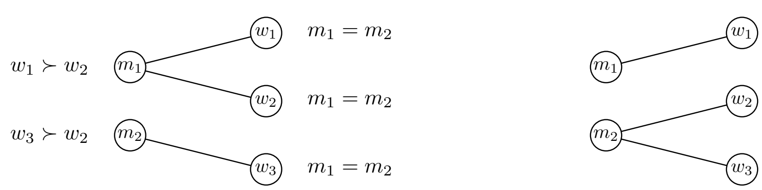

There are two Pareto-stable matchings $\mu_1 = \{(m_1, \{w_1, w_2\}), (m_2, w_3)\}$ and $\mu_2 = \{(m_1, w_1), (m_2, \{w_2, w_3\})\}$. In the implementation of the Pareto-improving draft mechanism, if there is a positive probability to output μ_2, then m_1 can misreport his preference to be $w_2 \succ w_1$, then the mechanism will always output μ_1 (although μ_2 is still a Pareto-stable assignment) and m_1 is better off. Similarly, if the mechanism has a positive probability to output μ_1, then m_2 can benefit by manipulation.

Is any implementation of Pareto-stable matching mechanisms strategyproof for students? It turns out that the answer is affirmative in our setting when courses have homogeneous preferences. The mechanism, called the *Pareto-improving serial dictatorship mechanism*, combines the random serial dictatorship mechanism and augmenting paths/cycles elimination. The mechanism is described as below (assume there are L priority levels for all students).

```
Pareto-improving serial dictatorship mechanism
```

For $\ell = 1, \ldots, L$: In each round ℓ, consider students in the ℓ-th priority level. Consider all these students one by one in a random order.

- Each considered student receives all of his/her most preferred courses among the remaining available courses (under the capacity constraint and breaking ties randomly).
- After the assignment, consider all students who have been assigned courses (including those in higher priority groups) and all courses, eliminate student-sided augmenting paths/cycles until there is none left.

Similar to the draft mechanism, here we also use student-sided augmenting paths/cyles elimination to derive Pareto efficiency. The key difference is that it assigns courses *all-at-once* while the draft mechanism assigns courses to students *one-at-a-time*.

THEOREM 3: *The Pareto-improving serial dictatorship mechanism outputs a Pareto-stable matching and is strategyproof for students.*

Draft versus dictatorship. The Pareto-improving draft and serial dictatorship mechanisms both satisfy the aforementioned "good" properties including pairwise and group stability, and one-sided and two-sided Pareto efficiency. In addition, the Pareto-improving serial dictatorship mechanism is strategyproof. However, the dictatorship mechanism may trigger the "callousness" phenomenon as described in Budish and Cantillon (2012). The following example illustrates how the draft mechanism can potentially solve the callousness issue by avoiding severely unfair allocations.

EXAMPLE 7: *There are two students m_1, m_2 each with a capacity of two and two courses w_1, w_2 each with unit capacity. m_1 has preference $w_1 \succ w_2$, and m_2 has preference $w_2 \succ w_1$. Both courses have homogeneous preferences $m_1 = m_2$. The Pareto-improving serial dictatorship mechanism assigns both courses to either* *student who has a higher ranking after the initial random tie breaking. But a fairer allocation would be to allocate each student one course, i.e., assign* w_1 *to* m_1 *and* w_2 *to* m_2*, which is exactly the outcome of the draft mechanism.*

As Budish and Cantillon (2012) documented about the HBS elective course allocation, the callous behavior harms efficiency in the sense that the welfare costs of using strategyproof dictatorship are much larger than the welfare costs of manipulability. As a consequence, Budish (2012) suggested the need for second-best alternatives to strategyproofness, e.g., incentive compatibility in large markets as discussed in Section III.D.

Incentive compatibility has been an important condition in various market design problems. However, one should also note that constraints always come with costs. As Budish and Cantillon (2012) showed for the course allocation data at HBS, on some simple measures of welfare the non-strategyproof draft mechanism outperforms the strategyproof dictatorship mechanism. As a special case of general dictatorship mechanisms, the *Pareto-improving serial dictatorship mechanism* also brings up callous behavior, which can be bad for welfare. Although the Pareto-improving draft mechanism solves the callousness issue in the *Pareto-improving serial dictatorship mechanism*, it is not strategyproof for students. This tradeoff between non-callousness and strategyproofness makes neither of them perfect for the NTU course allocation problem. In Section V, we compare the simulation results on both mechanisms to get a sense of which is better, if not perfect.

V. Simulations

We use NTU's course allocation data on PE and UE for three consecutive academic years: 2010, 2011 and 2012 to examine the effectiveness of introducing ties to students' preferences and to compare the Pareto-improving draft and dictatorship mechanisms with NTU's current mechanism. The data consist of (a) the number of course vacancies; (b) students' *strict* preferences for up to five PE and five UE courses, separately; and (c) course allocation results from NTU's current mechanism (after the manual adjustment process) (listed in Table 1). The input information for our simulations is specified as follows:

- Course preferences: we assume with little loss of generality that courses' homogeneous preferences for students are based only on the students' study year. That is, all courses strictly prefer final year students to penultimate year students, and so on. This assumption is for the simplicity of simulations.
- Course capacities: the above (a) provided by the data set.
- Student preferences: the above (b) provided by the data set. (Random ties are introduced to students' preferences to simulate the performance of the proposed mechanisms with indifferences in preferences.)
- Student capacities: either one or two, see the specific setup in the simulations described below.

Before introducing ties in students' preference lists, we first take a closer look at the distribution of the lengths of students' preference lists, which are depicted in Figure 2. For both PE and UE, it can be seen that the length of preferences is almost symmetrically distributed, and the average lengths are 3.00 and 3.07, respectively. For preference lists no longer than three, a maximum of two ties can be introduced. Further, in the recent course registration survey (see Figure 1 Question 3), 35 percent of students preferred two levels of preferences and 23 percent preferred three levels, representing a major portion of all surveyed students and over 76 percent of whom preferred to have ties.[13] Given the average lengths and survey information, in our simulations we consider two scenarios: two or three levels in students' preference lists.

- Two levels means that there is one strict '$\succ$' and the preference list of a student is divided into two levels at random.
- Three levels means that there are two strict '$\succ$' and the preference list of a student is divided into three levels at random.

Note that a student is indifferent between the courses in the same level. For example, if a student has preference $A \succ B \succ C \succ D \succ E$, then $A = B \succ C = D = E$ is a (random) realization with two levels and $A \succ B = C \succ D = E$ is a (random) realization with three levels.

[13]We note that the current limit of a maximum of five courses for PE/UE preference lists does not impose a significant constraint against students indicating all of the courses that they like, as the survey results show that 75 percent of the students are satisfied with the upper limit of five courses (see Figure 1 Question 4).

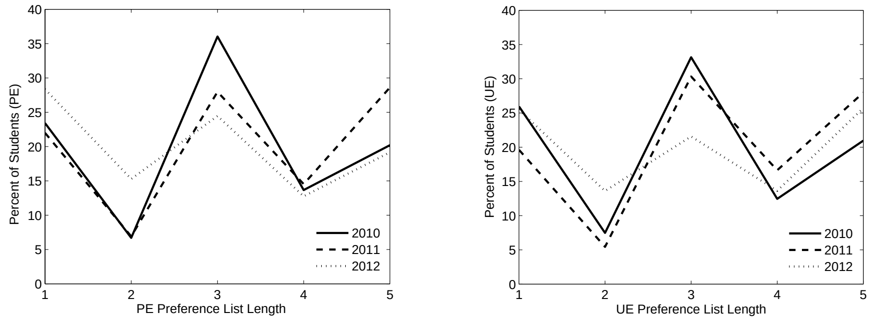

FIGURE 2. DISTRIBUTION OF THE LENGTHS OF STUDENTS' PE AND UE PREFERENCE LISTS.

Next we describe our simulation environments and results. Due to the randomness involved in introducing ties into students' preferences, five independent computations are performed for each of the simulation environments, and the percentage is the average of the five experiments. For the rest of the paper, percentages are measured as the total number of allocations over the total number of vacancies, unless otherwise specified.

A. Experimental Environment I

We consider three scenarios for simulations.

- Pure PE. We only consider PE courses; thus, we assume that each student has unit capacity and only consider his/her preference for PE courses.
- Pure UE. We only consider UE courses; thus, we assume that each student has unit capacity and only consider his/her preference for UE courses. Note that the first two scenarios are essentially many-to-one matchings.
- Combined PE+UE. Here all of the courses are pooled together. Note that if a course falls into both the PE and UE categories, recall that in the NTU's current system, two capacities of the course are specified manually for PE and UE, respectively. In our simulations here, we merge the two separate capacities, making it a single capacity for each such course. Each student still has two separate preferences, one for PE and one for UE; the capacity of the student is in the format of $1 + 1$, i.e., the student can get at most one PE course and at most one UE course, and it is possible that the student is allocated two

courses in total. Note that this scenario precisely captures the course allocation requirements of NTU, and is the one by which we quantitatively evaluate the performance of our mechanism, which introduces ties into students' preferences.

In the three scenarios, students essentially have unit capacities (in the third scenario, a student with two separate preferences can be treated as two students, among whom one has preferences consisting of PE only and the other has preferences consisting of UE only). In such cases, one can easily notice that the draft and dictatorship mechanisms are equivalent. In practice, we run simulations using both mechanisms and the results do appear to be the same as expected, except for some small deviations due to the random nature of the mechanisms. Therefore, we do not distinguish these minor differences caused purely by the randomness in the algorithms, and instead focus on the effects of allowing ties into students' preferences. The simulation results for the total number of allocations are shown in Table 2 and the corresponding statistics on the number of unassigned students for combined PE+UE are summarized in Table 3.[14] Next, we give a detailed explanation for the statistics in each column of Table 2 (using the data from 2010).

- Maximum matching. We need a generally reasonable benchmark to measure the performance of different matchings. (Note that the total number of students or course vacancies does not qualify, as no feasible matching can match the bound.) To this end, we consider a maximum cardinality matching (i.e., one has the maximum number of assigned pairs) from the set of all mutually acceptable pairs. The size of a maximum cardinality matching provides a theoretical upper bound on all feasible matchings. For PE, UE and combined PE+UE, the size of a maximum matching in 2010 is 7,140, 8,958 and 16,536, which takes 87.95 percent, 63.18 percent and 74.17 percent of the total vacancies, respectively.
- NTU's current mechanism. NTU's current mechanism is employed with respect to the given preferences. Note that the statistics here differ slightly from those in Table 1, as we consider slightly simplified preferences for the courses and do not include the manual adjustment process. It can be seen that in 2010,

[14]The percentages in this table are calculated as the total number of unassigned students over the total number of students.

Table 2—Experimental Results I: Total Number of Allocations

	Course	Maximum matching	NTU's current mechanism	Pareto-Stable without ties	Pareto-Stable with 3 levels	Pareto-Stable with 2 levels
2010	PE	7,140 87.95%	6,615 81.49%	6,647 81.88%	6,720 82.78%	6,852 84.41%
	UE	8,958 63.18%	8,142 57.43%	8,162 57.57%	8,295 58.51%	8,515 60.06%
	PE+UE	16,536 74.17%	**14,990** **67.23%**	**15,105** **67.75%**	**15,312** **68.68%**	**15,670** **70.28%**
2011	PE	7,725 76.55%	7,045 69.81%	7,150 70.85%	7,342 72.75%	7,368 73.01%
	UE	8,041 52.41%	7,495 48.85%	7,537 49.13%	7,626 49.71%	7,784 50.74%
	PE+UE	15,803 62.13%	**14,526** **57.11%**	**14,627** **57.51%**	**14,833** **58.32%**	**15,190** **59.72%**
2012	PE	7,308 78.54%	6,726 72.28%	6,762 72.67%	6,838 73.49%	7,005 75.28%
	UE	8,075 49.66%	7,556 46.47%	7,574 46.58%	7,672 47.18%	7,819 48.09%
	PE+UE	15,412 60.29%	**14,283** **55.87%**	**14,350** **56.13%**	**14,530** **56.84%**	**14,855** **58.11%**

about 81.49 percent of PE vacancies, 57.43 percent of UE vacancies and 67.23 percent of the combined PE+UE vacancies are allocated. We will compare the performance of our mechanisms to these statistics.

- Pareto-Stable without ties. In this column, the draft/dictatorship mechanism is employed with strict preferences for the students. In comparison with the allocations using NTU's current mechanism, for 2010 we see a 0.39 percent improvement for PE, a 0.14 percent improvement for UE allocations and a 0.52 percent improvement for combined PE+UE. Note that NTU's current mechanism also generates a Pareto-stable matching if considering the strict students' preferences. The marginal improvement in the number of allocations comes from the differences in the implementations of the algorithms.
- Pareto-Stable with 3 levels. In this column, the draft/dictatorship mechanism

is employed with three random levels on students' preferences. That is, each student's preference has at most three levels. Here we observe improvements of 1.29 percent, 1.08 percent and 1.45 percent for PE, UE and combined PE+UE, respectively, compared to NTU's current mechanism.

- Pareto-Stable with 2 levels. In this column, the draft/dictatorship mechanism is employed with two random levels on students' preferences. That is, each student's preference has at most two levels. We observe further improvements here of 2.92 percent, 2.63 percent and 3.05 percent for PE, UE and combined PE+UE, respectively. The last two columns show the simulation results after introducing ties into students' preferences, indicating that it can improve the overall efficiency of students.[15]

We also observe improvements in the total number of allocations after introducing ties into students' preferences for 2011 and 2012. In 2011, there are 2.94 percent, 0.86 percent and 1.21 percent improvements for PE, UE and combined PE+UE, respectively, after we divide preferences into three levels. The improvements for 2012 are 1.21 percent, 0.71 percent and 0.97 percent. Furthermore, we see more significant improvements of 3.20 percent, 1.89 percent and 2.61 percent for PE, UE and combined PE+UE, respectively, when we divide the preference lists into two levels, with improvements of 3.00 percent, 1.62 percent and 2.24 percent, respectively, in 2012.

In summary, for the combined PE+UE scenario, after introducing ties into students' preferences, over the three years we see an average improvement of 1.21 percent and 2.63 percent for three and two levels of preferences, respectively. This translates to roughly 292 or 639 more student-course assignments every year which, compared to the allocation results from NTU's current mechanism, significantly improves overall students' social efficiency. The statistics for combined PE+UE of the simulations are also depicted in Figure 3.

[15] Note that for the PE and UE scenarios, students have unit capacity, the allocations are therefore many-to-one matchings. For the column 'Pareto-Stable without ties', the simulations are equivalent to Erdil and Ergin's algorithm (Erdil and Ergin 2008), which computes a Pareto-stable matching for many-to-one markets with *one-sided* indifferences. For the last two columns 'Pareto-Stable with 3 and 2 levels', the simulations are equivalent to Erdil and Ergin's algorithm (Erdil and Ergin 2006), which computes a Pareto-stable matching for many-to-one markets with *two-sided* indifferences.

In addition to improving the total number of allocations, as the following table shows, we also observe a decrease in the total number of unassigned students of 2.87 percent, 2.18 percent and 2.42 percent for the three years, respectively. (Unassigned students refer to those who are not allocated to any course.)

Table 3—Experimental Results I: Number of Unassigned Students for PE+UE

	NTU's current mechanism	Pareto-Stable without ties	Pareto-Stable with 3 levels	Pareto-Stable with 2 levels
2010	1,866 13.85%	1,766 13.11%	1,679 12.46%	1,479 10.98%
2011	883 7.80%	859 7.59%	779 6.88%	636 5.62%
2012	1,007 8.73%	970 8.40%	885 7.67%	728 6.31%

B. *Experimental Environment II*

In the simulations described above for the combined PE+UE scenario, we explicitly set the capacities of students to be $1+1$. This assumption is in accordance with NTU's current system and helps us to compare the simulation results to the current mechanism. In practice, however, students may get more than one course from a category (either PE or UE) after the manual adjustment period. This fact is illustrated by the statistics in Table 1, where the number of allocations for UE is larger than the number of students in 2011. The 1+1 capacity is not a sharp constraint and is purely for equitability, resulting in significant inefficiency in course allocation. Thus, in the following simulations we assume that each student still has an overall capacity of two, but can be allocated any two courses (either one PE and one UE, or two PEs, or two UEs) from his/her preference. This change reflects practical situations in which students may have strong preferences for taking courses from one category (e.g., considering a student with 3 PE and 25 UE academic units left, then the student certainly prefers to get two UE courses rather than one PE and one UE). An overall capacity of two is imposed to allow for some degree of balance and fairness among students at different levels, by avoiding allocating almost all of the courses to senior students while junior students receive no courses.

When students have multi-unit capacities, implementations of the draft and dictatorship mechanisms will lead to different results. The main source of such differences is that in the draft mechanism students are allocated courses one-at-a-time while in the dictatorship mechanism students are allocated courses up to his/her capacity all-at-once. Thus, in addition to analyzing the efficiency improvements generated by introducing ties, we also focus on the comparisons between the Pareto-improving draft and dictatorship mechanisms.

Similar to the simulations described in the previous subsection, here we also consider three scenarios: pure PE, pure UE, and combined PE+UE. Given that students can now get two courses from the same category, and in practice they may have preferences for the two categories, in the simulations of combined PE+UE we randomly concatenate a student's PE and UE preferences into one list:

- If a student prefers category PE to UE, the new single preference is the PE preference followed by the UE preference.
- If a student prefers category UE to PE, the new single preference is the UE preference followed by the PE preference.
- If a student is indifferent between PE and UE, the new single preference is by randomly merging the PE and UE preferences (while keeping the same ranking for those PE and UE courses).

In our simulations, students are uniformly distributed between the above three types.

To introduce ties into students' preferences, for the PE and UE scenarios, the preference lists are randomly divided into two levels. For the combined PE+UE scenario, the concatenated preference lists are divided into four levels: one strict '$\succ$' in PE's preference, one strict '$\succ$' in UE's preference, and one strict '$\succ$' in the concatenation of the two preferences. If a student is indifferent between the two categories, the three strict '$\succ$' are placed at random.

Table 4 shows the simulation results in terms of the number of allocations from the Pareto-improving draft dictatorship mechanisms when students have a capacity of two and ties in preferences. In addition to the total number of allocations, we also compare the draft and dictatorship mechanisms by two other measures: average rank[16] and the total number of unallocated students. The statistics are given in Tables 5 and 6, respectively. Note that Table 6 only refers to the PE+UE scenario and the percentages are calculated as the total number of unassigned students over the total number of students.

Table 4—Experimental Results II: Total Number of Allocations

	2010		2011		2012	
Course	Draft	Dictatorship	Draft	Dictatorship	Draft	Dictatorship
PE	7,452 91.80%	7,386 90.98%	8,798 87.18%	8,690 86.11%	8,255 88.72%	8,185 87.96%
UE	10,046 70.86%	9,926 70.01%	10,479 68.30%	10,313 67.22%	10,399 63.95%	10,243 63.00%
PE+UE	**17,354** **77.83%**	**16,979** **76.15%**	**17,803** **70.00%**	**17,473** **68.70%**	**17,308** **67.70%**	**17,015** **66.56%**

Table 5—Experimental Results II: Average Rank

	2010		2011		2012	
Course	Draft	Dictatorship	Draft	Dictatorship	Draft	Dictatorship
PE	1.73	1.93	1.72	1.88	1.64	1.81
UE	1.96	1.96	1.83	1.95	1.78	1.95
PE+UE	1.85	1.85	1.72	1.79	1.66	1.74

From the statistics, we can see that for all of the scenarios, the draft mechanism outperforms the dictatorship mechanism in terms of the total number of allocations, the average rank and the total number of unassigned students, despite the fact that it is non-strategyproof. The dominant relationship is especially significant for the number of unassigned students. With the draft mechanism, only 5.06 percent, 3.15 percent and 3.90 percent of the students are unassigned any course

[16]The idea of using average ranks as a simple measure of welfare is motivated by Budish and Cantillon (2012). The average rank statistics here are calculated as the average rank of the courses in the student's assigned bundle based on the ranks in the *strict* preferences because it is impossible to get the real preferences with ties. Consequently, these average rank statistics are not the real average ranks. However, for the purpose of comparisons between two mechanisms, they are already sufficient.

Table 6—Experimental Results II: Number of Unassigned Students for PE+UE

	2010		2011		2012	
PE+UE	Draft	Dictatorship	Draft	Dictatorship	Draft	Dictatorship
Number	682	3,104	357	1,241	450	1,518
Percentage	5.06%	23.04%	3.15%	10.96%	3.90%	13.15%

for the three years, respectively, as compared to 23.04 percent, 10.96 percent and 13.15 percent with the dictatorship mechanism.

These statistics also demonstrate the negative effects on welfare that result from the "callousness" phenomenon in the dictatorship mechanism. The results support the conclusions of Budish and Cantillon (2012) concerning the role of strategyproofness in practical market design; that is, the callousness costs of a strategyproof dictatorship mechanism are much larger than the costs of manipulability in a draft mechanism (Budish 2012), especially in large marketplaces such as course allocation. Overall, if compared with the total allocation numbers from NTU's current mechanism in Table 2, there is approximately 10.60 percent, 12.89 percent and 11.83 percent improvement for the three years, respectively, with the dictatorship mechanism. These statistics are also illustrated in Figure 3.

VI. Concluding Remarks

We study the course allocation at NTU and formulate the problem as a many-to-many matching market with preferences. To improve overall efficiency, we consider introducing ties into students' preferences as a refinement to the current system. The fact that a stable outcome need not be Pareto-efficient with the presence of ties causes a loss in efficiency among the well-established stable solutions. We therefore employ the solution concept of Pareto stability as a refinement to the solution concept of stability, and establish an algorithm that computes a Pareto-stable matching for general many-to-many matching markets. We further propose two competing Pareto-stable matching mechanisms, i.e., Pareto-improving draft and dictatorship mechanisms, for the course allocation application.

With the designed draft and dictatorship mechanisms, we run simulations on the real course registration data from three academic years: 2010, 2011 and 2012.

Our results show significant improvement in efficiency when ties are introduced into students' preferences. In total, we can see up to 2,597 more course-student assignments for 2010 (3,263 for 2011 and 3,026 for 2012). This is equivalent to approximately 11.65 percent (12.83 percent for 2011 and 11.84 percent for 2012) improvement in total efficiency. These positive results, summarized in the following figure, call for changes to NTU's current course allocation system.

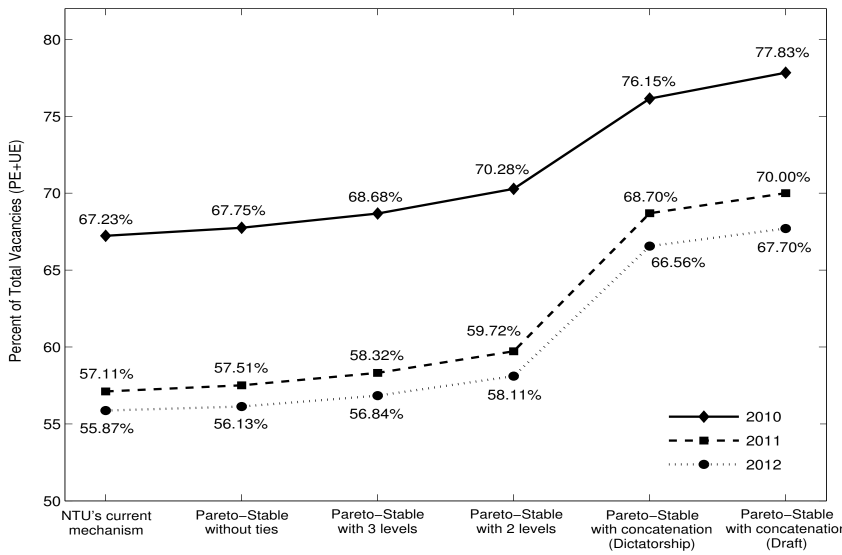

Figure 3. Summary of Simulation Results: Total Number of Allocations

Our simulation results comparing the draft and dictatorship mechanisms indicate that the Pareto-improving draft mechanism outperforms the dictatorship mechanism in terms of the total number of allocations, the average rank and the total number of unassigned students, despite the fact that the former is non-strategyproof. Our results echo the findings of Budish and Cantillon (2012), who suggested that strategyproofness has both benefits and costs.

Ties are a realistic condition occurring in many matching markets with preferences, especially when individuals have incomplete information. Our work, following the studies of Erdil and Ergin (2006) and Erdil and Ergin (2008), is devoted to improving social efficiency in the presence of ties. The models studied in these works can potentially be applied to other applications with a similar setup.

Our work considers a number of fundamental solution concepts, including pairwise and group stability, and one-sided and two-sided Pareto efficiency. We examine the existence and computation of these solution concepts. These results are of independent interest and may find applications in other many-to-many matching markets.

REFERENCES

Abdulkadiroğlu, Atila, Parag A. Pathak, and Alvin E. Roth. 2009. "Strategy-Proofness versus Efficiency in Matching with Indifferences: Redesigning the NYC High School Match." *American Economic Review* 99 (5): 1954-78.

Abdulkadiroğlu, Atila, and Tayfun Sönmez. 2003. "School Choice: A Mechanism Design Approach." American Economic Review 93 (3): 729-47.

Alkan, Ahmet. 2001. "On Preferences over Subsets and the Lattice Structure of Stable Matchings." Review of Economic Design 6 (1): 99-111.

Alkan, Ahmet. 2002. "A Class of Multipartner Matching Models with a Strong Lattice Structure." Economic Theory 19 (4): 737-46.

Blair, Charles. 1988. "The Lattice Structure of the Set of Stable Matchings with Multiple Partners." Mathematics of Operations Research: 13(4), 619-28.

Budish, Eric. 2012. "Matching 'versus' Mechanism Design." ACM SIGecom Exchanges 11 (2): 4-15.

Budish, Eric, and Estelle Cantillon. 2012. "The Multi-unit Assignment Problem: Theory and Evidence from Course Allocation at Harvard." American Economic Review 102 (5): 2237-71.

Chen, Ning, and Arpita Ghosh. 2011. "A Market Clearing Solution for Social Lending." In *Proceedings of the Twenty-Second International Joint Conference on Artificial Intelligence*, 152-57.

Cripps, Martin W., and Jeroen M. Swinkels. 2006. "Efficiency of Large Double Auctions." Econometrica: 74 (1), 47-92.

Echenique, Federico, and Jorge Oviedo. 2006. "A Theory of Stability in Many-to-Many Matching Markets." Theoretical Economics: 1 (2), 233-73.

Erdil, Aytek, and Haluk Ergin. 2006. "Two-Sided Matching with Indifferences." Manuscript.

Erdil, Aytek, and Haluk Ergin. 2008. "What's the Matter with Tie-breaking? Improving Efficiency in School Choice." American Economic Review: 98 (3), 669-89.

Gale, David, and Lloyd S. Shapley. 1962. "College Admissions and the Stability of Marriage." American Mathematical Monthly: 69 (1), 9-15.

Fudenberg, Drew, Markus M. Mobius, and Adam Szeidl. 2007. "Existence of Equilibria in Large Double Auctions." Journal of Economic Theory: 133 (1), 550-67.

Immorlica, Nicole, and Mohammad Mahdian. 2005. "Marriage, Honesty, and Stability." In *Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms*, 53-62.

Kominers, Scott D., Mike Ruberry, and Jonathan Ullman. 2010. "Course Allocation by Proxy Auction." In *Proceedings of the 6th International Conference on Internet and Network Economics*, 551-58.

Kojima, Fuhito, and Mihai Manea. 2010. "Incentives in the Probablistic Serial Mechanism." Journal of Economic Theory: 145 (1), 106-23.

Kojima, Fuhito, and Parag A. Pathak. 2009. "Incentives and Stability in Large Two-Sided Matching Markets." American Economic Review: 99 (3), 608-27.

Krishna, Aradhna, and M. Utku Ünver. 2008. "Improving the Efficiency of Course Bidding at Business Schools: Field and Laboratory Studies." Marketing Science: 27 (2), 262-82.

Martinez, Ruth, Jordi Masso, Alejandro Neme, and Jorge Oviedo. 2004. "An Algorithm to Compute the Full Set of Many-to-Many Stable Matchings." Mathematical Social Sciences 47 (2): 187-210.

Roberts, Donald J., and Andrew Postlewaite. 1976. "The Incentives for Price-Taking Behavior in Large Exchange Economies." Econometirca 44 (1): 115-27.

Roth, Alvin E.. 1984. "Stability and Polarization of Interests in Job Matching." Econometrica 52 (1): 47-57.

Roth, Alvin E., and Elliott Peranson. 1999. "The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design." American Economic Review 89 (4): 748-80.

Roth, Alvin E., and Marilda A. Oliveira Sotomayor. 1990. "Two-Sided Matching: A Study in Game-Theoretic Modeling and Analysis." New York: Cambridge University Press.

Roth, Alvin E., and John H Vande Vate. 1990. "Random Paths to Stability in Two-Sided Matching." Econometrica 58 (6): 1475-80.

Sönmez, Tayfun, and M. Utku Ünver. 2010. "Course Bidding at Business Schools." International Economic Review 51 (1): 99-123.

Sönmez, Tayfun, and M. Utku Ünver. 2011. "Matching, Allocation, and Exchange of Discrete Resources." In *Handbook of Social Economics*, edited by J. Benhabib, A. Bisin, and M. Jackson, 1A: 781-852. North-Holland.

Sotomayor, Marilda A. Oliveira. 1999. "Three Remarks on the Many-to-Many Stable Matching Problem." Mathematical Social Sciences 38 (1): 55-70.

Sotomayor, Marilda A. Oliveira. 2011. "The Pareto-Stability Concept is a Natural Solution Concept for Discrete Matching Markets with Indifferences." International Journal of Game Theory 40 (3): 631-44.

Appendix: For Online Publication

NTU Curriculum Structure and Course Registration Process

This section introduces the curriculum structure and course registration process of NTU. At NTU, an undergraduate needs to fulfill both the Major Requirement and the General Education Requirement (GER). The Major Requirement includes major core courses, which are compulsory courses to satisfy the program requirements, and major prescribed electives, which are courses for specialization in a particular degree program. The GER is the curriculum requirement for broadening study, which covers key fields of knowledge for all students. It constitutes about 25 percent to 40 percent of the total curriculum workload and is divided into 3 classes of studies:

1) GER CORE: these include courses related to Human Resources Management, Communication Skills and Singapore Studies.
2) GER Prescribed Electives (PE): the courses represent the key fields of knowledge broadly relevant to all professions and are categorized into 3 sub-areas of studies:
    a) Arts, Humanities and Social Sciences
    b) Business and Management
    c) Science, Technology and Society
3) GER Unrestricted Electives (UE): these are courses chosen by students to broaden their learning experience. They may cover any area offered by the various departments, including, e.g., modern languages, entrepreneurship, music, and drama courses.

A course can fall into different categories simultaneously. For example, the course Principle of Economics can be in both PE and UE. There is a minimum academic units requirement for each category of courses for students to fulfil to meet their graduation requirement. The curriculum structure is shown in Figure A1.

An online Student Automated Registration System (STARS) is currently used for course registration at NTU. The information of the courses (e.g., time schedules and vacancies) is first released. The registration takes place over three phases.

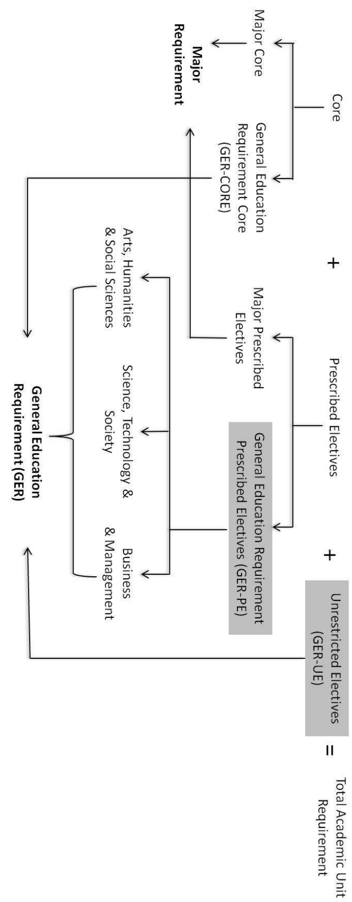

Figure A1. Curriculum Structure at NTU.

In the first phase, students register for the Major Requirement (major core and major prescribed electives) and GER CORE courses at their pre-specified date and time slot. These courses can be registered successfully as long as there are vacancies available (on a "first come, first serve" basis). Almost all students are able to register for their desired major courses and GER CORE courses in the first phase.[17]

The second phase decides allocations of PE and UE courses by a centralized mechanism described in Section I. The first two phases take place before

[17]A major reason is that most registrants of a major course are those from the department that offers the course; therefore, every department can easily manage their offered major courses for its own students.

a semester starts. During the first two weeks at the beginning of a semester, there is another and final phase where students can submit appeals for courses that they are keen to take, drop courses and add courses (provided vacancy availability). The appeals in this phase are handled manually by program coordinators from departments on a case by case basis. To ensure that certain special and urgent appeals are fulfilled, some courses may reserve a few vacancies for this final phase.

Responsive Preference with Lattice Structure

We give an example to show a responsive preference with a lattice structure. Assume that the preference of m over individuals is $P_m : (w_1 \succ w_2 = w_3 \succ w_4 = w_5)$ and its capacity is 3, then the lattice structure of m's responsive preference is shown in Figure B1. Each node in the figure denotes a feasible matching. For example, the node "{1,2,3}" means that m is matched with $\{w_1, w_2, w_3\}$.

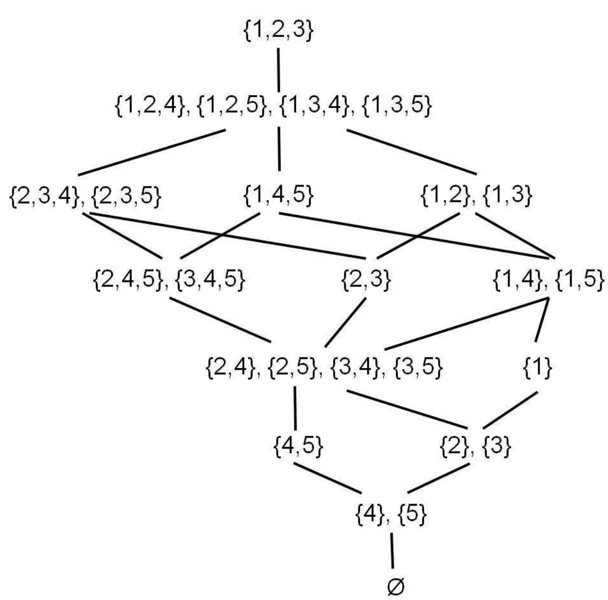

Figure B1. Lattice Structure of Responsive Preference.

Missing Proofs

C1. Proof of Claim 1

PROOF:

This can be seen by the following example: There are three students m_1, m_2, m_3 and two courses w_1, w_2 with unit capacity each. Their preferences are shown in Figure (a) below.

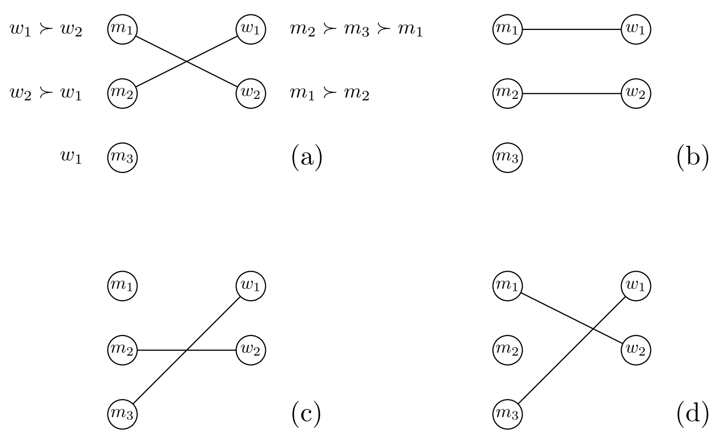

The first allocation $(m_1, w_2), (m_2, w_1)$ is pairwise stable. However, it can be seen that $\{m_1, m_2\}$ is not group stable as both of them can get better off by swapping the assigned courses (see Figure (b)). In this case, m_3 and w_1 form a blocking pair, which enforces the allocation to Figure (c). Now m_1 and w_2 form a blocking pair, which transforms the allocation to Figure (d), in which m_2 and w_1 are a blocking pair and the matching returns back to the first one. Hence, for the considered instance, it does not admit an allocation that is both pairwise stable and group stable.

C2. Proof of Proposition 2

PROOF:

Assume without loss of generality that W has homogeneous preferences over M and a Pareto stable assignment $\mu = (\mu_{mw})$ is not M-side group stable. Then there is a subset $S \subseteq M$ such that all members in S can strictly improve his allocation among reassignments inside S; denote the resulting new matching by μ'. Let

$S = \{m_1, m_2, \ldots, m_k\}$, and denote the assignment of each agent $x \in M \cup W$ in μ and μ' by $\mu(x)$ and $\mu'(x)$, respectively. Since μ is dominated by μ', we have $\mu'(m) \succ_m \mu(m)$ for all $m \in S$. Note that the assignments of all men that are not in S remain the same in μ and μ'.

Consider any man $m \in S$. Let $\mu(m) = \{w_1, w_2, \ldots, w_{c_m}\}$ with $w_1 \succeq_m w_2 \succeq_m \cdots \succeq_m w_{c_m}$ and $\mu'(m) = \{w'_1, w'_2, \ldots, w'_{c_m}\}$ with $w'_1 \succeq_m w'_2 \succeq_m \cdots \succeq_m w'_{c_m}$, where c_m is the capacity of m. We can insert a copy of $\emptyset$ if m is not fully matched in the two matchings. Then due to responsive preferences, we have $w'_i \succeq_m w_i$ for all i and at least one preference is strict.

As all women have the same preference, we can assume without loss of generality that the preference of women is complete over men (otherwise, those men who are unacceptable will never be matched in any feasible assignment). Consider the following two cases about the structure of the homogeneous preference of W over S.

Case 1. Women are indifferent among all the men in $S = \{m_1, m_2, \ldots, m_k\}$. We consider the exclusive-or structure of the two matchings μ and μ'.

We first show that no woman is worse off in μ'. Assume otherwise that there is a woman w_1 who is worse off in μ'. Since w_1 is indifferent between all men in S, we know that her number of assignments in μ' is less than that in μ, i.e., $|\mu'(w_1)| < |\mu(w_1)|$. Hence, there must exist a man m_1 such that m_1 breaks up the matching with w_1 in μ and is matched to a new woman w_2 in μ' in which $w_2 \succeq_{m_1} w_1$ (due to responsive preferences). Next consider w_2. If all men matched to w_2 in μ are already matched to her in μ', i.e., $\mu(w_2) \subset \mu'(w_2)$, then we know that w_2 does not exhaust her capacity in μ. Since μ is a stable matching, we know that m_1 is fully matched in μ and weakly prefers all his assigned parters to w_2; this implies, in particular, $w_2 =_{m_1} w_1$. Hence, as m_1 improves his assignment in μ', there must exist another woman w'_2 with $w'_2 \succ_{m_1} w_1$ such that $w'_2 \in \mu'(m_1) \setminus \mu(m_1)$ and w'_2 is fully matched in μ. For such a case, we switch the name of w_2 and w'_2. Therefore, there is a man m_2 such that m_2 breaks up the matching with w_2 in μ and is matched to a new woman w_3 in μ' in which $w_3 \succeq_{m_2} w_2$.

We continue with the argument. As the number of men is finite, eventually we will have a loop: $w_\alpha, m_\alpha, w_{\alpha+1}, m_{\alpha+1}, \ldots, w_\beta, m_\beta, w_\alpha$ in which m_i breaks the matching with w_i in μ and is matched to w_{i+1} in μ' for all i (where $w_{\beta+1} = w_\alpha$) and $w_{i+1} \succeq_{m_i} w_i$. (Note that it is possible that $\alpha = 1$, i.e., the loop goes back to the first woman w_1.) If at least one of the preferences is strict, then we have a Pareto improvement among these agents, this contradicts the fact that μ is Pareto stable. Hence, all these preferences are tight; in such a case, we can actually still use the old matchings in μ (i.e., remove all (m_i, w_{i+1}) and add (m_i, w_i) in μ'). Then we can continue with the same analysis on the exclusive-or structure of the two matchings μ and μ', and eventually derive a contradiction.

Therefore, we know that μ' is a Pareto improvement over μ (as all men in S are better off while no woman and any other man are worse off). This leads to a contradiction to the assumption that $\mu = (\mu_{mw})$ is Pareto efficient.

Case 2. There is at least one strict preference over two men in S, say, without loss of generality, $m_2 \succ m_1$. Similar to the above argument, there is a woman $w_3 \in \mu'(m_2) \setminus \mu(m_2)$ such that $w_3 \succ_{m_2} w$ for some $w \in \mu(m_2)$. As μ is a stable matching, w_3 must be fully matched in μ; thus, there is a man m_3 with $m_3 \succeq_{w_3} m_2$ (again, due stability of μ) who breaks the matching with w_3 in μ and is matched to a new woman w_4 in μ'. As m_3 also improves his assignment in μ', we can again use the same analysis as above to show that there is a Pareto improvement, which contradicts to the Pareto stability of μ (note that all women w_α considered in the process weakly prefers m_α to $m_{\alpha-1}$).

Therefore, when W have homogeneous preferences over M, a Pareto stable assignment must be M-side group stable.

C3. Proof of Claim 2

PROOF:

Without loss of generality, we assume that only women have ties in their preferences. Let μ be a stable matching and μ' be one derived from μ through Pareto improvement. Then for all $x \in M \cup W$, $\mu'(x) \succeq_x \mu(x)$, and at least one preference is strict.

Assume that μ' is not stable and (m, w) is a blocking pair. Let

$$\mu(m) = \{w_1, w_2, \ldots, w_{c_m}\} \text{ and } \mu'(m) = \{w'_1, w'_2, \ldots, w'_{c_m}\},$$

and let

$$\mu(w) = \{m_1, m_2, \ldots, m_{c_w}\} \text{ and } \mu'(w) = \{m'_1, m'_2, \ldots, m'_{c_w}\}.$$

(We can add $\emptyset$ to the list if one is not fully matched.) We can enumerate the indices such that $w'_i \succeq_m w_i$, for all $i = 1, 2, \ldots, c_m$ and $m'_j \succeq_w m_j$, for all $j = 1, 2, \ldots, c_w$. Note that $w_i =_m w'_i$ if and only if w_i and w'_i are the same woman. Since (m, w) is a blocking pair for μ', we know that $w \succ_m w'_{c_m}$ and $m \succ_w m'_{c_w}$.

Case 1: m and w are not matched in μ. Then $w \succ_m w'_{c_m} \succeq_m w_{c_m}$ and $m \succ_w m'_{c_w} \succeq_w m_{c_w}$, implying that (m, w) is a blocking pair for μ, a contradiction.

Case 2: m and w are matched in μ. Since the assignment of w is not worse off in μ' and $m \notin \mu'(w)$, there is $m' \in \mu'(w)$ such that $m' \notin \mu(w)$ and $m' \succeq_w m$. Consider m'; as all men have strict preferences and m' is not worse off in μ', there is $w' \in \mu(m')$ such that $w \succ_{m'} w'$. Hence, (m', w) is a blocking pair for μ, a contradiction.

Hence, μ' has no blocking pairs, and the claim follows.

C4. *Proof of Theorem 2*

PROOF:

Denote the matching computed by the mechanism by μ. By the rule of eliminating student-sided augmenting paths/cycles in the mechanism and the characterization of Pareto efficiency (Erdil and Ergin 2006), μ is immediately student-sided Pareto efficient. Then there is no (two-sided) augmenting path under the matching μ since no m (i.e., student) can be better off. We claim that there is no (two-sided) augmenting cycle either. Assume otherwise that there is an augmenting cycle $[m_1, w_2, m_2, \ldots, w_\ell, m_\ell, w_1, m_1]$, where $\mu_{m_k w_k} = 1$ and $\mu_{m_k w_{k+1}} = 0$ for all k (where $w_{\ell+1} = w_1$). Then according to the definition and the fact that no

m can get strict improvement, there must be some w getting better off; assume without loss of generality that w_1 gets better off, i.e., $m_\ell \succ m_1$. Due to the homogeneity property of all w's preferences and the fact that no w gets worse off, we can derive that $m_1 \succeq m_2 \succeq m_3 \succeq \cdots \succeq m_\ell \succ m_1$, which is a contradiction. Therefore, by Lemma 1, μ is (two-sided) Pareto efficient.

Next we show that μ is stable. The rule of the mechanism implies that for any two students m_1 and m_2, if m_1 has a higher priority than m_2, then m_1 does not envy any course assigned to m_2 (i.e., all courses assigned to m_1 are at least as good as any course assigned to m_2). Otherwise, when considering augmenting paths/cycles for the iteration of m_1, we would match a better course to m_1. Hence, if there is a blocking pair (m, w), where w strictly prefers m to one of her assignments m' (note that m' cannot be $\emptyset$ due to augmenting paths/cycles elimination at the iteration of m), i.e., $m \succ m'$. By the rule of the mechanism, m must be in a higher priority level than m'; and by above discussion, m does not envy any course assigned to m', which contradicts the assumption that m and w are a blocking pair. Hence, the mechanism always generates a stable matching.

This completes the proof of the theorem.

C5. Proof of Theorem 3

PROOF:

The proof of Pareto stability is the same as the one for Theorem 2. To the end of the strategyproofness, it can be seen that for each considered student, among the remaining available courses, we allocate him/her the best possible courses. Thus, the student has no incentive to lie. In the later augmenting paths/cycles eliminations, while the assignment of the student can be changed, a simple but critical invariant holds: Given the assignments of all previously considered students, we always allocate the best possible courses to the student. Therefore, it is a dominant strategy for the student to submit his/her true preference. This completes the proof of the theorem.

Details of the Algorithm

Note that in the algorithm, we always maintain the invariant that the algorithm contains no augmenting cycles. Why do we need such a condition, whereas it is allowed to have augmenting paths? Observe that the reason that a Pareto improvement may not preserve stability is that the path or cycle corresponding to the Pareto improvement contains a matched pair (m, w) where both m and w are also matched to a less preferred agent, say w' and m'. When the match (m, w) is removed in the reassignment process of the augmenting path/cycle, even though m and w could receive better partners in the path or cycle, they will prefer to be matched to each other instead of w' and m' respectively. For augmenting path, however, we can always start reassignment from one side of the path (say, the man), and stop proceeding along the path when we reach such a woman w (then (m', w) is unmatched and the process restarts). In this stability-preserving process, a woman becomes strictly better off. However, for the pair (m, w) in an augmenting cycle, we would need to release both (m', w) and (m, w') to preserve stability. That is, we would no longer have the monotonically improving property for women's assignments, which is critical to the analysis of the algorithm.

Note that in the algorithm, $\mu = (\mu_{mw})_{m \in M, w \in W}$ and $(d_m)_{m \in M}$ are global variables in both subroutines. The first subroutine, Increase-Cap, increases the virtual capacity of a man by one and does a number of reassignments to ensure the three invariants listed above (in particular, it guarantees that the assignment is stable for the increased virtual capacity vector). The second subroutine, Eliminate-Path, eliminates all possible augmenting paths to derive a Pareto-efficient assignment in a stability preserving fashion. After all augmenting paths have been eliminated, by Lemma 1, the returned assignment is Pareto-stable.

While the algorithm may look a bit complicated, the fact that no women ever get worse off in the process implies a simple, but critical, structure of the algorithm: we iteratively do a sequence of reassignments to improve women's assignments while preserving stability and containing no augmenting cycle. If at any moment in the algorithm a woman's assignment gets strictly improved, no matter at which stage the algorithm is, we terminate that thread immediately and go to Step (2) of the main algorithm to repeat the process given the current virtual capacity vector $\mathbf{d}$.

Next we describe the two subroutines in detail in the following subsections. (All discussions are with respect to the considered virtual capacity vector.) In the algorithm, for any (augmenting) cycle C and a pair $(m, w) \in C$, we use $C \setminus \{(m, w)\}$ to denote the path by removing pair (m, w) from C.

D1. Subroutine One: Capacity Increment

The first subroutine that increases virtual capacities of the men is the following.

INCREASE-CAP($\mathbf{d}$)

1) `Pick an arbitrary man` m `with` $d_m < c_m$
2) `Let` $d_m \leftarrow d_m + 1$, `i.e., increase the virtual capacity of` m `by one`
3) `Let` $S = \{w \mid (m, w)$ `is a blocking pair`$\}$
4) `Let` $T = \{w \in S \mid m$ `prefers` $w \succeq w'$ `for any` $w' \in S\}$
5) `If` $T = \emptyset$ `(i.e., there is no blocking pair),` <u>`return`</u>
6) `Otherwise`
   a) `If there exists` $w \in T$ `such that adding match` (m, w) `does not introduce any augmenting cycle`
      - `pick such a woman` w'
      - `add match` (m, w')
   b) `Otherwise`
      - `pick an arbitrary` $w' \in T$
      - `let` C `be a potential augmenting cycle by adding` (m, w')
      - `let` $P = \left[m \xrightarrow{C \setminus \{(m,w')\}} w'\right]$ `be the path from` m `to` w' `through` $C \setminus \{(m, w')\}$
      - `run` ELIMINATE-PATH(P)
   c) `If` w' `(defined either in Step (6.a) or (6.b)) is over-matched (i.e., matched to more than` $c_{w'}$ `neighbors)`
      - `let` m' `be a least preferred man matched to` w' `where deleting` (m', w') `does not introduce an augmenting cycle`
      - `delete match` (m', w')
      - `let` $d_{m'} \leftarrow d_{m'} - 1$
      - <u>`return`</u>
   d) `Otherwise,` <u>`return`</u>

When the virtual capacity of m is increased by one, there might be some blocking pairs, among which the subroutine tries to match m to one that he prefers most ($w' \in T$ in the above description). However, this could introduce potential augmenting cycles (Step 6(b)). Instead of matching m and w' directly, the subroutine considers a potential augmenting cycle C incurred by (m, w') and tries to do reassignments according the other path from m to w' along the cycle. Finally, if w' is over-matched, then we delete one of her least preferred assignments without incurring any augmenting cycles and delete the virtual capacity of that man by one. This guarantees that the assignment remains stable, and the assignment of w' strictly improves.

The existence of m' in Step 6(c) is guaranteed by the following lemma.

LEMMA 2: *Given a stable matching without augmenting cycles, for any woman w, let $S \subseteq M$ be the subset of men matched to w to whom w is least preferred. Then there is $m \in S$ such that deleting match (m, w) does not introduce any augmenting cycle.*

D2. Subroutine Two: Augmenting Path Elimination

Consider a given stable assignment, assume there is an augmenting path $P = [m_0, w_1, m_1, \ldots, w_\ell, m_\ell, w_{\ell+1}]$, where (m_i, w_i) is in the assignment and (m_i, w_{i+1}) is not. Note that it is possible that an individual x (either a man or a woman) or a pair (x, y) appears more than once in P. In this subsection, when we refer to an individual $x \in P$ or a pair $(x, y) \in P$, we denote the corresponding one at that position of P.

Before describing the subroutine, we will first consider a truncation process, which deletes some pairs in a given augmenting path according to different appearances of the same agent and will be used in the subroutine.

TRUNCATION.

For a given augmenting path P, we consider the following truncation function.

TRUNCATE-PATH(P)

1) **while** one of the following "if" conditions holds

- If there is m such that $P = [\dots, m, w_1, \dots, w_2, m, \dots]$ and m weakly prefers w_1 to w_2
  - truncate $P = [\dots, m, \cancel{(w_1, \dots, w_2, m,)} \dots]$
- If there is w such that $P = [\dots, w, m_1, \dots, m_2, w, \dots]$ and w weakly prefers m_2 to m_1
  - truncate $P = [\dots, w, \cancel{(m_1, \dots, m_2, w,)} \dots]$

2) Return path P

It can be seen that if TRUNCATE-PATH(P) is executed, by the rules of the truncation, no pair (x, y) can appear more than once after truncation. However, it is still possible that an individual appears more than once (e.g., when m strictly prefers w_2 to w_1, we do not truncate the two occurrences of m). The truncation process is necessary in our algorithm; in particular, it is important to the analysis of termination of the algorithm.

In TRUNCATE-PATH(P), if a truncation is executed at x (a man or a woman), we denote by $\Gamma(x)$ the truncated path. That is, $\Gamma(x) = [m, w_1, \dots, w_2, m]$ if $x = m$, and $\Gamma(x) = [w, m_1, \dots, m_2, w]$ if $x = w$.

We have the following observations.

PROPOSITION 3: *For any given augmenting path P, if a truncation is executed at x, then $\Gamma(x)$ forms a cycle and every individual involved is indifferent between its two neighbors in the cycle.*

PROOF:

We will only prove the claim for the first case when TRUNCATE-PATH(P) is executed at a man; the argument for the second case is similar. Assume that the given augmenting path $P = [\dots, m, w_1, \dots, w_2, m, \dots]$ is truncated between the two occurrences of m. By the rule of truncation, m weakly prefers w_1 to w_2; by the rule of augmenting path P, all individuals weakly prefer his/her unmatched neighbor to matched neighbor. Hence, $[m, w_1, \dots, w_2, m]$ forms a cycle and everyone is indifferent between its two neighbors (otherwise m strictly prefers w_1 to

w_2, it is an augmenting cycle, which contradicts to the fact that no augmenting cycle ever appears in the course of the algorithm).

PROPOSITION 4: *For any given augmenting path P, if a truncation is executed at x and x strictly prefers its one neighbor to the other for one occurrence of x in the truncation, then x still strictly prefers one neighbor to the other after truncation.*

PROOF:

We will only prove the claim when x is a man m; the argument for woman is similar. Consider the augmenting path $P = [\ldots, w_1, m, w_2, \ldots, w_3, m, w_4, \ldots]$ and a truncation is executed at m. Assume that m strictly prefers w_2 to w_1. Since P is an augmenting path, we know that m weakly prefers w_4 to w_3. By Proposition 3, m is indifferent between w_2 and w_3. Therefore, after truncation m strictly prefers one neighbor w_4 to the other w_1. The same argument holds if the strict preference occurs at the second occurrence m (*i.e.*, m strictly prefers w_4 to w_3).

LEMMA 3: *For any given augmenting path P,* Truncate-Path(P) *returns an augmenting path as well.*

PROOF:

Again we will only prove the claim for the first case of Truncate-Path(P) and the second case follows similarly. For the path $P = [\ldots, m, w_1, \ldots, w_2, m, \ldots]$ with the truncation for the middle of the two occurrences of m, if the first m is the beginning of path P, then certainly after truncation it is still a valid augmenting path. Otherwise, we can write P as $[\ldots, w_0, m, w_1, \ldots, w_2, m, w_3 \ldots]$ (note that the end of the path must be a woman). Notice that m weakly prefers w_1 to w_0, and w_3 to w_2. Further, we have m is indifferent between w_1 and w_2 by Proposition 3. Hence, m weakly prefers w_3 to w_0, which implies the desired result.

Elimination.

We next describe the subroutine to eliminate augmenting paths while preserving the three invariants listed at the beginning of the section. Note that for any augmenting path, its one side must be a man and the other side must be a woman. The subroutine starts from the man side and considers pairs one by one. Hence, for any man-woman pair in the path, the objective is to match them; and for any woman-man pair in the path, the objective is to unmatch them.

ELIMINATE-PATH(P)

1) Assume $P = [m^*, w_1, m_1, \ldots, w^*]$
2) Let $e = (m^*, w_1)$ be the first pair on path P
3) **while** $e \neq \emptyset$
   - If e is not a match (i.e., $e = (m, w)$)
     - if adding match (m, w) does not introduce an augmenting cycle
       a) add match (m, w)
       b) if w is not over-matched, <u>return</u>
       c) if w strictly prefers m to a current partners
          * let m' be a least preferred man matched to w where deleting (m', w) does not introduce an augmenting cycle (by Lemma 2, such m' exists)
          * delete match (m', w) and let $d_{m'} \leftarrow d_{m'} - 1$
          * <u>return</u> to Step 2 of the main algorithm ALG-PARETO-STABLE to run INCREASE-CAP
       d) else let e be the next pair after (m, w) in P
     - otherwise
       e) let $C = [m, w_1', m_1', \ldots, w_k', m_k', w, m]$ be such a potential cycle if adding (m, w)
       f) expand $P = \left[m^*, \ldots, m, w_1' \xrightarrow{C \setminus \{(m,w)\}} m_k', w, \ldots, w^*\right]$
       g) truncate $P = \left[m^*, \ldots, \text{TRUNCATE-PATH}\left(m, w_1' \xrightarrow{C \setminus \{(m,w)\}} m_k', w, \ldots, w^*\right)\right]$
       h) let e be the first pair returned by the TRUNCATE-PATH
   - If e is a match (i.e., $e = (w, m)$)
     - if deleting match (w, m) does not introduce an augmenting cycle
       i) delete match (w, m)
       j) let e be the next pair after (w, m) in P
     - otherwise
       k) run the above Steps (e,f,g,h)
          (switching the notations of m and w (except m^* and w^*))

The subroutine tries to add and delete matches one by one along pairs in the path P. If the current considered pair is a man-woman pair (i.e., $e = (m, w)$), the subroutines matches them if it does not introduce any augmenting cycle. If the assignment of w is strictly improved (i.e., the condition in Step (3.b) or (3.c) is satisfied), the subroutine terminates. Note that at this point the subroutine may not completely eliminate the augmenting path, however, the overall assignment of the woman gets strictly improved and the process restarts at the capacity increment stage. If matching m and w will introduce a potential augmenting cycle, instead of adding the match directly, the subroutine takes a "detour" and considers the other path from m to w along the cycle and expands it to the path P (Step 3(f); by the following Lemma 4, it is a valid expansion). Then the subroutine will do a truncation from m to the end of the path P and restarts the process by considering the first pair returned by the truncation (its first individual must be m). The subroutine performs similarly if the considered pair is a woman-man pair.

We first establish the following observations.

LEMMA 4: *The expansion of path P in Step (3.f) is a well-defined augmenting path.*

PROOF:

Let $P_1 = [m^*, w_1, m_1, \ldots, w', m]$ and $P_2 = [w, m', \ldots, w^*]$, where w' is the woman before m and m' is the man after w in P. Then the original augmenting path can be written as $P = [P_1, m, w, P_2]$. By the fact that P is an augmenting path, we know that m weakly prefers w to w' and w weakly prefers m to m'. Let $C' = \left[m_1' \xrightarrow{C\setminus\{m,w\}} w_k'\right]$, then the extended path (denoted by P') is $P' = [P_1, m, C', w, P_2]$. By the fact that C is an augmenting cycle if adding (m, w), we know that m weakly prefers w_1' to w and w weakly prefers m_k' to m. Therefore, m weakly prefers w_1' to w' and w weakly prefers m_k' to m'; this implies that the expanded path P' is a well-defined augmenting path.

We have the following key claim, which implies that the subroutine always terminates.

LEMMA 5: *The subroutine* ELIMINATE-PATH(P) *terminates in finite number of steps for any augmenting path P.*

D3. Proof of Lemma 5

In this section, we will prove the second subroutine Eliminate-Path always terminates.

PROPOSITION 5: *In the course of the subroutine* Eliminate-Path(P)*, for each considered pair e in the augmenting path P, we can always reach a different pair (i.e., redefine e by Step (3.d), (3.h), or (3.j)) with different starting individual. That is,*

- *if $e = (m, w)$, there is w' such that the subroutine matches (m, w') (not introducing any augmenting cycle) and next moves to $e = (w', \cdot)$;*
- *if $e = (w, m)$, there is m' such that the subroutine deletes (w, m') (not introducing any augmenting cycle) and next moves to $e = (m', \cdot)$.*

PROOF:

We will only prove the claim for the case when $e = (m, w)$; the same argument extends for $e = (w, m)$. Assume that adding $(m, w_1) \triangleq (m, w)$ introduces an augmenting cycle C_1 (otherwise, we are done); let w_2 be the other woman incident to m in C_1. Note that m weakly prefers w_2 to w_1 and w_2 weakly prefers m to her assignment in C_1. Next the subroutine expands path P with $\left[m, w_2 \xrightarrow{C_1 \backslash \{(m,w_1)\}} w_1\right]$, and consider adding (m, w_2). Again assume that it introduces an augmenting cycle C_2; let w_3 be the other woman incident to m in C_2. We may continue with this argument; if none of these matches can be added, then we get a loop $w_1, w_2, \ldots, w_r, w_{r+1} = w_1$, where adding (m, w_i) introduces an augmenting cycle C_i containing (m, w_{i+1}), for $i = 1, \ldots, r$. Note that m is indifferent between all $w_1, w_2, \ldots, w_r$. Then consider the following big cycle

$$C = \left[w_1 \xrightarrow{C_1 \backslash \{(m,w_1),(m,w_2)\}} w_2 \xrightarrow{C_2 \backslash \{(m,w_2),(m,w_3)\}} w_3 \cdots w_r \xrightarrow{C_r \backslash \{(m,w_r),(m,w_1)\}} w_1\right]$$

Note that C is available before the subroutine arrives at edge $(m, w) = (m, w_1)$ and it is possible that an edge appears more than once. For each w_i, $i = 2, \ldots, r + 1$, let m'_{i-1} and m_i be the other man (not m) incident to w_i in cycle C_{i-1} and C_i, respectively. Notice that each w_i weakly prefers m to m'_{i-1} and weakly prefers m_i to m, *i.e.*, w_i weakly prefers m_i to m'_{i-1}. Further, if one of the two preferences is strict, then w_i strictly prefers m_i to m'_{i-1}. Therefore, C is an augmenting cycle, a contradiction to the invariant that the algorithm will never produce any augmenting cycle in the process. (Note that the fact that there is an individual whose assignment can be strictly improved from C follows from the fact of augmenting cycles of each $C_1, \ldots, C_r$.)

We are now ready to prove the lemma.

PROOF OF LEMMA 5:

Note that if the condition in Step (3.b) or (3.c) is satisfied, *i.e.*, the assignment of a woman gets strictly improved, the subroutine returns and terminates. Hence, we will assume without loss of generality that in the course of ELIMINATE-PATH(P), all women involved are already fully-matched and are indifferent between their adjacent neighbors in path P.

Assume to the contrary that ELIMINATE-PATH(P) does not terminate. By Proposition 5, we know that the subroutine will not get stuck at any specific node. This implies that the subroutine will keep changing the statuses of pairs (*i.e.*, either matched or unmatched) through Step (3.a) and (3.i). Since the assignments of all women are kept at the same level, and the assignments of all men will not get worse off (by the definition of augmenting path)[18], we can divide the subroutine into *stages* where it moves from one stage to another if there is a man whose assignment get strictly improved. Since the subroutine does not terminate, it eventually gets into the last stage where no man will be able to improve his assignment. In other words, all individuals (men and women) are indifferent between their new assigned partner(s) and old partner(s) onwards.

Consider a moment when the subroutine is at the last stage, and let

$$P^* = [x^*, y^*, \ldots, w^*]$$

be the current remaining augmenting path (*i.e.*, after expansions and truncations in previous stages), where $e = (x^*, y^*)$ is the current considered pair (can be

[18]In the subroutine ELIMINATE-PATH, we will do a sequence of reassignments, e.g., first unmatch (w, m) then match (m, w'). Precisely speaking, the assignment of m first gets worse off then gets better off. Here saying m does not get worse off or gets strictly better off, we mean the overall assignment of m by combining these two consecutive reassignments.

either (m, w) or (w, m)) and w^* is the last woman of the augmenting path. Since it is guaranteed that the subroutine will never introduce any augmenting cycle, at this moment there is no augmenting cycle. In the rest of the proof we will restrict on the subroutine starting from this moment running on P^*.

For the initial augmenting path P^*, we set all pairs on it to be *unmarked*. In the process of the subroutine when P^* is updated, we mark/unmark pairs according to the following rules:

1) If the status of (x, y) is changed (become matched or unmatched), *mark* (x, y).

2) Recursively do the following: If a pair (x, y) is marked, *mark* all pairs in $\Gamma(y)$ (recall that $\Gamma(y)$ is the truncated path at a specific occurrence of y in the path P^*).

3) If P^* is expanded, *unmark* all expanded pairs.

Roughly speaking, the sign of a pair, marked or unmarked, denotes whether the subroutine has reached that pair or not in the path P^*. In particular, if the subroutine reaches to the last pair of P^*, all pairs have to be marked.

Let E^* denote a subset of pairs where $(x, y) \in E^*$ if in the process of running Eliminate-Path, the subroutine cannot change the status of (x, y) because otherwise it will bring a potential augmenting cycle. Note that (x, y) can be either (m, w) or (w, m). Certainly $E^* \neq \emptyset$. Let $(x_1, y_1), (x_2, y_2), \ldots, (x_\ell, y_\ell)$ be the order of pairs that are included into the subset E^* in the subroutine (note that $|E^*|$ is finite as the number of pairs is finite) and $C_1, C_2, \ldots, C_\ell$ be corresponding potential augmenting cycles. Note that $(x_i, y_i), (x_{i+1}, y_{i+1}) \in C_i$ for $i = 1, \ldots, \ell - 1$ (indeed, (x_{i+1}, y_{i+1}) is the reason that why the subroutine cannot move along with $\left[x_i \xrightarrow{C_i \setminus \{(x_i, y_i)\}} y_i\right]$ to reach y_i).

We have the following observations.

Claim 1. Consider any $(x, y) \in E^*$ and the moment when the subroutine is about to change its status but cannot do so because of a potential augmenting cycle C. Then x is indifferent between its two neighbors in C and the other neighbor right before it in P^*.

Proof. We will only prove the claim for the case $(x, y) = (m, w)$; the argument is similar when $(x, y) = (w, m)$. Let w_1 be the woman before m in the augmenting path P^* and w_2 be the other woman incident to m in C. By the definition of augmenting path P^* and augmenting cycle C, we know that m weakly prefers w to w_1 and w_2 to w; hence m weakly prefers w_2 to w_1. If m strictly prefers w_2 to w_1, then by the rule of Eliminate-Path, m is able to strictly improve his assignment, which contracts to the assumption that we are at the last stage where no one can improve his assignment anymore.

Claim 2. Consider any $(x, y) \in E^*$ and the moment when the subroutine is about to change its status but cannot do so because of a potential augmenting cycle C. Let y' be an individual who is able to strictly improve its assignment in $C = [x, \ldots, x', y', x'', \ldots, y, x]$ (by the above claim, $y' \neq x$). By the rule of the subroutine, we will expand the augmenting path to be

$$P^* = \left[\ldots, x \xrightarrow{C\setminus\{(x,y)\}} x', y', x'' \xrightarrow{C\setminus\{(x,y)\}} y, \ldots, w^*\right]$$

and all pairs between x and y are unmarked. Then all pairs after (x', y') (inclusive) in the current P^* are always unmarked from this moment through the course of the subroutine.

Proof. We will prove the claim for the first pair (x_1, y_1) added into E^*; the proof for the rest of pairs can be done in a similar way by induction. Initially all pairs in P^* are unmarked. The subroutine follows pairs in P^* one by one — changes their status and makes them marked — until the point when we get to (x_1, y_1). At this moment all pairs after $(x_1, y_1$ in P^* are still unmarked. Then the subroutine expands P^* according to the potential augmenting cycle $C_1 = [x_1, \ldots, x', y', x'', \ldots, y_1, x_1]$ and unmarks all expanded pairs.

We will first show that (x', y') is always unmarked. Since y is an individual who is able to strictly improve its assignment, we cannot change the status of (x', y') (*i.e.*, get it marked) directly, because otherwise its assignment will be strictly improved. Hence, the only way to mark (x', y') is through the second rule above by marking all pairs in a truncated cycle $\Gamma(z)$, where z is the node whose truncation contains (x', y'). By Proposition 3, we know that all individuals in

$\Gamma(z)$ are indifferent between their two neighbors; this implies that $(y', x'') \notin \Gamma(z)$. Since (x', y') and (y', x'') are consecutive pairs in P^*, the only way to separate them is to truncate at y', *i.e.*, $z = y'$. By Proposition 4, however, after such truncation, y' still strictly prefers its one neighbor to the other, which implies that it is still able to strictly improve its assignment, a contradiction.

Next consider any pair (x_0, y_0) after (x', y') in P^*, *i.e.*,

$$P^* = [\dots, x', y', x'', \dots, x_0, y_0, \dots, w^*]$$

since (x', y') is always unmarked, again the only way to mark (x_0, y_0) is through a truncated cycle. But that cycle has to include (x', y'), which is impossible.

We consider the following walk according to pairs $(x_1, y_1), (x_2, y_2), \dots, (x_\ell, y_\ell)$ in E^*: start from y_1 following the direction of $C_1 \setminus \{(x_1, y_1)\}$ until we get to y_2; next start from y_i following the direction of $C_i \setminus \{(x_i, y_i)\}$ until we get to y_{i+1} for $i = 2, \dots, \ell - 1$; finally start from y_ℓ following the direction of $C_\ell \setminus \{(x_\ell, y_\ell)\}$ until we get to the first y_k, where $(x_k, y_k) \in E^* \cap C_\ell$ (note that such (x_k, y_k) must exist, otherwise, the subroutine can reach y_ℓ, which contradicts to the above claim). Therefore, it forms a big cycle

$$C^* = \left[y_k \xrightarrow{C_k \setminus \{(x_k, y_k)\}} y_{k+1} \xrightarrow{C_{k+1} \setminus \{(x_{k+1}, y_{k+1})\}} y_{k+2} \to \cdots \to y_\ell \xrightarrow{C_\ell \setminus \{(x_\ell, y_\ell)\}} y_k \right]$$

Note that by the above claim, all pairs in the walk are unmarked. Further, it can be seen that C^* is an augmenting cycle, which contradicts to the fact that the algorithm never introduces an augmenting cycle. This completes the proof of the lemma.

D4. Analysis of the Algorithm

Again, the high level structure of the algorithm is to increase capacities of men and eliminate augmenting paths. While the algorithm may look involved, as the virtual capacity is not always monotonically increasing (e.g., in Step 6(c) of Increase-Cap and Step 3(c) of Eliminate-Path, we actually need to reduce the virtual capacities) and two subroutines may call each other, there is a simple, but crucial, idea behind the algorithm: the assignments of women keep improving (this is the exact reason that we do not want to introduce any augmenting cycle in the course of the algorithm). Therefore, at any moment of the algorithm, if a woman's assignment gets improved (e.g., Step 6(c) of Increase-Cap and Step 3(b), 3(c) of Eliminate-Path), the algorithm will abandon the current subroutine and restart the whole process (i.e., capacity increment and augmenting path elimination) starting from the current virtual capacity vector. Since every woman can improve her assignment at most n^2 times (as her capacity is at most n and every unit capacity can be improved at most n times), the whole algorithm will terminate.

It is easy to see that the three invariants listed at the beginning of the section are maintained in the course of the algorithm. Indeed, the last two (no augmenting cycle and women not worse off) hold trivially as they are guaranteed by the algorithm itself. For stability, in the subroutine Increase-Cap, when increasing the virtual capacity of m by one, we try to match m with a most preferred woman w where (m, w) forms a blocking pair. If w is not over-matched, then the resulting assignment is still stable. Otherwise, we delete a match (m', w) where m' is a least preferred man matched to w and reduce the virtual capacity of m' by one (Step (6.c) of Increase-Cap); this implies that the resulting assignment is still stable with respect to the new capacity vector. For the second subroutine Eliminate-Path, stability comes from the definition of augmenting path and the fact that when we delete a match (w, m), we know that m must be a least preferred man matched to w and w was over-matched (otherwise, when we add the match right before (w, m), the assignment of w gets strictly improved and the subroutine will run Step (3.b) or (3.c) to terminate). Therefore, the final returned assignment is stable.

When the algorithm Alg-Pareto-Stableterminates, by its rule there is no augmenting path. By the invariant that there is no augmenting cycle, we know that the returned assignment is Pareto-efficient. This yields the following result.

THEOREM 4: *The algorithm* Alg-Pareto-Stable *computes a Pareto-stable assignment in polynomial time.*