

Matching with Compatibility Constraints

The Case of the Canadian Residency Matching
Service

Muhammad Maaz

A thesis submitted for the degree of
Bachelor of Health Sciences (Honours)

Faculty of Health Sciences
McMaster University
Hamilton, Ontario, Canada
April 2020

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 binary 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 also investigate a dynamic game where unmatched positions are re-advertised without language restriction. The model can be generalized to other instances of the stable marriage problem.

Acknowledgement

I would like to thank Anastasios Papanastasiou, Assistant Professor at the Department of Economics at McMaster, for his supervision and mentorship during this thesis. This work was inspired by him sending me Gale and Shapley's 1962 paper as a fun reading exercise, which ultimately led me to explore matching theory more deeply.

Dedication

To my teachers.

Contents

1	Introduction	4
2	Literature Review	6
2.1	Beginnings	6
2.2	Recent developments	8
3	Model	10
3.1	Preliminaries	10
3.2	Deferred acceptance algorithm	11
3.3	Introducing compatibility constraints	11
4	Results	13
4.1	Stability	13
4.2	Existence of unmatched students or unfilled positions	14
4.3	Establishing an I -saturating stable matching	16
4.4	A dynamic game extension	18
5	Discussion and Conclusion	20

1 Introduction

The well-known paper by Gale and Shapley “College admission and the stability of marriage” introduced the deferred acceptance (DA) algorithm as a way of finding 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 and Japan. 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. For example, Kamada and Kojima studied the effects of government quotas on urban and rural doctors in Japan by investigating the properties of the DA algorithm under these constraints [5] [6].

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) [7], which uses a version of the DA algorithm¹. The unique constraint that exists is the that some spots are designated for French-speaking students in order to French services to the public. This is due to French’s status as the second official language of Canada [8]. 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 [8]. As well, French is often taught as a second language in English-speaking provinces like Ontario [8], while English is also taught in Francophone provinces.

¹The CaRMS actually runs four different matches [7]: 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 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 [9] - 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 plight of unmatched medical students [10], 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 as well (provincial governments are responsible for funding residency positions) [11]. 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 [11].

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) [12]. 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 [13] [14] [15] [16] [17] [18] [19]. News media have picked up on the plight 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 [20]. The frustration over the CaRMS has even spilled into the real world, with professional lobbying groups staging demonstrations outside the Ontario provincial legislature [11].

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 [11], 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 [9]. 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 [12], and Quebec graduates match to other provinces more than other province's students match to Quebec [11]. All in all, the plight of the unmatched is one of the most important issues

facing the Canadian medical community today.

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 other recent papers have focused on introducing real-world constraints into matching theory (see section 2.2), 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 author’s 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 Literature Review

2.1 Beginnings

Gale and Shapley’s seminal paper “College admissions and the stability of marriage” arguably founded the field of matching theory. In this paper, they introduced the stable marriage problem. The setup of the problem is that a population of people are separated into two disjoint partitions (men and women). How can we best “match” elements of these partitions to one another taking into account their individual preferences (or, in the marriage context, how can we best marry the men and women)? Gale and Shapley introduce the notion of stability as being a desirable property of a matching. A matching is *stable* if there does not exist a man and woman who are not matched to each other but who would prefer to be matched to each other [1]. It is important to realize that both parties must want to be matched to each other over their current partners. For example, if some man prefers some other woman over his wife, but that woman is content with her husband, then this is not an unstable match. In the marriage context, it is clear to see why instability would be undesirable.

The stable marriage problem was extended to a more general form called the college admissions problem by Gale and Shapley in the same paper [1]. In the college admissions problem, the two sets are students and colleges, where colleges have quotas of at least one student [1]. So, the one-to-one nature of the stable marriage problem is extended into a many-to-one problem,

as multiple students can be matched to one college. Gale and Shapley argue that the college admissions problem is essentially equivalent to the stable marriage problem, as it can be transformed into a one-to-one problem. Consider if a college has some quota q , then we can construct q “dummy” objects, each representing one seat at the college. Clearly, only one student can match to one of these seats, so it is a one-to-one matching problem now. However, there are limits to this intuition [21].

The question in both of these problems is: is it possible to find a stable match? Gale and Shapley propose the deferred acceptance (DA) algorithm and prove that DA always yields a stable match [1]. A corollary of this fact is that a stable match always exists. The first step in DA is to have each individual draw up a preference ranking of the other set (so, men rank the women and women rank the men). Rankings should be strict, such that there is no ambiguity or indifference [1]. It is not necessary for preferences to be complete. Then, have each man “propose” to the woman on the top of his list. If a woman has multiple proposals, she chooses the one she prefers the most in accordance with her preference list and rejects the others [1]. Rejected men propose to their next highest ranked woman, and the same procedure follows. The key caveat is that acceptances are tentative: if a woman receives a proposal in a later round which she prefers to someone who she has already said yes to, then she will renege on her previous suitor and take the newer one instead, hence the name *deferred* acceptance [1]. This continues until everyone is matched or there are no more proposals possible. In the college admissions extension, the algorithm is largely similar. However, when the students propose to the colleges, the college chooses the q highest ranked students from among the proposals, and reject the rest [1]. This reflects the many-to-one nature of the problem.

We described the algorithm with the men proposing or with the students proposing. However, it also works if the women propose or the colleges propose. This also yields a stable match [1]. Interestingly, it can be shown that the party that does the proposing yields the best stable match from their perspective, and the worst stable match from the other party’s perspective [22]. In other words, in the college admissions problem, the student-proposing algorithm yields the student-optimal match (the stable match that is the best for the students), while the college-proposing yields the college-optimal match. The student-optimal match is also the college-pessimal match (the worst stable match from the perspective of the colleges), and the college-optimal match is student-pessimal [22].

Beginning in the 1980’s, applications of matching theory and specifically DA flourished. One of the most notable examples is the market for medical

residents. 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 and had been using it [2] [3]. Roth relays that the formation of a centralized clearinghouse to handle such allocations was eventually formed to settle the chaos that began to emerge in the market. For example, exploding offers, which are offers that expired after short periods of time, became more and more prevalent, contributing to inefficient outcomes as medical students were pressured into making decisions without knowing the full scope of their possibilities. In the market for medical residents, the problem is matching medical students to hospitals. At the time, the NRMP was using hospital-proposing DA [3].

The theoretical aspects of matching theory began to be more robustly developed and studied during this time. Roth and Sotomayor laid out the foundational theorems and results of matching theory in their book "Two sided matching" [22]. Strategyproofness of DA also began to be studied. It was proved by Roth and Sotomayor that in the hospital-proposing algorithm, students have incentive to misrepresent their preferences, and vice versa for hospitals when running the student-proposing algorithm [22]. Matching theory attracted considerable interest from economists as well as mathematicians and computer scientists. It was even picked up in the medical literature. Persistent concerns about the NRMP in the medical literature were notably summarized by Williams [23], which attracted much interest and debate in the medical community. This eventually led to a redesign of the NRMP in 1999: Roth and Peranson modified the old NRMP algorithm, building upon its Gale-Shapley foundations [4] [3]. Some key changes were making it student-proposing instead of hospital-proposing, so that it is strategyproof for the students (students have no incentive to misrepresent their preferences) and allowing students to participate as couples [4]. The Roth-Peranson algorithm was also adopted in Canada, although the authors do not know when or how.

2.2 Recent developments

The theory behind matching has continued to see consistent progress. Kelso and Crawford introduced matching with money, which can be seen as a model of firms hiring workers [24]. Later, Hatfield and Milgrom unified and subsumed Kelso-Crawford and Gale-Shapley with their matching with contracts model [25]. Their model continues to be built upon and placed

upon more rigorous foundations by others [26].

Developments in matching theory tend to occur in tandem with practical considerations. For this reason, much of the recent literature has focused on the harder problem of *matching with constraints* [27]. 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 [28]. Clearly, by the Pigeonhole Principle [29], 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 [28]. This result was later restated by Roth as: in the resident-hospital matching market, any hospital with empty spots in some stable matching receives exactly the same set of residents in any stable matching [30]. 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 [5] [6] studied the Japanese medical residency match, which uses 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) [5], 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) [5]. They propose a *flexible deferred acceptance* algorithm that results in stability and respects regional quotas [6], 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 [6].

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 [31] [32]. 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 [31]. Recent work in the computer science litera-

ture has found that checking the mere existence of a feasible matching with floor constraints is \mathcal{NP} -complete [33]. It remains unclear whether such constraints are tractable, and what the definitions of concepts like individual rationality and stability would be in such situations [33].

3 Model

3.1 Preliminaries

As per Roth and Sotomayor [22], 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 $h \in H$.
- P is a list of preferences, as follows:
 - For each $i \in I$, P_i denotes the preferences of student i over $H \cup \{\emptyset\}$, 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 .
 - For each $h \in H$, P_h denotes the preferences of hospital h over $I \cup \{\emptyset\}$, whence, as with the students, we derive the strict preference relation \succ_h , which is defined similarly.

A **matching** is a function $\mu : H \cup I \rightarrow \mathcal{P}(H \cup I)$ such that [22]:

1. $\mu(h) \subseteq I \cup \emptyset$ such that $|\mu(h)| \leq q_h$ for all $h \in H$, meaning no hospital exceeds its quota,
2. $\mu(i) \subseteq H \cup \emptyset$ such that $|\mu(i)| \leq 1$ for all $i \in I$, meaning every student is only matched to one hospital or not at all.
3. $i \in \mu(h) \iff \mu(i) = \{h\}$ for all $h \in H$ and $i \in I$, meaning a student is matched to a hospital if and only if the hospital is matched to a set containing the student.

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.

We call a pair $(h, i) \in H \times I$ a **blocking pair** if *both* of the following two conditions hold [22]:

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

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

3.2 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 [22]:

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 preference ranking P_h , while a hospital h receiving less than q_h proposals accepts all of its proposals.

Step k . Any student i who was rejected at step $k - 1$ proposes to its next most preferred hospital. Each hospital h always takes its q_h top students, and rejects the others. If a hospital receives a proposal from a student which it prefers to some student that it shortlisted before step k , it will shortlist the new proposal and remove the less preferred student from its shortlist.

The algorithm terminates when there are no more rejections.

The algorithm also gives a stable matching if the hospitals propose [22], 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: use of DA only implies existence, not uniqueness, of a stable matching [1].

3.3 Introducing compatibility constraints

We build upon the basic model in section 3.1. Gale and Sotomayor define a pair (h, i) to be **compatible** if they are on each other's preference rankings [34]. We define a student-hospital pair as **incompatible** if they do not appear on each other's preference rankings: $i \notin P_h$ and $h \notin P_i$, where P_h is

the preference ranking of h and P_i is the preference ranking of i .³

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 well⁴. A student and hospital are only compatible if they share the same language, so an English-speaking student will only apply to, and rank, English hospitals, and French-speaking students will only rank French hospitals. Therefore, bilingual students can apply to both English and French hospitals⁵. In turn, hospitals only rank students that apply to them.

We can generalize the idea of such language incompatibilities to any sort of incompatibility based on some arbitrary binary characteristic. In general, we define a **matching with compatibility constraints problem** as a standard hospital-residents model as per section 3.1 with the following additional constraints:

- There is a binary characteristic which can only take one of two values. Let this set of values be $C = \{c_1, c_2\}$.
- All students $i \in I$ have 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 $I_{1,2}$.
- 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 i is compatible with a hospital h if and only if they share

³Note that preferences do not necessarily need to be complete in such matching problems.

⁴There 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: English and French.

⁵Note 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 rankings are *consistent* in two-sided matching problems, meaning that for some hospital h and student i , $h \in P_i$ if and only if $i \in P_h$ [35]. 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

the same characteristic, and they are incompatible if they do not. See figure 1 for a representation.

We can go back to the example that motivated this construction and apply this terminology. 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.

4 Results

4.1 Stability

Stability is an important consideration in matching markets. As Roth has shown, instability often leads to a collapse of matching markets [22]. 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 [34]. The following lemma will help us to establish stability.

Lemma 4.1. *The hospital-residents problem with compatibility constraints is an instance of the SMI problem.*

Proof. Let S be a finite set of residency seats. 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 . 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 . 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 4.1. *With compatibility constraints, the Gale-Shapley algorithm yields a stable matching.*

Proof. Gale and Sotomayor showed that the Gale-Shapley algorithm yields a stable matching for the SMI problem [34]. Combining this result with lemma 4.1 completes the proof. \square

Therefore, we have shown that even when compatibility constraints are introduced as per section 3.3, the Gale-Shapley algorithm still finds a stable matching.

4.2 Existence of unmatched students or unfilled positions

As touched upon in the introduction of the paper, a key issue in the CaRMS is that some students go unmatched, despite more residency spots than students. As well, many spots also go unfilled, largely in Quebec. With our matching with compatibility constraints framework, we can demonstrate that such a result is theoretically possible. The theorem below as well as its proof gives us some indication as to how and why unmatched students and unfilled spots arise simultaneously.

Theorem 4.1 (Unmatched Existence Theorem). *In a matching with compatibility constraints problem where the number of students is equal to the total number of residency spots - i.e. $|I| = \sum_{h \in H} q_h$ - then it is possible that some students will not be matched or hospital positions will be empty.*


Proof. It suffices to show an example where this happens. Consider the following example, in the language of language compatibility, with four students i_1, i_2, i_3, i_4 , where i_1 and i_2 are English-only students, i_3 is a bilingual student, and i_4 is a French-only student. Let there be four hospitals h_1, h_2, h_3, h_4 , each with quota $q_h = 1$, so that there are in total $\sum_{h \in H} q_h = 4$ hospital residency spots. Let the preference rankings of the hospitals and students be as follows:

$$\begin{array}{l|l}
i_1 : h_1 \succ h_2 & h_1 : i_3 \succ i_1 \succ i_2 \\
i_2 : h_2 \succ h_1 & h_2 : i_2 \succ i_1 \succ i_3 \\
i_3 : h_1 \succ h_3 \succ h_2 \succ h_4 & h_3 : i_3 \succ i_4 \\
i_4 : h_3 \succ h_4 & h_4 : i_3 \succ i_4
\end{array}$$

Now, we run the student-proposing deferred acceptance algorithm, in rounds:

1. Student i_1 proposes to hospital h_1 . Student i_2 proposes to hospital h_2 . Student i_3 proposes to hospital h_1 . Student i_4 proposes to hospital h_3 . Hospital h_1 accepts student i_3 and rejects i_1 . Hospital h_2 accepts student i_2 . And, hospital h_3 accepts i_4 .
2. Student i_1 proposes to h_2 . Hospital h_2 rejects i_1 .

The final matching that the algorithm yields is:



Student i_1 , who is an English-only speaker, is not matched to a hospital, while the Francophone hospital h_4 has its spot empty simultaneously. \square

Remark 1. Note that the Francophone hospital spot is unfilled in the above example, which is analogous to how the bulk of unfilled residency spots are in Quebec in the CaRMS.

Remark 2. Observe that if the constraints in the above example are lifted, then all positions and students would be matched.

Theorem 4.1 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 spots [22]. The proof of theorem 4.1 shows how introducing compatibility constraints leads this result to break down, even though in the example given students preferences'

are complete over *compatible* hospitals, and vice versa (this is as complete as preferences can be under compatibility constraints).

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 [11]. Remarkably, we can demonstrate that the problem of unmatched students persists.

Corollary 4.2. *In matching with compatibility constraints where the number of residency positions is greater than the number of students, it is possible for students to go unmatched or hospital positions to be unfilled or both.*

Proof. Let us build on the example given in the proof of theorem 4.1 by adding a fifth hospital h_5 (with one seat), which is Francophone, with preference rankings $h_5 : i_4 \succ i_3$, and updating the Francophone students' preference rankings to $i_3 : h_1 \succ h_3 \succ h_2 \succ h_4 \succ h_5$ and $i_4 : h_3 \succ h_4 \succ h_5$. It is clear to see that both of these students i_3 and i_4 are already matched to other hospitals during the progression of the algorithm, and so will have no need to propose to h_5 , so that h_5 is unfilled. The student i_1 remains unmatched as well. \square

Remark 3. *In the above example, consider instead if we added a fifth hospital h_5 (with one seat) which is Anglophone, with preference rankings $h_5 : i_1 \succ i_2 \succ i_3$. Let us update the Anglophone student's preferences to $i_1 : h_1 \succ h_2 \succ h_5$, $i_2 : h_2 \succ h_1 \succ h_5$ and $i_3 : h_1 \succ h_2 \succ h_5$. If we run the student-proposing algorithm, student i_1 , after being rejected by h_1 and h_2 , will propose to h_5 , and will be accepted by it. The other students' assignments remain the same as before, so that all students have been given positions. However, hospital h_4 still remains unfilled.*

4.3 Establishing an I -saturating stable matching

An I -saturating matching is defined as a matching in which, for all $i \in I$, $\mu(i) \neq \emptyset$ [36]. So, an I -saturating stable matching is such a matching that is also *stable*. Remark 3 gives some indication as to how we can vary the number of hospitals in order to match every student. When we added an Anglophone hospital, the number of Anglophone positions (three) was equivalent to the number of English speakers (i_1 , i_2 , and the bilingual student i_3). We can generalize this fact.

First, however, it is necessary to define a new term. With compatibility constraints, preferences are necessarily incomplete. However, we can define a weaker form of 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 can say that preferences are **compatibility-wise complete**. This is *as complete* as preferences can be under compatibility constraints. With this definition, we can show the following theorem.

Theorem 4.2. *In matching with compatibility constraints where preferences are compatibility-wise complete, if the number of Anglophone positions equals the number of English-speakers $|I_E| = \sum_{\forall h \in H_E} q_h$ and the number of Francophone positions equals the number of French-speakers $|I_F| = \sum_{\forall h \in H_F} q_h$, then the Gale-Shapley algorithm yields an I -saturating stable matching.⁶*

Proof. Assume for contradiction that there is a student i who does not have a position after running the DA algorithm. Let the number of students with the same characteristic as i , including i , be n . By assumption, the number of positions with i 's characteristic is also n . As preferences are compatibility-wise complete, the student i must have been rejected from every compatible hospital. Per the algorithm, a student is only rejected by a hospital if they are less preferable than some other student. Rejection can only happen if a hospital's shortlist is already full or a the student is removed to free up a position for a more preferable student. This means that every compatible hospital that rejected student i has filled its quota, meaning that n positions are filled. However, only $n - 1$ students are matched (every student with i 's characteristic except for i). By the Pigeonhole Principle [29], some student must have two positions. This is a contradiction. \square

Theorem 4.2 implies that, 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), 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. We can generalize this observation.

Corollary 4.3. *Assuming compatibility-wise completeness, the number of total positions required to guarantee every student a match is $|I_E| + |I_F|$.*

Proof. It follows from theorem 4.2, as we simply sum the required number of Anglophone positions $|I_E|$ and the required number of Francophone positions $|I_F|$. \square

⁶Note that English-speakers I_E includes English-only speakers and bilingual speakers, and French-speakers I_F includes French-only speakers and bilingual speakers.

Corollary 4.4. *If $I_E \cap I_F \neq \emptyset$, then the required number of positions to guarantee every student a match under compatibility-wise complete preferences is greater than the number of students I .⁷*

Proof. Observe that $|I| = |I_E| + |I_F| - |I_E \cap I_F| \Rightarrow |I_E| + |I_F| = |I| + |I_E \cap I_F|$. As $|I_E \cap I_F| > 0$, it follows that $|I_E| + |I_F| > |I|$. \square

The results of the above two corollaries shows the inefficiency introduced by compatibility constraints in matching markets. When such constraints are not present, having $\sum_{\forall h \in H} q_h = |I|$ suffices. However, when they are present, it necessitates more positions than $|I|$. As well, observe that with corollary 4.4, as the number of positions required is greater than the number of students, inevitably some positions will remain unfilled. This means 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 corollary 4.3 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.

4.4 A dynamic game extension

A policy maker might look at the unmatched students and positions and think about relaxing the assumptions we have made thus far about the language constraints of hospital positions. Consider an extension of the hospital-residents with compatibility constraints problem where there is a second match in which unfilled positions are re-advertised as “bilingual” with certain requirements. For example, an unfilled Francophone position might be re-advertised as open to all students regardless of language, on the condition that non-Francophones enrol in French classes concurrently alongside their residency. Then, unmatched students apply to these positions, and we use the Gale-Shapley DA algorithm once again. We may also allow previously matched students to forfeit their positions and re-participate in the match. Or, we may exogenously restrict them from participating, and only

⁷In the special case where there are no bilingual students, so that $I_E \cap I_F = \emptyset$, then the required number of positions to guarantee every student a position under compatibility-wise complete preferences is equal to the number of students $|I|$, which follows from corollary 4.3. Observe that when $I_E \cap I_F = \emptyset$, we effectively have two separate standard hospital-residents problems: one between I_E and H_E , and one between I_F and H_F .

allow the previously unmatched students can participate. Interestingly, we can show that the latter is simply a special case of the former.

A useful way of handling this game is to construct an SMI instance as per lemma 4.1: construct, for each hospital $h \in H$, q_h copies with the same preference relation P_h and place these copies in a finite set of residency seats S . As well, rewrite the preference relations for every student $i \in I$ by replacing every hospital h with its copies in S . We have reduced the problem to a one-to-one matching problem. With these preference relations, we can run the first match. After this run, we can update the preference relations in order to run the second match as follows.

First, identify which students wish to give up their seats in order to participate in the match. Vacate their seats and also open them up for applications. For all students $i \in I$ who wish to keep their seat, simply rewrite their preference relation to be $i : \mu(i) \succ \emptyset$, where $\mu(i)$ is their assigned seat from the first match. Conversely, for all residency seats $s \in S$ that are not given up, rewrite their preference relations as $s : \mu(s) \succ \emptyset$. Next, update the preference relations of the students without positions (whether they were unmatched a priori or chose to give up their seat) as they now apply to the newly open positions. Similarly, update the preference relations for these open positions. This is an instance of SMI between the sets I and S . Note that the way we have constructed it guarantees that those students and positions that choose to not participate in the second match will necessarily remain with their assignment from the first match.

Theorem 4.3. *In the second match, where previously unfilled positions are re-advertised as unrestricted, the Gale-Shapley algorithm yields a stable matching.*

Proof. Use the construction above to reduce the problem to SMI problem instance. By Gale and Sotomayor, Gale-Shapley yields a stable matching [34]. \square

We can see how the case where exogenously restricting the second match only to previously unmatched students is a special case of the above game, as it is merely the case that *all* students with positions are treated as per the construction above. Therefore, the following corollary is implied.

Corollary 4.5. *When students who were previously matched are exogenously restricted from participating in the second match, the Gale-Shapley algorithm yields a stable matching.*

Proof. Use the same construction as in the proof of theorem 4.3, but assign to *all* previously matched students i the preference relation $i : \mu(i) \succ \emptyset$, and all previously filled seats s the preference relation $s : \mu(s) \succ \emptyset$. This is a special case of theorem 4.3, and so Gale-Shapley yields a stable matching. \square

5 Discussion and Conclusion

In this paper we developed the matching with compatibility constraints model, where a binary 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 (ex. 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.

Lastly, we investigated an extension into a dynamic game where unfilled positions are re-advertised with relaxed restrictions in a second match. Whether this second match is only open to unmatched students, or previously matched students can renege on their assignment and participate it, Gale-Shapley yields a stable matching.

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 binary 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] 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.
- [6] 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.
- [7] Canadian Residency Matching Service. Home, Dec 2019.
- [8] Milton J Esman. The politics of official bilingualism in canada. *Political Science Quarterly*, 97(2):233–253, 1982.
- [9] 2019 main residency match by the numbers, Sep 2019.
- [10] Carms update: 40% fewer students go unmatched over last year.
- [11] AFMC. Reducing the number of unmatched canadian medical graduates: A way forward. Technical report, Association of Faculties of Medicine of Canada, Jan 2018.
- [12] C Ruth Wilson and Zachary N Bordman. What to do about the canadian resident matching service. *CMAJ*, 189(47):E1436–E1447, 2017.
- [13] 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.

- [14] Eugene YH Yeung. Canadian resident matching service: are we trying to reinvent the wheel? *CMAJ*, 2019.
- [15] Janice Willett. What to do about the canadian residency matching process. *CMAJ*, 2019.
- [16] Sarah L Silverberg and Kaylynn MH Purdy. Unmatched canadian medical graduates. *CMAJ*, 190(4):E118–E118, 2018.
- [17] Amit Persad. The overall culture of residency selection needs fixing. *CMAJ*, 190(14):E443–E443, 2018.
- [18] Amit Persad. The unmatched. *Canadian medical education journal*, 9(2):e89, 2018.
- [19] 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.
- [20] Franklin Warsh. Robert chu’s suicide sends a message medicine cannot ignore. *Toronto Star*, Jun 2017.
- [21] Alvin E Roth. The college admissions problem is not equivalent to the marriage problem. *Journal of economic Theory*, 36(2):277–288, 1985.
- [22] Alvin E Roth and Marilda Sotomayor. Two-sided matching. *Handbook of game theory with economic applications*, 1:485–541, 1992.
- [23] Kevin Jon Williams. A reexamination of the nrmp matching algorithm. national resident matching program. *Academic medicine: journal of the Association of American Medical Colleges*, 70(6):470–6, 1995.
- [24] Alexander S Kelso Jr and Vincent P Crawford. Job matching, coalition formation, and gross substitutes. *Econometrica: Journal of the Econometric Society*, pages 1483–1504, 1982.
- [25] John William Hatfield and Paul R Milgrom. Matching with contracts. *American Economic Review*, 95(4):913–935, 2005.
- [26] Orhan Aygün and Tayfun Sönmez. Matching with contracts: Comment. *American Economic Review*, 103(5):2050–51, 2013.
- [27] Fuhito Kojima. Recent developments in matching theory and its practical applications. In *Advances in Economics and Econometrics*:

- Eleventh World Congress of the Econometric Society, Econometric Society Monographs*, pages 138–175, 2015.
- [28] David G McVitie and Leslie B Wilson. Stable marriage assignment for unequal sets. *BIT Numerical Mathematics*, 10(3):295–309, 1970.
- [29] Tamara J Lakins. *The tools of mathematical reasoning*. American Mathematical Soc., 2016.
- [30] 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.
- [31] Yuichiro Kamada and Fuhito Kojima. Efficient matching under distributional constraints: Theory and applications. *American Economic Review*, 105(1):67–99, 2015.
- [32] 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.
- [33] 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.
- [34] David Gale and Marilda Sotomayor. Some remarks on the stable matching problem. *Discrete Applied Mathematics*, 11(3):223–232, 1985.
- [35] 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.
- [36] Alan Gibbons. *Algorithmic graph theory*. Cambridge university press, 1985.