

DEPARTMENT OF ECONOMICS
WORKING PAPER SERIES

2020-17



McMASTER UNIVERSITY

Department of Economics
Kenneth Taylor Hall 426
1280 Main Street West
Hamilton, Ontario, Canada
L8S 4M4

<http://www.mcmaster.ca/economics/>

Matching with Compatibility Constraints

The Case of the Canadian Medical Residency Match

Muhammad Maaz^{*1,2,3} and Anastasios Papanastasiou^{†2}

¹Faculty of Health Sciences, McMaster University

²Department of Economics, McMaster University

³Faculty of Medicine, University of Toronto

Abstract

The Canadian medical residency match has received considerable attention in the Canadian medical community as several students go unmatched every year. Simultaneously, several residency positions go unfilled, largely in Quebec, the Francophone province of Canada. The Canadian match is unique in that positions are designated with a language restriction, a phenomenon that has not been studied or described priorly in the matching literature. To study this phenomenon, we develop the model of matching with compatibility constraints, where based on a dual characteristic, a subset of students is incompatible with a subset of hospitals. We show that while the deferred acceptance algorithm still yields a stable matching, some desirable properties from standard two-sided matching are lost. For instance, we show that if the number of residencies exceeds the number of students, some students can yet go unmatched. We derive a lower bound for the number of English and Francophone residency positions such that every student is matched for all instances of (a form of) preferences. Our analysis suggests that to guarantee a stable match for every student, a number of positions at least equal to the population of bilingual students must be left unfilled. The model can be generalized to other instances of the stable marriage problem.

JEL codes: C78, D82

*Email: m.maaz@mail.utoronto.ca Phone: +1-647-712-4251

†Email: papanasa@mcmaster.ca

1 Introduction

The seminal paper by Gale and Shapley “College admission and the stability of marriage” introduced the deferred acceptance (DA) algorithm as a mechanism for establishing stable matchings in two-sided matching problems [1]. Since their paper, applications of DA have flourished, the most notable being the medical residency match. This application was motivated by Roth’s observation that the National Resident Matching Program (NRMP) in the United States, which is responsible for allocating medical school graduates to their post-graduate training (also called a residency), had independently arrived at the Gale-Shapley DA algorithm [2] [3]. In 1999, the DA algorithm was modified to include the ability for student couples to apply to match together. This modified algorithm is called the Roth-Peranson algorithm [4], and was adopted in many other countries, including Canada [5]. Since then, matching theory has remained a ripe field, both theoretically and practically, with the question of real-world *constraints* inspiring much of the matching work in the 21st century.

In Canada, medical students apply to be matched to postgraduate training (also called a residency) at a Canadian hospital through the Canadian Residency Matching Service (CaRMS) [5], which uses a version of the DA algorithm ¹. The unique constraint that exists is that some positions are designated for French-speaking students in order to provide French services to the public. This is due to French’s status as the second official language of Canada [6]. While this guarantees equal status for French and English in federal jurisprudence, some provinces also give French special status. The province of New Brunswick, for example, is officially bilingual, while the province of Quebec, Canada’s largest province, is officially unilingually French [6]. As well, French is often taught as a second language in English-speaking provinces like Ontario [6], while English is also taught in Francophone provinces.

¹The CaRMS actually runs four different matches [5]: 1. R-1: This is what graduating or graduated medical students apply to for their postgraduate training. 2. MSM: Medicine Subspecialty Match. This is for residents currently in an internal medicine program seeking to enter subspecialty training. 3. FM/EM: Family Medicine/Emergency Medicine. This is for residents who are currently in or have completed family medicine training and wish to pursue further training in emergency medicine. 4. PSM: Pediatric Subspecialty Match. This is for residents currently in a pediatric residency program who wish to pursue subspecialty training. In this paper, when we talk about the residency match, we are referring to the R-1 match.

According to CaRMS data, in the 2019 R-1 match, 103 out of 2984 Canadian medical graduates (CMG's) went unmatched - meaning that 96.5% did indeed obtain a residency position. While comparing favorably to other residency matching clearinghouses - for example, in the US, 79.6% of applicants to the NRMP are matched [7] - much attention in Canada has been drawn to the issue of unmatched medical residents. The Canadian Medical Association has increasingly been sounding the alarm over the issue of unmatched medical students [8], with the number of unmatched CMGs has been steadily increasing every year. Other professional organizations, like the Association of Faculties of Medicine of Canada (AFMC), have been lobbying the government to respond as well (provincial governments are responsible for funding residency positions) [9]. It is worth noting that unmatched medical students cannot practice medicine, despite nearly a decade in school, and are often left with little in terms of job prospects [9].

In the Canadian medical literature, much discussion has been ongoing as to what to do about the CaRMS. Wilson and Bordman, in a commentary in the Canadian Medical Association Journal, the preeminent general medical journal in Canada, declared that the CaRMS was "broken", citing the fact that 68 graduates went unmatched, while 64 residency positions were unfilled (including 56 in family medicine in the province of Quebec) [10]. This commentary attracted much discussion and replies in the subsequent months, including doctors, deans of medical schools, the CaRMS itself, and impassioned personal anecdotes from unmatched graduates [11] [12] [13] [14] [15] [16] [17]. News media have picked up on the problem of unmatched residents in recent years as well, with considerable coverage surrounding the tragic suicide of Dr. Robert Chu who went unmatched despite attempting to do so twice [18]. The frustration over the CaRMS has even spilled into the real world, with professional associations staging demonstrations outside the Ontario provincial legislature [9].

Wilson and Bordman's commentary, as well as match data analysis by the AFMC, demonstrated there was a seeming disconnect between the two sides of the matching market. There are more positions than graduates [9], which at first glance is a favorable situation. Again, comparing with the United States, there are indeed fewer positions than students in the NRMP, so the sub-100% match rate is perhaps easily explained away by that disparity [7]. However, in Canada, there are approximately 102 positions for every 100 medical graduates. In addition, it seems that unfilled residency positions tend to largely be in Quebec [10], and Quebec graduates match to other

provinces more than other province’s students match to Quebec [9]. All in all, the plight of the unmatched is one of the most important issues facing the Canadian medical community today.

1.1 Related literature

Observations of “undesirable” (from a policymaker’s perspective) matches yielded by current matching algorithms has led to work on possible modifications to the basic DA algorithm. This is not a new problem. As far back as 1970, McVitie and Wilson studied the stable marriage problem with unequal sets of men and women [19]. Clearly, by the Pigeonhole Principle [20], some elements will remain unmatched. McVitie and Wilson proved the Rural Hospital Theorem, which states that unmatched participants in one stable matching are unmatched in all stable matchings [19]. This result was later restated by Roth as, in the resident-hospital matching market, “any hospital that fails to fill all of its positions in some stable outcome will not only fill the same number of positions at any other stable outcome, but will fill them with exactly the same residents.” [21]. The theorem was termed the Rural Hospital Theorem on the basis that rural hospitals tend to have greater difficulty filling their residency positions as they are seen as less desirable than urban ones. From these early results, we can see that the idea of imbalances and disparities arising in matching markets is not new.

The aforementioned urban-rural disparity was observed in the data in countries that used centralized clearinghouses for their medical residents, and some countries became proactive in attempting to manipulate the matching algorithm in order to correct the imbalance. Kamada and Kojima [22] [23] studied the Japanese medical residency match, which uses the student-proposing DA. In response to public pressure about the lack of rural doctors, the Japanese government instituted regional quotas based on prefectures (government districts) [22], the idea being to set caps on how many residents may work in urban prefectures. Kamada and Kojima demonstrated that such tampering with the DA algorithm results in inefficiency and possible instability, as well as a lower match rate (fewer doctors overall receive positions) [22]. They propose a *flexible deferred acceptance* algorithm that results in stability and respects regional quotas [23], and show, through simulations, that while this still yields a lower match rate than normal DA, it does fill more positions than the Japanese implementation of regional quota DA [23].

The opposite problem of setting floor constraints instead of ceiling constraints is seemingly less tractable. Kamada and Kojima point out that floor constraints are likely much harder to use [24] [25]. For example, if no resident wants to be matched to a specific region, then individual rationality would be compromised, and even with an individually rational matching, stability is not guaranteed [24]. Recent work in the computer science literature has found that checking the mere existence of a feasible matching with floor constraints is \mathcal{NP} -complete [26]. It remains unclear whether such constraints are tractable, and what the definitions of concepts like individual rationality and stability would be in such situations [26].

Our paper’s contribution is thus twofold. From an economic theory point of view, we study a novel situation that has not been described in other well-studied matching markets in the literature. While there is a growing literature on introducing constraints into matching problems, these papers focus on other constraints, such as quotas. The situation described above in Canada, where due to language designations, a subset of students is incompatible with a subset of residency positions, has not been treated by other papers, to the authors’ knowledge. Secondly, with regards to the real world, given the intense scrutiny around the Canadian residency match, this paper aims to build a theoretical basis that can explain how and why the much-derided outcomes described above have arisen. On this basis, possible solutions to the problems affecting the CaRMS can be developed. This paper therefore serves as an extension of the theory of matching as well as an analysis of the CaRMS match.

2 Model

2.1 Preliminaries

As per Roth and Sotomayor [27], our hospital-residents model is a four-tuple $\langle H, I, q, P \rangle$:

- H is a finite set of hospitals.²

²Note this is purely semantics. Medical professionals may protest that in Canada it is actually universities that “host” residency positions, and have affiliations with hospitals which is where the resident would actually practice. This is true, however we are using “hospitals” as this is the standard terminology used in the matching literature.

- I is a finite set of students. The sets H and I are disjoint.
- q is a vector of hospital capacities: q_h for $h \in H$ gives the capacity of hospital $h \in H$.
- P is a collection of preference relations, such that:
 - For each $i \in I$, P_i denotes the preferences of student i over a subset of H , whence we derive the strict preference relation \succ_i ; so, $h_1 \succ_i h_2$ means that student i strictly prefers hospital h_1 to h_2 . A student i prefers any hospital $h_1 \in P_i$ to any hospital $h_2 \notin P_i$.
 - For each $h \in H$, P_h denotes the preferences of hospital h over a subset of I , whence, as with the students, we derive the strict preference relation \succ_h , which is defined similarly. Similarly, $i_1 \succ_h i_2 \forall h \in H, \forall i_1 \in P_h, \forall i_2 \notin P_h$.

Student i is said to be *acceptable* to hospital h if $i \in P_h$, and hospital h is acceptable to student i if $i \in P_h$. Note that since hospitals have a capacity of more than one, in reality they would have preferences between sets of students, not necessarily individual students. However, we will assume that hospitals have *responsive preferences*, meaning that replacing a less-preferred student with a more-preferred one, or filling a vacancy with an acceptable student (i.e. a student listed on its preferences) makes it better off [27].

A **matching** is a function $\mu : H \cup I \rightarrow \mathcal{P}(H \cup I)$ such that [27]:

1. No hospital exceeds its quota, with some positions possibly left unfilled: $\mu(h) \subseteq I \cup \{\emptyset\}$ such that $|\mu(h)| \leq q_h$ for all $h \in H$,
2. each student is matched to at most one hospital or not at all: $\mu(i) \subseteq H \cup \{\emptyset\}$ such that $|\mu(i)| \leq 1$ for all $i \in I$,
3. student i is matched to hospital h if and only if hospital h is matched to a set containing student i : $i \in \mu(h) \iff \mu(i) = \{h\}$ for all $h \in H$ and $i \in I$.

We call a pair $(h, i) \in H \times I$ a **blocking pair** if i and h are both acceptable to each other, and *both* of the following two conditions hold [27]:

1. $h \succ_i \mu(i)$, and,
2. either $i \succ_h i'$ for some $i' \in \mu(h)$, or, $|\mu(h)| < q_h$ and $i \succ_h \emptyset$

From the concept of a blocking pair we can define one of the central concepts in matching theory: stability. A matching μ is **stable** if and only if there do not exist any blocking pairs under μ [1].

2.2 Deferred acceptance algorithm

The current CaRMS configuration uses the Roth-Peranson algorithm, which is the student-proposing deferred acceptance algorithm [4]. As well, this is the algorithm that we will be analyzing in the context of matching residents to residencies throughout this paper. The **student-proposing deferred acceptance (DA) algorithm** is defined as follows [27]:

Step 1. Each student i proposes to its most preferred hospital. A hospital h receiving more than q_h proposals shortlists its q_h most preferred students according to its preferences P_h , and rejects the rest, while a hospital h receiving less than q_h proposals shortlists all of its proposals.

Step k . Any student i who was rejected at step $k - 1$ proposes to the hospital it prefers the most among the hospitals it applied to (i.e. hospitals in P_i) that hasn't rejected it yet. At each step, each hospital h takes the q_h top students from its shortlist and its proposers, and rejects the others.

The algorithm terminates when there are no more rejections. At termination, the matching is given by the shortlists of the hospitals in the most recent step.

The algorithm also gives a stable matching if the hospitals propose [27], although this can be a different matching than the one given by the student-proposing version. Note that it is possible for there to be stable matchings other than the one yielded by the DA algorithm [1].

2.3 Introducing compatibility constraints

We build upon the basic model in section 2.1. Our motivation for this model comes from the CaRMS language constraints. Namely, every student can be designated as either Anglophone, Francophone, or both (ie. bilingual). On the other hand, the set of hospitals can be partitioned into two disjoint sets on the basis of language as well.³ We say that student i **applies to**, or is an

³There is of course the situation that one hospital can have some English positions and some French positions. However, we can simply imagine this hospital as two different hospitals, one containing all the English positions, and one containing all the French positions. Therefore, the set of hospitals can always be partitioned into two disjoint sets:

applicant of, hospital h , if $h \in P_i^4$. Student i and hospital h are **compatible** only if they share the same language characteristic, and are incompatible otherwise. Thus, in our formulation, an English-speaking student can apply only to English hospitals, and French-speaking students can apply only to French hospitals, and bilingual students can apply to both English and French hospitals⁵. In addition, hospitals rank only the students that apply to them⁶.

We can generalize the idea of such language incompatibilities to any sort of incompatibility based on some arbitrary two-valued characteristic. In general, we define a **matching with compatibility constraints problem** as a standard hospital-residents model as per section 2.1 with the following additional constraints:

- There is a two-valued characteristic $C = \{c_1, c_2\}$.
- Each student $i \in I$ has the characteristic c_1 , c_2 , or both. Let the set of students with characteristic c_1 be denoted as I_1 , and the set of students c_2 be denoted I_2 , such that $I = I_1 \cup I_2$. Let the intersection of these sets $I_1 \cap I_2$ be denoted $I_{1,2}$. We restrict the sets $I_1 - I_{1,2}$ and $I_2 - I_{1,2}$ (the sets of students who only have c_1 and who only have c_2 , respectively) to be non-empty.
- There is a partition of hospitals H into two disjoint sets H_1 and H_2 , which correspond to the characteristics c_1 and c_2 .

English and French.

⁴This language of *applying* is from the real-world set-up of the residency match, where, when medical students seek residencies, they go through an application process entailing sending a CV, reference letters, and participating in an interview. At the end of the process, student submit a ranking to the CaRMS (or whichever centralized matching system) of the hospitals they applied to, and similarly hospitals rank the students that submitted applications to them according to the strength of their applications.

⁵This is the same as having French-only students prefer no match over a match with Anglophone hospitals, and vice versa for English-only students

⁶Note that in reality, it is the hospitals who impose such restrictions - for example, a hospital restricts its positions to French speakers. It does not necessarily follow that English-speaking students will not apply to French hospitals. However, Irving has shown that one can assume without loss of generality that preference are *consistent* in two-sided matching problems, meaning that for some hospital h and student i , $h \in P_i$ if and only if $i \in P_h$ [28]. Therefore, it follows that though these language restrictions are exogenously imposed by the hospitals, we can safely say that the students also do not apply to hospitals which would find them unacceptable due to language constraints.




Figure 1: Schematic of matching with compatibility constraints applied to the Anglophone(E)/Francophone(F) constraints in the CaRMS

- A student-hospital pair (h, i) is compatible if they share the same characteristic, and incompatible otherwise. A student is unacceptable to a hospital and a hospital is unacceptable to a student if they do not share the same characteristic. Hence, a student can only apply to compatible hospitals, however they may not apply to all. See figure 1 for a representation.

We now apply this terminology in the context of our example. Our characteristic set is $C = \{E, F\}$, where E is the English-speaking characteristic, and F denotes the French-speaking characteristic. English-only students $I_E - I_{E,F}$ are incompatible with the French hospitals H_F , while the French-only students $I_F - I_{E,F}$ are incompatible with the English hospitals H_E . This is shown in figure 1.

3 Results

3.1 Stability

Stability is an important consideration in matching markets. As Roth has shown, instability often leads to a collapse of matching markets [27]. In order to demonstrate stability, we can show that the matching with compatibility constraints is an instance of the stable marriage with incomplete preferences problem (SMI problem). First introduced by Gale and Sotomayor, an SMI problem is a one-to-one matching problem where preferences are not complete [29]. The following lemma will help us to establish stability.

Lemma 3.1. *The hospital-residents problem with compatibility constraints is an instance of the SMI problem.*

Proof. Let S be a finite set of residency positions. For every hospital $h \in H$ with quota q_h , construct q_h copies of h , each copy with the same preference relation as h , and each copy with capacity of 1. Place these copies in S . Rewrite the preference relations of every student $i \in I$ by replacing every hospital $h \in P_i$ with a list of the elements of S that were derived from h , arbitrarily breaking ties to maintain strict preferences. Now, the many-to-one sided matching problem between I and H has been translated into a one-to-one matching problem between I and S ; i.e. it is a stable marriage problem. Due to compatibility constraints, preferences are incomplete. Therefore, it is a stable marriage problem with incomplete preferences. \square

This result allows us to immediately establish stability, as follows.

Corollary 3.1. *With compatibility constraints, DA yields a stable matching.*

Proof. Gale and Sotomayor showed that the DA algorithm yields a stable matching for the SMI problem [29]. Combining this result with lemma 3.1 completes the proof. \square

Therefore, we have shown that even when compatibility constraints are introduced as per section 2.3, the DA algorithm still finds a stable matching.

3.2 Existence of unmatched students in stable matchings

As touched upon in the introduction of the paper, a key issue in the CaRMS is that some students go unmatched, despite more residency positions than students. As well, many positions also go unfilled, largely in Quebec. With our matching with compatibility constraints framework, we can demonstrate that such a result is theoretically possible with the following motivating example.

Consider a case where there are one English-only, one bilingual, and one French-only student. At first glance it seems that one should only need three positions, since there are only three students, say 2 Anglophone and 1 Francophone positions. But, the problem with this is that if the bilingual student places the Francophone position as first in his preferences, and

likewise the Francophone position does so to the bilingual student, they will be matched after running student-proposing DA. This leaves the French-only student without a position. On the other hand, if there are 1 Anglophone and 2 Francophone positions, then the bilingual student could out-compete the English-only student analogous to the above case, leaving the English-only student without a position.

This contrasts with the well-known result that when there are as many students as residency positions, and preferences are complete, then there are no unmatched students and no unfilled positions after running DA [27].

We can look further at the case where there are *more* residency positions than students. For example, in the CaRMS, there are about 102 positions for every 100 students [9]. Observe that in the example where the bilingual student ranks the Anglophone position first, and vice versa the Anglophone position ranks it first, then adding further Francophone positions does nothing to help the overall match rate, as the English-only student is still left without a position - and indeed leaves those Francophone positions unfilled. This mirrors the current situation in the CaRMS where English-only students seem to bear the brunt of the unmatched issue, while Francophone positions go unfilled.

However, now observe what would happen if there were 2 Anglophone and 2 Francophone positions. Then, even if the bilingual student gets matched to an Anglophone position, there is still one left over for the English-only student. Similarly, he cannot compete the French-only student out of a position because there is still one position left over for the French-only student. This example provides the motivation for the following section.

3.3 Establishing an I -saturating stable matching

An I -saturating matching is defined as a matching in which, for all $i \in I$, $\mu(i) \neq \emptyset$ [30]. So, an I -saturating stable matching is such a matching that is also *stable*.

As the motivating example above showed, it is insufficient to set the number of positions equal to the number of students. Consideration must be given to the number of Anglophone and Francophone positions individually. As well, the role of preferences is important. For example, with 2 Francophone and 1 Anglophone positions, if the bilingual student is matched to the Francophone position then no student will go unmatched. However, the issue is that a social planner choosing how many hospital positions to

have (which mimics the situation in Canada well, as funding for residency positions comes from the government) does not know a priori how the students will rank the hospitals nor how hospitals will rank students⁷. If only the very limited information of how many there are in each class is known, how many residency positions should be allocated, such that every student obtains a position *no matter* what ends up transpiring during the residency application process? In the vein of the motivating example, we will establish a necessary and sufficient condition such that no student is unmatched in all possibilities of (a form of) preferences.

First, we introduce a new definition for preference completeness. If every student has complete preferences over their respective compatible hospitals, and vice versa every hospital has complete preferences over their respective compatible students, then we say that preferences are **compatibility-wise complete**. This is *as complete* as preferences can be under compatibility constraints.

Next, we introduce some additional notation to make the statement easier to read. Let the set of English-only students be E , the set of French-only students be F , and the set of bilingual students be B , with sizes e , f , and b , respectively. These are all subsets of I , and we label their elements as: $E = \{i_1^E, i_2^E \dots i_e^E\}$, $F = \{i_1^F, i_2^F \dots i_f^F\}$, and $B = \{i_1^B, i_2^B \dots i_b^B\}$. Let the set of Anglophone hospitals be X and the set of Francophone hospitals be Y , with total quotas x and y , respectively. We restrict $e, f, b, x, y > 0$. Let the set of all possible compatibility-wise complete preferences be \mathbb{P} . Then, we can show the following result.

Theorem 3.1. *Every stable matching is I -saturating in all instances of compatibility-wise complete preferences if and only if $x \geq e + b$ and $y \geq f + b$. Formally: $(\forall P \in \mathbb{P})(\text{every student has a position in all stable matchings}) \Leftrightarrow (x \geq e + b) \wedge (y \geq f + b)$.*

Proof. We first prove the only if part of the statement.

Assume contradiction student i does not have a position in some stable matching. Let the number of students with the same characteristic as i , including i , be k . As preferences are compatibility-wise complete and the

⁷There are a host of factors that contribute to how hospitals rank applicants, including marks, reference letters, academic publications, and community service [31]. Similarly, there are a host of factors that contribute to how students rank hospitals, including prestige, reputation in a particular medical field (for example, students interested in trauma would like to go to premier trauma centres), family, and cost of living [32].

matching is stable, student i does not form a blocking pair with any of the k compatible hospitals. By assumption, the number of positions with i 's characteristic is greater than k , and because students cannot occupy more than one position, at most $k - 1$ positions with i 's characteristic are filled, and at least one position k_i compatible with student i is left unfilled. By compatibility-wise completeness, student i and the hospital with the position k_i form a blocking pair, and the matching is not stable.

Next, we prove the if part of the statement. Consider its contrapositive: $(x < e + b) \vee (y < f + b) \Rightarrow (\exists P \in \mathbb{P})(\text{there exists some student without a position in every stable matching})$. It suffices to show the existence of such a P , so we will use a constructive proof.

First, consider the case where $x < e + b$. Consider an instance P in which:

- Every Anglophone hospital: $i_m^B \succ i_n^E$ for all $m \leq b$ and for all $n \leq e$.
- Also, every Anglophone hospital: $i_m^A \succ i_n^A$ for all $m < n$.
- Every bilingual student: $h_x \succ h_y$ for all $h_x \in X$ and for all $h_y \in Y$.

We show that student i_e^E is left unmatched in every stable matching. Suppose not, so there is a stable matching μ with $x < e + b$ in which i_e^E is assigned to an Anglophone hospital h . From the definition of stability, h does not form a blocking pair with any bilingual or English-only student. Under the above preferences, this is only possible if all English and bilingual students are matched to some other English hospital that they prefer to h . This implies that μ assigns $e + b$ students to $x < e + b$ English positions, an impossibility. Thus, student i_e^E is left unmatched in every stable matching.

Along the same lines, we can show that when $y < f + b$ some students go unmatched in every stable matching for some P . For example, consider a set of compatibility-wise complete preferences in which:

- Every Francophone hospital: $i_m^B \succ i_n^F$ for all $m \leq b$ and for all $n \leq f$.
- Also, every Francophone hospital: $i_m^F \succ i_n^F$ for all $m < n$.
- Every bilingual student: $h_y \succ h_x$ for all $h_x \in X$ and for all $h_y \in Y$.

Following the same steps as in the first case, we obtain that student i_f^F is left unmatched in every stable matching. This completes the proof for the if part of the statement. □

In plainer words, theorem 3.1 means that, in order to guarantee every student a match in all (compatibility-wise complete) preference possibilities, the number of Anglophone positions needs to be at least equal to the number of English students (including bilingual students) and the number of Francophone positions needs to be at least equal to the number of French students (including bilingual students). For example, if we have 5 students (2 English-only, 2 French-only, and 1 bilingual), then in order to ensure that every student is matched (assuming compatibility-wise completeness), no matter what the preferences are, we would actually need 6 positions (3 Anglophone and 3 Francophone) instead of, as we might think at first glance, 5 positions for 5 students.

Note that we showed in theorem 3.1 that every student is matched in all stable matchings. Since the student-proposing DA algorithm specifically gives the student-optimal stable matching [27], theorem 3.1's condition is also a necessary and sufficient condition in the special case of the CaRMS (a fact which might be more useful for real-world applicability).

In general, from the condition $x \geq e + b$ and $y \geq f + b$ from theorem 3.1, the total number of positions $x + y = e + f + 2b$ is more than the number of students, $e + f + b$.⁸Of course, if the number of positions is greater than the number of students, this also means some positions will go unfilled. These two facts demonstrate, in some sense, the inefficiency introduced by compatibility constraints in matching markets. Our condition implies that there is inherently a trade-off for the policymaker deciding how many residency positions to fund: setting the number of residency positions in accordance with the lower bound of theorem 3.1 would mean that every student is matched, but would also mean some positions will be unfilled, which could be a waste of resources. The policymaker must therefore consider these two opposing goals: matching every student, or filling every residency position.

⁸Except for the degenerate case when $I_E \cap I_F = \emptyset$, meaning $b = 0$, we effectively have two separate standard hospital-residents problems: one between the English students and hospitals, and one between the French students and hospitals. Then, it suffices to have the number of Anglophone positions equal to the number of English students, and similarly for the number of Francophone positions equal to the number of French students.

4 Discussion and Conclusion

In this paper we developed the matching with compatibility constraints model, where a dual characteristic causes a subset of students to be incompatible with a subset of hospitals, in order to investigate the phenomenon of language restrictions in the Canadian medical residency match. This is, to the authors knowledge, the first paper to investigate this unique feature of the Canadian residency match, and use it to explain its present problems. Notably, we investigated theoretically how this could lead to the current issue in the CaRMS of unmatched students and unfilled positions. We showed that even when there are more residencies than students, as is the case in Canada, it is not guaranteed that every student is able to obtain a position.

We defined a weaker form of preference completeness, called compatibility-wise completeness, which is as complete as preferences can be under compatibility constraints. We then showed that when we assume compatibility-wise completeness (i.e. all English-speaking students apply to all English residencies), then we can guarantee every student obtaining a position by having the number of English positions equal to the number of English-speaking students and the number of French positions equal to the number of French-speaking students. Interestingly, the total required number of positions to guarantee this is greater than the number of students - which contrasts with the result in standard matching models that under complete preference relations, having positions equal in number to the students guarantees a match for everyone. Unfortunately, even given this guarantee, we cannot assuage the problem of unfilled residency positions.

The real-world applicability of this prescription may be limited as preferences in the real world are likely not compatibility-wise complete. There are significant logistical hurdles that applicants to residency positions must pass through for each application, including reference letters and interviews. Due to this, medical students in the CaRMS do not rank all hospitals with whom they are compatible. Taking into this account, the number of required residency positions to guarantee that every student matches is likely larger, albeit by an unknown amount, than what would be required under compatibility-wise complete preferences.

Ultimately, our model has implications for the CaRMS and analyzing its current issues that have received so much attention in the medical community. It's generalized formulation in terms of arbitrary two-valued characteristics allows it to be applied to any variant of one-to-one and many-to-one matching

situations. For example, in a marriage market, it could be used to analyze the effect of the existence of religious preferences. Future theoretical work could take this framework in numerous directions. As well, it would be interesting to see how the framework applied empirically to, for instance, the study of the CaRMS. It would be interesting to see how varying the number of Anglophone and Francophone positions affects the match rate by simulating the CaRMS. We leave it to future theoreticians and empiricists to build upon the results laid out in this paper.

References

- [1] David Gale and Lloyd S Shapley. College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1):9–15, 1962.
- [2] Alvin E Roth. The evolution of the labor market for medical interns and residents: a case study in game theory. *Journal of Political Economy*, 92(6):991–1016, 1984.
- [3] Alvin E Roth. The origins, history, and design of the resident match. *Jama*, 289(7):909–912, 2003.
- [4] Alvin E Roth and Elliott Peranson. The redesign of the matching market for american physicians: Some engineering aspects of economic design. *American economic review*, 89(4):748–780, 1999.
- [5] Canadian Residency Matching Service. Home. <https://www.carms.ca/>, Dec 2019.
- [6] Milton J Esman. The politics of official bilingualism in canada. *Political Science Quarterly*, 97(2):233–253, 1982.
- [7] National Resident Matching Program. 2019 main residency match by the numbers. National Resident Matching Program <http://www.nrmp.org/new-2019-main-residency-match-by-the-numbers>, Sep 2019.
- [8] Canadian Medical Association. Carms update: 40% fewer students go unmatched over last year. <https://www.cma.ca/carms-update-40-fewer-students-go-unmatched-over-last-year>, Apr 2019.
- [9] AFMC. Reducing the number of unmatched canadian medical graduates: A way forward. Technical report, Association of Faculties of Medicine of Canada, Jan 2018.
- [10] C Ruth Wilson and Zachary N Bordman. What to do about the canadian resident matching service. *Canadian Medical Association Journal*, 189(47):E1436–E1447, 2017.
- [11] Megan Sorokopud-Jones. The increase in unmatched canadian medical graduates: who is to blame and should we be concerned? *University of Manitoba Journal of Medicine*, 1(1), 2018.

- [12] Eugene YH Yeung. Canadian resident matching service: are we trying to reinvent the wheel? *Canadian Medical Association Journal*, 2019.
- [13] Janice Willett. What to do about the canadian residency matching process. *Canadian Medical Association Journal*, 2019.
- [14] Sarah L Silverberg and Kaylynn MH Purdy. Unmatched canadian medical graduates. *Canadian Medical Association Journal*, 190(4):E118–E118, 2018.
- [15] Amit Persad. The overall culture of residency selection needs fixing. *Canadian Medical Association Journal*, 190(14):E443–E443, 2018.
- [16] Amit Persad. The unmatched. *Canadian medical education journal*, 9(2):e89, 2018.
- [17] Geneviève Moineau. Association of faculties of medicine of canada response to: The unmatched by dr. amit persad. *Canadian medical education journal*, 9(2):e95–e96, 2018.
- [18] Franklin Warsh. Robert chu’s suicide sends a message medicine cannot ignore. *Toronto Star*, Jun 2017.
- [19] David G McVitie and Leslie B Wilson. Stable marriage assignment for unequal sets. *BIT Numerical Mathematics*, 10(3):295–309, 1970.
- [20] Tamara J Lakins. *The tools of mathematical reasoning*. American Mathematical Soc., 2016.
- [21] Alvin E Roth. On the allocation of residents to rural hospitals: a general property of two-sided matching markets. *Econometrica: Journal of the Econometric Society*, pages 425–427, 1986.
- [22] Yuichiro Kamada and Fuhito Kojima. Improving efficiency in matching markets with regional caps: The case of the japan residency matching program. *Discussion Papers, Stanford Institute for Economic Policy Research*, 1, 2010.
- [23] Yuichiro Kamada and Fuhito Kojima. Stability and strategy-proofness for matching with constraints: A problem in the japanese medical match and its solution. *American Economic Review*, 102(3):366–70, 2012.

- [24] Yuichiro Kamada and Fuhito Kojima. Efficient matching under distributional constraints: Theory and applications. *American Economic Review*, 105(1):67–99, 2015.
- [25] Yuichiro Kamada and Fuhito Kojima. Stability and strategy-proofness for matching with constraints: A necessary and sufficient condition. *Theoretical Economics*, 13(2):761–793, 2018.
- [26] Masahiro Goto, Atsushi Iwasaki, Yujiro Kawasaki, Ryoji Kurata, Yosuke Yasuda, and Makoto Yokoo. Strategyproof matching with regional minimum and maximum quotas. *Artificial intelligence*, 235:40–57, 2016.
- [27] Alvin E Roth and Marilda Sotomayor. Two-sided matching. *Handbook of game theory with economic applications*, 1:485–541, 1992.
- [28] Robert W Irving, David F Manlove, and Sandy Scott. The stable marriage problem with master preference lists. *Discrete Applied Mathematics*, 156(15):2959–2977, 2008.
- [29] David Gale and Marilda Sotomayor. Some remarks on the stable matching problem. *Discrete Applied Mathematics*, 11(3):223–232, 1985.
- [30] Alan Gibbons. *Algorithmic graph theory*. Cambridge university press, 1985.
- [31] Joshua Lakoff, Nicholas Cofie, Nancy Dalgarno, Kelly Howse, and Sylvia Heeneman. Analysis of factors affecting canadian medical students’ success in the residency match. *Canadian Medical Education Journal*, 2019.
- [32] Todd Dow, Connor McGuire, Emma Crawley, and Dafydd Davies. Application rates to surgical residency programs in canada. *Canadian Medical Education Journal*, 2019.