



MATCHING WITH COMPATIBILITY CONSTRAINTS: THE CASE OF THE CANADIAN MEDICAL RESIDENCY MATCH

Muhammad Maaz

McMaster University, Canada
University of Toronto, Canada
m.maaz@mail.utoronto.ca

Anastasios Papanastasiou

McMaster University, Canada
papanasa@mcmaster.ca

ABSTRACT

The Canadian medical residency match has received considerable attention in the medical community as several students go unmatched every year. Simultaneously, multiple residency positions go unfilled, largely in Quebec, the Francophone province of Canada. In Canada, positions are designated with a language restriction, a phenomenon that has not been described previously in the matching literature. We develop a model of matching with compatibility constraints, where, based on a dual-valued characteristic, a subset of students is incompatible with a subset of hospitals, and show how such constraints lead to inefficiency. We derive a lower bound for the number of Anglophone 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.

Keywords: Two-sided matching, CaRMS, matching with constraints.

JEL Classification Numbers: C78, D82.

We thank two anonymous referees for their helpful comments that led to substantial improvements of the paper. The authors are responsible for any remaining errors.

Copyright © Muhammad Maaz, Anastasios Papanastasiou / 4(1), 2020, 99–117.

1. INTRODUCTION

The seminal paper by Gale & Shapley (1962) introduced the deferred acceptance (DA) algorithm as a mechanism for establishing stable matchings in two-sided matching problems. Since then, applications of DA have flourished such as the well-known 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 (Roth, 1984, 2003). 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 (Roth & Peranson, 1999), and was adopted in many other countries, including Canada (Canadian Residency Matching Service, 2019). Since then, matching theory has remained a fertile 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) (Canadian Residency Matching Service, 2019), 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 the status of French as the second official language of Canada (Esman, 1982). 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 (Esman, 1982). As well, French is often taught as a second language in English-speaking provinces like Ontario, while English is also taught in Francophone provinces (Esman, 1982).

¹ The CaRMS actually runs four different matches (2019): 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 (2019), 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 were matched in 2019 (National Resident Matching Program, 2019) - much attention in Canada has been drawn to the issue of unmatched medical residents. The Canadian Medical Association (2019) has increasingly been sounding the alarm over the issue of unmatched medical students, as the number of unmatched CMGs steadily increasing every year. Other professional organizations, like the Association of Faculties of Medicine of Canada (2018), abbreviated as the AFMC, have been lobbying the government to respond as well (provincial governments are responsible for funding residency positions). 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 (Association of Faculties of Medicine of Canada, 2018).

In the Canadian medical literature, much discussion has been ongoing as to what to do about the CaRMS. Wilson & Bordman (2017), 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). This commentary attracted attention and responses in the subsequent months, including those from doctors, deans of medical schools, the CaRMS itself, and impassioned personal anecdotes from unmatched graduates (Sorokopud-Jones, 2018; Yeung, 2018; Willett, 2017; Silverberg & Purdy, 2018; Persad, 2018a,b; Moineau, 2018). 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 (Warsh, 2017). The frustration over the CaRMS has even spilled into some professional associations which staged demonstrations outside the Ontario provincial legislature (Association of Faculties of Medicine of Canada, 2018).

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 (Association of Faculties of Medicine of Canada, 2018), 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 ([National Resident Matching Program, 2019](#)). However, in Canada, there are approximately 102 positions for every 100 medical graduates ([Association of Faculties of Medicine of Canada, 2018](#)). In addition, it seems that unfilled residency positions tend to largely be in Quebec ([Wilson & Bordman, 2017](#)), and Quebec graduates match to other provinces more than other provinces' students match to Quebec ([Association of Faculties of Medicine of Canada, 2018](#)). All in all, the plight of the unmatched is one of the most important issues facing the Canadian medical community today.

We now review some related literature. Observations of “undesirable” (from a policymaker’s perspective) matches yielded by current matching algorithms led to work on possible modifications to the basic DA algorithm. This is not a new problem. As far back as fifty years ago, [McVitie & Wilson \(1970\)](#) studied the stable marriage problem with unequal sets of men and women. Clearly, by the Pigeonhole Principle ([Lakins, 2016](#)), some elements will remain unmatched. [McVitie & Wilson \(1970\)](#) proved the Rural Hospital Theorem, which states that unmatched participants in one stable matching are unmatched in all stable matchings. This result was later restated by [Roth \(1986\)](#) 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.” 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 & Kojima \(2010, 2012\)](#) 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) ([Kamada & Kojima, 2010](#)), the idea being to set caps on how many residents may work in urban prefectures. [Kamada & Kojima \(2010\)](#) 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). They proposed a *flexible deferred acceptance* algorithm that results in stability and respects regional quotas, 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 (Kamada & Kojima, 2012).

The opposite problem of setting floor constraints instead of ceiling constraints is seemingly less tractable. Kamada & Kojima (2010, 2012) point out that floor constraints are likely much harder to use. For example, if no student 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 (Kamada & Kojima, 2015). 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 (Goto et al., 2016). It remains unclear whether such constraints are tractable, and what the definitions of concepts like individual rationality and stability would be in such situations (Goto et al., 2016).

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 real world applications, 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 & Sotomayor (1992), our hospital-residents model is a four-tuple $\langle H, I, q, P \rangle$:

- H is a finite set of hospitals.²
- 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 .
- P is a collection of preference relations, such that:
 - For each $i \in I$, P_i denotes the preferences of student i over $H \cup \{\emptyset\}$, hence 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 .
 - For each $h \in H$, P_h denotes the preferences of hospital h over $I \cup \{\emptyset\}$, hence, as with the students, we derive the strict preference relation \succ_h , which is defined similarly.³

Student i is said to be *acceptable* to hospital h if $i \succ_h \emptyset$, and hospital h is acceptable to student i if $h \succ_i \emptyset$. Alternatively, we will use throughout this paper the terminology that student i **applies to**, or is an applicant of, hospital h , if $h \succ_i \emptyset$.⁴

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 makes it better off (Roth & Sotomayor, 1992).

A **matching** is a function $\mu : H \cup I \rightarrow \mathcal{P}(H \cup I)$ such that (Roth & Sotomayor, 1992):

² Note 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.

³ If for some student i , $\emptyset \succ_i h$, then i prefers being unmatched to being matched to h . If for some hospital h , $\emptyset \succ_h i$, then h prefers keeping some position unfilled rather than being matched to i .

⁴ This 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, students 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.

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 (Roth & Sotomayor, 1992):

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 (Gale & Shapley, 1962).

2.2. Deferred acceptance algorithm

The current CaRMS configuration uses the Roth-Peranson algorithm, which is the student-proposing deferred acceptance algorithm (Roth & Peranson, 1999). As well, this is the algorithm that we analyze in the context of matching residents to residencies throughout this paper. The **student-proposing deferred acceptance (DA) algorithm** is defined as follows (Roth & Sotomayor, 1992):

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 it finds acceptable) 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 (Roth & Sotomayor, 1992), 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 (Gale & Shapley, 1962).

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.⁵ 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 Anglophone hospitals, and French-speaking students can apply only to Francophone hospitals, and bilingual students can apply to both Anglophone and Francophone hospitals.⁶ In addition, hospitals rank only the students that apply to them.⁷

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:

- ⁵ There is, of course, the situation where one hospital can have some Anglophone positions and some Francophone positions. However, we can simply imagine this hospital as two different hospitals, one containing all the Anglophone positions, and one containing all the Francophone positions. Therefore, the set of hospitals can always be partitioned into two disjoint sets: English and French.
- ⁶ This is the same as having French-speaking-only students prefer no match over a match with Anglophone hospitals, and vice versa for English-speaking-only students.
- ⁷ Note 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 Francophone hospitals. However, Irving et al. (2008) have shown that one can assume without loss of generality that preferences are *consistent* in two-sided matching problems, meaning that for hospital h and student i , $h \succ_i \emptyset$ if and only if $i \succ_h \emptyset$. 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

- 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 .
- 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-speaking-only students $I_E - I_{E,F}$ are incompatible with the Francophone hospitals H_F , while the French-speaking-only students $I_F - I_{E,F}$ are incompatible with the Anglophone hospitals H_E . This is shown in Figure 1.

3. RESULTS

3.1. Stability

Stability is an important consideration in matching markets. As Roth & Sotomayor (1992) have shown, instability often leads to a collapse of matching markets. 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 & Sotomayor (1985), an SMI problem is a one-to-one matching problem where preferences are not complete, meaning that some hospital-student pairs are not mutually acceptable. The following lemma will help us to establish stability.

Lemma 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 (i.e., some hospital-student pairs are not mutually acceptable, per the definition of Irving et al. (2008)). Therefore, it is a stable marriage problem with incomplete preferences. \square

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

Corollary 1. *With compatibility constraints, DA yields a stable matching.*

Proof. Gale & Sotomayor (1985) showed that the DA algorithm yields a stable matching for the SMI problem. Combining this result with Lemma 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-speaking-only student, one bilingual student, and one French-speaking-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-speaking-only student analogous to the above case, leaving the English-speaking-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 (Roth & Sotomayor, 1992).

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 (Association of Faculties of Medicine of Canada, 2018). 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-speaking-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-speaking-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-speaking-only student. Similarly, they cannot compete the French-speaking-only student out of a position because there is still one position left over for the French-speaking

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 $i \in I$, $\mu(i) \neq \{\emptyset\}$ (Gibbons, 1985). Thus an I -saturating stable matching is a matching that is also *stable*.

As the illustrative 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 position, 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 students.⁸ 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 illustrative 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 finds all their respective compatible hospitals acceptable, and vice versa, i.e., every hospital finds all of their respective compatible students acceptable, then we say that preferences are **compatibility-wise complete**, or CW-complete. To use the language of Irving et al. (2008) and others who study the SMI problem, *complete preferences* would be satisfied if every hospital-student pair is acceptable to both hospital and student; due to compatibility constraints, this is generally not possible. However, CW-complete preferences are, in that sense, *as complete* as preferences can be under these constraints.

Next, we introduce some additional notation to make the statement easier

⁸ There are a host of factors that contribute to how hospitals rank applicants, including marks, reference letters, academic publications, and community service (Lakoff et al., 2020). 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 medicine would like to go to premier trauma centers), family, and cost of living (Dow et al., 2020).

to read. Let the set of English-speaking-only students be E , the set of French-speaking-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 CW-complete preferences be \mathbb{P} . Then, we can show the following result.

Theorem 1. *Every stable matching is I -saturating in all instances of CW-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 ‘if’ part of the statement: every stable matching is I -saturating in all instances of CW-complete preferences if $x \geq e + b$ and $y \geq f + b$.

Suppose to the contrary that student i does not have a position in some stable matching. Let the number of students with the same characteristic as i (note this includes bilingual students), including i , be k . Note that if i is bilingual, then k is the number of students that share at least one characteristic with i , and so k is the number of English-speaking-only, French-speaking-only, and bilingual students - i.e. *all* students.

As preferences are CW-complete and the matching is stable, student i does not form a blocking pair with any of its compatible hospitals. By assumption, the number of positions with i ’s characteristic is weakly 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 compatible with student i is left unfilled. By CW-completeness, student i and the hospital with that unfilled position form a blocking pair. Thus, the matching is not stable, which is a contradiction. Therefore, every student is matched in all stable matchings for all instances of CW-complete preferences.

Next, we prove the ‘only if’ part of the statement. Consider its contrapositive: $(x < e + b) \vee (y < f + b) \Rightarrow (\exists P \in \mathbb{P})(\text{there exists a stable matching at which some student is unmatched})$. It suffices to show the existence of such a matching μ for some 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 h : $i_m^B >_h i_n^E$ for all $m \leq b$ and for all $n \leq e$.

- Also, every Anglophone hospital h : $i_m^E \succ_h i_n^E$ for all $m < n$.
- Every bilingual student i : $h_x \succ_i h_y$ for all $h_x \in X$ and for all $h_y \in Y$.

We need to show that some student is left unmatched in some stable matching. 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-speaking-only student. Under the above preferences, this is only possible if all English-speaking and bilingual students are matched to some other Anglophone hospital that they prefer to h . This implies that μ assigns $e + b$ students to $x < e + b$ Anglophone 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 some stable matching for some P . For example, consider a set of CW-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, which implies the required result. This completes the proof for the ‘only if’ part of the statement. □

Theorem 1 shows that $(x \geq e + b) \wedge (y \geq f + b)$ is needed for every stable matching to be I -saturating. The next result establishes that this condition is also needed for the existence of an I -saturating matching in every instance of preferences.⁹

Lemma 2. *If there exists a stable matching in which every student is matched $\forall P \in \mathbb{P}$, then $(x \geq e + b) \wedge (y \geq f + b)$.*

⁹ We thank an anonymous referee for constructive suggestions on structuring the proof of Lemma 3.2.

Proof. We need to show $(\forall P \in \mathbb{P})$ (there is some stable matching in which every student is matched) $\implies (x \geq e + b) \wedge (y \geq f + b)$, equivalently, $(\nexists P \in \mathbb{P})$ (there is some student unmatched in every stable matching) $\implies (x \geq e + b) \wedge (y \geq f + b)$. It suffices to prove the contrapositive $(x < e + b) \vee (y < f + b) \implies (\exists P \in \mathbb{P})$ (there is some student unmatched in every stable matching). The result follows directly from the second part of our proof of Theorem 1, by observing that the argument is established for every stable matching. \square

In plainer words, Lemma 2 means that, in order to guarantee every student a match in all CW-complete preference possibilities, the number of Anglophone positions needs to be at least equal to the number of English-speaking students (including bilingual students) and the number of Francophone positions needs to be at least equal to the number of French-speaking students (including bilingual students). For example, if we have 5 students (2 English-speaking-only, 2 French-speaking-only, and 1 bilingual), then in order to ensure that every student is matched (assuming CW-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.

Such a requirement on positions would also guarantee, by Theorem 1, that every student is matched in every stable matching, when preferences are CW-complete. Since the student-proposing DA algorithm specifically gives the student-optimal stable matching (Roth & Sotomayor, 1992), this result also applies to the special case of the CaRMS, a fact which might be useful in policymaking.

From Theorem 1 and Lemma 2 we obtain the following result:

Corollary 2. *If there exists a stable matching under which every student is matched for all CW-complete preferences, at least as many positions as the number of bilingual students are left unfilled in every stable matching.*

Proof. If the assumption in the statement holds, by Lemma 2 $x \geq e + b$ and $y \geq f + b$ and by Theorem 1 exactly $f + e + b$ students are matched, and so $x + y - e - f - b \geq b$ positions are left unfilled, in every stable matching. \square

In general, from conditions $x \geq e + b$ and $y \geq f + b$ in Theorem 1 and Lemma 2, the total number of positions $x + y \geq e + f + 2b$ is greater than

the number of students, $e + f + b$ and some positions must remain unfilled.¹⁰ Corollary 2 shows that at least one position for every bilingual student is left unfilled, and demonstrates the inefficiency of introducing compatibility constraints in matching markets, in the sense that if language restrictions are lifted, all positions are filled and no students are left unmatched in every stable matching with complete preferences.¹¹ Our condition implies that there is an inherent 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 1 would mean that every student is matched, but would also mean some positions would be unfilled, which could be a waste of resources. The policymaker must therefore consider these two opposing goals: guaranteeing a match for every student, or filling every residency position.

4. CONCLUDING REMARKS

In this paper we developed the matching with compatibility constraints model, where a dual-valued 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 our best knowledge, the first paper to investigate this unique feature of the Canadian residency match and use it to explain its present problems under the lens of standard two-sided matching theory. Notably, we investigated theoretically how this could lead to the current issue of unmatched students and unfilled positions observed in the CaRMS. 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, or CW-completeness, which is as complete as preferences can be under compatibility constraints. We then showed that when we assume CW-completeness (i.e. all English-speaking students apply to all Anglophone

¹⁰ Except 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.

¹¹ Assuming that every student prefers to be matched to their least preferred hospital over going unmatched, and every hospital prefers to hire its least preferred student than leaving a position empty.

residencies), then we can guarantee every student obtaining a position by having the number of Anglophone positions equal to the number of English-speaking students and the number of Francophone 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 stating 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 solve the problem of unfilled residency positions. Rather surprisingly, the number of bilingual students leads to an increase in inefficiency, in the sense of unfilled positions in Corollary 2.

The real-world applicability of this prescription may be limited as preferences in the real world are not likely to be CW-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 to be larger, albeit by an unknown amount, than the requirement under CW-complete preferences.

Our model certainly has implications for the CaRMS and the analysis of the current issues that have received so much attention in the medical community. Its 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. In addition, 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

Association of Faculties of Medicine of Canada. (2018, Jan). *Reducing the number of unmatched Canadian medical graduates: A way forward* (Tech. Rep.). Association of Faculties of Medicine of Canada.

- Canadian Medical Association. (2019, Apr). *CaRMS update: 40% fewer students go unmatched over last year*. <https://www.cma.ca/carms-update-40-fewer-students-go-unmatched-over-last-year>.
- Canadian Residency Matching Service. (2019, Dec). *Home*. <https://www.carms.ca/>.
- Dow, T., McGuire, C., Crawley, E., & Davies, D. (2020). Application rates to surgical residency programs in Canada. *Canadian Medical Education Journal*, 11(3), e92–e100.
- Esman, M. J. (1982). The politics of official bilingualism in Canada. *Political Science Quarterly*, 97(2), 233–253.
- Gale, D., & Shapley, L. S. (1962). College admissions and the stability of marriage. *The American Mathematical Monthly*, 69(1), 9–15.
- Gale, D., & Sotomayor, M. (1985). Some remarks on the stable matching problem. *Discrete Applied Mathematics*, 11(3), 223–232.
- Gibbons, A. (1985). *Algorithmic graph theory*. Cambridge university press.
- Goto, M., Iwasaki, A., Kawasaki, Y., Kurata, R., Yasuda, Y., & Yokoo, M. (2016). Strategyproof matching with regional minimum and maximum quotas. *Artificial intelligence*, 235, 40–57.
- Irving, R. W., Manlove, D. F., & Scott, S. (2008). The stable marriage problem with master preference lists. *Discrete Applied Mathematics*, 156(15), 2959–2977.
- Kamada, Y., & Kojima, F. (2010). 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.
- Kamada, Y., & Kojima, F. (2012). 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.
- Kamada, Y., & Kojima, F. (2015). Efficient matching under distributional constraints: Theory and applications. *American Economic Review*, 105(1), 67–99.
- Lakins, T. J. (2016). *The tools of mathematical reasoning*. American Mathematical Soc.
- Lakoff, J., Howse, K., Cofie, N., Heeneman, S., & Dalgarno, N. (2020). Analysis of factors affecting Canadian medical students' success in the residency match. *Canadian Medical Education Journal*, 11(3), e43–e55.
- McVitie, D. G., & Wilson, L. B. (1970). Stable marriage assignment for unequal sets. *BIT Numerical Mathematics*, 10(3), 295–309.
- Moineau, G. (2018). Association of faculties of medicine of Canada response to: The unmatched by Dr. Amit Persad. *Canadian Medical Education Journal*, 9(2), e95–e96.
- National Resident Matching Program. (2019, Sep). *2019 main residency match by*

- the numbers*. National Resident Matching Program <http://www.nrmp.org/new-2019-main-residency-match-by-the-numbers>.
- Persad, A. (2018a). The overall culture of residency selection needs fixing. *Canadian Medical Association Journal*, 190(14), e443–e443.
- Persad, A. (2018b). The unmatched. *Canadian Medical Education Journal*, 9(2), e89–e92.
- Roth, A. E. (1984). 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.
- Roth, A. E. (1986). On the allocation of residents to rural hospitals: a general property of two-sided matching markets. *Econometrica: Journal of the Econometric Society*, 425–427.
- Roth, A. E. (2003). The origins, history, and design of the resident match. *Jama*, 289(7), 909–912.
- Roth, A. E., & Peranson, E. (1999). The redesign of the matching market for american physicians: Some engineering aspects of economic design. *American economic review*, 89(4), 748–780.
- Roth, A. E., & Sotomayor, M. (1992). Two-sided matching. *Handbook of game theory with economic applications*, 1, 485–541.
- Silverberg, S. L., & Purdy, K. M. (2018). Unmatched Canadian medical graduates. *Canadian Medical Association Journal*, 190(4), e118–e118.
- Sorokopud-Jones, M. (2018). The increase in unmatched Canadian medical graduates: Who is to blame and should we be concerned? *University of Manitoba Journal of Medicine*, 1(1).
- Warsh, F. (2017, Jun). Robert Chu's suicide sends a message medicine cannot ignore. *Toronto Star*. Retrieved from <https://www.thestar.com/opinion/commentary/2017/06/20/robert-chus-suicide-sends-a-message-medicine-cannot-ignore.html>
- Willett, J. (2017, Dec). What to do about the Canadian residency matching process [letter]. *Canadian Medical Association Journal*. Retrieved from <https://www.cmaj.ca/content/what-do-about-canadian-residency-matching-process>
- Wilson, C. R., & Bordman, Z. N. (2017). What to do about the Canadian resident matching service. *Canadian Medical Association Journal*, 189(47), e1436–1447.
- Yeung, E. Y. (2018, Sep). Canadian Resident Matching Service: are we trying to reinvent the wheel? [letter]. *Canadian Medical Association Journal*. Retrieved from <https://www.cmaj.ca/content/canadian-resident-matching-service-are-we-trying-reinvent-wheel>