Operations Research

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

Stable Matching with Proportionality Constraints

Thành Nguyen, Rakesh Vohra

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

OPERATIONS RESEARCH
Vol. 67, No. 6, November–December 2019, pp. 1503–1519
ISSN 0030-364X (print), ISSN 1526-5463 (online)

Crosscutting Areas

Stable Matching with Proportionality Constraints

Thành Nguyen,[a] Rakesh Vohra[b]

[a] Krannert School of Management, Purdue University, West Lafayette, Indiana 47907; [b] Department of Economics, University of Pennsylvania, Philadelphia, Pennsylvania 19104

Contact: nguye161@purdue.edu, http://orcid.org/0000-0003-4536-3908 (TN); rvohra@seas.upenn.edu (RV)

Received: January 11, 2017
Revised: July 1, 2018; March 25, 2019
Accepted: June 30, 2019
Published Online in Articles in Advance: September 3, 2019

Subject Classifications: games/group decisions: cooperative
Area of Review: Games, Information, and Networks

https://doi.org/10.1287/opre.2019.1909

Abstract. The problem of finding stable matches that meet distributional concerns is usually formulated by imposing side constraints whose "right-hand sides" are absolute numbers specified before the preferences or number of agents on the "proposing" side are known. In many cases, it is more natural to express the relevant constraints as proportions. We treat such constraints as soft but provide ex post guarantees on how well the constraints are satisfied while preserving stability. Our technique requires an extension of Scarf's lemma, which is of independent interest.

Funding: The research of the first author is supported in part by the National Science Foundation (NSF) [Grant 1728165]. The research of the second author is supported in part by the NSF [Grant 1763307], Foundations of Fair Data Analysis.

Keywords: stable matching • diversity • Scarf's lemma

1. Introduction

A number of school choice programs use student preferences and school priorities to find a stable match of students to schools. There is also a desire to satisfy side constraints motivated by equity and distributional considerations. It is usual to express them in terms of proportions. As an example, in 1989, the city of White Plains, New York, required each school to have the same proportions of Blacks, Hispanics, and "others," a term that includes Whites and Asians. The plan allowed for a discrepancy among schools of only 5%. Similarly, the 2003 Cambridge, Massachusetts, Public School District's goal for matching was for each grade in each school to be within a range of plus or minus 15 percentage points of the district-wide percentage of low- socioeconomic status (SES) students.[1]

In prior work, these constraints are expressed as absolute numbers. For example, a requirement that at least 10% of students in a school with a capacity of 100 belong to a particular SES becomes a constraint that at least 10 students in the school belong to the relevant SES. This assumes that each school is fully allocated. If the number of students is less than the number of slots, this clearly cannot be true. This can happen. Chicago Public Schools, the third largest in the United States, for example, saw a drop in enrollment from 426,215 in 2000 to about 350,535 in 2013. It classifies almost 50% of Chicago's public schools as half empty.[2] Even if the number of students exceeds the number of slots, it does not guarantee that a school is fully allocated. Student preferences and their outside options also matter.

In this paper, we consider both lower and upper bound constraints on the proportions of students from each category. Constraints on proportions have the advantage of not committing to an absolute number as a target. However, under proportionality constraints, stable matchings need not exist. Furthermore, deciding if there exists a stable matching with respect to proportionality is also NP-hard. (See Theorem 3 and Remark 4.)

Absent stability, participants who find a better match than the one offered by a clearinghouse will "vote with their feet," leading to the unraveling of the entire market. Our paper proposes a new solution that treats the diversity constraints as soft but provides guarantees on how well the constraints are satisfied ex post while preserving stability. Our general result shows that the violation of proportionality constraints at a school h is bounded by $\frac{2}{\#\text{ assigned students at } h}$. Thus, if a school accepts more than 100 students, the matching violates the diversity constraints by at most 2%. However, the violation increases as the number of accepted students decreases. This has a natural interpretation as smaller schools having "softer" diversity constraints. What determines whether a school receives a small or large number of accepted students are student preferences. The set of "small" schools cannot be determined a priori from capacity information alone. A large-capacity school that is unpopular could end up with a number of students that is well below its capacity. Thus, relaxing its proportionality constraints allows it to recruit more students. We also show that the error bound of $O\left(\frac{1}{\#\text{ assigned students}}\right)$ is unavoidable if we want to maintain stability.

The quality of our approximation guarantee *depends* on the resulting matching. One could easily satisfy the proportionality constraints exactly by assigning no students at all. This is not true in our case as the stable matching we return is "maximal" in a certain sense. We show that, in order to increase the number of students matched over our matching *without* making any student worse off, one must alter either capacities or the proportionality constraints.

The need for approximation in our setting is not driven by the stability requirement but the indivisibility of students. Scarf (1967) extends the notion of stability to *fractional* matchings. We use that definition to find a maximal (in the sense mentioned) stable fractional matching that satisfies the proportionality constraints. Our approximation results from rounding that fractional stable matching so as to preserve stability. The associated approximation guarantee can be interpreted as a measure of how close the *integral* matching is to the fractional one. In fact, for each school, the difference between the number of students assigned to it under the fractional and integer matching is at most one.

We contrast our result with prior work next.

1.1. Related Work

Prior attempts to incorporate diversity considerations in matching fall into one of four categories that we list. We give illustrative examples of papers in each category.[3]

1.1.1. Ceilings. Distributional concerns are modeled as ceilings on the number of agents of each type from the "proposing" side that can be accepted. Ceiling constraints are generally considered easy to accommodate (see Abdulkadiroglŭ and Sönmez 2003). Regional capacity constraints are ceilings that apply to subsets on the "accepting" side rather than just individual members. As long as the subsets have a laminarity property, satisfying these constraints as well as the stability is not difficult. See Fleiner and Kamiyama (2016) and Kojima et al. (2018) for examples. Ceiling constraints, however, can disadvantage minorities. This is discussed in Hafalir et al. (2013).

1.1.2. Floors. Instead of imposing ceilings, one imposes floors on the number of proposers of a particular type. Satisfying floors and stability is generally difficult to do. This is discussed in Biró et al. (2010) and Huang (2010), which also describe some solvable cases.

1.1.3. Set-Asides. Instead of ceilings and floors, one sets aside capacity for each subgroup and then runs a separate matching process for each subgroup. This approach generally produces inefficiencies and other perverse effects in the resulting matching (see Kojima 2012 and Ellison and Pathak 2016), which are subsequently addressed by adjusting the set-asides either dynamically or ex post (see Fragiadakis and Troyan 2016 and Aygun and Turhan 2016 for examples).

1.1.4. Modifying Priorities. Instead of focusing on floors and ceilings, one modifies the choice function on the accepting side so as to favor various groups. If the modified choice function is specified in the right way, the deferred acceptance algorithm (or some variant) finds a stable matching. However, there is no ex post guarantee on realized distribution.[4] An example of this approach can be found in Ehlers et al. (2014). In lieu of an ex post guarantee, some authors focus on priorities that produce distributions that are closest to a target distribution; see Erdil and Kumano (2012) and Echenique and Yenmez (2015) for examples.

In the first three cases, the relevant "right-hand sides" are quantities specified *before* agents on the proposing side make their participation decisions. This may over-constrain the problem because the number of proposers who will be matched is endogenous. In the fourth case, targeted groups are "favored" but no ex post guarantee is provided on the realized distribution.

Echenique and Yenmez (2015), for example, introduced several classes of choice functions that reflect the diversity constraints but also satisfy the substitutes property. The substitutes property guarantees the existence of a stable solution. In particular, each school is assumed to choose the set of students such that their distribution is "closest" to the target distribution. This approach does not provide *any* guarantee on the ex post distribution. Example 1 illustrates this.

Example 1. Assume three schools h_1, h_2, h_3, each with capacity 200 and 300 students divided into three equally sized groups: A, B, and C. School h_1 desires that at least half of its students come from group A and at least half from group B. School h_2 desires that at least half of its students come from group B and at least half from group C. School h_3 desires that at least half of its students come from group C and at least half from group A. Therefore, the ideal distribution for each school is 100–100 for the corresponding pair of groups. Student preferences are as follows: each student in A prefers h_1 to h_3; each student in B prefers h_2 to h_1; each student in C prefers h_3 to h_2.

Consider the gross substitute choice function generated by an ideal point, defined in Echenique and Yenmez (2015), in which the school chooses a subset of applying students such that their type distribution is closest in Euclidean distance to the ideal one. The deferred acceptance algorithm stops at the first iteration and assigns all 100 students of group A to h_1, 100 students of group B to h_2, and 100 students of group C to h_3.

The core idea in our approach is to relax the integrality constraints of the matching problem so that students can be fractionally allocated to schools. We then use an extension of the notion of stability to fractional matchings from Scarf (1967). We focus on finding fractional stable matchings. To do this, we need to generalize Scarf's lemma, which is of independent interest.[5] Subsequently, this fractional matching is rounded into an integral stable matching that only violates the proportionality constraints (but not the capacity constraints) in a limited way.

Given a fractional stable solution, rounding it into an integer one that preserves stability can be done via a network flow algorithm; hence, it is polynomial time. The problem of finding a fractional stable matching is polynomial parity argument on directed graphs complete. However, Tang et al. (2018) and Nguyen and Vohra (2018) have implemented algorithms for finding such fractional stable solutions for large-sized problems and report runtimes that are small enough to be practical.

Section 2 defines lower bound proportionality constraints and introduces the notions of bilateral and coalitional stability. Section 3 derives some attractive properties of these concepts and shows that stable solutions, however, might not exist. Section 4 introduces Scarf's lemma and extends it to problems with lower bound proportional constraints. Section 5 describes the algorithm. Section 6 analyzes the stability of the rounded solution. Section 7 extends the result to upper bound proportionality constraints and gives an explicit algorithm. We conclude in Section 8.

2. Proportionality Constraints and Stability Concepts

To describe the stable matching problem, we label the two sides of the market doctors and hospitals. Denote by H the set of hospitals and D the set of doctors. Each doctor $d \in D$ has a strict preference ordering $\succ_d$ over $H \cup \{\emptyset\}$, where $\emptyset$ denotes the outside option for each doctor. If $\emptyset \succ_d h$, we say that hospital h is not acceptable for d. Each hospital $h \in H$ has capacity $k_h > 0$ and a strict priority ordering $\succ_h$ over elements of $D \cup \{\emptyset\}$. If $\emptyset \succ_h d$, we say d is not acceptable for h.

A matching is an assignment of each doctor to a hospital or the doctor's outside option; each hospital is assigned an acceptable set of doctors that does not exceed its capacity. Given a matching μ, let $\mu(h)$ denote the subset of doctors matched to h and $\mu(d)$ denote the position that d obtains in the matching. Thus, μ satisfies

$$\begin{aligned} &\text{(i)} \quad \mu(d) \succ_d \emptyset \\ &\text{(ii)} \quad \text{if } d \in \mu(h) \text{ then } d \succ_h \emptyset \\ &\text{(iii)} \quad |\mu(h)| \le k_h. \end{aligned} \tag{1}$$

Next, we introduce the proportionality constraints for hospitals. For each hospital h, let $D^h := \{d : d \succ_h \emptyset, h \succ_d \emptyset\}$ be the set of doctors acceptable to h and who find h acceptable. Each D^h is partitioned into T_h sets: $D^h = D_1^h \cup D_2^h \cup .. \cup D_{T_h}^h$. Different hospitals can have different partitions. A doctor $d \in D_t^h$ is said to be of type t for hospital h. In the school choice context, in which hospitals correspond to schools and doctors to students; a type can represent an SES category. Allowing different schools to have different partitions allows schools the flexibility to use categories depending on the proximity of the student's residence to the school.[6]

The lower bound proportionality constraint at each hospital $h \in H$ is

$$\alpha_t^h \cdot |\mu(h)| \le |\mu(h) \cap D_t^h| \ \forall t = 1, \dots, T_h, \text{ where } 0 \le \alpha_t^h \le 1, \sum_t \alpha_t^h \le 1. \tag{2}$$

A matching satisfying (1) and (2) is called *feasible*.[7] Constraint (2) ensures that the proportion of doctors of each type in D^h who are matched to hospital h is above some threshold. These constraints don't need to hold for each hospital-type pair. This can be captured by setting $\alpha_t^h = 0$. Unlike floor constraints, the left-hand side of (2) is endogenous.

2.1. Bilateral Stability

Next, we introduce our new notion of bilateral stability. In the presence of (2), one needs to modify the usual notion of blocking to rule out blocking pairs that violate (2). A natural way to define (h, d) to be a blocking pair is that d prefers h to d's current match and either (i) h can accept d without violating its capacity and proportionality constraints, or (ii) h can replace a lower ranked doctor (according to $\succ_h$) with d so that h's capacity and proportionality constraints are not violated. This is a weak notion of stability that can lead to a matching that is "wasteful" as shown in the following example.

Example 2. Consider a single hospital h with capacity 100 and 100 doctors $d_1, \dots, d_{100}$. All doctors strictly prefer to be matched to h than remain unmatched, and the priority order of the hospital is $d_1 \succ_h d_2 \succ_h \cdots \succ_h d_{100}$. The set of doctors are divided into two subgroups, $D_1^h = \{d_1, d_3, \dots, d_{99}\}$ and $D_2^h = \{d_2, d_4, \dots, d_{100}\}$. The proportionality constraint is that at least 50% of the doctors in each subgroup are accepted.

Under this naive stability definition, the matching that assigns d_1, d_2 to h would be stable because no *single* doctor can form a blocking coalition with h because of the proportionality constraints. This is undesirable compared with the stable matching that assigns all doctors to h.

To overcome the waste of 98 positions in Example 2, we need to require each hospital not to waste positions if it can accept a set of doctors who demand it without violating the constraints. This means that one needs to allow for "coalitional" blocks that contain multiple students.

Therefore, we introduce a notion of bilateral stability that additionally requires hospitals to reach their "effective" capacity. We show that this stability notion implies coalitional stability and a nonwastefulness property. To this end, we define for each feasible matching a set of protected doctors and for each hospital its effective capacity. A protected doctor can never be rejected by any hospital to which the doctor is matched in favor of another doctor.

We start with the notion of waitlisted doctors.

Definition 1 (Waitlisted Doctor)**.** Given a feasible matching μ, a doctor d is *waitlisted* at h if d and h are mutually acceptable and either d is unmatched or $h \succ_d \mu(d)$.

Thus, the waitlisted doctors of a hospital h are those who prefer to be matched with h over their current outcome. In other words, each of these doctors would like to form a blocking coalition with h.

Fix a hospital h and suppose the set of doctors of type t at this hospital, D_t^h, does not contain any waitlisted doctors. Then h cannot increase the number of admitted doctors of type t because all such doctors are already matched to a more preferred hospital. In this case, the proportionality constraint corresponding to type t, $\alpha_t^h|\mu(h)| \leq |\mu(h) \cap D_t^h|$, implies that the total number of doctors that hospital h can accept is at most $\frac{1}{\alpha_t^h}|\mu(h) \cap D_t^h|$. This motivates the following definition of a hospital's effective capacity.

Definition 2 (Effective Capacity)**.** Consider a feasible matching μ and a hospital h. Let T_0 be the set of types t, such that D_t^h contains no waitlisted doctor. Denote hospital h's effective capacity with respect to μ by k_h^μ, where

$$k_h^\mu := \min\left\{k_h, \min_{t \in T_0} \frac{1}{\alpha_t^h}|\mu(h) \cap D_t^h|\right\}, \text{ and if}$$
$$T_0 = \emptyset \text{ or } \alpha_t^h = 0, \text{ then } k_h^\mu := k_h.$$

Remark 1. Given μ, k_h^μ is an upper bound on the number of positions that h can fill by accepting more waitlisted doctors without violating any proportionality constraints. Because μ is feasible, it satisfies both the capacity and the side constraints. Thus, it is clear that $|\mu(h)| \leq k_h^\mu$. Furthermore, from the preceding definition, if hospital h is not at its effective capacity with respect to μ, $|\mu(h)| < k_h^\mu$; then $|\mu(h)| < k_h$, and there is no $t \in T_0$ such that the proportionality constraint corresponding to D_t^h binds, that is, $|\mu(h)| = \frac{1}{\alpha_t^h}|\mu(h) \cap D_t^h|$.

In what follows, when μ is clear from context we omit the qualifier "with respect to μ" when referring to a hospital's effective capacity. Next, we define the types of doctors who are protected.

Definition 3 (Protected Type of Doctors)**.** Given a feasible matching μ, the set of type t doctors at hospital h is *protected* with respect to μ if (2) binds with respect to the effective capacity, that is

$$|\mu(h) \cap D_t^h| = \alpha_t^h \cdot k_h^\mu. \tag{3}$$

In what follows, if μ is clear from its context, we omit the qualifier "with respect to μ" when referring to protected doctors.

Example 3. We illustrate the definition of effective capacity and protected doctor using Figure 1. Consider the group $D_2^{h_1}$. Doctor d_7 is the only member of $D_2^{h_1}$ not matched to h_1, and this doctor prefers h_2 to h_1. Thus, the group $D_2^{h_1}$ does not contain any waitlisted doctor. This means that h_1 cannot admit more doctors from $D_2^{h_1}$. Together with the proportionality constraint for group $D_2^{h_1}$, it implies that hospital h_1 cannot admit more than $2/(\alpha_2^{h_1}) = 6$ doctors. The effective capacity of h_1 becomes six instead of its original capacity of eight.

As the effective capacity of h_1 is six, the proportionality constraints corresponding to $D_1^{h_1}$ and $D_2^{h_1}$ bind; thus, $D_1^{h_1}$ and $D_2^{h_1}$ are protected. If a type is protected, it means that the hospital cannot decrease the number of doctors of this type who are matched to it. Why? If the hospital decreases this number, it also needs to decrease the total number of doctors to satisfy the proportionality constraints.

To motivate the definition of stability, consider d_3, who prefers to be matched with h_1 rather than d_3's current match, h_2. As h_1 is at its effective capacity, h_1 must reject a doctor currently matched to h_1 in order to accept d_3. To satisfy the proportionality constraint,

Figure 1. (Color online) An Example of Effective Capacity and Protected Doctors

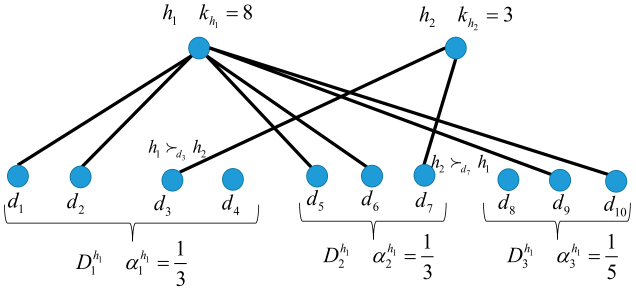

h_1 can only reject either a doctor of the same type as d_3 (d_1 or d_2) or an unprotected doctor (d_9 or d_{10}). The following definition requires this matching to be stable if h_1 has no incentive to replace d_3 with any of these doctors. That is, h_1 prefers each of d_1, d_2, d_9, and d_{10} to d_3.

Next, we introduce the notion of stability motivated by Example 3.

Definition 4 (Bilateral Stability)**.** A feasible matching μ is bilaterally stable if it satisfies the following two conditions:

1. Each hospital with a nonempty waitlist is at its effective capacity, that is, $|\mu(h)| = k_h^{\mu}$.
2. If d_a is on the waitlist of h, $d_r \in \mu(h)$, and $d_a \succ_h d_r$, then d_r is protected, and d_a and d_r are not of the same type.

The first condition ensures that no "in-demand" hospital can increase the number of doctors it accepts. The second condition ensures that, if h tries to reject d_r to accept a better d_a, then it violates the side constraint. This is because d_r is protected, and d_a is not of the same type as d_r; thus, by rejecting d_r, hospital h violates the side constraint of the group containing d_r.

In Definition 2, effective capacities can be fractional. As highlighted in Remark 3, this makes bilateral stability a very strong concept because, by condition 1 of Definition 4, if a matching is bilaterally stable, then effective capacities are integral. One can modify this definition, for example, by taking the floor function of the bound $\lfloor \frac{1}{\alpha_t^h} |\mu(h) \cap D_t^h| \rfloor$ without changing our main results. This is because our method starts from a fractional solution for which a fractional effective capacity is natural. We then round the fractional solution and change the parameters α_t^h accordingly so that the effective capacities are also integral.

2.2. Coalitional Stability

To define coalitional stability, we must specify each hospital's h preferences over subsets of D in a way that respects $\succ_h$ as well as its capacity and proportionality constraints and nothing more. This can be done via a choice function, $Choice_h(.) : 2^D \to 2^D$.

Definition 5. The choice function of h on a subset of acceptable doctors D^*, denoted $Choice_h(D^*)$, is a maximum cardinality subset of D^* that satisfies h's capacity constraints and proportionality constraints. If there are multiple such subsets, then $Choice_h(D^*)$ is the best one in the lexicographical order according to $\succ_h$.[8]

Given the choice function of the hospitals, next we consider the standard concept of coalitional stability in the many-to-one matching setting, which means that no group of doctors and possibly multiple hospitals can deviate from the matching to obtain better payoffs.

Definition 6 (Coalitional Stability)**.** A matching μ is coalitionally stable if, for every set of doctors D^* who prefer h to their current match, $Choice_h(\mu(h) \cup D^*) = \mu(h)$.

Remark 2. In contrast to the solution in Example 1, under the choice function of Definition 5, there is a unique coalitionally stable solution in which hospitals h_1, h_2, and h_3 get their 50 highest priority doctors from groups A, B, and C, respectively, and the remaining doctors are allocated to hospitals so that the proportionality constraints are satisfied.

3. Properties of Stable Matchings

In this section, we show that bilateral stability implies coalitional stability. Subsequently, we demonstrate that coalitionally stable matchings are nonwasteful and, in a certain sense, cannot be improved upon. This means that bilateral stable matchings inherit the same properties. These properties come at a cost. We show by an example that a stable matching need not exist. At the end of this section, we illustrate how our approach overcomes this problem.

Theorem 1. *If μ is a bilateral stable matching, then μ is also coalitionally stable.*

We need to show that, for any group of doctors D^* on the waitlist of h, $Choice_h(\mu(h) \cup D^*) = \mu(h)$. First notice that, because h is at its effective capacity, h cannot increase the number of doctors without violating the proportionality constraints. Thus,

$$|Choice_h(\mu(h) \cup D^*)| = |\mu(h)|.$$

Let $D_A := Choice_h(\mu(h) \cup D^*) \setminus \mu(h)$ be the set of accepted doctors. Let $D_R := \mu(h) \setminus Choice_h(\mu(h) \cup D^*)$ be the set of rejected doctors. Assume D_A and D_R are not empty. Let d_{min} be the lowest ranked doctor among D_R according to $\succ_h$. Because h breaks ties according to the lexicographical order, all doctors in D_A must be more preferred than d_{min}.

If d_{min} is unprotected, it contradicts the definition of bilateral stability because any $d_a \in D_A$ can replace d_{min} at h. If d_{min} is protected, then in order to satisfy this type's proportionality constraint, h needs to accept a doctor of the same type. Thus, there should be a $d_a \in D_A$ that is of the same type as d_{min}. In this case, we can replace d_{min} with a better doctor, d_a, which contradicts the definition of bilateral stability. Hence, D_A and D_R are empty, and thus, $Choice_h(\mu(h) \cup D^*) = \mu(h)$.

Remark 3. The converse of Theorem 1 is false. Consider a hospital with 99 positions and 50 men and 50 women applicants. The hospital needs to accept at least 50% men and 50% women. The coalitionally stable matching is to accept the best 49 men and the best 49 women. Under this matching, the hospital has both men and women on the waitlist and the hospital leaves one position unfilled. Hence, this matching is not bilaterally stable because the effective capacity of the hospital according to Definition 2 is 99. In our solution, we first find a fractional matching which assigns 49.5 best men and 49.5 best women to the hospital, which matches the hospital's effective capacity. We round the fractional solution to 50 men and 49 women (or 49 men and 50 women). The proportionality constraints then are modified to $\frac{50}{99}$ and $\frac{49}{99}$. Under these new constraints, the matching of 50 men and 49 women is bilaterally stable.

In the remainder of the paper, for short, we say that a matching is stable if it is bilaterally stable; otherwise, we specifically label it coalitionally stable.

The next theorem demonstrates the nonwastefulness and nonimprovability property of coalitionally stable matchings.

Theorem 2. *Given a feasible matching that is coalitionally stable, there is no other feasible matching (not necessarily stable) that assigns more doctors to hospitals such that no doctor is worse off.*

Theorem 2 shows that, if one would like to assign more doctors to hospitals, then some doctors will be worse off. Hence, there is a trade-off between the number of doctors assigned and the welfare of each doctor, and a coalitionally stable matching is on the efficient frontier of that trade-off.[9]

Proof of Theorem 2. Let μ be a feasible and stable matching. Assume μ' assigns more doctors to hospitals, and under μ', no doctors are worse off than in μ. Then there must be at least one hospital h that obtains more students under μ': $|\mu'(h)| > |\mu(h)|$.

Let $D^* := \mu'(h) \setminus \mu(h)$. Because no doctors are worse off under μ', each doctor in D^* prefers h to the doctor's match in μ. Observe that $\mu(h) \cup D^* = \mu(h) \cup \mu'(h)$ contains $\mu'(h)$, which is a set of doctors that satisfies the capacity and proportional constraints and has a larger cardinality than $\mu(h)$; thus, $Choice_h(\mu(h) \cup D^*) \neq \mu(h)$. This shows that μ is not coalitionally stable, a contradiction.

Next, observe that the choice function of Definition 5 violates the substitutes property; hence, a stable matching need not exist. This is captured in the following result.

Theorem 3. *A stable matching (bilateral or coalitional) need not exist.*

To see this, we consider an example similar to Example 1; that is, there are three schools h_1, h_2, h_3, each with capacity 200 and the students divided into three equally sized groups: A, B, and C. School h_1 desires that at least half of its students come from group A and at least half from group B. School h_2 desires that at least half of its students come from group B and at least half from group C. School h_3 desires that at least half of its students come from group C and at least half from group A. Student preferences are as follows: each student in A prefers h_1 to h_3; each student in B prefers h_2 to h_1; each student in C prefers h_3 to h_2. The only difference from Example 1 is that each group A, B, C contains 99 students instead of 100.

A stable matching does not exist. For a contradiction, suppose otherwise. Because of the proportionality constraints, each hospital accepts an even number of students, and thus, the total number of students accepted is even. Therefore, there is at least one student rejected. There cannot be two students from two different groups rejected because they can form a blocking coalition with at least one hospital. Assume, without loss of generality, only doctors from group C are rejected. This means that no doctors in group B are assigned to h_1 because otherwise, this doctor, together with the rejected doctor in group C, can form a blocking coalition with hospital h_2. Because of the proportionality constraint for h_1, no doctors from group A are assigned to h_1. Therefore, all 99 doctors from group A are assigned to h_3. This leads to a contradiction because the proportionality constraint at h_3 implies that all 99 doctors from group C are also assigned to h_3.

Remark 4. Biró et al. (2010) shows that deciding if a stable matching exists for a school admissions problem with a lower bound is NP-hard. In particular, in addition to its (upper) bound, each school has a lower bound. Call a school open if the number of students assigned to it is between the lower and the upper bound. A closed school cannot accept any students. A matching requires every school is either open or closed. A matching is stable in this context if, in addition to the standard blocking pair, it also allows for a blocking coalition of a closed school and a group of students of size at least the lower bound.

Theorem 1 of Biró et al. (2010) shows that the problem is NP-hard even when each school has an upper and lower bound of either (1, 1) or (3, 3). We reduce this problem to a matching problem with proportionality constraints as follows. For each school, introduce a dummy student who belongs to a special group and is only interested in that school. If the school's upper and lower bounds are one and one, respectively, then the school has a capacity of two, and it requires at least 50% of students from the special group and 50% from the "normal" students. If the school's upper and lower bounds are three and three, respectively, then the school has a capacity of four, and it requires at least 25% of students from the special group and 75% from the normal students. It is easy to see the correspondence between the school admissions problem with lower bound and the stable matching with proportionality constraints. This reduction shows that deciding if a coalitionally stable matching exists with proportionality constraints is NP-hard.

In the remainder of this paper, we overcome the problem of nonexistence in two steps. First, assume that the doctors are divisible. Using an extension of Scarf's lemma, we show the existence of a *fractional dominating solution* that satisfies the proportionality constraints. Domination is the fractional analog of stability. Second, we round the fractional solution to an integral one. Rounding violates the proportionality constraints but in a minimal way.

To illustrate these ideas, consider the example in the proof of Theorem 3. First, suppose doctors can be fractionally allocated to hospitals. In this example, there is a unique dominating solution that assigns to h_1, h_2, and h_3 the $49\frac{1}{2}$ best doctors from groups A, B, and C, respectively, and the remaining doctors are allocated wholly to each hospital. Using this allocation, the second step in our algorithm finds an integral one by rounding it to the nearest integers. In this example, $49\frac{1}{2}$ can be rounded to either 49 or 50, and thus, at least one hospital receives a distribution of students that slightly violates its proportionality constraint.

This example can be generalized so that each group A, B, and C contains $2k+1$ doctors. This shows that, to obtain a stable matching, one needs to violate the proportionality constraints, and the error bound of $O\left(\frac{1}{\#\text{ assigned students}}\right)$ is unavoidable.

4. Fractional Stable Matching

This section is the technical heart of the paper. We use Scarf's lemma to obtain a fractional stable matching. A direct application of the lemma does not accommodate (2). We derive what we call a "conic representation" of the lemma. This is both new and general enough to apply to other types of side constraints, but in this paper, we confine ourselves to proportionality constraints.

We first describe Scarf's lemma and show how it can be applied to matching.

4.1. Scarf's Lemma

Definition 7. Let $\mathcal{A}$ be an $m \times n$ nonnegative matrix with at least one positive entry in each row and column and let $b \in \mathbb{R}^m_+$ be a positive vector. Associated with each row i of $\mathcal{A}$ is a strict ranking $\succ_i$ over the columns in $\{1 \le j \le n : \mathcal{A}_{ij} > 0\}$. Let $\mathcal{P} = \{x : x \ge 0, \mathcal{A}x \le b\}$. We say $x \in \mathcal{P}$ *dominates* column j if there exists a row i such that

- $\mathcal{A}_{ij} > 0$ and the constraint i binds, that is, $(\mathcal{A}x)_i = b_i$;
- $k \succ_i j$ for all columns $k \ne j$ such that $\mathcal{A}_{ik}x_k > 0$.

In this case, we say that x dominates column j via row i.

Theorem 4 (Scarf 1967). *There exists an extreme point of $\mathcal{P}$ that dominates every column of $\mathcal{A}$.*

To understand the connection of domination to stability, it is helpful to consider the matching problem without side constraints. For each $d \in D \cup \{\emptyset\}$ and $h \in H \cup \{\emptyset\}$, let $x(d,h) = 1$ if we assign d to h and zero otherwise. Now, each doctor $d \in D$ can be assigned to at most one hospital:

$$\sum_{h \in H \cup \emptyset} x(d,h) \le 1 \ \forall d \in D. \tag{4}$$

Each hospital h can be assigned at most k_h doctors:

$$\sum_{d \in D \cup \emptyset} x(d,h) \le k_h \ \forall h \in H. \tag{5}$$

Each inequality (4) inherits the order that doctor d, that is, $\succ_d$, has over $H \cup \{\emptyset\}$. Each inequality (5) inherits the priority ordering that hospital h, that is, $\succ_h$, has over $D \cup \{\emptyset\}$. As follows, system (4) and (5) along with a nonnegativity restriction on the x variables satisfies the conditions of Scarf's lemma.

Now, as is well known, every nonnegative extreme point of (4) and (5) corresponds to a matching. By Scarf's lemma, one of these extreme points, x^*, say, is dominating. To show that the matching implied by x^* is stable, suppose a pair (d^*, h^*) such that $x^*(d^*, h^*) = 0$. We show that (d^*, h^*) cannot be a blocking pair. By domination, there must exist a binding constraint from (4) and (5). Either it is indexed by d^* or h^*, say, d^*. Then

$$\sum_{h \in H \cup \{\emptyset\}} x^*(d^*, h) = 1.$$

As $x^*(d^*, h^*) = 0$, it follows that exactly one $h' \in H \cup \{\emptyset\}$ exists such that $x^*(d^*, h') = 1$. Further, by domination, $h' \succ_{d^*} h^*$, which means (d^*, h^*) cannot be a blocking pair.

The side constraints in (2) can be written as

$$\alpha^h_t \cdot \sum_{d \in D} x(h,d) \le \sum_{d \in D^h_t} x(h,d) \ \forall t = 1, \ldots, T_h, \ \forall h \in H. \tag{6}$$

It is tempting but incorrect to append inequality (6) to (4) and (5) and invoke Scarf's lemma. If we rewrite (6) in the form $\mathcal{A}x \le b$, the relevant inequalities have negative coefficients, and the corresponding coordinates of b are zero. Therefore, Scarf's lemma does not apply. Also, the condition of stability in our setting is now endogenous and depends on the effective capacity of a hospital. Because of this, it is not clear how one can apply Scarf's lemma directly as in Nguyen and Vohra (2018).

4.2. Conic Representation

We need another approach to determine a dominating solution of (4) and (5) that satisfies (6). We exploit the fact that the constraints in (6) form a polyhedral cone. Therefore, any point in the cone can be expressed as a nonnegative linear combination of its generators. We give a high-level description first.

Consider the problem of finding a dominating solution satisfying resource constraints $\mathcal{A}x \le b$ and side constraints $\mathcal{M}x \ge 0$. The set $\{x \in \mathbb{R}^n_+ | \mathcal{M}x \ge 0\}$ is a polyhedral cone and can be rewritten as $\{\mathcal{V}z | z \ge 0\}$, where $\mathcal{V}$ is a finite nonnegative matrix. The columns of $\mathcal{V}$ correspond to the *generators* of the cone $\{x \in \mathbb{R}^n_+ | \mathcal{M}x \ge 0\}$. The "trick" is to apply Scarf's lemma to $\mathcal{Q} = \{z \ge 0 : \mathcal{A}\mathcal{V}z \le b\}$. To do so, we need to endow each row of $\mathcal{A}\mathcal{V}$ with an ordering so that domination with respect to this system corresponds to stability.

Figure 2 gives a geometric illustration of (4)–(5). The polytope in Figure 2 corresponds to the matching polytope (4) and (5). The inequalities (6) are represented by the cone. The conic version of Scarf's lemma gives us a fractional dominating solution, x^*, say, that is inside the cone but on the boundary of the matching polytope. In particular, there is no other dominating point in the matching polyhedron that vector dominates x^*. In other words, x^* maximizes a suitable positive weighted sum of doctor's utilities. Our rounding algorithm, described in the next section, rounds x^* into an integral solution on the boundary of the polytope but possibly outside the cone.

Next, we show how to determine the generators of the cone associated with (6).

4.2.1. Generators of a Cone.

The following is standard (see Nemhauser and Wolsey 1988). See Figure 3 for an illustration.

Lemma 1. *For any matrix $\mathcal{M}$, if the set $\{x \in \mathbb{R}^n_+ | \mathcal{M}x \ge 0\}$ contains a nonzero vector, there exists a finite set of nonnegative vectors $\mathcal{V}$ such that this set can be expressed as $\{\sum_{v_i \in \mathcal{V}} z_i v_i | z_i \ge 0\}$. The set of vectors, $\mathcal{V}$, are called the generators of $\{x \in \mathbb{R}^n_+ | \mathcal{M}x \ge 0\}$.*

Figure 2. Geometric Presentation of Conic Scarf's Lemma

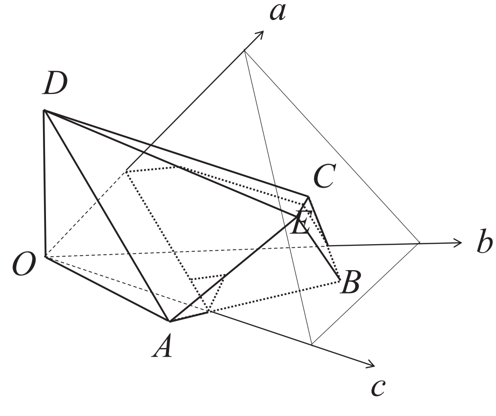

Note. Cone (O, a, b, c) intersects with polytope $(0, A, B, C, D, E)$.

Figure 3. Cone and Its Generators

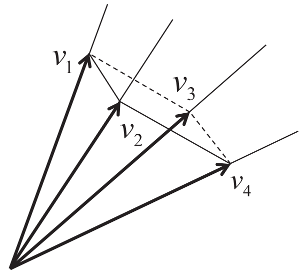

The proportionality constraints are of the form $\mathcal{M}x \ge 0$. To determine the generators of (6), fix a hospital $h \in H$ and focus on

$$\alpha^h_t \cdot \sum_{d \in D} x(h,d) \le \sum_{d \in D^h_t} x(h,d) \;\; t = 1, \ldots, T_h. \tag{7}$$

It is straightforward to see that the generators can be described in this way:

1. Select one doctor from each D^h_t and call it d_t.
2. Select an extreme point of the system

$$\sum_{t=1}^{T_h} v(d_t, h) = 1, \;\; \alpha^h_t \le v(d_t, h) \;\forall t = 1, \ldots, T_h.$$

An extreme point can be determined using the following two-step procedure.

a. Choose an index $t^* \in \{1, \ldots, T_h\}$ and set $v(d_{t^*}, h) = 1 - \sum_{t \ne t^*} \alpha^h_t \ge \alpha^h_{t^*}$.

b. For all $t \ne t^*$, set $v(d_t, h) = \alpha^h_t$.

As there are T_h types of doctors and each type contains $|D^h_t|$ doctors, the number of generators associated with hospital h can be as large as $T_h \cdot \prod_t |D^h_t|$.[10]

Let $\mathcal{V}_h$ be the set of generators associated with hospital h. Each $v \in \mathcal{V}_h$ has T_h nonzero coordinates and can be interpreted as a probability vector. Thus, from the point of view of each $h \in H$, each $v \in \mathcal{V}_h$ can be seen as a lottery over doctors in D.

We illustrate with an example.

Example 4. Suppose $H = \{h_1, h_2\}$ and two doctors $d_1 \in D^{h_1}_1$ and $d_2 \in D^{h_1}_2$. Hospital h_2 considers all doctors to be the same type, that is, $D^{h_2}_1 = \{d_1, d_2\}$.

The only proportionality constraints are imposed on h_1: the number of type 1 doctors should be at least one third of the total number of doctors assigned to h_1. This constraint can be written as

$$\frac{1}{3}[x(d_1, h_1) + x(d_2, h_1)] \le x(d_1, h_1).$$

The set of generators for this constraint are

$$\mathcal{V}_{h_1} = \{(1/3, 2/3); (1, 0)\} := \{v_1, v_2\}.$$

We can interpret $v_1 = (1/3, 2/3)$ to mean assigning d_1 and d_2 to h_1 with probability 1/3 and 2/3, respectively.

All solutions satisfying the proportionality constraints at h_1 can be expressed as a linear combination of v_1 and v_2.

There are no proportionality constraints imposed on h_2. This is the same as setting $\alpha_1^{h_2} = 0$. The set of generators for h_2 are

$$\mathcal{V}_{h_2} = \{(1,0),(0,1)\} := \{v_3, v_4\}.$$

We interpret $v_3 = (1,0)$ to mean assigning d_1 to h_2 with probability one and d_2 to h_2 with probability zero.

4.3. Conic Version of Scarf's Lemma

Associated with each hospital $h \in H$ is a set $\mathcal{V}_h$ of generators. Let $\mathcal{V}$ be the matrix whose columns correspond to the generators in $\bigcup_{h\in H}\mathcal{V}_h$. Let $\mathcal{A}$ be the constraint matrix associated with (4) and (5). Each row of the matrix $\mathcal{A}\cdot\mathcal{V}$ corresponds to an element of either D or H. The columns of $\mathcal{A}\cdot\mathcal{V}$ correspond to the set of generators. For each $h \in H$, a column in $\mathcal{A}\cdot\mathcal{V}$ that corresponds to $v \in \mathcal{V}_h$ will have a "1" in the h^{th} row and $v(d,h)$ in the d^{th} row. All other entries in that column will be zero. Let $z \in \mathbb{R}_+^{|\bigcup_{h\in H}\mathcal{V}_h|}$ be a nonnegative weight vector on the set of generators. The constraints $\mathcal{A}\cdot\mathcal{V}\cdot z \le b$ can be interpreted as follows:

- For each hospital h, the total weight of generators in $\mathcal{V}_h$ is at most k_h.
- For each doctor d, the weight of generators that assigns d to a hospital is at most one.

Example 5. Consider Example 4. Suppose $k_h = 2$. The polyhedron $\mathcal{Q}$ is displayed here.

$$\begin{array}{c} \\ d_1\\ d_2\\ h_1\\ h_2\end{array}\begin{array}{c}\begin{array}{cccc} v_1 & v_2 & v_3 & v_4\end{array}\\ \begin{bmatrix} 1/3 & 2/3 & 1 & 0\\ 2/3 & 1/3 & 0 & 1\\ 1 & 1 & 0 & 0\\ 0 & 0 & 1 & 1\end{bmatrix}\end{array} \cdot z \le \begin{bmatrix}1\\1\\2\\2\end{bmatrix}.$$

To invoke Scarf's lemma, we need each row of $\mathcal{A}\mathcal{V}$ to have a strict ordering over the columns, that is, generators, in its support. The support of each generator corresponds to one hospital and a coalition of doctors at most one of each type.

1. For each $h \in H$, we order the generators in $\mathcal{V}_h$ lexicographically. Given $v, v' \in \mathcal{V}_h$, among the doctors who are assigned by v, v' with positive probability to h, let d_1 and d_1' be the lowest ranked doctors (according to $\succ_h$). If $d_1 \succ_h d_1'$, then h ranks v over v'. We write this as $v \succ_h v'$. If $d_1 = d_1'$, we compare $v(d_1,h)$ and $v'(d_1,h)$. If $v(d_1,h) > v'(d_1',h)$, then h ranks v over v', that is, $v \succ_h v'$. If it is the reverse, then $v' \succ_h v$. If $v(d_1,h) = v'(d_1',h)$, move to the next lowest ranked doctor in each generator and so on. Because $v \ne v'$, this procedure must terminate in an unambiguous ordering.

2. For each $d \in D$ and any $v, v' \in \cup_{h\in H}\mathcal{V}_h$, we rank v above v' if v assigns d to a higher ranked hospital (according to $\succ_d$) than v' does. We write this as $v \succ_d v'$. If $v, v' \in \mathcal{V}_h$ for some $h \in H$, then $v \succ_d v'$ if $v(d,h) > v'(d,h)$ and the reverse otherwise. If $v(d,h) = v'(d,h)$, we order v and v' in the same way that h would.

Example 6. Continuing Example 4, let $h_1 \succ_d h_2$ for all $d \in D$ and $d_1 \succ_h d_2$ for all $h \in H$. The order of each element of $D \cup H$ over the set of generators is displayed here.

$$\begin{array}{c} \\ d_1\\ d_2\\ h_1\\ h_2\end{array}\begin{array}{c}\begin{array}{cccc} v_1 & v_2 & v_3 & v_4\end{array}\\ \begin{bmatrix} 1/3 & 2/3 & 1 & 0\\ 2/3 & 1/3 & 0 & 1\\ 1 & 1 & 0 & 0\\ 0 & 0 & 1 & 1\end{bmatrix}\end{array} \cdot z \le \begin{bmatrix}1\\1\\2\\2\end{bmatrix};\ order: \begin{array}{l} v_1 \succ_{d_1} v_2 \succ_{d_1} v_3\\ v_2 \succ_{d_2} v_1 \succ_{d_2} v_4\\ v_2 \succ_{h_1} v_1\\ v_3 \succ_{h_2} v_4.\end{array}$$

Consider the order for h_1 on v_1, v_2. Because they both assign d_1 and d_2 to h_1, we need to compare the probability of assigning d_2, which is the worst doctor for h_1. Because v_1 assigns d_2 with higher probability, $v_2 \succ_{h_1} v_1$.

Consider the order for d_1 on v_1, v_2, and v_3. Because v_1, v_2 assigns d_1 to h_1 and v_3 assigns d_1 to h_2, thus, d_1 prefers both v_1 and v_2 to v_3. Between v_1 and v_2, the one that assigns with a lower probability is better, and thus, $v_1 \succ_{d_1} v_2$.

Remark 5. By Theorem 4, there exists a dominating solution of $\mathcal{Q}$, call it z^*. Let

$$\mathcal{V}^* = \{v \in \mathcal{V} : z_v^* > 0\} \tag{8}$$

be the set of generators with positive z^* weight. Denote by $\mathcal{V}_h^*$ the subset of generators in $\mathcal{V}^*$ that assigns doctors to h and similarly denote by $\mathcal{V}_d^*$ the subset of generators in $\mathcal{V}^*$ that assign d to a doctor. Because z^* is a dominating solution, for every generator v that assigns a group of doctors $d_1, \ldots, d_{T_h}$ to h, one of the following must be true:

- The constraint $\mathcal{Q}$ corresponding to h binds. That is, h is fully allocated, and h ranks all the generators in $\mathcal{V}_h^*$ over v.
- There is a $d_i \in \{d_1, \ldots, d_{T_h}\}$ such that the constraint in $\mathcal{Q}$ corresponding to d_i binds, and d_i ranks all the generators in $\mathcal{V}_{d_i}^*$ over v.

In Example 6, the following is a dominating solution (that we interpret later):

$$z_{v_1}^* = 1; z_{v_2}^* = 1; z_{v_3}^* = 0, z_{v_4}^* = 0.$$

We can recover the corresponding matching x^* by setting $x^* := \mathcal{V}z^*$:

$$x^*(d_1,h_1) = 1; x^*(d_2,h_1) = 1; x^*(d_1,h_2) = 0; x^*(d_2,h_2) = 0.$$

Notice, in this case, the matching is integral. If x^* is integral, it corresponds to a stable matching. In general,

x^* is fractional. In the next section, we provide an algorithm to convert x^* to an integral solution.

In the remainder of the paper, we refer to x^* as a *fractional stable* solution.

5. Algorithm

In the previous section, we showed how to obtain a fractional matching from a dominating solution. In particular, we let z^* be a dominating solution and set $x^* = \mathscr{V}z^*$. We show here how to round x^* into an integer $\bar{x}$ that satisfies (4) and (5) and almost satisfies (6).

5.1. Rounding

Lemma 2. *Given x^*, there exists integral $\bar{x}$ such that*

- $x^*(d,h) = 0 \Rightarrow \bar{x}(d,h) = 0.$
- $\lfloor \sum_{h\in H} x^*(d,h)\rfloor \le \sum_{h\in H} \bar{x}(d,h) \le 1\ \forall d \in D.$
- $\lfloor \sum_{d\in D} x^*(d,h)\rfloor \le \sum_{d\in D} \bar{x}(d,h) \le k_h\ \forall h \in H.$
- $\left\lfloor \sum_{d\in D_t^h} x^*(d,h)\right\rfloor \le \sum_{d\in D_t^h} \bar{x}(d,h) \le \left\lceil \sum_{d\in D_t^h} x^*(d,h)\right\rceil$

$\forall t = 1,\dots,T_h,\ \forall h \in H.$

Furthermore, $\bar{x}$ can be found by a polynomial time algorithm.

Lemma 2 shows that we can always round a matching x^* to $\bar{x}$ such that capacities at the hospitals are not violated, and the number of doctors for each type is rounded either up or down to the closest integral number.[11] This is essentially the best integer solution that can be hoped for.

We show that the problem of finding $\bar{x}$ can be formulated as the problem of finding a feasible flow in a network, all of whose arc capacities are integral. Integrability of $\bar{x}$ follows immediately.

Introduce a source node σ, a sink node τ, one node for each $d \in D$, $h \in H$, and D_t^h. For each $d \in D$, there is an arc directed from σ to d with an upper bound arc capacity of one. For each d, there is an arc directed to D_t^h if $d \in D_t^h$ and $x^*(d,h) > 0$ with an upper bound arc capacity of one and a lower bound of $\lfloor \sum_{h\in H} x^*(d,h)\rfloor$. For each D_t^h, there is an arc directed to h with an upper bound arc capacity of

$$\left\lceil \sum_{d\in D_t^h} x^*(d,h)\right\rceil$$

and a lower bound arc capacity of

$$\left\lfloor \sum_{d\in D_t^h} x^*(d,h)\right\rfloor.$$

For each $h \in H$, there is an arc directed from h to τ with an upper bound arc capacity of k_h and a lower bound arc capacity of $\lfloor \sum_{d\in D} x^*(d,h)\rfloor$.

Note that x^* is a feasible flow in this network, so we know that a feasible integer flow exists.

5.2. Modifying α

Denote by $\bar{\mu}$ the matching associated with $\bar{x}$. Because of rounding, the proportionality constraints might be violated. We need to change α to make $\bar{\mu}$ feasible. In particular, consider a group D_t^h.

- If, in the fractional solution, $\sum_{d\in D_t^h} x^*(d,h) = \alpha_t^h \sum_{d\in D^h} x^*(d,h)$, then let

$$\bar{\alpha}_t^h = \frac{\sum_{d\in D_t^h} \bar{x}(d,h)}{\sum_{d\in D^h} \bar{x}(d,h)}. \tag{9}$$

- If $\sum_{d\in D_t^h} x^*(d,h) > \alpha_t^h \sum_{d\in D^h} x^*(d,h)$ but $\sum_{d\in D_t^h} \bar{x}(d,h) < \alpha_t^h \sum_{d\in D^h} \bar{x}(d,h)$, then also let $\bar{\alpha}_t^h$ be as before. Otherwise, $\bar{\alpha}_t^h = \alpha_t^h$.

With this, our main result is the following.

Theorem 5. *Given the fractional stable matching x^*, let $\bar{\mu}$ be the matching obtained from x^* via Lemma 2. Then $\bar{\mu}$ is feasible and stable for the instance $(\{\succ_d\}_{d\in D}, \{\succ_h\}_{h\in H}, \{\bar{\alpha}^h\}_{h\in H})$.*

The proof is given in Section 6. In the following, we show proximity bounds for the new matching.

5.3. Proximity Bounds

We can use Lemma 2 to quantify the closeness of $\bar{\mu}$ to x^*. By Lemma 2,

$$|\bar{\mu}(h)| \in \left\{\left\lfloor \sum_{d\in D} x^*(d,h)\right\rfloor, \left\lceil \sum_{d\in D} x^*(d,h)\right\rceil\right\}.$$

Thus, we never violate the capacity constraint of h. Furthermore, the rounding bound also implies that

$$\left||\bar{\mu}(h)| - \sum_{d\in D} x^*(d,h)\right| \le 1\ \forall h \in H.$$

By Lemma 2, $|\bar{\mu}(h)\cap D_t^h| \in \left\{\left\lfloor \sum_{d\in D_t^h} x^*(d,h)\right\rfloor, \left\lceil \sum_{d\in D_t^h} x^*(d,h)\right\rceil\right\}.$

$$\text{Hence}\quad \left||\bar{\mu}(h)\cap D_t^h| - \sum_{d\in D_t^h} x^*(d,h)\right| \le 1\ \forall D_t^h.$$

In fact, when the proportionality constraint associated with D_t^h binds, we can say something more:

$$\left||\bar{\mu}(h)\cap D_t^h| - \alpha_t^h \sum_{d\in D} x^*(d,h)\right| \le 1$$

because

$$|\bar{\mu}(h)\cap D_t^h| \in \left\{\left\lfloor \alpha_t^h \sum_{d\in D} x^*(d,h)\right\rfloor, \left\lceil \alpha_t^h \sum_{d\in D} x^*(d,h)\right\rceil\right\}.$$

Notice that the total number of doctors assigned to any h differs by at most one from the original fractional quantity. The same is true for the number of doctors of a particular type. The proportions, however, can behave quite differently. Using these proximity bounds, it is straightforward to argue that

$$|\alpha_t^h - \bar{\alpha}_t^h| = \left|\alpha_t^h - \frac{|\bar{\mu}(h) \cap D_t^h|}{|\bar{\mu}(h)|}\right| \le \frac{2}{1+\sum_{d\in D} x^*(d,h)} \quad \forall D_t^h.$$

Of course, if h were fully allocated under x^*, this bound would reduce to $\frac{2}{1+k_h}$. If the proportionality constraint for D_t^h binds, this bound improves to

$$\frac{1+\alpha_t^h}{1+\sum_{d\in D} x^*(d,h)}.$$

In all cases, the closeness of the realized proportions, $\frac{|\bar{\mu}(h)\cap D_t^h|}{|\bar{\mu}(h)|}$ to α_t^h, depend upon the size of $|\bar{\mu}(h)|$, the number of doctors matched to h. If $|\bar{\mu}(h)|$ is small, even small changes in the number of doctors assigned to h can have a large effect on the relevant proportions. If large, then a change in one doctor more or less will have a negligible effect on the relevant proportions.

6. Stability of $\bar{\mu}$

Recall that our algorithm starts from a dominating solution, z^*, which is a weight vector of the generators. The algorithm converts z^* to $x^* := \mathcal{V}z^*$, which is a fractional matching between doctors and hospitals. The solution x^* is then rounded to an integral solution $\bar{x}$. We denote $\bar{\mu}$ to be the corresponding matching.

We now show that $\bar{\mu}$ is stable with respect to $\bar{\alpha}$. We prove this by contradiction. Assume that it is not stable. We construct a generator that is not dominated by z^*. We use the notation of $\mathcal{V}^*, \mathcal{V}_h^*$ defined in Remark 5.

A group D_t^h *reaches its lower bound* in x^* if the corresponding proportionality constraint binds, that is,

$$\sum_{d\in D_t^h} x^*(d,h) = \alpha_t^h \sum_{d\in D} x^*(d,h).$$

The following observations are helpful in the proof.

Observation 1. If a group D_t^h reaches its lower bound in x^*, then it also reaches its lower bound with respect to $\bar{x}$ and the modified $\bar{\alpha}$ defined in Equation (9).

Observation 1 comes directly from the definition of $\bar{\alpha}$.

Observation 2. If a group D_t^h reaches its lower bound in x^*, then, for every $v \in \mathcal{V}_h^*$ and $d \in D_t^h$ such that $v(h,d) > 0$, $v(h,d) = \alpha_t^h$.

Observation 2 comes from the way we construct the generators. In particular, for all generators $v \in \mathcal{V}_h^*$, if $v(h,d) > 0$, then $v(h,d) \ge \alpha_t^h$. Thus, if D_t^h reaches its lower bound in x^*, then there cannot be a $v \in \mathcal{V}_h^*$ such that $v(h,d)$ is strictly greater than α_t^h.

Observation 3. If d is a waitlisted doctor at h under $\bar{\mu}$, then any generator in $\mathcal{V}_h$ that assigns d to h cannot be dominated via d.

If d is assigned by $\bar{\mu}$ to a hospital h' such that $h >_d h'$, this means that there is a generator $v \in \mathcal{V}_{h'}^*$. This implies that any generator assigning d to h is ranked above v. If d is unassigned under $\bar{\mu}$, it means that the constraint of d does not bind, that is, $\sum_{h\in H} x^*(d,h) < 1$. Hence, the constraint corresponding to d in $\mathcal{Q}$ also does not bind. Thus, no generator can be dominated at this constraint.

Proof of Theorem 5. Suppose $\bar{\mu}$ is not stable. This means that either h does not reach its effective capacity or there exists d_r currently matched with h that can be exchanged for a higher priority doctor d_a who is on h's waitlist. There are two cases in which d_r can be replaced by d_a. Either they are of the same type or of a different type but d_r is not protected. We construct a generator that is not dominated by z^*, which leads to a contradiction.

Case 1. h is not at its effective capacity under $\bar{\mu}$.

From Remark 1, this means that $\sum_{d\in D} \bar{x}(d,h) < k_h$. However, because of the rounding procedure, this implies that $\sum_{d\in D} x^*(d,h) < k_h$. Thus, no generator can be dominated at h. It remains to create a generator $v \in \mathcal{V}_h$ that is not dominated via any doctor. This leads to a contradiction because z^* is a dominating solution.

Let $\{i_1, \ldots, i_k\}$ be the set of types that contain waitlisted doctors. Choose one waitlisted doctor from each type to be part of the generator v. Let them be $d_{i_1}, \ldots, d_{i_k}$. Also, let $v(d_{i_1},h) = 1 - \sum_{t\neq i_1} \alpha_t^h$ and $v(d_{i_2},h) = \alpha_{i_2}^h, \ldots, v(d_{i_k},h) = \alpha_{i_k}^h$. By Observation 3, the generator that we are constructing cannot be dominated at $d_{i_1},..,d_{i_k}$.

Denote the remaining set of types by $\{i_{k+1}, \ldots, i_{T_h}\}$. These types do not contain any waitlisted doctor because h is not at its effective capacity; according to Remark 1, their side constraints do not bind in the matching $\bar{\mu}$. According to Observation 1, this means that the side constraints of these types do not bind in the fractional solution x^*. Because of Observation 2, this means that, for each type $t \in \{i_{k+1}, \ldots, i_{T_h}\}$, there exists $v_t \in \mathcal{V}_h^*$ that assigns a doctor $d_t \in D_t^h$ to h with probability higher than α_t^h, that is, $v_t(d_t,h) > \alpha_t^h$. Let d_t be part of the generator v, and let $v(d_t,h) = \alpha_t^h$. By the way the preference order is defined for doctors, d_t prefers this new generator to v_t. Thus, v cannot be dominated at any doctor. This contradicts the fact that z^* is a dominating solution.

Case 2a. d_r and d_a are of the same type.

Let $v_r \in \mathcal{V}_h^*$ be a generator that assigns d_r to h. Let v_a be the generator obtained from v_r by assigning d_a to h instead of assigning d_r to h with the same probability. Clearly, because $d_a > d_r$, v_a is ranked above v_r by h and all doctors of different types than d_a. Because of Observation 3, v_a is not dominated at d_a. Thus, z^* does not dominate v_a, a contradiction.

Case 2b. $d_a \in D_a^h$ and $d_r \in D_r^h$ are not of the same type, and d_r is not protected under $\bar{\mu}$.

Because d_r is not protected under $\bar{\mu}$, in the fractional solution, the side constraint of D_r^h does not bind (does not reach its lower bound). Among all the doctors d whose side constraints do not bind and $x^*(d,h) > 0$, let d_{min} be the least preferred doctor according to h. Assume $d_{min} \in D_{min}^h$. Clearly, $d_a \succ_h d_r \geq d_{min}$. If d_a and d_{min} are of the same type, we return to Case 2a. Assume, therefore, that they are of different types.

Let $v_{min} \in \mathcal{V}_h^*$ be a generator that assigns d_{min} to h. Because $x^*(d_{min}, h) > 0$, such a v_{min} exists. There might be several such generators; if so, choose one with the highest probability of assigning d_{min} to h, that is, the highest $v_{min}(d_{min}, h)$. We construct $v_a \in \mathcal{V}_h$ by modifying v_{min} such that v_a is dominated by neither h nor any doctor.

a. v_a assigns d_a to h with probability $1 - \sum_{t \neq a} \alpha_t^h$. By Observation 3, v_a cannot be dominated via d_a.

b. Because the side constraint D_{min}^h does not bind in the fractional solution, there exists $v'_{min} \in \mathcal{V}_h^*$ that assigns a doctor $d'_{min} \in D_{min}^h$ to h with probability higher than α_{min}^h ($d'_m in$ and d_{min} can coincide). Let v_a assign d'_{min} to h with probability α_{min}^h. Thus, v_a does not dominate at d'_{min}.

c. v_a assigns the same doctor in the remaining groups as in v_{min} with probability α_t^h for group D_t^h. With this choice, the sum of the components of v added up over the doctors is one.

The set of doctors assigned by v_a and v_{min} are different only in D_a^h and D_{min}^h. To compare v_a and v_{min}, we only need to compare these doctors. First, notice that $d_a \succ d_{min}$. Now, if $d'_{min} \neq d_{min}$, then $d'_{min} \succ_h d_{min}$ because of the choice of d_{min}. Thus, that h ranks the generators according to the lexicographical order; therefore, v_a is better than v_{min} because it replaces d_{min} with a better doctor. For the case $d'_{min} = d_{min}$, notice that v_a assigns d_{min} with probability α_{min}^h, which is less than the probability of v_{min}. Therefore, h also prefers v_a to v_{min}. Hence, v_a cannot be dominated via h.

Furthermore, the doctors break ties among generators according to h's lexicographical order; v_a cannot be dominated via any doctor in the remaining groups.

Hence, we conclude that $v_a \succ_h v_{min}$ and cannot be dominated by z^*, which is a contradiction.

7. Lower Bounds and Upper Bound

In some cases, the proportion of individuals of a particular type matched to a school or hospital is constrained to fall within some interval. To accommodate this, we extend the earlier analysis to include both lower and upper bound proportionality constraints. Using the notation from prior sections, we consider the following constraints.

$$\alpha_t^h \cdot |\mu(h)| \leq |\mu(h) \cap D_t^h| \leq \beta_t^h \cdot |\mu(h)| \; \forall t = 1, \ldots, T_h, \\ \text{where } 0 \leq \alpha_t^h \leq \beta_t^h \leq 1, \sum_t \alpha_t^h \leq 1 \leq \sum_t \beta_t^h. \tag{10}$$

Call a matching that satisfies (10) feasible. If we choose $\beta_t^h = 1$ for all h and t, we recover (2). We maintain the same notation as before, and departures are noted as they arise.

We develop an algorithm to find a stable matching that slightly violates the proportionality constraints (10). The main result is in Theorem 7. We first define the notion of stability in Section 7.1. Section 7.2 describes the set of cone generators used in the algorithm presented in Section 7.3. The main proof to show that the matching we obtain by this algorithm is stable is given in Section 7.4.

7.1. Stability

The presence of upper and lower bounds on the relevant proportions requires a modification of the definition of a hospital's effective capacity with respect to μ. To see why, fix a hospital h and a subset $S \subset \{1, .., T_h\}$ of the types at that hospital. The upper bounds for all types in $\{1, .., T_h\} \setminus S$ induce a lower bound on the percentage of doctors whose type is in S. Specifically, the number of doctors with types in S needs to be at least a $1 - \sum_{t \notin S} \beta_t^h$ fraction of all the doctors assigned to h. That is,

$$\left| \mu(h) \cap \left(\bigcup_{s \in S} D_s^h \right) \right| \geq \left(1 - \sum_{t \notin S} \beta_t^h \right) \cdot |\mu(h)|. \tag{11}$$

Given a matching μ, if, among the doctors in $\cup_{s\in S} D_s^h$, there are no waitlisted doctors, then, h cannot hope to increase the number of admitted doctors with types in S. Therefore, from (11), the effective capacity of h is at most

$$\frac{|\mu(h) \cap (\cup_{s \in S} D_s^h)|}{1 - \sum_{t \notin S} \beta_t^h}.$$

This motivates the following extension of Definition 2.

Definition 8 (Effective Capacity). Consider a feasible matching μ and a hospital h. Let T_0 be the set of types t, such that D_t^h contains no waitlisted doctor. Let

$$bound_1 := \min_{t \in T_0} \frac{1}{\alpha_t^h} |\mu(h) \cap D_t^h|;$$

$$bound_2 := \min_{S \subset T_0} \frac{1}{1 - \sum_{t \notin S} \beta_t^h} \left| \mu(h) \cap \bigcup_{s \in S} D_s^h \right|.$$

The effective capacity of hospital h with respect to μ, denoted by k_h^μ, is $\min\{k_h, bound_1, bound_2\}$. If $T_0 = \emptyset$, $bound_1$ and $bound_2$ are set to infinity. Similarly, if $\alpha_t^h = 0$ and $1 - \sum_{t \notin S} \beta_t^h = 0$, then $bound_1$ and $bound_2$ are set to infinity, respectively.

As before, when μ is clear from the context, we omit its mention when referring to a hospital's effective capacity.

Remark 6. k_h^μ is an upper bound on the number of slots that h can fill by accepting more waitlisted doctors without violating the side constraints. Because μ is feasible, it satisfies the capacity and the side constraints. Thus, it is clear that $|\mu(h)| \le k_h^\mu$.

From Definition 8, if h is not at its effective capacity, $|\mu(h)| < k_h^\mu$, the following three statements hold.

1. $|\mu(h)| < k_h$.
2. There is no $t \in T_0$ such that the lower bound proportionality constraint corresponding to D_t^h binds, that is, $|\mu(h)| = \frac{1}{\alpha_t^h}|\mu(h) \cap D_t^h|$.
3. There is no $S \subset T_0$ such that all upper bound proportionality constraints for types $t \notin S$ bind. That is, $|\mu(h)| = \frac{1}{\beta_t^h}|\mu(h) \cap D_t^h|$ for all $t \notin S$.

We must also extend Definition 3:

Definition 9 (Protected and Surplus Doctors)**.** Given a feasible matching μ, a doctor of type t is protected at $h \in H$ with respect to μ if the lower bound proportionality constraint associated with D_t^h binds with respect to its effective capacity.

A doctor of type t is surplus at h with respect to μ if the upper bound proportionality constraint associated with type D_t^h binds with respect to the hospital's effective capacity.

As before, when μ is clear from context, we omit the qualifier "with respect to μ."

As in Definition 3, if a type is protected (surplus) at h, it means that h cannot be matched to fewer (more) doctors of this type without reducing the number of positions at h.

Now, consider a hospital h and two doctors $d_a \succ_h d_r$. Assume that d_r is currently matched with h, and d_a is waitlisted at h. This means that h has an incentive to exchange d_r for d_a. The definition of bilateral stability allows for such a blocking coalition but requires that, if h does so, it violates the side constraints. This means that h must either decrease the number of protected doctors or increase the number of surplus doctors. Specifically, we have the following definition.

Definition 10 (Bilateral Stability)**.** A feasible matching μ is called bilaterally stable if the following two conditions hold:

1. Every hospital with a nonempty waitlist is at its effective capacity, that is, $|\mu(h)| = k_h^\mu \ \forall h \in H$.
2. For any $d_a, d_r \in D$ for which d_a is waitlisted for h, $\mu(d') = h$, and $d_a \succ_h d_r$, then d_a, d_r are of different types, and either d_a is surplus or d_r is protected.

The first condition does not permit a hospital to increase its intake. The second says permitting h to replace d_r with d_a will violate (10).

As in Section 2.2, we show that (bilateral) stability implies coalitional stability. A hospital's choice function in the presence of side constraints is defined next.

Definition 11. The choice function of h on a subset of acceptable doctors D^*, denoted $Choice_h(D^*)$, is the subset of D^* with the largest cardinality that satisfies the capacity constraints of h and the proportionality constraints. If there are multiple such subsets, then $Choice_h(D^*)$ is the best one in the lexicographical order according to $\succ_h$.

With this, we have the following theorem.

Theorem 6. *Let μ be a stable matching. Then, for any group of doctors D^* on the waitlist of h, $Choice_h(\mu(h) \cup D^*) = \mu(h)$.*

The proof of Theorem 6 is analogous to that of Theorem 1 and is omitted.

7.2. Cone Generators

We take the same steps as before. The first is to determine the generators of (10). Fix a hospital $h \in H$ and focus on

$$\alpha_t^h \cdot \sum_{d \in D} x(h,d) \le \sum_{d \in D_t^h} x(h,d) \le \beta_t^h \cdot \sum_{d \in D} x(h,d) \quad t = 1, \ldots, T_h. \tag{12}$$

The generators are the extreme points of the system

$$\sum_{t=1}^{T_h} v(d_t, h) = 1, \ \alpha_t^h \le v(d_t, h) \le \beta_t^h \ \forall t = 1, \ldots, T_h. \tag{13}$$

It is easy to see that an extreme point of (13) can be determined using the following algorithm.

1. Select one doctor from each D_t^h and call it d_t.
2. Choose an ordering of the selected doctors and call it σ.
3. Set $v(d_t, h) = \alpha_t^h$ for $i = 1, \ldots, T_h$.
4. In the order selected, increase the value of each $v(d_t, h)$ as much as possible (up to β_t^h) until the remaining mass of $1 - \sum_{t=1}^{T_h} \alpha_t^h$ is exhausted.

With this algorithm, consider a generator for hospital h. If we order the T_h nonzero components in the order that they are selected by the algorithm, they are of the form

$$\beta_{i_1}^h, .., \beta_{i_k}^h, \gamma, \alpha_{i_{k+2}}^h, \ldots, \alpha_{T_h}^h,$$

where $\gamma = 1 - \beta_{i_1}^h - \ldots - \beta_{i_k}^h - \alpha_{i_{k+2}}^h - \ldots - \alpha_{T_h}^h$.

Denote the resulting extreme point by $\{v^\sigma(d, h)\}_{d \in D, h \in H}$. Keep in mind that it is possible for two distinct orderings to give rise to the same extreme point. For subsequent arguments, it is useful to distinguish between the two and, hence, the need to record the order.

Definition 12. A generator $v^\sigma \in \mathcal{V}_h$ *contains* doctor $d \in D$ if $v^\sigma(d, h) > 0$. The *order* of a doctor d contained in generator $v^\sigma \in \mathcal{V}_h$ is the order of d in σ and is denoted $\sigma(d)$. The order of d is undefined if the generator does not contain d.

Example 7. Consider the example in Figure 1. Assume $\beta_1^{h_1} = \beta_2^{h_1} = \beta_3^{h_1} = .45$.

We describe one generator of this system.

1. Select $d_2 \in D_1^{h_1}$, $d_6 \in D_2^{h_1}$, $d_8 \in D_3^{h_1}$.
2. Let σ be the order (d_6, d_2, d_8).
3. Set $v^\sigma(d_2,h_1) = 1/3$, $v^\sigma(d_6,h_1) = 1/3$, $v^\sigma(d_8,h_1) = 1/5$. The remaining coordinates are set to zero.
4. The remaining mass is $1 - 1/3 - 1/3 - 1/5 = 2/15$, which is distributed in the order of σ. This gives $v^\sigma(d_6, h_1) = .45, v^\sigma(d_2,h_1) = .35, v^\sigma(d_8,h_1) = 1/5 = .2$.

We say the generator v^σ contains d_2, d_6, d_8. The order of the doctors in this generator are $\sigma(d_2) = 2$, $\sigma(d_6) = 1$, and $\sigma(d_8) = 3$.

7.3. Algorithm

7.3.1. Ranking of Columns in $\mathcal{AV}$. We consider the conic version of Scarf's lemma as in Section 4.3. The system $\mathcal{A} \cdot \mathcal{V} \cdot z \le b$ is constructed in the same way as in Section 4.3. Each column of $\mathcal{A} \cdot \mathcal{V}$ corresponds to a generator v^σ, and each row of $\mathcal{A} \cdot \mathcal{V}$ corresponds to a constraint for either a doctor or a hospital.

We now describe how each agent in $D \cup H$ *ranks* the columns of $\mathcal{AV}$. (We use the word "rank" to distinguish between the ordering over the columns of $\mathcal{AV}$ and $\mathcal{A}$.)

- h ranks two generators $v^\sigma, \bar{v}^{\sigma'} \in \mathcal{V}_h$ according to the lowest ranked doctor (according to $>_d$) contained in each of them. If the lowest ranked doctor of both v^σ and $\bar{v}^{\sigma'}$ are the same, say d_{min}, then break ties by comparing $\sigma(d_{min})$ and $\sigma'(d_{min})$. The lower the order, the less preferred. If they are equal, move to the second worst doctor contained in each and so on.
- d compares two generators $v^\sigma, \bar{v}^{\sigma'}$ that contain d according to the hospital that each assigns d to using $>_d$. If $v^\sigma, \bar{v}^{\sigma'}$ both assign d to the same hospital h, break ties by comparing $\sigma(d)$ and $\sigma'(d)$. Specifically, if $\sigma(d) > \sigma'(d)$, then v^σ is ranked above $\bar{v}^{\sigma'}$. If $\sigma(d) = \sigma'(d)$, then d uses h's ordering over the generators to break the tie.

7.3.2. Scarf's Algorithm and Rounding. We use the algorithm in Scarf (1967) to derive a dominating $z^* \in \mathcal{Q}$. Set $x^* = \mathcal{V}z^*$, and use Lemma 2 to round x^* into an integer $\bar{x}$ that satisfies (4) and (5) and almost satisfies (6). Let $\bar{\mu}$ be the corresponding matching.

Define $\bar{\alpha}$ and $\bar{\beta}$.

- If $\sum_{d\in D_t^h} x^*(d,h) = \alpha_t^h \sum_{d\in D^h} x^*(d,h)$, then let

$$\bar{\alpha}_t^h = \frac{\sum_{d\in D_t^h} \bar{x}(d,h)}{\sum_{d\in D^h} \bar{x}(d,h)}. \tag{14}$$

- If $\sum_{d\in D_t^h} x^*(d,h) > \alpha_t^h \sum_{d\in D^h} x^*(d,h)$ but $\sum_{d\in D_t^h} \bar{x}(d,h) < \alpha_t^h \sum_{d\in D^h} \bar{x}(d,h)$, then let $\bar{\alpha}_t^h$ be as in (14). Otherwise, $\bar{\alpha}_t^h = \alpha_t^h$.
- Similarly, if $\sum_{d\in D_t^h} x^*(d,h) = \beta_t^h \sum_{d\in D^h} x^*(d,h)$, then let

$$\bar{\beta}_t^h = \frac{\sum_{d\in D_t^h} \bar{x}(d,h)}{\sum_{d\in D^h} \bar{x}(d,h)}. \tag{15}$$

- If $\sum_{d\in D_t^h} x^*(d,h) < \beta_t^h \sum_{d\in D^h} x^*(d,h)$ but $\sum_{d\in D_t^h} \bar{x}(d,h) > \beta_t^h \sum_{d\in D^h} \bar{x}(d,h)$, then also let $\bar{\beta}_t^h$ be as in (15). Otherwise, $\bar{\beta}_t^h = \beta_t^h$.

An argument similar that in Section 5.3 yields the following proximity bounds for $\bar{\alpha}$ and $\bar{\beta}$:

$$|\alpha_t^h - \bar{\alpha}_t^h| = \left|\alpha_t^h - \frac{|\bar{\mu}(h) \cap D_t^h|}{|\bar{\mu}(h)|}\right| \le \frac{2}{1 + \sum_{d\in D} x^*(d,h)} \quad \forall D_t^h,$$

and

$$|\beta_t^h - \bar{\beta}_t^h| = \left|\beta_t^h - \frac{|\bar{\mu}(h) \cap D_t^h|}{|\bar{\mu}(h)|}\right| \le \frac{2}{1 + \sum_{d\in D} x^*(d,h)} \quad \forall D_t^h.$$

Our main result is the following.

Theorem 7. *$\bar{\mu}$ is feasible and stable for the instance $(\{>_d\}_{d\in D}, \{>_h\}_{h\in H}, \{\bar{\alpha}^h\}_{h\in H}, \{\bar{\beta}^h\}_{h\in H})$.*

7.4. Stability of $\bar{\mu}$

Recall that $\mathcal{V}^* = \{v^\sigma \in \mathcal{V} : z^*_{v^\sigma} > 0\}$. For each hospital h, the set of generators in $\mathcal{V}^*$ associated with hospital h is denoted $\mathcal{V}_h^*$.

A group D_t^h *reaches its lower bound* in x^* if

$$\sum_{d\in D_t^h} x^*(d,h) = \alpha_t^h \sum_{d\in D} x^*(d,h).$$

Similarly, D_t^h *reaches its upper bound* in x^* if

$$\sum_{d\in D_t^h} x^*(d,h) = \beta_t^h \sum_{d\in D} x^*(d,h).$$

We use Observation 3 and the following one in the proof.

Observation 4. Fix a hospital h and consider a partition of the set of types into two groups such that group 1 contains $D_{i_1}^h, \ldots, D_{i_k}^h$ and group 2 contains $D_{i_{k+1}}^h, \ldots, D_{i_{T_h}}^h$ if, for all doctors d in group 2 and all $v^\sigma \in \mathcal{V}_h^*$, $\sigma(d) \ge k+1$. Then, either each member of group 1 reaches its upper bound or each member in group 2 reaches its lower bound.

Notice that D_t^h reaches its lower bound in x^* if and only if, for all generator $v^\sigma \in \mathcal{V}_h^*$ that contains $d \in D_t^h$, $v^\sigma(h,d) = \alpha_t^h$. Similarly, D_t^h reaches its upper bound in

x^* if and only if for all generator $v^\sigma \in \mathcal{V}_h^*$ containing $d \in D_t^h$, $v^\sigma(h,d) = \beta_t^h$.

According to the algorithm producing the generators, if we order the T_h nonzero components in the order that they are selected by the algorithm, they are of the form

$$\beta_{i_1}^h, .., \beta_{i_k}^h, \gamma, \alpha_{i_{k+2}}^h, \ldots, \alpha_{T_h}^h,$$

where $\gamma = 1 - \beta_{i_1}^h - \ldots - \beta_{i_k}^h - \alpha_{i_{k+2}}^h - \ldots - \alpha_{T_h}^h$.

Because of the assumption that, for all doctors d in group 2 and all $v^\sigma \in \mathcal{V}_h^*$, $\sigma(d) \geq k+1$, all doctors in group 2 are always selected after all doctors in group 1. Thus, either all the types in group 1 reach their upper bound, and if not, it is because the remaining mass of $1 - \sum_t \alpha_t^h$ is exhausted, and therefore, all types in group 2 reach their lower bound.

Proof of Theorem 7. We are now ready to prove Theorem 7. Suppose, for a contradiction, that $\bar{\mu}$ is not stable. Let $d_a, d_r \in D$ be such that $\bar{\mu}(d_r) = h$, $\bar{\mu}(d_a) \neq h$, and $d_a \succ_h d_r$, that is, d_a is waitlisted at h.[12] There are two cases to consider. In the first, h does not reach its effective capacity. In the second, we can exchange d_r for d_a without violating (10).

Our goal is to construct a generator that is not dominated by z^*, which is a contradiction to the domination of z^*.

Case 1. h does not reach its effective capacity.

From Remark 6, this means that $|\bar{\mu}(h)| < k_h$. Because the rounding does not violate hospital capacity, the capacity constraint at h does not bind in the fractional solution, that is, $\sum_{d \in D} x^*(d,h) < k_h$. Thus, no generator can be dominated at h. It remains to create a generator $w^{\sigma'} \in \mathcal{V}_h$ that is not dominated via any doctor. This leads to a contradiction because z^* is a dominating solution.

The first step is to choose σ'. Order the types so that all the types that contain a waitlisted doctor come first and choose one waitlisted doctor from each of the types to be part of the generator. Let $\{i_1, \ldots, i_k\}$ be the set of these types. By Observation 3, the generator that we are constructing cannot be dominated by the constraints at these doctors.

Denote the remaining set of types by $\{i_{k+1}, \ldots, i_{T_h}\}$. If there is a doctor $d \in D_{i_{k+1}}^h \cup \ldots \cup D_{i_{T_h}}^h$ and a $v^\sigma \in \mathcal{V}_h^*$ such that $\sigma(d) \leq k$, then take this doctor to be the next in the order σ'. Hence, $\sigma'(d) > \sigma(d)$. Therefore, $w^{\sigma'}$ cannot be dominated via d. Repeat until we cannot find such a doctor. Without loss of generality, suppose this happens at the first instance. According to Observation 4, either $D_{i_1}^h, \ldots, D_{i_k}^h$ reaches its upper bound in z^* or $D_{i_{k+1}}^h, \ldots, D_{i_{T_h}}^h$ reaches its lower bound in z^* because there is no waitlisted doctor in $D_{i_{k+1}}^h \cup \ldots \cup D_{i_{T_h}}^h$. Because of the rounding and modifying of α, this means that the corresponding constraints in the rounded matching $\bar{\mu}$ also bind. However, these types define the effective capacity at h, which contradicts the fact that h is not at its effective capacity.

By this argument, we create a generator $w^{\sigma'}$ not dominated via any doctor. This contradicts the fact that z^* is a dominating solution.

Case 2. d_a can be exchanged for d_r.

- We argue that d_a and d_r are not of the same type. Suppose they are of the same type. Let $v^\sigma \in \mathcal{V}_h^*$ be a generator containing d_r. Let $w^\sigma \in \mathcal{V}_h$ be obtained from v^σ by shifting the probability weight from d_r to d_a but keeping the same order. Generator w^σ is ranked above v^σ by h and all doctors of types that differ from d_a and d_r. Moreover, because d_a is a waitlisted doctor, this generator cannot be dominated via d_a. Thus, this new generator is not dominated by z^*, a contradiction.
- Given that $d_a \in D_a^h, d_r \in D_r^h$ are not of the same type, we argue that either D_a^h reaches its upper bound or D_r^h reaches its lower bound in z^*.[13] Suppose, for a contradiction, otherwise. Then, D_a^h has not reached its upper bound, and D_r^h has not reached its lower bound in z^*.

Among all doctors d that have a type that has not reached its lower bound in z^*, and $x^*(d,h) > 0$ (doctor d_r is a member of this set), let d_{min} be the least preferred according to $\succ_h$. Let D_{min}^h be the set of doctors of the same type as d_{min}. Let $v^\sigma \in \mathcal{V}_h^*$ be a generator containing d_{min} such that d_{min}'s order, $\sigma(d_{min})$, is as small as possible. Such a generator exists because $x^*(d,h) > 0$.

Claim 1. d_{min}, d_a are of different types; furthermore, if $d'_a \in D_a^h$ is contained in a generator $v^{\sigma'}$ (d'_a may be the same as d_a), then $\sigma(d_{min}) > \sigma(d'_a)$.

If d_{min} and d_a are of the same type, we could, in v^σ, shift the probability weight from d_{min} to d_a. This produces a generator that is not dominated via h because $d_a \succ_h d_r \succeq_h d_{min}$. It is clearly not dominated via d_a. Finally, it is not dominated via any doctor other than $\{d_a, d_{min}\}$ in the two generators, who face a tie and break it in h's favor.

If σ orders d_{min} before the type of d'_a, that is, $\sigma(d_{min}) < \sigma(d'_a)$, then switch the order of these two types, and if $d'_a \neq d_a$, shift the probability weight from d'_a to d_a. This new generator is not dominated via d_a because of Observation 3. Second, because we have switched the order of d_{min} and d'_a, this new generator is not dominated via d_{min}. Third, because $d_a \succ_h d_r$ for all other doctors and for h, the new generator is ranked above v^σ. Therefore, it cannot be dominated.

Let D^* be the set of doctors who belong to types that have not reached their lower bounds in z^*, and their order in v^σ is at least the order of $\sigma(d_{min})$. See Figure 4. Because of Claim 1, $D_a^h \cap D^* = \emptyset$.

Figure 4. (Color online) The Generator v

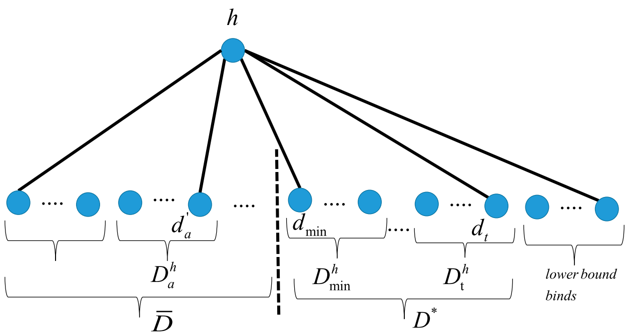

Suppose, among the doctors in D^*, there is a $d_t \in D_t^h$ whose order under a different generator $\bar{v}^{\sigma'} \in \mathcal{V}_h^*$ is no larger than $\sigma(d_{min})$. Notice that d_t could be of the same type as d_{min}, but $d_t \neq d_{min}$ because d_{min} is the least preferred under $>_h$.

Create a new generator from $\bar{v}^{\sigma'}$ by switching the order of D_t^h and D_{min}^h and let d_t be part of this new generator (if d_t, d_{min} are the same type, replace d_{min} with d_t). The new generator is not dominated via d_t because d_t ranks it above $\bar{v}^{\sigma'}$. It is also ranked above v^{σ} by h because d_{min} has a higher order and, thus, is also ranked higher by all other doctors because they break ties according to $>_h$. Hence, this new generator is not dominated.

We are left with the case that the order of each doctor in D^* in every generator in $\mathcal{V}_h^*$ is at least $\sigma(d_{min})$. This means that, for all generators in $\mathcal{V}_h^*$, doctors in D^* are ordered after $\bar{D}$, where $\bar{D}$ is the set of doctors whose types are ordered before D_{min}^h under σ. Similar to the argument in Observation 4, this means that either $\bar{D}$ reaches its upper bound or D^* reaches its lower bound. However, this is impossible because $D_{min}^h \in D^*$ was chosen such that it does not reach its lower bound, and $D_a^h \in \bar{D}$ is assumed to not reach its upper bound.

This concludes the proof.

8. Conclusion

It is common to require that a matching satisfy a variety of distributional goals. These are sometimes expressed as lower or upper bounds on the proportion of agents of a particular type being matched. This paper is the first that we are aware of to address this problem. It uses a novel extension of Scarf's lemma to identify a stable matching that approximately satisfies such proportionality constraints. In addition, ex post bounds on the deviation between the realized and desired proportions are provided.

Acknowledgments

The authors thank the associate editors and the anonymous referees for their insightful comments and suggestions.

Endnotes

[1] See Case Studies of School Choice and Open Enrollment in Four Cities, http://www.thecoweninstitute.com.php56-17.dfw3-1.websitetestlink.com/uploads/Case-Studies-on-Choice-and-Enrollment-11-2011-1490715209.pdf, last accessed August 6, 2019.

[2] Other school districts facing steep enrollment declines are Buffalo, Philadelphia, Columbus (Ohio), Pittsburgh, Cleveland, Detroit, and Kansas City.

[3] Besides these, some (e.g., Ágoston et al. 2018 and Gonczarowski et al. 2018) have used heuristic and integer programming approaches. However, no theoretical guarantees on performance are available.

[4] If one is not careful, there is also a "circularity" problem, in that stability is defined with respect to the modified choice function.

[5] One could formulate the problem of finding a stable matching as an integer program and start with a fractional solution to it. However, such fractional solutions are not guaranteed to be stable in the sense of Scarf.

[6] We assume the categories are disjoint. Our model can be extended to capture overlapping categories. The proportionality constraints associated with overlapping categories are cone constraints. The approximation guarantees depend on the structure of these categories.

[7] Our results extend to the case with both upper *and* lower bounds on the proportions of each type to be matched as well. This is described in Section 7. That section also describes an explicit algorithm for determining the matching.

[8] If $\alpha_t^h = 0$ for all h and t, this choice function reduces to being responsive: for any set $D^* \subset D$, hospital h's choice from D^* consists of the (up to) k_h highest priority doctors among the feasible doctors in D^*.

[9] It is possible that another stable matching may have fewer agents worse off and a larger number of assigned doctors. This is not uncommon when the substitutes assumption fails as it does here. This is because the rural hospital theorem does not hold.

[10] The number of types, T_h, is typically a small constant. Hence, the number of generators is polynomial in the number of doctors.

[11] This is similar to theorem 3 in Budish et al. (2013).

[12] d_a, d_r denote for the doctor to accept and the doctor to reject, respectively.

[13] Because of the rounding procedure and the modifying of α's, this implies that either d_a is at surplus or d_r is protected at the rounded matching $\bar{\mu}$ with the modified $\bar{\alpha}$'s, $\bar{\beta}$'s. This is what we need to prove.

References

Abdulkadiroglŭ A, Sönmez T (2003) School choice: A mechanism design approach. *Amer. Econom. Rev.* 93(3):729–747.

Ágoston KC, Biró P, Szántó R (2018) Stable project allocation under distributional constraints. *Oper. Res. Perspect.* 5:59–68.

AygunO, TurhanB (2016) Dynamic reserves in matching markets: Theory and applications. Working paper, Bogazici University, Istanbul.

Biró P, Fleiner T, Irving RW, Manlove DF (2010) The college admissions problem with lower and common quotas. *Theoret. Comput. Sci.* 411(34):3136–3153.

Budish E, Che Y-K, Kojima F, Milgrom P (2013) Designing random allocation mechanisms: Theory and applications. *Amer. Econom. Rev.* 103(2):585–623.

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

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

Ellison G, Pathak PA (2016) The efficiency of race-neutral alternatives to race-based affirmative action: Evidence from Chicago's exam schools. NBER Working Paper 22589, National Bureau of Economic Research, Cambridge, MA.

Erdil A, Kumano T (2012) Prioritizing diversity in school choice. Working paper, Washington University, St. Louis.

Fleiner T, Kamiyama N (2016) A matroid approach to stable matchings with lower quotas. *Math. Oper. Res.* 41(2):734–744.

Fragiadakis D, Troyan P (2016) Improving matching under hard distributional constraints. *Theoret. Econom.* 12(2):863–908.

Gonczarowski YA, Kovalio L, Nisan N, Romm A (2019) Matching for the Israeli mechinot gap year: Handling rich diversity requirements. *Proc. 2019 ACM Conf. Econom. Comput.* (ACM, New York), 321–321.

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

Huang C-C (2010) Classified stable matching. Charikar M, ed. *Proc. 21st Annual ACM-SIAM Sympos. Discrete Algorithms* (SIAM, Philadelphia), 1235–1253.

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

Kojima F, Tamura A, Yokoo M (2018) Designing matching mechanisms under constraints: An approach from discrete convex analysis. *J. Econom. Theory* 176:803–833.

Nemhauser GL, Wolsey LA (1988) *Integer and Combinatorial Optimization* (Wiley-Interscience, New York).

Nguyen T, Vohra R (2018) Near-feasible stable matchings with couples. *Amer. Econom. Rev.* 108(11):3154–3169.

Scarf HE (1967) The core of an *n*– person game. *Econometrica* 35(1): 50–69.

Tang D, Nguyen T, Subramanian V, Vohra R (2018) Computational testing of Scarf algorithm for near feasible stable matching with couples. Working paper, University of Michigan, Ann Arbor.

Thành Nguyen is an associate professor at the Krannert School of Management of Purdue University. His research lies at the intersection of optimization and game theory. He is especially interested in analyzing contracts and negotiations in supply chain networks and developing algorithms for market design problems.

Rakesh Vohra is the George A. Weiss and Lydia Bravo Weiss University Professor, Department of Economics and the Department of Electrical and Systems Engineering, University of Pennsylvania. His primary research interests are auction theory and game theory.