Something New in Medical Residency Matching Markets

Zhuojin Yu

November 2022

School of Economics

University of Sydney

Thesis submitted in partial fulfilment of the award course requirements of the Bachelor of Advanced Studies (Honours)

Supervised by Dr Onur Kesten

STATEMENT OF ORIGINALITY

I hereby declare that this submission is my own work and to the best of my knowledge it contains no material previously published or written by another person. Nor does it contain any material which has been accepted for the award of any other degree or diploma at the University of Sydney or at any other educational institution, except where due acknowledgment is made in this thesis.

Any contributions made to the research by others with whom I have had the benefit of working at the University of Sydney is explicitly acknowledged.

I also declare that the intellectual content of this study is the product of my own work and research, except to the extent that assistance from others in the project's conception and design is acknowledged.

[REDACTION] [redacted]

Zhuojin Yu

6th November 2022

ACKNOWLEDGEMENT

I would like to start by thanking my supervisor – Dr Onur Kesten. He offered me guidance on my research directions and always supported my rational and irrational conjectures. Without his constant encouragement and uplifting reassurance, I cannot begin to image finishing the thesis in my wildest dream.

And I cannot express my thanks enough to everyone I interviewed from NSW Allocation Committee -- Susan Sellars, David Cross, David Roberts, David Schmidt, Gretchen Buck, Julia Blackford, Noah Mitchell, and Josephus Paya. Without your help, I literally have no start and no research topic for my thesis. An especially loud thanks to Susan and Josephus, my first interviewer and my most inspiring interviewer. Josephus's selfless sharing of his expertise in computer science (still not sure how much I understand that) made my thesis possible.

Then I would love to thank my honours cohort. Thanks for their forbearance in my relentless complaining. Thanks for their time and brain cells to discuss my thesis and understand my confusion towards many, many things. I am so grateful that we are together in the stress, anxieties, struggles, hope, hard-work, happiness, thrills and excitement. I enjoy this year with you more than I can tell, every minute of it.

A massive thank-you for our heart-warming professors, admin staff and PhD students for their support and listening. I take shelter and grow in the environment they created.

And, of course, my family, especially my grandma and my mom. First of all, thanks for their hard-working so I don't have to worry about anything financially. I share every detail and every draft I wrote with them even though they all hate economics and can't understand a word of English. My grandma and mom supported me emotionally in a way that I feel so appreciative and lost, dreading what I am supposed to do if I lose them one day.

And lastly, I would like to thank a scarlet fallen leaf stuck in the traffic-light poll, a spider web entangled among flowers, a sun that washes the sky to blue, a breeze in my hair, a quiet café playing my favourite music, the Quad standing clear-cut from the maroon sky from the Fisher window, a stone I picked outside of SSB, Taylor

Swift's music…. All of you made my thesis a lasting memory, an irreplaceable experience, and I cannot think of a better way to spend my 21st year on this planet.

ABSTRACT

Worldwide medical residency markets commonly employ variants of the two-sided central clearinghouse designed by Roth and Peranson in 1999. In the NSW physiotherapy residency matching market, a one-sided and computationally efficient matching mechanism is used – the Kuhn-Munkres algorithm. The mechanism is new for medical matching markets, with no publicly known application and no existing literature. A crucial contribution of the thesis is presenting the algorithm and starting a discussion around the Kuhn-Munkres algorithm in matching. The thesis models the iterative working of the Kuhn-Munkres algorithm. I show that the Kuhn-Munkres algorithm is rank-efficient, outcome unfair, procedurally fair and not strategy-proof. Comparing the Roth-Peranson and Kuhn-Munkres algorithms on efficiency, fairness and incentive properties, the thesis concludes that there is no settled winner between the two algorithms. The competition eventually comes down to the trade-off between cost reductions and market complexities.

Table of Contents

1 Introduction

Medical graduates take entry-level positions in the residency as their first step to becoming a doctor. Residency programs are designed to offer a diverse experience, facilitating graduates to apply, consolidate and expand their knowledge within a supportive framework (NSW Physiotherapy Allocation, 2022). In physiotherapy, residency programs are outstandingly important. Physiotherapy is a growing profession with sustained pressure on workforce supply. Australian Physiotherapy Association (2021, general demographics and characteristics section, para.4) reported that “the ratio is 145 physiotherapists per 100 000 people in major cities, this drops to 89 per 100 000 in regional areas and 46 per 100 000 in remote and very remote areas”. Besides the labour supply shortage, an aging population continuously raises the demand for physiotherapy (IBISWorld, 2022). This labour market pressure calls for a well-designed matching process, positioning each graduate to their most suitable program.

NSW uses a central clearinghouse to implement the matching process – from medical graduates to hospital positions (NSW Physiotherapy Allocation, 2022). Centralized implementation alone doesn’t guarantee a well-designed matching process. Only an organized central clearinghouse can promote the proper functioning of the matching. First, an organized central clearinghouse can prevent much chaos in decentralized markets. In the 1920s, hospitals expected US medical students to sign residency contracts in their sophomore year, even before finalising their majors (Roth, 1984). Moreover, an offer was only active for a day or even an hour (Roth, 1984). The centralized clearinghouse adopted in the 1950s resolved both issues (Roth, 1984). Second, a central clearinghouse can accommodate changing market trends and needs. In 1999, Roth and Peranson (1999) modified the US medical residency matching for couples’ wishes to work in the same vicinity. Scotland also admitted couples’ preferences over pairs of hospitals since 2009 (Irving, 2011). In 2008, Japan introduced regional quotas to promote geographical equality for the distribution of new residents (Kamada & Kojima, 2015). Finally, unorganized central clearinghouses will not survive. When US and UK first implemented centralized mechanisms, both were abandoned quickly due to incentive problems and prearranged agreements outside the centralized process (Roth, 1984; Roth, 1991). Voluntary participation in the US dropped significantly in the 1970s due to the mechanism's incompetence in accommodating couples' preferences. The US gastrointestinal centralized market crashed wholly and suddenly due to a lack of confidence from students (McKinney, Niederle & Roth, 2005). Therefore, researchers should constantly scrutinise a central clearinghouse's functioning and properties.

NSW established the physiotherapy residency matching mechanism 40 years ago (NSW Physiotherapy Allocation, 2022). Unfortunately, there is neither public information nor economic literature about the NSW physiotherapy matching market. By interviewing many current Allocation Committee members[1], I gathered that the market had undergone significant transformations. Initially, the mechanism used pure GPA-based priority matching. A higher GPA meant a more elevated chance of getting into top preferences. Later, universities decided not to disclose graduates' GPAs, so the mechanism switched to random allocation. Several unlucky university medal winners ended up unmatched, urging a mechanism change. This luck-based mechanism was replaced by a 2-stage matching process – allocation and recruitment. First, the Allocation Committee assessed all applicants and allocated desirable applicants to hospitals for interviews. Participating hospitals admitted applicants who successfully passed the recruitment stage.

The current mechanism was enacted two years ago, which overall inherited the 2-stage process. Every applicant submits a rank-order list strictly preferring all hospitals (NSW Physiotherapy Allocation, 2022). Then the allocation process begins. In the new mechanism, the Allocation Committee only allocates, not assesses, applicants. The process assigns all applicants to hospitals for interviews without rejecting anyone in this process. The allocation quota for each hospital is determined by the rule below:

---

[1] Susan Sellars, David Cross, David Roberts, David Schmidt, Gretchen Buck, Julia Blackford, Noah Mitchell and Josephus Paya

> The number of candidates allocated to each hospital is determined by how many First Year Graduate positions are taken by that hospital relative to the overall number of FYG positions and the total number of candidates. For example, if a hospital takes 12 FYG positions and there are 120 total FYG positions they represent 10% of the program and will take 10% of the applicants. If there are 250 applicants, this hospital will have 25 7 candidates allocated. (NSW Physiotherapy Allocation, 2022, the allocation process section, para.2)

Therefore, the number of interview positions equals the number of applicants, guaranteeing each applicant one place. After receiving assigned applicants, each hospital manages the recruitment process individually. The recruitment process incorporates selection criteria, interviews, reference and employment requirements checks (NSW Physiotherapy Allocation, 2022). In December, hospitals inform their final decisions to the Allocation Committee (NSW Physiotherapy Allocation, 2022).

The thesis will solely focus on the allocation process since the recruitment process is just a merit-based selection. However, no published sources inform applicants what mechanism is in use. The only description of the allocation mechanism is

> The candidates within the pool are submitted to the OPPM computer program, which allocates all applicants to a hospital, based on giving the highest number of applicants as close as possible to their highest preference (NSW Physiotherapy Allocation, 2022, the allocation process section, para.1)

In an interview with the NSW Allocation Committee website admin[2], he told me that the allocation process utilizes the Kuhn-Munkres algorithm (Paya, private conversation). Interestingly, to the best of my knowledge, no publicly known medical matching markets harness the Kuhn-Munkres algorithm (abbreviated as KM

[2] Much thanks to our NSW Allocation Committee website admin, Josephus Paya

hereafter). Moreover, failing to find any matching literature about the algorithm, KM seems to be a black box to economists. The thesis correspondingly aims to unravel the mysterious KM from an economics perspective for the first time. There are three puzzles this thesis seeks to resolve: How does KM work; What are KM's properties; Is KM working well compared to the existing algorithms?

KM is an iterative way to assign as many applicants as possible to their top preferences. I propose that KM is rank efficient, outcome unfair, procedurally fair and not strategy-proof. One of the most prevalent mechanisms used in the worldwide medical matching market is the algorithm designed by Roth and Peranson (1999) (abbreviated as RP hereafter). Regarding efficiency, KM consistently achieves a stronger efficiency than RP both theoretically and practically. KM is always outcome unfair, while RP is outcome fair in simple markets. KM and RP are both always procedurally fair, theoretically and practically. Theoretically, KM and RP are not strategy-proof but with some protection against strategy behaviours. In practice, KM and RP are both unavoidably subject to applicants' gaming.

At the first glance, it may seem that KM fares better than RP on efficiency, worse than RP on fairness and that equally as RP on incentive properties. However, one should not take the results at face value. KM is a matching mechanism in one-sided simple matching markets, and RP is for two-sided complex matching markets. To address the differences, I restrict theoretical discussions to simple markets, while RP's complications in complex markets are referred to as practical comparisons. It is therefore concluded that there is no easy win or lose between RP and KM, propelling future exploration of KM.

The rest of the paper is structured as follows. Section 2 introduces the background on KM, other mechanisms and algorithms in medical matching markets. Section 3 sets up the model and explains how KM functions. Section 4 investigates KM's efficiency, fairness and incentive properties. Section 5 discusses the theoretical and practical comparisons between KM and RP. Section 6 concludes and discusses further research directions.

2 Background

2.1 KM Evolvement and Application

KM is a combinatorial optimization algorithm initially designed to solve the personnel-assignment problem. The personnel-assignment problem tries to find an optimal way to assign N personnel to N jobs (Kuhn, 1955). Each personnel has a numerical rating on each job (Kuhn, 1955). Kuhn (1955) described the aim of KM as maximizing the sum of personnel's ratings for their assigned jobs. While Flood (1956) argued that the optimal assignment by KM should minimize the sum of the ratings if the ratings are costs incurred for jobs. Munkres (1957) proved the equivalence of KM's two functions. KM can be used to reach a maximum or a minimum of the sum of ratings assigned (Munkres, 1957).

Kuhn designed the algorithm in 1955. Flood (1956) simplified Kuhn's method, producing a set of possible solutions. To produce a unique deterministic algorithm outcome, Munkres (1957) manufactured a process based on the transfer series in Kuhn's original approach. All these attempts make KM a polynomial-time algorithm, dramatically reducing the maximum number of operations required for a solution from $N!$ to $\frac{11N^3+12N^2+31N}{6}$ (Munkres, 1957). KM's computational efficiency is remarkable compared to the time complexity required for locating a similar solution in two-sided markets. Gusfield and Irving (1989) proposed an optimal stable solution that minimizes the sum of agents' ratings on the other side of the market. They found that at most $(N^4 logN)$ operations are needed for a solution, which is almost computationally impossible for a large market (Gusfield and Irving, 1989).

The computational efficiency of KM predicts its wide application. "Munkres assignment can be applied to TSP (travelling-salesman problem), pattern matching, track initiation, data correlation, and (of course) any pairwise assignment application" (Bevilacqua Research Corporation, 2022, para. 5, point 14). KM's implementation is especially prevalent in energy management (Mirzaeinia & Hassanalian, 2019; Sanseverino et al., 2015). The University of Texas Southwestern students proposed to optimize medical rotation schedules using KM (MacLean et al., 2020). The existing research scrutinizes how KM is applied. This paper bridges the gap by asking – how and how well does KM function?

2.2 Worldwide Medical Residency Matching Mechanisms

2.2.1 the United States (NRMP[3]; APPIC[4]) and Canada (CaRMS[5])

The US medical residency, US professional clinical psychology internship and Canada residency matching markets all use RP. Applicants register online and apply to hospitals for interviews, after which applicants and hospitals both submit strict preference lists (see the details in NRMP, APPIC and CaRMS). In addition to traditional one-by-one rankings, the markets allow couples' preferences over pairs of hospitals and applicants' preferences over pairs of complement positions (Roth & Peranson, 1999). RP inputs applicants' and hospitals' preferences and outputs a set of matched applicant-hospital pairs. RP is the applicant-proposing deferred acceptance mechanism without preferences with complementarities (i.e. the two cases mentioned) (Roth & Peranson, 1999). To address the complications in complex markets, RP identifies and eliminates blocking pairs[6] caused by an applicant with complement preferences before processing the next applicant (Roth & Peranson, 1999). This technique is based on the instability-chaining model – blocking pairs are tackled once at a time until there is none left (Roth & Vande Vate, 1990). If the blocking pairs start to cycle, RP proceeds with other applicants, forcing the termination of a cycle[7] (Kojima, Pathak & Roth, 2013). RP (Roth & Peranson, 1999) works in the following way:

(1) Match a single applicant to a hospital using the applicant-proposing deferred acceptance mechanism
(2) When the applicant has complement preferences between positions and displaces more than one applicant or one position or both, solve blocking

[3] National Resident Matching Program
[4] Association of Psychology Postdoctoral and Internship Centers
[5] Canadian Resident Matching Service
[6] A blocking pair is an applicant-hospital pair, which both prefer each other than their current matching.
[7] Exactly how the force rumination of a cycle works is publicly unavailable (Kojima, Pathak & Roth, 2013)

pairs created by the applicant one at a time until all bocking pairs created are eliminated

(3) Repeat (1) and (2) until all singles are processed

(4) Match a couple to a pair of hospitals using the applicant-proposing deferred acceptance mechanism

(5) When the couple displace more than one applicant or one position or both, solve blocking pairs created by the couple one at a time until all bocking pairs created are eliminated

(6) Repeat (4) and (5) until all couples are processed

2.2.2 Japan (JRMP[8])

Japan's residency matching mechanism operates identically as in the US after a central adjustment to all hospital's quotas (see the details in JRMP). In 2008, the market introduced regional caps to curtail geographical inequality in the distribution of doctors, restricting the number of applicants assigned to a specific area (Kamada & Kojima, 2015). The hospitals' quota within the same prefecture will be reduced proportionately to its original capacity if their sum exceeds the regional cap (Kamada & Kojima, 2015).

2.2.3 the United Kingdom (UKFP[9])

Before introducing the matching mechanisms, there are a few notable differences in the UK matching market from the US. First and foremost, it doesn't require hospitals to submit preferences (Biró, Irving & Schlotter, 2011). Hospitals' preferences are derived from the master list, which is a strict ranking of students based on their scores (Biró, Irving & Schlotter, 2011). So no student is expected to apply individually for interviews. Second, participation in this residency program is compulsory (Roth, 1991). Third, couples are not allowed to submit preferences over pairs of foundation schools and individual programmes. Couples need to submit two identical preference lists, which will be linked in the matching (UK Foundation Programme, 2022). Finally, the matching is divided into two stages – national and local matching. The national matching process allocates a student to a foundation

[8] Japan Residency Matching Program
[9] The United Kingdom Foundation Programme

school[10], while the local matching process allocates a student to a hospital within that foundation school (UK Foundation Programme, 2022). Therefore, once allocated to a particular foundation school, students need to submit another strict preference list of individual programmes in that foundation school (see the details in UK Foundation Programme, 2022).

The national matching process uses serial dictatorship (Kamada & Kojima, 2015). Students' scores determine the priority order (UK Foundation Programme, 2022). When students' scores result in a tie, decile scores and then SJT test scores are utilized to break the tie (UK Foundation Programme, 2022). If both scores still result in a tie, the tie will be broken randomly (UK Foundation Programme, 2022). This priority order of students is the master list and will be carried into the local matching process (Biró, Irving & Schlotter, 2011). The local matching process is managed by each foundation school differently. This paper only introduces the matching mechanism in the largest regional market – Scotland.

Students assigned to the Scotland foundation school submit a strict preference list of 10 programs (Irving, 2011). Hospitals' preferences are derived from the master list, and couples' preferences are constructed consistently[11] from their individual preference lists (Biró, Irving & Schlotter, 2011). The matching mechanism is a variant of RP, using a different processing sequence of agents in the algorithm (Biró, Irving & Schlotter, 2011). The Scotland algorithm first processes singles and randomizes the order of couples and displaced students by complement preferences (Biró, Irving & Schlotter, 2011). When the blocking pairs cycle, this algorithm reverts tp another master list which may be different due to random tie-breaking (Biró, Irving & Schlotter, 2011). The algorithm is repetitively implemented to select the matching with the minimum number of blocking pairs (Biró, Irving & Schlotter, 2011).

2.2.4 Israel

Unlike all mentioned markets, Israel's internship matching market is one-sided. Only students need to submit strict preference lists (Roth & Shorrer, 2015). Each hospital's quota is centrally determined by the Ministry of Health (Bronfman et al.,

[10] One region usually has one foundation school. In total, there are 19 foundation schools in the UK.
[11] See how couples' preferences are constructed in (Biró, Irving & Schlotter, 2011)

2018). The mechanism unfolds in three stages. First, the mechanism runs random serial dictatorship repeatedly (Bronfman et al., 2015). Averaging over all assignments, an expected probability vector for all positions is computed for all interns (Bronfman et al., 2015). Second, the mechanism uses linear programming to maximize the sum of interns' utilities for their assignments subject to several constraints[12] (Bronfman et al., 2015). The mechanism assumed a common utility function for all interns, constructed from a survey (Bronfman et al., 2015). Finally, the probability matrix is decomposed into a convex combination of deterministic assignments (Bronfman et al., 2018). The final assignment is chosen randomly per the weights on those deterministic assignments (Bronfman et al., 2015). Israel's internship matching market is of the closest setting to NSW – both one-sided matching market and hospital quotas are determined centrally. But Israel's matching mechanism varies significantly from KM, stressing the novelty of KM in matching literature and the importance of an investigation in KM.

2.3 Algorithms in Matching

Although KM is a brand-new application, researchers have already attempted operational algorithms as answers to problems in medical matching markets. However, all proposed algorithms currently remain on the paper. KM, as an algorithm already in use, serves as a theoretical and empirical connection between operational research and matching literature. Section 2.3.1 outlines the upfront challenge of using an algorithm in a matching market. Section 2.3.2 describes the current algorithm's attempts to solve the problems.

2.3.1 Hard to find an efficient algorithm

An algorithm greatly automates the matching process if solvable in a reasonable amount of time. Unfortunately, finding an efficient algorithm has proved to be difficult, if not impossible, in two-sided complex matching markets. Ronn (1990)

---

[12] (1) hospital quota constraint (2) couples consistency constraint – a guarantee to match into the same hospital (3) do not harm: an agent cannot be worse off than the expected probability vector (Bronfman et al., 2015)

proved that the problem of if a stable matching[13] exists is NP-complete[14], even when all hospitals have one quota and only couples present in the market. Biro, Manlove and Mcbride (2014) proved NP-completeness under further stronger restrictions than Ronn's assumptions, assuming couples' and hospitals' preferences lists are at most of length 2. The problem remains NP-complete even if there is only one pair of hospitals on couples' preference lists (Manlove, Mcbride & Trimble, 2017). When hospitals' preferences are derived from strict master lists (like in the UK), and all hospitals only have one quota, finding a stable matching is still NP-complete (Biró, Irving & Schlotter, 2011). McDermid and Manlove (2010) shifted the attention to restricting the preference domain of couples. Under the restriction, assuming each applicant's preference list has a length of at most 3 and hospitals have at most 2 quotas, the existence of a stable matching is still NP-complete (McDermid & Manlove, 2010). The computational complexity of finding a desirable solution is still prominent in one-sided matching markets. It is proved that decomposing a stochastic assignment matrix is NP-hard[15] with couples (Bronfman et al., 2018). That is, the problem takes, at best polynomial time operations. Although applied in a theoretically simple one-sided matching market, KM is an efficient algorithm for a desirable solution.

2.3.2 Attempts to find an efficient algorithm for a desirable outcome

The first method researchers try is to restrict the market complexities. Although an efficient algorithm is unlikely to exist, these NP-completeness results provide meaningful boundaries to search for a polynomial-time solvable algorithm in two-sided matching markets (Manlove, Mcbride & Trimble, 2017). Manlove et al. (2007) assumed no complement preferences and applied the constraint programming algorithm, obtaining better time and space complexities than the deferred-

---

[13] A stable matching is a desirable outcome where (1) no agent prefers unmatched to its current partner (2) no applicant-hospital pair such that both prefer each other than its current partner (Roth, 1982).

[14] A class of computational problem in which no efficient algorithm is found now, that is, can be solved in polynomial time complexity; NP stands for nondeterministic polynomial – Wikipedia, 2022

[15] NP-hard is a class of problems that are at least as hard as NP-problems – Wikipedia, 2022.

acceptance mechanism. They pointed out that variants of the markets (including couples) can be readily tackled by adding side constraints (Manlove et al., 2007). McDermid and Manlove (2010) proposed a polynomial-time algorithm that always finds a stable matching after restricting couples' preferences and hospital quotas (McDermid & Manlove, 2010).

The second method researchers try is a solution that approximates the desired aim, albeit not reaching the goal exactly. In complex two-sided markets, researchers admit solutions minimizing the number of blocking pairs -- maximum cardinality stable matching (Biro, Manlove and Mcbride, 2014; Manlove, Mcbride & Trimble, 2017). Biro, Manlove and Mcbride (2014) put forward the first integer programming algorithm, while Manlove, Mcbride and Trimble (2017) proposed a constraint programming algorithm. Constraint programming on average is faster than integer programming in finding this most stable solution, and they both admit a very low number of blocking pairs (Manlove, Mcbride & Trimble, 2017). In a one-sided ordinal market with couples, Bronfman et al. (2018) suggested a polynomial-time algorithm for approximately decomposing the target stochastic matrix.

The third method researchers try is to apply efficient existing algorithms directly to matching markets. A stable matching can be found in complex two-sided markets using SAT and IP algorithms (Drummond, Perrault & Bacchus., 2015). SAT and IP solvers are efficient when the deferred-acceptance algorithm fails to find a stable matching (Drummond, Perrault & Bacchus., 2015). Answer set programming is another off-the-shelf algorithm (de Clercq et al., 2016). It is easily adaptable to matching markets with couples and preferences with ties, locating a fair and maximum cardinality stable matching (de Clercq et al., 2016). Boehmer, Heeger and Nirdermeier (2022) found XP algorithm can find a stable matching, considering preferences and with ties and applicants' preferences change over time. This list is by no means exhaustive but as strong evidence suggesting the close relation between operational research and matching markets. KM, another off-the-self solver, fills the gap as an efficient algorithm for one-sided matching markets with strict preferences and no couples.

3 Model

3.1 Model Setup

A one-sided assignment problem is a tuple $\langle A, H, P, X \rangle$. KM is a mapping $\mu: P \rightarrow X$. $A = \{a_1, a_2, \dots a_i, \dots, a_n\}$ is a finite set of applicants. $H = \{h_{11}, h_{12}, \dots, h_{jk} \dots, h_{mn}\}$ is a finite set of hospital positions. There are m hospitals and n positions. h_{jk} is the kth position among all positions belonging to hospital j. The Allocation Committee assigns each hospital a quota such that the number of positions equals the number of applicants (NSW Physiotherapy Allocation, 2022). So the set of H is of the same size n as the set of A. $P_{n\times n}$ is the applicants' preference matrix. Each cell $p_{a_i h_{jk}}$ denotes a_i's preference for position h_{jk}. The $a_i th$ row of P is the applicant a_i's preference for all hospital positions -- $p_{a_i} = \cup_{h_{jk}} p_{a_i h_{jk}}$. The $h_{jk}th$ column of P is all applicants' preferences for the position h_{jk} -- $p_{h_{jk}} = \cup_{a_i} p_{a_i h_{jk}}$. All applicants' preferences p_{a_i} are restricted on the domain Θ that contains all possible strict rankings over hospitals. One thing to be careful, applicants' preferences over hospitals are strict, not over positions. All positions belonging to the same hospital have the same preference number assigned, despite applicants submitting a strict ranking over hospitals. That is,

$$(3.1.1)\ \forall a_i \in A, \forall h_{jk}, h_{j'k'} \in H, \forall p_{a_i} \in \Theta,\ (p_{a_i h_{jk}} = p_{a_i h_{j'k'}}) \Leftrightarrow (j = j')$$

$X_{n\times n}$ is the assignment matrix on the domain Ω satisfying

$$(3.1.2)\ \forall a_i \in A, \forall h_{jk} \in H,\ x_{a_i h_{jk}} = \{0,1\}$$

$$(3.1.3)\ \forall a_i \in A, \forall h_{jk} \in H,\ \textstyle\sum_{a_i} x_{a_i h_{jk}} = \sum_{h_{jk}} x_{a_i h_{jk}} = 1$$

The KM assignment X also satisfies

$$(3.1.4)\ \textstyle\sum_{a_i} \sum_{h_{jk}} x_{a_i h_{jk}} p_{a_i h_{jk}} = \text{minimum}$$

3.2 General Intuition

In X, $x_{a_i h_{jk}} = 1$ represents that a_i is assigned to the position h_{jk}. To satisfy (3.1.3) in a deterministic assignment, only one applicant is assigned to one position, and one position only admits one applicant. Correspondingly, the set of final assignments is independent – no two of assignments lie in the same line of X (Munkres, 1957).

Condition (3.1.4) means that KM minimises the sum of applicants' preferences towards their assigned positions (Munkres, 1957). This goal is achieved by first assigning as many applicants as possible to their 1st positions. If all applicants consist of an independent set of assignments, then KM will terminate. If not, KM seeks to assign as many applicants as possible to their 1st and 2nd positions. KM continues in this manner until the set of final independent assignments includes all applicants. Therefore, KM assigns the maximum number of applicants to their first l positions independently, which is how KM "give(s) the highest number of applicants as close as possible to their highest preference" (NSW Physiotherapy Allocation, 2022, the allocation process section, para.1).

To be specific, KM iteratively assigns one more applicant to their first l preferences. If the number of applicants that can be assigned to their first l positions is already maximal, then the algorithm relaxes the constraint by the smallest possible value q. KM then tries to assign more applicants to their first $(l + q)$ positions. The number of assignments is bounded above n, and the constraints can be relaxed maximally to their first n positions, so the algorithm must terminate after a finite number of steps ((Munkres, 1957). The process of increasing assignments is captured by the series of independently assigned agents Z_k of increasing size N_k. KM terminates when $N_k = N$. The process of relaxing constraints is captured by the series of modified preference matrices P_k of decreasing entries' sum Q_k.

3.3 Step-by-Step Breakdown

This section presents the working of the algorithm that the NSW residency matching market now employs, which is a variant of KM for rectangular matrices (Paya, private communication). The general working of KM is demonstrated by the flow chart below. Steps 1 and 2 are preliminary steps trying to assign every applicant to their first preferences. Step 3 checks how many applicants are assigned. Step 4 decides if the set of independently assigned applicants is maximal given current constraints. If the set doesn't reach the maximum, then step 5 will assign one more applicant. If the set already reaches the maximum, then step 6 will relax current constraints. KM inputs the preference matrix P and outputs P_k. The starred zeros in P_k become 1 and all other entries become 0, producing the assignment matrix X.

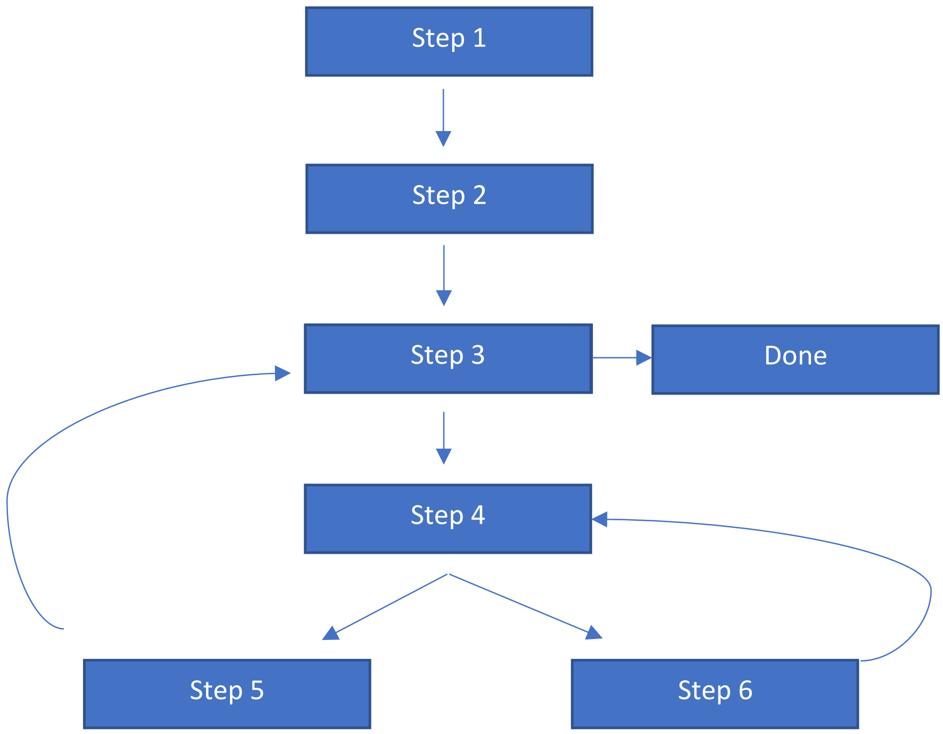

Step 1 Subtract the row minima for each row from P. Go to Step 2 (Bevilacqua Research Corporation, 2022)

To minimise the sum of preferences, KM tries to assign every applicant to their first choice. P becomes P^1 such that

$$(3.3.1)\ \forall a_i \in A, \forall h_{jk} \in H,\ p_{a_i h_{jk}} = 1 \Leftrightarrow {p^1}_{a_i h_{jk}} = 0$$

Step 2 Find a zero in P^1. If there is no starred zero in its row or column, star it. Repeat for all zeros in the matrix. Go to Step 3. (Bevilacqua Research Corporation, 2022)

An entry ${p^1}_{a_i h_{jk}} = 0$ means assigning the applicant a_i to position h_{jk} is a candidate assignment. Not all candidacy is realized as final assignments because KM requires the set of assignments to be independent. So, we need final assignments denoted by a unique symbol with no two in the same row or column -- starred zeros. One crucial detail is that KM always locates a zero by searching the first-row first column, first-row second column…, and so on before going to the next row. This detail impacts the fairness property of KM.

Step 3 Cover all columns containing a starred zero. If N columns are covered, KM terminates. Otherwise, go to Step 4. (Bevilacqua Research Corporation, 2022)

Covering a column is a way to keep track of the number of independent assignments. A column only has one starred zero. So, the number of columns covered equals the number of starred zeros, which is the number of assignments. When there are n assignments, KM terminates. One desirable but improbable case is that KM terminates after steps 1-3, when every applicant can be assigned to their first preferences.

Step 4 Find a noncovered zero and prime it. If there is no starred zero in its row, go to Step 5. Otherwise, cover the row and uncover the column with the starred zero. Repeat until all zeros in the matrix are covered. Go to Step 6[16]. (Bevilacqua Research Corporation, 2022)

This step is judging if the number of applicants assigned to their first l positions reaches the maximum. In the first case, when a row has a primed zero but no starred zero, it means that this applicant can be assigned but is now unassigned. So the series Z_k can continue to grow. In the second case, the maximal number of applicants are assigned to their first l positions, so KM resorts to constraint relaxation. Munkres (1957) proved why in the second case, the set of starred zeros is maximal. By construction, a starred zero is covered by exactly one line. Thus, the number of independent starred zeros cannot be more than the number of covered lines containing all zeros (Munkres, 1957). Following a mathematical theorem by König, the set of starred zeros describes the maximum number of independent assignments (Munkres, 1957).

Step 5 Construct a sequence Z. The first entry z_1 is the primed zero in Step 4. z_2 is the starred zero in z_1's column, if there is any. z_3 is the primed zero in z_2's row (there must be one). The series stops when z_{2k-1} has no starred zero in its column. In Z, star all primed zeros and unstar all

[16] As a special case, the matrix can have no uncovered zero. Then proceed to step 6.

starred zeros. Erase all primes and uncover all lines, go back to Step 3. (Bevilacqua Research Corporation, 2022)

Each iteration of Step 5 increases the size of Z_k from n_k to $(n_k + 1)$. This step assigns exactly one additional applicant to their first l preferences, without displacing any assigned applicants. Z starts by assigning an unassigned applicant, who may snatch an assigned applicant's current position. The assigned applicant will be reallocated to another position within her l rankings, which may again snatch another assigned applicant's current position. When an applicant is assigned to a vacant position without displacing anyone, Z terminates. Z describes a unique sequence of transfers. Because at most one primed zero is in a row, and at most one starred zero is in a column (Munkres, 1957). A starred zero must have a primed zero in its row, but a primed zero may not have a starred zero in its column. Therefore, Z must continue at a starred zero and stop at a primed zero. After increasing the assignment by one, KM goes back to step 3 to check how many applicants are assigned now.

Step 6 Find the smallest uncovered value q. Subtract q to every element in all uncovered columns. Add q to every element in all covered rows[17]*. Go back to Step 4.*

This step relaxes the constraint by the smallest value possible, so no optimal assignments will be skipped (Flood, 1956). Since it is impossible to assign all applicants to their first l preferences, KM tries to assign everyone to their first $(l + q)$ preferences. All zeros in the matrix are once-covered, so they will not be altered by updating P_k (Munkres, 1957) . Once a zero is created, it remains a zero. That's why KM assigns applicants to their first $(l + q)$ preferences, not just the $(l + q)th$ preference. After relaxing the constraint, KM goes back to step 4 to check if there is any room for additional assignments now.

3.4 An Example

Suppose the initial preference matrix P is

[17] The operations in Step 6 are identical to subtracting q from every uncovered element and adding q to every element covered by a column and a row.

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	1	1	2	2	3
$A2$	1	1	2	2	3
$A3$	1	1	2	2	3
$A4$	2	2	3	3	1
$A5$	2	2	3	3	1

$H1$ and $H2$ are two positions belonging to the same hospital. $H3$ and $H4$ belong to the same hospital. $H5$ is a hospital with only one quota. Yellow cells denote covered elements.

Step 1 Prepare all applicants' first preferences for assignments.

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	0	0	1	1	2
$A2$	0	0	1	1	2
$A3$	0	0	1	1	2
$A4$	1	1	2	2	0
$A5$	1	1	2	2	0

Step 2 Assign applicants independently to their first preferences.

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	0^*	0	1	1	2
$A2$	0	0^*	1	1	2
$A3$	0	0	1	1	2
$A4$	1	1	2	2	0^*
$A5$	1	1	2	2	0

Step 3 Only three applicants are assigned, so KM continues.

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	0^*	0	1	1	2
$A2$	0	0^*	1	1	2
$A3$	0	0	1	1	2
$A4$	1	1	2	2	0^*
$A5$	1	1	2	2	0

Step 4 There is no starred zero uncovered, so KM considers relaxing the goal of only assigning applicants to their first preferences.

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	0^*	0	1	1	2
$A2$	0	0^*	1	1	2
$A3$	0	0	1	1	2
$A4$	1	1	2	2	0^*
$A5$	1	1	2	2	0

Step 6 Now, applicants A1, A2 and A3 can be assigned to their first or second preferences.

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	0^*	0	0	0	2
$A2$	0	0^*	0	0	2
$A3$	0	0	0	0	2
$A4$	1	1	1	1	0^*
$A5$	1	1	1	1	0

Step 4 Identify a sequence where $A3$ could be assigned.

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	0^*	0	$0'$	0	2
$A2$	0	0^*	0	0	2
$A3$	0	0	0	0	2
$A4$	1	1	1	1	0^*
$A5$	1	1	1	1	0

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	0^*	0	$0'$	0	2
$A2$	$0'$	0^*	0	0	2
$A3$	0	0	0	0	2
$A4$	1	1	1	1	0^*
$A5$	1	1	1	1	0

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	0^*	0	$0'$	0	2
$A2$	$0'$	0^*	0	0	2
$A3$	$0'$	0	0	0	2
$A4$	1	1	1	1	0^*
$A5$	1	1	1	1	0

Step 5 A3 becomes assigned.

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	0	0	0^*	0	2
$A2$	0	0^*	0	0	2
$A3$	0^*	0	0	0	2
$A4$	1	1	1	1	0^*
$A5$	1	1	1	1	0

Step 3 Four applicants are assigned, so KM continues.

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	0	0	0^*	0	2
$A2$	0	0^*	0	0	2
$A3$	0^*	0	0	0	2
$A4$	1	1	1	1	0^*
$A5$	1	1	1	1	0

Step 4 Given the current preference matrix, KM already assigns the maximum number of applicants. So constraints need to be further relaxed.

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	0	0	0^*	$0'$	2
$A2$	0	0^*	0	0	2
$A3$	0^*	0	0	0	2
$A4$	1	1	1	1	0^*
$A5$	1	1	1	1	0

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	0	0	0^*	$0'$	2
$A2$	0	0^*	$0'$	0	2
$A3$	0^*	0	0	0	2
$A4$	1	1	1	1	0^*
$A5$	1	1	1	1	0

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	0	0	0^*	$0'$	2
$A2$	0	0^*	$0'$	0	2
$A3$	0^*	$0'$	0	0	2
$A4$	1	1	1	1	0^*
$A5$	1	1	1	1	0

Step 6 A4 and A5 can be assigned to all their preferences.

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	0	0	0^*	$0'$	3
$A2$	0	0^*	$0'$	0	3
$A3$	0^*	$0'$	0	0	3
$A4$	0	0	0	0	0^*
$A5$	0	0	0	0	0

Step 4 Find that A5 can be assigned.

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	0	0	0^*	$0'$	3
$A2$	0	0^*	$0'$	0	3
$A3$	0^*	$0'$	0	0	3
$A4$	$0'$	0	0	0	0^*
$A5$	$0'$	0	0	0	0

Step 5 Execute the transfer and assign A5.

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	0	0	0	0^*	3
$A2$	0	0	0^*	0	3
$A3$	0	0^*	0	0	3
$A4$	0	0	0	0	0^*
$A5$	0^*	0	0	0	0

For completeness, the final assignment matrix is

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	0	0	0	1	0
$A2$	0	0	1	0	0
$A3$	0	1	0	0	0
$A4$	0	0	0	0	1
$A5$	1	0	0	0	0

4 Properties

Proposition 1 KM is rank efficient.

4.1.1 Definition of Rank efficiency

Featherstone (2020) defined a new concept *rank efficiency* in the probabilistic matching market. The thesis defines the concept in the deterministic assignment market. The rank distribution of an assignment X is

(4.1.1) $\forall a_i \in A, \forall h_{jk} \in H, \forall l \in \{1,2,\dots,m\},\ N^x(l) = \sum_{a_i} \sum_{h_{jk}} x_{a_i h_{jk}} 1\{p_{a_i h_{jk}} \leq l\}$

The above equation characterises the number of applicants getting their lth preference or better in X. Another assignment Y is said to rank-dominate X if

(4.1.2) $\forall l \in \{1,2,\dots,m\}, N^y(l) \geq N^x(l)$

(4.1.3) $\exists l' \in \{1,2,\dots,(m-1)\}, N^y(l') > N^x(l')$

In KM, every applicant is assigned, so everyone receives her mth preference or better. That is,

(4.1.4) $N^y(m) = N^x(m)$

So l' can only range from 1 to $(m-1)$. An assignment X is rank efficient if it is not rank-dominated by any other assignment on the domain Ω (satisfying conditions 3.1.2 and 3.1.3). We say a mechanism is rank efficient if and only if its outcome X is rank efficient.

4.1.2 Proof of Rank efficiency

Suppose X is KM's assignment. The sum of applicants' preferences towards their assignments in X is

$$(4.1.5)\ 1 \times N^x(1) + 2 \times [N^x(2) - N^x(1)] + \cdots + m \times\ [N^x(m) - N^x(m-1)] =$$

$$(4.1.6)\ -\sum_{l=1}^{(m-1)} N^x(l) +\ m \times\ N^x(m)$$

By the same logic, the sum of preferences assigned for another assignment $Y \in \Omega$ is

$$(4.1.7)\ -\sum_{l=1}^{(m-1)} N^y(l) +\ m \times\ N^y(m)$$

From equation (4.1.4), we know that equation (4.1.7) can be replaced by

$$(4.1.8)\ -\sum_{l=1}^{(m-1)} N^y(l) +\ m \times\ N^x(m)$$

Suppose Y rank-dominates X, then

$$(4.1.9)\ -\sum_{l=1}^{(m-1)} N^y(l) < -\sum_{l=1}^{(m-1)} N^x(l)$$

So, KM's assignment X achieves a larger sum of preferences assigned than the assignment Y, contradicting that KM always minimises the sum of preferences assigned. KM's rank efficiency is also evident from the transfer processes. The algorithm iteratively assigns the maximum number of applicants to their first l preferences, leaving no room for rank dominance.

4.1.3 Rank efficiency and Pareto efficiency

Featherstone (2020) argued that in deterministic matching markets, rank efficiency implies Pareto efficiency but not vice versa. The intuition is that "Pareto concepts never insist on hurting one participant to help many, while rank efficiency often does so" (Featherstone, 2020, p. 2). This thesis presents proof of the efficiency concepts' comparison in deterministic markets[18].

Lemma 1 An assignment $Y \in \Omega$ that Pareto-dominates $X \in \Omega$ must rank-dominate X.

---

[18] See the proof for ordinal markets in Featherstone (2020)

Denote $p_{a'\mu^x(a')} = l_{a'}{}^x$ and $p_{a'\mu^y(a')} = l_{a'}{}^y$. Suppose $l_{a'}{}^x > l_{a'}{}^y$. No applicant is worse off in Y than X and a' is better-off in Y than X, so

$$(4.1.10)\ \forall l \in \{1,2,\dots,m\}, N^y(l) \geq N^x(l)$$

$$(4.1.11)\ \exists l_{a'}{}^y \in \{1,2,\dots,m\}, N^y(l_{a'}{}^y) > N^x(l_{a'}{}^y)$$

Lemma 2 An assignment $Y \in \Omega$ that rank-dominates $X \in \Omega$ may not Pareto-dominate X.

Suppose applicants a' and a'' are assigned to different preferences, while all other applicants are as well off in Y and X. Since Y rank-dominates X, we must have

$$(4.1.12)\ (l_{a'}{}^x + l_{a''}{}^x) - (l_{a'}{}^y + l_{a''}{}^y) > 0$$

Rearranging the above equation,

$$(4.1.13)\ (l_{a'}{}^x - l_{a'}{}^y) - (l_{a''}{}^y - l_{a''}{}^x) > 0$$

This inequality is satisfied when $(l_{a'}{}^x - l_{a'}{}^y) > 0$ and $(l_{a''}{}^y - l_{a''}{}^x) < 0$. In this case, a' is better-off and a'' is worse-off in Y than X. The presence of a worse-off applicant negates any possibility that Y Pareto-dominates X.

Therefore, an assignment is more easily rank-dominated than Pareto-dominated. Rank efficiency is stronger than Pareto efficiency, even in deterministic matching markets.

Proposition 2 KM is outcome unfair.

4.2.1 Definition of outcome fairness

In two-sided matching markets, fairness is defined as no justified envy -- if an applicant prefers a position to her current assignment, then that position must not prefer her to one of its current assignments. There is no direct parallel of such a fairness definition in one-sided markets. It is unreasonable to declare the unfairness of a mechanism because some applicants get better assignments than others, so criteria are required to order applicants. In the NSW residency matching mechanism, the rows of the preference matrix are ordered by their application time (Paya, private communication). The earlier an applicant submits her application, the lower row she is in. The row position will affect the final assignment. So, I define outcome fairness of an assignment as -- an early applicant never receiving a worse assignment than a late applicant with identical preferences. A mechanism is outcome-fair if and only if its assignment is outcome-fair for all applicants with identical preferences. Stated formally, if

$$(4.2.1)\ \forall h_{jk} \in H,\ p_{a'h_{jk}} = p_{a''h_{jk}} = p_{h_{jk}}$$

(4.2.2) a' is in a lower row than a''

then,

$$(4.2.3)\ \sum_{h_{jk}} x_{a'h_{jk}} p_{h_{jk}} \leq \sum_{h_{jk}} x_{a''h_{jk}} p_{h_{jk}}$$

This definition dictates that handing in the application early is never a disadvantage. If applicants with identical preferences cannot all get into the same hospital, the assignment that a late applicant wins the competition should not be warranted.

4.2.2 Proof of outcome fairness

It is easy to verify that KM is outcome unfair. In the example given in 3.4 (green cells denote final assignments),

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	1	1	2	2	3
$A2$	1	1	2	2	3
$A3$	1	1	2	2	3
$A4$	2	2	3	3	1
$A5$	2	2	3	3	1

$A1$, $A2$ and $A3$ have identical preferences. $A3$, who hands in her application last among the three, receives a better allocation than $A1$ and $A2$. $A4$ and $A5$ also have identical preferences. However, the early applicant $A4$ gets a better allocation than $A5$. How does this occur? Why applying early is sometimes disadvantageous and sometimes not?

Proposition 3 KM is procedurally fair.

4.3.1 Definition of procedural fairness

This irregular occurrence of disadvantages towards early applicants needs to be formalised by another fairness concept. Then we need to answer the question – why are the disadvantages not permanent?

KM's step 2 tries to assign all applicants to their first preferences, starting with early applicants. So late applicants remain unassigned if there is not enough quota for all applicants with identical preferences. To assign these applicants, early applicants are involved in transfers in step 5. They will be relocated to other zeros in their own rows, possibly newly created zeros from relaxing the constraint (strictly worse positions). $A1$ and $A2$ are displaced from their first preference for $A3$ to be assigned, while $A3$ snatches their freed-up position. This is why early applicants are disadvantaged.

However, transfers in KM are constructed based on the column order of the preference matrix, not based on applicants' rankings. This is the result of a crucial detail -- KM always locates a zero by searching the first-row first column, first-row second column…, and so on before going to the next row. Accordingly, an early applicant is relocated to the zero in her lowest uncovered column (the lowest uncovered column just before covering the row, not the first column), which need not be a worse position than the current assignment. Early applicants tend to involve in transfers, but involving in a transfer is not equivalent to a worse assignment (see appendix 8.1 for an example of where an applicant is better off in a transfer). $A5$'s lowest uncovered zero is in column 1, so she is assigned to her second preference, while $A4$'s assignment is untouched. That's why an early applicant may not be disadvantaged[19].

Whether an early applicant is disadvantaged depends on which positions have zeros and where those zeros are. Different preferences contribute to the irregular occurrence of an early application's disadvantage. Hence, the new fairness definition should consider fairness before any preference is realized.

The paper proposes a procedural fairness concept of an assignment – an early applicant never receives a worse assignment than a late applicant with identical preferences *on average, given all possible preferences*. Bolton, Brandts & Ockenfels (2005) argued that allocation fairness and procedural fairness are distinctly different.

---

[19] One can get a better allocation than her current assignment in a transfer (see the Appendix for an example)

Procedural fairness is crucial in that it levels the playing field, is ex-ante acceptable and also how modern societies deem fairness (Bolton, Brandts & Ockenfels, 2005). We say a mechanism is procedurally fair if and only if its assignment is procedurally fair for all applicants with identical preferences. As stated formally, if (4.2.1) and (4.2.2) are satisfied, then

$$(4.3.1)\ \forall p_{h_{jk}} \in \Theta, \mathrm{E}\left(\sum_{h_{jk}} x_{a'h_{jk}} p_{h_{jk}}\right) \leq \mathrm{E}\left(\sum_{h_{jk}} x_{a''h_{jk}} p_{h_{jk}}\right)$$

4.3.2 Proof of procedural fairness

To get the expected value of an applicant's preference towards her assignments, we need to randomly draw preferences from Θ and apply the law of large numbers for the sample average to approximate the expected value. The thesis, therefore, resorts to computer simulations. Suppose all applicants have identical preferences $p_{h_{jk}}$. For different realizations of $p_{h_{jk}}$, the averages of each row's rankings for their assignments asymptotically constitute the expected rank distribution of X. If the expected rank distribution is lower in the right tail than in the left tail, meaning late applicants indeed get better assignments than early applicants on average. That's when KM is procedurally unfair.

Applicants' preferences are randomly drawn from Θ, a space containing all random permutations over $[1,23]$. 23 is chosen because there are 23 hospitals in the NSW matching market (Paya, private communication). $p_{h_{jk}}$ is one random realization of Θ in each simulation. For simplicity, one hospital only has one quota, so there are 23 applicants. The expected rank distribution is the average of 500 simulations.

Figure 1 illustrates the resulting expected rank distributions. In graph (a), the expected rank distribution is lower in the left tail (lower rows) than in the right tail (higher rows). One may be tempted to conclude that early applicants get an advantage on average. However, the exact opposite conclusion will be inferred by investigating graph (b). This is because different random number generators are used in plotting the two rank distributions. Given the random number generator zero (graph (c)), there is no systematic difference between early and late applicants on average. These results corroborate our theoretical prediction – assignments are unfair to early applicants in some $p_{h_{jk}}$ and not unfair in other $p_{h_{jk}}$. Thus, in the graph (d), I shuffle the random number generators with each simulation[20]. The expected rank distribution exhibits no systematic favour towards late applicants. Therefore, KM is procedurally fair.

(a) Lower rows get better preference

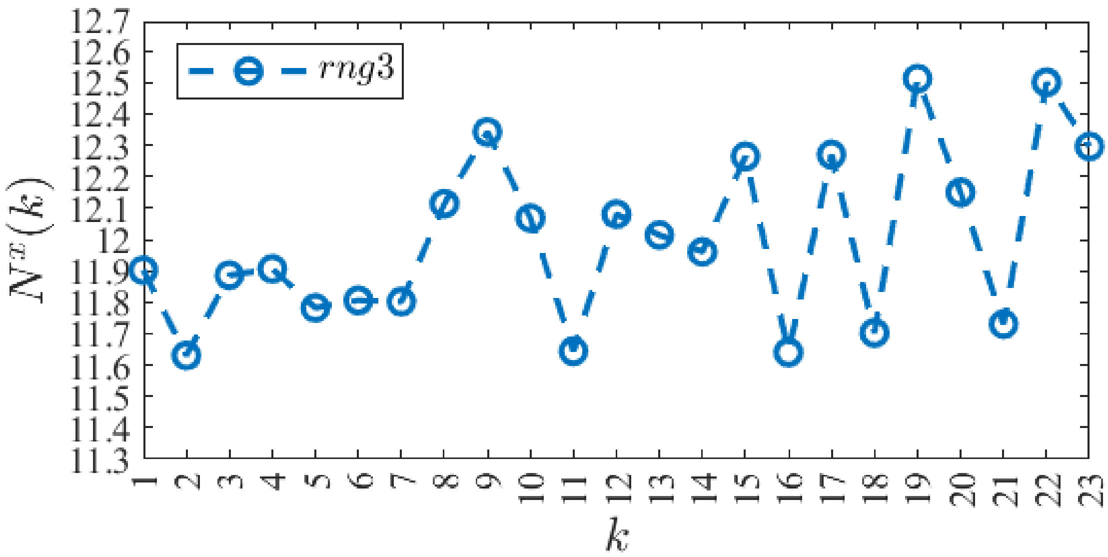

(b) Higher rows get better preference

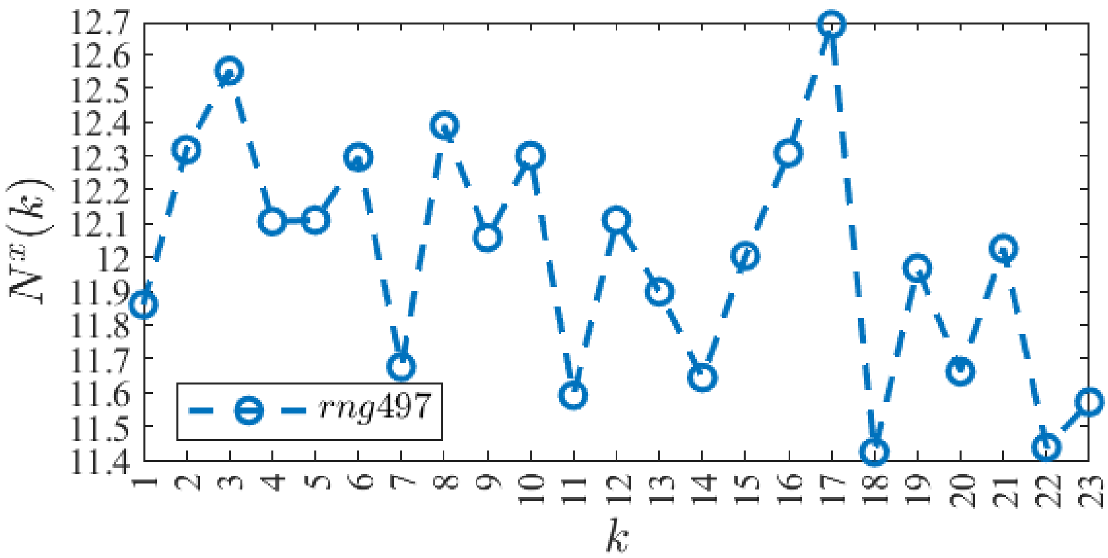

(c) No systematic differences for a random number generator

[20] Much thanks for Josephus Paya for suggesting me this idea of shuffling random number generators

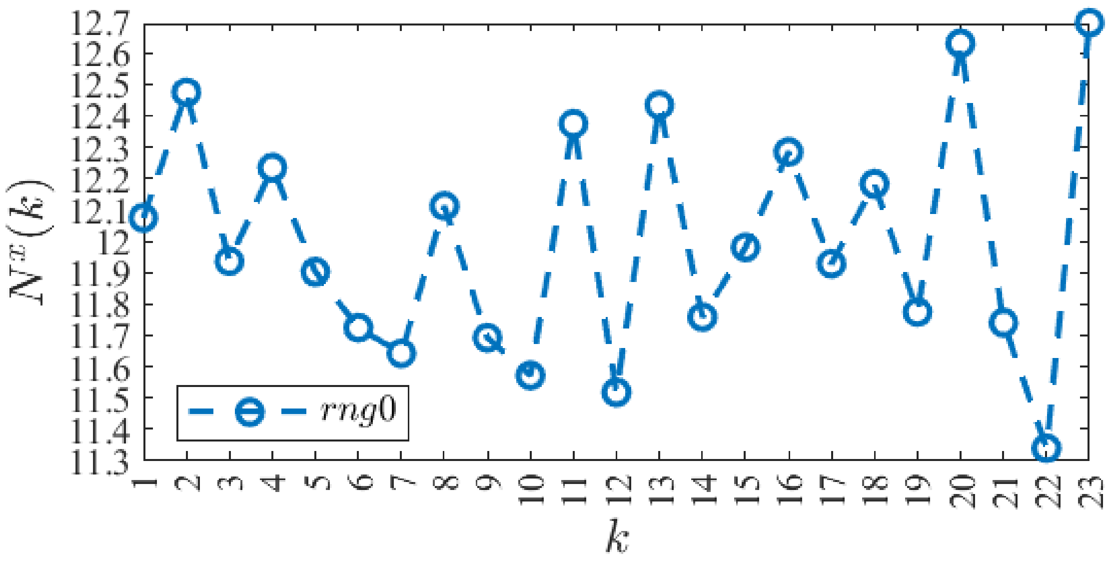

(d) No systematic differences for shuffling random number generators

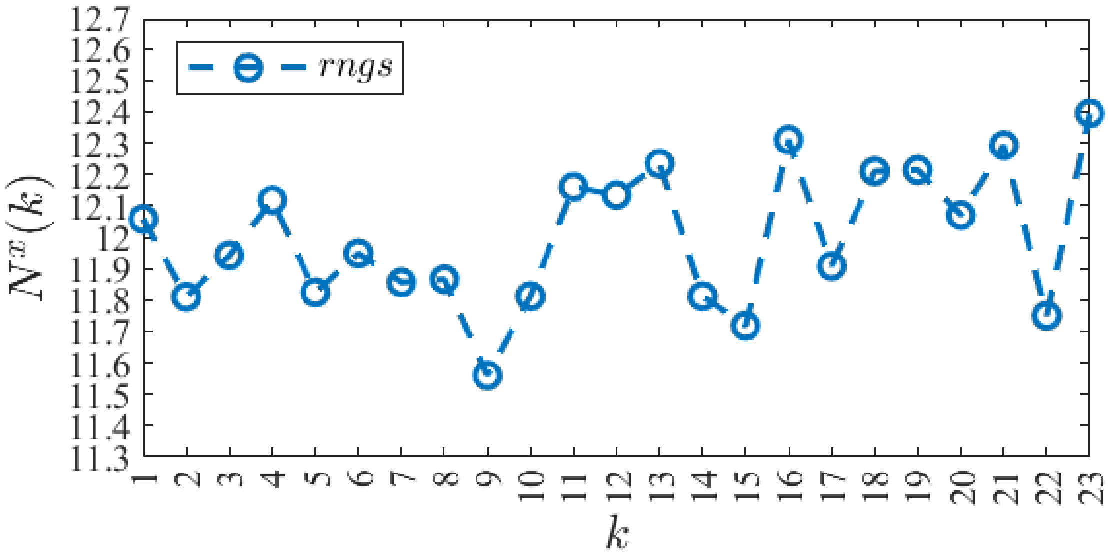

Figure 1 expected rank distributions with all identical preferences

The paper conducts two robustness checks for the result. First, the above simulations pertain to a very limited case where all applicants have the same preference. Figure 2 relaxes this assumption, in which applicants' preferences are completely random. There may or may not be any identical preferences among applicants. When the random number generator varies with each simulation, the expected ranks assigned for early and late applicants are not systematically different. KM remains procedurally fair.

Second, I relax the assumption that all hospitals have one quota. In the 2021 residency matching market, there were 23 hospitals and 233 applicants. Assuming all 233 applicants have identical preferences, figure 3 shows the resulting expected rank distribution in a large market. Each hospital's quota matches the 2021 data and is ordered identically as in the NSW residency matching system (Paya, private communication). Figure 3 illustrates very few variations in the ranks assigned for lower and higher rows on average. The result that KM is procedurally fair remains valid in a large market.

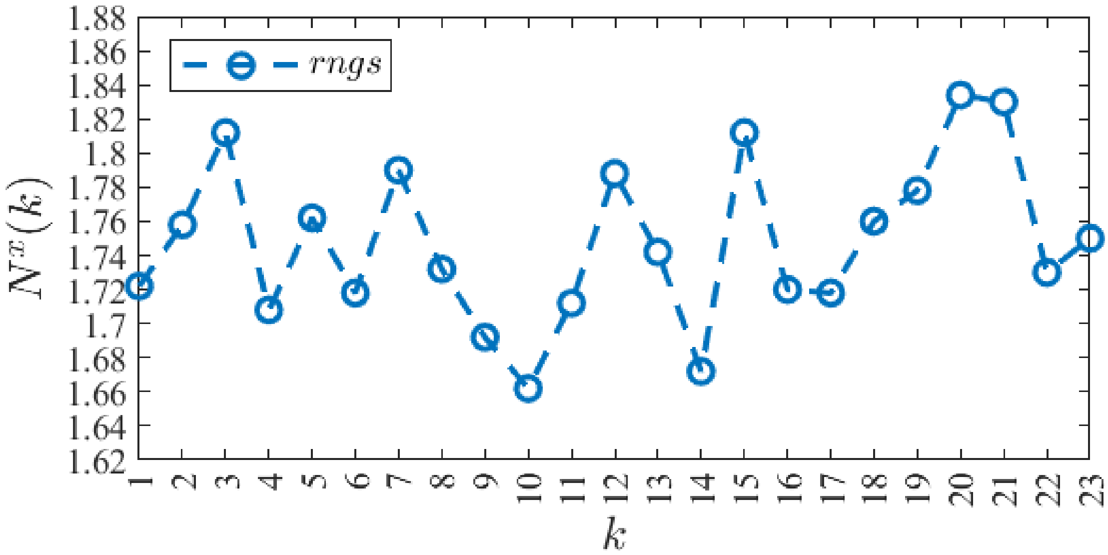

Figure 2 The expected rank distribution with random preferences

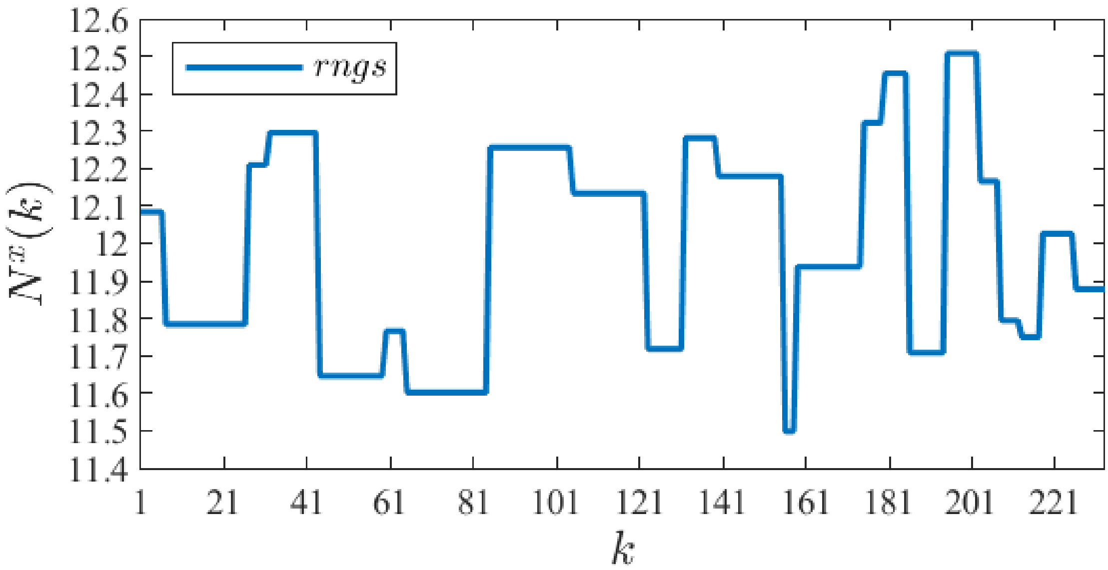

Figure 3 The expected rank distribution in a large market

Proposition 4 KM is not strategy-proof.

4.4.1 Definition of strategy-proofness

We say an assignment is strategy-proof if it is a dominant strategy for all applicants to report their true preferences $p_{a_i h_{jk}}$. That is, an assignment is strategy-proof if

$$(4.4.1)\ \forall a_i \in A, \forall h_{jk} \in H, \forall p_{a_i h_{jk}}, p'_{a_i h_{jk}} \in \Theta, \forall x, x' \in \Omega,$$

$$\sum_{h_{jk}} x_{a_i h_{jk}} p_{a_i h_{jk}} \leq \sum_{h_{jk}} x'_{a_i h_{jk}} p'_{a_i h_{jk}}$$

We say a mechanism is strategy-proof if and only if its assignment is strategy-proof. Roth (2008) stressed that strategy-proofness eradicates the costs of gathering information on others' preferences. Insurance for truth-telling levels the playing field by eliminating any advantages towards sophisticated agents (Basteck & Mantovani, 2018). Applicants' strategizing behaviours could incur great efficiency loss, which is especially important for KM (Rosenfeld & Hassidim, 2020). KM's strength of being rank efficient will be undermined without insurance against applicants' gaming behaviours.

4.4.2 Proof of strategy-proofness

KM is not strategy-proof. Featherstone (2020) proved that "no ordinal assignment is both rank efficient and strategy-proof" (proposition 10, p. 28). In our example 3.4, if $A1$ reports her preferences as $\{1; 1; 3; 3; 2\}$ (her true preferences are $\{1; 1; 2; 2; 3\}$), then the assignment becomes (the new assignment covered in green; assignment under truthful reporting in squares)

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	1	1	2	2	3
$A2$	1	1	2	2	3
$A3$	1	1	2	2	3
$A4$	2	2	3	3	1
$A5$	2	2	3	3	1

$A1$ improves from her second preference to her first preference after gaming.

4.4.3 A gaming strategy – change cut-offs

4.4.3.1 Formulating the gaming strategy

In the example above, $H4$ is $A1$, $A2$ and $A3$'s true second preference. KM incurs the same cost by assigning any of the three to $H4$. $A1$ games by listing $H4$ as her last preference, so $A1$ becomes costlier for $H4$ than $A2$ and $A3$. To minimise the sum of preferences assigned, KM would not consider assigning $A1$ in $H4$. The example hints that a strategizing applicant needs to be a cheap enough applicant for the one position she desires while a costly enough applicant for all other positions.

Based on the example's intuition, the thesis proposes a profitable gaming strategy – change cut-offs. The gaming applicant lists her desirable position as the first preference, making sure she is considered for that position. Then she states her preferences by just falling out of all other positions' cut-offs for cheap enough assignments. To formalise the gaming strategy, how to find the cut-offs for the remaining positions? And how to define just falling out of the cut-offs?

To find the cut-offs, the strategy uses Flood (1956)'s method. Suppose the gaming applicant a_g's desirable position is h_g. Now remove a_g and h_g out of the market, remaining $(N-1)$ applicants and positions. Flood (1956) constructed three steps[21] that locate all the possible assignments achieving the minimised preferences' sum in the remaining market. A position's cut-off is the largest preference number among all zeros in its column. By way of illustration, suppose the true preference matrix is

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	1	2	3	4	5
$A2$	1	2	3	4	5
$A3$	4	5	2	3	1
$A4$	1	2	4	5	3
$A5$	2	3	4	5	1

If $A1$ is the gaming applicant, then the remaining market will be

	$H2$	$H3$	$H4$	$H5$
$A2$	2	3	4	5
$A3$	5	2	3	1
$A4$	2	4	5	3
$A5$	3	4	5	1

[21] See the detailed steps in the appendix

The reduced market's preference matrix[22] for possible assignments is

	$H2$	$H3$	$H4$	$H5$
$A2$	0	0	0	4
$A3$	4	0	0	1
$A4$	0	1	1	2
$A5$	1	1	1	0

The cut-offs for positions $H2$ to $H5$ are $\{2; 3; 4; 1\}$. Any applicant who lists $H2$ as their second preference or better is considered for $H2$. Any applicant who lists $H3$ as their third preference or better is considered for $H3$.

To guarantee falling out of the cut-offs, the gaming applicant states all other positions as her kth preference where k is larger than the cut-offs by at least 1. In this example, $A1$ would report her preference as $\{1; 3; 4; 5; 2\}$. Readers can check that $A1$ ends up assigned to her first preference, which is better than her third preference under truth-telling. I call this gaming strategy changing cut-offs – list the desirable position as the first preference; list remaining positions behind the cut-offs by at least one.

4.4.3.2 Relations between change cut-offs and profitable deviations

Lemma 1 Changing cut-offs is a sufficient condition for profitable deviation.

Proof Define the minimum sum of preferences in the remaining market as R, given all others are truthful. The gaming applicant is a_g and her desirable position is h_g. If $\mu^x(a_g) = h_g$, the preference sum for the market will be

$$(4.4.2)\ R + p_{a_g h_g}$$

Careful readers may question would R change after introducing a_g back in the market. Applicant a_g is more costly than any applicant within the cut-offs for the remaining positions, so R cannot decrease. Applicant a_g is not assigned to any of the remaining positions. So if R increases, then KM's sum of assigned preferences will not be minimised. Therefore, R must remain unmoved.

[22] See the intermediate steps in the appendix

If a_g doesn't get her desirable position h_g, it may generate a series of assignment changes. Our focuses are only on who is assigned h_g and where is a_g assigned. Suppose$\mu^y(a_s) = h_g$ and $\mu^y(a_g) = h_f$. Suppose $\mu^x(a_f) = h_f$. a_f and a_s could be the same applicant, when a_g and a_f swap their assignments. When a_g causes a chain of assignment changes for remaining positions, a_f and a_s are different applicants. The preference sum for the market is

$$\text{(4.4.3) } R' - p_{a_f h_f} + p_{a_g h_f} + p_{a_s h_g}$$

Since $p_{a_f h_f}$ is within the cut-off for position h_f, $p_{a_g h_f} - p_{a_f h_f} \geq 1$. $p_{a_g h_g} = 1, so\ (p_{a_s h_g} - p_{a_g h_g}) \geq 0$. When a_g is introduced back into the market, all remaining positions' cut-offs cannot be lower since the new applicant is more costly than all applicants within the current cut-offs. So $(R' - R) \geq 0$.

Therefore,

$$\text{(4.4.4) } R' - p_{a_f h_f} + p_{a_g h_f} + p_{a_s h_g} - \left(R + p_{a_g h_g}\right) =$$

$$\text{(4.4.5) } (R' - R) + (p_{a_g h_f} - p_{a_f h_f}) + (p_{a_s h_g} - p_{a_g h_g}) \geq 1$$

Assigning a_g to h_g achieves a smaller sum of preferences assigned than any other assignments, so a_g's deviation must be profitable.

Lemma 2 Changing cut-offs is not a necessary condition for profitable deviation.

Proof Even when an applicant can profitably deviate from truth-telling, changing cut-offs may not be feasible. We showed that in example 3.4,

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	1	1	2	2	3
$A2$	1	1	2	2	3
$A3$	1	1	2	2	3
$A4$	2	2	3	3	1
$A5$	2	2	3	3	1

$A1$ gets a better assignment after misrepresenting her preference (move from the red to the green cell). One can check that the cut-offs for $H2$ to $H5$ are $\{2; 3; 3; 1\}$. However, $A1$ cannot list $H3$ and $H4$ as her fourth choice because there are only three hospitals in the market. The gaming strategy is impossible to implement for such cut-offs.

Then how does $A1$ get a better assignment by misstating her preference? The trick is that a zero is a possible candidate for an assignment, while a starred zero is an assignment. To avoid an unwanted assignment, one doesn't have to eliminate zeros in the remaining positions. Preventing a zero from becoming a starred zero in all remaining positions is enough. Changing cut-offs is a radical method of destroying zeros (candidacy) for all positions other than the desirable one. In example 3.4, $A1$ is already a candidate for $H3$ before $A3$ is assigned. To assign $A3$, $A1$'s candidacy is realized and thus ends up in her second preference. When $A1$ lies that $H3$ is her third choice, $A1$ falls in the cut-off for $H3$ only after $A3$ is assigned. Although $A1$ falls in the cut-off for position $H3$, the candidacy is never realized due to the timing of the transfer.

4.4.3.3 Implication of the changing cut-offs on asymptotic incentive properties

The existence of such a gaming strategy suggests that KM is asymptotically not strategy-proof. The feasibility of changing cut-offs grows as the market size grows, so applicants face larger room for profitable deviation in larger markets. The overarching assumption is that preferences are more heterogeneous when the market size increments. For heterogeneous preferences, different positions will be prioritized by different applicants and thus have low cut-offs. Accordingly, it is more likely that every position's cut-off can be added by at least one, compared to a high cut-off.

4.4.4 Nash Equilibrium of KM

Since KM is proven not strategy-proof, the thesis investigates the Nash equilibrium of KM. Changing cut-offs is a best response when it is feasible. A natural question to ask is – if every applicant adopts the strategy of changing cut-offs, is it a Nash equilibrium? Unfortunately, no. Even if this gaming strategy always exists when all other applicants are truthful, changing cut-offs converge to be infeasible. When an applicant changes her preference, she necessarily lists some unwanted positions more behind than her true preferences to fall out of the cut-offs. By doing that, those positions' cut-offs are pushed higher than before. For another applicant disliking the same positions, this applicant must list the positions further behind to fall out of the cut-offs. The process continues until there is no more room for listing a position as less preferred to fall out of the cut-offs.

Regarding finding the Nash equilibrium, an upfront challenge is that there is an exploding number of strategies available for applicants. Even restricted to a $3 * 3$ market (three hospitals and three applicants), there are 6^3 possible combinations of reported preferences. The paper outlines two preliminary results and discusses further research directions in the conclusion.

Lemma 1 There are multiple Nash equilibria for KM.

Suppose the preference matrix is

	$H1$	$H2$	$H3$
$A1$	1	2	3
$A2$	1	3	2
$A3$	1	3	2

In such a simple market, there are nine Nash equilibria (see the appendix for details). The non-uniqueness of Nash equilibrium is because a class of reported preferences are trivially the same in terms of getting one's best possible position. The relevant positions for one's final assignment are the columns containing zeros. If preferences are different in a way that none of the zeros is impacted, then the best available position achieved is identical. For example, if one is considered for assignments to her first three preferences, then she is free to swap her preferences after her third position.

Different types of applicants have different relevant positions for their final assignments. There are two types of applicants – one never involved in a transfer and the other at least involved in one. For the first type, an applicant is assigned to her lowest uncovered column, which is her unique, relevant position. For the second type, an applicant is kicked out of a starred zero column to a primed zero column, which are her relevant positions. However, the relevant positions are jointly determined by one's own preference and all others' preferences for those positions. The best response could switch from one type of applicant to another, confusing the applicant's best response by type. The Nash equilibrium of KM is still an open question.

Lemma 2 In a Nash equilibrium, there could be efficiency loss compared to truthful reporting.

For example, the truthful assignment (the green-covered cell denotes assignment) is

	$H1$	$H2$	$H3$	$H4$
$A1$	1	2	3	4
$A2$	1	3	2	4
$A3$	1	3	2	4
$A4$	1	2	3	4

One of its Nash equilibria is

	$H1$	$H2$	$H3$	$H4$
$A1$	1	2	3	4
$A2$	2	1	3	4
$A3$	1	3	2	4
$A4$	1	2	3	4

The sum of preferences in terms of true preferences moves from 9 to 10. In this Nash equilibrium, some applicants (in this example, $A1$) are worse off than under truth-telling, whose efficiency loss outweighs the efficiency gain from better-off applicants (in this example, $A2$).

5 Discussion of RP and KM

This section compares RP, the most used medical residency matching mechanism worldwide, with KM on their efficiency, fairness, and incentive properties. I make two choices to simplify the comparisons. First, all preferences are strict, avoiding the complications of preferences with ties (Irving, Manlove & Scott, 2008; Erdil & Ergin, 2008). Second, the discussions divide theoretical and practical comparisons. This is because KM and RP are applied in theoretically different markets -- KM in simple one-sided markets and RP in complex two-sided markets. RP could be worse than KM on certain properties due to market complementarities instead of its inherent deficiency. Comparing both mechanisms in simple markets constitute theoretical discussions. Comparing both mechanisms in action, taking market complementarities and participants' behaviours into account, constitute practical discussions.

KM is theoretically and practically more efficient than RP (section 5.1). Section 5.2 KM is theoretically outcome unfair while RP is theoretically outcome fair, but both outcome-unfair in practice (section 5.2). KM and RP are both theoretically and practically procedurally fair (section 5.3). KM and RP shield against some degree of gaming theoretically, while both are subject to strategizing practically (section 5.4). These results should still be carefully treated. RP considers two-sided preferences while KM only applies to one-sided markets, so definitions for efficiency, fairness and strategy-proofness are not directly comparable. The answer to which is better boils down to the interactions among desirable properties, transaction costs and market complementarities.

5.1 Which is more efficient

5.1.1 Theoretical Comparison

RP in simple markets is the applicant-proposing deferred-acceptance mechanism (Roth & Peranson, 1999). The deferred-acceptance mechanism is Pareto efficient (Roth, 2008). The thesis proved that KM is rank efficient, which is strong than Pareto efficient. Is the deferred-acceptance mechanism also rank efficient? That is, does the deferred-acceptance mechanism minimise the sum of applicants' preferences for assigned hospitals? Featherstone (2020) proved that rank efficiency could only be achieved by a certain class of linear-programming mechanisms. Therefore, KM is better than RP in efficiency in simple many-to-one markets.

5.1.2 Practical Comparison

Roth (2008) stated that a matching is stable if and only if it is in the core. If a matching is in the core, then it is Pareto efficient (Roth, 2008). So the efficiency of RP in complex markets depends on its stability. In a market with complementarities between positions and couples, stable matching may not exist (Roth & Peranson, 1999). Klaus and Klijn (2005) defined one genre of couples' preferences as weakly responsive – an improvement for one is an improvement for the couple. They proved that if all positions are acceptable, stable matching must exist in a market containing only weakly responsive couples (Klaus & Klijn, 2005). However, Kojima, Pathak and Roth (2013) argued that weakly responsive preferences are impractical, especially for couples seeking matchings in geographically close hospitals. Even if couples' preferences are responsive, a stable matching may not exist if not all positions are acceptable and compatible (McDermid & Manlove, 2010). If a stable matching doesn't exist, then RP will not be Pareto efficient in practice.

However, RP always outputs a stable matching in 1987 and 1993-1996 NRMP matchings, indicating that stable matchings are likely to exist (Roth & Peranson, 1999). Kojima, Pathak and Roth (2013) discovered that "a stable matching exists when there are relatively few couples and preference lists are sufficiently short to market size" (p. 1585). Does this asymptotic existence of a stable outcome mean RP's outcome is stable? Sadly, RP may fail to find a stable matching even when one exists because of the limitations in computational complexity (Biro, Manlove & Mcbride, 2014). RP is an incomplete method, only attempting to find a stable matching without insurance to find one (Drummond, Perrault & Bacchus, 2015). Consequently, RP could be unstable and inefficient in practice.

KM always minimises the sum of applicants' preferences assigned, so it is rank efficient in practice. Practically, KM achieves better results than RP in efficiency.

5.2 Which is more outcome fair

5.2.1 Theoretical Comparison

In two-sided matching markets, outcome fairness is referred to as no justified envy. Roth (1982) proved that the deferred-acceptance mechanism produces no blocking pairs. That is, all applicants only envy those who are more preferred by hospitals, resulting in all justifiable envies. Section 4.2 showed that KM is outcome unfair. Therefore, KM is worse than RP on outcome fairness theoretically.

5.2.2 Practical Comparison

Whether justified envy exists in RP's outcome is equivalent to whether RP's outcome is stable. As we discussed in 5.1.2, stable matchings asymptotically exist in complex markets, but RP doesn't guarantee finding that stable matching (Drummond, Perrault & Bacchus, 2015). So RP is outcome unfair in practice.

The thesis proved that KM is outcome unfair, but not for all applicants with identical preferences. Like in example 3.4, the early applicant $A4$ gets a better assignment than $A5$. Could KM in NSW be outcome-fair for most applicants? In a KM transfer, early applicants are reallocated from lower columns to higher columns. Accordingly, KM is unfair to early applicants with a certain types of preferences – more-preferred positions in the lower columns and less-preferred positions in the higher columns. These applicants are worse-off in a transfer. Now the question becomes, are such applicants prevalent in NSW?

Figure 4 shows all hospitals' quotas according to the column order and the number of applicants listing them as their first choice in NSW (Paya, private conversation). The grey line depicts the differences between quotas and first preferences for all hospitals. The lower the grey line is, the more popular a hospital is relative to its capacity. As illustrated, more popular hospitals are in the higher columns of the preference matrix. Many early applicants' more-preferred positions, which are the popular positions, reside in the higher columns of the preference matrix. For many applicants in NSW residency matching markets, their preference type is opposite to the type of disadvantaged preferences, so disadvantages for early applicants are seemingly small. It is therefore concluded that albeit KM and RP are both outcome-unfair practically, KM's room for unfairness is small in NSW.

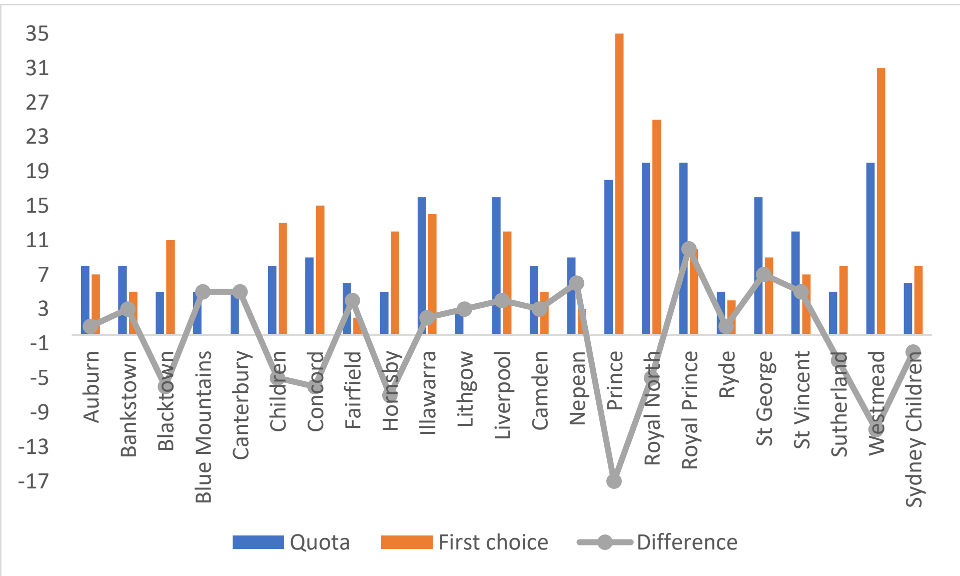

Figure 4 NSW participating hospitals' popularity

5.3 Which is more procedurally fair

5.3.1 Theoretical Comparison

For RP, I define procedural fairness as – the order of admission into the algorithm is independent of final matchings. That is, an applicant only envies another because she is less preferred by a hospital, not because she is processed later in the algorithm. In the deferred-acceptance mechanism, a hospital accepts or rejects a proposal solely based on its preferences, regardless of the proposal's order. So RP is theoretically procedurally fair, the same as KM (the proof in section 4.3).

5.3.2 Practical Comparison

Roth and Peranson (1999) simulated the effects of different admission orderings of applicants on the matching outcomes. Most of the matchings are identical, using different processing sequences (Roth and Peranson, 1999). In the most affected two out of eight cases, approximately 0.05% of applicants end up in different hospitals (Roth and Peranson, 1999). Therefore, RP is procedurally fair practically.

Although KM is proven to be procedurally fair, applicants might think otherwise. Bolton, Brandts and Ockenfels (2005) stated that "only observable and verifiable unbiased procedures would seem credibly fair" (p. 1073-1074). In NSW, applicants are completely uninformed about the matching mechanism. Applicants may perceive the secrecy as a disguise for unfair procedures. Hence, even though RP and KM remain both procedurally fair practically, KM's procedural fairness may be unreliable for NSW applicants.

5.4 Which is more strategy-proof

5.4.1 Theoretical Comparison

Roth (1984) proved that no stable matching exists that is strategy-proof for all applicants. Luckily, the number of profitable gaming strategies is limited. Hospitals are enticed to truncate their quotas and pre-arrange matches with applicants (Sonmez, 1997; Sonmez, 1999). In a symmetric information environment, truncation stochastically dominates all other gaming strategies (Roth & Rothblum, 1999). Ehlers (2004) extended the assumption of symmetric information and concluded that all other gaming strategies are still stochastically dominated by a broadened truncation strategy. Ehlers (2008) further proved that truncation by the non-proposing side is the best and only profitable gaming strategy in the deferred-acceptance mechanism. However, the limited type of profitable gaming strategy doesn't shield against the large degree of gaming (Coles & Shorrer, 2014). As markets grow larger, the optimal degree of truncation approaches 100% if beliefs about others' preferences are uniformly distributed (Coles & Shorrer, 2014). Section 4.4 argued that KM is not strategy-proof. Therefore, neither RP nor KM is strategy-proof.

What if we take a step back and only require some protection against strategic behaviours, not devoid of all strategic behaviours? Roth and Rothblum (1999) found that truth-telling is an approximate Nash Bayesian equilibrium in one-to-one markets, assuming only the prior distribution of preferences is common knowledge. Kojima and Pathak (2009) extended the conclusion asymptotically to many-to-one markets under complete information – truth-telling is an approximate Nash equilibrium. When the market is large enough, the manipulated rejection chains by truncation will be absorbed somewhere by a vacant position (Kojima & Pathak, 2009). Experimental evidence also suggested that most agents involved are truthful as the market size increases (Agarwal, 2015). Castillo and Dianat (2016) restricted their experiments to no unmatched agents and complete information. They found that most agents are truthful for fear of over-truncating and losing a potential match (Castillo & Dianat, 2016).

As for KM, Featherstone (2020) proved that under a symmetric-belief environment and no outside options, truth-telling is a best response in rank-value mechanisms. Rank-value mechanisms are linear programming mechanisms with the same social value placed on all applicants' kth preference (Featherstone, 2020). KM values all applicants' kth preference k, so KM is a rank-value mechanism. Correspondingly, applicants' best responses are truth-telling under the above restrictions. Whether truth-telling can constitute a Nash equilibrium of KM remains an open question. But KM and RP both bear some protection against strategic manipulations.

An additional strength of KM is its insulation from policymakers' gaming. Researchers proposed the notion of policymakers' strategy-proofness as no incentives to change the matching outcomes (Hakimov & Raghavan, 2022). To achieve strategy-proofness for policymakers, they stressed the importance of transparency and verifiability of an algorithm (Hakimov & Raghavan, 2022). Although KM is not transparent, it successfully eliminates the incentive for policymakers to tinker (Featherstone, 2020). Featherstone (2020) observed that crucial criteria for policymakers to assess the matching outcome is the number of applicants assigned to their first or top three preferences. Policymakers tinker with the matching outcome to assign more applicants to their top preferences, sabotaging the properties of an algorithm (Featherstone, 2020). KM already assigns the maximum number of applicants as close to their top preferences as possible, so it is immune to any such gaming from policymakers.

We conclude that although KM and RP are not strategy-proof in simple markets, they both have some protection against gaming under certain restrictive assumptions. Other than shielding participants' strategic manipulations, KM preserves the integrity of the outcome from policymakers.

5.4.2 Practical Comparison

In complex two-sided markets, there is no stable mechanism that is strategy-proof for every couple, even assuming all couples have responsive preferences (Klaus & Klijn, 2005). Couples can pretend to be singles as a way of profitable deviation from truth-telling (Klaus, Klijn & Masso, 2007). Therefore, an identical conclusion holds theoretically and practically – KM and RP are not strategy-proof.

Then would both mechanisms shield against some gaming behaviours in practice, as in the theoretical discussion? Roth and Perason (1999) proposed that RP's manipulation space is limited as the market size increases. Using computer simulations, they found that 0.01% of hospitals have a successful chance of manipulation. Immorlica and Mahdian (2005) also concluded that RP has a small gaming space asymptotically, generalizing preferences from uniform distributions to random preferences. Given the positive result in large markets, RP seems to have protections against strategic behaviours.

However, whether an agent engages in strategic manipulations or not depends on her beliefs, not on definite profitability from strategizing (Featherstone, 2020). From survey data in NRMP2015, 24% of senior and 47% of independent applicants admitted that they didn't report their preferences truthfully (Castillo & Dianat, 2016). Many applicants admitted that they modified their preferences based on the likelihood of successful matchings (Echenique, Wilson & Yariv, 2016). There is also experimental evidence suggesting that applicants skipped potential matchings because of a perceived low chance of getting into their top preferences, even after understanding how RP works (Echenique, Wilson & Yariv, 2016). Applicants' gaming is also evident from NRMP matching outcomes. Agarwal (2015) argued that applicants used similar standards to judge hospitals, so applicants' preferences should be fairly homogeneous. From NRMP2015, close to half of the applicants are matched to their highest preferences, and more than 85% are matched to their top four preferences (Agarwal, 2015; Echenique, Wilson & Yariv, 2016). These field data indicate that applicants' reported preferences are diverse, contradicting the assumption of truthful homogeneous preferences. Castillo and Dianat (2016) argued that manipulation is not limited to preference reporting. Agents are already strategizing in applying for interviews, and they only apply for attainable hospitals based on the probability of successful matchings (Castillo & Dianat, 2016).

As for KM, there is no behavioural research investigating applicants' strategic behaviours. A natural thought is that applicants avoid any gaming because they don't know what mechanism is in place in NSW, let alone how to game it. There is no overlapping generation of applicants[23], so applicants cannot use past matching data to predict others' preferences (Rosenfeld & Hassidim, 2020). However, empirical data from NSW2021 may suggest non-truthful reporting. In 2021 NSW residency matching, 69.1% are assigned their top preference, and 96.57% are assigned their top three preferences (Paya, private communication). Assuming applicants' preferences are homogeneous due to similar criteria in assessing hospitals, the data serves as evidence for gaming behaviours. Rosenfeld and Hassidim (2020) pointed out that for an obscure and complicated mechanism, applicants tend to engage in conjecture gaming, which could describe applicants' behaviours in NSW. Without any further examination, if applicants engage in gaming in NSW is uncertain. We

[23] An applicant can only apply to NSW Residency Allocation once.

conclude that applicants may be prone to strategic manipulations in practice both for KM and RP.

6 Concluding Remarks

NSW physiotherapy residency matching uses the Kuhn-Munkres algorithm, a new mechanism for worldwide matching markets in practice and in literature. A major contribution of the thesis is finding the mechanism and investigating its properties, hopefully inspiring future discussions. Investigating the properties of KM is important due to its potentially wide application. KM is computationally efficient in a one-sided matching market, so policymakers can significantly reduce transaction costs for two-sided markets after converting the market to one-sided and applying KM. The question is, does KM achieve satisfactory properties in a one-sided market? KM's assignment is rank efficient, which is even stronger than Pareto efficiency. However, KM is lacking in fairness and strategy-proofness. Applying early could become a disadvantage to applicants with a specific types of preferences. Applicants are not incentivized to tell their true preferences. A sufficient but not necessary gaming strategy for profitable deviation is found, and the feasibility of the strategy is increasing as the market size grows. KM doesn't tick all the boxes for efficiency, fairness and truthful revelations, so is it not useful?

The thesis then takes a step back. KM is worth considering if it can achieve better properties than existing mechanisms. Now the question is, is KM better than RP? Theoretically, KM is more efficient but less fair than RP, while both shield against participants' gaming behaviours to a certain degree. Practically, KM is more efficient than RP, while both are likely to be fair but subject to applicants' gaming. A bonus for using KM is KM's insulation against policymakers' attempts to modify the assignments. We still cannot conclude that KM should replace RP. RP is applicable to markets with complementarities between positions and preferences, while KM is not. Eventually, policymakers need to balance market complementarities and the need for cost reductions.

The thesis serves as an opening for research around KM. One direction for future researchers could be to formalize fairness. Section 5.2.2 points out that early applicants with preferences prioritizing the positions in the lower columns in the preference matrix could be disadvantaged. The precise cut-off between a low and high column and the cut-off between early and late applicants are both undecided. The transfers depend on the hospital quotas, the preference matrix's row and column order and all applicants' preferences. How the cut-offs are determined given interactions among all those factors is an open question.

Another future research topic is the Nash equilibrium of KM. In the markets I tried, all Nash equilibria generate identical payoffs for each applicant. Although the payoffs could be different from truth-telling assignments, they are the same across all Nash equilibria. If this result is true, then applicants will coordinate on a deterministic set of welfare outcomes. I found no counter-examples, but if this conjecture is correct requires further research, especially when the number of strategies compounds in large markets.

Lastly, the compatibility of KM with labour market trends should be examined. As stated in the introduction, physiotherapy is an occupation facing a severe labour market shortage. If all hospitals' quotas grow sufficiently large for all applicants, then all applicants will be assigned to their top preferences. If all hospitals' quotas remain relatively steady, but the number of applicants grows, then the unfairness and incentive issues are expected to exacerbate. KM will necessarily involve more applicants in transfers, and the feasibility of changing cut-offs will grow. Besides the market size changes, how KM accommodates couples' preferences is an open question. Two adjacent rows are not guaranteed to be in the hospital. If couples' preferences just collapse in one row and deduct two quotas from a hospital for identical assignments, it will create many problems. If the couples' preferences are the type that will be worse off in transfers, then it is better to apply individually because one applicant could involve in one less transfer. If couples' preferences are the type that will be better off in transfers, then it is unfair to other applicants because the separate application could leave one position for other applicants with identical preferences. KM is now applied in a one-sided market without any complementarities, so its adaptability in complex markets is worth further investigation.

7 References

Agarwal, N. (2015). An Empirical Model of the Medical Match. The American Economic Review, 105(7), 1939–1978. https://doi.org/10.1257/aer.20131006

APA | Workforce report shows physiotherapy in demand. (2021, October 1). Austrian.physio; Australian Physiotherapy Association. https://australian.physio/inmotion/workforce-report-shows-physiotherapy-demand

Association of Psychology Postdoctoral and Internship Centers. (2022). APPIC - Association of Psychology Postdoctoral and Internship Centers. https://www.appic.org/

Basteck, C. & Mantovani, M. (2018). Cognitive ability and games of school choice. Games and Economic Behavior, 109, 156–183. https://doi.org/10.1016/j.geb.2017.12.011

Bevilacqua Research Corporation. (2022, October 28). The algorithm work*shop*. https://brc2.com/the-algorithm-workshop/

Biró, P., Irving, R. & Schlotter, I. (2011). Stable matching with couples: An empirical study. The ACM Journal of Experimental Algorithmics, 16, 1.1–1.27. https://doi.org/10.1145/1963190.1970372

Biró, P., Manlove, D. F. & McBride, I. (2014). The Hospitals / Residents Problem with Couples: Complexity and Integer Programming Models. Experimental Algorithms, 8504, 10–21. https://doi.org/10.1007/978-3-319-07959-2_2

Boehmer, N., Heeger, K. & Niedermeier, R. (2022). Deepening the (Parameterized) Complexity Analysis of Incremental Stable Matching Problems. *Leibniz International Proceedings in Informatics*, *21*, 21:1-21:15. https://drops.dagstuhl.de/opus/volltexte/2022/16819/pdf/LIPIcs-MFCS-2022-21.pdf

Bolton, G. E., Brandts, J. & Ockenfels, A. (2005). Fair Procedures: Evidence from Games Involving Lotteries. The Economic Journal (London), 115(506), 1054–1076. https://doi.org/10.1111/j.1468-0297.2005.01032.x

Bronfman, S., Hassidim, A., Afek, A., Romm, A., Shreberk, R., Hassidim, A. & Massler, A. (2015). Assigning Israeli medical graduates to internships. Israel Journal of Health Policy Research, 4(1), 6–6. https://doi.org/10.1186/2045-4015-4-6

Bronfman, S., Alon, N., Hassidim, A. & Romm, A. (2018). Redesigning the Israeli Medical Internship Match. ACM Transactions on Economics and Computation, 6(3-4), 1–18. https://doi.org/10.1145/3274646

Canadian Resident Matching Service. (2022, February 16). R-1 main residency match. CaRMS. https://www.carms.ca/match/r-1-main-residencymatch/?gclid=Cj0KCQjwqoibBhDUARIsAH2OpWhRPFovUoklHnHh9HTatP6cE49gllLQ2qUEb0KEH3jILIHb3u1y1bYaAuUfEALw_wcB

Castillo, M. & Dianat, A. (2016). Truncation strategies in two-sided matching markets: Theory and experiment. Games and Economic Behavior, 98, 180–196. https://doi.org/10.1016/j.geb.2016.06.006

Coles, P. & Shorrer, R. (2014). Optimal truncation in matching markets. Games and Economic Behavior, 87, 591–615. https://doi.org/10.1016/j.geb.2014.01.005

De Clercq, S., Schockaert, S., De Cock, M. & Nowe, A. (2016). Solving stable matching problems using answer set programming. Theory and Practice of Logic Programming, 16(3), 247–268. https://doi.org/10.1017/S147106841600003X

Drummond, J., Perrault, A. & Bacchus, F. (2015). SAT Is an Effective and Complete Method for Solving Stable Matching Problems with Couples. IJCAI. https://www.ijcai.org/Abstract/15/079

Echenique, F., Wilson, A. J. & Yariv, L. (2016). Clearinghouses for two-sided matching: An experimental study. Quantitative Economics, 7(2), 449–482. https://doi.org/10.3982/QE496

Edmonds, J. & Karp, R. M. (2003). Theoretical Improvements in Algorithmic Efficiency for Network Flow Problems. In Combinatorial Optimization — Eureka, You Shrink (pp. 31–33). Springer Berlin Heidelberg. https://doi.org/10.1007/3-540-36478-1_4

Ehlers, L. (2004). In search of advice for participants in matching markets which use the deferred-acceptance algorithm. Games and Economic Behavior, 48(2), 249–270. https://doi.org/10.1016/j.geb.2003.09.007

Ehlers, L. (2008). Truncation Strategies in Matching Markets. Mathematics of Operations Research, 33(2), 327–335. https://doi.org/10.1287/moor.1070.0294

Erdil, A. & Ergin, H. (2008). What's the matter with tie-breaking?: Improving efficiency in school choice. The American Economic Review, 98(3), 669–689. https://doi.org/10.1257/aer.98.3.669

Featherstone, C. R. (2020). Rank Efficiency: Modeling a Common Policymaker Objective. https://doi.org/https://clayton-featherstone.github.io/

Flood, M. M. (1956). The Traveling-Salesman Problem. Operations Research, 4(1), 61–75. https://doi.org/10.1287/opre.4.1.61

Gusfield, D. & Irving, R. W. (1989). The stable marriage problem : structure and algorithms. MIT Press.

Hakimov, R. & Raghavan, M. Improving Transparency in School Admissions: Theory and Experiment (March 31, 2022). Available at SSRN: https://ssrn.com/abstract=3572020 or http://dx.doi.org/10.2139/ssrn.3572020

IBISWorld - Industry Market Research, Reports, and Statistics. (2022, March 17). Www.ibisworld.com; IBISWorld. https://www.ibisworld.com/au/industry/physiotherapy-services/619/

Immorlica, N. & Mahdian, M. (2005). Marriage, honesty, and stability. Proceedings of the Sixteenth Annual ACM-SIAM Symposium on Discrete Algorithms, 53–62.

Irving, R. (2011). Matching Practices for Entry-labor Markets – Scotland. matching in practice. http://www.matching-in-practice.eu/wp-content/uploads/2013/05/MiP_-Profile_No.3.pdf

Irving, R. W., Manlove, D. F. & Scott, S. (2008). The stable marriage problem with master preference lists. DISCRETE APPLIED MATHEMATICS, 156(15), 2959–2977. https://doi.org/10.1016/j.dam.2008.01.002

Japan Residency Matching Program. (2022). *JRMP*. https://www.jrmp.jp/

Kamada, Y. & Kojima, F. (2015). Efficient Matching under Distributional Constraints: Theory and Applications. The American Economic Review, 105(1), 67–99. https://doi.org/10.1257/aer.20101552

Klaus, B. & Klijn, F. (2005). Stable matchings and preferences of couples. Journal of Economic Theory, 121(1), 75–106. https://doi.org/10.1016/j.jet.2004.04.006

Klaus, B., Klijn, F. & Masso, J. (2007). Some things Couples always wanted to know about Stable Matchings (but were afraid to ask). Review of Economic Design, 11(3), 175–184. https://doi.org/10.1007/s10058-006-0017-9

Kojima, F. & Pathak, P. A. (2009). Incentives and stability in large two-sided matching markets. The American Economic Review, 99(3), 608–627. https://doi.org/10.1257/aer.99.3.608

Kojima, F., Pathak, P. A. & Roth, A. E. (2013). MATCHING WITH COUPLES: STABILITY AND INCENTIVES IN LARGE MARKETS. The Quarterly Journal of Economics, 128(4), 1585–1632. https://doi.org/10.1093/qje/qjt019

König, D. (1916). Über Graphen und ihre Anwendung auf Determinantentheorie und Mengenlehre. Mathematische annalen, 77(4), 453–465. https://doi.org/10.1007/BF01456961

Kuhn, H. W. (2005). The Hungarian method for the assignment problem. Naval Research Logistics, 52(1), 7–21. https://doi.org/10.1002/nav.20053

MacLean, M. T., Lysikowski, J. R., Rege, R. V., Sendelbach, D. M. & Mihalic, A. P. (2020). Optimizing Medical Student Clerkship Schedules Using a Novel

Application of the Hungarian Algorithm. Academic Medicine, 96(6), 864–868. https://doi.org/10.1097/ACM.0000000000003676

Manlove, D. F., McBride, I. & Trimble, J. (2017). “Almost-stable” matchings in the Hospitals / Residents problem with Couples. Constraints : an International Journal, 22(1), 50–72. https://doi.org/10.1007/s10601-016-9249-7

Manlove, D. F., O’Malley, G., Prosser, P. & Unsworth, C. (2007). A Constraint Programming Approach to the Hospitals / Residents Problem. Integration of AI and OR Techniques in Constraint Programming for Combinatorial Optimization Problems, 4510, 155–170. https://doi.org/10.1007/978-3-540-72397-4_12

McDermid, E. J. & Manlove, D. F. (2010). Keeping partners together: algorithmic results for the hospitals/residents problem with couples. Journal of Combinatorial Optimization, 19(3), 279–303. https://doi.org/10.1007/s10878-009-9257-2

McKinney, C. N., Niederle, M. & Roth, A. E. (2005). The Collapse of a Medical Labor Clearinghouse (and Why Such Failures Are Rare). The American Economic Review, 95(3), 878–889. https://doi.org/10.1257/0002828054201323

Mirzaeinia, A. & Hassanalian, M. (2019). Minimum-cost drone-nest matching through the Kuhn-Munkres algorithm in smart cities: Energy management and efficiency enhancement. Aerospace, 6(11), 125–. https://doi.org/10.3390/aerospace6110125

Munkres, J. (1957). Algorithms for the Assignment and Transportation Problems. Journal of the Society for Industrial and Applied Mathematics, 5(1), 32–38. https://doi.org/10.1137/0105003

National Resident Matching Program. (2022). *The Match*. NRMP. https://www.nrmp.org/

NP-hardness. (2022, July 15). Wikipedia, the free encyclopedia. https://en.wikipedia.org/wiki/NP-hardness

NSW Physiotherapy *Allocation Committee (2022). NSW Physiotherapy Allocation - Home. https://www.nswphysioallocation.com.au/*

Ronn, E. (1990). NP-complete stable matching problems. Journal of Algorithms, 11(2), 285–304. https://doi.org/10.1016/0196-6774(90)90007-2

Rosenfeld, A. & Hassidim, A. (2020). Too smart for their own good: Trading truthfulness for efficiency in the Israeli medical internship market. Judgment and Decision Making, 15(5), 727–740.

Roth, A. E. (1982). The Economics of Matching: Stability and Incentives. Mathematics of Operations Research, 7(4), 617–628. https://doi.org/10.1287/moor.7.4.617

Roth, A. E. (1984). The Evolution of the Labor Market for Medical Interns and Residents: A Case Study in Game Theory. The Journal of Political Economy, 92(6), 991–1016. https://doi.org/10.1086/261272

Roth, A. E. (1991). A Natural Experiment in the Organization of Entry-Level Labor Markets: Regional Markets for New Physicians and Surgeons in the United Kingdom. The American Economic Review, 81(3), 415–440.

Roth, A. E. (2008). What Have We Learned from Market Design. The Economic Journal (London), 118(527), 285–310. https://doi.org/10.1111/j.1468-0297.2007.02121.x

Roth, A. E. & Peranson, E. (1999). The Redesign of the Matching Market for American Physicians: Some Engineering Aspects of Economic Design. The American Economic Review, 89(4), 748–780. https://doi.org/10.1257/aer.89.4.748

Roth, A. E. & Rothblum, U. G. (1999). Truncation Strategies in Matching Markets-in Search of Advice for Participants. Econometrica, 67(1), 21–43. https://doi.org/10.1111/1468-0262.00002

Roth, A. E. & Shorrer, R. I. (2015). The redesign of the medical intern assignment mechanism in Israel. Israel Journal of Health Policy Research, 4(1), 11–11. https://doi.org/10.1186/s13584-015-0014-y

Roth, A. E. & Vande Vate, J. H. (1990). Random Paths to Stability in Two-Sided Matching. Econometrica, 58(6), 1475–1480. https://doi.org/10.2307/2938326

Sanseverino, E. R., Ngoc, T. N., Cardinale, M., Li Vigni, V., Musso, D., Romano, P. & Viola, F. (2015). Dynamic programming and Munkres algorithm for optimal photovoltaic arrays reconfiguration. Solar Energy, 122, 347–358. https://doi.org/10.1016/j.solener.2015.09.016

Sonmez, T. (1997). Manipulation via Capacities in Two-Sided Matching Markets. Journal of Economic Theory, 77(1), 197–204. https://doi.org/10.1006/jeth.1997.2316

Sonmez, T. (1999). Can Pre-arranged Matches Be Avoided in Two-Sided Matching Markets? Journal of Economic Theory, 86(1), 148–156. https://doi.org/10.1006/jeth.1999.2507

Trang, L. H., Viet, H. H., Van Hoai, T. & Hao, T. X. (2019). A Searching for Strongly Egalitarian and Sex-Equal Stable Matchings. Advances in Intelligent Systems and Computing, 935, 1100–1111. https://doi.org/10.1007/978-3-030-19063-7_87

UK Foundation Programme. (2022, October 31). https://foundationprogramme.nhs.uk/

8 Appendix

8.1 An example where an applicant is better off in a transfer

$A4$ moves from her second preference to her first preference in the last transfer. The initial preference matrix is

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	3	2	2	1	1
$A2$	2	3	3	1	1
$A3$	3	2	2	1	1
$A4$	2	3	3	1	1
$A5$	2	3	3	1	1

Below presents the process of KM.

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	2	1	1	0	0
$A2$	1	2	2	0	0
$A3$	2	1	1	0	0
$A4$	1	2	2	0	0
$A5$	1	2	2	0	0

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	2	1	1	0^*	0
$A2$	1	2	2	0	0^*
$A3$	2	1	1	0	0
$A4$	1	2	2	0	0
$A5$	1	2	2	0	0

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	1	0	0	0^*	0
$A2$	0	1	1	0	0^*
$A3$	1	0	0	0	0
$A4$	0	1	1	0	0
$A5$	0	1	1	0	0

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	1	$0'$	0	0^*	0
$A2$	0	1	1	0	0^*
$A3$	1	0	0	0	0
$A4$	0	1	1	0	0
$A5$	0	1	1	0	0

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	1	$0'$	0	0^*	0
$A2$	$0'$	1	1	0	0^*
$A3$	1	$0'$	0	0	0
$A4$	0	1	1	0	0
$A5$	0	1	1	0	0

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	1	0	0	0^*	0
$A2$	0	1	1	0	0^*
$A3$	1	0^*	0	0	0
$A4$	0	1	1	0	0
$A5$	0	1	1	0	0

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	1	0	$0'$	0^*	0
$A2$	0	1	1	0	0^*
$A3$	1	0^*	0	0	0
$A4$	0	1	1	0	0
$A5$	0	1	1	0	0

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	1	0	$0'$	0^*	0
$A2$	$0'$	1	1	0	0^*
$A3$	1	0^*	0	0	0
$A4$	0	1	1	0	0
$A5$	0	1	1	0	0

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	1	0	$0'$	0^*	0
$A2$	$0'$	1	1	0	0^*
$A3$	1	0^*	$0'$	0	0
$A4$	$0'$	1	1	0	0
$A5$	0	1	1	0	0

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	1	0	0	0^*	0
$A2$	0	1	1	0	0^*
$A3$	1	0^*	0	0	0
$A4$	0^*	1	1	0	0
$A5$	0	1	1	0	0

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	1	0	$0'$	0^*	0
$A2$	0	1	1	0	0^*
$A3$	1	0^*	0	0	0
$A4$	0^*	1	1	0	0
$A5$	0	1	1	0	0

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	1	0	$0'$	0^*	0
$A2$	0	1	1	$0'$	0^*
$A3$	1	0^*	0	0	0
$A4$	0^*	1	1	0	0
$A5$	0	1	1	0	0

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	1	0	$0'$	0^*	0
$A2$	0	1	1	$0'$	0^*
$A3$	1	0^*	$0'$	0	0
$A4$	0^*	1	1	0	0
$A5$	0	1	1	0	0

In the transfer below, $A4$ is relocated from her current assignment (second preference) to her lowest uncovered zero (her first preference).

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	1	0	$0'$	0^*	0
$A2$	0	1	1	$0'$	0^*
$A3$	1	0^*	$0'$	0	0
$A4$	0^*	1	1	$0'$	0
$A5$	$0'$	1	1	0	0

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	1	0	0^*	0	0
$A2$	0	1	1	0	$0*$
$A3$	1	0^*	0	0	0
$A4$	0	1	1	0^*	0
$A5$	0^*	1	1	0	0

The final assignment is

	$H1$	$H2$	$H3$	$H4$	$H5$
$A1$	0	0	1	0	0
$A2$	0	0	0	0	1
$A3$	0	1	0	0	0
$A4$	0	0	0	1	0
$A5$	1	0	0	0	0

8.2 Flood' method

Step 1 Subtract the smallest element from all elements P.

Step 2 Find the smallest number of lines that covered all zeros from step 1. If there are N lines, terminate the process and the zeros will constitute a set of independent assignments. If not, go to step 3.

Step 3 Find the smallest uncovered element. Subtract that element from every uncovered column. Add that element from every covered row. Go back to step 2.

8.3 Flood' method for the example in section 4.3

The reduced market's preference matrix is

	$H2$	$H3$	$H4$	$H5$
$A2$	2	3	4	5
$A3$	5	2	3	1
$A4$	2	4	5	3
$A5$	3	4	5	1

Step 1 Subtract 1 *from all elements.*

	$H2$	$H3$	$H4$	$H5$
$A2$	1	2	3	4
$A3$	4	1	2	0
$A4$	1	3	4	2
$A5$	2	3	4	0

Step 2 Only one column is covered. Proceed to step 3.

	$H2$	$H3$	$H4$	$H5$
$A2$	1	2	3	4
$A3$	4	1	2	0
$A4$	1	3	4	2
$A5$	2	3	4	0

Step 3 Subtract 1 *from every uncovered column. Go back to step 2.*

	$H2$	$H3$	$H4$	$H5$
$A2$	0	1	2	4
$A3$	3	0	1	0
$A4$	0	2	3	2
$A5$	1	2	3	0

Step 2 Three columns are covered. Proceed to step 3.

	$H2$	$H3$	$H4$	$H5$
$A2$	0	1	2	4
$A3$	3	0	1	0
$A4$	0	2	3	2
$A5$	1	2	3	0

Step 3 Subtract 1 *from every uncovered column. Go back to step 2.*

	$H2$	$H3$	$H4$	$H5$
$A2$	0	1	1	4
$A3$	3	0	0	0
$A4$	0	2	2	2
$A5$	1	2	2	0

Step 2 Two columns and one row are covered. Proceed to step 3.

	$H2$	$H3$	$H4$	$H5$
$A2$	0	1	1	4
$A3$	3	0	0	0
$A4$	0	2	2	2
$A5$	1	2	2	0

Step 3 Subtract 1 *from every uncovered column. Add* 1 *from every covered row. Go back to step 2.*

	$H2$	$H3$	$H4$	$H5$
$A2$	0	0	0	4
$A3$	4	0	0	1
$A4$	0	1	1	2
$A5$	1	1	1	0

Step 2 Four lines are covered. The process terminates.

	$H2$	$H3$	$H4$	$H5$
$A2$	0	0	0	4
$A3$	4	0	0	1
$A4$	0	1	1	2
$A5$	1	1	1	0

8.4 Nash Equilibria in section 4. 4 lemma 1

	$H1$	$H2$	$H3$
$A1$	1	2	3
$A2$	1	3	2
$A3$	1	2	3

	$H1$	$H2$	$H3$
$A1$	1	2	3
$A2$	3	1	2
$A3$	1	2	3

	$H1$	$H2$	$H3$
$A1$	1	3	2
$A2$	2	3	1
$A3$	1	3	2

	$H1$	$H2$	$H3$
$A1$	1	3	2
$A2$	3	2	1
$A3$	1	3	2

	$H1$	$H2$	$H3$
$A1$	2	1	3
$A2$	1	2	3
$A3$	1	2	3

	$H1$	$H2$	$H3$
$A1$	2	1	3
$A2$	1	3	2
$A3$	1	2	3

	$H1$	$H2$	$H3$
$A1$	2	1	3
$A2$	3	1	2
$A3$	1	2	3

	$H1$	$H2$	$H3$
$A1$	3	1	2
$A2$	1	3	2
$A3$	1	2	3

	$H1$	$H2$	$H3$
$A1$	3	1	2
$A2$	2	3	1
$A3$	1	3	2