Ranked-to-Match: The Effects of Early Matching in the NRMP

Peter Doe

June 8, 2023

Abstract

We study a behavioral model of early matching within the context of the National Resident Matching Program. In our model, two hospitals choose to give early offers to doctors prior to a stable match. Some doctors have a behavioral preference to early match while others do not. We show that the less-desirable program benefits from the option to make early offers. Our results provide a theoretical foundation for behavior widely documented within the medical ethics and graduate medical education literature and confirm beliefs commonly held by residency program directors.

1 Introduction

The National Resident Matching Program (henceforth, NRMP) is a leading example of success in market design. Since 1952 the NRMP has brokered matches between medical school graduates and residency programs. Its persistence is viewed as a consequence of the stability of the implemented match (Roth, 1991) and the timing of the availability of information in the market. Recently, research in medical ethics and medical education has identified a potential source of market failure via coercive *post-interview communication* (henceforth, PIC); residency programs are successfully arranging matches through coercive PIC with doctors prior to the operation of the NRMP match.

The unraveling of the medical resident market to a date before the NRMP operates has welfare implications that we examine in this paper.

We focus on communication by hospital residency programs that either states how a hospital ranks a doctor, asks a doctor how she ranks the hospital, or implies that a positive rank by the hospital of the doctor depends upon a commitment from the doctor to rank the hospital first on her list.[1] Such PIC is a violation of the Match Agreement. We refer to PIC involving these questions or statements as *coercive*, and connect unraveling with coercive PIC. Early matches (which are a form of unraveling) are either formal offers for outside-of-match positions or informal agreements to mutually top-rank each other. Numerous free-form responses of surveyed doctors support a direct connection between coercive PIC and early matches. For instance, one surveyed doctor stated

> "Many program directors explicitly stated that my position on their rank list depended on postinterview communication. That a commitment to rank them first would increase my chances of matching at their program." (Williams et al., 2019)

As discussed in section 1.2, PIC is the object of numerous studies in the medical ethics and medical graduate education literatures. In these studies, doctors consistently report changes to their submitted Rank-Order List (ROL) as a result of PIC. That coercive PIC results in early matches (and hence unraveling) is well-documented by this literature.

In our model two hospitals hold a common preference over doctors, but the doctors' preferences are uncertain. A hospital that makes many early offers risks that too many doctors may accept, and hence not enough positions remain for doctors who only participate in the NRMP match. Conversely, a hospital that makes too few early offers allows desirable doctors to be poached by the other hospital in the early match process. When choosing to offer early matches to doctors, hospitals are forced to weigh the uncertain response of the doctors against the risk of desirable doctors accepting early match offers from another hospital. This forces hospitals to hedge by

[1] For conciseness, we refer to hospital residency programs simply as *hospitals* and medical seniors as *doctors*; additionally, we use masculine pronouns for hospitals and feminine pronouns for doctors.

making early offers to doctors that the hospital does not always desire to prevent these doctors from accepting early offers from another hospital. Because the less-desirable hospital matches to lower-ranked doctors, the less-desirable hospital is able to poach lower-ranked doctors that the more-desirable program has not made early offers to but may desire.

Our main result is that less-desirable programs benefit from early matching at the expense of more-desirable programs. The utility transfer occurs because the less-desirable program successfully poaches doctors near the bottom of more-desirable program's accepted doctors. As the effectiveness of coercive PIC increases, so does the transfer from the more-desirable program to the less-desirable program.

We interpret this result as a theoretical justification for commonly held beliefs within the medical education field. For instance, in a survey of program directors in internal medicine about the impact of a policy banning formal early matches (the All-In policy), Adams et al. (2012) states

> "The most commonly cited concern [about the All-In policy] was that smaller, nonuniversity programs and those in geographically less-desirable areas would suffer in recruitment."

Our interpretation is that banning formal early matches makes coercive PIC less effective. Program directors predicted that reducing the effectiveness of coercive PIC would transfer qualified doctors from lower-ranked to higher-ranked hospitals. We establish that coercive PIC is detrimental to more-desirable hospitals. We also find that the responsiveness of the doctors to coercive PIC has an effect on the distortion caused.

1.1 Institutional Background

Medical students spend four years in medical school before beginning residency. Residency programs provide medical graduates with hands-on experience in particular fields of medicine. A residency program lasts between three and seven years, and completion of the first year is required to become a licensed physician within the United States. In the 2022-2023 admissions cycle there were 48,156 applicants for 40,375 residency positions (Program, 2023).

The application cycle begins in September of the fourth year when the Electronic Residency Application System (ERAS) opens. Applicants have already completed the Step One and Step Two exams, and the Step Three exam is not completed until after matches are announced in March. In late September the applications are simultaneously released to residency programs. In early October, Medical Student Performance Evaluations are released by the medical schools to residency programs. Shortly after this, residency programs contact applicants to arrange interviews, with the majority of interviews concluded by the end of the year.

Each student and residency program submits a Rank-Order List (ROL) to the NRMP, with the deadline for both sides set in February. The NRMP computes a slightly modified version of the doctor-proposing deferred acceptance algorithm.[2] On Match Day in March, the NRMP announces to doctors and residency programs whom they have been matched with. The match outputted by the NRMP is a legally binding contract between doctors and residency programs.

For doctors who did not match and for residency programs with vacant seats remaining, the NRMP facilitates the Supplemental Offer and Acceptance Program (SOAP). SOAP is held for about a week shortly after the matches have been announced, and is colloquially referred to as the "scramble."

We draw attention here to the revelation of information to doctors and residency programs. After a doctor completes an interview at a residency program, there is no more information about the program that the doctor learns.[3] Similarly, after a residency program interviews a doctor, the only institutional feature by which the program may (potentially) gain information would be through the release of the Fall semester grades, but we have come across no discussion of these as

[2]The modifications have been shown by Roth and Peranson (1999) to have a minimal impact on the match achieved.

[3]Some doctors are invited back to the residency program for a "Second Look," but the surveys discussed in Subsection XX indicate that most doctors do not find these opportunities to be informative.

factors in the formulation of the ROL. In essence, all relevant information to make a ROL based on a doctor's or program's merits is available after interviews are conducted.

1.2 Literature Review

Our paper bridges the literature on inappropriate PIC within medical ethics and medical graduate education to the literature on unraveling within market design. We view a primary contribution of this paper to be the introduction of the discussion between medical researchers to researchers of market design. An overview of some relevant work within each literature is given below.

Studies on inappropriate PIC have been conducted within numerous medical specialties. The vast majority of studies are unincentivized surveys administered either to prospective or current residents or to program directors of residency programs, with a response rate of roughly 50% common among these surveys. Many address both coercive PIC and other kinds of inappropriate PIC, such as questions by hospitals concerning marital status, family plans, religion, and other interviews.

Surveys of (prospective) residents (either administered by a medical school to its students or by a residency to applicants) generally find substantial levels of inappropriate PIC. Fields surveyed include Pediatrics (Opel et al., 2007), Emergency Medicine (Thurman et al., 2009), Dermatology (Sbicca et al., 2010, 2012), Radiation Oncology (Holliday et al., 2015; Berriochoa et al., 2016), Orthopedic Surgery (Camp et al., 2016), Internal Medicine (Cornett et al., 2017; Williams et al., 2019; Swan and Baudendistel, 2014), and Integrated Vascular Surgery (Fereydooni et al., 2022). Additionally, numerous surveys are not specific to one field: Anderson et al. (1999); Pearson and Innes (1999); Miller et al. (2003); Jena et al. (2012); Monir et al. (2021). Several other studies find similar patterns within other matches: in the Urology match (Teichman et al., 2000; Sebesta et al., 2018; Handa et al., 2021) and in the military match (Ratcliffe et al., 2012).

There is not a standard template for the questions in these surveys. However, a few themes are apparent in the results. Doctors commonly report ($\sim$ 15% of respondents) being told they are "ranked-to-match." Less frequently ($\sim 5\%$ of respondents), doctors report being offered incentives to match early. Quite commonly ($\sim 25\%$ of respondents), doctors report changing their rank as a result of inappropriate (although not always coercive) PIC. Although the data quality is low due to the nature of the studies, the overall picture from these surveys is that coercive PIC occurs in some but not all cases, and that it has a meaningful impact on the final match.

Surveys to program directors are less frequent but still cover a wide spectrum of medical fields. Fields surveyed include General Surgery (Anderson and Jacobs, 2000), Family Practice (Carek et al., 2000), Dermatology (Sbicca et al., 2010), Obstetrics and Gynecology (Curran et al., 2012; Frishman et al., 2014), Internal Medicine (Chacko et al., 2018), Urology (Farber et al., 2019), and Otolaryngology (Harvey et al., 2019). Additionally, Grimm et al. (2016) surveys program directors without reference to a specific field. Program directors commonly report ($\sim 50\%$ of respondents) that they feel doctors have made an informal commitment to rank the hospital first. Understandably, surveys of program directors imply that substantially less inappropriate PIC is initiated by programs than implied by surveys of doctors

Economists have traditionally studied the causes of unraveling rather than its effect. The three primary causes of unraveling identified in the literature are the following. First, several authors identify the instability of the mechanism as a cause (Roth, 1991; Sönmez, 1999). Second, others identify uncertainty over the preferences of one or both sides as a cause (Roth and Xing, 1994; Li and Rosen, 1998; Hałaburda, 2010; Niederle et al., 2013; Ambuehl and Groves, 2020). The uncertainty that these models impose is on a common quality that is uncertain for one or both sides of the market but will be revealed. Third, a few authors identify costs as a source of unraveling; Damiano et al. (2005) examines search frictions as a source of unraveling, and Echenique and Pereyra (2016) identifies discounting combined with strategic complementarities as a source.

This paper is most similar to the papers which identify costs as a cause of unraveling. From the institutional background provided in Subsection 1.1, the mechanism is stable and there is no uncertainty over the quality of either side of the market that will be revealed before the match is arranged. We depart from the prior literature by assuming that unraveling occurs and focusing on its welfare effects. The only other authors to address early matching in the context of the NRMP is Ashlagi et al. (2023), who find that an institutionalized early match would have negative welfare consequences.

2 Model

2.1 Basics

There are two hospitals H and L (High and Low), each with capacity Cap. There is a continuum of doctors with scores distributed on the interval $[0, 1]$. Doctors types are further differentiated along two dimensions. First, each doctor has a *hospital preference* of H or L. Second, each doctor has a *match preference* of early (ER) or late (LA). We define $\Theta = [0, 1] \times \{H, L\} \times \{\text{ER}, \text{LA}\}$ as the doctor space, and let $f_{\text{Full}} : \Theta \times \Omega \to \mathbb{R}_+$ be the density of doctors conditioned on the state ω.[4] For simplicity, we assume that score, hospital preference, and match preference are mutually independent, and that the state ω only enters f_{Full} through the hospital preference term. Hence we have

$$f_{\text{Full}}(x, I, P \mid \omega) = f(x) \cdot h(I \mid \omega) \cdot p(\text{P})$$

with the normalization

$$h(H \mid \omega) + h(L \mid \omega) = 1$$

$$p(\text{ER}) + p(\text{LA}) = 1$$

Early-match preference doctors have preferences of either $H \succ L \succ \emptyset$ or $L \succ H \succ \emptyset$, (doctors with hospital preference H or L, respectively) where $\emptyset$ denotes an outcome in which that doctor unmatched. Furthermore, we assume that ER doctors prefer matching have a behavioral bias

[4] Note that $\int_\Theta f(\theta) d\theta \neq 1$ in general.

	ER	LA
H	$H \succ L \succ \emptyset$ prefers to match early	$H \succ L \succ \emptyset$ prefers to match late
L	$L \succ H \succ \emptyset$ prefers to match early	$L \succ H \succ \emptyset$ prefers to match late

Figure 1: Doctor preferences conditioned on type

toward a guaranteed partner. That is, given the opportunity to match to her second-favorite hospital with certainty or to partake in the NRMP match, the ER doctor prefers the former.

Late-match preference doctors have preferences of either $H \succ L \succ \emptyset$ or $L \succ H \succ \emptyset$ (doctors with hospital preference H or L, respectively). Furthermore, late-match preference doctors *never* reveal their preferences if asked. We interpret this as LA doctors complying fully with the Match Agreement statements about persuasion.[5]

We assume that $p(\text{ER}) = r$ for $0 < r < 1$. We call r the *responsiveness* of the market because it reflects how likely a doctor is to respond to an offer to match before the stable match is arranged.

Hospital preferences over doctors are given by the following utility function. If μ is a matching, then the utility of hospital $I \in \{H, L\}$ is

$$u_I(\mu \mid \omega) = \int_\Theta \mu_\theta(I)\theta_1 f_{\text{Full}}(\theta \mid \omega) d\theta$$

We denote the expected utility:

$$u_I(\mu) = \mathbb{E}_\omega \left[\int_\Theta \mu_\theta(I)\theta_1 f_{\text{Full}}(\theta \mid \omega) d\theta \right]$$

[5]An alternate interpretation is that LA doctors are unsure of their own preferences, or have yet to conclude interviews, etc.

A *matching* is a function $\mu : \{H, L\} \times \Theta \to \{0, 1\}$ such that $\mu_\theta(H) + \mu_\theta(L) \leq 1$. A matching μ is *stable* if there no $I \in \{H, L\}$ and $\theta, \theta' \in \Theta$ such that $\theta_1 > \theta'_1$, $\mu_\theta(I) = 0$, $\mu_{\theta'}(I) = 1$, $I \succ_{\theta_1} J$ for $\mu_\theta(J) = 1$.

2.2 Uncertainty

We assume that there is uncertainty over the aggregate hospital preferences of the doctors. We model this by making the state $\omega \in \{\omega_H, \omega_L\}$ a random variable such that

$$h(H \mid \omega_H) = 1$$

$$h(H \mid \omega_L) = \frac{1}{2}$$

If $\omega = \omega_H$, then we say that H is *popular* and L is *unpopular*.[6] Otherwise, we say that L is popular and H is unpopular. Which hospital is popular is unknown to both hospitals, but both hospitals share the belief

$$\Pr(H \text{ is popular}) = \Pr(\omega = \omega_H) = \frac{1}{2}$$

We emphasize here that the only uncertainty facing a doctor is the aggregate preferences of the *other* doctors; every doctor knows her own preferences (whether H or L is preferred and whether the doctor is ER or LA) and $f(\cdot)$ and $p(\cdot)$.

To summarize, hospitals are uncertain over doctors' preferences. This uncertainty is both individual (hospitals do not know θ's preferences) and aggregate (hospitals do not know which hospital is popular). Each hospital and each doctor, however, is certain of his or her own preferences.

[6]We choose 1 and $\frac{1}{2}$ here to simplify the analysis. When $h(H \mid \omega_L) \neq \frac{1}{2}$, finding the equilibrium of the model becomes substantially more difficult and counter intuitive. We consider that the clarity of exposition is worth the simplification. We conjecture that our results still hold on a portion of the parameter space when this requirement is relaxed.

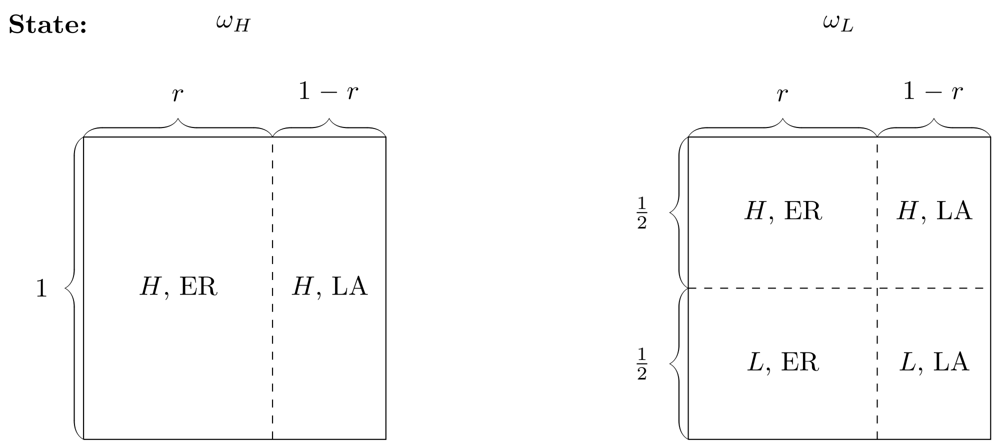

Figure 2: Distribution of doctor types conditioned on the state ω

2.3 Structure

There are two phases. First, in the *early-matching* phase, each hospital I chooses a Lebesgue-measurable set Π_I to make private offers to, with $\Pi = (\Pi_H, \Pi_L)$. We use λ to denote Lebesgue measure. Each doctor type θ observes the set of offers made to her. Doctors of type θ then responds by choosing a function $\rho_\theta(I \mid \Pi)$ such that

$$\rho_\theta(I \mid \Pi) \in \{0, 1\} \qquad \forall I \in \{S, L\}$$

$$\rho_\theta(H \mid \Pi) + \rho_\theta(L \mid \Pi) \leq 1$$

The first line requires that θ gives a binary response, and the second line requires that θ accept at most one early offer. We note that ρ is a matching.

If a doctor accepts an offer, it is a binding commitment between the hospital and the doctor; the doctor is matched with the hospital and the hospital sets a new capacity $\mathrm{Cap}_I(\rho, \omega) = \mathrm{Cap}_I - \int_\Theta \rho_\theta(I \mid \Pi) f_{\mathrm{Full}}(\theta \mid \omega) d\theta$. The density of remaining doctors is $g(\theta \mid \rho, \omega) = f_{\mathrm{Full}}(\theta \mid \omega) \cdot (1 - \rho_\theta(H \mid \Pi) - \rho_\theta(L \mid \Pi))$.

Second, in the *late-matching* phase, the doctor-optimal stable match $\mu(\cdot \mid g, \mathrm{Cap})$ is implemented among the remaining doctors and the remaining hospital capacities based on agents' true preferences

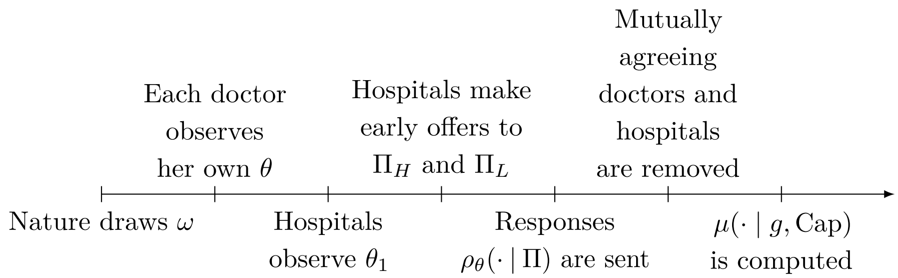

Figure 3: Timeline of events

regardless of private offers in the late-matching phase.[7] In this case (See Azevedo and Leshno (2016) for more details), each hospital sets a cutoff (H sets q_H and L sets q_L) and the match correspondence is

$$\mu_{(x,H,\mathrm{P})}(H \mid g, \mathrm{Cap}) = \mathbb{1}\{x \geq q_H\}$$

$$\mu_{(x,H,\mathrm{P})}(L \mid g, \mathrm{Cap}) = \mathbb{1}\{x \geq q_L \text{ and } x < q_H\}$$

and similarly for (x, L, P).

In the event that for some I $\int_\Theta \rho_\theta(I \mid \Pi) f_{\mathrm{Full}} d\theta > \mathrm{Cap}$, we truncate $\rho_\theta(I \mid \Pi)$ to $\tilde{\rho}_\theta(I \mid \Pi)$ defined as follows. Define x such that $\int_{\{\theta \mid \theta_1 \geq x\}} \rho_\theta(I \mid \Pi) f_{\mathrm{Full}} d\theta = \mathrm{Cap}$. Then

$$\tilde{\rho}_\theta(I \mid \Pi) = \begin{cases} \rho_\theta(I \mid \Pi) & \theta_1 \geq x \\ 0 & \text{otherwise} \end{cases}$$

However, the hospital I that exceeded capacity in the early match is penalized by having 3 deducted from his utility, where 3 is chosen so that neither hospital would choose to exceed capacity.

2.4 Equilibria

Each doctor responds deterministically based on her type θ to offers made. If $\theta_3 = \mathrm{LA}$ (the doctor is has a preference for late matching), then $\rho_\theta(\cdot \mid \Pi) = 0$ (the doctor rejects all offers). If

[7] The condition that hospitals only need rank accepting doctors first is crucial. Otherwise, the offers in the late match phase would be more difficult to analyze.

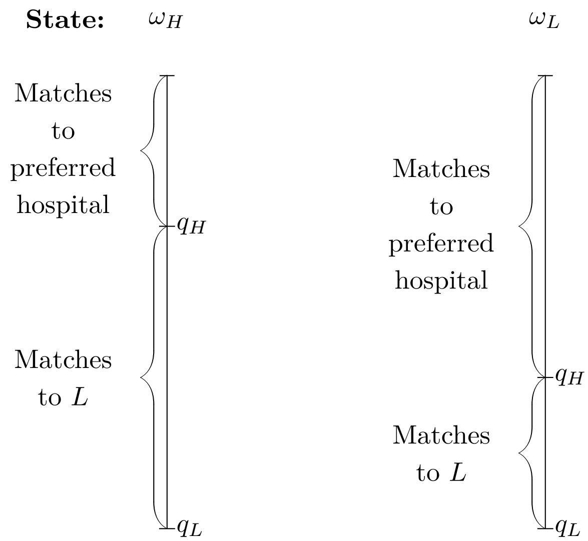

Figure 4: Cut-off structure

$\theta_3 = \text{ER}$ (the doctor has a preference for early matching), then the doctor accepts the offer from her most-preferred hospital (if she receives an offer). Formally, $\rho_\theta(\theta_2 \mid \Pi) = \Pi_{\theta_2}(\theta_1)$ and $\rho_\theta(I \mid \Pi) = 1$ $\Pi_{\theta_2}(\theta_1) = 0$ and $\Pi_I(\theta_1) = 1$ for $I \neq \theta_2$, and $\Pi_I(\theta_1) = 0$ otherwise.

Let

$$\text{LATE}_\theta(\cdot \mid \Pi, \omega) = \mu\Big(\ \cdot\ \mid g\big(\ \cdot\ \mid \rho_\theta(\cdot \mid \Pi), \omega\big), \text{Cap}\big(\rho_\theta(\cdot \mid \Pi), \omega\big)\Big)$$

A strategy for hospital I is a Lebesgue-measurable set $\Pi_I \subseteq [0,1]$. To clarify the exposition, we provide the following definition to reduce the number of equilibria. In words, it states that if two strategies produce the same matching in both phases except on a set of doctors of measure zero, then the two strategies are the same.

Definition 1. *Two strategies* Π *and* Π' *are said to be outcome-equivalent if* $\rho_\theta(I \mid \Pi) + LATE(I \mid \Pi) = \rho_\theta(I \mid \Pi') + LATE(I \mid \Pi')$ *for almost every* θ *for both* $I \in \{H, L\}$ *and both* $\omega \in \{\omega_H, \omega_L\}$.

Our equilibrium concept is a pure-strategy Nash equilibrium. An *equilibrium* is a vector Π such that

$$\Pi_I \in \arg\max_{\Pi'_I} \mathbb{E}_\omega u\Big(\rho(I \mid \Pi) + \text{LATE}(I \mid \Pi, \omega)\Big)$$

Although pure-strategy equilibria are not guaranteed to exist, in our analysis we prove its existence in this model.

3 Analysis

Our analysis has two main results. In theorem 1 we prove that there is a unique equilibrium, and that this equilibrium is a cutoff equilibrium. In theorem 2 we prove that hospital L's utility is increasing in r. There is a transfer of utility from hospital H to hospital L that increases in magnitude as r increases.

Formally, our results are the following. First, if Π_I is outcome-equivalent to $\tilde{\Pi}_I = [\pi_I, 1]$, then we say that Π_I is a *cutoff strategy*. For simplicity, we refer cutoff strategies by the cutoff π.

Let

$$\pi^*_H(r) = 1 - \frac{4\text{Cap}}{3+r}$$
$$\pi^*_L(r) = 1 - \frac{6\text{Cap}}{3+r}$$

We depict $\pi^*(r)$ in Figure 4. Since the cutoffs in the match phase are deterministic given Π and the state ω, we write $q_H(\Pi \mid \omega)$ and $q_L(\Pi \mid \omega)$.

Our first result is that $\pi^*(r)$ is an equilibrium for every $r \in [0, 1]$, and is unique.

Theorem 1. *$\pi^*(r)$ is the unique equilibrium.*

Despite the flexibility in choosing Π, each hospital only uses cutoff strategies. Perhaps counterintuitively, in Lemma 1 we show that *every* best-response is a cutoff strategy.

To prove Theorem 1, we use backward induction. The critical step is the proof that if Π_I is a best-response, then Π_I is a cutoff strategy. We state this in Lemma 1:

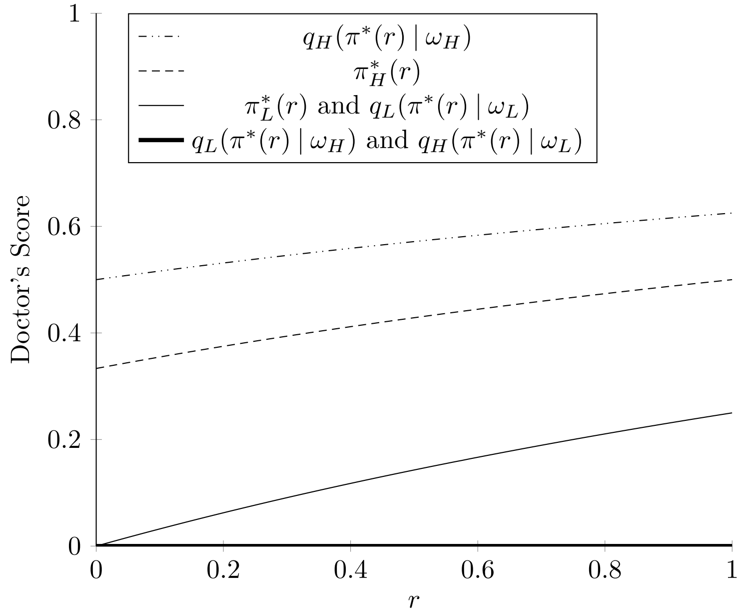

Figure 5: Equilibrium structure when Cap = 1/2. Hospital H makes an offer to doctors above $q_H(\pi \mid \omega_L)$ in the late match and makes an early offer to doctors above π_H in the early match.

Lemma 1. *If Π_I is a best response to Π_J, then there exists some Π_I^* that is outcome-equivalent to Π_I such that*

$$\Pi_I^* = [\pi_I, 1]$$

where $q_I(\Pi \mid \omega_J) \leq \pi_I \leq q_I(\Pi \mid \omega_J)$.

If there are multiple π that satisfy the lemma, we take π to be the smallest one.

Lemma 1 narrows the scope of possible deviations Π_I' tremendously. Because q is a deterministic function π, finding equilibria becomes equivalent to checking possible $\pi_I \in [0, 1]$.

The main difficulty in proving Lemma 1 is that early offers made by I to doctors in Π_J have different yields based on the state ω, whereas early offers made to doctors *not* in Π_J always result in perfect yield. A generic strategy Π_I could produce a variety of combinations of yield based on the state ω that cutoff strategies cannot replicate.

The insight that resolves this difficulty is that when hospital I drops doctors from Π_I, these doctors either accept early offers from hospital J (giving hospital I access to doctors on the margin of q_J) or are still available (so hospital I is just as well off). The proof of Lemma 1 leverages this in case 3.

Here we provide a sketch of the proof of Lemma 1. Toward a contradiction, we suppose that the lemma does not hold. Then there are two sets of doctors A and B with positive measure such that B is above A with regards to θ_1. The three cases are:

1. If B is above $q_I(\pi^* \mid \omega_I)$, then hospital I should always make early offers to doctors in B.

2. If B is below $q_I(\pi^* \mid \omega_I)$ and below $q_J(\pi^* \mid \omega_J)$, then hospital I should drop doctors in A. This follows because if $\omega = \omega_I$, hospital I does not want these doctors, and if $\omega = \omega_J$, then hospital J will not make regular offers to A (and any doctor hospital J gains from hospital I dropping A lets hospital I gain a doctor on the margin of $q_J(\pi^* \mid \omega_J)$).

3. If B is below $q_I(\pi^* \mid \omega_I)$ and above $q_J(\pi^* \mid \omega_J)$, then hospital I should drop some doctors in A and admit some doctors in B such that $q_I(\pi^* \mid \omega_I)$ remains the same. This works because, in state ω_I, hospital I shifts toward higher scoring doctors without affecting the cutoffs. If $\omega = \omega_J$, then hospital I is better off because any doctor in A that also received an early offer from hospital J now accepts the early offer from hospital J, making hospital J worse off. Because the set of matched doctors is the same (all doctors with $\theta_1 \in [1 - 2\text{Cap}, 1]$), we see that any loss to hospital J is a gain for hospital I.

These cases complete the proof.

Using Lemma 1, proving Theorem 1 involves tedious algebra and case work. The proof mostly consists of guessing the order of the cutoffs q and π. The proof is relegated to the appendix.

In our second theorem, we see that hospital H's utility decreases in r.

Theorem 2. *$u_H(\pi^*(r))$ is decreasing in r.*

The proof of Theorem 2 relies only on calculating a derivative, so we relegate it to the supplementary *Mathematica* notebook.

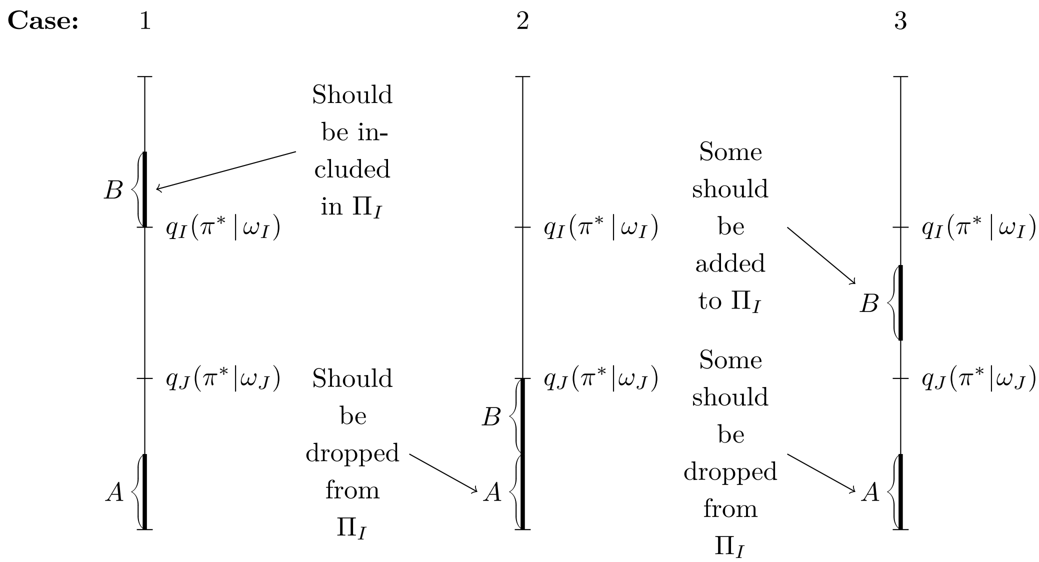

Figure 6: Illustration of the proof of Lemma 1. Hospital I makes offers to doctors in A but not in B. In Case 1, hospital I should always make offers to doctors in B because regardless of the state ω hospital I *always* wants these doctors. In Case 2, hospital I should never make offers to doctors in A because regardless of the state ω hospital I *never* wants these doctors. In Case 3, hospital I should replace offers made to doctors in A with offers made to doctors in B because hospital I prefers doctors in B and if the state is ω_J then every doctor in A is still available (or is early matched to hospital J, which hurts hospital J and thus helps hospital I).

Because the same set of doctors is accepted for every r and no hospital exceeds capacity, it follows that hospital L's utility is increasing when hospital H's utility is decreasing. Hence hospital L prefers larger r.

4 Discussion

We connect Theorem 2 to the commonly held belief that the All-In policy would harm less competitive programs. In our model, r reflects how attractive early offers are to doctors. If hospitals are unable to make contractual agreements with doctors prior to the NRMP, then some doctors who could have been persuaded to accept a early offer will not accept it. Similarly, stronger restrictions on PIC may prevent hospitals from as effectively persuading doctors to accept an early match. We interpret these as lowering the responsiveness of the market, r. Theorem 2 predicts that the less competitive hospitals would be harmed by these policies.

Our model does not allow us to test different proposals to reduce early matching, but it does provide predictions about the number of early offers made and the hospitals that benefit from early matches. One counterintuitive observation from π^* is that as the doctors become less responsive, *more* early offers are made. Policies intended to reduce the rate of early matches are likely to increase the number of early offers made even though the number of early matches decreases.

We see Theorem 1 as a technical advancement that could be useful for solving extensions of the current model. The main difficulty with establishing that best-responses are only in cutoff strategies is that non-cutoff strategies are able to flexibly combine doctors included in other hospitals early offer sets to produce complex lotteries over the early match ρ. Theorem 1 demonstrates that when there are just two hospitals, this flexibility is unnecessary. The two-hospital assumption makes the analysis tractable.

Two extensions are immediately apparent. The first is to increase the number of hospitals. Extending Lemma 1 to this case may prove impossible, and we conjecture that hospitals will have a cutoff for every combination of other hospitals. The key insight from the proof of Lemma 1 is that when hospital I makes early offers to doctors below $q_J(\pi \mid \omega_J)$ (hospital J's cutoff when J is popular), these early offers are only beneficial to hospital I in the state that J is popular; hence the early offers to these doctors are unnecessary. It is not apparent how viable this lemma is with more than two hospitals.

In our model, hospital L gains utility from early offers because hospital L has lower cutoffs in the late match than hospital H. Hospital L does not need to hedge his bets concerning q_L because q_L is always below π_H. For a multi-hospital extension, it is unclear how this property would translate. It would be fruitful to further understand the properties of a hospital that cause it to gain from early offers.

A second extension is to consider the welfare of the doctors. It would be interesting to examine which types of doctors are impacted most by early offers, and whether some doctors are likely to gain from the early offer system.

References

M. Adams, T. B. Morrison, S. Call, A. J. Halvorsen, J. Moore, M. Lucarelli, S. Angus, and F. S. McDonald. Changes in the national residency matching program policy: are internal medicine program directors "all-in"? *Journal of Graduate Medical Education*, 4(2):148–153, 2012.

S. Ambuehl and V. Groves. Unraveling over time. *Games and Economic Behavior*, 121:252–264, 2020.

K. D. Anderson and D. M. Jacobs. General surgery program directors' perceptions of the match. *Current surgery*, 57(5):460–465, 2000.

K. D. Anderson, D. M. Jacobs, and A. V. Blue. Is match ethics an oxymoron? *The American journal of surgery*, 177(3):237–239, 1999.

I. Ashlagi, E. Love, J. I. Reminick, and A. E. Roth. Early vs single match in the transition to residency: Analysis using nrmp data from 2014 to 2021. *Journal of Graduate Medical Education*, 15(2):219–227, 2023.

E. M. Azevedo and J. D. Leshno. A supply and demand framework for two-sided matching markets. *Journal of Political Economy*, 124(5):1235–1268, 2016.

C. Berriochoa, M. C. Ward, M. A. Weller, E. Holliday, A. Kusano, C. R. Thomas Jr, and R. D. Tendulkar. Applicant interview experiences and postinterview communication of the 2016 radiation oncology match cycle. *International Journal of Radiation Oncology* Biology* Physics*, 96(3):514–520, 2016.

C. L. Camp, P. L. Sousa, A. D. Hanssen, M. D. Karam, G. J. Haidukewych, D. A. Oakes, and N. S. Turner. Orthopedic surgery applicants: What they want in an interview and how they are influenced by post–interview contact. *Journal of surgical education*, 73(4):709–714, 2016.

P. J. Carek, K. D. Anderson, A. V. Blue, and B. E. Mavis. Recruitment behavior and program directors: How ethical are their perspective about the match process? *FAMILY MEDICINE-KANSAS CITY-*, 32(4):258–260, 2000.

K. M. Chacko, S. Reddy, M. Kisielewski, S. Call, L. L. Willett, and S. Chaudhry. Postinterview communications: two surveys of internal medicine residency program directors before and after guideline implementation. *Academic Medicine*, 93(9):1367–1373, 2018.

P. A. Cornett, C. Williams, R. L. Alweis, J. McConville, M. Frank, B. Dalal, R. I. Kopelman, V. P. Luther, A. B. O'connor, and E. A. Muchmore. Problematic communications during 2016 fellowship recruitment in internal medicine. *Journal of community hospital internal medicine perspectives*, 7(5):277–281, 2017.

D. S. Curran, P. B. Andreatta, X. Xu, C. E. Nugent, S. R. Dewald, and T. R. Johnson. Postinterview communication between obstetrics and gynecology residency programs and candidates. *Journal of Graduate Medical Education*, 4(2):165–169, 2012.

E. Damiano, H. Li, and W. Suen. Unraveling of dynamic sorting. *Review of Economic Studies*, 72: 1057–1076, 2005.

F. Echenique and J. S. Pereyra. Strategic complementarities and unraveling in matching markets. *Theoretical Economics*, 11(1):1–39, 2016.

N. J. Farber, C. J. Neylan, A. Srivastava, S. S. Pandya, A. Kaplan, K. D. Radadia, E. A. Singer, and S. E. Elsamra. The urology match process and the limited value of post-interview communication for program directors. *Urology*, 128:23–30, 2019.

A. Fereydooni, J. L. Ramirez, K. L. Morrow, S. Sorondo, J. T. Lee, D. M. Coleman, and V. Chandra. Interview experience, post-interview communication, and gender-based differences in the integrated vascular surgery residency match. *Journal of Vascular Surgery*, 75(1):316–322, 2022.

G. N. Frishman, K. A. Matteson, J. L. Bienstock, K. E. George, T. Ogburn, P. N. Rauk, P. F. Schnatz, and L. A. Learman. Postinterview communication with residency applicants: a call for clarity! *American journal of obstetrics and gynecology*, 211(4):344–350, 2014.

L. J. Grimm, C. S. Avery, and C. M. Maxfield. Residency postinterview communications: more harm than good? *Journal of Graduate Medical Education*, 8(1):7–9, 2016.

H. Hałaburda. Unravelling in two-sided matching markets and similarity of preferences. *Games and Economic Behavior*, 69(2):365–393, 2010.

N. Handa, A. X. Vo, J. Guo, M. K. Keeter, and S. J. Kielb. The impact of stricter guidelines on post-interview contact, restricted interview topics and differences by gender in the urology match. *Urology Practice*, 8(1):143–148, 2021.

R. S. Harvey, K. J. Kovatch, M. E. Prince, K. M. Malloy, and M. C. Thorne. Trends in post-interview communication practices. *The Laryngoscope*, 129(3):607–612, 2019.

E. B. Holliday, C. R. Thomas Jr, and A. S. Kusano. Integrity of the national resident matching program for radiation oncology: national survey of applicant experiences. *International Journal of Radiation Oncology* Biology* Physics*, 92(3):525–531, 2015.

A. B. Jena, V. M. Arora, K. E. Hauer, S. Durning, N. Borges, N. Oriol, D. M. Elnicki, M. J. Fagan, H. E. Harrell, D. Torre, et al. The prevalence and nature of postinterview communications between residency programs and applicants during the match. *Academic Medicine*, 87(10): 1434–1442, 2012.

H. Li and S. Rosen. Unraveling in matching markets. *American Economic Review*, pages 371–387, 1998.

J. B. Miller, D. C. Schaad, R. A. Crittenden, N. E. Oriol, and C. MacLaren. Communication between programs and applicants during residency selection: effects of the match on medical students' professional development. *Academic Medicine*, 78(4):403–411, 2003.

R. L. Monir, K. Michaudet, J. G. Monir, K. P. Rahmanian, C. Michaudet, L. A. Cooper, and H. Harrell. Impact of match violations on applicants' perceptions and rankings of residency programs. *Cureus*, 13(1), 2021.

M. Niederle, A. E. Roth, and M. U. Ünver. Unraveling results from comparable demand and supply: An experimental investigation. *Games*, 4(2):243–282, 2013.

D. Opel, R. Shugerman, H. McPhillips, W. S. Swanson, S. Archibald, and D. Diekema. Professionalism and the match: a pediatric residency program's postinterview no-call policy and its impact on applicants. *Pediatrics*, 120(4):e826–e831, 2007.

R. Pearson and A. Innes. Ensuring compliance with nrmp policy. *Academic Medicine*, 74(7):747–8, 1999.

N. R. M. Program. Results and data 2023 main residency match®, 2023. https://www.nrmp.org/wp-content/uploads/2023/05/2023-Main-Match-Results-and-Data-Book-FINAL.pdf.

T. A. Ratcliffe, S. J. Durning, A. B. Jena, T. Grau, A. R. Artino Jr, V. M. Arora, K. E. Hauer, N. Borges, N. Oriol, D. M. Elnicki, et al. Postinterview communication between military residency applicants and training programs. *Military medicine*, 177(suppl_9):54–60, 2012.

A. E. Roth. 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*, pages 415–440, 1991.

A. E. Roth and E. Peranson. The redesign of the matching market for american physicians: Some engineering aspects of economic design. *American economic review*, 89(4):748–780, 1999.

A. E. Roth and X. Xing. Jumping the gun: Imperfections and institutions related to the timing of market transactions. *The American Economic Review*, pages 992–1044, 1994.

J. A. Sbicca, E. S. Gorell, M. H. Kanzler, and A. T. Lane. The integrity of the dermatology national resident matching program: results of a national study. *Journal of the American Academy of Dermatology*, 63(4):594–601, 2010.

J. A. Sbicca, E. S. Gorell, D. H. Peng, and A. T. Lane. A follow-up survey of the integrity of the dermatology national resident matching program. *Journal of the American Academy of Dermatology*, 67(3):429–435, 2012.

E. M. Sebesta, M. J. Lipsky, M. Nunez, K. L. Cooper, and G. M. Badalato. The national resident matching program code of conduct: what is the perceived degree of compliance during the urology match process? *Urology*, 122:37–43, 2018.

T. Sönmez. Can pre-arranged matches be avoided in two-sided matching markets? *Journal of Economic theory*, 86(1):148–156, 1999.

E. C. Swan and T. E. Baudendistel. Relationship between postinterview correspondence from residency program applicants and subsequent applicant match outcomes. *Journal of Graduate Medical Education*, 6(3):478–483, 2014.

Teichman, K. D. Anderson, M. M. Dorough, C. R. Stein, S. A. Optenberg, and I. M. Thompson. The urology residency matching program in practice. *The Journal of urology*, 163(6):1878–1887, 2000.

R. J. Thurman, E. Katz, W. Carter, J. Han, E. Kayala, N. McCoin, and A. B. Storrow. Emergency medicine residency applicant perceptions of unethical recruiting practices and illegal questioning in the match. *Academic Emergency Medicine*, 16(6):550–557, 2009.

C. M. Williams, R. L. Alweis, A. B. O'Connor, B. Dalal, D. Rai, A. Abdullah, R. Kopelman, P. Cornett, M. O. Frank, V. P. Luther, et al. Inappropriate communication during internal medicine fellowship recruitment: a mixed-methods analysis. *The American journal of medicine*, 132(6):770–775, 2019.

5 Appendix

Proof of Lemma 1

Proof. Toward a contradiction, suppose that Π_I is a best response to Π_J and that Π_I is not outcome-equivalent to Π_I^* for any π_I. Then there are sets A and B such that $\lambda(A) > 0$, $\lambda(B) > 0$, and $\inf B \geq \sup A$. WLOG there are three cases:

1. $\inf B > q_I(\Pi \mid \omega_I)$: In this case, consider the deviation

$$\Pi_I^\beta = \Pi_I \cup \{b \in B \text{ s.t. } b \geq \beta\}$$

   Because $r > 0$ there exists some β^* such that $\beta^* \geq q_I(\Pi_I^{\beta^*}, \Pi_J \mid \omega_I)$ and $\lambda([\beta^*, 1] \cap B) > 0$.

   If $q_I(\Pi_I^{\beta^*}, \Pi_J \mid \omega_I) = q_I(\Pi \mid \omega_I)$ for every such B and β, then the supposition that Π_I is not outcome-equivalent to some Π_I^* is violated, a contradiction. Hence, take B and β^* such that $q_I(\Pi_I^{\beta^*}, \Pi_J \mid \omega_I) \neq q_I(\Pi \mid \omega_I)$.

Because $\Pi_I^* \supset \Pi_I$, it follows that

$$\int_\Theta \rho_\theta(I \mid \Pi_I^{\beta^*}, \Pi_J) f_{\text{Full}}(\theta \mid \omega) d\theta > \int_\Theta \rho_\theta(I \mid \Pi) f_{\text{Full}}(\theta \mid \omega) d\theta$$

Hence, $q_I(\Pi_I^{\beta^*}, \Pi_J \mid \omega_I) > q_I(\Pi \mid \omega_I)$.

- If $\omega = \omega_I$, then I has traded doctors in the interval $[q_I(\Pi \mid \omega_I), q_I(\Pi_I^{\beta^*}, \Pi_J \mid \omega_I)]$ for doctors above β^*.
- If $\omega = \omega_J$, then I has traded doctors in the interval $[q_I(\Pi \mid \omega_J), q_I(\Pi_I^{\beta^*}, \Pi_J \mid \omega_J)]$ for doctors above β^*.

Hence, I is strictly better off in both states, a contradiction. This also establishes that $\pi_I \leq q_I(\Pi \mid \omega_I)$.

2. $\sup B < q_I(\Pi \mid \omega_I)$ and $\sup B < q_J(\Pi \mid \omega_J)$: In this case, consider the deviation

$$\Pi_I^\alpha = \Pi_I \backslash \{a \in A \text{ s.t. } a < \alpha\}$$

Because $r > 0$ there exists some α^* such that $\lambda([0, \alpha^*] \cap A) > 0$ and $q_I(\Pi_I^{\alpha^*}, \Pi_J | \omega_I) \neq q_I(\Pi | \omega_I)$. and $\sup A < q_J(\Pi_I^{\alpha^*}, \Pi_J \mid \omega_I)$.

If $q_I(\Pi_I^{\alpha^*}, \Pi_J \mid \omega_I) = q_I(\Pi \mid \omega_I)$ for every such A and α, then the supposition that Π_I is not outcome-equivalent to some Π_I^* is violated, a contradiction. Hence, take A and α^* such that $q_I(\Pi_I^{\alpha^*}, \Pi_J \mid \omega_I) \neq q_I(\Pi \mid \omega_I)$.

Because $\Pi_I^* \subset \Pi_I$, it follows that

$$\int_\Theta \rho_\theta(I \mid \Pi_I^{\alpha^*}, \Pi_J) f_{\text{Full}}(\theta \mid \omega) d\theta < \int_\Theta \rho_\theta(I \mid \Pi) f_{\text{Full}}(\theta \mid \omega) d\theta$$

Hence, $q_I(\Pi_I^{\alpha^*}, \Pi_J \mid \omega_I) < q_I(\Pi \mid \omega_I)$.

- If $\omega = \omega_I$, then I has traded doctors below α^* for doctors in the interval $[q_I(\Pi_I^{\alpha^*}, \Pi_J | \omega_I), q_I(\Pi | \omega_I)]$. This is a strict improvement.

- If $\omega = \omega_J$, then every doctor that I forgoes (doctors (θ_1, I, ER) such that $\theta_1 < \alpha^*$, $\theta_1 \in A$, and $\theta_1 \in \Pi_J$) is either still available (if J overfills his capacity in the early match) or implies that I can acquire a doctor in $[q_J(\Pi_I^{\alpha^*}, \Pi_J \mid \omega_I), q_J(\Pi \mid \omega_I)]$. This is a weak improvement.

Hence, I is strictly better off, a contradiction.

3. $\sup B < q_I(\Pi \mid \omega_I)$ and $\inf B > q_J(\Pi \mid \omega_J)$: In this case, consider the deviation

$$\Pi_I^{\alpha,\beta} = \Big(\Pi_I \cup \{b \in B \text{ s.t. } b \geq \beta\}\Big) \backslash \{a \in A \text{ s.t. } a < \alpha\}$$

There exists some α^* and β^* such that $\lambda([0, \alpha^*] \cap A) > 0$, $\lambda([\beta^*, 1] \cap B) > 0$, and $q_I(\Pi_I^{\alpha^*,\beta^*}, \Pi_J | \omega_I) = q_I(\Pi \mid \omega_I)$.

It follows that

$$\int_\Theta \rho_\theta(I \mid \Pi_I^{\alpha^*,\beta^*}, \Pi_J) f_{\text{Full}}(\theta \mid \omega) d\theta = \int_\Theta \rho_\theta(I \mid \Pi) f_{\text{Full}}(\theta \mid \omega) d\theta$$

Observe also that

$$\int_\Theta \rho_\theta(I \mid \Pi_I^{\alpha^*,\beta^*}, \Pi_J) f_{\text{Full}}(\theta \mid \omega_J) d\theta \leq \int_\Theta \rho_\theta(I \mid \Pi) f_{\text{Full}}(\theta \mid \omega_J) d\theta$$

with strict inequality if $\lambda(A \cap [0, q_J(\Pi \mid \omega_J)]) > 0$. Hence $q_J(\Pi_I^{\alpha^*,\beta^*}, \Pi_J \mid \omega_J) \leq q_J(\Pi \mid \omega_J)$.

- If $\omega = \omega_I$, then I has traded doctors below α^* for doctors above β^*. This is a strict improvement.
- If $\omega = \omega_J$, observe that $q_J(\Pi_I^{\alpha^*,\beta^*}, \Pi_J \mid \omega_J) \leq q_J(\Pi \mid \omega_J)$. Hence, J acquires weakly worse doctors under $\Pi_I^{\alpha^*,\beta^*}$ than Π_I. Because the every doctor who was matched under Π is matched under $\Pi_I^{\alpha^*,\beta^*}, \Pi_J$, it follows that I is weakly better off.

Hence, I is strictly better off, a contradiction.

In all three cases, a contradiction was reached. Hence, Π_I is outcome-equivalent to $[\pi_I, 1]$ for some π_I.

To establish the final claim that $\pi_I \geq q_I(\Pi \mid \omega_J)$, suppose toward a contradiction that $\pi_I < q_I(\Pi \mid \omega_J)$. Then hospital I could instead respond with $\pi_I + \epsilon$ for $\epsilon > 0$. For ϵ small enough, $\pi_I + \epsilon < q_I(\pi_I + \epsilon, \Pi_J \mid \omega_J)$. Hospital I prefers to not be matched with doctors with scores in $[\pi_I, \pi_I + \epsilon]$ and instead be matched with doctors in the range $[q_I(\pi_I, \Pi_J \mid \omega_J), q_I(\pi_I + \epsilon, \Pi_J \mid \omega_J)]$ or $[q_I(\pi_I, \Pi_J \mid \omega_I), q_I(\pi_I + \epsilon, \Pi_J \mid \omega_I)]$. Hence, $\pi_I + \epsilon$ is a better response, a contradiction. □

Lemma 2. *If π_I is a best-response to π_J such that $\pi_I < \pi_J \leq q_J(\pi \mid \omega_J)$, then $\pi_I = q_I(\pi \mid \omega_I)$.*

Proof. Suppose toward a contradiction that $\pi_I < \pi_J \leq q_J(\pi \mid \omega_J)$ and $\pi_I \neq q_I(\pi \mid \omega_I)$. If $\pi_I > q_I(\pi \mid \omega_I)$ a contradiction is immediate because hospital I can profitably deviate to $\pi_I - \epsilon$ for some $\epsilon > 0$ that is small such that $\pi_I > q_I(\pi_I - \epsilon, \pi_J \mid \omega_I)$. If $\pi_I < q_I(\pi \mid \omega_I)$ then observe that $\pi_J < q_J(\pi \mid \omega_J)$ by lemma 1. Consider the deviation by hospital I to $\pi_I + \epsilon$ for $\epsilon > 0$ small such that $\pi_I + \epsilon \leq q_I(\pi_I + \epsilon, \pi_J \mid \omega_I)$

- If $\omega = \omega_I$, then hospital I is strictly better off as he has exchanged doctors in $[\pi, \pi_I + \epsilon]$ for doctors in $[q_I(\pi_I + \epsilon, \pi_J \mid \omega_I), q_I(\pi \mid \omega_I)]$, a strict improvement.
- If $\omega = \omega_J$, then hospital I is weakly better off because every doctor in $[\pi, \pi_I + \epsilon]$ is still available to hospital I.

Hence, the deviation by hospital I to $\pi_I + \epsilon$ is a strict improvement. □

Corollary 1. *If π is a equilibrium such that $\pi_L < \pi_H$, then $\pi_L = q_L(\pi \mid \omega_L)$.*

Corollary 2. *There is no equilibrium π such that $\pi_L > \pi_H$.*

Proof. Suppose toward a contradiction that $\pi_L > \pi_H$. Then $q_L(\pi \mid \omega_L) > \pi_L$ by Lemma 1. Lemma 2 implies that $\pi_H = q_H(\pi \mid \omega_H)$. But by definition, $q_L(\pi \mid \omega_L) < q_H(\pi \mid \omega_H)$, a contradiction. □

Proof of Theorem 1

Proof. The proof is as follows. First, we conjecture that for some π following holds

$$q_H(\pi \mid \omega_L) = q_L(\pi \mid \omega_H) \leq \pi_L \leq q_L(\pi \mid \omega_L) \leq \pi_H < q_H(\pi \mid \omega_H) \qquad (*)$$

Under $(*)$ we can calculate q as a function of π. We then solve for the first-order conditions to find π^* and see that $(*)$ holds under π^*. We then check if either hospital has a profitable deviation $\tilde{\pi}$. When checking for profitable deviations, we use Lemmas 1 and 2 to make the search tractable.

To show uniqueness, we again use Lemmas 1 and 2, and also Corollaries 1 and 2, to rule out many possible alternative equilibria π'. We then check any remaining π' by hand.

We use the computer software *Mathematica* to assist with algebraic manipulation. The commands executed can be found in the Supplementary Materials.

First, observe that $(*)$ implies

$$q_H(\pi \mid \omega_L) = 1 - 2\text{Cap}$$

$$q_L(\pi \mid \omega_H) = 1 - 2\text{Cap}$$

$$q_L(\pi \mid \omega_L) = \frac{1 - 2\text{Cap} + \pi_H r - 2\pi_L r}{1 - r}$$

$$q_H(\pi \mid \omega_H) = \frac{1 - \text{Cap} - \pi_H r}{1 - r}$$

We then calculate

$$u_H(\pi) = \frac{\text{Cap}^2(6 - 8r) + r(1 + 3(-2 + \pi_H)\pi_H - 2(-2 + \pi_L)\pi_L - 2(\pi_H - \pi_L)^2 r) + 8\text{Cap}(-1 + (1 + \pi_H - \pi_L)r)}{8(-1 + r)}$$

$$u_L(\pi) = \frac{\text{Cap}^2(10 - 8r) + r(-1 - 3(-2 + \pi_H)\pi_H + 2(-2 + \pi_L)\pi_L + 2(\pi_H - \pi_L)^2 r) + 8\text{Cap}(-1 + (1 - \pi_H + \pi_L)r)}{8(-1 + r)}$$

We see that $u_H(\pi)$ is concave in π_H, and $u_L(\pi)$ is concave in π_L. Hence, we need only solve the first-order conditions, which yields:

$$\pi_H^*(r) = 1 - \frac{4\text{Cap}}{3 + r}$$

$$\pi_L^*(r) = 1 - \frac{6\text{Cap}}{3 + r}$$

Under $\pi^*(r)$ we see that $(*)$ holds.

We now check for profitable deviations.

Suppose (toward a contradiction) that $\tilde{\pi}_L$ is a profitable deviation for hospital L that is a best-response to π_H^*. Consider the following cases:

- If $\tilde{\pi}_L > \pi_L^*$, then we observe $q_L(\tilde{\pi}_L, \pi_H^* \mid \omega_L) < q_L(\pi^* \mid \omega_L) < \tilde{\pi}_L$, a contradiction to Lemma 1.

- If $\tilde{\pi}_L < \pi_L^*$, then we observe $q_L(\tilde{\pi}_L, \pi_H^* \mid \omega_L) > q_L(\pi^* \mid \omega_L) > \tilde{\pi}_L$, a contradiction to Lemma 2.

Hence, hospital L has no profitable deviations.

Suppose (toward a contradiction) that $\tilde{\pi}_H$ is a profitable deviation for hospital H that is a best-response to π_l^*. Consider the following cases:

- If $\tilde{\pi}_H > \pi_H^*$ such that $\tilde{\pi}_H > q_H(\tilde{\pi}_H, \pi_L^* \mid \omega_H)$, then this is a contradiction to lemma 1.

- If $\tilde{\pi}_H > \pi_H^*$ such that $\tilde{\pi}_H \leq q_H(\tilde{\pi}_H, \pi_L^* \mid \omega_H)$, then observe that $q_L(\tilde{\pi}_H, \pi_L^* \mid \omega_L) \leq \tilde{\pi}_H$. Thus $(*)$ continues to hold. This is a contradiction to π_H^* maximizing $u_H(x, \pi_L^*)$ under $(*)$.

- If $\tilde{\pi}_H < \pi_H^*$ such that $\tilde{\pi}_H > \pi_L^*$, then observe that $q_L(\tilde{\pi}_H, \pi_L^* \mid \omega_L) < \pi_L^*$. This case is considered in the cell "DEVIATION 1" in the supplementary *Mathematica* notebook. There, we derive q given π_L^* and $\tilde{\pi}_H$. We then derive hospital H's utility, show that it is concave in $\tilde{\pi}_H$, and that the first order conditions are satisfied for $\tilde{\pi}_H > \pi_H^*$. Hence, the best-response must be $\tilde{\pi}_H = \pi_H^*$, a contradiction.

- If $\tilde{\pi}_H < \pi_H^*$ such that $\tilde{\pi}_H \leq \pi_L^*$ and $\tilde{\pi}_H \geq q_H(\tilde{\pi}_H, \pi_L^* | \omega_L)$, then observe that $q_L(\tilde{\pi}_H, \pi_L^* | \omega_L) = q_L(\tilde{\pi}_H, \pi_L^* \mid \omega_H)$ and $q_H(\tilde{\pi}_H, \pi_L^* \mid \omega_L) \geq q_L(\tilde{\pi}_H, \pi_L^* \mid \omega_L)$. We note that the reversal in the order of $q_L(\tilde{\pi}_H, \pi_L^* \mid \omega_L)$ and $q_H(\tilde{\pi}_H, \pi_L^* \mid \omega_L)$ occurs precisely when $\tilde{\pi}_H$ and π_L^* switch order. This case is considered in the cell "DEVIATION 2" in the supplementary *Mathematica* notebook. There, we derive q given π_L^* and $\tilde{\pi}_H$. We then derive hospital H's utility, show that it is concave in $\tilde{\pi}_H$, and that the first order conditions are satisfied for $\tilde{\pi}_H > \pi_H^*$. Hence, the best response must be $\tilde{\pi}_H = \pi_L^*$, a contradiction.

- If $\tilde{\pi}_H < \pi_H^*$ such that $\tilde{\pi}_H \leq \pi_L^*$ and $\tilde{\pi}_H < q_H(\tilde{\pi}_H, \pi_L^* \mid \omega_L)$, then observe by Lemma 1 this cannot be a best response, a contradiction.

Hence, π^* is an equilibrium.

To show uniqueness, we note that by corollaries 1 and 2 (and lemma 1), we need only consider other equilibria π' of the form $\pi'_L = \pi'_H$. Suppose (toward a contradiction) that such a π' is an equilibrium. Then by Lemma 1 $q_H(\pi' \mid \omega_H) \geq \pi'_H \geq q_H(\pi' \mid \omega_L)$ and $q_L(\pi' \mid \omega_L) \geq \pi'_L \geq q_L(\pi' \mid \omega_H)$. Observe that $q_H(\pi' \mid \omega_L) = q_L(\pi' \mid \omega_L)$. This implies that either $q_H(\pi' \mid \omega_H) = 1 - 2\text{Cap}$ or $q_L(\pi' \mid \omega_L) = 1 - 2\text{Cap}$. We then observe that $q_L(\pi' \mid \omega_L) \leq q_H(\pi' \mid \omega_H)$. Thus, $q_L(\pi' \mid \omega_L) = 1 - 2\text{Cap}$. Hence, $\pi'_H = \pi'_L = 1 - 2\text{Cap}$.

In the cell "ALTERNATE EQUILIBRIUM BEST RESPONSE" in the attached *Mathematica* notebook, we show that the deviation

$$\pi^*_H(r) = 1 + 2\text{Cap}\Big(-1 + \frac{1}{3 - 2r}\Big)$$

is a profitable deviation, a contradiction. Hence, π^* is the unique equilibrium, and the theorem is proved.

$\square$